ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed
Computing

Computational Thinking

Roger Wattenhofer

wattenhofer@ethz.ch

Fall 2020

Contents

Introduction

Computation is everywhere, but what is computation actually? In this lecture we will discuss the power and limitations of computation. Computational thinking is about understanding machine intelligence: What is computable, and how efficiently?

Understanding computation lies at the heart of many exciting scientific, social and even philosophical developments. Computational thinking is more than programming a computer, it means thinking in abstractions. Consequently, computational thinking has become a fundamental skill for everyone, not just computer scientists. For example, functions which can easily be computed but not inverted are at the heart of understanding data security and privacy. Machine learning on the other hand has given us fascinating new tools to teach machines how to learn function parameters. Thanks to clever heuristics, machines now appear to be capable of solving complex cognitive tasks. To give just one more example: How can we design the best electronic circuit for a given problem? In this class, we study various problems together with the fundamental theory of computation.

While computation is predominantly an engineering discipline, some of our insights will be highly surprising. One may claim that this class is going to discuss the three biggest computational surprises. Discussing these surprises will add a philosophical touch to the class.

The weekly lectures will be based on blackboard discussions and coding demos, supported by a script and coding examples. The course uses Python $\rightarrow$ notebook as a programming language. Python is popular and intuitive, a programming language that looks and feels like human instructions. The lecture will feature weekly exercises, on paper and in Python.

Have fun!

Chapter 1

Algorithms

The term "computer" used to be a job description for a person doing the same tedious computations over and over, hopefully without error. When electrical computers became available, these human computers often transitioned to become computer programmers. Instead of doing the computations themselves, they told the computer what to do.

Definition 1.1 (Algorithm). *An **algorithm** is a sequence of computational instructions that solves a class of problems. Often the algorithm computes an output for a given input, i.e., a mathematical function.*

Remarks:

- While the number of algorithms is theoretically unlimited, surprisingly many problems can be solved with just a few algorithmic paradigms that we will review in this chapter. A simple yet powerful algorithmic concept is recursion. Let us start with an example.

1.1 Recursion

You have won an *all-you-can-carry* run through an electronics store. The rules are simple: Whatever you manage to carry, you can have for free. Being well-prepared you bring a high-capacity backpack to the event. Which items should you put into your backpack such that you can carry the maximum possible value out of the store?

Problem 1.2 (Knapsack). *An item is an object that has a* `name`*, a* `weight` *and* $\rightarrow$ notebook *a* `value`*. Given a list of* `items` *and a knapsack with a weight* `capacity`*, what is the maximal value that can be packed into the knapsack?*

Remarks:

- An algorithm solving Knapsack computes a function; the inputs are the set of possible `items` and the `capacity` limit of the knapsack, the output is the maximal possible `value`.

- A simple way to solve Knapsack is to check for every `item` whether it should be packed into the knapsack or not, expressed as the following recursion:

```python
def knapsack(items, capacity):
    if len(items) == 0:
        return 0
    first, *rest = items
    take = 0
    if first.weight <= capacity:
        take = knapsack(rest, capacity-first.weight) + first.value
    skip = knapsack(rest, capacity)
    return max(take, skip)
```

Algorithm 1.3: A recursive solution to Knapsack.

Remarks:

- Algorithm 1.3 may look like pseudo-code, but really is correct Python.

- In Lines 7 and 8, the algorithm calls itself. This is called a recursion.

Definition 1.4 (Recursion). *An algorithm that splits up a problem into sub-problems and invokes itself on the sub-problem is called a **recursive algorithm**. A **recursion** ends when reaching a simple base case that can be solved directly. Also, see Definition 1.4.*

Remarks:

- In mathematics, we find a similar structure in some prominent inductive functions such as the Fibonacci function.

- Recursive algorithms are often easy to comprehend, but not necessarily fast.

- How can we measure "fast"?

Definition 1.5 (Time Complexity). *The **time complexity** of an algorithm is the number of basic arithmetic operations $(+, -, \times, \div, etc.)$ performed by the algorithm with respect to the size n of the given input.*

Remarks:

- Each variable assignment, `if` statement, iteration of a `for` loop, comparison (`==`, `<`, `>`, etc.) or `return` statement also counts as one basic arithmetic operation, and so do function calls (`len()`, `max()`, `knapsack()`).

- Unfortunately, there is no agreement on how the size of the input should be measured. Often the input size n is the number of input items. If input items get large themselves (e.g., the input may be a single but huge number), n refers to the number of bits needed to represent the input.

- We are usually satisfied if we know an approximate and asymptotic time complexity. The time complexity should be a simple function of n, just expressing the biggest term as n goes to infinity, ignoring constant factors. Such an asymptotic time complexity can be expressed by the "big O" notation.

Definition 1.6 ($\mathcal{O}$-notation). *The $\mathcal{O}$-**notation** is used to denote a set of functions with similar asymptotic growth. More precisely,*

$$\mathcal{O}(f(n)) = \left\{ g(n) \,\middle|\, \lim_{n \to \infty} \frac{g(n)}{f(n)} < \infty \right\}.$$

Remarks:

- In other words, $\mathcal{O}(f(n))$ is the set of functions $g(n)$ that asymptotically do not grow much faster than $f(n)$.

- For example, $\mathcal{O}(1)$ includes all constants and $\mathcal{O}(n)$ means "linear in the input size n".

- In other words, the $\mathcal{O}$-notation is quite crude, but nevertheless useful, both in theory and practice.

- Other useful asymptotic notations are $\Omega()$ for lower bounds, but also $o(), \omega(), \Theta()$, etc.

Lemma 1.7. *The time complexity of Algorithm 1.3 is $\mathcal{O}(2^n)$.*

Proof. Each call of the `knapsack()`-procedure performs constantly many basic arithmetic operations itself and makes (at most) two additional calls to the `knapsack()`-procedure. Hence, it suffices to count the total number of `knapsack()`-invocations. We get 1 invocation on the first item, at most 2 on the second, 4 on the third, ..., and 2^{n-1} on the last. Hence, there are less than 2^n invocations of the `knapsack()`-function. $\square$

Remarks:

- The time complexity of Algorithm 1.3 is exponential in the number of items. Even if there were only $n = 100$ items to be evaluated, the currently fastest supercomputer in the world would take $2^{100} \text{ ops}/(148 \cdot 10^{15} \text{ ops/s}) \approx 271\,000$ years to compute our `knapsack` function. So for many realistic inputs, Algorithm 1.3 is not usable. We need a better approach!

1.2   Greedy

What about sorting all the items by their `value`-to-`weight` ratio, and then simply greedily packing them!?

```
1  def knapsack(items, capacity):
2    items.sort(key=lambda item : -item.value/item.weight)
3    value = 0
4    for item in items:
5      if item.weight <= capacity:
6        capacity -= item.weight
7        value += item.value
8    return value
```

Algorithm 1.8: A naive greedy algorithm for Knapsack.

Remarks:

- Algorithm 1.8 is fast, with a time complexity of $\mathcal{O}(n \log n)$, just for calling the sorting function in Line 2. So a large input is no problem.

- Also, the output of Algorithm 1.8 often seems reasonable. However, Algorithm 1.8 does not solve Knapsack optimally. For example, assume a capacity 6 knapsack, two items each with value 3 and weight 3, and one higher-ratio item with value 5 and weight 4.

- Can we gain a speed-up from first sorting the elements?

1.3 Backtracking

Definition 1.9 (Backtracking). *A **backtracking** algorithm solves a computational problem by constructing a candidate solution incrementally, until either a solution or a contradiction is reached. In case of a contradiction, the algorithm "backtracks" (i.e. reverts) its last steps to a state where another solution is still viable. For efficient backtracking algorithms, there are two main ingredients:*

- ***Look-ahead:*** *We order the search space such that the most relevant solutions come up first.*

- ***Pruning:*** *We identify sub-optimal paths early, allowing to discard parts of the search space without explicitly checking.*

Remarks:

- Algorithm 1.3 was an inefficient backtracking algorithm.

- Our look-ahead idea is to sort the items by value-to-weight ratio as in Algorithm 1.8.

- The algorithm prunes the solution space if it cannot possibly achieve the best solution so far.

```
1  def knapsack(items, capacity):
2    items.sort(key=lambda item: -item.value/item.weight)
3    return bt(items, capacity, 0)
4
5  def bt(items, capacity, missing):
6    if len(items) == 0:
7      return 0
8    first, *rest = items
9    if first.value / first.weight * capacity < missing:
10     return 0  # branch is worse than the best previous solution
11   take = 0
12   if first.weight <= capacity:
13     take = bt(rest, capacity-first.weight, missing-first.value)
14     take += first.value
15   skip = bt(rest, capacity, max(take, missing))
16   return max(take, skip)
```

Algorithm 1.10: An efficient backtracking solution to Knapsack.

Remarks:

- The missing parameter is the additional value that is required to surpass the previously best solution.

- The time complexity of Algorithm 1.10 is still $\mathcal{O}(2^n)$ in the worst case. Can we do better?

1.4 Dynamic Programming

Definition 1.11 (Dynamic Programming). ***Dynamic programming (DP) is** a technique to reduce the time complexity of an algorithm by utilizing extra memory. To that end, a problem is divided into sub-problems that can be optimized independently. Intermediate results are stored to avoid duplicate computations.*

Remarks:

- Knapsack can be solved with dynamic programming. To that end, we store a value matrix V where V[i][c] is the maximum value that can be achieved with capacity c using only the first i items.

```
1  def knapsack(items, capacity):
2    n = len(items)
3    V = zero matrix of size (n+1)×(capacity+1)
4    for item i in items:
```

```
5       for c in range(capacity+1):
6           V[i+1][c] = max(V[i][c-item.weight] + item.value, V[i][c])
7       return V[n][capacity]
```

Algorithm 1.12: A dynamic programming solution to Knapsack.

Remarks:

- Note that Algorithm 1.12 is not correct Python. Line 3 is just pseudo-code, far from actual Python notation. Line 4 could be Python, but unfortunately needs an extra `enumerate()` function.

- Line 6 is incorrect: If `item.weight > c`, `c-item.weight` becomes negative. The programmer of Algorithm 1.12 assumed that accessing a negative index of an array returns 0; however, most programming languages return an error. We can fix Line 6 by adding the conditional expression `if c >= item.weight else 0` to the second term of the `max()` function.

- The time complexity of Algorithm 1.12 is $\mathcal{O}(n \cdot \texttt{capacity})$. In Definition 1.5 we postulated that the time complexity should be a function of n. So the DP approach only makes sense when `capacity` is a natural number with $\texttt{capacity} < 2^n/n$. .

Definition 1.13 (Space Complexity). *The **space complexity** of an algorithm is the amount of memory required by the algorithm, with respect to the size n of the given input.*

Remarks:

- As for Definition 1.5, we are usually satisfied if we know the approximate (asymptotic) space complexity.

- Also, the amount of memory can be measured in bits or memory cells.

- The space complexity of Algorithm 1.12 is $\mathcal{O}(n \cdot \texttt{capacity})$.

- For reasonably small `capacity`, Algorithm 1.12 is faster than Algorithms 1.3–1.10, but is it correct?

Lemma 1.14. *Assuming that all `items` have integer `weights`, Algorithm 1.12 solves Knapsack correctly.*

Proof. We show the correctness of each entry in the matrix `V` by induction. As a base case, we have `V[0] = [0, ..., 0]` since without item, no value larger than 0 can be achieved. For the induction step, assume that `V[i]` correctly contains the maximum values that can be achieved using only the first `i` items. When we set a value `V[i+1][c]`, we can either include the item `i+1` or select the optimal solution for Knapsack with capacity using only the first `i` items. Algorithm 1.12 stores the `max()` of these two values in `V[i+1][c]` (for all `c` $\in$ $\{0, \ldots, \texttt{capacity}\}$), which is thus optimal.

Hence, the value `V[n][capacity]` contains the maximum value that can be achieved with the weight `capacity`, using any combination of the n items. □

Remarks:

- Line 6 of Algorithm 1.12 is typical for dynamic programming algorithms: either the previous best solution can be improved, or it remains unchanged. This is called Bellman's principle of optimality.

- The computation order of Algorithm 1.12 is important. For example, we can only compute the entry `V[i+1][c]` once we have computed both `V[i][c-item.weight]` and `V[i][c]`.

- The sub-problem dependencies can be visualized as a dependency graph. In order to apply dynamic programming, this graph must be a directed acyclic graph (DAG).

- Algorithm 1.12 is a so-called *bottom-up* dynamic programming algorithm as it begins computing the entries of matrix `V` starting with the simple cases.

- :But do we really need to compute the entire matrix `V`?

Definition 1.15 (Memoization). ***Memoization** generally refers to a technique that avoids duplicate computations by storing intermediate results.*

```
1   def knapsack(items, capacity, memo={}):
2       index = (len(items), capacity)
3       if index in memo:
4           return memo[index]
5       if len(items) == 0:
6           return 0
7       first, *rest = items
8       take = 0
9       if first.weight <= capacity:
10          take = knapsack(rest, capacity-first.weight, memo)
11          take += first.value
12      skip = knapsack(rest, capacity, memo)
13      memo[index] = max(take, skip)
14      return memo[index]
```

Algorithm 1.16: A top-down DP solution to Knapsack.

Remarks:

- Memoization can be used to implement *top-down* DP algorithms.

- This is not so different from what we started out with in Algorithm 1.3!

- We only changed Line 1 and added Line 2 to set up memoization, which is then used in Lines 3–4 and 13–14.

- Top-down DP is inheriting the best of recursion and bottom-up DP. Consequentially, the time complexity of Algorithm 1.16 is

$$\mathcal{O}(\min(2^n, n \cdot \texttt{capacity})).$$

- So far we have learned a family of related algorithmic techniques: recursion, backtracking, dynamic programming, and memoization. Together, this family can help solving many demanding algorithmic problems.

- However, there are powerful algorithmic paradigms beyond this family of techniques, for instance linear programming.

1.5   Linear Programming

So far, we were only considering unsplittable items. However, for liquid goods, Knapsack can be solved quickly using a greedy method (Algorithm 1.8). What if we had more than one constraint?

Problem 1.17 (Liquid Knapsack). *A* `beverage` *has a* `name`*, a* `value` *per liter* → notebook *and a preparation* `time` *per liter. Given t hours to prepare for a party and a fridge with a storage* `capacity`*, what is the maximal value that can be prepared and stored in the fridge?*

Remarks:

- With more than one constraint, the greedy method does not work.

- However, this problem has a nice property: the objective and the constraints are linear functions of the quantity of each prepared beverage. We call such problems *linear programs*.

Definition 1.18 (Linear Program or LP).   *A **linear program (LP)** is an* → notebook *optimization problem with n variables and m linear inequalities*

$$
\begin{array}{rrrrrr}
a_{11}x_1 + & a_{12}x_2 + & \ldots + & a_{1n}x_n & \leq & b_1 \\
a_{21}x_1 + & a_{22}x_2 + & \ldots + & a_{2n}x_n & \leq & b_2 \\
\vdots & \vdots & & \vdots & & \vdots \\
a_{m1}x_1 + & a_{m2}x_2 + & \ldots + & a_{mn}x_n & \leq & b_m
\end{array}
$$

We are interested in finding a point $\boldsymbol{x} = (x_1, \ldots, x_n)^T$ with $x_i \geq 0$ respecting all these constraints and maximizing a linear function

$$f(\boldsymbol{x}) = c_1 x_1 + c_2 x_2 + \cdots + c_n x_n$$

*where a_{ij}, b_i and c_i are given real-valued parameters. We call the point $\boldsymbol{x}$ an **optimum** of the LP.*

Remarks:

- There is also a short hand notation using linear algebra

$$\max\{\boldsymbol{c}^T \boldsymbol{x} \mid \mathbf{A}\boldsymbol{x} \leq \boldsymbol{b}, \boldsymbol{x} \geq 0\},$$

where $\mathbf{A}$ is the matrix with entries a_{ij} and $\boldsymbol{b}$ and $\boldsymbol{c}$ the vectors given by the b_i and c_i, respectively.

- In general, if you have the problem of maximizing or minimizing a linear function under constraints that are linear (in)equalities, there is a way to formulate it in above canonical form. For instance, a constraint $\boldsymbol{a}^T \boldsymbol{x} = b$ can be rewritten as a combination of $\boldsymbol{a}^T \boldsymbol{x} \leq b$ and $\boldsymbol{a}^T \boldsymbol{x} \geq b$ which itself can be rewritten as $-\boldsymbol{a}^T \boldsymbol{x} \leq -b$. Also, minimizing a linear function with coefficients $c_1, \ldots, c_n$ is the same as maximizing a linear function with coefficients $-c_1, \ldots, -c_n$.

- It is possible to model some functions which do not look linear at first sight. For example, minimizing an objective function $f(x) = |x|$ can be expressed as $\min\{t | x \leq t, -x \leq t\}$.

Definition 1.19 (Feasible Point). *Given an LP, a point is **feasible** if it is a solution of the set of constraints.*

Remarks:

- Geometrically, the set of feasible points of an LP corresponds to an n-dimensional convex polytope. The hyperplanes bounding the polytope are given by the restricting inequalities.

- Polytopes are a generalization of 2D polygons to an arbitrary number of dimensions. Convexity, however, deserves a more formal definition.

Definition 1.20 (Convex Set). *A set of points in $\mathbb{R}^n$ is **convex** if for any two points of the set, the line segment joining them is also entirely included in the set.*

Lemma 1.21. *The set of feasible points of an LP is convex.*

Proof. Given two feasible points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, any point in the line segment joining them can be written as $\boldsymbol{x}_1 + \lambda(\boldsymbol{x}_2 - \boldsymbol{x}_1)$ for $\lambda \in [0,1]$. For any constraint $\boldsymbol{a}^T \boldsymbol{x} \leq b$, we compute

$$\boldsymbol{a}^T[\boldsymbol{x}_1 + \lambda(\boldsymbol{x}_2 - \boldsymbol{x}_1)] = (1-\lambda)\boldsymbol{a}^T \boldsymbol{x}_1 + \lambda \boldsymbol{a}^T \boldsymbol{x_2} \leq (1-\lambda)b + \lambda b = b.$$

□

Definition 1.22. *Given an LP, we call **polytope** the set of feasible points.* → notebook *A constraint $\boldsymbol{a}^T \boldsymbol{x} \leq b$ is **tight** at $\boldsymbol{x}$ if $\boldsymbol{a}^T \boldsymbol{x} = b$. For an LP with n variables, feasible points activating n (resp. n − 1) linearly independent constraints are called the **nodes** (resp. **edges**) of the polytope. Each edge links two nodes $\boldsymbol{x}_1, \boldsymbol{x}_2$ with n − 1 common activating constraints; we say that the two nodes $\boldsymbol{x}_1, \boldsymbol{x}_2$ are **neighbors**.*

Remarks:

- A polytope can be *unbounded*, i.e. infinitely large. In some cases, it is even possible to have an infinitely large solution, e.g./ $\max\{x|x \geq 0\}$. Following our definition, the LP does not admit an optimum in this case.

- In order to solve an LP, one has to find a point in the polytope that maximizes our objective function $f(\boldsymbol{x})$.

Theorem 1.23. *If the polytope of an LP is bounded, then at least one node of the polytope is an optimum of the LP.*

Proof. For any value y that the objective function can take, the set of points reaching this value is given by the hyperplane $\boldsymbol{c}^T\boldsymbol{x} = y$. We can find an optimum of the LP by sliding this hyperplane until the boundary of the polytope is reached, which happens at some node of the polytope.   □

Remarks:

- One popular method exploiting Theorem 1.23 for solving LPs is the simplex algorithm. The idea is simple: starting from a node of the LP polytope, greedily jump to a neighboring node having a better objective until you cannot improve the solution anymore.

```
1  def simplex(polytope, f, x):
2    for y in neighbors(x, polytope):
3      if f(y) > f(x):
4        return simplex(polytope, f, y)
5    return x
```

Algorithm 1.24: Simplex Algorithm.

Remarks:

- While the simplex algorithm performs well in practice, there are instances where its time complexity is exponential in the size of the input. Other LP algorithms known as interior point methods are provably fast.

- In practice, we do not build and store the whole polytope of the LP, as the polytope could have an exponential number of nodes! Instead, we represent a node as a set of tight constraints. To find its neighbors, we remove a constraint of the set, add another constraint and check if the point is feasible.

- The node returned by the simplex algorithm is better than any neighboring node by construction, but how can we convince ourselves that no other point anywhere in the feasible polytope is better?

Definition 1.25 (Local Optimum). *A feasible node $\boldsymbol{x}$ is a **local optimum** if $f(\boldsymbol{x}) \geq f(\boldsymbol{y})$ for any neighboring node $\boldsymbol{y}$.*

Remarks:

- In contrast to a local optimum, an optimum (from Definition 1.18) is called *global optimum*.

- While it is easy to find a local optimum, finding a global optimum is often difficult. However, it turns out that every local optimum of an LP is also a global optimum!

Theorem 1.26. *The node $\boldsymbol{x}^*$ returned by the simplex algorithm is an optimum.*

Proof. Let us consider the hyperplane $\boldsymbol{c}^T\boldsymbol{x} = f^*$, where $f^* = \boldsymbol{c}^T\boldsymbol{x}^*$. We know that all the neighbors of node $\boldsymbol{x}^*$ are on the side $\boldsymbol{c}^T\boldsymbol{x} \leq f^*$. Since the polytope is convex, we know that the whole polytope must be on this side of the hyperplane. Hence no node x' in the polytope can be on the side $\boldsymbol{c}^T\boldsymbol{x} > f^*$, and hence the node $\boldsymbol{x}^*$ is a global optimum.   □

Figure 1.27: Illustration of Theorem 1.26. The neighbors of $\boldsymbol{x}^*$ are x_1 and x_2.

Remarks:

- So we have seen that every local optimum of an LP is also a global optimum. This important property in optimization is true for convex functions in general!

Definition 1.28 (Convexity). *A function $f : D \to \mathbb{R}$ is convex if for any $\boldsymbol{x}, \boldsymbol{y} \in D$ and any $\lambda \in [0,1]$: Any point $\boldsymbol{z} = \lambda\boldsymbol{x} + (1-\lambda)\boldsymbol{y}$ is in D as well. In addition $f(\boldsymbol{z}) \leq \lambda f(\boldsymbol{x}) + (1-\lambda)f(\boldsymbol{y})$.*

If f is differentiable twice, then f is convex if and only if its hessian $\nabla^2 f$ is positive semidefinite, i.e., the hessian has non-negative eigenvalues.

Example 1.29. *The following functions are convex:*

- $f(x) = x^2$

- $f(\boldsymbol{x}) = ||\boldsymbol{x}||_2$ *(Euclidean norm)*

- $f(p_1, \ldots, p_k) = \sum_{p_i} p_i \log p_i$ *(reverse entropy)*

- $f(A) = -\log \det(A)$ *for a positive semidefinite matrix* $\mathbf{A}$

Remarks:

- We call Algorithm 1.24 with x being any node of the polytope. But wait, how do we find such a node?! It turns out that we can construct an auxiliary LP:

Definition 1.30 (Phase 1 LP). *Given an LP*

$$\max\{c^T x \mid \mathbf{A}x \leq b, x \geq 0\},$$

we build the so-called **phase 1 LP** *by replacing every constraint* $a_i^T x \leq b_i$ *with* $a_i^T x - y_i \leq b_i$, *introducing a new artificial variable* y_i. *If we minimize all artificial variables* y_i, *we get:*

$$\max\{-\mathbf{1}^T y \mid \mathbf{A}x - \mathbf{I}y \leq b, x \geq 0, y \geq 0\}.$$

Lemma 1.31. *Setting each* $x_i = 0$ *and each* $y_i = \max(0, -b_i)$ *yields a feasible node of the phase 1 LP.*

Proof. With each original variable $x_i = 0$, each constraint is reduced to $-y_i \leq b_i$, which is satisfied when $y_i = \max(0, -b_i)$.

Also, this point is a node of the polytope: Algebraically, a point is a node if at least n linearly independent constraints are tight at this point. The constraint $x_i \geq 0$ is tight for each original variable x_i and either $a_i x - y_i \leq b_i$ or $y_i \geq 0$ is tight for each artificial variable y_i, depending on the sign of b_i. Thus, the number of tight constraints is at least equal to the number of variables, and this point is a node of the polytope. $\square$

Lemma 1.32. *If the original LP is feasible, then the phase 1 LP will find a feasible node.*

Proof. If the original LP is feasible, then its polytope is not empty, i.e., there exists a feasible node x in the original LP. Together with $y = 0$, node x is also feasible in the phase 1 LP. Since $\max\{-\mathbf{1}^T y\} = \min\{\text{sum}(y)\}$ is optimal for $y = 0$, node x is optimal in the phase 1 LP. With Theorem 1.26, we know that the phase 1 LP will find such a node x. $\square$

```
def solveLP(A, b, c):
    x,y = simplex(polytope([A -I], b), -1, (0, max(0,-b)))
    if sum(y) == 0:
        return simplex(polytope(A, b), c, x)
    else:
        return 'no solution'
```

Algorithm 1.33: Two-phase simplex algorithm to solve LPs.

Remarks:

- Algorithm 1.33 is the complete procedure to solve an LP. This process is often called the *two-phase simplex algorithm*.

- In Python, one can solve an LP using the function `linprog` from the module `scipy.optimize`.

1.6 Linear Relaxation

Despite linear programming covering a broad class of problems, we are often confronted with *discrete* tasks. In practice we often need an integer solution.

Definition 1.34 (Integer Linear Programming or ILP). *An* **integer linear program (ILP)** *is an LP in which all variables are restricted to integers.*

Remarks:

- In a lot of combinatorial problems, variables are restricted to just two values $\{0, 1\}$. Such variables are called *indicator* ("to be or not to be") variables. We call such programs *binary ILPs*.

- Apart from LP and ILP, there exist many other optimization techniques: Mixed Integer Linear Programming (MILP) with both integer and continuous variables, Quadratic Programming (QP), Semidefinite Programming (SDP), ...

Problem 1.35 (ILP Knapsack). *We can model Knapsack (Problem 1.2) with capacity c and n items of value v_i and weight w_i as a binary ILP, using indicator variables x_i:*

$$\begin{aligned} maximize \quad & \sum v_i \, x_i \\ subject\ to: \quad & \sum w_i \, x_i \leq c \\ & x_i \in \{0, 1\}. \end{aligned}$$

Remarks:

- Unlike LPs, no efficient algorithm solving ILPs is known.

- It is tempting to relax the constraints $x_i \in \{0, 1\}$ to $0 \leq x_i \leq 1$, apply the simplex algorithm, and round the possible solution to the nearest feasible point.

Definition 1.36 (Linear Relaxation). *Given a binary ILP, we construct the* **linear relaxation** *of the LP by replacing the constraint $x \in \{0, 1\}^n$ with the constraint $0 \leq x_i \leq 1$.*

Remarks:

- However, in general, there is no guarantee that a linear relaxation finds the optimum.

- In the case of Knapsack, the solution of the linear relaxation is similar to Algorithm 1.8. All items i with a high value-to-weight ratio will get an indicator variable $x_i = 1$, all items with a low value-to-weight ratio will get an indicator variable $x_i = 0$. The critical item(s) in the middle will get a non-integer indicator variable which we must round down to 0 to get a valid solution. This solution can be arbitrarily bad, as the best (highest value-to-weight ratio) item might already be too heavy; we might end up without any object in the knapsack.

- However, a linear relaxation sometimes has the same optimum as its ILP. In particular, this is for some class of constraint matrices, the so-called totally unimodular matrices.

- A matrix is totally unimodular if every square submatrix has determinant -1, 0 or $+1$.

Problem 1.37 (Assignment Problem). *Given a list of customers and a list* $\rightarrow$ notebook *of cabs, how to match customers to cabs in order to minimize the total waiting time?*

Algorithm 1.38. *This problem can be modeled as an ILP. We denote the* $\rightarrow$ notebook *waiting time of customer i for cab j by $w_{i,j}$. Also, we introduce a set of indicator variables $x_{i,j}$ describing the assignment: $x_{i,j} = 1$ if and only if customer i is assigned to cab j. We get:*

$$\begin{aligned}
minimize\ & \sum_{i,j} x_{i,j} w_{i,j} \\
subject\ to: & \sum_j\ x_{i,j}\ &= 1 & \quad for\ each\ customer\ i \\
& \sum_i\ x_{i,j}\ &\leq 1 & \quad for\ each\ cab\ j \\
& \quad x_{i,j}\ &\in \{0,1\}
\end{aligned}$$

This ILP can be solved optimally with linear relaxation: the constraint matrix is totally unimodular.

1.7 Flows

Graphs and flows are useful algorithmic concepts, related to LPs and linear relaxations.

Definition 1.39 (Graph). *A graph G is a pair (V, E), where V is a set of **nodes** and $E \subseteq V \times V$ is a set of **edges** between the nodes. The number of nodes is denoted by n and the number of edges by m.*

Remarks:

- A *directed* graph $G = (V, E)$ is a graph, where each edge has a direction, i.e., we distinguish between edges (u, v) and (v, u). If all edges of a graph are undirected, then the graph is called *undirected*.

- In a directed graph, we note $\mathtt{in}(u)$ (resp. $\mathtt{out}(u)$) the set of edges entering (resp. leaving) node u.

- A *weighted* graph $G = (V, E, \omega)$ is a graph, where $\omega : E \to \mathbb{R}$ assigns a weight $\omega(e)$ for each edge $e \in E$.

- Weights can for instance be used for delay $d(e)$ or capacity $c(e)$ of an edge.

- In the rest of this chapter, we consider capacitated directed graphs.

- Consider a company that wants to optimize the flow of goods in a transportation network from their factory to a customer.

Definition 1.40 (Flow). *Formally, an s-t-**flow** from a source node s to a target node t is given as a function $f : E \to \mathbb{R}$ such that*

$$f(u, v) \leq c(u, v) \qquad for\ all\ (u, v) \in E \qquad (capacity\ constraints)$$

$$\textstyle\sum_{e \in \mathtt{in}(u)} f(e) = \sum_{e \in \mathtt{out}(u)} f(e)\ for\ all\ u \in V \setminus \{s, t\}\ (flow\ conservation)$$

*We call the total flow reaching t the **value** of f, i.e. $|f| = \sum_{(u,t) \in E} f(u, t)$.*

Problem 1.41 (Max-Flow). *What is the maximum flow that can be established between a* `source` *and a* `target` *node in a network?*

Remarks:

- Max-Flow can be written as an LP maximizing the value of the flow.

- Flows are also useful to model discrete data. Imagine traffic flow for example: every road as some capacity and at each intersection, every car getting in is expected to eventually get out!

- Fortunately, we can use the linear relaxation of the ILP and be guaranteed to have the optimal solution!

Theorem 1.42 (Integral Flow Theorem). *If the capacity of each edge is an integer, then there exists a maximum flow such that every edge has an integral flow.*

Proof. Assume you have an optimal but non-integral flow. If there is a path from s to t with every edge being non-integral, we can increase the flow on that path, so our original flow was not optimal. Hence, there cannot be a non-integral path from s to t.

Let u be a node adjacent to an edge e with non-integral flow. Then u needs at least another edge e' with non-integral flow because of flow conservation at node u. We can follow these non-integral edges. Since they cannot include both s and t, we must find a cycle C of non-integral edges. All edges in C can both change their flow by $\pm\varepsilon$, without changing the flow from s to t. We change the flow of all edges in C until a first edge in C has integral flow. Now we have one edge less with non-integral flow. If there is still an edge with non-integral flow, we repeat this procedure, until all edges have integral flow. $\square$

Remarks:

- Thanks to Theorem 1.42, we can solve a discrete maximum flow problem with the linear relaxation of the ILP formulation and the simplex algorithm!

- There are also more efficient algorithms, known as augmenting paths algorithms.

Definition 1.43 (Augmenting Path). *We define an **augmenting path** as a path from s to t such that the flow of each edge does not reach its capacity.*

Remarks:

- We can find an augmenting path in linear time, using a recursive algorithm!

```python
def find_augmenting_path(u, t, G, flow, visited):
  visited.insert(u)
  for v in G.neighbors(u):
    if v is not in visited and flow[u, v] < G.capacity[u, v]:
      path = find_augmenting_path(v, t, G, flow, visited)
      if len(path) > 0 or v == t:
        path.append((u, v))
        return path
  return []
```

→ notebook

Algorithm 1.44: Find augmenting path

Remarks:

- If the network has an augmenting path, then none of the edge of this path is at full capacity and we can add some flow on this path. This gives us a greedy algorithm: Find an augmenting path, push as much flow as possible on this path, then try again. This is known as the *Ford-Fulkerson* algorithm.

```python
def max_flow(s, t, G):
  while there is an augmenting path:
    visited = set()
    path = find_augmenting_path(s, t, G, visited):
    flow = update(G, flow, path)
  return flow # no augmenting path anymore
```

→ notebook

Algorithm 1.45: Ford-Fulkerson algorithm

Chapter Notes

The word algorithm is derived from the name of Muhammad ibn Musa al-Khwarizmi, a Persian mathematician who lived around AD 780–850. Some algorithms are as old as civilizations. A division algorithm was already used by the Babylonians around 2500 BCE [2]. Analyzing the time efficiency of recursive algorithms can be a difficult task. An easy but powerful approach is given by the master theorem [1]. Linear programming is an old concept whose origins lie in solving logistic problems during World War 2. Back in the days, the term

programming meant optimization, and not *coding*. This chapter was written in collaboration with Henri Devillez and Roland Schmid.

Bibliography

[1] Jon Louis Bentley, Dorothea Haken, and James B. Saxe. A general method for solving divide-and-conquer recurrences. *SIGACT News*, 12(3):36–44, September 1980.

[2] Jean-Luc Chabert, editor. *A History of Algorithms*. Springer Berlin Heidelberg, 1999.

Chapter 2

Complexity

In Chapter 1 we learned some nifty algorithmic techniques. However, sometimes we are dealing with a problem where none of these techniques seem to work! Should we give up? And what do we do after giving up?

2.1 Optimization vs. Decision

Definition 2.1 (Computational Problem). *A **computational problem** is defined as a (possibly infinite) set of inputs X and a (possibly empty) set of outputs Y for each input $x \in X$.*

Definition 2.2 (Optimization Problem). *In an **optimization problem**, given an input $x \in X$ and a quality measure function $q(x, y)$, the task is to find an output $y \in Y$, such that $q(x, y)$ is optimal.*

Remarks:

- We have seen several optimization problems in the previous chapter: What is the maximal value that can be packed into a knapsack? What is the optimal value of flow in a given network?

- We can also rewrite any linear program (Definition 1.18) as an optimization problem: we define the input to the problem as $x = (A, b, c)$, the set of admissible outputs as $Y = \{y | Ay \leq b\}$, and the function to be optimized as $f(x, y) = c^T y$.

Definition 2.3 (Decision Problem). *In a **decision problem**, given an input (x, d) and a quality measure function $q(x, y)$, where $x \in X$ and d is a decision value, the task is to check whether there exists a solution $y \in Y$ with $q(x, y) \geq d$. The output of the decision problem is a boolean value.*

Remarks:

- Note that we can reformulate any optimization problem as a decision problem. The corresponding decision problem asks a Yes/No question about quality of a solution. For example: Can we pack items of total value 100 into a knapsack? Is there a flow of value at least 5 in a given network?

Lemma 2.4. *By solving the optimization problem, we find a solution for the corresponding decision problem.*

Proof. Assume we have a decision problem, where we want to verify whether a problem has a solution of at least quality q. We can then run an algorithm for the corresponding optimization problem and check whether the solution is better or worse than q. ∎

Remarks:

- The opposite direction of Lemma 2.4 is usually also possible: if we have an algorithm for a decision problem, we can run it for many different decision values in order to determine the optimal value. This way, we can solve the corresponding optimization problem, but it is not clear how efficiently. The decision problem is therefore intuitively simpler than the optimization problem.

- In the first part of this chapter, we will focus on decision problems.

- In the previous chapter, we have discussed a measure for discussing the running time of algorithms – the time complexity (Definition 1.5). Algorithms for Knapsack seemed to have exponential time complexity, while network flow and matching were polynomial.

- When studying the complexity of computers, we are first and foremost interested in whether a problem can be solved in polynomial time. All problems that can be solved in polynomial time are considered "easy", and grouped together in a class.

Definition 2.5 (Complexity Class P). *The **complexity class** P contains all decision problems that can be solved in polynomial time, i.e. where we know an algorithm whose running time is in $\mathcal{O}(\mathsf{poly}(n))$, where n is the size of the input.*

Remarks:

- There are many problems that belong to class P: sorting, matrix multiplication, max value in an array, etc.

- In order to show that a problem is in P, we usually give an explicit algorithm and analyze the running time directly. There is also an indirect way of showing that a problem is in P: if we can show that some problem known to be in P can be used to solve our problem. Then our problem is "at most as difficult" as a problem in P and is therefore also in P.

Definition 2.6 (Polynomial Reduction). *Let $A = (x_A, q_A)$ and $B = (x_B, q_B)$ be two decision problems, and Y_A, Y_B denote the two admissible sets for the respective problems. We say that there exists a polynomial reduction from problem A to problem B, short $A \leq B$, if there exist a function $r : (X_A, Y_A) \to (X_B, Y_B)$ such that $f(x, y)$ has quality q_A if and only if $f(r(x, y))$ has quality q_B. In addition, there is an algorithm that can compute r in $\mathcal{O}(\mathsf{poly}(n))$, where n is the size of input A.*

Lemma 2.7. *The sorting problem has a polynomial reduction to the max value problem.*

Proof. Assume that the maximum value of an array with n values can be computed in polynomial time. We can compute the maximum value iteratively n times, and thus sort the array in polynomial time. $\square$

Remarks:

- We can use polynomial reductions between problems in order to determine classes of problems that have roughly the same difficulty.

Problem 2.8 (Clique). *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the **clique (decision) problem** we want to know whether there exists a subset of k nodes in G such that there is an edge between any two nodes.*

Problem 2.9 (Independent Set). *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the **independent set problem** we want to know whether there exists a subset of k nodes in G such that there is **no** edge between any two nodes in the subset.*

Theorem 2.10. *Clique $\leq$ Independent Set and Independent Set $\leq$ Clique.*

Proof. There is a simple function that maps one problem to the other: Let $K = V \times V$ denote all possible edges in a graph G. Given the graph $G = (V, E)$ in Clique, remove all edges E of this graph and add all edges in $K - E$. Then, all cliques in G will become independent sets in the new graph, and vice versa. The transformation can be computed in polynomial time. $\square$

Problem 2.11 (Vertex Cover). *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the **vertex cover problem** we want to know whether there exists a subset S of k nodes in G such that every edge in E has at least one of its two adjacent nodes in S.*

Theorem 2.12. *Vertex Cover $\leq$ Independent Set and Independent Set $\leq$ Vertex Cover.*

Proof. Given a any independent set S, all other nodes $V - S$ will be a vertex cover. The nodes in S are not connected to each other by definition, so the nodes $V - S$ must cover all edges. Likewise, if the nodes $V - S$ cover all edges, then the nodes S cannot have any edges between them, so S is an independent set.

If S is a maximal independent set, the vertex cover $V - S$ must be minimal: If the vertex cover was not minimal, we could find a bigger independent set, so S was not maximal. And vice versa. $\square$

Lemma 2.13. *The relation $\leq$ is transitive, i.e., if $A \leq B$ and $B \leq C$ then $A \leq C$.*

Proof. Let A, B and C be computational problems, such that $A \leq B$ and $B \leq C$. Then there exists an algorithm that can reduce problem A to C in $\mathcal{O}(\text{poly}(\text{poly}(n))) = \mathcal{O}(\text{poly}(n))$ time. $\square$

Remarks:

- Therefore, there are polynomial reductions between Clique and Vertex Cover as well. If you know how to solve one of these problems in polynomial time, you can solve the other two in polynomial time as well! Unfortunately, nobody knows whether these problems (and many others!) are in P.

- Using polynomial reductions, we can define a hierarchy of different problems.

- Even though we could not find an efficient (polynomial) algorithm for Knapsack as of now, we can try to see whether a simpler decision problem can be solved efficiently: Assume a friend (magically) suggests a solution to a seemingly difficult problem. Can you at least efficiently verify whether the proposed solution is admissible and gives the desired result?

Definition 2.14 (NP). *The class of non-deterministic polynomial problems (NP) contains decision problems, for which, given an input x and a possible solution y, the problem whether $q(x, y)$ has a certain quality can be verified in time $\mathcal{O}(\text{poly}(n))$, where n is the size of the input x.*

Remarks:

- The solution y in Definition 2.14 is also called certificate or witness.

- The term "non-deterministic" means that an algorithm is allowed to use parallelism and randomness. A non-deterministic algorithm can in parallel evaluate all possible solutions.

- The class NP contains a large set of computational problems. In fact, it is difficult to imagine a problem not in NP.

Lemma 2.15. *Knapsack is in NP.*

Proof. For Knapsack, a solution is a set of items. In order for Knapsack to be in NP, we would be given a set of items and a total value, for example 100. Verifying whether a given set S of items has a total value of 100 is just a sum. We also need to make sure that S is admissible. This can be also done in polynomial time. So Knapsack is in NP. $\square$

Lemma 2.16. *Clique is in NP.*

Proof. Given a graph and a proposed set of k nodes for a clique, we need to check whether the k proposed nodes are connected. This can be done by going through the list of all edges and counting the edges between the k nodes. If we counted $k \cdot (k + 1)/2$ edges, the proposed node set forms a clique. $\square$

Lemma 2.17. $P \subseteq NP$.

Proof. Every problem in P is automatically in NP. For problems in P, we can assume that the solution is first computed in Definition 2.14. $\square$

Lemma 2.18. *Let B be a problem in* NP *and A be a problem with $A \leq B$. Then, A is also in* NP.

Proof. Given an input x_A and a solution y_A to problem A, there exists a function r that transforms the pair (x_A, y_A) to a pair (x_B, y_B) of problem B. Since B is in NP, we can verify in polynomial time whether the pair (x_B, y_B) is a valid solution for B and therefore also whether (x_A, y_A) is valid for A. A is in NP, because r and the verification can be both computed in polynomial time in the input size of A.   □

Remarks:

- Observe that all three problems (Clique, Independent Set, Vertex Cover) are in the class NP, and they seem to be more difficult than the problems in P. Are there even more difficult problems? Is Knapsack more difficult? What is the most difficult problem?!

Definition 2.19 (NP-hard)**.** *A decision problem H is called* NP-*hard, if there exists a polynomial reduction from every problem in* NP *to H.*

Remarks:

- In other words, an NP-hard problem is at least as difficult as *all* the problems in NP.

- NP-hard problems do not necessarily have to be in NP.

- Are there also NP-hard problems in NP?

Definition 2.20 (NP-complete)**.** *A decision problem is called* NP-*complete if it is* NP-*hard and it is contained in* NP.

Lemma 2.21. *Let B be an* NP-*complete problem and C be a problem for which exists a polynomial reduction $B \leq C$. Then, C is* NP-*hard.*

Proof. Since B is NP-complete, there exists a polynomial reduction from every problem in NP to B. By using Lemma 2.13, we can show there also exists a polynomial reduction from every problem in NP to C. Therefore, C is NP-hard.   □

Remarks:

- Can we identify such an NP-complete problem? If yes, we could use polynomial reductions to show that there are many more problems contained in the NP-complete class. We will postpone this question to Section 2.2.

- Figure 2.22 visualizes some of the complexity classes and how they are related to each other. The figure just represents our current knowledge of how these classes relate to each other. There are still many open questions in complexity theory that could change the landscape once an answer is found. Also, there are many more classes.

- It is generally believed that there are some problems which are in NP but not in P. Or simply:

Conjecture 2.23. $\mathsf{P} \neq \mathsf{NP}$.
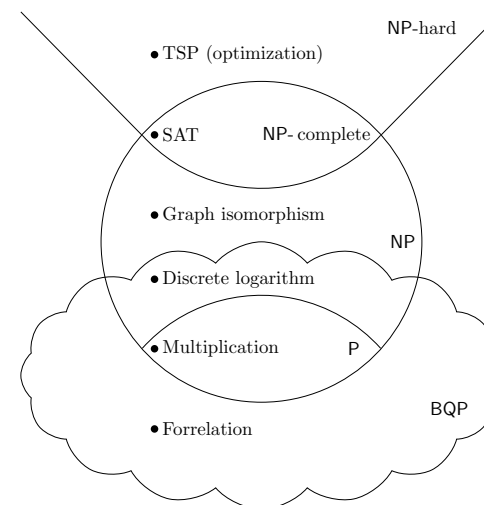
Figure 2.22: A complexity "pet" zoo: Diagram with different complexity classes and sample problems inside each class. The boolean satisfiability problem (SAT) will be introduced in the next section; the optimization version of the traveling salesperson problem (TSP) will be introduced in Section 2.4; the discrete logarithm and the graph isomorphism problems will be discussed in Chapter 3.

Remarks:

- The P versus NP question is one of the most important open scientific problems. There were many proof attempts, however, so far, without success.

- Another open question is whether quantum computers will allow to solve all difficult problems efficiently.

Definition 2.24 (BQP)**.** *A decision problem is in the class of **bounded-error quantum polynomial time** (BQP) problems, if it can be solved on a quantum computer in polynomial time and if for any input the probability to compute the wrong answer is at most $1/3$.*

Remarks:

- It is known that some problems from NP are also in BQP.

- The relation between the class of NP-complete problems and problems from BQP is unknown.

- It has been shown that there is a problem in BQP that is not in NP – the forrelation problem. In this problem, the question is whether a

boolean function correlates to a Fourier transform of another boolean function. This result implies that only a quantum computer could solve (or even just verify) Forrelation efficiently.

2.2 Boolean Formulas

Definition 2.25 (Boolean Formula). *Let $\boldsymbol{x} = (x_1, \ldots, x_n)$ be a vector of boolean variables (True/False, respectively 1/0). A **boolean formula** f consists of boolean variables and logical operations AND, OR and NOT. In particular, f maps a given assignment of variables to True or False (equivalently, 1 or 0). We say that a boolean formula is **satisfiable** if there exists an assignment of variables $\boldsymbol{x}$, such that $f(\boldsymbol{x}) = 1$. We will call x_i and its negation $\neg x_i$ **literals**. Literals that are all connected by a logical AND (OR) operation we will call an **AND-clause** (OR-clause).*

Definition 2.26 (Conjunctive Normal Form CNF). *A boolean formula f is in **conjunctive normal form (CNF)**, if it consists of OR-clauses that are connected by an AND operation. A CNF formula is satisfiable, if there is an assignment of variables such that all OR-clauses of the formula are satisfied.*

Remarks:

- Any boolean formula can be written as CNF.

Problem 2.27 (SAT). *In the boolean satisfiability problem (SAT), the task is to determine whether a given CNF formula is satisfiable.*

Problem 2.28 (3-SAT). *3-SAT is a special case of SAT, where the given formula is a CNF that has exactly three literals in each OR-clause.*

Theorem 2.29. *SAT is in NP.*

Proof. Assume we are given an arbitrary boolean formula and an assignment of its variables that acts as a solution. We need to verify in polynomial time whether the given assignment is True. Since the input (formula and assignment) has the size of the boolean formula, the verification can be done efficiently. □

Remarks:

- Intuitively, SAT seems difficult to solve. With n variables, there are 2^n possible assignments. If only one assignment satisfies the boolean formula, we are trying to find a needle in a haystack.

- The satisfiability problem with DNF (disjunctive normal form) formulas, on the other hand, is in P. In a DNF formula, the operations AND and OR are swapped. Therefore, only one AND-clause has to be satisfied in order to satisfy the whole formula. To do so, it is sufficient to verify that there is an AND-clause that contains no variable in its negated and non-negated form.

- Note that CNF $\leq$ DNF is not true. While every CNF can be expressed as DNF, the DNF may be exponentially bigger than the CNF, so the reduction function r in Definition 2.6 is not polynomial.

Theorem 2.30 (Cook/Levin Theorem). *SAT is NP-complete.*

Proof. The proof of this theorem is beyond the scope of this script and hence omitted here. The theorem is usually proven by using an alternative definition of NP-completeness via non-deterministic Turing machines. Using this definition one can show that it is possible to encode a Turing Machine as a SAT formula in polynomial time. □

Lemma 2.31. *3-SAT is NP-complete.*

Proof. Observe that 3-SAT $\leq$ SAT holds since 3-SAT is a special case of SAT, and therefore, according to Lemma 2.18, 3-SAT is in NP. The other direction also holds, that is SAT $\leq$ 3-SAT: In order to show this, OR-clauses that contain more than three literals need to be split into smaller clauses, while clauses with less literals need to be filled up. Here, we will only consider some small examples. A clause with two literals $(x_1 \vee x_2)$ can simply be replaced by $(x_1 \vee x_2 \vee y) \wedge (x_1 \vee x_2 \vee \neg y)$, where y is a new variable. A clause of four literals $(x_1 \vee x_2 \vee x_3 \vee x_4)$ we can rewrite as $(x_1 \vee x_2 \vee z) \wedge (x_3 \vee x_4 \vee \neg z)$, where z is a new variable. □

Remarks:

- Not all instances of SAT are hard. In 2-SAT, every clause contains exactly two literals. In contrast to 3-SAT, a solution of 2-SAT can be found in polynomial time.

- In order to show that some of the previously introduced problems (e.g. Clique or Vertex Cover) are also NP-complete, we will search for polynomial-time reductions from SAT (or 3-SAT) to these problems.

Theorem 2.32. *Clique is NP-hard.*

Proof. We will prove that Clique is NP-hard by using Lemma 2.21 and showing the reduction SAT $\leq$ Clique. We therefore will model a boolean formula as a graph. Given a boolean formula as a CNF, we construct a graph G as follows: Each OR-clause of the form $(\neg x_1 \vee x_2 \vee x_3)$ in the boolean formula forms a cluster of 3 nodes. Each literal from the clause corresponds to a node in the cluster. In the sample clause, we label the corresponding three nodes $\neg x_1$, x_2 and x_3. The nodes inside the same cluster will not be connected. We will also not connect nodes that represent the same variable in a negated and non-negated state, i.e. x and $\neg x$ will not be connected. Each other pair of nodes in G we connect by an edge. Note that is construction is polynomial in the size of the boolean formula.

Next, we show that a given boolean formula of k OR-clauses is satisfiable if and only if the corresponding graph G has a clique of size k. Let us first assume that a given boolean formula is satisfiable, that is, there is an assignment of $\boldsymbol{x}$, such that in each clause there is a literal that evaluates to True. Select one such variable from each of the k clauses. Note that in the graph G, all these variables will be connected. This is because the corresponding nodes are all in different clusters and because nodes with labels x and $\neg x$ cannot both be True in the boolean assignment. Therefore, these k nodes form a clique in G.

Now assume that there exists a clique of size k in G. Note that this clique cannot contain nodes from the same cluster, as such nodes would not be connected. By definition of G, the clique will also not contain nodes with labels x

and $\neg x$. Therefore, the clique of size k will have exactly one node from each cluster, and each variable will be represented in either its negated or its non-negated form. By setting the values of the node labels to True, we will receive an assignment of variables that satisfies the corresponding boolean formula. □

Remarks:

- Note that the opposite direction, Clique $\leq$ SAT, does not directly follow from above proof. Not all graphs G can be reached with the construction. In order to show Clique $\leq$ SAT, one would have to show that every graph G can be transformed into a boolean formula in polynomial time.

- In Lemma 2.16, we showed that Clique is in NP. Therefore, Clique is also NP-complete.

- Also Independent Set and Vertex Cover are NP-complete, since there are reductions in both directions between Clique and these problems, which we have shown in Theorem 2.10 and Theorem 2.12.

- In other words, all these problems are the most difficult problems in NP. If you could solve one of them in polynomial time, you could solve all problems in NP!

- What about Knapsack? We need yet another problem.

Problem 2.33 (Subset Sum). *Given a number $s \in \mathbb{N}$ and a set of n natural numbers $x_1, \ldots, x_n$, does there exist a subset of these numbers that has a sum of s.*

Theorem 2.34. *Subset Sum is NP-complete.*

Proof. It is easy to verify whether a given subset of elements has the desired sum s. Therefore, Subset Sum is in NP. In order to show that Subset Sup is NP-hard, the reduction 3-SAT $\leq$ Subset Sum is used. This reduction is non-trivial and will be omitted in this script. □

Theorem 2.35. *The Knapsack (decision) problem is NP-complete.*

Proof. Subset Sum is a special case of Knapsack: let the natural numbers be items with values $x_1, \ldots, x_n$ and weights $x_1, \ldots, x_n$, and let the knapsack have capacity s. If items of value s can be packed into the knapsack, then the corresponding items form a subset of sum s. Therefore, Subset Sum $\leq$ Knapsack. Since Subset Sum is NP-complete, Knapsack has to be NP-hard. In Lemma 2.15, we already showed that Knapsack is in NP. □

Remarks:

- In fact, thousands of interesting problems in many different areas are NP-complete: Achromatic Number, Battleship, Cut, Dominating Set, Equivalence Deletion, . . . , Super Mario Bros, Tetris, . . . , Zoe.

- Figure 2.36 below shows the reductions of NP-complete problems that we discuss in this chapter.

Figure 2.36: Reductions between NP-complete problems

2.3   Boolean Circuits

Definition 2.37 (Boolean Circuit). *Let $n, m \in \mathbb{N}$. A **boolean circuit** is a directed acyclic graph (DAG) with n boolean **input nodes** and m **output nodes**. The input nodes have no incoming edges, while the output nodes have no outgoing edges. The nodes of the graph that are not input nodes represent logical operations (AND, OR, NOT) and are called **gate nodes**. We will label the corresponding nodes $\wedge, \vee$ and $\neg$. Each NOT gate has in-degree 1. The in- and out-degrees of the other gates do not have to be bounded. We further define the **size** of the circuit to be its total number of gates, and the **depth** of the circuit to be the length of the longest (directed) path from an input to an output node.*

A boolean circuit C represents a mapping $C : \{0,1\}^n \mapsto \{0,1\}^m$. The value of an input node is the value of the corresponding input bit of C. The value of each other gate is determined recursively by applying the logic operation that this gate represents to the values received through incoming edges. The output of the boolean circuit are the values of all output nodes.

Remarks:

- Boolean circuits provide a nice tool for understanding the complexity of computation. They are limited in computation power, yet may help to answer the P $\neq$ NP conjecture.

- In the literature, it is often assumed that AND and OR gates have a bounded number of input values, or that the circuit has a bounded depth.

- In the following, we will consider boolean circuits that have only one output value (True or False).

Problem 2.38 (Circuit-SAT). *Let $C : \{0,1\}^n \mapsto \{0,1\}$ be a boolean circuit with only one output value. In **Circuit-SAT**, the task is to determine whether a given circuit C has an input vector $\boldsymbol{z} \in \{0,1\}^n$, such that $C(\boldsymbol{z}) = 1$.*

Theorem 2.39. *Circuit-SAT is in NP.*

Proof. We will start by showing that Circuit-SAT is in NP by applying Lemma 2.18 and showing that Circuit-SAT $\leq$ SAT. We therefore need to construct a boolean formula that represents the gates of the given circuit. First, we introduce a new variable g_i representing gate i for each gate of the circuit. We next differentiate between the three possible gates: If g_j is a NOT gate of gate g_i, then we add the clauses $(g_i \wedge \neg g_j) \vee (\neg g_i \wedge g_j)$ to the formula. If g_k is an AND gate with inputs from two gates, g_i and g_j, we add the following clauses to the formula: $(g_k \wedge g_i \wedge g_j) \vee (\neg g_k \wedge g_i \wedge \neg g_j) \vee (\neg g_k \wedge \neg g_i \wedge g_j) \vee (\neg g_k \wedge \neg g_i \wedge \neg g_j)$. In the last case, if g_k is an OR gate with inputs from g_i and g_j, we add $(g_k \wedge g_i \wedge g_j) \vee (g_k \wedge g_i \wedge \neg g_j) \vee (g_k \wedge \neg g_i \wedge g_j) \vee (\neg g_k \wedge \neg g_i \wedge \neg g_j)$ to the formula. Observe that an AND or an OR gate with k inputs can be replaced by $k - 1$ gates with two inputs by increasing the depth of the circuit.

Therefore, a boolean formula is satisfiable if and only if the corresponding circuit evaluates to 1. The construction was polynomial in the input size, and the resulting boolean formula can be transformed into a CNF, hence Circuit-SAT is in NP. $\square$

Theorem 2.40. *Circuit-SAT is NP-complete.*

Proof. In the previous theorem we have shown that Circuit-SAT is in NP. In order to show that Circuit-SAT is NP-hard, we will reduce SAT to Circuit-SAT, i.e. SAT $\leq$ Circuit-SAT. Let each variable of the boolean formula be an input to the circuit. For each negated variable, we can add a NOT gate. Then, for each clause in the formula, we add an AND gate that has all variables from the clause as input to the gate. And, finally, we add one OR gate that has values from all AND gates as its input. This gives a polynomial construction of a boolean formula as a circuit thus showing that Circuit-SAT is NP-complete. $\square$

Problem 2.41 (Minimum Circuit Size Problem or MCSP)**.** *Given a boolean function $f : \{0,1\}^n \mapsto \{0,1\}$, the task of the **Minimum Circuit Size Problem (MCSP)** is to determine whether f can be represented by a boolean circuit of size s, i.e. a circuit with s gates.*

Remarks:

- There are several approaches of how boolean formulas can be simplified. Often the laws from boolean algebra are applied (Distributive law, Idempotent law, Identity law, Complement law, DeMorgan's law, Karnaugh maps). These minimization rules however do not guarantee that the resulting boolean formula uses the minimum number of gates.

- How difficult is MCSP?

Theorem 2.42. *MCSP is in NP.*

Proof. We will first show that for every n, there exists a boolean function $f : \{0,1\}^n \mapsto \{0,1\}$ which cannot be computed by a circuit of size $2^n/(4n)$. Note that there exist 2^{2^n} possible functions f. Each circuit of size s with in-degree 2 can be represented by at most $3s \log s$ bits. Assume for contradiction that every boolean function can be represented by a circuit of size $s = 2^n/(4n)$. There exist at most $2^{3s \log s}$ such circuits. Observe that

$$2^{3s \log s} = 2^{3 \cdot 2^n/(4n) \log(3 \cdot 2^n/(4n))} \leq 2^{3n \cdot 2^n/(4n)} < 2^{2^n}$$

holds, meaning that not all functions can be computed by a circuit of size $s = 2^n/(4n)$.

The first part showed that the input to MCSP is at least of size $2^n/(4n)$. In order for MCSP to be in NP, we would have to efficiently check that the circuit has the desired size s and that it computes a given boolean function. Since the truth table of the given boolean function has size $(n+1)2^n$, it is possible to compute the truth table for the given circuit in a time that is polynomial in the input size. $\square$

Remarks:

- Is MCSP also NP-hard? As of now, there is no answer to this question. In fact, there is evidence that this result might be difficult to establish.

- Some complexity classes are defined with respect to boolean circuits.

Definition 2.43 (Class AC^0)**.** *The class AC^0 contains all decision problems that can be decided by a family of circuits $\{C_n\}$, where C_n has a constant depth and the size of the circuit is polynomial in n.*

Remarks:

- Note that we need to consider a circuit family, because the input size of each circuit is fixed, while the input to a problem can vary.

- Note that we do not allow the boolean circuits to save output bits or use recursive operations.

- There exist many other classes based on the boolean circuits: The class NC^0 is defined in the same way as AC^0, but the number of incoming edges to the AND and the OR gate is restricted to only two edges. Both classes can be extended to versions that allow a non-constant depth of the circuit: AC^i and NC^i have depth $\mathcal{O}(\log^i n)$.

2.4 Solving Hard Problems

We have seen that many important computational problems are known to be NP-hard. So what do we do if we encounter such a hard problem?

Remarks:

- We will now consider optimization problems instead of decision problems. This often makes sense in practice, where you usually want to find the best solution to some problem, instead of just deciding if the problem has a solution with a specific cost.

- So how can we approach a hard problem in practice? So far, we have mostly considered exact algorithms.

Definition 2.44 (Exact Algorithm)**.** *An **exact algorithm** always returns the optimal solution to the problem.*

Remarks:

- We have seen that for NP- complete problems we do not know if there is an exact algorithm with polynomial running time.

- Even when the worst-case running time of an algorithm is exponential, one can often use tricks to considerably reduce this running time in practice, e.g. by using look-ahead or pruning techniques as in Definition 1.9.

- What if we also accept weaker-than-optimal solutions, but instead, we require our algorithm to always have a reasonable running time?

Definition 2.45 (Heuristic). *A **heuristic** is an algorithm that is guaranteed to have a polynomial running time, at the cost of returning a suboptimal solution.*

Remarks:

- One example for a heuristic is our greedy algorithm for Knapsack, Algorithm 1.8.

- Good heuristics are usually based on insights into the structure or the behavior of the problem. Some heuristics are known to usually provide good solutions in practice.

- However, while heuristics are often good on some inputs, they may return very weak solutions on other inputs.

- In general, it is better to have algorithms which provide guarantees that they always return a solution which is reasonably close to the optimum. We want to understand algorithms with such guarantees.

2.5 Vertex Cover Approximation

Let us first revisit the optimization version of Vertex Cover: we want to find a vertex cover $S \subseteq V$ in an input graph $G = (V, E)$, with $|S|$ as small as possible. Can we come up with an algorithm where we can prove that the size of the returned vertex cover is always "reasonably close" to the optimum?

Remarks:

- A natural approach is the greedy method of Algorithm 2.46.

```
1  def VertexCover_Greedy_Naive(G):
2      S = ∅
3      while E ≠ ∅:
4          select an arbitrary edge (u, v) ∈ E
5          S = S ∪ {u}
6          remove all edges from E that are adjacent to u
7      return S
```

Algorithm 2.46: Naive greedy algorithm for Vertex Cover.

Lemma 2.47. *Algorithm 2.46 returns a vertex cover.*

Proof. An edge from E is only removed when one of its adjacent nodes is included in the set S. □

Remarks:

- However, the size of this vertex cover can be very far from the optimal size.

Theorem 2.48. *The size $|S|$ of the vertex cover returned by Algorithm 2.46 can be an $n - 1$ factor larger than the optimum.*

Proof. Consider a "star" graph with n nodes and $n-1$ edges, where $n-1$ distinct leaf nodes u_1, u_2, ..., u_{n-1} are connected by an edge to the same center node v. Whenever Algorithm 2.46 selects an edge (u_i, v) in this graph, it can happen that it always chooses the leaf node u_i. In this case, the algorithm only removes a single edge (u_i, v) in each step, and the algorithm lasts for $n - 1$ iterations of the loop. In the end, the algorithm returns the vertex cover $S = \{u_1, ..., u_{n-1}\}$, which consists of $n - 1$ nodes.

On the other hand, the set $\{v\}$ is an optimal vertex cover of cost $c^* = 1$ in this graph, so we have $\frac{|S|}{c^*} = n - 1$. □

Remarks:

- It is a natural idea to try to improve Algorithm 2.46 with a clever tie-breaking rule: for example, to always select the higher-degree adjacent node of the chosen edge, or the highest-degree node altogether among the available nodes. This solves the star. However, there are more complicated counterexamples which show that even in this case, the solution we obtain can be a $\log n$ factor worse than the optimum.

- However, there is a slightly different greedy approach that provides much better guarantees.

```
1  def VertexCover_Greedy(G):
2      S = ∅
3      while E ≠ ∅:
4          select an arbitrary edge (u, v) ∈ E
5          S = S ∪ {u, v}
6          remove all edges from E that are adjacent to u or v
7      return S
```

Algorithm 2.49: Greedy algorithm for Vertex Cover.

Theorem 2.50. *Algorithm 2.49 always returns a vertex cover with $|S| \leq 2 \cdot c^*$.*

Proof. Algorithm 2.49 returns a correct vertex cover: an edge is only removed from E if at least one of its adjacent nodes is inserted into S.

Assume that the algorithm runs for c iterations of the main loop, i.e. it selects c different edges (u_i, v_i), and inserts both adjacent nodes u_i and v_i of these edges into S for $i \in \{1, ..., c\}$. Note that none of these edges (u_i, v_i) share a node, because all edges adjacent to either u_i or v_i are removed when (u_i, v_i) is selected (in graph theory, they refer to such a set of edges as a *matching*). This means that in the optimal vertex cover S^*, each vertex can only be adjacent to at most one of the edges (u_i, v_i); thus in order to cover all the edges (u_i, v_i), we already need at least c nodes. As a result, we get that $c^* = |S^*| \geq c$.

On the other hand, our algorithm returns a vertex cover of size $|S| = 2 \cdot c$, so we have $|S| = 2 \cdot c \leq 2 \cdot c^*$. □

Remarks:

- In a graph of $\frac{n}{2}$ independent edges, the solution returned by Algorithm 2.49 is indeed 2 times worse than the optimum, so our analysis is tight.

- This algorithm was somewhat counter-intuitive; one might expect Algorithm 2.46 to be more efficient. However, as we have seen, Algorithm 2.49 is always at most a factor 2 away from the optimum c^*, while Algorithm 2.46 can be an $(n-1)$-factor away.

- Algorithms that are proven to always be within an α factor from the optimum are called *approximation algorithms*. Theorem 2.50 shows that Algorithm 2.49 is a 2-approximation algorithm for Vertex Cover.

- For our formal definition of this notion, we assume a minimization problem.

Definition 2.51 (Approximation Algorithm). *We say that an algorithm $\mathcal{A}$ is an α-**approximation algorithm** if for every possible input of the problem, it returns a solution with a cost of at most $\alpha \cdot c^*$, where c^* is the cost of the optimal solution (the one with minimal cost).*

Remarks:

- The definition is similar for maximization problems. In this case, we denote the value of the optimal solution (the one with highest value) by v^*. For an approximation algorithm, we require that the returned solution always has a value of at least v^*/α.

- This section only considers approximations algorithms that run in polynomial time. That is exactly the point of these algorithms: to find a reasonably good solution without having an unreasonably high running time.

Definition 2.52. *We say that a problem is α-approximable if there exists a polynomial-time algorithm $\mathcal{A}$ that is an α-approximation algorithm for the problem.*

Remarks:

- In general, α can be any constant value with $\alpha > 1$, or it can also be a function of n (the size of the input), e.g. $\alpha = \log n$ or $\alpha = n^{3/4}$.

- In many cases, we can already get some approximability results from the most trivial algorithms. Independent Set naturally has $v^* \leq n$, since we can find a solution of value 1 by simply outputting a set that consists of an arbitrary single node. This shows that Independent Set is (at least) n-approximable.

- This allows us to classify hard problems, based on how well their optimum solution can be approximated. E.g. the complexity class of problems that are α-approximable for some constant α is called APX. Since Theorem 2.50 has $\alpha = 2$, Vertex Cover is in APX.

- For Vertex Cover, there is currently no $2 - \varepsilon$ approximation algorithm known for any $\varepsilon > 0$, so the algorithm above is indeed the best we have. Also, it is proven that no approximation better than 1.3606 is possible at all unless $P = NP$. This means that unless $P = NP$, we are unable to approximate the best solution of this problem arbitrarily well (in polynomial time).

2.6 Bin Packing Approximation

Let us now look at some other problems that are α-approximable to some constant α. We continue with the bin packing problem, which is similar to Knapsack.

Problem 2.53 (Bin Packing). *We have a set of n items of size $x_1, ..., x_n$, and* $\rightarrow$ notebook *an unlimited number of available bins, each having a capacity of B. A set of items fits into a bin if its total size is at most B. Our goal is to put all of the items into bins, using the smallest possible number of bins.*

Remarks:

- You can easily imagine immediate applications of this, e.g. packing files on disks.

- In case of this problem, c^* denotes the number of bins that are used in the optimal solution.

- In Bin Packing, a simple greedy heuristic already allows us to obtain a 2-approximation.

```
def FirstFit(items, B):
    for each item in items:
        place item in the first bin where it still fits
        if item does not fit into any bin:
            open a new bin, and insert item into the new bin
```

Algorithm 2.54: First Fit algorithm for Bin Packing.

Lemma 2.55. *In Algorithm 2.54, at most 1 bin is filled at most half.*

Proof. Assume that there are at least 2 bins b_1 and b_2 that are filled at most half. This means that b_1 still has free space of at least $\frac{B}{2}$, and b_2 only contains items of size $x_i \leq \frac{B}{2}$. However, in this case, the items sorted into b_2 would also fit into b_1. This contradicts the First Fit algorithm, which only opens a new bin when the next item does not fit into any of the previous bins. $\square$

Theorem 2.56. *Algorithm 2.54 is a 2-approximation.*

Proof. Assume that the First Fit algorithm uses m bins. Due to Lemma 2.55, at least $m-1$ of these bins are filled to a capacity of more than $\frac{B}{2}$. This implies

$$\sum_{i=1}^{n} x_i > (m-1) \cdot \frac{B}{2}.$$

On the other hand, we know that

$$c^* \geq \frac{\sum_{i=1}^{n} x_i}{B},$$

since all the items have to be sorted into a bin, and every bin can only take at most B. The two inequalities imply

$$B \cdot c^* > (m-1) \cdot \frac{B}{2},$$

and thus

$$2 \cdot c^* > m - 1.$$

Since c^* and m are integers, this implies $2 \cdot c^* \geq m$, proving our claim. $\square$

Remarks:

- One can even improve on this algorithm by first sorting the elements in decreasing order, and then applying the same First Fit rule. With only a slightly more detailed analysis, one can show that this improved algorithm achieves an approximation ratio of 1.5.

- Can we get a better approximation ratio than 1.5? Unfortunately not, unless we have $\mathsf{P} = \mathsf{NP}$. One can show that getting a better than 1.5 approximation is already an NP-hard problem.

- To prove this, we present a reduction to Partition.

Problem 2.57 (Partition). *In the Partition Problem, given a set of n items of positive size $x_1, ..., x_n$, we want to decide if we can partition them into two groups such that the total sum is equal in the two groups.*

Remarks:

- Partition is a special case of Subset Sum where $s = \frac{1}{2} \cdot \sum_{i=1}^{n} x_i$.

- It is known that this special case is still NP-complete.

Theorem 2.58. *Subset Sum $\leq$ Partition.*

Proof. The reduction is not too complex, but a bit boring, so we do not discuss it here. $\square$

Theorem 2.59. *It is NP-hard to solve Bin Packing with an α-approximation ratio for any constant $\alpha < \frac{3}{2}$.*

Proof. Given an input of Partition with integers $x_1, ..., x_n$, we convert it into a Bin Packing problem: we consider items of size $x_1, ..., x_n$, and we define $B = \frac{1}{2} \cdot \sum_{i=1}^{n} x_i$.

If the original Partition problem is solvable, then this Bin Packing problem can also be solved with 2 bins. In this case, an α-approximation algorithm with $\alpha < \frac{3}{2}$ must always return a solution with $m \leq c^* \cdot \alpha < 2 \cdot \frac{3}{2} = 3$ bins; since m is an integer, this means a solution with $m = 2$ bins.

On the other hand, if the items cannot be partitioned into two sets of size B, then the algorithm can only return a solution with $m \geq 3$ bins.

Hence an α-approximation algorithm for Bin Packing would also solve Partition in polynomial time: we can just solve the converted Bin Packing problem, and if $m = 2$, then output 'Yes', whereas if $m \geq 3$, then output 'No'. $\square$

Remarks:

- Thus, Bin Packing is a problem that can be approximated to some constant $\alpha = 1.5$, but not arbitrarily well, i.e. not for any $\alpha > 1$.

2.7 TSP Approximation

Let us now consider another natural graph question.

Problem 2.60 (Hamiltonian Cycle). *The input is a graph $G = (V, E)$. In the **Hamiltonian cycle problem** we want to know whether there exists a cycle in G which contains each node of the graph exactly once.*

Remarks:

- Hamiltonian cycle was among the first problems that were shown to be NP-complete.

- In practice, cities often have different distances between each other.

Problem 2.61 (Traveling Salesperson or TSP). *The input is a clique graph* → notebook *$G = (V, E)$ (there is an edge between any two nodes of G), with positive edge weights, i.e. a function $d : E \to \mathbb{R}^+$. The goal of **TSP** is to find a Hamiltonian cycle in G where the total weight of the edges contained in the cycle is minimal.*

Theorem 2.62. *In general graphs, TSP is NP-hard to approximate to any constant factor α.*

Proof. We provide a reduction to Hamiltonian Cycle, which is NP-complete.

Assume that we have an α-approximation algorithm for TSP. Given an input graph $G = (V, E)$ for Hamiltonian Cycle, we turn this into an instance of TSP on V. For each pair of nodes u and v, we define the weight of edge (u, v) in the TSP in the following way:

- if $(u, v) \in E$, then we assign $d(u, v) = 1$ in the TSP,

- if $(u, v) \notin E$, then we assign $d(u, v) = \alpha \cdot n + 1$ in the TSP.

If the initial graph had a Hamiltonian cycle, then the resulting TSP has a cycle of total weight n as the optimal solution. In this case, our approximation algorithm is guaranteed to return a solution of cost at most $\alpha \cdot n$. On the other hand, if there was no Hamiltonian cycle in the original graph, then the optimal TSP cycle must contain at least one edge of weight $\alpha \cdot n + 1$, and thus $c^* \geq \alpha \cdot n + 1$.

Hence an α-approximation returns a solution of cost at most $\alpha \cdot n$ if and only if the original graph had a Hamiltonian cycle; thus running such an approximation algorithm allows us to solve the Hamiltonian cycle problem. This shows that finding a polynomial-time α-approximation of TSP for *any* constant α is NP-hard. □

Remarks:

- This means that in case of general TSP, the best we can hope for is an approximation ratio that depends on n.

- On the other hand, there is a special case of the problem that does allow a constant-factor approximation.

Problem 2.63 (Metric TSP)**.** *In Metric TSP, the distances between any three nodes v_1, v_2, v_3 must satisfy the triangle-inequality: $d(v_1, v_2) \leq d(v_1, v_3) + d(v_3, v_2)$.*

Remarks:

- Note that the triangle inequality is a very realistic assumption in many real-world applications.

- This is another possible approach to solving hard problems in practice: maybe we can show that our actual problems are restricted to only a special case of the original problem, and that this special case is computationally more tractable.

- A Metric TSP already has a 2-approximation algorithm.

```
1  def Tree_Based_TSP(G)
2      find a Minimum Spanning Tree T in G
3      form a sequence of nodes P_0 by traversing all nodes of T
4      simplify P_0 to P by only keeping the first occurrence of
       ↪  each node
5      return P
```

Algorithm 2.64: Tree-based approximation for TSP.

Remarks:

- The Minimum Spanning Tree (MST) is, intuitively speaking, the subset of edges with smallest total weight which already connects the whole graph, i.e. there is a path between any two nodes through these edges. An MST in a graph can easily be found in polynomial time with some simple algorithms.

Theorem 2.65. *Algorithm 2.64 is a 2-approximation.*

Proof. See also Figure 2.66. Let t^* denote the total cost of the MST. Note that if we delete a single edge from any Hamiltonian cycle, we get a spanning tree, so t^* must be smaller than the cost of any Hamiltonian cycle. This implies $t^* \leq c^*$.

Now consider the cost of the Hamiltonian cycle returned by Algorithm 2.64. Since the traversal visits each edge of the spanning tree exactly twice, the total cost of P_0 is $2 \cdot t^*$. By only keeping the first occurrence of each node in P_0, we create "shortcuts": whenever we delete a node from P_0, we only shorten our tour due to the triangle inequality. As such, for the total costs we have

$$cost(P) \leq cost(P_0) = 2 \cdot t^* \leq 2 \cdot c^*.$$

□

Figure 2.66: A weighted complete graph as an input to TSP (a), an MST in this graph (b), the tour P_0 obtained from this MST, starting from an arbitrary node (c), and the simplified tour P after removing repeated occurrences (d).

Remarks:

- Being slightly more clever even allows a 1.5-approximation.

- For TSP, it is often reasonable to assume the even more special case of Euclidean TSP, where the nodes of the graph corresponds to specific points in a plane, for example. That is, each node v is associated with two coordinates v_x and v_y, and the weight of the edge between u and v is the Euclidean distance of u and v, i.e. $\sqrt{(u_x - v_x)^2 + (u_y - v_y)^2}$.

2.8 FPTAS: Knapsack

We have now seen many problems that are α-approximable for a specific constant α. Sometimes we can even get arbitrarily close to the optimum in polynomial time: there exists a $(1 + \varepsilon)$-approximation for any $\varepsilon > 0$.

Definition 2.67 (FPTAS). *We say an algorithm $\mathcal{A}$ is a fully polynomial-time approximation scheme (FPTAS) if for any $\varepsilon > 0$*

- *$\mathcal{A}$ is a $(1 + \varepsilon)$-approximation algorithm, and*

- *the running time of $\mathcal{A}$ is polynomial in both n and $\frac{1}{\varepsilon}$.*

Remarks:

- While an FPTAS is not as good as a polynomial-time exact algorithm, it is almost as good: for any desired error ε, we can efficiently find a solution that is only ε away from the optimum.

- There is also a slightly weaker notion of PTAS (polynomial-time approximation scheme), where we only require the running time to be polynomial in n, but not in $\frac{1}{\varepsilon}$. That is, the running time is still polynomial in n for any fixed $\varepsilon > 0$, but it might increase very quickly in $\frac{1}{\varepsilon}$; for example, it includes a factor of $n^{2^{1/\varepsilon}}$. These algorithms are not as useful in practice: if we want to get very close to the optimum by selecting a small ε, the running time still becomes unreasonably large.

- As an example for a nicely approximable problem, we revisit the Knapsack problem, and present an FPTAS for it.

- Since we now study the problem from more of a mathematical than a programming-based perspective, we introduce a shorter notation for the inputs of the problem.

Definition 2.68 (Knapsack notation). *We will use C to denote the capacity of our Knapsack, and we denote the value and weight of the i^{th} item as v_i and w_i, respectively. We denote the number of items by n as before. Finally, let us use $v_{\max}$ to denote the value of the highest value item, i.e. $v_{\max} := \max v_i$.*

Remarks:

- In case of this problem, we use v^* to denote the maximal value we can fit into the knapsack.

- Note that if any item has $w_i > C$, then it can never fit into the knapsack, so we might as well remove such items. We will assume that $v_{\max}$ denotes the maximal value of an item that still fits into the knapsack.

- Recall that we have discussed a dynamic programming solution for Knapsack, where each cell $V[i][c]$ of our DP table stored the maximum value that can be achieved with capacity c using only the first i items. The starting point of our FPTAS algorithm will be a slightly different variant of this DP method: in our new table, $W[i][v]$ will store the weight of the lowest-weight subset of items $1, ..., i$ that has a total value of v.

- One can show that this table W can also be computed with a similar dynamic programming method to Algorithm 1.12. As before, the table has n rows. The number of columns is now at most $n \cdot v_{\max}$, which is

an upper bound on the value of any subset; we have n items, and each of them has value at most $v_{\max}$. From this alternative DP table, we can find the optimum v^* by taking the maximal v value in the table where $W[n][v] \leq C$.

- The main difference from the original DP algorithm is that now each column expresses the *value* of a specific subset of items, instead of the *weight* of a specific subset of items as before.

- We can now use this alternative DP algorithm to develop an FPTAS for Knapsack. The main idea is to slightly round up the values of the items, which results in some inaccuracy for the algorithm, but it also reduces the number of columns in our DP table.

```
1  def Knapsack_FPTAS(items, C):
2      k = ε · v_max / n
3      consider the modified Knapsack problem where item i has
       ↪  weight w_i and value v̂_i = ⌈v_i/k⌉ · k
4      run the alternative DP algorithm on this problem, with
       ↪  columns only corresponding to multiples of k
```

Algorithm 2.69: FPTAS algorithm for Knapsack.

Lemma 2.70. *Algorithm 2.69 has a running time of $\mathcal{O}\left(n^3 \cdot \frac{1}{\varepsilon}\right)$.*

Proof. Given the parameter ε, the algorithm introduces a scaling parameter $k = \varepsilon \cdot v_{\max} / n$, and defines the *rounded-up value* of item i as $\hat{v}_i = \lceil \frac{v_i}{k} \rceil \cdot k$. This means that whichever subset of nodes we select, the total value of the subset is a multiple of k.

On the other hand, any subset of the items has a total value of $n \cdot v_{\max}$ at most. If each column represents a multiple of k, then the table has at most $n \cdot v_{\max} / k = n^2 \cdot \frac{1}{\varepsilon}$ columns. The table still has n rows, so the total running time of the DP algorithm with these rounded values is $O(n^3 \cdot \frac{1}{\varepsilon})$. $\qquad\square$

Theorem 2.71. *Algorithm 2.69 is an FPTAS.*

Proof. The polynomial running time has already been established in Lemma 2.70. Let us use S to denote the set of items chosen by Algorithm 2.69. Since we round the values to multiples of k, the value of each item has been overestimated by k at most, so $\hat{v}_i - k \leq v_i$. Hence for the total value of items in S we have

$$\sum_{i \in S} v_i \geq \sum_{i \in S} (\hat{v}_i - k) = \sum_{i \in S} \hat{v}_i - |S| \cdot k = \sum_{i \in S} \hat{v}_i - |S| \cdot \frac{\varepsilon \cdot v_{\max}}{n} \geq \sum_{i \in S} \hat{v}_i - \varepsilon \cdot v_{\max}.$$

To simplify this last expression even more, we make two more observations. First, since the value of each item is only rounded upwards in this modified problem, and our algorithm finds the optimal solution, we have $\sum_{i \in S} \hat{v}_i \geq v^*$. Also, choosing the item with maximal value $v_{\max}$ is always a valid solution, so we also have $v_{\max} \leq v^*$. Hence for the total value in S we have

$$\sum_{i \in S} \hat{v}_i - \varepsilon \cdot v_{\max} \geq v^* - \varepsilon \cdot v^* = (1 - \varepsilon) \cdot v^*.$$

Note that this shows $v^* / v(S) \leq \frac{1}{1-\varepsilon}$, while our original FPTAS definition requires a slightly different relation: that $v^* / v(S) \leq (1 + \varepsilon)$. However, since we now have an $\alpha = \frac{1}{1-\hat{\varepsilon}}$-approximation algorithm for any $\hat{\varepsilon} > 0$, we can obtain $\alpha = (1 + \varepsilon)$ by choosing $\hat{\varepsilon} = \frac{\varepsilon}{1+\varepsilon}$. □

2.9 Non-Approximability: Independent Set

The most difficult problems are those that cannot be approximated to any constant α, but only to a factor $\alpha = f(n)$ depending on n, e.g. $\alpha = \log n$ or $\alpha = \sqrt{n}$. This means that as the size of the problem becomes larger, the difference between the optimum and our solution also grows larger.

Remarks:

- Remember the Independent Set problem? It turns out that this is one of the most difficult problems from this perspective: approximating it to any factor that is slightly smaller than n would already imply that $\mathsf{P} = \mathsf{NP}$.

Theorem 2.72. *For any positive constant $\delta > 0$, it is NP-hard to approximate the Independent Set problem to a ratio of $n^{1-\delta}$.*

Proof. The proof of this theorem is more involved than our previous claims, so we do not discuss it here. The idea of the proof is built on the notion of so-called probabilistically checkable proof systems. □

Remarks:

- This means that for any approximation algorithm $\mathcal{A}$, there are some graphs where $\mathcal{A}$ returns very weak solutions.

- This $n^{1-\delta}$ factor is indeed huge. For example, it might be that the largest independent set contains $n^{0.99}$ nodes, but $\mathcal{A}$ returns an independent set of only 1 or 2 nodes. As such, an approximation algorithm like this is not useful in practice.

- The same result also holds for the Clique problem.

- Recall from Theorem 2.12 that as a decision problem, Independent Set $\leq$ Vertex Cover; however, we have seen that Vertex Cover is 2-approximable. This shows that even if the decision version of two problems are equivalent, the approximability of the optimization version may differ significantly.

- So while all NP-complete decision problems have the same difficulty in theory, their optimization variants can behave very differently in practice.

Chapter Notes

The first formal definition of computation was provided by Alan Turing in 1936 [20]. In this paper, Turing presents an automatic machine (now known as

Turing machine) that can compute certain classes of numbers. We will discuss the Turing machine in more detail in the last chapter. The birth of computational complexity as a field is attributed to Hartmanis and Stearns for their paper "On the computational complexity of algorithms" [14]. They proposed to measure time with respect to the input size and showed that some problems can only be solved if more time is given.

The first problem that was shown to be in P was the maximum matching problem, see Edmonds [9]. In another paper, Edmonds introduced a notion that is equivalent to the class NP [10]. In 1971, Cook [6] showed that SAT was NP-complete. Two years later, independently of Cook's result, Levin [17] showed that six problems were NP-complete, the so-called universal search problems. Both authors have formulated the famous P versus NP problem in computer science. The name NP was however given to the class a year later by Karp [16], who proved that 21 (combinatorial) problems were NP-complete. He thereby first used the technique of polynomial reductions in these proofs. Since then, thousands of problems in different areas of science have been shown to be NP-hard. Also many new complexity classes have been proposed in the literature. Different "zookeepers" around the world keep track of the newly introduced complexity classes in their complexity zoos [1]. The question $\mathsf{P} \neq \mathsf{NP}$ still remains open. In 2000, the P versus NP problem was announced as one of the seven Millennium Prize Problems by the Clay Mathematics Institute [4]. Solving P versus NP (or BQP versus NP) will be a major milestone in science, in case of equality ($\mathsf{BQP} = \mathsf{NP}$) there will be amazing practical consequences.

Circuit theory is almost 200 years old, being studied by electrical engineers and physicists. In the 1940s, the invention of semiconductor devices and later the transistor further boosted the quest for understanding circuits. The computational model of boolean circuits was introduced by Claude Shannon [19] in 1949, who showed that most boolean functions require circuits of exponential size (see Theorem 2.42). It was a natural step to restrict the circuits in size and depth and consider problems that are still computable on restricted circuits. Furst, Saxe and Sipser [11] for example showed that Parity is not in AC^0, but in NC^1. The first connection between circuits and Turing machines has been made by Savage [18]. It is generally believed that proving bounds using boolean circuits is easier than by using Turing machines.

While all NP-hard problems have a superpolynomial runtime in the worst case, there are numerous possible tricks and optimizations that can make these algorithms viable on special cases of the problem, even for very large inputs. For example, there are yearly SAT-competitions where SAT-solver algorithms try to decide the satisfiability of SAT formulas from real-life applications, with many of these formulas containing millions of variables and tens of millions of clauses [2].

Heuristic solutions to hard problems have also been extensively studied. In particular, there is a wide range of so-called metaheuristics, which are general approaches and techniques for developing a heuristic solution to a problem. This includes some simpler approaches like local search or gradient descent, and also some more sophisticated ones like simulated annealing, tabu search or genetic algorithms [13]. Note that many techniques in machine learning are also developed for this purpose: to provide heuristic solutions to hard problems.

The 2-approximation algorithm for Vertex Cover has long been known, discovered by both Gavril and Yannakakis independently [12]. The FPTAS algo-

rithm for Knapsack has also been around for a long time, and it is one of the most popular FPTAS examples [21].

Bin Packing has also been studied for multiple decades [12]. A recent analysis of the First Fit heuristic is available by Dósa and Sgall, proving the even stronger result that First Fit is a 1.7-approximation [8]. The variant which first sorts the items in decreasing order is known to even provide a $\frac{11}{9}$-approximation up to an additive constant, and this bound is known to be tight [7].

For Travelling Salesperson, a slightly improved version of our approximation algorithm is due to Christofides, which improves the approximation ratio to only 1.5 [5]. In the special case of an Euclidean TSP, there even exists a (rather complicated) PTAS algorithm, which received a Gödel prize [3].

The inapproximability result for Independent Set was the final result of a line of lower bounds in the early 2000s; John Hastad also received a Gödel prize for this result [15, 22].

This chapter was written in collaboration with Darya Melnyk and Pál András Papp.

Bibliography

[1] Complexity zoo. `https://complexityzoo.uwaterloo.ca/Complexity_Zoo`.

[2] The international SAT competition. `http://www.satcompetition.org`.

[3] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, September 1998.

[4] James A Carlson, Arthur Jaffe, and Andrew Wiles. *The millennium prize problems*. American Mathematical Society Providence, RI, 2006.

[5] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.

[6] Stephen A Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158, 1971.

[7] György Dósa. The tight bound of first fit decreasing bin-packing algorithm is ffd(i) ¡ 11/9 opt(i)+6/9. In *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies*, pages 1–11, Berlin, Heidelberg, 2007.

[8] György Dósa and Jiri Sgall. First Fit bin packing: A tight analysis. In *30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, volume 20 of *LIPIcs*, pages 538–549, Dagstuhl, Germany, 2013.

[9] Jack Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B*, 69(125-130):55–56, 1965.

[10] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.

[11] Merrick Furst, James B Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical systems theory*, 17(1):13–27, 1984.

[12] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.

[13] Fred W Glover and Gary A Kochenberger. *Handbook of metaheuristics*, volume 57. Springer Science & Business Media, 2006.

[14] Juris Hartmanis and Richard E Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965.

[15] Johan Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 627–636. IEEE, 1996.

[16] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

[17] Leonid Anatolevich Levin. Universal sequential search problems. *Problemy peredachi informatsii*, 9(3):115–116, 1973.

[18] John E Savage. Computational work and time on finite machines. *Journal of the ACM (JACM)*, 19(4):660–674, 1972.

[19] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[20] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.

[21] Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.

[22] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 681–690, New York, NY, USA, 2006.

Chapter 3

Cryptography

In Chapter 2 we learned that some functions are really hard to compute. This might seem like terrible news. On the flip side, it enables modern cryptography!

3.1 Encryption

We start with the oldest problem in cryptography: How can we send a secret message?

Definition 3.1 (Perfect Security). *An encryption algorithm has perfect security, if the encrypted message reveals no information about the plaintext message to an attacker, except for the possible maximum length of the message.*

Remarks:

- If an encryption algorithm offers perfect security, any plaintext message of the same length could have generated the given ciphertext.

- Sometimes perfect security is also called information-theoretic security.

- Is there an algorithm that offers perfect security?

```
1   # m = plaintext message Alice wants to send to Bob
2   # k = random key known by Alice and Bob, with len(k) = len(m)
3   # c = ciphertext, the encrypted message m
4
5   def encrypt_otp_Alice(m, k)
6       Alice sends c = m ⊕ k to Bob   # ⊕ = XOR
7
8   def decrypt_otp_Bob(c, k)
9       Bob computes m' = c ⊕ k
```

Algorithm 3.2: One Time Pad

Remarks:

- In cryptography, it's always Alice and Bob, with a possible attacker Eve.

Theorem 3.3. *Algorithm 3.2 is correct.*

Proof. $m' = c \oplus k = (m \oplus k) \oplus k = m$. $\qquad\square$

Theorem 3.4. *Algorithm 3.2 has perfect security.*

Proof. Given a ciphertext c, for every plaintext message m there exists a unique key k that decrypts c to m, that is $m = c \oplus k$. Therefore, if k is uniformly random, every plaintext is equally likely and thus, ciphertext c reveals no information about plaintext m. $\qquad\square$

Remarks:

- Algorithm 3.2 only works if the message m has the same length as the key k. How can we encrypt a message of arbitrary length with a key of fixed length?

- Block ciphers process messages of arbitrary length by breaking them into fixed-size blocks and operating on each block.

```
1   # m, k, c as defined earlier, now with len(k) << len(m)
2
3   def encrypt_ECB(m, k)
4       Split m into r len(k)-sized blocks m_1, m_2, ..., m_r
5       for i = 1, 2, 3, ..., r:
6           c_i = m_i ⊕ k
7       c = c_1; c_2; ...; c_r   # ; stands for concatenation
8       return c
```

Algorithm 3.5: Electronic Code Book

Remarks:

- In Algorithm 3.5, blocks of the same plaintext result in the same ciphertext, because the same key k is reused to encrypt every block. Furthermore, reusing the same key reveals information about m_1 and m_2: Suppose you have two messages m_1, m_2 encrypted with the same key k, resulting in c_1, c_2. We now have $c_1 \oplus c_2 = (m_1 \oplus k) \oplus (m_2 \oplus k) = m_1 \oplus m_2$. So, reusing the same key k in Algorithm 3.2 is insecure.

- But there are better block based encryptions. AES (Advanced Encryption Standard) is the current state of the art.

- For encryption, Alice and Bob need to agree on a key k first! While this may be feasible for, e.g., secret agents, it is quite impractical for everyday usage.

3.2 Key Exchange

How to agree on a common secret key in public, if you never met before?

Definition 3.6 (Primitive Root). *Let $p \in \mathbb{N}$ be a prime. Then $g \in \mathbb{N}$ is a primitive root of p if the following holds: For every $y \in \mathbb{N}$, with $1 \leq y < p$, there is an $x \in \mathbb{N}$ such that $g^x = y \mod p$.*

→ notebook

```
1  # p = publicly known large prime number
2  # g = publicly known primitive root of p
3
4  def Diffie_Hellman_Alice():
5      Pick a random secret key a ∈ {1, 2, ..., p − 1}
6      Send k_a = g^a  mod p to Bob
7      Receive k_b from Bob
8      Calculate k = (k_b)^a  mod p
9
10 def Diffie_Hellman_Bob():
11     # same as Alice, swapping all a, b.
```

→ notebook

Algorithm 3.7: Diffie-Hellman Key Exchange

Theorem 3.8. *In Algorithm 3.7, Alice and Bob agree on the same key k.*

Proof. Everything mod p, we have

$$k = (k_b)^a = \left(g^b\right)^a = g^{b \cdot a} = g^{a \cdot b} = \left(g^a\right)^b = (k_a)^b = k.$$

$\square$

Remarks:

- Algorithm 3.7 does not have perfect security, but instead only computational security.

Definition 3.9 (Computational Security). *An algorithm has computational security, if it is secure against any adversary with polynomial computational resources.*

Remarks:

- The definition of security differs from one cryptographic primitive to another (e.g., encryption, signatures, etc.).

- The computational security of Algorithm 3.7 is based on the difficulty of the discrete logarithm.

Problem 3.10 (Discrete Logarithm or DL). *Given a prime $p \in \mathbb{N}$, a primitive root g of p, and $y \in \mathbb{N}$ with $1 \leq y < p$, find an $x \in \mathbb{N}$ such that $g^x = y \mod p$.*

Problem 3.11 (Decisional Diffie-Hellman or DDH). *Given a prime $p \in \mathbb{N}$, a primitive root g of p, and $g^a, g^b, g^c \in \mathbb{N}$ with $1 \leq g^a, g^b, g^c < p$, decide if $c = a \cdot b$.*

Lemma 3.12. *DDH $\leq$ DL.*

Proof. We just compute the discrete logarithms of g^a, g^b, g^c to decide if $c = a \cdot b$

$\square$

Remarks:

- The discrete logarithm assumption states it is infeasible to solve DL given computationally bounded resources.

- We have no proof that DL is hard, but there is no known efficient algorithm.

- Conversely, modular exponentiation can be done in polynomial time using repeated squaring.

  → notebook

- The decisional Diffie-Hellman assumption states it is difficult to solve DDH.

Lemma 3.13. *Algorithm 3.7 is secure against passive adversaries under the DDH assumption.*

Proof. There is a polynomial-time reduction from breaking passive security of Diffie-Hellman key exchange to violating the DDH assumption. Intuitively, the additional information obtained for the key $k = g^{a \cdot b}$ by an eavesdropper in Algorithm 3.7, can be used to distinguish whether $c = a \cdot b$. We omit the proof as it is advanced, but similar security proofs will be introduced in Section 3.8. $\square$

Remarks:

- Passive security means that any eavesdropping adversary cannot extract from the ciphertext significantly more information for the encrypted message than someone with access only to the message length.

- What about stronger adversaries?

Definition 3.14 (Man in the Middle Attack). *A man in the middle attack is defined as an adversary Eve deciphering or changing the messages between Alice and Bob, while Alice and Bob believe they are communicating directly with each other.*

Theorem 3.15. *The Diffie-Hellman Key Exchange from Algorithm 3.7 is vulnerable to a man in the middle attack.*

Proof. Assume that Eve can intercept and relay all messages between Alice and Bob. That alone does not make it a man in the middle attack, Eve needs to be able to decipher or change messages without Alice or Bob noticing. Indeed, Eve can emulate Alice's and Bob's behavior to each other, by picking her own a', b', and then agreeing on common keys $g^{a \cdot b'}$, $g^{b \cdot a'}$ with Alice and Bob, respectively. Thus, Eve can relay all messages between Alice and Bob while deciphering and (possibly) changing them, while Alice and Bob believe they are securely communicating with each other. $\square$

Remarks:

- It is a bit like concurrently playing chess with two grandmasters: If you play white and black respectively, you can essentially let them play against each other by relaying their moves.

- How do we fix this? One idea is to personally meet in private first, exchange a common secret key, and then use this key for secure communication. However, having a key already completely defeats the purpose of a key exchange algorithm.

- Can we do better? Yes, with public key cryptography.

3.3 Public Key Cryptography

Definition 3.16 (Public Key Cryptography). *A public key cryptography system uses two keys per participant: A public key k_p, to be disseminated to everyone, and a secret (private) key k_s, only known to the owner. A message encrypted with the public key of the intended receiver can be decrypted only with the corresponding secret key. Also, messages can be digitally signed; a message verifiable with a public key must have been signed with the corresponding secret key.*

Remarks:

- Popular public key cryptosystems include RSA, Elliptic Curve Cryptography, etc.

- We study a public key cryptosystem based on the discrete logarithm problem.

- In Diffie-Hellman Key Exchange algorithm (Algorithm 3.7), Alice picked a secret number a and computed a public number $k_a = g^a \mod p$ which Alice sent to Bob. We use the exact same idea in Algorithm 3.17 to generate a pair of public and secret keys (k_p, k_s).

```
1   # p,g as defined earlier
2
3   def generate_key():
4       Pick a random secret key k_s ∈ {1, 2, ..., p − 1}
5       k_p = g^{k_s}  mod p
6       return k_p, k_s
```

Algorithm 3.17: Key Generation

3.4 Digital Signatures

Definition 3.18 (Digital Signature Scheme). *A digital signature scheme is a triple of algorithms:*

- *A key generation algorithm that outputs a public/secret key pair k_p, k_s.*

- *A signing algorithm that outputs a digital signature σ on message m using a secret key k_s.*

- *A verification algorithm that outputs* True *if the signature σ on the message m is valid using the public key k_p of the signer, and* False *otherwise.*

A digital signature scheme should be correct, and unforgeable.

Definition 3.19 (Correctness). *A signature scheme is correct if the verification algorithm on input σ, m, k_p returns* True *only if σ is the output of the signing algorithm on input m, k_s.*

Definition 3.20 (Unforgeability). *A signature scheme is unforgeable if no adversary can produce a valid message-signature pair without receiving it from external sources.*

Remarks:

- All algorithms (key generation, signing, and verification) should be efficient, i.e., computable in polynomial time.

- Digital signatures offer *authentication* (the receiver can verify the origin of the message), *integrity* (the receiver can verify the message has not been modified since it was signed), and *non-repudiation* (the sender cannot falsely claim that they have not signed the message).

- Widely known signature schemes are ElGamal, Schnorr, and RSA.

```
1    # p, g, m as defined earlier
2    # h = cryptographic hash function like SHA256
3    # k_p, k_s = Alice's public/secret key pair
4    # s, r = the signature sent by Alice
5
6    def sign_Alice(m, k_s):
7        Pick a random x ∈ {1, 2, ..., p − 1}
8        r = g^x  mod p
9        s = x · h(m) − k_s · r  mod p − 1
10       return s, r   # signature σ = (s, r)
11
12   def verify_Bob(m, s, r, k_p):
13       return r^{h(m)} == k_p^r · g^s  mod p
```

Algorithm 3.21: ElGamal Digital Signatures

Theorem 3.22. *The ElGamal digital signature scheme (Algorithms 3.17, 3.21) is correct.*

Proof. The algorithm is correct, meaning that a signature generated by (an honest) Alice will always be accepted by Bob. That is because,

$$k_p^r \cdot g^s = g^{r \cdot k_s} \cdot g^{x \cdot h(m) - k_s \cdot r} = g^{r \cdot k_s + x \cdot h(m) - k_s \cdot r} = g^{x \cdot h(m)} = r^{h(m)} \mod p$$

.   □

Remarks:

- The random variable x in Line 7 is often called a *nonce* – a number only used once.

- Writing "$\texttt{mod } p-1$" in Line 9 is not a typo. In the exponent, we always compute modulo $p-1$, since that will make sure that values larger than $p-1$ will be truncated because of Fermat's Little Theorem (see below, Theorem 3.23).

- The function $h()$ in Line 9 is a so-called cryptographic hash function.

Theorem 3.23 (Fermat's Little Theorem). *Let p be a prime number. Then, for any $x \in \mathbb{N}$: $x^p = x \mod p$. If x is not divisible by p, then $x^{p-1} = 1 \mod p$.*

Definition 3.24 (Cryptographic Hash Function). *A cryptographic hash function is a function that maps data of arbitrary size to a bit array of a fixed size (the **hash value** or **hash**). A cryptographic hash function is easy to compute but hard to invert.*

Remarks:

- A hash function is deterministic: the same message always results in the same hash value.

- Being hard to invert can be formalized:

Definition 3.25 (Collision Resistance). *Weak collision resistance: Given an input x, it is difficult to find a different input x' such that $h(x) = h(x')$. Strong collision resistance: It is difficult to find two different values x and x' such that $h(x) = h(x')$.*

Definition 3.26 (One-Way). *A function that is practically infeasible to invert is called one-way.*

Remarks:

- SHA2, SHA3 (Secure Hash Algorithm 2/3), RIPMED, and BLAKE are some example families of cryptographic hash functions. SHA256 is a specific implementation of the SHA2 construction which outputs a 256 bit output for arbitrary sized inputs. Earlier constructions like MD5 or SHA1 are considered broken/weak now.

Lemma 3.27. *One-way $\leq$ strong collision resistance.*

Proof. Suppose a hash function h that is not one-way. Now, let us choose some random value x and compute $h(x)$. Since we can invert $h(x)$, it is highly likely we get another value x', such that $x' \neq x$ (collision). That is because a hash function maps data of arbitrary size to fixed-size values, hence there are many collisions (but they are typically hard to find). Thus, h is not strong collision resistant.   □

Lemma 3.28. *Existence of one-way functions $\Rightarrow \mathsf{P} \neq \mathsf{NP}$.*

Proof. Suppose there is a function f that is one-way. We define a decision problem (defined in Chapter 2) as follows. Given an input $(\bar{x}, y)$, decide whether there is an x with $\bar{x}$ a prefix of x, such that $f(x) = y$. This decision problem is not in P, otherwise for a given y one can compute in polynomial time an x such that $f(x) = y$ and thus f would not be one-way. This would be done by extending the prefix (starting from the empty string) character by character until we find x. This decision problem is in NP, since given $(\bar{x}, y)$ as input and x as a solution, one can check in polynomial time that $f(x) = y$ and $\bar{x}$ is a prefix of x. This means, that we have a decision problem that is in NP but not in P and thus $\mathsf{P} \neq \mathsf{NP}$.   □

Remarks:

- The existence of one-way functions is still an open problem.

Theorem 3.29. *ElGamal signatures are unforgeable under the discrete logarithm assumption.*

Proof. To forge a signature, a malicious Bob must either find a collision in the hash function, $h(m) = h(m') \mod p-1$, or extract Alice's secret key k_s. Therefore, Bob must either break the collision resistance property of the cryptographic hash function or solve DL. Both problems are assumed to be hard.

Note that Alice must chose an x uniformly at random for *each* signature, and make sure no information on x is leaked. Otherwise, security can be compromised. In particular, if Alice uses the same nonce x and secret key k_s to sign two different messages, Bob can compute k_s.   □

Remarks:

- Why do we need the cryptographic hash function in Algorithm 3.21?

Theorem 3.30. *ElGamal signatures without cryptographic hash functions are vulnerable to existential forgery.*

Proof. Let s, r be a valid signature on message m. Then, $(s', r') = (sr, r^2)$ is a valid signature on message $m' = rm/2$ (as long a s either m or r is even), because

$$k_p^{r'} \cdot g^{s'} = \left(g^{k_s}\right)^{r^2} \cdot g^{s \cdot r} = g^{r^2 \cdot k_s} \cdot g^{r \cdot (x \cdot m - r \cdot k_s)} = g^{r^2 \cdot k_s + r \cdot x \cdot m - r^2 \cdot k_s} =$$
$$g^{r \cdot x \cdot m} = (g^x)^{r \cdot m} = \left(r^2\right)^{m/2} = (r')^{m'} \mod p.$$

□

Remarks:

- Existential forgery is the creation of at least one message-signature pair (m, s), when m was never signed by Alice.

- Craig Wright used Satoshi Nakamoto's key in Bitcoin and signed a random message attempting to impersonate the famous creator of Bitcoin. However, when Wright was asked to sign "I am Satoshi" he could not deliver!

- Similarly to digital signatures, message authentication codes are used to ensure a message received by Bob is indeed sent by Alice. However, MACs are symmetric, i.e., they are generated and verified using the same secret key.

Definition 3.31 (Message Authentication Code or MAC). *A message authentication code is a bitstring that accompanies a message. It can be used to verify the authenticity of the ciphertext in combination with a secret authentication key k_a (different from k) shared by the two parties.*

Remarks:

- Eve should not be able to change the encrypted message and/or the MAC, and get Bob to believe that Alice sent the encrypted message.

- Algorithm 3.32 shows a hash based MAC construction.

```
1   # m, k_p, c as defined earlier
2   # k_a = key to authenticate c
3
4   def encrypt_then_MAC(m, k_p, k_a):
5       c = encrypt(m, k_p)
6       a = h(k_a; c)
7       return c, a
```

Algorithm 3.32: Hash Based Message Authentication Code

Remarks:

- Bob accepts a message c only if he calculates $h(k_a; c) = a$.

- With some hash functions (e.g., SHA2), it is easy to append data to the message without knowing the key and obtain another valid MAC. To avoid these attacks, in practice we use $h(k_a; h(k_a; c))$.

- Now Alice and Bob can securely communicate over the insecure communication channels of the internet, due to the known public keys.

- But how does Bob know that Alice's public key really belongs to Alice? What if it is really Eve's key? Quoting Peter Steiner: "*On the Internet, nobody knows you're a dog.*"

3.5 Public Key Infrastructure

"*Love all, trust a few.*" – William Shakespeare

What can we do, unless we personally meet with everyone to exchange our public keys? The answer is trusting a few, in order to trust many.

Remarks:

- Let's say that you don't know Alice, but both Alice and you know Doris. If you trust Doris, then Doris can verify Alice's public key for you. In the future, you can ask Alice to vouch for her friends as well, etc.

- Trust is not limited to real persons though, especially since Alice and Doris are represented by their keys. How do you know that you give your credit card information to a shopping website, and not some infamous Nigerian princess Eve? You probably don't know the owner of the shopping website personally.

Definition 3.33 (Public Key Infrastructure or PKI). *Public Key Infrastructure (PKI) binds public keys with respective identities of entities, like people and organizations. People and companies can register themselves with a certificate authority.*

Definition 3.34 (Certificate Authority or CA). *A certificate authority is an entity whose public key is stored in your hardware device, operating system, or browser by the respective vendor like Apple, Google, Microsoft, Mozilla, Ubuntu, etc.*

Remarks:

- A certificate is an assertion that a known real world person, with a physical postal address, a URL, etc. is represented by a given public key, and has access to the corresponding secret key.

- You can accept a public key if a certificate to that effect is signed by a CA whose public key is stored in your device.

- CA's whose public keys are stored in your device are also called root CA's. Sometimes, there are intermediate CA's whose certificates are signed by root CA's, and who can sign many other end-user certificates. This enables scaling, but also introduces vulnerabilities.

- If a CA's secret key is compromised by a malicious actor, they can sign themselves a certificate saying that they are someone else (say, Google), and then impersonate Google to innocent browsers which trust this CA. A CA's key can be revoked if this happens, or CA's keys can have shorter expiry times.

- Another problem is that your own set of root certificates might be compromised, e.g., if malicious software replaces your browser's root certificates with fakes.

3.6 Transport Layer Security

To communicate securely over the internet, we simply combine the cryptographic primitives we learned so far!

Remarks:

- Alice and Bob don't want Eve to be able to read their messages. Therefore, they encrypt their messages using block based encryption (see Section 3.1).

- For the encryption algorithm, they need to agree on a secret key using a key exchange protocol (see Section 3.2).

- When Alice receives a message, how can she be sure that the message hasn't been modified on the way from Bob to her? Alice and Bob use message authentication (see Section 3.4) to ensure integrity of the communication.

- Let's assume that Alice hasn't met Bob in person before. How can she be sure that she is really communicating with Bob and not with Eve? She would ask Bob to authenticate himself (see Sections 3.4, 3.5).

Protocol 3.35 (Transport Layer Security, TLS)**.** *TLS is a network protocol in which a client and a server exchange information in order to communicate in a secure way. Common features include a bulk encryption algorithm, a key exchange protocol, a message authentication algorithm, and lastly, the authentication of the server to the client.*

Remarks:

- TLS is the successor of Secure Sockets Layer (SSL).

- HTTPS (Hypertext Transfer Protocol Secure) is not a protocol on its own, but rather denotes the usage of HTTP via TLS or SSL.

3.7 Public Key Encryption

Public key or asymmetric encryption schemes allow users to send encrypted messages directly.

Definition 3.36. *A public key encryption scheme is a triple of algorithms:*

- *A key generation algorithm that outputs a public/secret key pair k_p, k_s.*

- *An encryption algorithm that outputs the encryption c of a message m using the receiver's public key k_p.*

- *A decryption algorithm that outputs the message m using the secret key k_s.*

Remarks:

- The key generation algorithm for ElGamal encryption scheme is the same as in ElGamal signatures (Algorithm 3.17).

```
1   # p, g, m, k_p, k_s  as defined earlier
2
3   def encrypt(m, k_p):
4       Pick a random nonce x ∈ {1, 2, ..., p − 1}
5       c_1 = g^x  mod p
6       c_2 = m · k_p^x  mod p
7       return c_1, c_2   # encryption c = (c_1, c_2)
8
9   def decrypt(c_1, c_2, k_s):
10      m' = c_2 · c_1^{k_s·(p−2)}  mod p
11      return m'
```

Algorithm 3.37: ElGamal Encryption Algorithm

Theorem 3.38. *ElGamal encryption scheme (Algorithms 3.17, 3.37) is correct.*

Proof. Alice can recover the message: $m' = c_2 \cdot c_1^{k_s \cdot (p-2)} = \left(m \cdot k_p^x\right) \cdot g^{x \cdot k_s \cdot (p-2)} = m \cdot \left(k_p^x\right)^{p-1} = m$. The last step uses Theorem 3.23. $\square$

3.8 Security of PK Encryption

For each scheme studied so far, we have seen that they are correct. But are they secure? How do we know that there is no simple attack on a public-key encryption scheme?

Remarks:

- In cryptography, there are two popular ways to prove a scheme is secure: simulation-based security and game-based security. In this section we will focus on game-based security. In Section 3.10 we will discuss simulation-based security.

- There are various game-based security models for asymmetric public-key encryption schemes, most prominently the IND-CPA model.

Definition 3.39 (Indistinguishability under Chosen Plaintext Attack or IND-CPA)**.** *Consider the following game between an adversary (trying to break the security) and a challenger (challenging the adversary), where the adversary is a probabilistic polynomially-bounded algorithm.*

1. *The challenger generates a key pair (k_s, k_p) based on a security parameter n (e.g., a key size in bits), and publishes k_p to the adversary, while k_s remains secret.*

2. *The adversary performs $O(\mathsf{poly}(n))$ computations (encryptions or other operations).*

3. *The adversary sends two distinct plaintexts of his choice m_0, m_1 ($m_0 \neq m_1$) to the challenger.*

4. *The challenger selects a bit $b \in \{0, 1\}$ uniformly at random and sends the challenge ciphertext $c = \mathsf{encrypt}(k_p, m_b)$ to the adversary.*

5. *The adversary performs any number of additional computations.*

6. *Finally, the adversary outputs a guess for the value of b.*

An encryption scheme is IND-CPA secure if any probabilistic polynomial time adversary has only a negligible "advantage" over random guessing.

Remarks:

- A negligible "advantage" wins above game with probability $\frac{1}{2} + \varepsilon$.

- Intuitively, an encryption scheme is IND-CPA secure if the adversary does not learn any additional information on how to decrypt a message.

- IND-CPA security is equivalent to semantic security, where an adversary that sees the ciphertext has no advantage against an adversary that does not see the ciphertext.

- Does IND-CPA security hold for deterministic encryption schemes?

Theorem 3.40. *All deterministic public-key encryption schemes fail IND-CPA security.*

Proof. Consider any deterministic encryption scheme. The adversary can always win the IND-CPA game: First, the adversary picks two messages m_0, m_1 and encrypts them to c_0 and c_1 respectively using the public key k_p. Then, when the adversary receives the ciphertext c from the challenger, she simply compares c to c_0 and c_1 and always outputs the correct bit! $\qquad \square$

Remarks:

- For instance, the traditional RSA encryption scheme fails IND-CPA security as it is deterministic. To prove RSA security, RSA-OAEP is used.

- Therefore, only randomized encryption schemes are of interest, like ElGamal.

Theorem 3.41. *ElGamal is secure in the IND-CPA model under the DDH assumption.*

Proof. We will show that if ElGamal is not IND-CPA secure, then the DDH assumption does not hold (DDH $\leq$ IND-CPA ElGamal). In particular, given an adversary A that gains a non-negligible advantage ε over random guessing in the IND-CPA game, we can construct an adversary (efficient algorithm) B

that given $g^a, g^b, g^c \in \mathbb{N}$ with $1 \leq g^a, g^b, g^c < p$, can decide if $c = a \cdot b$ (DDH) with probability $1/2 + \varepsilon$.

To that end, algorithm B feeds (g, p, g^a) as input to algorithm A (g^a is the public key for A). When the adversary sends the two messages m_0, m_1 to the challenger, algorithm B (posing as the challenger) chooses b randomly from $\{0, 1\}$ and returns to A the ciphertext $c = (g^b, g^c \cdot m_b)$. Algorithm A then returns a value b'. If $b = b'$, algorithm B returns $\mathtt{True}$, else $\mathtt{False}$.

In case g^a, g^b, g^c is a DDH tuple, i.e., $a \cdot b = c$, we expect algorithm A to produce the correct output with a good probability, i.e., at least with advantage ε. Therefore, adversary B will guess correctly with probability at least $1/2 + \varepsilon$. On the other hand, if g^a, g^b, g^c is a random tuple, then the ciphertext will have an unpredictable structure and thus adversary A will return the correct value with probability $1/2$ (random guess). Therefore, adversary B can distinguish between the DDH tuple and the random tuple with at least a non-negligible advantage ε, which violates the DDH assumption. $\qquad \square$

Remarks:

- What about stronger notions of game-based security? There are two widely used models, Indistinguishability under Chosen Ciphertext Attack (IND-CCA) and Indistinguishability under Adaptive Chosen Ciphertext Attack (IND-CCA2).

- Both IND-CCA and IND-CCA2 are defined similarly to IND-CPA, but in addition the adversary is given access to a decryption oracle which decrypts arbitrary ciphertexts at the adversary's request.

- In IND-CCA the adversary can query the decryption oracle only until he receives the challenge. On the contrary, in IND-CCA2, the adversary can also query the decryption oracle after he has received the challenge – but he cannot query the challenge for decryption.

Definition 3.42 (IND-CCA). *Same as IND-CPA (Definintion 3.39) except step 2 which is replaced by "The adversary performs $O(\mathsf{poly}(n))$ computations or queries to the decryption oracle".*

Definition 3.43 (IND-CCA2). *Same as IND-CPA (Definintion 3.42) except step 5 which is replaced by "The adversary performs any number of additional computations or queries to the decryption oracle".*

Theorem 3.44. *ElGamal is not secure in the IND-CCA2 model.*

Proof. Consider the following scenario: The adversary sends the challenger the two messages m_0, m_1 and receives the ciphertext c which corresponds to the encryption of one of the messages m_b. Then, the adversary picks a random message m and queries the decryption oracle with the ciphertext $c' = c \cdot \mathsf{encrypt}(k_p, m)$. The decryption oracle cannot refuse to respond to the query since the input is different to the challenge, and returns $m_b \cdot m$. Next, the adversary simply removes m and recovers the challenge message. $\qquad \square$

Remarks:

- IND-CCA security of ElGamal is an open problem.

- The reason ElGamal is not IND-CCA2 secure is that it is multiplicatively homomorphic.

- Any homomorphic encryption scheme is not IND-CCA2 secure.

Definition 3.45 (Homomorphic Encryption Schemes). *An encryption scheme is said to be homomorphic under an operation $*$ if $E(m_1 * m_2) = E(m_1) * E(m_2)$.*

Remarks:

- In other words, we can dicrectly compute with encrypted data!

- $m_1 * m_2$ and $E(m_1) * E(m_2)$ indicates that ciphertexts and messages can both be operated upon using the same operation. This depends on the representation of ciphertexts, and is not always precisely defined. In the case of ElGamal encryption's homomorphism, we use pairwise vector multiplication to multiply ciphertexts: $E(m_1) \cdot E(m_2) = (c_{11}, c_{12}) \cdot (c_{21}, c_{22})^T$.

Lemma 3.46. *ElGamal encryption scheme (Algorithms 3.17, 3.37) is homomorphic under modular multiplication.*

Proof. We refer the encryption of message m with public key k_p, large prime p, generator g, and a random nonce x as $E(m) = (c_1, c_2) = (g^x, m \cdot k_p^x)$

$$E(m_1) \cdot E(m_2) = (g^{x_1}, m_1 \cdot k_p^{x_1}) \cdot (g^{x_2}, m_2 \cdot k_p^{x_2})^T$$
$$= (g^{x_1 + x_2}, (m_1 \cdot m_2)k_p^{x_1 + x_2}) = E(m_1 \cdot m_2)$$

$\square$

Remarks:

- Not every public encryption scheme is homomorphic under all operations. If an encryption scheme is homomorphic only under some operations, it's called a *partial homomorphic encryption scheme*. For example, we have:

  - Modular multiplication: ElGamal cryptosystem, RSA cryptosystem.

  - Modular addition: Benaloh cryptosystem, Pallier cryptosystem.

  - XOR operations: Goldwasser–Micali cryptosystem.

- There are fully homomorphic encryption schemes that support all possible functions, like Craig Gentry's lattice-based cryptosystem.

- Homomorphic encryption is used in electronic voting schemes to sum up encrypted votes.

- A close cryptographic notion to homomorphism is malleability.

Definition 3.47 (Malleability). *An encryption scheme is malleable if, given the encryption c of a message m, one can generate another ciphertext c' which decrypts to a related message $m' = f(m)$, where f is a known function, without knowing or learning m.*

Lemma 3.48. *ElGamal encryption scheme is malleable.*

Proof. An attacker can change (c_1, c_2) to (c_1, c_2') where $c_2' = z \cdot c_2 \mod p$. With everything mod p

$$m' = c_2' \cdot s'^{p-2} = (z \cdot m \cdot s) \cdot s^{p-2} = z \cdot m \cdot s^{p-1} = z \cdot m$$

Resulting in a valid encryption of $z \cdot m$. $\square$

Remarks:

- Malleability and indistinguishability are closely related. For instance IND-CCA2 is equivalent to non-malleability under the same attack.

- What other problems can we solve using crypto? The answer is surprisingly many! In the next sections we will discuss some of the most exciting cryptographic primitives beyond TLS.

3.9 Commitment Schemes

Suppose Alice and Bob want to play XOR over the telephone. Both choose a random bit. If the chosen bits are the same, Alice wins; if they are different, Bob wins. They do not trust each other, therefore they need a scheme that ensures an unbiased result that is verifiable by both parties. For this purpose they can use so-called commitment schemes.

Remarks:

- In a commitment scheme, Alice locks (commits) her bit in a box and gives the box to Bob. Bob reveals his bit, and only then Alice provides Bob with the key to open the box to reveal her bit.

- Commitment schemes are the digital analog of a safe.

Definition 3.49 (Commitment Scheme). *A commitment scheme is a two-phase interactive protocol between Alice, the sender, and Bob, the receiver. A commitment scheme must be correct, binding and hiding.*

- ***Commit phase:*** *Alice commits to a message m by producing a public commitment c and a secret decommitment d. Alice sends c to Bob.*

- ***Reveal phase:*** *Alice sends m and d to Bob. Bob verifies that the message m corresponds to the commitment c.*

Definition 3.50 (Hiding). *A commitment scheme is hiding if Bob cannot extract any information about the committed message before the reveal phase.*

Definition 3.51 (Binding). *A commitment scheme is binding if Alice cannot change her commitment.*

Definition 3.52 (Correctness). *If both Alice and Bob follow the protocol, then Bob always returns* `True` *in the reveal phase.*

Remarks:

- Is there a simple way to create a commitment scheme? How about using hash functions?

```
1  # m,h as defined earlier
2  # n = security parameter
3
4  def commit_Alice(m):
5      Pick a random n-bit string r
6      Send c = h(r; m) to Bob
7
8  def reveal_Bob(c, m, r):  # Bob receives m, r from Alice
9      return c == h(r; m)
```

Algorithm 3.53: A Simple Commitment

Theorem 3.54. *Any cryptographic hash function can produce a (computationally binding and computationally hiding) commitment scheme.*

Proof. Algorithm 3.53 demonstrates how to extract a commitment scheme from a cryptographic hash function. The probability that some random r', m' exist such that $h(r'; m') = h(r; m)$ is 2^{-n}. To find such an r', m' is infeasible for a computationally bounded Alice, hence the commitment scheme is computationally binding. The scheme is also computationally hiding, but we omit the proof because it is complex. □

Remarks:

- What if Alice or/and Bob are computationally unbounded?

```
1  # p, g, m, n as defined earlier
2  # y = a random value in {1, 2, ..., p − 1}
3  # x with y = g^x mod p is unknown
4
5  def commit_Alice(m):
6      Pick a random r ∈ {1, 2, ..., p − 1}
7      c = g^m · y^r mod p
8      Send c to Bob
9
10 def reveal_Bob(m, c, r):  # Bob receives m, c, r from Alice
11     return c == g^m · y^r mod p
```

Algorithm 3.55: Pedersen Commitment

Theorem 3.56. *Pedersen commitments are correct.*

Proof. Given m, c, r, y, Bob can verify $c = g^m \cdot y^r \mod p$. Thus, the Pedersen commitment scheme is correct. □

Theorem 3.57. *Pedersen commitments are perfectly hiding.*

Proof. Given a commitment c, every message m is equally likely to be the committed message to c. That is because given m, r and any m', there exists r' such that $g^m \cdot y^r = g^{m'} \cdot y^{r'} \mod p$, specifically $m' + x \cdot r' = m + x \cdot r \mod p - 1$, where $y = g^x \mod p$. □

Theorem 3.58. *Pedersen commitments are computationally binding.*

Proof. Towards contradiction, suppose Pedersen commitments are not computationally binding. This means that a "polynomial" Alice can successfully compute two pairs m, r and m', r' such that (everything modulo p)

$$g^m \cdot y^r = g^{m'} \cdot y^{r'}$$
$$\Leftrightarrow g^{m'(p-2)} \cdot y^{r(p-2)} \cdot g^m \cdot g^r = g^{m'(p-2)} \cdot y^{r(p-2)} \cdot g^{m'} \cdot y^{r'}$$
$$\Leftrightarrow g^{m+m'(p-2)} \cdot y^{r+r(p-2)} = g^{m'+m'(p-2)} \cdot y^{r'+r(p-2)}$$
$$\Leftrightarrow g^{m+m'(p-2)} \cdot (y^r)^{(p-1)} = (g^{m'})^{(p-1)} \cdot y^{r'+r(p-2)}$$
$$\Leftrightarrow g^{m+m'(p-2)} = y^{r'+r(p-2)}$$

Then, Alice can compute (using the extended Euclidean Algorithm) the multiplicative inverse of $r' + r(p - 2)$, i.e., Alice can find a z such that $z(r' + r(p - 2)) = 1 \mod p - 1$. Consequently, Alice can compute the discrete logarithm $\log_g y = x = z(m + m'(p-2))$ since $g^{z(m+m'(p-2))} = y \mod p$. But this contradicts the discrete logarithm assumption, therefore such an adversary does not exist. □

Remarks:

- If Alice sends both c, r as a commitment to Bob, Pedersen commitments can be perfectly binding and computationally hiding.

- But why compromise at all? Ideally, we want a perfectly hiding and perfectly binding commitment scheme.

Theorem 3.59. *A commitment scheme can either be perfectly binding or perfectly hiding but not both.*

Proof. Given a commitment c, a computationally unbounded adversary can simply generate the commitments for every m until finding one that outputs c. In a perfectly binding scheme c uniquely identifies m. Hence, the scheme is not perfectly hiding. □

Remarks:

- Commitment schemes have important applications in several cryptographic protocols, such as zero-knowledge proofs, and multiparty computation.

3.10 Zero-Knowledge Proofs

Peggy and Victor play Where's Waldo. Can Peggy prove she found Waldo without revealing Waldo's location to Victor?

Remarks:

- In the physical world, Peggy can cover the picture with a large piece of cardboard that has a small, Waldo-shaped hole in its center. She can then place the cardboard such that only Waldo is visible through the hole and therefore prove to Victor she has found Waldo without revealing any information regarding Waldo's location.

- In Zero-Knowledge Proofs (ZKP), the *prover*, Peggy, wants to convince the *verifier*, Victor, of the knowledge of a secret without revealing any information about the secret to Victor.

Definition 3.60 (Zero-Knowledge Proof). *A pair of probabilistic polynomial time interactive programs P, V is a zero-knowledge proof if the the following properties are satisfied:*

- **Completeness:** *If the statement is true, then an honest verifier V will be convinced by an honest prover P.*

- **Soundness:** *If the statement is false, a cheating prover P cannot convince the honest verifier V that it is true, except with negligible probability.*

- **Zero-knowledge:** *If the statement is true, no verifier V learns anything beyond the statement being true.*

Remarks:

- Soundness concerns the security of the verifier, and zero-knowledge the security of the prover.

- Examples of ZKPs are the Hamiltonian cycle (Problem 2.60) for a large graph, or the Schnorr protocol (also known as $\Sigma-$protocol).

```
1   # n = security parameter
2   # G = large graph
3
4   def ZKP_HamiltonianCycle(G):
5       repeat n times:
6           Peggy creates graph H, isomorphic to G
7           Peggy commits to H   # using a commitment scheme
8           Victor tosses a coin c = [heads, tails]
9           if c == heads:
10              Victor asks Peggy for the G → H mapping
11              Peggy returns the mapping and the decommitment of H
12              Victor verifies the mapping and the commitment
```

```
13          else:
14              Victor asks Peggy for the Hamiltonian cycle
15              Peggy returns the decommitment of the cycle only
16              Victor verifies the cycle and the commitment
```

Algorithm 3.61: Hamiltonian Cycle ZKP

Remarks:

- A graph H is isomorphic to G if H is just like G except that all the nodes (and edges) have random names. Think of G as a "symmetric-looking" graph, e.g., every node in G has the same number of neighbors.

- If c is tails, the commitment should allow Victor to verify the cycle. This can be done by, for example, committing to every edge (or lack thereof) separately. Then, Victor only decommits the edges that form the cycle.

Theorem 3.62. *Algorithm 3.61 is complete.*

Proof. If Peggy knows a Hamiltonian cycle in G, she can easily satisfy Victor's demand in both cases. In case c is heads, Peggy returns the isomorphism (the renaming of G's nodes in H) and reveals the H she committed to. In case c is tails, Peggy can easily construct and return a Hamiltonian cycle in H by applying the isomorphism to the cycle in G. □

Theorem 3.63. *Algorithm 3.61 is sound.*

Proof. If Peggy does not know the information, she can guess which question Victor will ask and generate either a graph H isomorphic to G (in which she does not know a Hamiltonian cycle) or a Hamiltonian cycle for an unrelated graph H'. If she does not know a Hamiltonian cycle for G she cannot do both. Therefore, Peggy's chance of fooling Victor is 2^{-n}, where n is the number of rounds the protocol is repeated. □

Theorem 3.64. *Algorithm 3.61 is zero-knowledge, if the commitment is perfectly-hiding.*

Proof. The main idea is that Peggy's answers do not reveal the original Hamiltonian cycle in G. In each round, Victor only learns H's isomorphism to G or that there exists a Hamiltonian cycle in the commited graph. Thus, the information remains unknown since Victor would need both answers for a single round to discover the cycle in G. But how can we be sure that that Victor learns absolutely nothing?

To prove zero-knowledge, we can show that the whole transcript that Victor sees can be generated by Victor himself without the help of Peggy. For each round, if c is heads, Victor can generate a commitment of a isomorphism of G and then simply reveal it. If c is tails, Victor can generate a commitment of a fully connected graph and reveal a Hamiltonian cycle. The transcript generated by Victor is indistinguishable from the transcript that is produced by interacting

with Peggy (provided that commitments are perfectly-hiding). In other words, the transcript that was produced by interacting with Peggy helps Victor just as much as the transcript generated by himself. Since the same transcript can be generated without even knowing the Hamiltonian cycle of G, the real transcript seen by Victor cannot contain any information on the Hamiltonian cycle in G. □

Remarks:

- For all realistic purposes, it is infeasible to break the soundness of a zero-knowledge proof with a reasonable number of rounds.

- One classic three-round protocol that exhibits the properties of a zero-knowledge proof is the Schnorr protocol.

- In Schnorr's protocol, Peggy wants to prove to Victor she knows x, the discrete logarithm of a publicly known value $y = g^x \mod p$ without revealing any information about x.

```
1   # g, p, x, y as defined earlier
2
3   def ZKP_Schnorr():
4       Peggy picks a random r ∈ {1, 2, ..., p − 1}
5       Peggy sends t = g^r mod p to Victor
6       Victor picks a random challenge c ∈ {1, 2, ..., p − 1}
7       Victor sends c to Peggy
8       Peggy sends to Victor u = r + x · c mod p − 1
9       Victor verifies g^u == t · y^c mod p
```

Algorithm 3.65: Schnorr ZKP

Theorem 3.66. *Algorithm 3.65 is complete.*

Proof. If both Peggy and Victor are honest, we have $g^u = g^{r+x \cdot c} = g^r \cdot (g^x)^c = t \cdot y^c \mod p$. □

Theorem 3.67. *Algorithm 3.65 is zero-knowledge.*

Proof. Victor can generate an indistinguishable communication transcript from the actual communication transcript, without Peggy's help, as follows: Given g, p, y, Victor randomly picks $c', u' \in \{1, 2, \ldots, p - 1\}$, and outputs $(t', c', u') = (g^{u'} \cdot y^{c' \cdot (p-2)}, c, u)$. Note that $t' \cdot y^{c'} = g^{u'} \cdot y^{c' \cdot (p-2)} \cdot y^{c'} = g^{u'} \cdot (y^{c'})^{p-1} = g^{u'}$, so the verification holds. The communication transcript between Peggy and Victor is $(t, c, u) = (g^r, c, r + x \cdot c)$, where $r, c \in \{1, 2, \ldots, p - 1\}$ are randomly picked. Thus, the two transcripts are indistinguishable. By the same reason as in the previous example, the transcript seen by Victor cannot contain any information about the discrete logarithm of g^x. □

Remarks:

- Proving soundness for Algorithm 3.65 is complicated and omitted.

- The analysis of Algorithm 3.65 is very subtle, when the challenge space is very big (not polynomial). In particular, the protocol is zero-knowledge only when the verifier is honest. This and many other subtleties pose the main challenge for the mathematically precise definition of zero-knowledge.

- Can we design non-interactive zero-knowledge proofs?

- The Fiat-Shamir heuristic is a technique for converting an interactive proof of knowledge into a non-interactive proof of knowledge by replacing the challenge with the outcome of a known cryptographic hash function. For instance, the non-interactive version of Schnorr ZKP protocol is the following: Peggy generates $t = g^x \mod p$ and uses a cryptographic hash function h to calculate the challenge $h(t)$. Peggy, then, calculates the proof of knowledge $u = r + x \cdot h(t)$ and makes t, u public. As a result Peggy can calculate the challenge on her own while anyone can verify the proof.

- Fiat-Shamir enables the creation of digital signature schemes from an interactive zero-knowledge proof.

3.11 Threshold Secret Sharing

How does a company share its vault passcode among its board of directors so that at least half of them have to agree to opening the vault?

Definition 3.68 (Threshold Secret Sharing). *Let $t, n \in \mathbb{N}$ with $1 \leq t \leq n$. An algorithm that distributes a secret among n participants such that t participants need to collaborate to recover the secret is called a (t,n)-threshold secret sharing scheme.*

```
1    # s = secret real number to be shared
2    # t = threshold number of participants to recover the secret
3    # n = total number of participants
4
5    def distribute(s, t, n):
6        Generate t − 1 random a_1, ..., a_{t-1} ∈ R
7        Obtain a polynomial f(x) = s + a_1 x + ... + a_{t-1} x^{t-1}
8        Generate n distinct x_1, ..., x_n ∈ R \ {0}
9        Send (x_i, f(x_i)) to participant P_i
10
11   def recover(x = [x_0, x_1, ..., x_t], y = [f(x_0), f(x_1), ..., f(x_t)]):
12       f = lagrange(x, y)
13       return f(0)
```

Algorithm 3.69: Shamir's (t, n) Secret Sharing Scheme

Theorem 3.70. *Algorithm 3.69 is correct.*

Proof. Any t shares will result in the reconstruction of the same polynomial, hence the secret will be revealed. $\qquad\square$

Theorem 3.71. *Algorithm 3.69 has perfect security.*

Proof. A polynomial of degree $t-1$ can be defined only by t or more points. So, any subset $t-1$ of the n shares cannot reconstruct a polynomial of degree $t-1$. Given less than t shares, all polynomials of degree $t-1$ are equally likely; thus any adversary, even with unbounded computational resources, cannot deduce any information about the secret if they have less than t shares. $\quad\square$

Remarks:

- Note that for numerical reasons, in practice modulo p arithmetic is $\rightarrow$ notebook used instead of real numbers.

- What happens if a participant is malicious? Suppose during recovery, one of the t contributing participants publishes a wrong share $(x_i', f(x_i'))$. The $t-1$ honest participants are blocked from the secret while the malicious participant is able to reconstruct it. To prevent this, we employ *verifiable* secret sharing schemes.

Definition 3.72 (Verifiable Secret Sharing or VSS). *An algorithm that achieves threshold secret sharing and ensures that the secret can be reconstructed even if a participant is malicious is called verifiable secret sharing.*

Remarks:

- Typically, a secret sharing scheme is verifiable if auxiliary information is included that allows participants to verify their shares as consistent.

- VSS protocols guarantee the secret's reconstruction even if the distributor of the secret (the dealer) is malicious.

- So far, we assumed a dealer knows the secret and all the shares. However, we want to avoid trusted third parties and distribute trust. A strong cryptographic notion towards this direction is multiparty computation.

3.12   Multiparty Computation

Alice, Bob, and Carol are interested in computing the sum of their income without revealing to each other their individual income.

```python
# a, b, c = Alice's, Bob's and Carol's income

def Sum_MPC():
    Alice picks a large random number r
    Alice sends to Bob m_1 = a + r
    Bob sends to Carol m_2 = b + m_1
    Carol sends to Alice m_3 = c + m_2
    Alice computes s = m_3 - r
    Alice shares s with Bob and Carol
```

Algorithm 3.73: Computation of the Sum of 3 Parties' Income

Theorem 3.74. *Algorithm 3.73 is correct, meaning the output is the desired sum.*

Proof. The output of the algorithm is $m_3 - r = c + m_2 - r = c + b + m_1 - r = c + b + a + r - r = a + b + c$. $\qquad\square$

Theorem 3.75. *Algorithm 3.73 keeps the inputs secret.*

Proof. Bob receives $r + a$, hence no information is revealed concerning Alice's income as long as r is large enough. In addition, both Carol and Alice cannot deduce any information about the individual incomes as they are obfuscated. $\square$

Remarks:

- Algorithm 3.73 is an example of secure 3-party computation.

- The generalization of this problem to multiple parties is known as multiparty computation.

Definition 3.76 (Multiparty Computation or MPC). *An algorithm that allows n parties to jointly compute a function $f(x_1, x_2, \ldots, x_n)$ over their inputs $x_1, x_2, \ldots, x_n$ while keeping these inputs secret achieves secure multiparty computation.*

Remarks:

- Formal security proofs in MPC protocols are conducted in the *real/ideal world paradigm*.

Definition 3.77. *The real/ideal world paradigm states two worlds: In the ideal world, there is an incorruptible trusted third party who gathers the participants' inputs, computes the function, and returns the appropriate outputs. In contrast, in the real world, the parties exchange messages with each other. A protocol is secure if one can learn no more about each participant's private inputs in the real world than one could learn in the ideal world.*

Remarks:

- In Algorithm 3.73, we assume all participants are honest. But what if some participants are malicious?

- In MPC, the computation is often based on secret sharing of all the inputs and zero-knowledge proofs for a potentially malicious participant. Then, the majority of honest parties can assure that bad behavior is detected and the computation continues with the dishonest party eliminated or her input revealed.

Chapter Notes

In 1974, Ralph Merkle designed Merkle Puzzles [15], the first key exchange scheme which works over an insecure channel. In Merkle Puzzles, the eavesdropper Eve's computation power can be at most quadric to Alice's and Bob's computational power. This quadratic difference is not enough to guarantee security in practical cryptographic applications. In 1976, Diffie and Hellman introduced a practically secure key exchange scheme over an insecure channel [6].

Diffie Hellman key exchange [6], Schnorr zero-knowledge proofs [19], ElGamal signature and encryption schemes [7] all rely on the hardness of the discrete logarithm problem [3]. So far we have been conveniently vague in our choice of a group, but the discrete logarithm problem is solvable in polynomial-time when we choose an inappropriate group. To avoid this, we can select a group that contains a large subgroup. For example, if $p = 2q + 1$ and q is prime, there is a subgroup of size q, called the quadratic residues of p, which is often used in practice.

Another frequently employed hard problem is integer factorization [12]. The RSA cryptosystem [18], developed in 1977 at MIT by Ron Rivest, Adi Shamir, and Leonard Adleman, depends on integer factorization. RSA was also the first public-key encryption scheme that could both encrypt and sign messages.

A trapdoor one-way function is a function that is easy to compute, difficult to invert without the trapdoor (some extra information), and easy to invert with a trapdoor [6, 25]. The factorization of a product of two large primes, used in RSA, is a trapdoor function. While selecting and verifying two large primes and multiplying them is easy, factoring the resulting product is (as far as known) difficult. However, if one of the prime numbers is given as a trapdoor, then it is easy to compute the other prime number. There are no known trapdoor one-way functions based on the difficulty of discrete logarithms (either modulo a prime or in a group defined over an elliptic curve), because there is no known "trapdoor" information about the group that enables the efficient computation of discrete logarithms. In general, a digital signature scheme can be built by any trapdoor one-way function in the random oracle model [14].

A random oracle [1] is a function that produces a random output for each query it receives. It must be consistent with its replies: if a query is repeated, the random oracle must return the same answer. Hash functions are often modeled in cryptographic proofs as random oracles. If a scheme is secure assuming the adversary views some hash function as a random oracle, it is said to be secure in the random oracle model.

Secure digital signature schemes are unforgeable. There are several versions of unforgeability. For instance, Schnorr signatures, a modification of ElGamal signatures, are existentially unforgeable against adaptively chosen message attacks (EUF-CMA) [20]. In the adaptively chosen message attack, the adversary wants to forge a signature for a particular public key (without access to the corresponding secret key) and has access to a signing oracle, which receives messages and returns valid signatures under the public key in question. The proof that Schnorr digital signatures are EUF-CMA is based on the proof that the Schnorr zero-knowledge proof is sound.

Zero-knowledge proofs are a complex cryptographic primitive; formally defining zero-knowledge proofs was a delicate task that took 15 years of research [2, 10]. One key application for zero-knowledge proofs is in user identification schemes. Another recent one is in cryptocurrencies, such as Monero [23].

The concept of information-theoretically secure communication was introduced in 1949 by American mathematician Claude Shannon, the inventor of information theory, who used it to prove that the one-time pad system was secure [22]. Secret sharing schemes are information theoretically secure. Verifiable secret sharing was first introduced in 1985 by Benny Chor, Shafi Goldwasser, Silvio Micali and Baruch Awerbuch [5]. Thereafter, Feldman introduced a practical verifiable secret sharing protocol [9] which is based on Shamir's secret sharing scheme [21] combined with a homomorphic encryption scheme. Verifiable secret sharing is important for secure multiparty computation to handle active adversaries.

Multiparty computation (MPC) was formally introduced as secure two-party computation (2PC) in 1982 for the so-called Millionaires' Problem, a specific problem which is a Boolean predicate, and in general, for any feasible computation, in 1986 by Andrew Yao [24, 26]. MPC protocols often employ a cryptographic primitive called oblivious transfer.

An oblivious transfer protocol, originally introduced by Rabin in 1981 [17], allows a sender to transfer one of potentially many pieces of information to a receiver, while remaining oblivious as to what piece of information (if any) has been transferred. Oblivious transfer is complete for MPC [11], that is, given an implementation of oblivious transfer it is possible to securely evaluate any polynomial time computable function without any additional primitive! An "1-out-of-n" oblivious transfer protocol [8, 16, 13] is a generalization of oblivious transfer where a receiver gets exactly one database element without the server (sender) getting to know which element was queried, and without the receiver knowing anything about the other elements that were not retrieved. A weaker version of "1-out-of-n" oblivious transfer, where only the sender should not know which element was retrieved, is known as Private Information Retrieval [4].

This chapter was written in collaboration with Zeta Avarikioti, Tejaswi Nadahalli, Ard Kastrati, and Klaus-Tycho Foerster.

Bibliography

[1] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and communications security*, pages 62–73, 1993.

[2] Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, 1991.

[3] Dan Boneh. The decision diffie-hellman problem. In *International Algorithmic Number Theory Symposium*, pages 48–63. Springer, 1998.

[4] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 41–50. IEEE, 1995.

[5] Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 383–395. IEEE, 1985.

[6] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.

[7] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.

[8] Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. *Communications of the ACM*, 28(6):637–647, 1985.

[9] Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science (sfcs 1987)*, pages 427–438. IEEE, 1987.

[10] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 285–306. 2019.

[11] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer–efficiently. In *Annual international cryptology conference*, pages 572–591. Springer, 2008.

[12] Thorsten Kleinjung, Kazumaro Aoki, Jens Franke, Arjen K Lenstra, Emmanuel Thomé, Joppe W Bos, Pierrick Gaudry, Alexander Kruppa, Peter L Montgomery, Dag Arne Osvik, et al. Factorization of a 768-bit rsa modulus. In *Annual Cryptology Conference*, pages 333–350. Springer, 2010.

[13] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious prf with applications to private set intersection. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 818–829, 2016.

[14] Leslie Lamport. Constructing digital signatures from a one-way function. Technical report, 1979.

[15] Ralph C Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21(4):294–299, 1978.

[16] Moni Naor and Benny Pinkas. Oblivious polynomial evaluation. *SIAM Journal on Computing*, 35(5):1254–1281, 2006.

[17] Michael O Rabin. How to exchange secrets with oblivious transfer.

[18] Ronald L Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.

[19] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *Conference on the Theory and Application of Cryptology*, pages 239–252. Springer, 1989.

[20] Yannick Seurin. On the exact security of schnorr-type signatures in the random oracle model. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 554–571. Springer, 2012.

[21] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[22] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[23] Nicolas van Saberhagen. Monero whitepaper. Technical report, 2013.

[24] Andrew C Yao. Protocols for secure computations. In *23rd annual symposium on foundations of computer science (sfcs 1982)*, pages 160–164. IEEE, 1982.

[25] Andrew C Yao. Theory and application of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (SFCS 1982)*, pages 80–91. IEEE, 1982.

[26] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 162–167. IEEE, 1986.

Chapter 4

Data & Storage

A computer does more than just computation. In particular, a computer can also store and retrieve large amounts of data efficiently. In this chapter, we want to understand some of the key ingredients when it comes to dealing with data and storage.

4.1 Dictionary

We manage a library and want to be able to quickly tell whether we carry a given book or not. We need the capability to *insert*, *delete*, and *search* books.

Definition 4.1 (Dictionary). *A **dictionary** is a data structure that manages* → notebook
*a set of **objects**. Each object is uniquely identified by its **key**. The relevant operations are*

- search: *find an object with a given key*

- insert: *put an object into the set*

- delete: *remove an object from the set*

Remarks:

- There are alternative names for dictionary, e.g. key-value store, associative array, map, or just set.

- If the dictionary only offers search, it is called *static*; if it also offers insert and delete, it is *dynamic*.

- When discussing the algorithms, we will often ignore that we actually have a set of objects, each of which is identified by a unique key, and just talk about the set of keys. With regard to the library example, books are globally uniquely identified by a key called ISBN. Whenever we say we insert/delete/search a key, we can just drag the key's object along.

- The classic data structure for dictionaries is a binary search tree.

Definition 4.2 (Binary search tree). *A **binary search tree** is a rooted tree, where each node stores a key. Additionally, each node may have a pointer to a left and/or right child tree. For all nodes, if existing, the nodes in the left child tree store smaller keys, and those in the right child tree store larger keys.*

```python
def search(self, key): # self is current node, initially root
    if key < self.key:
        if self.left is None: return None
        else: return self.left.search(key)
    elif key > self.key:
        if self.right is None: return None
        else: return self.right.search(key)
    return self.val
```

→ notebook

Algorithm 4.3: Search Tree: Search

Remarks:

- The cost of searching in a binary search tree is proportional to the *depth* of the key, which is the distance between the node with the key and the root.

- There are search trees called *splay trees* that keep frequently searched keys close to the root for quick access. On the other hand, there may be rarely accessed keys deep in a splay tree.

- Using *balanced search trees*, we can maintain a dictionary with worst-case logarithmic depth for all keys, and thus worst-case logarithmic cost per insert/delete/search operation.

- Is there a way to build a dictionary with less than logarithmic cost and with keys that cannot be ordered?

4.2 Hashing

In this section we use hashing to implement an efficient dictionary.

Definition 4.4 (Universe, Key Set, Hash Table, Buckets). *We consider a **universe** U containing all possible keys. We want to maintain a subset of this universe, the **key set** $N \subseteq U$ with $|N| =: n$, where $|N| \ll |U|$. We will use a **hash table** M, i.e. an array M with m **buckets** $M[0], M[1], \ldots, M[m-1]$.*

Remarks:

- The standard library of almost every widely used programming language provides hash tables, sometimes by another name. In C++, they are called unordered_map, in Python dictionary, in Java HashMap.

- The translation from virtual memory to physical memory uses a piece of hardware called *translation lookaside buffer* (TLB), which is a hardware implementation of a hash table. It has a fixed size and acts like a cache for frequently looked up virtual addresses.

- Compilers make use of hash tables to manage the symbol table.

Definition 4.5 (Hash Function)**.** *Given a universe U and a hash table M, a* **hash function** *is a function $h : U \to M$. Given some key $k \in U$, we call $h(k)$ the* **hash** *of k.*

Remarks:

- A hash function should be fast to compute and distribute hashes nicely, e.g. $h(k) = k \bmod m$ for a key $k \in \mathbb{N}$; in contrast to Chapter 3, we do not care whether a hash function is one-way.

- If we use ISBN mod m as our library hash function, can we insert/delete/search books in constant time?!

- What if two keys $k \neq k'$ have $h(k) = h(k')$?

Definition 4.6 (Collision)**.** *Given a hash function $h : U \to M$, two distinct keys $k, k' \in U$ produce a* **collision** *if $h(k) = h(k')$.*

Remarks:

- Since keys may experience collisions, the key must be stored in the bucket.

- There are competing objectives we want to optimize for when hashing. On the one hand, we want to make the hash table small since we want to save memory. On the other hand, small tables will have more collisions. How likely is it to get a collision for a given n and m?

Theorem 4.7 (Birthday Problem)**.** *If we throw a fair m-sided dice $n \leq m$ times, let D be the event that all throws show different numbers. Then D satisfies*

$$\mathbb{P}[D] \leq \exp\left(-\frac{n(n-1)}{2m}\right).$$

Proof. We have that

$$\mathbb{P}[D] = \frac{m}{m} \cdot \frac{m-1}{m} \cdot \ldots \cdot \frac{m-(n-1)}{m} = \prod_{i=0}^{n-1} \frac{m-i}{m}$$

$$= \prod_{i=0}^{n-1}\left(1 - \frac{i}{m}\right) = \exp\left(\sum_{i=0}^{n-1} \ln\left(1 - \frac{i}{m}\right)\right)$$

We can use that $\ln(1+x) \leq x$ for all $x > -1$ and the monotonicity of e^x:

$$\mathbb{P}[D] = \exp\left(\sum_{i=0}^{n-1} \ln\left(1 - \frac{i}{m}\right)\right) \leq \exp\left(\sum_{i=0}^{n-1} -\frac{i}{m}\right) = \exp\left(-\frac{n(n-1)}{2m}\right)$$

$\square$

Remarks:

- Theorem 4.7 is called the "birthday problem" since traditionally, people use birthdays for illustration: In order to have a chance of at least 50% that two people in a group share a birthday, we only need 23 people.

- If we insert more than roughly $n \approx \sqrt{m}$ keys into a hash table, the probability of a collision approaches 1 quickly. In other words, unless we are willing to use at least $m \approx n^2$ space for our hash table, we will need a good strategy for resolving collisions.

- Theorem 4.7 assumes totally random hash functions — for non-random distributions of hashes, we might have more collisions. In particular, if we fix a hash function, then we can always end up with a key set N that suffers from many collisions. E.g., if many books have an ISBN that ends in 000, then ISBN mod 1000 is a terrible hash function.

- Maybe we can use modulo, but with a different m?

- However, for any hash function there are bad key sets.

- On the other hand, for every key set there are good hash functions! How do we efficiently pick a good hash function, i.e. one that is likely to distribute hashes well?

Definition 4.8 (Universal Family)**.** *Let $\mathcal{H} \subseteq \{h : U \to M\}$ be a family of hash functions from U to M. If for all pairs of distinct keys $k \neq k' \in U$, the probability of a collision is $\mathbb{P}[h(k) = h(k')] \leq \frac{1}{m}$ when we choose $h \in \mathcal{H}$ uniformly, then $\mathcal{H}$ is called a* **universal family (of hash functions)***.*

Remarks:

- In other words: if we choose a hash function from a universal family, we can expect the hashes to be distributed well, regardless of the key set.

- We cannot just pick a random function from U to M because there are $|M|^{|U|}$ many, so we need $|U| \log |M|$ bits to encode such a random function. That is even more bits than keys in our huge universe U.

Lemma 4.9. *For prime m, $j \in \{0, \ldots, m-1\}$, $\delta \in \{1, \ldots, b-1\}$, and $b \leq m$, the linear function $f_\delta(j) := j \cdot \delta \bmod m$ is a bijection.*

Proof. Suppose there are integers j, j' with $0 \leq j < j' < m$ such that $f_\delta(j) = f_\delta(j') \bmod m$. Then, there exists an integer k such that $j' \cdot \delta = j \cdot \delta + km \Leftrightarrow \delta(j' - j) = km$, which implies that $m \mid \delta(j' - j)$. Since δ is relative prime to m, we have $m \mid (j' - j)$, thus $j = j' \bmod m \Rightarrow j = j'$. $\square$

Theorem 4.10 (Universal Hashing)**.** *Let m be prime and $s \in \mathbb{N}$. Let $U = \{0, \ldots, b-1\}^s$ and let $M = \{0, \ldots, m-1\}$ with $b \leq m$. For a key $k = (k_0, \ldots, k_{s-1}) \in U$ and coefficient tuple $j = (j_0, \ldots, j_{s-1}) \in \{0, \ldots, m-1\}^s$, define*

$$h_j(k_0, \ldots, k_{s-1}) = \sum_{i=0}^{s-1} j_i \cdot k_i \bmod m.$$

Then $\mathcal{H} := \{h_j : j \in \{0, \dots, m-1\}^s\}$ is a universal family of hash functions.

Proof. Let $(k_0, \dots, k_{s-1}) = k \neq k' = (k'_0, \dots, k'_{s-1}) \in U$. Using $\delta_i = k'_i - k_i$ we get

$$h_j(k') - h_j(k) = \sum_{i=0}^{s-1} j_i \cdot k'_i - \sum_{i=0}^{s-1} j_i \cdot k_i = \sum_{i=0}^{s-1} j_i \cdot \delta_i \mod m.$$

The terms with $\delta_i = 0$ are 0, and so we can ignore them. Let $X := \{i \in \{0, \dots, m-1\} : \delta_i \neq 0\}$ be the non-empty set of the indices of the non-zero terms, with $x = |X|$. There are m^x possibilities to choose the coefficients $\{j_i : i \in X\}$. So how many of these possibilities will have a conflict, i.e., a sum of 0? If we choose $x - 1$ coefficients arbitrarily, then the last coefficient can always render the sum 0, as according to Lemma 4.9 the last coefficient can add any value $\{0, \dots, m-1\}$ to the sum. In other words, m^{x-1} sums are 0. Therefore, our chance of randomly picking an j that produces a collision is $\frac{m^{x-1}}{m^x} = \frac{1}{m}$. ☐

Remarks:

- Theorem 4.10 gives us a general method for picking hash functions from a universal family in an efficient manner. We simply choose a prime number m and uniformly at random some factors $j_0, \dots, j_r$. Thus, we can represent our hash function as the tuple $(m, j_0, \dots, j_r)$.

- In practice, hash tables perform really well, and if we detect that we had bad luck in choosing our hash function, we just choose a new one and rebuild our table with the new function — this is called *rehashing*.

4.3 Static Hashing

How can we state the tradeoff between space and collisions more precisely?

Definition 4.11 (Number of Collisions). *Given a hash function $h : U \to M$ and a key set $N \subseteq U$, define the number of collisions that h produces on N as*

$$C(h, N) := |\{\{k, k'\} \subseteq N : k \neq k', h(k) = h(k')\}|.$$

Lemma 4.12 (Space vs. Collisions). *We want a hash function h which produces less than c collisions, i.e. $C = C(h, N) < c$, with $n = |N|$. If h comes from a universal family, a hash table size $m = \lceil \frac{n(n-1)}{c} \rceil$ works.*

Proof. There are $\binom{n}{2}$ pairs of distinct keys in N, and each pair produces a collision with probability at most $1/m$ since h is chosen from a universal family. Using the linearity of expectation we can bound the number of expected collisions:

$$\mathbb{E}[C] \leq \binom{n}{2} \cdot \frac{1}{m} = \frac{n(n-1)}{2m}.$$

We choose m such that $2 \cdot \mathbb{E}[C] \leq c$:

$$2 \cdot \mathbb{E}[C] \leq c \Leftrightarrow \frac{n(n-1)}{m} \leq c \Leftrightarrow \frac{n(n-1)}{c} \leq m.$$

Now we use the Markov inequality which states that for any random variable X that only takes on non-negative integer values, we have $\mathbb{P}[X \geq k \cdot \mathbb{E}[X]] \leq \frac{1}{k}$. Hence, we have $Pr[C < 2 \cdot \mathbb{E}[C] \leq c] \geq \frac{1}{2}$. In other words, to find a hash function with $C < c$ we only need to sample 2 hash functions in expectation.

☐

Remarks:

- According to Theorem 4.12, if we want no collisions, we set $c = 1$ and choose $m = \lceil \frac{n(n-1)}{1} \rceil = n(n-1)$.

- Similarly, if we can tolerate n collisions, we find that a hash table of size $m = n - 1$ suffices.

- Algorithm 4.13 defines perfect static hashing, which applies the result of Theorem 4.12.

<!-- → notebook -->

```python
# Compute primary hash table M and secondary hash tables M_i
def perfect_static_hashing(N): # N = fixed set (list) of keys
    n = len(N)
    M = [None]*n
    h = None
    h2 = [None]*n
    while True:
        h = from_universal_family(n) # h: N -> range(n)
        if C(h,N) < n: break # C counts the number of collisions
    for i in range(n):
        N_i = [x for x in N if h(x) = i]
        n_i = len(N_i)
        if n_i > 1:
            M[i] = [None]*(n_i*(n_i-1))
            while True:
                h2[i] = from_universal_family(len(M[i]))
                if C(h2[i], N_i) < 1: break
    return M, h, h2
```

Algorithm 4.13: Perfect Static Hashing

Remarks:

- In the first stage, we find a hash function h with at most n collisions in linear space according to Theorem 4.12.

- In the second stage, we find a hash function h_i per bucket i without collisions by using an amount of space that is quadratic in the number of keys in the bucket n_i as per Theorem 4.12.

Theorem 4.14 (Perfect Static Hashing). *Algorithm 4.13 returns a collision-free data structure of with a total size (M and all M_i) less than $3n$.*

Proof. The data structure is collision-free because of the *if* condition in the inner *while* loop. Bucket i is of size $n_i(n_i - 1) = 2\binom{n_i}{2}$, holds n_i keys that produced $\binom{n_i}{2}$ collisions. Since the total number of collisions is less than n (*if* condition in the first *while* loop), and each bucket is of size twice its collisions, the total size of all buckets is less than $2n$. The size of the primary table M is n, so together the total size of the data structure is less than $n + 2n = 3n$. □

Remarks:

- We now have a hashing algorithm that can be built in linear space and expected linear time, and offers worst-case constant time search for a static set N.

- But what about a dynamic dictionary?

4.4 Dynamic Hashing

Definition 4.15 (Hashing with Chaining). *In **hashing with chaining**, every bucket $M[i]$ stores a pointer to a secondary data structure that manages all keys k with $h(k) = i$. Insertion, search, and deletion of k are all relegated to those data structures. In the simplest implementation, we use a list for each bucket.*   → notebook

Remarks:

- Algorithm 4.13 is an instance of hashing with chaining with the M_i being the secondary data structures managing the buckets.

- The Java standard library uses hashing with chaining to resolve collisions.

Definition 4.16 (Load Factor). *The fraction $\frac{n}{m} =: \alpha$ is called the **load factor** of the hash table.*

Remarks:

- The performance of all three operations (insert/delete/search) depends on the load factor for all collision resolution strategies.

- Hashing with chaining allows for a load factor $\alpha > 1$ since the size of the table is the number of secondary data structures; performance deteriorates with growing α.

- If we use linked lists as secondary structures and use a hash function chosen from a universal family, the cost for an unsuccessful search is $1 + \alpha$ in expectation, while that for a successful search is roughly $1 + \frac{\alpha}{2}$ in expectation.

- If we use one of the strategies of this section and α grows too large, we should rehash with a bigger m in order to remain efficient. In the Java standard library, if a hash table surpasses a load factor of 0.75, it is rehashed into a hash table with twice the size of the old one.

Definition 4.17 (Hashing with Probing). *In **hashing with probing**, keys are stored directly in the hash table. The sequence $(h_i(k) \bmod m)_{i \geq 0}$ is called the **probing sequence** of k, and each step of the iteration is a **probe**.*

```
def search(self, k):
    i = 0
    while i < self.size:
        j = self.h(i, k) % self.size   # j = h_i(k)
        key_value = self.M[j]
        if key_value == None: return None
        elif key_value[0] == k: return key_value[1]
        i += 1
    return None
```

Algorithm 4.18: Hashing with Probing: Search

Remarks:

- Algorithm 4.18 defines how to search for a key in hashing with probing. The search is *successful* if the *elif* condition is satisfied. If the preceding *if* is satisfied (or *None* returned at the end) the search is *unsuccessful*.

- To insert a key, we adapt Algorithm 4.18: with an unsuccessful search where *if* was satisfied, we insert in the empty bucket. Therefore, the cost of an insert is roughly the cost of an unsuccessful search. An unsuccessful search where the loop finished triggers a rehash.

- Table 4.19 describes three different types of hashing with probing, each together with the approximate time that a successful or unsuccessful search takes in expectation. More generally, linear probing uses some linear function $h_i(k) = h(k) + ci$ for some $c \neq 0$, and quadratic probing uses some quadratic function $h_i(k) = h(k) + ci + di^2$ with $d \neq 0$. As long as we guarantee that $h_i(k)$ is integer for all $i \in [m]$, the constants c and d can be rational.

Probing	$h_i(k)$	$\approx$ cost successful	$\approx$ cost unsuccess.
Linear	$h(k) + i$	$\frac{1}{2}\left(1 + \frac{1}{(1-\alpha)^2}\right)$	$\frac{1}{2}\left(1 + \frac{1}{1-\alpha}\right)$
Quadratic	$h(k) + i^2$	$\frac{1}{1-\alpha} + \ln\frac{1}{1-\alpha} - \alpha$	$1 + \ln\frac{1}{1-\alpha} - \frac{\alpha}{2}$
Double hashing	$h_1(k) + i \cdot h_2(k)$	$\frac{1}{1-\alpha}$	$\frac{1}{\alpha}\ln\left(\frac{1}{1-\alpha}\right)$

Table 4.19: Different types of hashing with probing together with the expected number of probes per search, where α is the load factor of the table. For hashing with probing, we need $\alpha \leq 1$ since we must have $n \leq m$. Each of h, h_1, h_2 is a hash function drawn from a universal family.

Remarks:

- What is the reason for the different costs in Table 4.19?

- Linear probing suffers from so-called *primary clustering*: the probing sequences are simplistic; if two probing sequences meet in the same bucket j, they will continue to collide in buckets $j + 1, j + 2, \ldots$

- Quadratic probing does not suffer from primary clustering, but it is subject to *secondary clustering*: if two keys have the same hash, then their probing sequences are identical.

- The form of quadratic probing defined in Table 4.19 has one additional issue: the probing sequence of a key does not necessarily cover the whole table. Assume $m = 7$ buckets and $h(k) = 0$, then the probing sequence of k is $(0, 1, 4, 2, 2, 4, 1)$ — buckets $3, 5, 6$ do not appear.

- Double hashing does not suffer from either version of clustering. One can show that if the hash functions h_1 and h_2 used in double hashing are independently drawn from a universal family, then double hashing performs as well as an idealized hash function that assigns hashes uniformly at random.

4.5 Cuckoo Hashing

So far, the cost of all operations for dynamic key sets has been given in expected time cost. There are algorithms that allow us to do better and give us worst case guarantees on some of the operations. Two widely known possibilities to achieve this are called *dynamic perfect hashing* and *cuckoo hashing*.

→ notebook

```
1  Cuckoo hashing uses two hash tables, and as such provide two
   ↪  possible buckets M₁[h₁(k)] or M₂[h₂(k)] for each key k.
2  If one of the two buckets is empty, simply place k there.
3  If both buckets are occupied by other keys, k is anyway inserted
   ↪  in one of the two possible buckets, replacing some key k₁
   ↪  that currently resides in this bucket.
4  The kicked out key k₁ moves to its other bucket, potentially
   ↪  kicking out the currently resident key k₂; this is repeated
   ↪  recursively until an empty bucket is found.
5  If this recursion loops or takes too long (logarithmic in the
   ↪  table size), the hash table is rebuilt using two new hash
   ↪  functions.
```

Algorithm 4.20: Cuckoo Hashing: Insert

Remarks:

- Search and delete only need to check two buckets to figure out whether a given key is in the table, and so those operations are worst case constant time.

- One can show that the expected insert cost in *cuckoo hashing* is constant as long as the load factor α is below 0.5.

- Cuckoo hashing gets its name from cuckoo birds: they lay their eggs into the nests of other birds, and once the cuckoo chicks hatch, they push the other eggs/chicks out of the nest.

- The idea behind cuckoo hashing is to use the *"power of two choices"*, which can be roughly described as: if you can choose between two resources and use the one that is less busy, you gain efficiency.

- To adapt perfect static hashing to a dynamic setting where we can also handle inserts and deletions, all we have to do is choose the size of M_i twice as large as in Algorithm 4.13, and rehash appropriately: Whenever $C(h_i, N_i) > 0$ for some bucket i, we rehash that bucket until there are no collisions. Once some bucket reaches $n_i^2 \approx |M_i|$ due to insertions, we rehash the entire table. This leaves us with expected constant time insert and delete, and worst case constant time search. To keep the table linear-sized, we rehash everything after every m updates (inserts or deletes).

4.6 Key-Value Databases

Definition 4.21 (Key-Value Database System)**.** *The concept of dictionaries is used in key-value database systems. The server maintains the dictionary and clients can insert and query the stored data using the keys.*

Remarks:

- Popular key-value databases are Redis and Memcached. They are often used for caching in web services. Dynamically generated documents or results of queries to other databases can be stored temporarily to allow fast access to often requested data.

- The data is often kept in main memory to speed up the access and only duplicated to disk to recover the database in case of a system failure.

- Depending on the used database, different data types can be stored in the value. This can be an integer, a string, or even an array.

- Document databases are an extension of simple key-value database systems. The value has to be in a format that the database understands, such as a JSON or XML document. These databases allow queries on the content of the documents. MongoDB and CouchDB are popular document databases.

4.7 Relational Databases

However, most databases offer queries beyond simple key searches. Questions like "What is the movie with the largest cast?" or "How many directors have directed more than ten movies?" should be answered without first writing a new program. Relational databases can store large amounts of *structured* data and answer possibly complex questions about it.

Definition 4.22 (Table, Row, Column, Database). *A **table** consists of **rows**, so that each row (data record) contains the same **fields**, i.e., kinds of entries. When the rows of a table are written line by line, the fields form the **columns** of the table. Each column is referred to by a descriptive name, and is associated with the type of the respective field, e.g., integer, floating point, string, or a date. A **database** is a collection of tables.*

Remarks:

- In the database context, tables are also called *relations*, because the entries in each row are related to each other, namely by belonging to the same row.

<div align="center">

movies

title	director	year
12 Angry Men	Sidney Lumet	1957
Raiders of the Lost Ark	Steven Spielberg	1981
War of the Worlds	Steven Spielberg	2005
Manos: The Hands of Fate	Harold P. Warren	1966

⋮

</div>

Figure 4.23: A database containing a single table called "movies" storing the title, director, and year of release for each movie.

Remarks:

- Databases as we study them are accessed using the so-called *structured query language* (SQL). Thus they are referred to as SQL or *relational* databases.

- MySQL and PostgreSQL are two popular open source SQL database systems.

- SQL database systems typically run as a daemon process on some server. Client applications connect to the server and authenticate themselves via username and password. Therefore, multiple users accessing the same database may result in concurrency issues. Some form of concurrency control is necessary!

- Other database systems are tailored to single-user processing. They relieve developers from the burden of implementing efficient data structures for relational data. SQLite is one such example, and is used, e.g., in Firefox, Chrome, Android, Adobe Lightroom, and Windows 10.

4.8 SQL Basics

Definition 4.24 (SQL Data Types). *SQL defines the following types of columns.*
- *CHARACTER(m) and CHARACTER VARYING(m) for fixed and variable length strings of (maximum) length m,*
- *BIT(m) and BIT VARYING(m) for fixed and variable length bit strings of (maximum) length m,*
- *NUMERIC, DECIMAL, INTEGER, and SMALLINT for fixed point and integer numbers,*
- *FLOAT, REAL, and DOUBLE PRECISION for floating point numbers,*
- *DATE, TIME, and TIMESTAMP for points in time, or*
- *INTERVAL for ranges of time.*

Remarks:

- The range of each type includes the special value NULL. Note that NULL is different from the string 'NULL', the empty string, and from the number 0 (zero). NULL indicates that the row has *no value* for the corresponding field.

- Many database systems implement more types, e.g., geographic coordinates, IP addresses, geometric objects, or large integers.

- All SQL statements end with a semicolon. The SQL language is case insensitive, but by convention keywords are often typed in upper case.

- The SQL-92 specification is over 600 pages long, newer versions of the standard even longer. To add insult to injury there are lots of vendor specific "SQL dialects", i.e., modifications and extensions. However, the basic set of commands for creating, manipulating, and querying tables are largely the same across database implementations.

CREATE DATABASE *database-name***;**
Additional parameters allow to set database-specific options, e.g., user-based permissions, or default character sets for text strings. How a database is opened depends on the implementation.

CREATE TABLE *table-name* **(***field-name type, field-name type, . . .***);**
To enforce that all rows have a value for a particular field, one can add NOT NULL to the type when creating the table. Fields have a default value, which is NULL if not specified by adding DEFAULT *value* to the type description.

Remarks:

- There are also GUI and web-based client applications (that execute locally or on an http-server, respectively) and offer access to the database in a more intuitive manner than the classic command line tools. Examples for PostgreSQL are pgAdmin, DataGrip and DBeaver.

- Such tools are especially helpful for creating the databases and tables and often support multiple database systems. They also feature importing data from various formats, e.g., CSV files, instead of using SQL statements to populate the tables.

INSERT INTO *table-name* (*field-name*, . . .) **VALUES** (*value*, . . .);   → notebook

Values must be listed in the same order as the corresponding field names. When a field name (and thus its value) is omitted the field's default value is assumed. When the list of field names is omitted the field's values must be listed in the same order that was used when creating the table. To insert more than one row in one statement, multiple rows may be separated by commas.

```
SELECT * FROM movies;
SELECT * FROM movies WHERE director = 'Spielberg, Steven';
SELECT title FROM movies WHERE year BETWEEN 1990 AND 1999;
SELECT * FROM movies WHERE title IS NULL OR director IS NULL;
SELECT title, director FROM movies WHERE title LIKE '%the%';
```
→ notebook

Listing 4.25: Querying the movies table.

SELECT *field-name*, . . . **FROM** *table-name* **WHERE** *condition*;

Lists all specified fields of all rows in the table that fulfill the condition. The special field * lists all fields. The WHERE condition may be omitted to list the whole table. A condition can include comparisons ($<, >, =, <>$) between fields constants. The special value NULL can be tested with IS NULL. Conditions can be joined using parenthesis and logic operators like AND, OR, and NOT. Strings can be matched with patterns using **field-name LIKE pattern**. In the pattern, an underscore (_) matches a single character, whereas % matches arbitrarily many.

```
SELECT MIN(year) FROM movies;
SELECT AVG(year) FROM movies WHERE director='Lumet, Sidney';
SELECT COUNT(*) FROM movies;
SELECT COUNT(DISTINCT director) FROM movies;
```
→ notebook

Listing 4.26: Aggregation with SQL.

SELECT *aggregate*, . . . ;

Functions for aggregation include AVG to compute the average of a certain field, MIN and MAX for the minimum and maximum value, SUM for the sum of a field, and COUNT to count the number of occurrences. In an aggregation, the keyword DISTINCT indicates that only distinct values should be considered.

```
SELECT director, COUNT(title) FROM movies GROUP BY director;
SELECT director, COUNT(title) FROM movies GROUP BY director
HAVING COUNT(title) > 10;
SELECT year, director, COUNT(title) FROM movies
GROUP BY director, year
ORDER BY year DESC, director ASC;
```
→ notebook

Listing 4.27: Grouping and sorting.

SELECT *field-name*|*aggregate*, . . . **GROUP BY** *field-name*,. . . ;

Aggregations may be partitioned using the group-by clause. Similar to before, the query result can only include aggregates and fields by which the result is partitioned.

Since WHERE clauses are applied before GROUP BY the result of aggregations cannot appear in them. When the result should be conditioned on the result of an aggregation, a HAVING clause can be used.

SELECT . . . **ORDER BY** *field-name*,. . . ;

After each field-name, the keyword ASC or DESC can be used to determine ascending or descending sorting order, respectively.

```
UPDATE movies SET title = 'Star Wars Episode IV: A New Hope'
WHERE title = 'Star Wars';
DELETE FROM movies WHERE title = '';
```

Listing 4.28: Updating and removing rows.

UPDATE *table* **SET** *field-name* = *value*,. . . **WHERE** *condition*;

Updates the specified fields in all rows fulfilling the condition.

DELETE FROM *table-name* **WHERE** *condition*;

Removes all rows fulfilling the condition from the table.

4.9   Modeling

The way our example table from Figure 4.23 is designed results in lots of duplicate data—the director's name is stored anew for each row, and two directors with the same name cannot be distinguished. The situation worsens when we

want to store the cast of each movie. In other words, the way we modeled our data can be improved. *Entity-Relationship* (ER) diagrams are a tool to find good representations for data.

Definition 4.29 (Entity-Relationship Diagram). *Rectangles denote* **entities** *(tables), and diamonds with edges to entities indicate* **relations** *between those entities. On such an edge, the number 1 or the letter n denotes whether the corresponding entity takes part once or arbitrarily many times in the relation. Entities and relations can have* **attributes** *(columns) with a name, drawn as ellipses. Italicised attributes are* **key attributes** *which must be unique for each such entity.*
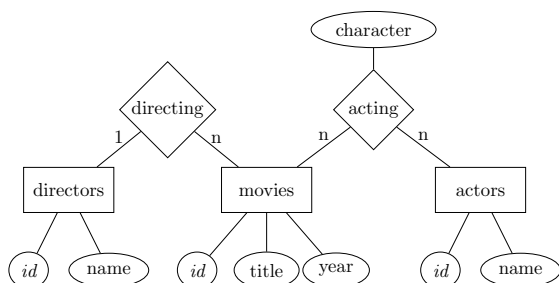

Figure 4.30: Model for a movie database. Movies and directors are in a 1-to-n relation: Each movie is directed by 1 director, and a director may work on many movies. Movies and actors are in a n-to-n relation, which has an additional attribute: An actor may appear in many movies, and each appearance is associated with a character in that movie, played by that actor.

Remarks:

- It is standard practice to assign a so-called *key attribute*, often named *id*, to every entity.

- What do ER diagrams have to do with SQL? Primarily, ER diagrams are for conceptually modeling the kind of data and relations one wishes to store. They can be translated into databases. Each entity corresponds to a table with the corresponding attributes as columns. An n-to-n relation is represented by a table with columns for each attribute, and a column for the key attribute of each entity in the relation.

- A close relative of the ER diagram is the Unified Modeling Language (UML). UML is used to represent the tables of a database (or classes of object oriented software) accurately, with detailed information, e.g. fields.

actor			acting		
id	name		actor_id	character	movie_id
1	Harrison Ford		1	Indy	2
2	Tom Cruise		2	Ray Ferrier	3
⋮			⋮		

Figure 4.31: The actor table and a table capturing the acting relation.

Remarks:

- The same scheme can be used for 1-to-1 and 1-to-n relations. However, one may also include the relation in the table storing the entity on the 1-side.

directors		movies			
id	name	id	title	year	director_id
1	Sidney Lumet	1	12 Angry Men	1957	1
2	Steven Spielberg	2	Raiders of the Lost Ark	1981	2
3	Harold P. Warren	3	War of the Worlds	2005	2
⋮		4	Manos: The Hands of Fate	1966	3
		⋮			

Figure 4.32: The movie and director tables using the new database layout. The director table simply maps ids to director names. Since the directing relationship is 1-to-n, it can be represented by adding a column to the movies table that stores the director for each movie.

Remarks:

- Similarly, a 1-to-1 relation can be turned into an attribute of one of the entities.

- Tables dedicated to capturing relations are often called *join* tables.

4.10 Joins

How can we access the data, which is now scattered across multiple tables?

```
SELECT movie.title, director.name AS director, movie.year
FROM movie
INNER JOIN director ON movie.director_id = director.id;
```

Listing 4.33: Example query that returns the table depicted in Figure 4.34.

SELECT . . .
   FROM *left-table* **INNER JOIN** *right-table* **ON** *condition*;
   Returns all rows that can be formed from a row in the left-table and a
   row in the right-table that satisfy the specified condition.

movie.title	director	movie.year
12 Angry Men	Sidney Lumet	1957
Raiders of the Lost Ark	Steven Spielberg	1981
War of the Worlds	Steven Spielberg	2005
Manos: The Hands of Fate	Harold P. Warren	1966

⋮

Figure 4.34: The result returned by the query in Listing 4.33.

Remarks:

- In a query, one can create aliases for field and table names using the
  AS keyword, see Listing 4.33.

- The result of a JOIN clause can be ordered, fields can be aggregated
  and grouped, and conditions can be added using WHERE clauses.

- For example, we can combine joins and aggregations to answer our
  initial question of which movie has the largest cast.

```
SELECT movie.title, COUNT(*) AS cast_size
FROM acting INNER JOIN movie ON acting.movie_id = movie.id
GROUP BY movie.id ORDER BY cast_size DESC LIMIT 10;
```

→ notebook

Listing 4.35: Finding the 10 movies with the largest cast.

Remarks:

- The query from Listing 4.35 uses a LIMIT clause to return only the
  ten first entries of the sorted results.

- An INNER JOIN where the condition is TRUE returns the Cartesian
  product of both tables. This special case can also be obtained with
  the CROSS JOIN clause.

- An inner join will only return those rows of one table that have a
  matching row (that satisfies the condition) in the other table. For
  example, in Listing 4.33, a director with id 5 would not appear in the
  result if there are no movies which have director_id=5.

- If you want unmatched rows to appear in the result, you need to use
  an OUTER JOIN.

```
SELECT movie.title, director.name AS director, movie.year
FROM movie
RIGHT OUTER JOIN director ON movie.director_id = director.id;
```

→ notebook

Listing 4.36: Example query that returns the table depicted in Figure 4.37.

movie.title	director	movie.year
12 Angry Men	Sidney Lumet	1957
Raiders of the Lost Ark	Steven Spielberg	1981
War of the Worlds	Steven Spielberg	2005
Manos: The Hands of Fate	Harold P. Warren	1966
NULL	Jon Doe	*NULL*

⋮

Figure 4.37: The result returned by the query in Algorithm 4.36. The right outer
join includes all rows from the inner join (see Figure 4.34) and, additionally, all
entries from the directors table for which there is no matching entry in the
movies table. In our example, "director" Jon Doe has not directed any movies,
hence the movie title and year column are filled with *NULL* values.

SELECT . . .
   FROM *left-table* **LEFT|RIGHT|FULL OUTER JOIN** *right-table*
   ON *condition*;
   Returns all rows from the inner join. In addition, a LEFT or RIGHT
   OUTER JOIN also returns all rows from the left or right table that have no

matching row on the opposite table, respectively. The fields in unmatched rows that cannot be filled from the other table are filled with NULL values. A FULL OUTER JOIN returns both of the above.

Remarks:

- A LEFT OUTER JOIN in Listing 4.36 would include the movies with no director instead of the directors who have not directed any movie.

- Queries may use more than one JOIN clause.

```
SELECT movie.title
FROM actor INNER JOIN acting
ON acting.actor_id = actor.id AND actor.name = 'Ford, Harrison'
RIGHT OUTER JOIN movie ON acting.movie_id = movie.id
WHERE acting.actor_id IS NULL;
```
→ notebook

Listing 4.38: Finding all movies that Harrison Ford did not appear in.

Remarks:

- The conditions for the first join in Listing 4.38 ensure that only movies with Harrison Ford are taken into account for the second OUTER JOIN. That second join in turn delivers all movies that cannot be matched, yielding a NULL entry for the actor_id for movies without Harrison Ford.

4.11 Keys & Constraints

What is stopping us from inserting a row in the acting table that contains an actor_id or a movie_id that does not exist? Or from creating a director with a duplicate id?

Definition 4.39 (Key). *In a table, a column (or set of columns) is a **unique key** if the corresponding values uniquely identify the rows within the table. The **primary key** of a table is a designated unique key. A **foreign key** is a column (or set of columns) that references the primary key of another table.*

Remarks:

- SQL databases can automatically enforce these constraints. For example, a row containing a foreign key can only be inserted if it references an existing primary key. Vice versa, a row may only be removed if its primary key is not referenced by any foreign key.

ALTER TABLE *table*
    ADD CONSTRAINT UNIQUE (*field-name,...***);**
    Any two rows must differ in at least one of the specified fields.
→ notebook

ALTER TABLE *table* **ADD PRIMARY KEY (***field-name,...***);**
    Sets the specified fields as the primary key for the table. Any two rows must differ in at least one of the specified fields. The entries in these fields must not be NULL.

ALTER TABLE *left-table* **ADD FOREIGN KEY (***field-name,...***)**
    REFERENCES right-table;
    Ensures that the values in the specified fields in the left table are the primary key of a row in the right table.

Remarks:

- Constraints for new tables can also be set using CREATE TABLE.

- Other ALTER TABLE queries add different constraints (e.g., checking that an integer field contains only certain values), remove constraints, and change the name, type or default value of fields.

- To ensure that checking constraints and searching for data is fast, database systems rely on *index* data structures.

4.12 Indexing

Definition 4.40 (Index). *In the database context, an **index** is a data structure that speeds up searching for rows with specific values.*

Remarks:

- Without an index data structure, rows with a specific value can only be found by scanning through the whole table.

- Earlier in the chapter you learned how hash tables can retrieve the row associated with a key quickly. Many database systems implement hash tables as one possible index data structure.

```
CREATE INDEX directorid ON director USING HASH (id);
```

Listing 4.41: Adding a hash table index to our database.

Remarks:

- The director associated with a movie is now found quickly when performing a join.

- Some database systems automatically create index data structures to speed up queries that involve frequently used fields.

- Index data structures have a name—"directorid" in Listing 4.41. This is for referencing it later, e.g., if one decides to delete the index.

- Hash tables scatter the data across the storage (volatile or persistent), and it is likely that every access incurs overhead. Many database queries require scanning through ranges of the data sequentially. For example, when searching the movies from 2000–2005. Thus, accessing supposedly closeby rows requires accessing items at many different places.

- B+ trees are a data structure designed to minimize the amount of I/O operations for both searching and scanning.

```
CREATE INDEX movieyear ON movies USING BTREE (year);
```

Listing 4.42: Adding a B+ tree index to our database.

Definition 4.43 (B+ Tree). *A **B+ Tree** of order b is a rooted search tree mapping **keys** to **rows**. B+ trees are **balanced**, i.e., all leaf nodes are at the same depth.*

Every non-leaf node has between $\lfloor b/2 \rfloor$ and b children. A non-leaf node v with k children contains exactly $k - 1$ keys, in sorted order. The ith key of a non-leaf node v is identical to the smallest key in the subtree rooted at v's $(i + 1)$st child.

Leaf nodes contain all keys inserted into the tree, together with a pointer which points to the row associated with that key. Every leaf has pointers to at least $\lfloor (b - 1)/2 \rfloor$ and at most $b - 1$ table rows. Additionally, every leaf w has a pointer which points to its next sibling w'.


Figure 4.44: Example B+ tree of order $b = 4$.

Remarks:

- The root node is a special case—it may have as little as 2 children if it is not a leaf. If it is a leaf, it may have only one pointer to a table row.

- The order b is sometimes called *branching factor*. To reduce the number of necessary I/O operations, b is chosen so that all data necessary to store a node is the size of one page on disk.

- Finding the row for some key k in a B+ tree works similar to a binary search tree.

- When inserting a key k, we have to check if the leaf v that should contain k is already full (i.e. already contains $b - 1$ keys). In that case v, and possibly predecessors of v that contain too many keys, need to be split using B+SplitUp.

- If the root node is split into two nodes v, v', then a new root r containing key k and v and v' as children is created, and the recursion stops.

- Inserting a key k is now performed by first making room using B+SplitUp if necessary, and then inserting k at the leaf.

```
1  def B+Insert(self, k, r):
2      # Search for k to find the leaf v at which k must be inserted
3      v = self.search(k)
4      if len(v) == b-1:
5          B+SplitUp(v, k)
6          Replace child of key k with row r in node v
7      else:
8          Insert key k with row r into node v
```

Algorithm 4.45: B+Insert

Remarks:

- Vice versa, when deleting a key, nodes with too few keys need to be filled up or removed from the tree.

- The properties of B+ trees ensure that every node has at least one sibling. Thus, the merge operation always has two nodes to work with.

- If no keys can be "borrowed" from a sibling, the merge may propagate until the last two children of the root node are merged into one node. In that case the root node is replaced by the merged node, decreasing the height of the tree by 1.

```
1  def B+Delete(self, k):
2      # Search for k to find the leaf v containing k
3      v = self.search(k)
4      Remove k from v
```

```
5    if len(v) < (b-1)/2:
6      B+MergeUp(v)
```

Algorithm 4.46: B+Delete

Remarks:

- The height of a B+ tree is changed only when inserting a new or removing an old root node. Therefore, all leaf nodes are always at the same depth, thus ensuring the *balanced* property.

- A B+ tree containing n keys has height at most $O(\log_b n)$.

- It may happen that many nodes contain as little as $b/2$ keys, wasting memory and I/O operations. B* trees ensure that nodes contain at least $\frac{2}{3}b$ keys by cleverly "trading" entries with neighboring nodes when they contain too many or too few keys.

4.13 Transactions

Definition 4.47 (Transaction). *A database **transaction** is a sequence of statements that is executed atomically, i.e. the transaction appears to be instantaneous to an observer.*

Remarks:

- Why would we need transactions? Consider a bank managing customer's accounts using a database system. Alice wants to calculate the liquid assets, and Bob wants to make a money transfer:

```
-- Alice's statement:
SELECT SUM(balance) FROM accounts;
-- Bob's statements:
UPDATE accounts SET balance=balance-100 WHERE customer = 'Bob';
UPDATE accounts SET balance=balance+100 WHERE customer = 'Jim';
```

Listing 4.48: Concurrency issues in databases.

Remarks:

- Assuming that the database system uses multiple threads or processes to process queries, Alice's query may be CHF 100 short.

- To execute the queries atomically, both Alice and Bob can use transactions.

BEGIN TRANSACTION; *statement₁*; . . . ; END TRANSACTION;
Executes the statements atomically.

Remarks:

- One way to implement transactions is to keep track of all fields read from and written to (the read- and write-set, respectively). Then, before a transaction ends, the database system checks whether another transaction wrote to any value in the read-set. If the read-set is unchanged, the write-set can be applied atomically, e.g., by using a global lock.

- SQL offers different so-called isolation levels. The isolation level defines when writes of one transaction become visible to others. The above technique implements the *repeatable reads* level, ensuring that read values were committed before and are not written by another transaction.

- Consider some transaction A that selects all years between 1999 and 2004. What happens if another transaction B concurrently inserts an entry for the year 2000? In the *repeatable reads* isolation level, A may not see B's data if B's insert is scheduled *after* A read all other entries for the year 2000, and A would still be allowed to finish. *Repeatable reads* do not ensure atomicity . . .

- The highest isolation level is called *serializable*. This level ensures that the transactions behave "as if they were executed in some sequential order", possibly at the cost of low concurrency.

4.14 Programming with Databases

How do you write an application that relies on a SQL database to store data? Should you construct the necessary SQL statements by manipulating strings, send them to the SQL server, and then parse the result?

Remarks:

- Writing such a SQL client is one possibility, but this is error-prone: The compiler used for the application will not be able to detect errors made in the SQL statements. Moreover, the declarative SQL most likely does not mix well with the programming language chosen for the application.

- One way to mitigate these issues in object oriented programming languages is object/relational mapping.

Definition 4.49 (Object/Relational Mapping). ***Object/Relational Mapping (ORM)*** *is a design pattern used in object oriented programming to store objects in and retrieve them from relational (SQL) databases.*

Remarks:

- In the simplest case, an ORM simply maps a class to a table. An object then corresponds to a row, and the object's attributes correspond to the row's fields.

- The ORM takes care of storing and retrieving object in the database and performs type conversions where necessary. It provides object oriented abstractions for database queries involving WHERE and other clauses. ORMs also remove boilerplate code, i.e., setting up the SQL connection, error handling, data conversion, etc.

- Popular ORMs include SQLAlchemy for Python, ActiveRecord for Ruby, Hibernate for Java, and the Entity Framework for .NET.

- The ORM needs to know how it should translate between objects and rows. For that, many ORM implementations allow to specify the database layout using object oriented methods. Many ORM mappers also support creating the database using the object oriented specification.

- Some concepts from object oriented programming are difficult to model with database concepts, and vice versa. The problems arising from combining these two paradigms are called the *Object-relational impedance mismatch*.

Chapter Notes

Dictionaries based on search trees are useful for providing additional operations such as nearest neighbor queries or range queries, where we want to find all keys in a certain range. Binary search trees were first published by three independent groups in 1960 and 1962 (for references, see Knuth [18]). The first instance of a self-balancing search tree that guarantees logarithmic cost for insert/search/delete is the AVL-tree, named so after its inventors Adelson-Velski and Landis [1]. For multidimensional keys, e.g. geometric data or images, there are specialized tree structures such as kd-trees [3] or BK-trees [5].

Hashing has a long history and was initially used and validated based on empirical results. One of the first publications was Peterson's 1957 article [20] where he defined an idealized version of probing and empirically analyzed linear probing. Universal hashing was introduced two decades later by Carter and Wegman in 1979 [6]. Perfect static hashing was invented in 1984 by Fredman et al. [13] and is sometimes also referred to as FKS hashing after its inventors. Its dynamization by Dietzfelbinger et al. took another decade until 1994 [12]. A comprehensive study on perfect hashing by Czech et al. was compiled in 1997 [11]. Cuckoo hashing is a comparatively recent algorithm; it was introduced by Pagh and Rodler in 2001 [19].

There have been a number of other developments regarding hashing since the late 1970s; for an overview, see Knuth [18], in particular the section on History at the end of chapter 6.4. For a neat visualization of hashing with probing, see [14] online.

The power of two choices paradigm has found widespread application and analysis in load balancing scenarios. It was initially studied from the perspective of a balls-into-bins game where we want to minimize the maximum number of balls in any bin, and to do this we can pick two random bins and put the next ball into the least full of the two bins. Richa et al. [21] compiled an excellent survey on the earliest sources and numerous applications of this paradigm.

In 1970, Edgar F. Codd proposed the relational database model [9] while working at IBM research. Later in the 70s, another group at IBM developed SQL's predecessor SEQUEL (Structured English QUEry Language) [7]. After being renamed SQL due to trademark issues, it was standardized by the ISO in 1987 and later revised [15]. Other companies started developing relational database systems, and nowadays there are many SQL databases implementing different feature sets to choose from.

Around the same time, ER diagrams were conceived as a modeling tool [4, 8]. The Unified Modeling Language (UML), first standardized by the ISO in 1995 [16] and revised in 2012, also includes diagrams that model databases.

B Trees were invented in 1970 [2] for use in file systems. Many variants were studied, among them B* Trees [17], in which at most $1/3$ of the memory is unused instead of $1/2$ for B Trees. People soon realized that (also for file systems) scanning subsequent rows is an important operation. B+ Trees require at most one I/O operation to find the next element, cf. [17, 10].

Techniques from database systems can also be found in other areas of computer science. Transactions as a parallel programming model have been adotped for other programming languages under the term *transactional memory*. Ideas developed to ensure that database transactions appear atomic w.r.t. writing data to disk were adopted by general purpose file systems under the name *journaling*.

This chapter was written in collaboration with Georg Bachmeier and Jochen Seidel.

Bibliography

[1] M Adelson-Velskii and Evgenii Mikhailovich Landis. An Algorithm for the Organization of Information. *Doklady Akademii Nauk USSR*, 146(2):263–266, 1962.

[2] R. Bayer and E. McCreight. Organization and maintenance of large ordered indices. In *Proceedings of the 1970 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET '70, 1970.

[3] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975.

[4] A. P. G. Brown. Modelling a real world system and designing a schema to represent it. In *IFIP TC-2 Special Working Conference on Data Base Description*, 1975.

[5] W. A. Burkhard and R. M. Keller. Some approaches to best-match file searching. *Commun. ACM*, 16(4):230–236, 1973.

[6] J.Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143 – 154, 1979.

[7] Donald D. Chamberlin and Raymond F. Boyce. Sequel: A structured english query language. In *Proceedings of the 1974 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET '74. ACM, 1974.

[8] Peter Pin-Shan Chen. The entity-relationship model—toward a unified view of data. *ACM Trans. Database Syst.*, 1976.

[9] E. F. Codd. A relational model of data for large shared data banks. *Commun. ACM*, 1970.

[10] Douglas Comer. The ubiquitous B-Tree. *ACM Comput. Surv.*, 1979.

[11] Zbigniew J. Czech, George Havas, and Bohdan S. Majewski. Perfect hashing. *Theoretical Computer Science*, 182(1 - 2):1 – 143, 1997.

[12] Martin Dietzfelbinger, Anna Karlin, Kurt Mehlhorn, Friedhelm Meyer auf der Heide, Hans Rohnert, and Robert E. Tarjan. Dynamic perfect hashing: Upper and lower bounds. *SIAM J. Comput.*, 23(4):738–761, 1994.

[13] Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with 0(1) worst case access time. *J. ACM*, 31(3):538–544, 1984.

[14] David Galles. Closed hashing. `https://www.cs.usfca.edu/~galles/visualization/ClosedHash.html`. Accessed: 2017-03-29.

[15] International Organization for Standardization. Information technology – Database languages – SQL – part 1: Framework (SQL/Framework), 2011. ISO/IEC 9075-1.

[16] International Organization for Standardization. Information technology – Object Management Group Unified Modeling Language (OMG UML) – Part 1: Infrastructure, 2012. ISO/IEC 19505-1.

[17] Donald E. Knuth. *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., 1973.

[18] Donald E. Knuth. *The Art of Computer Programming, Volume 3: (2Nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1998.

[19] Rasmus Pagh and Flemming Friche Rodler. *Algorithms — ESA 2001: 9th Annual European Symposium Århus, Denmark, August 28–31, 2001 Proceedings*, chapter Cuckoo Hashing, pages 121–133. Springer Berlin Heidelberg, 2001.

[20] W. W. Peterson. Addressing for random-access storage. *IBM J. Res. Dev.*, 1(2):130–146, 1957.

[21] Andrea W Richa, M Mitzenmacher, and R Sitaraman. The power of two random choices: A survey of techniques and results. *Combinatorial Optimization*, 9:255–304, 2001.

Chapter 5

Machine Learning

So far, we told the computer exactly what to do: every problem was solved by a specific algorithm. However, in the real world, we might have to deal with messy data in order to understand its underlying function. In other words, it may be difficult to separate function from noise. Maybe a computer can do part of the job, and learn some of the parameters of a function? Welcome to machine learning!

5.1 Linear Regression

Definition 5.1 (Dataset, Input, Output). *A **dataset** D is a set of n tuples (x, y) sampled from an unknown function*

$$f : \ x \mapsto y$$

*We call $x \in X$ an **input** and $y \in Y$ the corresponding **output** of f.*

Remarks:

- We want to learn a function $\hat{f}$ such that $\hat{f}(x) \approx f(x)$.

- Since we learn $\hat{f}$ from the dataset D, we may write $\hat{f}_D$.

- For example, if the dataset consists of points on a line, we choose $\hat{f}(x) = w_1 x + w_0$ and determine the parameters $w_i \in \mathbb{R}$.

- If the points in D do not line up perfectly, our approximation function $\hat{f}$ will have some error. Even if f is truly linear, there can still be some random noise that introduces error, e.g., measurement error. → notebook

Definition 5.2 (Approximation Error). *The **approximation error** $Err(x)$ denotes the deviation of $\hat{f}$ from the unknown function f at some input x:*

$$Err(x) = f(x) - \hat{f}(x).$$

Remarks:

- For our linear function this amounts to $Err(x) = y - (w_1 x + w_0)$.

- We want to minimize this error over the entire dataset D. Hence, we choose $\hat{f}$ according to an objective function as follows.

Definition 5.3 (Squared Error Loss). *The **loss function** is used to determine the real-valued parameters $\boldsymbol{w} = (w_0, w_1, ...)^T$ according to the dataset D. There are several options, the most common being the **squared error loss** function:*

$$L(\hat{f}, D) = \sum_{(x,y) \in D} Err(x)^2.$$

Remarks:

- By squaring the error, we ensure each term is positive. Squaring also weighs large errors more highly.

- Another natural choice for a loss function is the absolute error: $\rightarrow$ notebook

$$L_{abs}(\hat{f}, D) = \sum_{(x,y) \in D} |Err(x)|.$$

- However, the squared error loss is often preferred as it has both a closed-form solution and is differentiable. How do we find a solution for our linear function?

Lemma 5.4. *Let $\bar{x} = \frac{1}{n} \sum_D x$ and $\bar{y} = \frac{1}{n} \sum_D y$ be the average input and output of the dataset D. For a linear function $\hat{f}(x) = w_1 x + w_0$, the squared error loss is minimal for*

$$w_1^* = \frac{\sum_{(x,y) \in D}(y - \bar{y})(x - \bar{x})}{\sum_{(x,y) \in D}(x - \bar{x})^2}, \; w_0^* = \bar{y} - w_1^* \bar{x}.$$

*We call these weights the **ordinary least-square** (OLS) estimates, as they minimize the squared error loss.*

Proof. For our linear function, the squared error loss amounts to

$$L(\hat{f}, D) = \sum_{(x,y) \in D} (y - (w_1 x + w_0))^2.$$

We find the minimum loss by differentiating $L(\hat{f}, D)$ with respect to $\boldsymbol{w}$:

$$\frac{\partial L}{\partial w_0} = \sum_{(x,y) \in D} -2 \left(y - (w_1 x + w_0) \right) \stackrel{!}{=} 0$$

$$\iff \sum_{(x,y) \in D} (y - w_1 x - w_0) = 0$$

$$\iff \sum_{(x,y) \in D} (y - w_1 x) = n w_0$$

$$\iff w_0 = \frac{1}{n} \sum_{(x,y) \in D} y - w_1 \frac{1}{n} \sum_{(x,y) \in D} x$$

$$\iff w_0 = \bar{y} - w_1 \bar{x}$$

$$\frac{\partial L}{\partial w_1} = \sum_{(x,y) \in D} -2x \left(y - (w_1 x + w_0) \right) \stackrel{!}{=} 0$$

$$\iff \sum_{(x,y) \in D} x(y - w_1 x - \bar{y} + w_1 \bar{x}) = 0$$

$$\iff \sum_{(x,y) \in D} x(y - \bar{y}) = w_1 \sum_{(x,y) \in D} x(x - \bar{x}).$$

Note that $\sum_D y = n\bar{y} = \sum_D \bar{y}$, so $\sum_D \bar{x}(y - \bar{y}) = 0$ and similarly we get $\sum_D \bar{x}(x - \bar{x}) = 0$. After subtracting "0" from both sides, we obtain

$$\sum_{(x,y) \in D} (x - \bar{x})(y - \bar{y}) = w_1 \sum_{(x,y) \in D} (x - \bar{x})(x - \bar{x})$$

$$\iff w_1 = \frac{\sum(y - \bar{y})(x - \bar{x})}{\sum(x - \bar{x})^2}.$$

$\square$

Remarks:

- This also works if the input is not a single scalar x but a whole vector $\boldsymbol{x}$. This is known as linear regression.

Definition 5.5 (Linear Regression, Features, Weights). *With **linear regression**, we search for a function $\hat{f}$ of the form*

$$\hat{f}(\boldsymbol{x}) = \sum_{i=1}^{d-1} w_i x_i + w_0,$$

*where the $x_i \in \mathbb{R}$ are the **features** of the input. The parameters $w_i \in \mathbb{R}$ are called the **weights** and need to be determined. We can write this in vector notation as*

$$\hat{f}(\boldsymbol{x}) = \boldsymbol{w}^T \boldsymbol{x},$$

where $\boldsymbol{x} = (1, x_1, x_2, ..., x_{d-1})^T$ with an additional 1 to incorporate the weight w_0.

Remarks:

- Let us find the weights $\boldsymbol{w}$. In order to describe the closed-form solution concisely, we use a matrix notation where $\mathbf{X}$ is a matrix of the input features (the so-called *design matrix*). $\mathbf{X}$ has n rows, each row represents a feature vector $\boldsymbol{x} = (1, x_1, x_2, ..., x_{d-1})$. The outputs are given as a vector $\boldsymbol{y}$ of length n, where each value has a corresponding row in $\mathbf{X}$.

Theorem 5.6. *The ordinary least-square (OLS) estimates for the weight parameters $\boldsymbol{w}$ of a linear regression model are given by*

$$\boldsymbol{w}^* = \left(\mathbf{X}^T \mathbf{X} \right)^{-1} \mathbf{X}^T \boldsymbol{y}.$$

Proof. The squared error loss is $L(\hat{f}, D) = \sum_{(\boldsymbol{x},y) \in D} \left(y - \boldsymbol{w}^T \boldsymbol{x}\right)^2$ which can be rewritten in matrix form as

$$L(\hat{f}, D) = (\boldsymbol{y} - \mathbf{X}\boldsymbol{w})^T (\boldsymbol{y} - \mathbf{X}\boldsymbol{w}).$$

Again, we can differentiate with respect to $\boldsymbol{w}$ to find the optimal weights:

$$\frac{\partial L}{\partial \boldsymbol{w}} = -\left(\mathbf{X}^T(\boldsymbol{y} - \mathbf{X}\boldsymbol{w})\right)^T - (\boldsymbol{y} - \mathbf{X}\boldsymbol{w})^T \mathbf{X} \overset{!}{=} \boldsymbol{0}^T$$
$$\iff -\mathbf{X}^T(\boldsymbol{y} - \mathbf{X}\boldsymbol{w}) - \mathbf{X}^T(\boldsymbol{y} - \mathbf{X}\boldsymbol{w}) = \boldsymbol{0}$$
$$\iff \mathbf{X}^T(\boldsymbol{y} - \mathbf{X}\boldsymbol{w}) = \boldsymbol{0}$$
$$\iff \mathbf{X}^T \boldsymbol{y} = \left(\mathbf{X}^T \mathbf{X}\right) \boldsymbol{w}$$
$$\iff \boldsymbol{w} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \boldsymbol{y}.$$

$\square$

Remarks:

- We must be careful when differentiating in matrix form, as we are differentiating sums.

- We assume in this proof that $\mathbf{X}^T \mathbf{X}$ is invertible. This is for example the case when $\mathbf{X}$ has full column rank, but might not be the case in general.

- We could have derived the result in Lemma 5.4 in matrix form as well. To see that the results are the same, just expand $(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \boldsymbol{y}$. Alternatively, if you consider the training data to be normalized (i.e. $\bar{x} = \bar{y} = 0$), then $w_0 = 0$ and $w_1 = \frac{\sum xy}{\sum x^2} = (\sum xx)^{-1} \sum xy$.

- What if the relation between $\boldsymbol{x}$ and y is not just linear?

5.2 Feature Modeling

Definition 5.7 (Feature). *A **feature** of the input can be any real-valued term depending on the input variables.*

Remarks:

- Note that we previously used the feature vector $\boldsymbol{x} = (1, x_1, x_2, ..., x_{d-1})^T$. However, a feature can be more complex. For example, if we know in advance that the quantity $\sin x_1 \sqrt{x_2}$ is a good term to approximate $f(\boldsymbol{x})$, then we can include this term as a feature in our linear regression model. Also splines, radial basis functions or wavelets work out of the box.

- When modeling a problem, we often start with an "educated guess" about which family of functions $\mathcal{F}$ is well-suited to model the unknown function f. We then restrict ourselves to find the best $\hat{f} \in \mathcal{F}$.

Definition 5.8 (Model). *We call the family of functions $\mathcal{F}$ chosen to approximate f a **model**. The function $\hat{f} \in \mathcal{F}$ is found by fitting the parameters of the model to "best" represent the dataset D.*

Remarks:

- For instance, we might want to restrict $\hat{f} \in \mathcal{F}$ to polynomials (of $\rightarrow$ notebook degree m), yielding

$$\hat{f}(x) = w_0 + w_1 x + w_2 x^2 + w_3 x^3 + ... + w_m x^m.$$

- This is called *polynomial regression*, even though it is just a special case of linear regression (and is solved the same way).

- Note that the linear regression method can handle any model $\mathcal{F}$, as long as the (unknown) parameters w_i are linear coefficients.

- For multi-dimensional input we can add all combinations of the variables up to the m^{th} power. For polynomial regression with degree m and input dimension d, this gives $\left(\binom{d}{m}\right) = \binom{d+m-1}{m}$ features! E.g. for $m = 2$, $d = 3$, we have the 6 features $\{1, x_1, x_2, x_1^2, x_2^2, x_1 x_2\}$.

- What if some input variables are not continuous? For example, we might have a categorical input variable, such as a city name. Should we encode the city as a single variable taking values 1, 2 and 3? This would establish an inherent ordering and scaling between the cities, which would be unreasonable in many cases. One way to deal with such inputs is to use so-called one-hot encoding.

Definition 5.9 (One-Hot Encoding). *A **one-hot encoding** of a categorical input variable taking k values is a vector representation of length k consisting of 0's and a single 1 indicating the corresponding category.*

Remarks:

- Note that using a one-hot encoding may increase the dimension of the model significantly as each category gets its own coefficient w_j.

- See Figure 5.10 for an example.

x_i
London
Budapest
Zurich
-
London

x_{i1}	x_{i2}	x_{i3}
0	1	0
1	0	0
0	0	1
0	0	0
0	1	0

Figure 5.10: One-hot encoding of a categorical variable that can take 3 values.

5.3 Generalization & Overfitting

What if we include more and more features? For instance, we can add more input features, higher degree terms and other engineered features. Eventually we might have more features than data points, and a linear model will be able to fit the training data perfectly.

Lemma 5.11. *Given $\mathbf{X} \in \mathbb{R}^{n \times d}$ with $d \geq n$, there is at least one solution $\boldsymbol{w}$ to $\mathbf{X}\boldsymbol{w} = \boldsymbol{y}$ for all $\boldsymbol{y} \in \mathbb{R}^n$, if and only if $\mathbf{X}$ has rank n.*

Proof. This is a standard result from Linear Algebra. The proof goes along the lines of: $\mathbf{X}$ has rank $n \iff$ there are n linearly independent columns of $\mathbf{X}$, but n linearly independent columns in $\mathbb{R}^{n \times d}$ form a basis of $\mathbb{R}^{n \times n}$ so (setting w for all other columns to 0) $\mathbf{X}\boldsymbol{w} = \boldsymbol{y}$ has a unique solution for all $\boldsymbol{y} \in \mathbb{R}^n$. $\quad\square$

Remarks:

- For example powers of a feature, $\{1, x_i, x_i^2, x_i^3, ...\}$ are linearly independent, so polynomial regression with high enough degree will always be able to fit training data perfectly. This is called *polynomial interpolation*.

- What is the problem with adding too many features? Overfitting. The fitted model will not generalize well to unseen data. Ultimately, our goal is to minimize the expected error over *all possible* data.

Definition 5.12 (Expected Loss). *We assume all our data (including any unseen data) comes from some unknown distribution $(\boldsymbol{x}, y) \sim P(X, Y)$, where X denotes the entire input space and Y the output space. Then the* ***expected loss*** *is defined as*

$$L(\hat{f}) = \mathbb{E}_{\boldsymbol{x},y}\left[L(\hat{f}, \boldsymbol{x})\right] = \int L(\hat{f}, \boldsymbol{x})\, dP(\boldsymbol{x}, y),$$

where $L(\hat{f}, \boldsymbol{x})$ is the loss incurred by $\hat{f}$ at data point $(\boldsymbol{x}, y)$.

Remarks:

- Expected loss is also referred to as *risk*.

- Unfortunately, we cannot calculate the expected loss as we do not know the probability distribution P. Thus we also cannot directly minimize it.

- So far, we have instead minimized the loss on our dataset D. This is called the empirical loss (or empirical risk).

Definition 5.13 (Empirical Loss). *We estimate the expected loss by the empirical loss on a dataset D, given by*

$$\hat{L}_D(\hat{f}) = \frac{1}{n} \sum_{(\boldsymbol{x},y) \in D} L(\hat{f}, \boldsymbol{x}).$$

Remarks:

- The empirical loss is exactly the (normalized) total loss from Definition 5.3.

- By the law of large numbers $\hat{L}_D(\hat{f}) \to L(\hat{f})$ for any fixed $\hat{f}$ almost surely as $n \to \infty$. Therefore the more data we have, the closer the empirical loss will be to the expected loss and so, the closer $\hat{f}$ can be to the true function f by minimizing the empirical loss.

- But how do we know how well our $\hat{f}$ performs on the rest of the domain X?

- We could take the empirical loss as our estimate. Unfortunately, since we fit our model to this data, it will inherently underestimate the expected loss.

- We could find new data for evaluation, but we usually just have one dataset to work with. However, we can sample a subset D_t from our dataset D and only train our model with this subset, while reserving the rest for evaluation.

Definition 5.14 (Train-Evaluation Split). *Partitioning a dataset D into two disjoint subsets D_t and D_e (typically 80% to 20%) is called a* ***train-evaluation split***.

Remarks:

- This is also called a *train-test* split. For ease of notation we will stick with evaluation.

Definition 5.15 (Training Loss, Evaluation Loss). *We define the* ***training loss*** *as*

$$\hat{L}_t(\hat{f}) = \frac{1}{|D_t|} \sum_{(\boldsymbol{x},y) \in D_t} L(\hat{f}, \boldsymbol{x}).$$

Similarly, we define the ***evaluation loss*** *as*

$$\hat{L}_e(\hat{f}) = \frac{1}{|D_e|} \sum_{(\boldsymbol{x},y) \in D_e} L(\hat{f}, \boldsymbol{x}).$$

Remarks:

- Note that $\hat{f}$ depends on the training data, so $\hat{f} = \hat{f}_{D_t}$.

- We can use the evaluation dataset D_e to estimate how well the function $\hat{f}_{D_t}$ generalizes to new data.

Definition 5.16 (Overfitting, Underfitting). *A model is* ***overfitting*** *when it fits* $\rightarrow$ notebook *the training dataset D_t too well, learning random patterns/noise that will not be present in new unseen data D_e. The model will not generalize well. Conversely, a model is* ***underfitting*** *when it is not expressive enough to approximate f. A more complex model $\mathcal{F}'$ should be tried.*

Remarks:

- $\hat{L}_e(\boldsymbol{w}) \gg \hat{L}_t(\boldsymbol{w})$ is a clear indication of overfitting.

- See Figure 5.17 for examples of overfitting and underfitting.

Figure 5.17: Examples of underfitting and overfitting

Remarks:

- There are several ways we can counter overfitting. We could try a simpler model, gather more training data, or introduce *regularization*.

- High training and evaluation errors could be a sign of underfitting. However, high errors could also indicate that the data is inherently noisy/random and cannot be fitted well.

- By repeating the train-evaluation split process multiple times we can get a more accurate evaluation of how well our model generalizes, and whether it is overfitting or underfitting. This is known as *cross validation*.

Definition 5.18 (Cross Validation). *In k-fold cross validation, we randomly partition D into k equal sized subsets. We train a model on the union of $k-1$ of these subsets. Then we evaluate our model on the last, withheld subset. This is repeated k times, with each subset used $k-1$ times as part of the training data and once as evaluation data. This gives k evaluation scores, which are averaged to produce a single evaluation metric.*

Remarks:

- There are other types of cross validation. For example in *Monte Carlo cross validation*, D_e is randomly sampled from D each time.

- Cross validation can also be used for *model selection*: We first do a train-evaluation split and then use cross validation on the training set D_t to select our model $\mathcal{F}$. The best model is then evaluated on the as yet unseen evaluation data.

- The mean and variance of the error across splits can tell us more about the sources of error in our model.

5.4 Bias-Variance Tradeoff

Definition 5.19 (Bias, Variance). *$\boldsymbol{Bias}$ is the expected error of a model, that is, how much the model $\mathcal{F}$ deviates from the target value on average. Formally,*

$$Bias^2[\mathcal{F}] = \mathbb{E}_{\boldsymbol{x}}\left[\mathbb{E}_{D,\boldsymbol{w}}\left[f(\boldsymbol{x}) - \hat{f}(\boldsymbol{x})\right]^2\right],$$

where the expectation is taken over $\boldsymbol{x}$ and D. The $\boldsymbol{variance}$ of a model is the variance in its predictions when training on different random datasets. Formally, with $\bar{f}(\boldsymbol{x}) = \mathbb{E}_{D,\boldsymbol{w}}\left[\hat{f}(\boldsymbol{x})\right]$:

$$Var[\mathcal{F}] = \mathbb{E}_{\boldsymbol{x},D,\boldsymbol{w}}\left[\left(\hat{f}(\boldsymbol{x}) - \bar{f}(\boldsymbol{x})\right)^2\right].$$

Remarks:

- If the bias is large, we know that our model $\mathcal{F}$ cannot approximate f well and we should consider a more expressive model, i.e., a more general family of functions $\mathcal{F}$. If on the other hand $f \in \mathcal{F}$, then the bias will be zero.

- Note that random noise will not contribute to the bias. On average (taking the expectation over D), the effects due to noise will cancel out. The bias really only covers systematic errors because our model is not able to match f exactly (even with the best weights). See Figure 5.20 (left).

- The variance tells us how sensitive our model is to the specific D. If the model is very sensitive, then $\boldsymbol{w}$ and ultimately the predictions will vary greatly depending on the dataset, in other words it will overfit. More noise will increase the variance. See Figure 5.20 (right).

- The simplest possible model is a constant, $\hat{f} = c$. This always has zero variance. Its predictions never change.


Figure 5.20: Fitting polynomials of increasing degree to the cosine function within the range $[-5, 5]$. The first order model (left) is too simple and has a high bias. On the other hand the 11^{th} order model (right) is too complex and has a high variance. The 6^{th} order model (middle) has low bias and low variance. In each iteration $f(x) = 8\cos(x)$, 25 inputs were sampled uniformly from $[-7, 7]$ and y was calculated with additional Gaussian noise. Predictions were then plotted on the interval $[-5, 5]$.

Theorem 5.21 (Bias-Variance Decomposition). *We can decompose the mean squared error (MSE) of a model into its squared bias and its variance:*

$$\text{MSE}(\mathcal{F}) = Bias^2[\mathcal{F}] + Var[\mathcal{F}]$$

where the mean squared error (or expected squared error) equals:

$$\text{MSE}(\mathcal{F}) = \mathbb{E}_{\boldsymbol{x}, D, \boldsymbol{w}}\left[\left(y - \hat{f}(\boldsymbol{x})\right)^2\right]$$

Remarks:

- In general our dataset will not cover the whole input space and the data will contain some noise. This means bias or variance or both will generally be greater than zero.

- However, there is an inherent tradeoff between bias and variance. A more complex model will be able to approximate f better, giving lower bias. But this also allows it to fit noise better leading to higher variance, since the noise it fits is random. On the other hand a simpler model will generally have higher bias and lower variance.


Figure 5.22: Fitting polynomials of increasing degree to the cosine function within the range $[-5, 5]$. On the left we see that as the degree (complexity) of the model increases, the bias decreases and the variance increases. The plot on the right shows the bias variance tradeoff frontier.

- If we have high variance, regularization is one way of reducing the total error. Regularization focuses on decreasing the variance at the potential expense of increasing the bias.

5.5 Regularization

How about including all the features we may possibly want, but charging a cost for each non-zero weight w_i? This should help us reduce overfitting to noise, but still allows us to fit f well. This is what Lasso does.

Definition 5.23 (Lasso Regression). *Lasso regression minimizes*

$$\min_{\boldsymbol{w}} \left\{ \frac{1}{n} \sum_{(\boldsymbol{x}, y) \in D} (y - \boldsymbol{w}^T \boldsymbol{x})^2 + \lambda \sum_{i=0}^{d-1} |w_i| \right\}$$

Remarks:

- The regularization parameter λ weighs the parameter cost versus the MSE loss. It is a so-called *hyperparameter* that can be tuned.

Definition 5.24 (Hyperparameter). *A **hyperparameter** is a parameter that controls the training process and whose value has to be chosen in advance.*

Remarks:

- The degree m in polynomial regression was also a hyperparameter.

- Cross validation can be used for hyperparameter tuning.

- For each hyperparameter you can define a set of values. Exhaustively iterating through all combinations of these values is called *grid search*.

Remarks:

- Lasso introduces bias into the regression solution by guiding the weights to be close to zero. This can reduce variance considerably relative to the OLS solution, see Figure 5.25.


Figure 5.25: Fitting a polynomial of high degree to a cubic function with and without regularization. We see that variance and total error drops when regularization is added.

Remarks:

- When using regularization, features should be normalized (subtract mean, divide by standard deviation). This ensures that the coefficient of a feature is not influenced by its magnitude.

- The target vector $\boldsymbol{y}$ should also be centered around 0, so that the intercept w_0 does not count towards the total cost.

- Ridge has the same objective, but Ridge uses the L^2 norm for the cost of the weights, whereas Lasso uses the L^1 norm. In other words, we replace $|w_i|$ with w_i^2.

- Ridge has a closed form solution, just like OLS. However Lasso does not, so we have to solve it differently, for example by *gradient descent*.

5.6 Gradient Descent

The OLS method had a closed form solution (Lemma 5.4 and Theorem 5.6), but often we do not have a closed form solution to our optimization problem. An alternative is to minimize the loss function using gradient based methods. If we can calculate the gradient of the loss function with respect to the weights, then we can perturb the weights in the right direction to decrease the loss. We can repeat this process until reaching a minimum.

Definition 5.26 (Gradient Descent). *Given a loss function $L(\hat{f}(\boldsymbol{w}), D)$, repeatedly perform the update*

$$w_j := w_j - \alpha \frac{\partial}{\partial w_j} L(\hat{f}, D)$$

*simultaneously for $j = 0, 1, ..., d$, where hyperparameter α is the **learning rate**.*

Remarks:

- The partial derivative tells us which weights to decrease or increase. If $\frac{\partial}{\partial w_j} L(\hat{f}, D)$ is positive, then the loss is increasing in w_j, so we decrease w_j to lower the loss.

- The gradient, i.e., the vector of partial derivatives, $\nabla L(\hat{f}(\boldsymbol{w}), D)$ points in the direction of steepest ascent, with the negation pointing in the direction of steepest descent.

- A (loss) function can have multiple minima, and gradient descent may well converge to a local minimum rather than finding the global minimum. The minimum reached depends on the initial starting point and the learning rate α. See Figure 5.27.


Figure 5.27: Gradient descent with different initialization points. On the (left) the initialization is favorable and we reach the global minimum of the polynomial function. However on the (right) gradient descent gets stuck in a local minimum.

Remarks:

- The *learning rate* α is a hyperparameter that controls the size of each update step. If the learning rate is too high, the algorithm might jump beyond the optimum $\boldsymbol{w}$; if it is too low the algorithm will be very slow to converge. See Figure 5.28 for examples. Often a decaying learning rate is used for efficiency at the beginning and accuracy towards the end of training.

Figure 5.28: Gradient descent with different learning rates for finding the minimum of a quadratic function. On the (left) the learning rate is very low so convergence is slow and on the (right) the learning rate is too high so gradient descent diverges.

Remarks:

- Gradient descent is very similar to Newton's method for optimization, but in Newton's method the learning rate is not a hyperparameter, but 1 over the second derivative of the loss function, $\frac{\partial^2}{\partial w_j^2} L(\hat{f}, D)$. Newton's method often converges in fewer steps, but unfortunately the second derivative is usually hard to calculate.

- Any differentiable loss function can be optimized with gradient descent. If the loss function is convex (see Definition 1.28), then the global minimum will eventually be reached (with a suitable learning rate).

- When the training dataset is large, it is often too costly to calculate $\frac{\partial}{\partial w_j} L(\hat{f}, D)$ over the whole dataset D just to make a single update step. In practice the training data is often shuffled and split into minibatches D_i (subsets of equal size) and $\frac{\partial}{\partial w_j} L(\hat{f}, D_i)$ is used in the update step. This is called *stochastic gradient descent* (SGD).

- Even though we have a closed form solution for minimizing the loss in linear regression, let's derive the corresponding update rule to see how this works.

Theorem 5.29 (LMS Rule). *The **Least Mean Squares** (LMS) update rule for linear regression is given by*

$$w_j := w_j + \alpha \left(y - \boldsymbol{w}^T \boldsymbol{x} \right) x_j$$

And the batch update rule is given by

$$w_j := w_j + \alpha \sum_{(\boldsymbol{x},y) \in D} \left(y - \boldsymbol{w}^T \boldsymbol{x} \right) x_j$$

Proof. We derive the batch update rule. Recall the squared error loss function

$$L(\hat{f}, D) = (\boldsymbol{y} - \mathbf{X}\boldsymbol{w})^T (\boldsymbol{y} - \mathbf{X}\boldsymbol{w}).$$

And recall from Theorem 5.6 that the derivative with respect to $\boldsymbol{w}$ is given by

$$\frac{\partial L}{\partial \boldsymbol{w}} = -2\mathbf{X}^T (\boldsymbol{y} - \mathbf{X}\boldsymbol{w})$$

Substituting into the gradient descent formula, we get the update rule

$$w_j := w_j + 2\alpha \left(\mathbf{X}^T (\boldsymbol{y} - \mathbf{X}\boldsymbol{w}) \right)_j$$

$$= w_j + 2\alpha \sum_{(\boldsymbol{x},y) \in D} \left(y - \boldsymbol{w}^T \boldsymbol{x} \right) x_j$$

If we scale the learning rate by a half we get the LMS batch update rule. And taking a single training sample for D we get the single update rule. □

5.7 Logistic Regression

So far we have discussed how to learn a function $f : x \mapsto y$. when $y \in \mathbb{R}$. In this section we introduce a method for binary classification, where y is either 0 or 1, i.e., $y \in \{0, 1\}$. Similar to linear regression, we learn a linear function $\boldsymbol{w}^T \boldsymbol{x}$, but now we classify all samples with $\boldsymbol{w}^T \boldsymbol{x} > 0$ as 1, all samples with $\boldsymbol{w}^T \boldsymbol{x} < 0$ as 0. In other words we find a hyperplane to separate the classes. There are different approaches to do so, we present a statistically motivated approach called logistic regression. The key idea is simple: We take the output of the linear function and squash it into the range $[0, 1]$. We treat this squashed output as a probability. The more sample $\boldsymbol{x}$ is in the direction of vector $\boldsymbol{w}$, the higher the probability that sample $\boldsymbol{x}$ belongs to class 1.

Definition 5.30 (Binary Logistic Regression). *We want to find an approxi-* *mation $\hat{f}(\boldsymbol{x}) \approx f(\boldsymbol{x}) = y \in \{0, 1\}$. We choose the following form for $\hat{f}$:*

$$\hat{f}(\boldsymbol{x}) = \psi(\boldsymbol{w}^T \boldsymbol{x}) = \frac{1}{1 + \exp(-\boldsymbol{w}^T \boldsymbol{x})}$$

*where ψ is called the **logistic** or **sigmoid** function.*

Remarks:

- The logistic function ψ squashes inputs from $[-\infty, \infty]$ into the range $[0, 1]$. Other functions, such as the *probit link function*, could be used instead.

- $\hat{f}(\boldsymbol{x})$ can be seen as an estimate of the probability that $y = 1$. Therefore, we usually classify $\boldsymbol{x}$ as positive (1) if $\hat{f}(\boldsymbol{x}) > 0.5$ and as negative (0) if $\hat{f}(\boldsymbol{x}) < 0.5$.

- This defines a *decision boundary* $\psi(\boldsymbol{w}^T \boldsymbol{x}) = 0.5$, or $\boldsymbol{w}^T \boldsymbol{x} = 0$, where everything on one side of the boundary is classified as positive and everything on the other side as negative. Samples on the boundary can be classified arbitrarily as positive or negative. The decision boundary is linear in the features.

- As in linear regression, we again can add higher order features, e.g., x^2 or $sin(x)$, to learn non-linear decision boundaries.

- How do we find the optimal values for $\boldsymbol{w}$? We could minimize the squared error loss (Definition 5.3). However, this would not guarantee predictions between 0 and 1 and we could not interpret the predictions as probabilities.

- Instead we apply the logistic function, interpret the output as a probability and choose the model that maximizes the likelihood of generating exactly the labels in our training data.

Definition 5.31 (Bernoulli Likelihood function).

$$\mathcal{L}(\boldsymbol{w}) = \prod_{(\boldsymbol{x},y) \in D} P(y \mid \boldsymbol{x}, \boldsymbol{w}) = \prod_{(\boldsymbol{x},1) \in D} \psi(\boldsymbol{w}^T \boldsymbol{x}) \prod_{(\boldsymbol{x},0) \in D} (1 - \psi(\boldsymbol{w}^T \boldsymbol{x}))$$

Remarks:

- $\mathcal{L}(\boldsymbol{w})$ is the probability of observing the vector of outputs $\boldsymbol{y}$ given input data $\mathbf{X}$ and parameters $\boldsymbol{w}$. Intuitively, we want to choose $\boldsymbol{w}$ such that we maximize this probability, i.e., we want to choose the parameter values that make our observations the most likely to occur. This is called Maximum Likelihood Estimation (MLE).

Lemma 5.32 (Logistic Regression Loss Function or Log Loss). *Assuming the y's in $\boldsymbol{y}$ are independent and identically distributed Bernoulli with parameters $p = \hat{f}(\boldsymbol{x}) = \psi(\boldsymbol{w}^T \boldsymbol{x})$, maximizing $\mathcal{L}(\boldsymbol{w})$ is equivalent to minimizing the following loss function:*

$$L(\hat{f}, D) = -\frac{1}{n} \sum_{(\boldsymbol{x},y) \in D} \left[y \log(\hat{f}(\boldsymbol{x})) + (1 - y) \log(1 - \hat{f}(\boldsymbol{x})) \right]$$

Proof.

$$\mathcal{L}(\boldsymbol{w}) = \prod_{(\boldsymbol{x},1) \in D} \psi(\boldsymbol{w}^T \boldsymbol{x}) \prod_{(\boldsymbol{x},0) \in D} (1 - \psi(\boldsymbol{w}^T \boldsymbol{x}))$$

$$= \prod_{(\boldsymbol{x},y) \in D} (\psi(\boldsymbol{w}^T \boldsymbol{x}))^y (1 - \psi(\boldsymbol{w}^T \boldsymbol{x}))^{1-y}$$

$$= \prod_{(\boldsymbol{x},y) \in D} \hat{f}(\boldsymbol{x})^y (1 - \hat{f}(\boldsymbol{x}))^{1-y}$$

In practice we maximize the logarithm of the likelihood function because the product simplifies to a sum, which prevents numerical problems and makes differentiation simpler. Taking the logarithm does not change the optimal $\boldsymbol{w}$:

$$\log \mathcal{L}(\boldsymbol{w}) = \sum_{(\boldsymbol{x},y) \in D} \left[y \log(\hat{f}(\boldsymbol{x})) + (1 - y) \log(1 - \hat{f}(\boldsymbol{x})) \right]$$

$$= -n \cdot L(\hat{f}, D)$$

Therefore $\operatorname{argmax}_{\boldsymbol{w}} \mathcal{L}(\boldsymbol{w}) = \operatorname{argmin}_{\boldsymbol{w}} L(\hat{f}, D)$.

□

Remarks:

- The scaling factor n does not change the optimum $\boldsymbol{w}$, but it averages the value of the loss across samples, hence allowing for better comparison across models.

- If $y = 1$ and $\hat{f} = 1$, then $L = -\log(\hat{f}) = -\log(1) = 0$, i.e, our classifier is correct with perfect "confidence", and hence the loss is zero. If on the other hand $y = 1$ and $\hat{f} = 0$, then $L = -\log(\hat{f}) = -\log(0) \approx \infty$, i.e., our classifier is wrong with perfect "confidence", which incurs very high cost. Similarly if $y = 0$ and $\hat{f} = 0$, then $L = 0$, and if $y = 0$ and $\hat{f} = 1$, then $L = \infty$.

- Unfortunately, there is in general no closed form solution for logistic regression. However, we can use gradient descent (Definition 5.26) to minimize $L(\hat{f}, D)$. But first we need to calculate the gradient.

Lemma 5.33 (Gradient of the Log Loss). *The gradient of $L(\hat{f}, D)$ from Lemma 5.32 with respect to w_j is given by*

$$\frac{\partial L}{\partial w_j} = \frac{1}{n} \sum_{(\boldsymbol{x}, y) \in D} \left[\hat{f}(\boldsymbol{x}) - y\right] \cdot x_j$$

Proof. First we calculate the derivative of the logistic function:

$$\frac{d}{dz}\psi(z) = \frac{d}{dz}\left[\frac{1}{1 + e^{-z}}\right]$$
$$= \frac{e^{-z}}{(1 + e^{-z})^2}$$
$$= \frac{1}{1 + e^{-z}} \cdot \frac{e^{-z}}{1 + e^{-z}}$$
$$= \psi(z) \cdot (1 - \psi(z))$$

Using this we can calculate the derivative of the loss for a single training sample, $L(\hat{f}, \boldsymbol{x})$, with respect to w_j:

$$\frac{\partial}{\partial w_j} L(\hat{f}, \boldsymbol{x}) = -\frac{\partial}{\partial w_j}\left[y \log(\hat{f}(\boldsymbol{x})) + (1 - y)\log(1 - \hat{f}(\boldsymbol{x}))\right]$$
$$= -\frac{\partial}{\partial w_j}\left[y \log(\psi(\boldsymbol{w}^T\boldsymbol{x})) + (1 - y)\log(1 - \psi(\boldsymbol{w}^T\boldsymbol{x}))\right]$$
$$= -\left[\frac{y}{\psi(\boldsymbol{w}^T\boldsymbol{x})} - \frac{1 - y}{1 - \psi(\boldsymbol{w}^T\boldsymbol{x})}\right] \cdot \frac{\partial}{\partial w_j}\psi(\boldsymbol{w}^T\boldsymbol{x}) \qquad \text{chain rule}$$
$$= -\left[\frac{y}{\psi(\boldsymbol{w}^T\boldsymbol{x})} - \frac{1 - y}{1 - \psi(\boldsymbol{w}^T\boldsymbol{x})}\right] \cdot x_j \cdot \psi(\boldsymbol{w}^T\boldsymbol{x})(1 - \psi(\boldsymbol{w}^T\boldsymbol{x})) \quad \text{chain rule}$$
$$= -\left[\frac{y - \psi(\boldsymbol{w}^T\boldsymbol{x})}{\psi(\boldsymbol{w}^T\boldsymbol{x})(1 - \psi(\boldsymbol{w}^T\boldsymbol{x}))}\right] \cdot x_j \cdot \psi(\boldsymbol{w}^T\boldsymbol{x})(1 - \psi(\boldsymbol{w}^T\boldsymbol{x}))$$
$$= -\left[y - \psi(\boldsymbol{w}^T\boldsymbol{x})\right] \cdot x_j$$
$$= \left[\hat{f}(\boldsymbol{x}) - y\right] \cdot x_j$$

And since differentiation and finite summation are interchangeable, i.e.,

$$\frac{d}{dx}\sum g(x) = \sum \frac{d}{dx}g(x)$$

we get the gradient for the total loss, $L(\hat{f}, D)$, as:

$$\frac{\partial L(\hat{f}, D)}{\partial w_j} = \frac{1}{n}\sum_{(\boldsymbol{x}, y) \in D}\left[\hat{f}(\boldsymbol{x}) - y\right] \cdot x_j$$

$\square$

Remarks:

- We can then use gradient descent (Def. 5.26) to update the model parameters.

- We can also add lasso or ridge regularization to prevent our model from overfitting (Def. 5.23).

- What if y can take on more than two values? A straightforward way to extend binary logistic regression to $k > 2$ classes is to train k separate logistic regression models, one for each class. We then choose the class with the highest probability score. This is a general method for extending binary classifiers to multinomial problems, and is referred to as One versus Rest (OvR).

- There are other ways to extend logistic regression to the multinomial case. For example, we can use the softmax function instead of the sigmoid function.

Definition 5.34 (Softmax Regression). ***Softmax regression*** *with $k \geq 2$ classes chooses the following functional form for $\hat{f}$:*

$$\hat{f}(\boldsymbol{x})_i = \sigma(\boldsymbol{w}_{(1)}^T\boldsymbol{x}, \boldsymbol{w}_{(2)}^T\boldsymbol{x}, ..., \boldsymbol{w}_{(k)}^T\boldsymbol{x})_i = \frac{\exp(\boldsymbol{w}_{(i)}^T\boldsymbol{x})}{\sum_{j=1}^{k}\exp(\boldsymbol{w}_{(j)}^T\boldsymbol{x})} \qquad \text{for } i = 1, \ldots, k$$

*where σ is called the **softmax function**.*

Remarks:

- $\hat{f}(\boldsymbol{x})$ is a vector of length k, with elements summing up to 1. It can be seen as a vector of probabilities, with $\hat{f}(\boldsymbol{x})_i \approx \mathbb{P}(f(\boldsymbol{x}) = i)$.

- A different set of linear weights, $\boldsymbol{w}_{(i)}$, is learned for each class i. Since the probabilities sum up to 1, one set of weights is "redundant".

- For $k = 2$, softmax regression reduces exactly to binary logistic regression with $\boldsymbol{w} = \boldsymbol{w}_{(2)} - \boldsymbol{w}_{(1)}$.

- We also get the same form for the loss function:

Lemma 5.35 (Softmax Regression Loss Function). *Assuming the y's in Y are independent and identically distributed categorical random variables with parameters $p_i = \sigma(\boldsymbol{w}_{(i)}^T \boldsymbol{x})$ for $i = 1, \ldots, k$, then maximizing the likelihood $\mathcal{L}(\boldsymbol{w})$ is equivalent to minimizing the following loss function:*

$$L(\hat{f}, D) = -\frac{1}{n} \sum_{(\boldsymbol{x}, y) \in D} \sum_{i=1}^{k} \left[\mathbb{1}\{y = i\} \log(\hat{f}(\boldsymbol{x})_i) \right]$$

where $\boldsymbol{w}$ is the set of all weights $\{\boldsymbol{w}_{(i)}\}_{i=1,\ldots k}$ and $\mathbb{1}\{\cdot\}$ is the indicator function.

Remarks:

- Note that for $k = 2$ this loss is identical to the loss for logistic regression in Lemma 5.32, only the notation has been changed to use the indicator function.

- In order to learn non-linear decision boundaries with logistic regression, we need to do feature engineering. This can be challenging and time consuming, and it also makes the resulting model more difficult to interpret. In the following we introduce a different type of model that addresses these issues: Decision Trees.

5.8 Tree-Based Methods

So far we have considered regression models based on the form $\boldsymbol{w}^T \boldsymbol{x}$ where the output is essentially a weighted sum of the input features, potentially squashed by a sigmoid function. Binary decision trees introduce hierarchy: We keep partitioning the input space into smaller regions. In order to make a prediction $\hat{f}(\boldsymbol{x})$, we simply start at the root of the decision tree and apply the decision rules until we reach a leaf, which then determines the output value.

```
1  def predict(self, x): # self is current node, initially root
2      if self.is_leaf():
3          return self.value
4      if x[self.feature] ≤ self.threshold:
5          self.left.predict(x)
6      else:
7          self.right.predict(x)
```

Algorithm 5.36: Decision tree algorithm.

Remarks:

- For classification, in Line 3 we return the majority class of the leaf. For regression, we return the average value of the leaf.

- Decision trees can learn non-linear decision boundaries and are easy → notebook to interpret and visualize.

- Decision trees are binary trees that contain nodes V, where each internal (non-leaf) node has an associated splitting rule. Every node v corresponds to a subset $D_v \subseteq D$.

- How are splitting rules defined?

Definition 5.37 (Decision Tree Splitting Rule). *A **splitting rule** of an internal node v is given by a tuple (i, t), where i is the index of a feature and t is a threshold value. A split defines **left** and **right** subsets*

$$D_{v,l} = \{\boldsymbol{x} \mid x_i \leq t\}, \quad D_{v,r} = \{\boldsymbol{x} \mid x_i > t\}$$

Remarks:

- Unlike everything so far, a splitting rule here is based on a single feature value and not a linear combination of features. As such all splits are axis-aligned.

- Now we know how to use a decision tree to make predictions, but how do we build a decision tree in the first place? We do so by finding good splitting rules. And we find good splitting rules by minimizing a loss function of course!

Definition 5.38 (Regression tree loss: MSE). *For node v with samples D_v the mean squared error (MSE) loss is defined as:*

$$L(D_v) = \frac{1}{|D_v|} \sum_{y \in D_v} (y - \bar{y})^2$$

where the prediction $\bar{y}$ of a node is the average of the target values of all samples in D_v.

Remarks:

- The MSE is the variance of the target value. The aim is to create splits that lower the total variance in the leaves.

- For classification the loss function is a measure of purity. We call a node with all samples belonging to the same class perfectly pure. We aim to find splits that successively increase the purity of the nodes. The most common measures of purity are *entropy* and *Gini impurity*.

- We now have a measure of loss for each node, but we need to combine the loss of the left and right subsets to decide what the next split should be. One could consider many things here: minimizing the total loss, minimizing the maximum loss, making balanced splits, etc. One natural choice is to minimize the weighted average loss.

Definition 5.39 (CART loss function). *To find the best split (i^*, t^*) at node v, CART minimizes the following loss function:*

$$L(i, t) = \frac{|D_{v,l}|}{|D_v|} L(D_{v,l}) + \frac{|D_{v,r}|}{|D_v|} L(D_{v,r})$$

where the loss $L(\cdot)$ measures the impurity or error of the resulting left and right subsets.

Remarks:

- The weights ensure that all training samples have equal contribution. For example, for regression trees the weighted MSE loss reduces to the mean squared error over D_v. This is not the same as $L(D_v)$, since we use the two new mean values $\bar{y}_l$ and $\bar{y}_r$ to calculate the errors.

- CART trees are built recursively using binary splits, with every split minimizing above loss function.

- One could keep splitting until all the leaves contain single samples, like a binary search tree (Definition 4.2). This would give 100% accuracy on the training data, but would not generalize well to new data. Stopping criteria can be used to halt splitting. Typical stopping criteria are: maximum tree depth, minimum number of samples in a node, minimum decrease in the value of the loss function.

- Instead of stopping criteria, it is generally better to grow a large tree and then prune it back. This way we might add some useful additional splits and have the chance to remove less useful splits. Even if all additional splits are not helpful, we can still remove them when pruning. In essence pruning allows us to see some steps into the future, before finalizing our tree.

- CART is a greedy algorithm that finds good solutions, but is unlikely to find the optimal solution. Finding the optimal tree (e.g. a minimum depth decision tree) is NP-hard.

- One big advantage of decision trees is their ease of interpretability, but they suffer from high variance and overfit easily. A common technique to improve decision trees is to use many trees in an ensemble.

Definition 5.40 (Bootstrap Sample). *A bootstrap sample D_b is obtained by drawing n samples from dataset D uniformly with replacement.*

Remarks:

- In expectation, every D_b contains $1 - 1/e = 63.2\%$ samples from D.

- What happens if we draw many bootstrap samples and use each to train a separate model?

Definition 5.41 (Bootstrap Aggregating or Bagging). *Bagging is an ensemble algorithm where q models are trained on q bootstrap samples D_b. The models are combined either by averaging the outputs (regression) or by majority vote (classification).*

Remarks:

- Bagging ensembles can be built with any base learners, including decision trees. In general, an ensemble method can turn many low-bias high-variance base learners into a single low-bias low-variance model.

- The use of bootstrap samples de-correlates the individual learners, i.e., they will generally make different mistakes which can be averaged out.

- If we use decision trees as our base learners we can do even better. What if we add even more randomness to bagging by subsampling the features that can be chosen for finding the best splits?

Definition 5.42 (Random Forest). *A random forest is a bagging ensemble of q decision trees. The learners $\hat{f}_b$ are trained in such a way that at every node only a random subset of the features can be used for finding the best split.*

Remarks:

- The random subsampling of features for each split further de-correlates the individual trees, making random forests powerful models that can achieve both relatively low bias and low variance.

5.9 Evaluation

How can we know whether our models are performing well? For regression we can simply measure the models MSE (or absolute error) on the evaluation set (Definition 5.15). However, for classification, the value of the loss function is not intuitive; it does not directly tell us how good our model is at classifying samples. To get a sense of the performance of a classifier, the most important tool is the confusion matrix.

Definition 5.43 (Confusion Matrix). *A confusion matrix, also known as error matrix, visualizes the performance of a classifier on a given dataset. Rows represent the actual class labels, and columns contain the predictions of the classifier.*

Predicted label

	p	**n**
p′	True Positive TP	False Negative FN
n′	False Positive FP	True Negative TN

(True label)

Remarks:

- From the confusion matrix we can derive different metrics to summarize the performance of a classifier:

accuracy	$\text{ACC} = \frac{TP+TN}{P+N}$
positive predictive value (precision)	$\text{PPV} = \frac{TP}{TP+FP}$
true positive rate (recall)	$\text{TPR} = \frac{TP}{TP+FN}$
false positive rate	$\text{FPR} = \frac{FP}{FP+TN}$
F1 score	$F_1 = 2 \cdot \frac{PPV \cdot TPR}{PPV+TPR}$

Remarks:

- Confusion matrices and the derived metrics can be used with more than two classes. However, the derived metrics can then be computed in several different ways. The *macro* method first computes the metrics for each class and then averages these values. The *micro* method aggregates the predictions from all samples and computes the metrics directly.

- Another very common way of evaluating binary classifiers ($k = 2$) is the Receiver Operator Characterstic (ROC) Curve and the derived Area Under Curve (AUC) metric.

Definition 5.44 (Receiver Operator Characteristic (ROC) Curve). *The ROC Curve plots the TPR against the FPR. Given a binary classifier $\hat{f}$, one can order the data points by $\hat{f}(\boldsymbol{x})$, and then plot the TPR and FPR for every possible classification threshold $\tau_{pr} \in [0,1]$. The resulting graph is called the ROC Curve.*

Remarks:

- τ_{pr} is the threshold for which if $\hat{f}(\boldsymbol{x}) \geq \tau_{pr}$ we classify a sample as positive, and otherwise as negative.

- Trade-off between TPR and FPR: If we want 100% TPR (recall), we can simply classify every sample as positive. However, we would then also (wrongly) classify every negative sample as positive, leading to a high FPR. An increase in the TPR generally leads to an increase in the FPR.

- A perfect classifier would achieve $TPR = 1$ at $FPR = 0$, i.e., the curve would "hug" the top left corner.

- The ROC curve of a random classifier follows the diagonal, i.e., $TPR = FPR$.

- The area under the ROC curve (AUC) is often used as a convenient summary of a classifier's performance

- Intuitively, the AUC metric tells us the following: Given a random negative sample $\boldsymbol{x}_N$ and a random positive sample $\boldsymbol{x}_P$, what is the probability that $\hat{p}(\boldsymbol{x}_P) > \hat{p}(\boldsymbol{x}_N)$, i.e, the classifier will assign higher probability to the random positive sample than to the negative sample.

Chapter Notes

This chapter was written in collaboration with Gino Brunner and Béni Egressy.

Bibliography

Chapter 6

Neural Networks

Computers are better than humans at playing Chess, Go, Poker, Dota, or Starcraft. They compose pop songs, write fiction stories, draw paintings, replace actors in movies, and drive vehicles. Whenever a computer does something mind-boggling, you can bet that a neural network is involved. Neural networks have become fascinating function approximators. How so? At their core, neural networks are based on simple linear mappings, combined with non-linear activation functions and gradient descent. So conceptually neural networks are not so different from our discussions in Chapter 5. But size matters! The biggest neural networks have up to 175 billion weights. So training needs data, hardware and patience.

6.1   Nodes and Networks

Definition 6.1 (Node). *A node (or neuron) is a computing unit v that produces an activation value y. The node v first calculates an affine transformation on $\boldsymbol{x} \in \mathbb{R}^d$, then applies an activation function $\sigma : \mathbb{R} \to \mathbb{R}$:*

$$y = \sigma(\boldsymbol{w}^T \boldsymbol{x}),$$

where $\boldsymbol{x}$ is an input vector and $\boldsymbol{w} \in \mathbb{R}^d$ are learned weights. We call $z = \boldsymbol{w}^T \boldsymbol{x}$ the pre-activation value. Like in Chapter 5 (Definitions 5.5 and 5.30), we assume that w_0 is integrated into $\boldsymbol{w}$, i.e., $\boldsymbol{w} = (w_0, w_1, \ldots, w_{d-1})^T$, and $\boldsymbol{x}$ includes an additional constant 1, i.e., $\boldsymbol{x} = (1, x_1, \ldots, x_{d-1})^T$.

Remarks:

- In the literature, the intercept w_0 is sometimes referred to as "bias" b, and kept separate from $\boldsymbol{w}$, i.e., $y = \sigma(\boldsymbol{w}^T \boldsymbol{x} + b)$. This naming complicates the vector notation; it may also be confusing since we used the term bias for a model property in Section 5.4.

- The activation function σ can take many different forms. Some nodes may simply use the identity as activation function, i.e., $\sigma(z) = z$. Most nodes apply non-linear activation functions in order to allow the model to approximate non-linear functions, e.g. the sigmoid function $\psi(\boldsymbol{w}^T \boldsymbol{x}) = \frac{1}{1+\exp(-\boldsymbol{w}^T \boldsymbol{x})}$ of Definition 5.30.

- In order to allow for gradient-based training, the activation function must be differentiable.

- We combine many neural nodes into a network:

Definition 6.2 (Neural Network). *A neural network is a directed acyclic graph (DAG) formed by a set of nodes V that are connected by a set of directed edges E. The input $\boldsymbol{x}$ of the network is stored by the n input nodes V_i (with no incoming edges). The output $\boldsymbol{y}$ of the network will be computed by the m output nodes V_o (with no outgoing edges). All other nodes (with incoming and outgoing edges) are called hidden nodes V_h. We have $V_i + V_h + V_o = V$. Note that we use the letters x and y to refer to both, the input and output of the whole network as well as the input and output of a single node. The we will use subscripts if the usage is not clear from the context.*

In neural networks the function is computed in the forward direction: The input $\boldsymbol{x}_v$ of each node v is the vector of computed outputs $\boldsymbol{y}$ of its DAG predecessor nodes. Then, v computes its own output as $y_v = \sigma(\boldsymbol{w}_v^T \boldsymbol{x}_v)$. Hence then nodes must be processed in DAG order.

Given an input $\boldsymbol{x} \in \mathbb{R}^n$, a neural network as a whole then approximates a function $f(\boldsymbol{x}) : \mathbb{R}^n \to \mathbb{R}^m$ by calculating $\hat{f}(\boldsymbol{x}) \approx f(\boldsymbol{x})$.

```
1   # V = Vi ∪ Vh ∪ Vo  = network
2   # v.x = v's input, the output of v's DAG input-nodes
3   def forward(V):
4       for v in Vh ∪ Vo (in DAG order):
5           v.y = v.σ(v.w,v.x)   # Definition 6.1
6       return Vo
```

Algorithm 6.3: Feed-forward computation in DAG.


Figure 6.4: Example of a neural network for an input $\boldsymbol{x}$ and an output $\boldsymbol{y}$ with three input nodes, eight hidden nodes and two output nodes.

Remarks:

- There exist cyclic neural networks as well, see Definition 6.25.

- In neural networks, nodes generally follow a structure that can be represented with layers.

Definition 6.5 (Multi-Layer Perceptron or MLP). *The nodes are often orga-
nized in layers. The first layer are the input nodes V_i, the last layer the output
nodes V_o. Each hidden node is in a layer l, and all nodes of layer l have the
same input, namely the outputs y_v of all nodes v of layer $l-1$. Layered networks
are known as Multi-Layer Perceptrons, where perceptron is an older name for
node.*

Remarks:

- Layering will help us to speed up computation, as the pre-activation
  of a whole layer can be computed with a single matrix-vector multi-
  plication $z = \mathbf{W} \cdot x$, where the matrix $\mathbf{W}$ is composed of the weight
  rows w, z is the vector of pre-activation values in the nodes and x are
  the output values of the previous layer (with the additional 1).

- The number of layers determines the depth of the network. A "deep"
  network is a neural network with multiple layers.

- Can a neural network compute/approximate any function?

6.2 Power and Limitations

Theorem 6.6 (Universal Approximation Theorem). *Given a continuous func-
tion $f : \mathbb{R}_+ \to \mathbb{R}$ (for simplicity, input $x \geq 0$), there exists a neural network $\hat{f}$
with one hidden layer that approximates f arbitrarily well. That is*

$$|f(x) - \hat{f}(x)| < \varepsilon \text{ for all } x \geq 0 \text{ and } \varepsilon > 0.$$

Proof. We construct a neural network with sigmoid non-linearities in the hidden
nodes v_i ($i > 0$) and a single linear output node v_o.

The single output node v_o computes function $\hat{f}(x) = w^T y$, where w are v_o's
weights, and y_i are the outputs produced by the hidden nodes. As usual, w
includes an additional value w_0, and y starts with a additional constant 1. We
choose $w_0 = f(0)$, such that $\hat{f}(0) = f(0)$ even without any hidden nodes. Every
hidden node v_i computes $y_i = g_i(x) = \psi (\kappa \cdot x + b_i)$, where $\kappa \to \infty$ is a large
constant. While the sigmoid function $\psi(z) = \frac{1}{1+\exp(-z)}$ from Definition 5.30 is
a smooth step function, $\kappa \to \infty$ will make that step sharp.

The construction is inductive. We start out with $x = 0$, hence $\hat{f}(x) = f(x)$.
As long as the difference between $\hat{f}(x)$ and $f(x)$ is less than ε we keep growing
x. As soon as $|f(x) - \hat{f}(x)| \geq \varepsilon$, we introduce a new hidden node v_i. The value
b_i of the hidden node v_i is representing the current position x, as $b_i = -\kappa \cdot x$.
This makes sure that v_i will introduce a new step in $\hat{f}$ right at the current x.
The weight w_i of v_o for the new input y_i is set as $w_i = f(x) - \hat{f}(x)$. This corrects
the output of $\hat{f}$ for the newly accumulated error. If $f(x)$ was increasing, then
a correcting $+\varepsilon$ step is added, if $f(x)$ was decreasing, a correcting $-\varepsilon$ step is
added. In both cases, we again get $\hat{f}(x) \approx f(x)$. Figure 6.7 visualizes the effects
of the parameters b_i and w_i. $\square$

Figure 6.7: Effect of weights on the sigmoid function.

Remarks:

- Our proof is a simplified version of the original. The full theorem
  is more general, also applicable to continuous functions with inputs
  from a multi-dimensional compact set. There exist also versions using
  other activation functions than sigmoid.

- The theoretical construction given in the proof is not used in practice,
  as it would lead to numerical instabilities ($\kappa \to \infty$).

- What is the minimum size of a neural network in order to approximate
  a function f?

Definition 6.8 (VC-dimension). *The Vapnik–Chervonenkis dimension is a
measure of the complexity/capacity of a model $\mathcal{F}$. Given an arbitrary input
set X of size n. Assume that each input $x_i \in X$ maps to a binary classification
output $y_i \in \{0,1\}$. If for all 2^n possible outputs there exist a function $f \in \mathcal{F}$
that correctly classifies all inputs, i.e., $f(x_i) = y_i$ for all $x_i \in X$, then the model
$\mathcal{F}$ has (at least) VC-dimension n. The VC dimension is therefore defined as the
size of the largest set for which there exists a classifier $f \in \mathcal{F}$ for any labeling
of the set.*

Example 6.9. *Let $\mathcal{F}$ be the family of interval classifiers given by*

$$f_{\alpha,\theta}(x) = \begin{cases} 1 & \text{for } x \in [\theta, \theta + \alpha] \\ 0 & \text{otherwise} \end{cases}$$

*The VC dimension of this classifier is at least 2, as for 2 data points x_1 and
x_2 we can always choose θ and α to either include none, one or both of the
points. Therefore any labeling of 2 points can be represented by an interval
classifier $f \in \mathcal{F}$. However, for 3 points x_1, x_2 and x_3 with $x_1 < x_2 < x_3$
we cannot define an interval that includes x_1 and x_3 but not x_2. The labeling
$y_1 = 1, y_2 = 0, y_3 = 1$ can therefore not be represented and we conclude that the
VC dimension of $\mathcal{F}$ is 2.*

Example 6.10. *Let $\mathcal{F}$ be the family of trees with n leaves which classify a scalar
input $x \in \mathbb{R}$, i.e. the n leave nodes are labeled with either 0 or 1. Traversing this
tree from the root to the leaves we can see that each internal node divides the
remaining interval into two sub-intervals. This means that the whole tree divides
the real number line into n distinct intervals whose boundaries we can choose
by the threshold weights of the nodes in the tree. We can therefore construct
a classifier from $\mathcal{F}$ that surrounds each data point in a set of n points with a*

correctly labeled interval to reproduce all possible labelings (as we are also free to choose the label on the leave nodes). The VC dimension of $\mathcal{F}$ is therefore at least n. On the other hand, the VC dimension cannot be larger than n. Consider $n + 1$ ordered data points with alternating labels. It follows that no two points may be in the same interval. Thus we would need $n + 1$ many intervals, which is impossible with n leaves.

Remarks:

- As can be seen from the examples, the VC dimension is not a measure specific to neural networks, but can be applied to any model $\mathcal{F}$.

- With a sigmoid activation function, a neural network with w weights has at least VC-dimension w^2. A network with a simple sign activation function on the other hand has at most VC-dimension $w \log w$.

- The VC dimension can be generalized in various ways, e.g. beyond binary classification, or by restricting the possible outputs.

- However, the promise of a neural network is not only to approximate any function, but rather to *learn* how to approximate any function. It can be shown via various reductions (Definition 2.6) that learning is NP-hard (Definition).

6.3 Training Neural Networks

During training, neural networks learn to automatically extract features from the raw input: In the forward computation the representation of the data held by the network becomes progressively closer to the value of the approximated function. Therefore, neural networks are effectively feature extractors.

Definition 6.11 (Feature Extractor ϕ)**.** *A feature extractor ϕ is a function that transforms raw input data $\boldsymbol{x}$ into features. These features represent the initial data in a way that simplifies approximating a function f.*

Remarks:

- In Section 5.2 we manually engineered ϕ. In Sections 5.3 and 5.4 we then learned how to tell whether we did a good job.

- Neural networks on the other hand automatically learn ϕ and thus, no manual feature extraction is needed. The idea is that every additional hidden layer of the network represents the data more abstractly than the previous one. With every layer, the representation of the data is less like input $\boldsymbol{x}$ and more like output $\boldsymbol{y}$.

- Why and how exactly this works is not well understood – this is the *mystique* of deep neural networks.

- Training a neural network is similar to training a linear regression node with gradient descent (Chapter 5): We calculate the gradient of the loss with respect to the network parameters and adjust the parameters accordingly. This is called backpropagation.

Definition 6.12 (Backpropagation)**.** *Backpropagation is an algorithm that computes the gradient of the loss L with respect to the parameters W of the neural network. In the DAG representation of a neural network, each node computes its pre-activation value z as a weighted sum over its input values $\boldsymbol{x}$. By the chain rule, we can calculate the error of the output nodes with respect to z as*

$$\frac{\partial L(\hat{f}, D)}{\partial z} = \frac{\partial L}{\partial y} \cdot \frac{\partial y}{\partial z}.$$

Given this gradient, and using $\boldsymbol{z} = \boldsymbol{w}^T \boldsymbol{x}$, we can calculate the error with respect to the weights w_i and the error with respect to the node's inputs x_i as

$$\frac{\partial L}{\partial w_i} = \frac{\partial L}{\partial z} \cdot x_i, \text{ and } \frac{\partial L}{\partial x_i} = \frac{\partial L}{\partial z} \cdot w_i.$$

The gradients with respect to the weights $\boldsymbol{w}$ can then be used for updating the weights, while the gradient with respect to the inputs can be aggregated to pass the gradient $\frac{\partial L}{\partial y}$ to the preceding nodes in the network. Concretely, each node in the network adjusts its own weights $\boldsymbol{w}$ based on the error signal $\frac{\partial L}{\partial w}$ and tells its input nodes x_i to adjust by backpropagating $\frac{\partial L}{\partial x_i}$. The backpropagation algorithm is given in Algorithm 6.13.

```
1   # v.x = v's input, stored during forward computation
2   # v.z = pre-activation value, stored during forward computation
3   # v.err = initial error of the node, set to ∂L/∂y for output nodes
4   # v.err = 0 initially for hidden and input nodes
5   # v.w_grad = gradient vector of node v
6   # v.prev = list of indices of v's input nodes
7   def backward(V):
8       for v in V (in reversed DAG order):
9           v.z_err = v.err * ∂v.y/∂v.z   # ∂y/∂z is the gradient of σ(z)
10          v.w_grad = v.z_err * v.x   # v.w_grad and v.x are vectors
11          for v_i in v.prev:
12              v_i.err += v.z_err * v.w[v_i]
13      return [v.w_grad for v in V]
```

Algorithm 6.13: Backpropagation Algorithm

Remarks:

- Many libraries backpropagate gradients with a simple function call.

- Memory may be problem, since we need to memorize $\boldsymbol{x}$ and z at every node.

- Backpropagation is only the method for computing the gradient, while another algorithm, such as stochastic gradient descent (Definition 5.26), is used to perform learning using this gradient.

- Backpropagating gradients, i.e., applying the chain rule, means performing a number of multiplications. For deep neural networks this can lead to numerical issues.

Definition 6.14 (Vanishing Gradients). *Gradient descent updates can stagnate due to vanishing gradients, i.e., gradients that are close to 0. These can occur during backpropagation for different reasons:*

- *The activation function $\sigma(\cdot)$ saturates, i.e. the gradient $\frac{\partial y}{\partial z} \approx 0$. In this case, the gradient $\frac{\partial L}{\partial w_i}$ of all weights w_i and the backpropagated gradients $\frac{\partial L}{\partial x_i}$ of the node will also be close to zero. The node stops learning.*

- *The summation in Line 12 of Algorithm 6.13 can accidentally become 0 (when terms cancel each other).*

Definition 6.15 (Exploding Gradients). *The gradient calculations with backpropagation can also lead to devastatingly large gradients, called exploding gradients. Note that gradient descent only converges for sufficiently small gradient steps and too large gradients can lead to divergence. Reasons are:*

- *Some large weight $|w_i| \gg 1$ boosts the backpropagated error.*

- *The summation in Line 12 of Algorithm 6.13 can accidentally become large because of many positive (or negative) terms.*

Remarks:

- Vanishing or exploding gradients can propagate in a network, i.e. nodes with vanishing or exploding gradients can backpropagate the problem to their predecessor nodes.

- Simple yet generally effective solutions include reducing the number of layers or clipping the gradients (for the exploding case). A related solution is to normalize the activation values, as done in techniques such as *batch normalization* or *layer normalization*.

- Another effective solution for the vanishing gradient problem is the introduction of skip connections (connect nodes which are not in neighboring layers) which provide an additional path for the flow of information. During backpropagation this helps the gradients to continuously flow backwards, even if they vanish in certain points of the network.

- An additional option for mitigating the vanishing gradients problem is to use activation functions that do not saturate in both directions (positive and negative). The best-known of such activation functions is the Rectified Linear Unit activation (ReLU).

Definition 6.16 (ReLU). *The Rectified Linear Unit activation is defined as:*

$$\sigma(x) = \max(0, x)$$

Remarks:

- The ReLU activation function introduces a nonlinear transformation that remains very close to linear (piecewise linear with only two pieces) and does not saturate in the positive direction.

- Remarkably, ReLU is non-differentiable, which violates Definition 6.1. In fact, for gradient-based learning it is enough if the subderivatives of the function exist (and for ReLU they do).

- The gradient $\frac{\partial y}{\partial z}$ of the ReLU activation is 0 when $x \leq 0$, 1 otherwise. This simplicity speeds up the computation of backpropagation.

- ReLU activation is the current default choice for neural networks. However, a large number of related activation functions exist, which modify some aspects of the function, e.g., leakyReLU has a non-zero slope for values smaller than 0. Some other functions are: PReLU, GeLU, SeLU, Maxout, etc.

- Another design choice that impacts the performance of the model is the initialization scheme.

Definition 6.17 (Initialization scheme). *Rule that determines the initial parameter values W of a neural network, i.e., the values before training starts.*

Remarks:

- As seen in Figure 5.27, when the loss function is non-convex (has multiple local minima), starting the learning process at different points can lead to different solutions.

- Stochastic initialization is a good default for initializing the parameters of a neural network. These schemes give random initial values (with some constraints) to the parameters of the network in order to "break the symmetry", i.e., to prevent that nodes with the same input and same activation converge to the same values during optimization.

- The loss landscape of neural networks is complex, with a large number of local minima. Surprisingly, converging to a local minimum during training is good enough for a neural network to perform well usually.

6.4 Practical Considerations

The complex loss landscape of neural networks makes the learning process significantly more complicated than in classical machine learning models. Therefore, sophisticated learning algorithms (also called optimizers) that build on top of Stochastic Gradient Descent (Section 5.6) are used in practice. There is no consensus on which of the existing algorithms is best.

Remarks:

- Adam optimizer is currently considered a good default. It belongs to the family of adaptive learning rate algorithms, which adapt the learning rate for each parameter individually throughout the course of learning.

- Other popular optimizers include SGD with Momentum, RMSProp, and linear learning rate decay.

- The learning scheme/rate is probably the most important hyperparameter in neural networks. Finding an appropriate learning rate can produce a dramatic improvement in the performance of the network.

- Neural networks often have a remarkably large amount of hyperparameters, which do have a strong impact on the performance of the model. Although tuning hyperparameters is more and art than a science there are automatic hyperparameter optimization algorithms that can help in this process.

Definition 6.18 (Hyperparameter Optimization Algorithm). *A hyperparameter optimization algorithm is an algorithm that wraps the learning algorithm of a model and chooses its hyperparameters, hiding this choice from the user.*

Remarks:

- When there are few hyperparameters to set, a common approach is Grid Search as discussed in Definition 5.24. The main problem of Grid Search is that the computational cost grows exponentially with the number of hyperparameters, which makes it expensive for large neural networks.

- An alternative is Random Search: the hyperparameter values are samples from a uniform distribution in a certain interval. Random search converges faster to an optimum.

- A large number of hyperparameter optimization algorithms exist using techniques such as evolutionary algorithms, Bayesian optimization or population-based-training.

- Hyperparameter optimization algorithms often have their own hyperparameters, such as the range of values that have to be explored. Fortunately, these secondary hyperparameters are easier to set in the sense that similar secondary hyperparameters can lead to acceptable performance in a wide range of tasks.

6.5 Regularization

Everything discussed so far portraits neural networks as powerful function approximators. Neural networks can approximate any continuous functions and even functions of high complexity, e.g., functions from a function class with high VC-dimension. The issue with this is that neural networks tend to overfit.

To give an intuitive explanation why this is the case, recall the bias-variance trade-off from the previous chapter. There, we saw that polynomials of too high degree yield a high variance which leads to a bad generalization performance. Now, the universal approximation theorem states that a sufficiently large neural network can approximate any continuous function. Hence, a sufficiently large neural network can also approximate any polynomial. Without any restrictions, the variance of a neural network can be very high and the generalization performance very poor.

- To prevent it, classical parameter norm penalty can be applied, like the L2 (ridge) and L1 (lasso) penalties seen in Definition 5.23. These penalties are applied by including the penalty term in the loss function of the model, exactly the same as in Section 5.5.

- Furthermore, there are some other regularization techniques specific to neural networks.

Definition 6.19 (Dropout). *Dropout is a regularization technique: for each sample at each training iteration, we set the output y of each node to zero with probability p. After training has completed we do not drop nodes anymore as we want to use the full capacity of the network. However, we multiply each activation where dropout was applied with $1 - p$. This is done to keep the activation on the same level as it was in expectation during training.*

Remarks:

- Effectively, dropout trains a different model at each iteration, where all models share the non-zeroed parameters. For large networks there is no risk that dropout breaks the information flow between the input and the output of the network.

- Dropout reduces the inter-dependencies between nodes in the network, which helps the model to learn more robust features, and also reduces overfitting.

- Dropout is computationally cheap and can be applied to any model that uses distributed representations and that is trained with gradient descent, i.e., any neural network.

- Dropout reduces overfitting but does not completely eliminate the problem. Luckily, dropout can be easily combined with other regularization strategies, for example, with early stopping.

Definition 6.20 (Early Stopping). *Early stopping is a regularization strategy that returns to the parameter setting that produces the lowest validation error. In early stopping, training terminates when the best recorded validation error does not improve for a predefined number of epochs; this number is called patience.*

Remarks:

- Early stopping can be understood as an efficient algorithm for selecting the number of training steps, which is a hyperparameter.

- The cost of early stopping in terms of computation is that the validation needs to be run periodically after each epoch and that at least one copy of the parameters needs to be stored in memory. These costs are however small and generally do not cause any limitation.

- Early stopping does not affect the learning dynamics, can be used in conjunction with other regularization strategies, and is easy to implement.

6.6   Advanced Layers

While neural networks can theoretically learn any function, large networks (with too many weights) often struggle to converge to good solutions. Often we can use knowledge about the underlying problem to reduce the number of weights substantially. Some early successes of neural networks were achieved in image processing, where methods from classical computer vision were adapted to neural networks in the form of convolutions.

Definition 6.21 (Convolutional Neural Network or CNN). *A convolutional neural network is a neural network layer that works on structured input data such as the pixels of an image. A CNN applies the same function (the same weights) to all neighborhoods of the input layer.*


Figure 6.22: Convolution operation for a 4×4 input image $\boldsymbol{X}$, a 3×3 filter $\boldsymbol{W}$, and a 2×2 output $\boldsymbol{Y}$. The filter $\boldsymbol{W}$ slides over the image and for each position of the filter, a value y_i is calculated as the dot-product of the filter and the sub-matrix of $\boldsymbol{X}$ that it covers. E.g., $y_1 = w_1 x_1 + w_2 x_2 + \cdots + w_9 x_{11}$.

Example 6.23. *We want to detect vertical edges in images. Vertical edges can be found calculating the convolution (with symbol $\circledast$) of the image and a vertical*

Sobel filter, given by the matrix $\boldsymbol{W}$:

$$\boldsymbol{W} = \begin{pmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{pmatrix}$$

As an example, consider the 3×5 greyscale image $\boldsymbol{X}$ where each pixel takes a value from 0 to 255 represented by the following matrix:

$$\boldsymbol{X} = \begin{pmatrix} 80 & 92 & 163 & 234 & 230 \\ 85 & 98 & 237 & 233 & 232 \\ 83 & 96 & 236 & 235 & 231 \end{pmatrix}$$

To compute the convolution operation, the filter $\boldsymbol{W}$ slides over the image $\boldsymbol{X}$. Generally, we require that the output $\boldsymbol{Y} = \boldsymbol{X} \circledast \boldsymbol{W}$ is of the same size as the input $\boldsymbol{X}$, for this, we need to pad the matrix $\boldsymbol{X}$. To prevent the padding from creating artificial edges, we extend the image by copying the border pixels $\boldsymbol{X'}$:

$$\boldsymbol{X'} = \begin{pmatrix} 80 & 80 & 92 & 163 & 234 & 230 & 230 \\ 80 & 80 & 92 & 163 & 234 & 230 & 230 \\ 85 & 85 & 98 & 237 & 233 & 232 & 232 \\ 83 & 83 & 96 & 236 & 235 & 231 & 231 \\ 83 & 83 & 96 & 236 & 235 & 231 & 231 \end{pmatrix}$$

In a convolution, the filter $\boldsymbol{W}$ slides over the image $\boldsymbol{X}$. For each position of the filter on the padded image (see Figure 6.22), one element of $\boldsymbol{Y}$ is calculated as the sum of the element-wise multiplication of the overlap. As the filter slides horizontally and vertically, this operation is repeated to obtain the values of each element of $\boldsymbol{Y}$. The resulting matrix $\boldsymbol{Y}$ is:

$$\boldsymbol{Y} = \begin{pmatrix} -49 & -401 & -561 & -196 & 13 \\ -51 & -540 & -551 & -52 & 10 \\ -52 & -611 & -552 & 20 & 13 \end{pmatrix}$$

The large values in the second and third column of matrix $\boldsymbol{Y}$ indicate that there is a strong gradient (variation) between those columns in the image, i.e., the image has a vertical dark/light edge. This edge is less pronounced in the first row.

Remarks:

- In this example the filter $\boldsymbol{W}$ was given. A CNN does not know these weights, but only the information that $\boldsymbol{W}$ is a 3×3 convolution filter between two layers. With appropriate training data, the CNN will learn the weights of $\boldsymbol{W}$ by using backpropagation.

- Sometimes these learned patterns correspond to those that we as humans consider meaningful, e.g., edges. However, usually the patterns extracted by CNNs are not understandable. The power of CNNs reside in learning complex patterns that humans would not be able to design.

- Discrete convolutions can be performed beyond two dimensions. For example, an audio signal can be represented by the signal intensity at discrete time steps. In that case, a 1-dimensional convolution can be applied over the time dimension to filter the signal.

- In general, to apply convolutions to a given input, the input has to be structured as a tensor.

Definition 6.24 (Tensor). *A tensor of order d is a d-dimensional array. A tensor generalizes vectors (1-dimensional) and matrices (2-dimensional). The shape of a tensor is a list of d integers defining the size of each dimension of the tensor.*

Remarks:

- An image, represented by its RGB (red, green, blue) pixel values can $\rightarrow$ notebook be naturally represented as a tensor of order 3 and shape [3, height, width]. An index $(0, x, y)$ into this tensor yields the intensity of red in the pixel at location (x, y).

- Note that the memory requirements of higher order tensors can be high. E.g., an order 5 tensor with 100 values in each dimension stores 100^5 values, which, given a floating point precision of 32 bits, requires 40 GB of memory.

- We have seen that CNNs exploit translation invariance in the structure of the data. Are there other such structural biases we can exploit? For example, what if we want a neural network to remember important features over several time steps of a sequential input?

Definition 6.25 (Recurrent Neural Network or RNN). *In contrast to feed-forward neural networks, a recurrent neural network operates with time steps t. Each time step t gets an input $\boldsymbol{x}_t$ and a state $\boldsymbol{s}_t$. It outputs an updated state $\boldsymbol{s}_{t+1}$ and an output $\boldsymbol{y}_t$. More formally, we define the mappings*

$$\boldsymbol{y}_t = \hat{g}(\boldsymbol{x}_t, \boldsymbol{s}_t)$$

$$\boldsymbol{s}_{t+1} = \hat{h}(\boldsymbol{x}_t, \boldsymbol{s}_t)$$

where $\{\boldsymbol{x}_t\}_{t=0}^{\tau}$ and $\{\boldsymbol{y}_t\}_{t=0}^{\tau}$ are the input and corresponding output sequence of length τ and $\hat{g}(\cdot)$ and $\hat{h}(\cdot)$ are differentiable functions with learnable parameters. The initial state $\boldsymbol{s}_0$ can be a vector of learnable parameters, or simply initialized to $\boldsymbol{0}$.

Remarks:

- There are several ways of how to define $\hat{g}(\cdot)$ and $\hat{h}(\cdot)$, from simple linear projections to complex combinations of operations to combine the given inputs.

Example 6.26. *Consider that we want to solve the multi-path problem of wireless transmission, i.e., given a signal we wish to filter delayed copies from the signal. To do this online, i.e., while the signal is received, we have to remember*

the current input signal to filter a similar pattern later. We therefore seek to train an RNN to remember a given input for a few time steps and then reproduce it for filtering purposes. E.g., given the input signal $[5, 10, 0, 1.5, 3.5, 1]$ we want the RNN to output $[5, 10, 0, 0, 0, 0]$. This can be achieved if we initialize the state $\boldsymbol{s}_0$ to $\boldsymbol{0}$ and parametrize $\boldsymbol{s}_{t+1} = \hat{h}(x_t, \boldsymbol{s}_t) = \boldsymbol{W}\boldsymbol{s}_t + \boldsymbol{w}_h \cdot x_t$ with $\boldsymbol{W}$ and $\boldsymbol{w}_h$ as

$$\boldsymbol{W} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -0.1 & -0.3 & 0 & 0 \end{pmatrix} \qquad \boldsymbol{w}_h = [0, 0, 0, 1]^T$$

where $\boldsymbol{W}$ shifts the state and filters new inputs and $\boldsymbol{w}_h$ reads in the new symbol. The readout is given by $y_t = \hat{g}(x_t, \boldsymbol{s}_t) = \boldsymbol{w}_g^T \boldsymbol{s}_t + x_t$ with

$$\boldsymbol{w}_g = [-0.1, -0.3, 0, 0]^T$$

Note that this parametrization simply reads the input symbol into the state $\boldsymbol{s}_t$, propagates the symbol for some time steps and then subtracts it from a later input.

Remarks:

- Instead of these given weights, neural networks will learn $\hat{g}$ and $\hat{h}$ when trained on real world signals. This is particularly useful if the input is a vector of multiple correlated noisy signals and a simple remember-and-reproduce solution is sub-optimal.

- Earlier we discussed that neural networks are directed acyclic graphs (DAGs). But RNNs are cyclic, as the state from step t gets fed back to the network in step $t + 1$. How can we train such a network?

- The solution is to copy the network τ times, i.e., unroll the cycle. This yields one long DAG where the state $\boldsymbol{s}_t$, calculated as intermediate output of one copy, is fed into the next copy. The calculated gradients for each copy are then summed to update the parameters. This is called *backpropagation through time (BPTT)*. Note that this can lead to vanishing/exploding gradients as we are essentially trying to train a network of depth τ.

- RNNs that are commonly used today are *Gated Recurrent Units (GRUs)* and *Long Short Term Memories (LSTMs)*. These address the issue of vanishing/exploding gradients in their definition of $\hat{g}$ and $\hat{h}$. The resulting architectures implement ideas similar to that of skip connections in feed-forward neural networks, albeit historically GRUs and LSTMs came long before people started talking about skip connections in MLPs and CNNs.

- As apparent from the equations in the definition above, RNNs are inherently sequential. They can process one input only after the previous input has been processed. This is slower than approaches that can process the whole sequence in parallel (such as CNNs).

- Both, CNNs and RNNs take advantage of *weight sharing*. In CNNs, the same weights (filters) are applied to all locations of the image. In RNNs, the same functions $\hat{g}(\cdot)$ and $\hat{h}(\cdot)$ (with the same learnable weights) are applied to all time steps t.

- What if we do not want to apply the same function everywhere? More specifically, what if only a selection of the input is of interest? Can we design an architecture that favors solutions which select features from the input instead of using the whole input? Can we index the input in a differentiable way?

Definition 6.27 (Attention). *Attention is a method to aggregate inputs in a selective manner. Given n input vectors $\{\boldsymbol{x}_i\}_{i=0}^{n-1} \in \mathbb{R}^d$, each input vector is projected into a key vector $\boldsymbol{k}_i \in \mathbb{R}^{d_k}$ and a value vector $\boldsymbol{v}_i \in \mathbb{R}^d$. The projection is done by two learned matrices $\boldsymbol{W}_k \in \mathbb{R}^{d \times d_k}$ and $\boldsymbol{W}_v \in \mathbb{R}^{d \times d}$. Additionally, attention learns a query vector $\boldsymbol{q} \in \mathbb{R}^{d_k}$. For each input vector an attention score s_i is calculated as the dot-product between the query $\boldsymbol{q}$ and the key vector $\boldsymbol{k}_i$:*

$$s_i = \boldsymbol{q} \cdot \boldsymbol{W}_k \boldsymbol{x}_i = \boldsymbol{q} \cdot \boldsymbol{k}_i$$

The attention scores are normalized by a softmax (Definition 5.34) operation and each normalized score is multiplied by its corresponding value vector $\boldsymbol{v}_i = \boldsymbol{W}_v \boldsymbol{x}_i$. The results are added to produce the attention output $\boldsymbol{y} \in \mathbb{R}^d$:

$$\boldsymbol{y} = \frac{\sum_{i=0}^{n-1} \exp(s_i) \cdot \boldsymbol{v}_i}{\sum_{j=0}^{n-1} \exp(s_j)}$$

Remarks:

- The attention mechanism described here is called dot-product attention. This is the most common type of attention, but other variants exist as well.

Example 6.28. *Consider that we want to find in what position a sequence of $n = 5$ integers contains the value "3". Our input is $\boldsymbol{x} = [7, -5, -2, 3, 4]^T$. We use a ReLU for the keys: $\boldsymbol{k}_i = ReLU [x_i - 3, 3 - x_i]^T$. And we use a one-hot encoding for the values $\boldsymbol{v}_i$. For example, $x_1 = 7$ gets key and value*

$$\boldsymbol{k}_1 = [4, 0]^T, \boldsymbol{v}_1 = [1, 0, 0, 0, 0]^T.$$

Using $\boldsymbol{q} = [-1, -1]^T$ we get the scores $\boldsymbol{s} = [-4, -8, -5, 0, -1]^T$. Plugging the scores and values into the softmax gives

$$\boldsymbol{y} = [0.013, 0.000, 0.004, 0.718, 0.264]^T.$$

Because of the softmax, the output is not quite as clean as one might hope, since "$4 \approx 3$".

Remarks:

- Thanks to weight sharing, attention can processes input vectors of any length.

- A neural network can consist of multiple attention aggregations to select multiple (potentially different) inputs.

- If the scores are calculated based on the inputs, i.e., $s_i = f_i(\{\boldsymbol{x}_i\}_{i=0}^{k-1})$ for some functions f_i, attention is also referred to as *self-attention*, as the input "attends" to itself.

- Attention architectures yield the state-of-the-art performance in natural language processing tasks, as most of the time some words are more important than others to understand a sentence.

- All architectures presented in this section, i.e., CNNs/RNNs and Attention, take some domain knowledge to tailor the neural network to a specific purpose. This is also referred to as an *inductive bias*.

6.7 Architectures

In the previous section we introduced several building blocks that give different inductive biases. Let us now see how different losses and architectures can be combined to solve advanced computational challenges.

Definition 6.29 (Autoencoder). *An autoencoder is a neural architecture formed by two neural networks: an encoder $f_{enc}(\cdot)$, which encodes the input $\boldsymbol{x} \in \mathbb{R}^n$ into a representation $\boldsymbol{z} \in \mathbb{R}^m$ called latent code; and a decoder $f_{dec}(\cdot)$, which decodes the latent code back into an approximation of the original input, i.e., $f_{dec}(f_{enc}(\boldsymbol{x})) \approx \boldsymbol{x}$. The representation is often designed to compress information by setting $m \ll n$. Autoencoders minimize a loss term named "reconstruction loss", that represents the difference between the output and the input.*

Remarks:

- The encoder and the decoder can be any type of neural network, e.g., MLPs, CNNs, RNNs or a combination of them.

- An example of reconstruction loss is the L^2 error:

$$L = \frac{1}{|D|} \sum_{\boldsymbol{x} \in D} ||\boldsymbol{x} - f_{dec}(f_{enc}(\boldsymbol{x}))||_2^2$$

- A schematic depiction of an autoencoder is shown in Figure 6.30.

Remarks:

- The requirement $m \ll n$ is not a necessity. We can also design autoencoders to have a specific structure in the latent code, or a latent code tailored to a given purpose through an additional loss.

Figure 6.30: Schematic view of an autoencoder.

- The main advantage of autoencoders is that they learn the latent representation (or code) in an unsupervised manner, i.e., without labels. This makes them a versatile architecture that can be used in a wide range of problems, such as dimensionality reduction, compression, data denoising or unsupervised feature extraction.

- There exist many different types of autoencoders, with different losses, architectural elements or with additional inductive biases.

- Besides compression and encoding, neural networks can also be used for data generation.

Definition 6.31 (Generative Adversarial Network or GAN). *GANs are a class of deep generative models in which two neural networks are trained simultaneously, while competing in a two-player minimax game. The generator's $f_{gen}(\boldsymbol{r})$ tries to produce realistic synthetic samples from a random input $\boldsymbol{r}$. The binary classifier called discriminator $f_{dis}(\boldsymbol{x})$ estimates whether a sample $\boldsymbol{x}$ is real or synthetic. The goal of the generator is to maximize the probability that the discriminator makes a mistake on $f_{dis}(f_{gen}(\boldsymbol{r}))$.*

Remarks:

- As in the case of autoencoders, the discriminator and generator can be any type of neural network.

- During training, the discriminator improves its ability to recognize synthetic samples while the generator learns to produce increasingly realistic samples to deceive the discriminator. In this adversarial setting, the equilibrium is reached when the generator produces realistic samples such that the discriminator cannot distinguish whether they are real or synthetic.

- The architecture of a vanilla GAN is shown in Figure 6.34.

Remarks:

- GANs achieved remarkable results in image generation. In particular, they can generate realistic-looking pictures and videos, which has raised concerns about malicious uses of these models to generate deepfakes.

- A GAN is a fully automated Turing Test with generator = testee, and discriminator = tester.

Figure 6.32: GAN architecture for generation of synthetic images of celebrities. For each sample, the discriminator needs to decide whether its input corresponds to a real or to a synthetic celebrity.

Example 6.33. *Faceswap-GAN is a popular implementation of a trained model for deepfake generation. At high level, this model uses an autoencoder as generator. Given an image of a human face, it produces a segmentation mask as well as the reconstructed input image. A segmentation mask is a representation of an image that delineates the most important objects in the image; in the case of human faces, these are the eyes, nose, ears, etc. Roughly speaking, an arbitrary image of a face A can be combined with the segmentation mask of another face B in order to generate an image that replaces the features of image B with those of A, i.e., a deepfake. This is what the model does at inference time.*

During training, Faceswap-GAN used a discriminator that determines whether an input image is a real face or a deepfake, as well as some other advanced methods such as a perceptual loss that improve image quality.


Figure 6.34: Architecture of Faceswap-GAN.

Remarks:

- Although originally conceived for generative tasks such as denoising, reconstruction or data generation, GANs have proved useful in other domains such as supervised learning, semi-supervised learning or reinforcement learning.

- Deep learning is applied in many different areas and consequently, there is a wide range of architectures. We list some promising architectures in Table 6.35:

Name	Description	Purpose
Transformer	A family of sequence-to-sequence models based on the self-attention operation.	State of the art in natural language processing (NLP) and gaining relevance in other areas. Recently received a lot of media attention with models like BERT and GPT-3.
Variational Auto-Encoder (VAE)	A generative model that tries to match the data distribution by enforcing a simplified posterior distribution in the latent space of an auto-encoder.	Data generation and posterior approximation in inherently stochastic models.
Siamese Network/ Contrastive Learning	The same network applied to different inputs, trained to recover a notion of similarity in the output representation.	Unsupervised representation learning (e.g., recovering an approximately metric space), authentication, hashing and matching
Graph Neural Network	Replicating a network on all nodes of a graph and incorporating operations for message exchanges with neighbors on the graph.	Graph/node/edge classification, community detection in social network and predicting protein/molecular interactions
Implicit Network	Training a neural network to recover a value from an index.	Data compression, super-resolution, image in-painting
Hyper Network	A neural network that outputs the weights of another neural network.	Mode abstraction and meta learning

Table 6.35: A glossary of promising architectures.

6.8 Reinforcement Learning

A neural network and its corresponding loss have to be end-to-end differentiable in order to apply gradient descent. So what if a problem is not differentiable? What if we want to find an optimum in a sequential setting, like an optimal sequence of decisions to reach a desired goal in an environment?

Definition 6.36 (Markov Decision Process or MDP). *A Markov decision process formally defines an environment. An MDP is a 5-tuple (S, A, T, R, s_0),*

where S is a set of states, A is a set of possible actions, $T : S \times A \to S$ is a state transition function, which describes the next state based on current state and action. However, T could also be probabilistic, i.e., $T : S \times A \to S \times [0, 1]$. $s_0 \in S$ (or $s_0 : S \to [0, 1]$) is an initial state (distribution). Finally, R is a reward function. Rewards can be given when reaching certain states ($R : S \to \mathbb{R}$), or when taking the right action in a state, $R : S \times A \to \mathbb{R}$.


Figure 6.37: An example Markov decision process to figure out whether and when one should study or party in front of an exam.

Problem 6.38. *Figure 6.37 shows a simple example of an MDP with 6 states $S = \{u_0, u_1, u_2, u_3, fail, pass\}$ and two possible actions $A = \{study, party\}$. The states fail and pass are terminal states, and represent whether the agent fails or passes the exam, respectively. The agent starts in state u_0, i.e., $s_0 = u_0$. The transition probabilities for choosing the 'party' action (red) or 'study' action (black) are shown in the figure. Every time the agent chooses to party, it receives a reward of +1. If the agent chooses to study, it receives a reward of -1 (studying is painful). At the terminal states, the agent gets a reward of +10 for passing the exam and -10 in the failing it. Intuitively, the states u_1 and u_3 represent states where the agent has learned something. These states are more likely to be reached when studying rather than partying, and from these states the agent is more likely to pass the exam. How should the agent act? We can calculate the solution backward from the terminal states by filling in a 2×4 matrix Q giving the quality of each action in each of the non-terminal states. In u_3 taking the action 'study' will yield an expected reward of $Q[u_3, study] = -1 + 0.7 \cdot 10 + 0.3 \cdot (-10) = 3$. Similarly, choosing 'party' in this state yields an expected reward of $Q[u_3, party] = +1 + 0.3 \cdot 10 + 0.7 \cdot (-10) = -3$. In all earlier states we can assume that we take the action with higher expected reward in later states and thereby calculate the remaining values recursively as given in Table 6.39.*

Definition 6.40 (Policy). *A policy $\pi : S \times A \to [0, 1]$ describes how the agent acts in the environment, i.e., how likely it will take an action $a \in A$ in a given state $s \in S$.*

State	study	party
u_3	3	-3
u_2	-1.9	-5.1
u_1	0.53	0.57
u_0	-1.171	-0.159

Table 6.39: Expected reward Q for each action in each non-terminal state. It's best to party in states u_0 and u_1, and best to study in states u_2 and u_3. Note that filling in this table is dynamic programming (Definition 1.11).

Remarks:

- Given an MDP and a policy, the state distribution of the agent after τ steps, i.e., how likely it is that the agent is in a given state after τ actions can be calculated.

- The goal in reinforcement learning is to find a good policy π, that is, a policy that accumulates positive rewards.

- More formally, we want to find the optimal policy π^* that maximizes the expected cumulative γ-discounted reward:

$$\pi^* = \max_\pi \mathbb{E}\left[\sum_{t=0}^\infty \gamma^t R(s_t)\right]$$

  where $\gamma \in [0,1]$ is a discount factor that weighs immediate returns relative to future returns, where s_t is the state at time step t, respectively. The expectation is taken over actions sampled from the policy and states sampled from the transition distribution $P(\cdot|s,a)$ given the state s and action a.

- To find such a policy, we need to know how valuable each state is to a given policy.

Definition 6.41 (Value Function). *A value function $V_\pi : S \to \mathbb{R}$ is a policy specific function that given a state returns the expected cumulative discounted reward of the policy starting in state s_t.*

$$V_\pi(s_t) = \mathbb{E}\left[\sum_{\tau=t}^\infty \gamma^{\tau-t} R(s_\tau)\right]$$

Definition 6.42 (Action-Value Function or Q-Function). *An action-value function or Q(uality)-function $Q_\pi : S \times A \to \mathbb{R}$ is a policy specific function that given a state s and an action a returns the expected cumulative discounted reward of taking action a in state s and following policy π thereafter. That is*

$$Q_\pi(s,a) = R(s) + \mathbb{E}\left[\gamma V_\pi(s')\right]$$

where the expectation is over states s' sampled according to $T(s'|a,s)$.

Remarks:

- The value function can also be defined in terms of the Q function as

$$V_\pi(s) = \mathbb{E}_{a \sim \pi}\left[Q_\pi(s,a)\right].$$

- Given the general definition of value and action-value function, we can get a better understanding of what we want to find: the optimal policy π^*

- Note that V_{π^*} gives us for each state $s \in S$ the maximal expected cumulative reward that can be achieved when starting in state s.

- Further, if we are given Q_{π^*} it is easy to derive the optimal policy by simply taking the action that maximizes Q_{π^*}, i.e.

$$\pi^*(a|s) = \begin{cases} 1 \text{ if } a = \text{argmax}_{a'} Q^{\pi^*}(a',s) \\ 0 \text{ else} \end{cases}$$

  In other words, the optimal policy in an MDP is deterministic!

- Knowing that we can derive the optimal policy from a quantity which requires the optimal policy might be a bit "recursive", but we can nevertheless try:

```
1   # S = states
2   # A = actions
3   # T = transitions
4   # R = rewards
5   def value_iteration(S, A, R, T):
6       V = zero vector of size len(S)
7       Q = zero matrix of size len(S) × len(A)
8       while Q not converged:
9           for s in S:
10              for a in A:
11                  Q[s][a] = R(s) + γ∑_{s'∈S} T(s'|a,s) * V[s']
12              V[s] = max(Q[s])   # max over a in A
13      return Q
```

Algorithm 6.43: Value iteration

Remarks:

- If the MDP has cycles and $\gamma \to 1$, then this algorithm may not converge.

Lemma 6.44. *If we are given an MDP that can be represented as a DAG, value iteration converges in one iteration if we process the states $s \in S$ in reversed DAG order.*

Proof. Reversed DAG order means we will start at the terminal states and propagate the cumulative rewards back to the initial states. By induction, we always propagate the maximal achievable value by the max operation in Line 12 of Algorithm 6.43. Therefore, after one iteration, all states have the optimal value V_{π^*} assigned. □

Remarks:

- However, in many real world applications, the state space is just hilariously large. The game Go for instance has $3^{19 \times 19} \approx 10^{172}$ possible states, so it is infeasible to compute or even store the whole Q-table. We can however train a neural network to approximate how likely a given position is to lead to a win. In combination with a bit of look-ahead planning (recursion) a neural network was able to beat the world champion.

Chapter Notes

While the beginnings of artificial neural networks go back to the 1940s [5], deep learning only became widely adapted and efficient in recent years with the use of GPUs to run computations in parallel. However, many of the theoretical investigations and architectures presented here have been known for quite some time by now. The universal approximation capability of neural networks was first shown for sigmoid non-linearities [1] and later generalized to other nonlinearities [3]. Even before it was shown that learning various functions is NP-complete [4, 6]. This is still an active research area, e.g. [2]. The VC Dimension discussion is even older [7]. We summarize further milestones achieved by neural networks in the table below.

Year	Name	Milestone
1989	MNIST	Handwritten digit classification
2005	DARPA	Self-driving car challenge: 212km in 7h
2012	AlexNet	Image classification breakthrough
2014	Deepface	Human level performance in face recognition
2014	DQN	Superhuman performance in many Atari games
2016	AlphaGo	Beats Champion in Go
2017	Waymo	Fully autonomous self-driving on public roads
2018	Obvious	Sells art generated by a GAN for $432,500
2020	OpenAI	GPT-3 model can create poetry, code, etc.
2020	AlphaFold	Achieves 90% in CASP protein folding

Table 6.45: Neural Network Milestones

This chapter was written in collaboration with Damian Pascual and Oliver Richter.

Bibliography

[1] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, Dec 1989.

[2] Surbhi Goel, Adam Klivans, Pasin Manurangsi, and Daniel Reichman. Tight hardness results for training depth-2 relu networks. 2021.

[3] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251 – 257, 1991.

[4] Stephen Judd. On the complexity of loading shallow neural networks. *Journal of Complexity*, 4(3):177 – 192, 1988.

[5] Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, Dec 1943.

[6] Nimrod Megiddo. On the complexity of polyhedral separability. *Discrete Computational Geometry*, 3:325–337, 1988.

[7] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

Chapter 7

Computability

Computability was pioneered by Alan Turing and Kurt Gödel. Turing probably committed suicide by eating an apple he poisoned with cyanide. Gödel on the other hand had an obsessive fear of being poisoned with food; when his wife was hospitalized, he refused to eat, and eventually starved to death. Now it's your turn to study this intoxicating subject.

7.1 Undecidability

In the previous chapters, we have analyzed various computational tasks. While some functions were easy to compute efficiently, others were difficult; in these cases, we have focused on approximations or heuristics. However, given enough time and resources, could we always find the solution to a problem?

Problem 7.1 (Halting problem)**.** *Given a program P and an input x to P, does $P(x)$ halt (stop running) after a finite amount of time?*

Remarks:

- Can we write a Python program that solves Problem 7.1? Somewhat surprisingly, the input of our program is also a program (plus an input parameter of this program).

- Naturally, we must somehow encode the input program. There are various ways to do this. We can, for example, consider the whole code of the program as a long string of text, encoding each character in this string with a byte.

- The halting problem is sometimes easy to solve, for example, in case of the simple programs in Algorithm 7.2. But is it always?

```
1   def P_1(x):
2       print("Hello, world!")
3       return
4
```

```
5   def P_2(x):
6       x = 1
7       while x > 0:
8           x += 1
9       return
```

Algorithm 7.2: Example programs; P_1 is halting, P_2 is not.

Definition 7.3 (Undecidable)**.** *We say that a problem is undecidable if no algorithm solves the problem in finite time for every possible finite input.*

Remarks:

- This should be surprising! Definition 7.3 does not say that the runtime increase will be exponential or even double exponential in the input size, but that the problem really cannot be solved in *any* finite amount of time.

Theorem 7.4. *The halting problem is undecidable.*

Proof. Assume for the sake of contradiction that there exists a program $P_H(P, x)$ that successfully solves the halting problem in finite time for any input. Test program $P_T(x)$ then takes a bitstring as an input and calls P_H as a subroutine:

```
1   def P_T(x):
2       if P_H(x, x) == True:
3           while True: pass      #loop forever
4       else: return
```

Algorithm 7.5: Test program $P_T(x)$.

P_T interprets its input x as a program encoding, and calls the halting solution P_H on program x with input x. Since we assumed that P_H can always solve this problem in finite time, Line 2 is evaluated in finite time for any x.

Since P_T is also a program, it has a bitstring encoding τ according to our encoding scheme. What happens when we call $P_T(\tau)$? Note that this means that we are calling $P_H(\tau, \tau)$ as a subroutine, i.e. querying whether the program described by τ (that is, P_T) halts on input τ.

- If $P_H(\tau, \tau)$ is true, then P_T goes in to an infinite loop according to our code, and never halts. Hence $P_H(\tau, \tau)$ should be false instead!

- If $P_H(\tau, \tau)$ is false, then P_T immediately terminates according to our code, so $P_H(\tau, \tau)$ should be true!

We get a contradiction in both cases, so the halting problem is undecidable. □

Remarks:

- We assumed that the encodings of a program and its input are simply bitstrings. This is also close to practice. What if a bitstring is an invalid program, not respecting the syntax of programs? We could argue that in this case, the program simply halts (with an error).

- Of course Theorem 7.4 only holds in general: we cannot solve the halting problem correctly in finite time for *every* algorithm-input pair. Some specific algorithm-input pairs, e.g. the simple examples in Algorithm 7.2, can be decided easily.

- Also, a program P that actually halts is easy as well: we just run/simulate P, which will eventually halt. The halting problem is only undecidable because of programs P that do not halt. In this case it is difficult to distinguish if P is still running because it did not reach the halting point yet, or because it is never going to halt.

- There is a name for this weaker kind of decidability that is only required to work in one of the two cases (unlike our original concept of decidability in Definition 2.3).

Definition 7.6 (Semi-decidable). *We say that a problem is semi-decidable if there exists an algorithm $\mathcal{A}$ such that*

- *if the answer is* True, *then $\mathcal{A}$ outputs* True *in finite time,*

- *if the answer is* False, *then $\mathcal{A}$ either outputs* False *in finite time or keeps running indefinitely.*

Theorem 7.7. *The halting problem is semi-decidable.*

Remarks:

- Given the undecidability of halting, is there an easy way to show that some other problems are also undecidable? Yes, we can use reductions again (Definition 2.6). Given a problem Π, we can show that if Π was decidable, then the halting problem would also be decidable. This implies that Π is also undecidable. One slight difference from Definition 2.6, however, is that for this argument, we do not need the reductions to run in polynomial time.

- Here is an example for such a reduction.

Problem 7.8 (Mortality problem). *Given a program P, is it true that $P(x)$ halts for any possible input x?*

Remarks:

- This is different from the halting problem, but only so much: instead of a specific input, we now want to know if P halts on *every* input.

Theorem 7.9. *The mortality problem is undecidable.*

Proof. Assume that we have a program P_M that takes a program description as an input, and solves the mortality problem in finite time. Then given a specific input for the halting problem (a program P and an input x), consider the following program:

```
1  def P_T(y):
2      if y == x:
3          run program P(x)
4      else:
5          return
```

Algorithm 7.10: Another testing program.

Now let us run our mortality solution P_M on the encoding of P_T. We know that P_T certainly terminates in finite time for any input different from x. This implies that $P_M(P_T)$ is true if and only if $P(x)$ halts; hence a solution for the mortality problem allows us to solve the halting problem on P and x. Since halting is undecidable, such a solution cannot exist, so we have a contradiction. $\square$

Remarks:

- Well, that was not so surprising; after all, mortality is a close relative of the halting problem.

- How about problems that are a far cry from halting? We need a clean and simple theoretical definition of what we mean by a program, algorithm or function. Let us make a brief detour into abstract machine models and computation theory.

7.2 The Turing Machine

What kind of building blocks do we need to obtain a simple theoretical model of a machine that can, intuitively speaking, do the same computations as a real computer?

Remarks:

- First we need some abstract *states* that represent the current state of our program, and we need to describe the transitions between these states. Such a set of states with predefined transition rules is known as a finite automaton.

- We also need some memory to store data. We usually assume that a memory consists of cells; our program can read data from these cells, write data into these cells, and move between these cells to be able to access any of them. We will now consider a *tape* of cells that is infinite in both directions, i.e. the cells can be enumerated by integers $..., -2, -1, 0, 1, 2, ...$ (from $-\infty$ to ∞).

- We also have a *tape pointer* that points to a specific tape cell at each point in time, indicating that this is the tape cell that we can currently read/write. Initially the tape pointer points to cell 0.

- What kind of data can we write onto this tape? We assume that we have an *alphabet* Σ of possible symbols, and we can write exactly one symbol into each cell. In the simplest case, this alphabet can be binary, i.e. $\Sigma = \{0, 1\}$. We usually use some extra symbol, e.g. $\perp$ for the cells that we consider empty, and we use $\Sigma := \Sigma \cup \{\perp\}$.

- These building blocks define a famous theoretical model of computation.

Definition 7.11 (Turing Machine or TM). *A Turing Machine has a finite set of states S, and a two-way infinite tape. Initially the machine is in a specified starting state $s_0 \in S$, the tape has some symbols on it (the input), and the tape pointer points to cell 0 of the tape.*

In each discrete time step, depending on the current state s and the current tape cell content σ, the machine executes the following steps:

- *change to another state $s' \in S$,*

- *write the tape, i.e. change the content σ of the current tape cell to any symbol $\sigma' \in \Sigma$,*

- *possibly move the tape pointer one step to the left or one step to the right.*

Formally a TM is defined by a function $(S, \Sigma) \to (S, \Sigma, m)$, where $m \in \{left, right, stay\}$ indicates the movement of the tape pointer.

With a TM we usually also select a specific halting (accepting) state $s_h \in S$. We say that the TM accepts an input if, when executed on this input, the TM eventually enters state s_h.

Remarks:

- We assume that computation is over whenever the halting state s_h is reached: the machine does not do anything (i.e. never changes the state/tape content/tape pointer) after this point, and the current tape content is considered to be the *output* of the computation.

- While a TM acts as a model of computation, we can also interpret it as a function: it converts an input (the initial content of the tape) to an output (final content of the tape).

- However, the function is not complete: for some inputs, the output may be undefined, since just like a Python program, a TM can easily run forever and never halt.

- How can we do actual computations in this abstract setting? Let us see an example for a simple operation: incrementing an integer.

Example 7.12 (Incrementation with a TM). *Given a positive integer input x, our task is to increment x by 1. We assume that the input x is given in a binary representation, starting with least significant bit (LSB) first at cell 0, and going until cell $\lfloor \log_2 x \rfloor$. The tape pointer starts at cell 0, and empty cells of the tape are marked with a $\perp$.*

Lemma 7.13. *Example 7.12 can be solved on a simple TM with 2 states.*

Proof. The formal process of incrementing a binary number is as follows. We start going from the LSB to the most significant bit (MSB) until we encounter a 0, and we change every 1 to a 0 during the process. When we first find a 0, we change it to a 1 (or if we have left the MSB, we add an extra 1 to the front), and the incrementation is done. We have to translate this process to states and transitions.

This can be done with two states s_0 and s_h. The starting state is s_0, this is where the execution begins; s_h is a halting state where none of the transitions do anything. The transitions from s_0 are defined as follows:

Transitions from state s_0			
Read	Write	Pointer	Next state
$1 \longrightarrow$	0	right	s_0
0 or $\perp \longrightarrow$	1	stay	s_h

This ensures that the machine enters the halting state exactly when the incrementation is finished, i.e. when the tape contains $x + 1$ in the same binary representation. $\square$

Remarks:

- Describing a TM for more complex computations can be some work, since it usually requires a higher number of states and transitions. Even in case of our incrementation example, if the number is in a reversed representation (i.e. starting with MSB), we already need an extra state to first move to the end of the input and then start processing the input from the other direction.

- The definition of the halting problem on TMs is as follows: given a description of a TM and an input (initial tape content), decide if this TM ever goes into the halting state. Note that this is only a reformulation of our original halting problem, so the same proof shows that this problem is undecidable.

- There also various other versions of TMs, e.g. a TM that has multiple tapes, and it can read/write these tapes simultaneously in every step. One can show that this is equivalent to the single-tape setting in terms of computability.

- There is one important concept in computation that this basic machine model cannot capture: randomization. In order to model that, we need to slightly extend the machine model.

Definition 7.14 (Randomized TM or RTM). *In a Randomized Turing Machine, each transition is replaced by a set of available transitions, and a probability distribution over these transitions. In each step, a transition is chosen at random according to this probability distribution.*

Remarks:

- For example, a program on an RTM might have a state where it moves left on the tape with 50% probability, and moves right with 50% probability.

- Another formulation of RTMs is to take a deterministic TM, and add an extra tape of infinite random bits to the machine. The TM then reads bits from this extra tape, and (possibly) executes different transitions based on the next random bit.

- Randomness is a useful tool. However, strictly speaking, randomness does not increase the power of the machine in terms of computability. If we can solve a decision problem Π on an RTM in finite time, then we can also solve Π on a regular TM, by enumerating and simulating all the possible randomized outcomes. Note that this might increase the running time drastically.

- Our argument uses an important assumption: that we can use a TM to simulate the execution of another TM. We also need this property when expressing the proof of Theorem 7.4 in a TM-based context.

- Luckily, this is possible:

Theorem 7.15 (Universal TM). *There exists a Universal Turing Machine which receives the encoding of another TM T (i.e., a program encoded as a string) and an input x to T on its tape, and simulates the behavior of T on x.*

Remarks:

- This is somewhat similar to a real-world Von Neumann computer architecture, where source code, constants and inputs of a computation are stored in the same memory.

- So how close are TMs to real computers? The fact that our program moves between a finite number of states is pretty realistic. What is unusual, however, is that we can only move in memory one step at a time.

- More realistic machine models do exist:

Definition 7.16 (RAM Machine). *A RAM Machine is a model of computation that has explicit registers (instead of only cells) which can store integer values. The machine is capable of addressing these registers indirectly through pointers (instead of moving only sequentially between them).*

Remarks:

- While RAM machines seem more expressive, they are in fact equivalent to TMs: any program on a RAM machine can also be simulated on a TM.

- Since TMs are simpler, we usually stick to TMs. We say that a problem Π is computable if a TM can compute Π. This means that we can essentially use TMs to define the general notion of an algorithm.

Theorem 7.17 (Church-Turing Thesis). *Any real-world algorithm or computation can be translated into an equivalent computation on a TM.*

Remarks:

- Intuitively, we can imagine an algorithm as a computation we can do with pen and paper, using a finite set of rules. This describes both our notion of real-world programs and the set of computations that are doable on a TM.

- The term Church-Turing thesis is often used differently in different contexts. The version shown above is a slightly informal phrasing, relating TMs to real-world computations. Sometimes the Church-Turing thesis is not considered as a theorem, but rather as a definition of the term algorithm.

- The Church-Turing Thesis allows us to classify models of computation: we consider a model "complete" if it can be used to run any algorithm according to this definition.

Definition 7.18 (Turing-complete). *We say that a model of computation is Turing-complete if it can simulate any TM.*

Remarks:

- Naturally, TMs are Turing-complete. And so are RAM-machines, RNNs (Definition 6.25), and all popular programming languages, e.g. Python, C++ or Java.

- While a futuristic quantum computer surely is an impressive vision, it only allows us speed up the computation of problems. As such, quantum computers are also not "more" than Turing-complete: they can solve the same computational problems as a (Randomized) TM. In fact, there is no known computational model that can compute more than a TM!

- It is also not so easy to think of reasonable computational models that can execute a smaller subclass of computations than TMs. Some examples for such models are finite automata (essentially TMs without tape) or regular expressions.

- How about very different models that do not look anything like computers or programming languages? Can we also use them to do computations?

7.3   Computing on Grids

Definition 7.19 (Tile). *A tile is a 1×1 square. Each side of the tile (left, top, right, and bottom) has a specific color. We assume that we are not allowed to rotate or flip tiles.*

Figure 7.20: Example of a tile

Definition 7.21 (Correct Tiling). *Two tiles can be placed next to each other if their touching side has the same color. E.g. if tile t_1 is red on its top and tile t_2 is red on its bottom, then we can place t_2 directly above t_1.*

Remarks:

- Our goal is to tile a given shape (e.g. a 5×4 rectangle) with a set of tiles such that the tiling is correct.

Problem 7.22 (Tiling). *Assume we are given a set of n tiles, and we can take an arbitrary number of copies of each of these tiles. Does there exist a correct tiling of the entire infinite plane with our given set of tiles?*

Remarks:

- At first glance, this problem seems to have no connection to TMs or the halting problem.

- The most simple correct tiling one can imagine is a periodic tiling, when the same pattern of tiles keep repeating.

Definition 7.23 (Periodic Tiling). *We say that a tiling of the plane is periodic if there exist positive integers w, h such that for every pair of coordinates $i, j \in \mathbb{Z}$, grid square (i, j) has the same tile as grid squares $(i + w, j)$ and $(i, j + h)$.*

Remarks:

- In a periodic tiling, we tile a $w \times h$ rectangle R such that the tiling within rectangle R is correct, and it is also correct to place two such rectangles R next to each other top/bottom or left/right. Then we can cover the entire plane with copies of R.

Lemma 7.24. *If we know that a periodic tiling exists, we can find it in finite time.*

Proof. We take every possible rectangle size $w \times h$, in an increasing ordering according to the sum $w + h$: first 1×1, then 1×2 and 2×1, then 1×3, 2×2 and 3×1, and so on. For each such size, we can try all possible tilings in each of these rectangles, and check their correctness.

If there exists a periodic tiling with a rectangle of size $w \times h$, then we try at most $(w + h)$ total sizes, and thus at most $(w + h)^2$ rectangle shapes before reaching $w \times h$. Each such shape has at most $n^{(w+h)^2}$ possible tilings. Since $(w + h)^2 \cdot n^{(w+h)^2}$ is a finite number, the algorithm indeed terminates in finite time. $\square$

Remarks:

- Unfortunately, this does not answer the question whether the tiling problem is decidable in general. There are sets of tiles where a tiling of the entire plane is possible, but only in a fashion that is not periodic.

- To settle the question of decidability, we show that these tilings are in fact a surprisingly expressive model: we can use them to simulate any TM. This property will allow a reduction to the halting problem.

Theorem 7.25. *Some tile sets are Turing-complete: tilings can simulate the run of any TM on any input x.*

Proof. To outline the main idea of the proof, we will assume a slightly simpler setting: that we only need to tile the bottom half of the plane, i.e. below the origin. This is only for convenience; with further tricks, the same proof method can be extended to the entire plane.

The main idea of the proof is that each row describes the complete state of the tape of a TM in a specific time step, with the top row corresponding to time step 0, the row immediately below corresponding to time step 1, and so on. We can design the tiles carefully such that given a specific row (i.e. current configuration of the TM), the only possible tiling of the row directly below is the next configuration of the TM.

For example, for each tape symbol σ, we can create a tile that has color σ on the top and bottom, and a special default color (say, white) on the left and right. This already allows us to automatically copy the content of the tape into the row below. Furthermore, we use special tiles to keep track of the tape pointer and the current state: the tape cell with the pointer will also be marked with the current state. If, for example, we have a transition from state q_1 to q_2 that also replaces a 0 by a 1 on the tape and moves the tape pointer one step to the right, then we can describe this behavior with the tiles shown in Figure 7.26.


Figure 7.26: Example tiles to simulate a transition of the TM.

The number of tiles we require altogether is only a function of the tape alphabet size and the number of states and transitions in the TM, and thus it is a finite number. By defining some special tiles for the first row, one can also ensure that the only possible tiling of the first row is to have the input x on the tape, the tape pointer at position 0 and the TM in its initial state s_0 (note that this is a major technical step that we do not discuss here).

With such a set of tiles, if the TM halts on x in k steps, then this allows us to tile the first k rows of the plane, and then no tiling will be possible for the

$(k+1)^{\text{th}}$ row (since there is no following configuration of the TM). On the other hand, if the TM runs forever, then there is always a next configuration, so the tile set allows us to tile the entire lower half of the plane. □

Remarks:

- So we can do actual computations with a set of tiles!

Theorem 7.27. *The tiling problem is undecidable.*

Proof. Consider an instance of the halting problem with TM T and input x. As we have seen in Theorem 7.25, we can create a set of tiles that correspond to running T on x, and a tiling with this set is possible if and only if T halts on x. Thus solving this tiling problem allows us to decide whether T halts on x; however, the halting problem is undecidable, so the tiling problem must be undecidable, too. □

Remarks:

- There are other models of computation that work on grids. A popular example is Game of Life.

Definition 7.28 (Game of Life or GoL). *In Game of Life, each cell has two states, black (alive) and white (dead), and the update rule is as follows:*

- *If the cell is black: if it has exactly 2 or 3 black neighbors among its 8 neighboring cells, it remains black, otherwise it becomes white.*

- *If the cell is white: if it has exactly 3 black neighbors among its 8 neighboring cells, it becomes black, otherwise it remains white.*

Remarks:

- These simple rules create a surprisingly wide range of patterns. There are stable configurations which keep their shape without changing; there are oscillators that keep repeating a few specific configurations periodically; there are "gliders" that exhibit a similar periodicity but also slowly move through the grid in the meantime. There are more complex patterns that repeatedly create smaller oscillators or gliders.

- These constructions can then be used to form gadgets on a higher abstraction level: we can create logical AND and OR gates, and ultimately a finite automaton. These tools then allow us to simulate the behavior of a TM in GoL, similarly to the tiles before.

Theorem 7.29. *Game of Life is Turing-complete.*

Remarks:

- As a result, we can also formulate some undecidable problems in this model.

Problem 7.30 (GoL Reachability). *Given an initial configuration c, the task is to decide if another configuration c' will ever occur.*

Theorem 7.31. *GoL Reachability is undecidable.*

Remarks:

- GoL is a in fact a special case of a widespread model of computation on grids called Cellular Automaton.

Definition 7.32 (Cellular Automaton or CA). *A Cellular Automaton consists of a (two-dimensional) grid of cells, where each cell is in a specific state. In each iteration, every cell (concurrently and independently) changes its state based on the current states of the cells in its immediate neighborhood.*

Remarks:

- If we denote the set of states by S, then a CA is essentially described by a function $f : S \times S^N \to S$ (with N denoting the size of the neighborhood). Each cell executes this function in each round to obtain the state in the next round.

- Since GoL is a special case of Cellular Automata, CAs in general are also Turing-complete.

- CAs can model various processes in natural sciences, ranging from Physics to Biology, with the cells of the automaton representing anything from chemical molecules to actual (biological) cells.

- There are many other areas (beyond halting and grids) where we can find undecidable problems. To mention another surprising example: Given a couple of $k \times k$ matrices with integer entries, it is undecidable if they can be multiplied in some order, possibly with repetitions, such that we obtain the zero matrix as a result.

7.4 Post Correspondence Problem

Finally, we discuss some variants of the so-called Post Correspondence Problem. This problem is an interesting conclusion to our whole lecture: it demonstrates that seemingly similar problems can easily have a completely different complexity.

Problem 7.33 (PCP). *We have a set of dominoes, where each domino (α, β) has two words written on the domino: one word α on the top, and one word β on the bottom. Can we make a sequence of these dominoes, possibly with repetitions, such that the concatenation of words on top is the same as the concatenation of words on the bottom?*

Remarks:

- Given a finite alphabet of symbols Σ, a *word* is a finite string formed from these letters, possibly with repetitions. A concatenation of words $\alpha_1, \alpha_2, \ldots$ is the word obtained by writing these words after each other in this order.

Theorem 7.35. *PCP is undecidable.*

Figure 7.34: Example solution of a PCP: both the top and the bottom string is *caabbda*. The sequence consists of 4 dominoes, using one of the dominoes twice.

Proof. The proof is quite technical, so we only outline the main idea. Similarly to tiles, we can use dominoes to simulate the running of a TM. The concatenated string will describe the history of the run of the TM as a list of subsequent configurations. The bottom string is always "one step ahead" the top string in this computation; thus by defining an appropriate domino for each possible transition in the TM, we can ensure that the next configuration is always a valid follow-up to the current configuration. If the TM reaches a terminal state, then some extra dominoes ensure that the top string can catch up to the bottom string; this way the two strings become identical, and thus we have a valid sequence of dominoes that solves the PCP problem.

Since such a PCP solution exists if and only if the TM halts. Since the halting problem is undecidable, the PCP problem is also undecidable. □

Remarks:

- However, there is a simple algorithm that terminates in finite time if the answer is yes.

Theorem 7.36. *PCP is semi-decidable.*

Proof. We can enumerate all possible domino sequences based on their length in increasing order. Then if there exists a solution with a domino sequence of length k for some finite number k, then until reaching this sequence, we check at most

$$n + n^2 + ... + n^k \leq k \cdot n^k$$

possible sequences. Since each such check takes at most $O(k)$ time, we can find the solution in finite time. □

Remarks:

- The PCP problem is often used in reductions when analyzing problems related to formal languages.

- In terms of the number of dominoes used, the best known method to simulate a TM requires 5 different dominoes. This shows that if we have the correct 5 dominoes, then the problem is undecidable. However, what happens if we restrict the problem to less than 5 dominoes?

Theorem 7.37. *With only 1 domino, the PCP problem is decidable in polynomial time.*

Proof. In this case, any sequence consists of a specific number of repetitions of our single domino. In this case, the top and bottom strings are only identical if our single domino has the same word on the top and bottom side (and in this case, a single instance of the domino already provides a solution). We can easily check this in linear time: we only need to read the two words and compare them. □

Theorem 7.38. *With only 2 dominoes, the PCP problem is decidable.*

Proof. The proof of this claim is quite involved, so we do not discuss it here. □

Remarks:

- With 3 or 4 dominoes, it is still an open question whether PCP is undecidable or not.

- Another possible modification is to restrict the size of the alphabet. More specifically we need an alphabet of at least $|\Sigma| \geq 2$ letters for undecidability. If Σ only consists of a single character, then the problem becomes decidable.

Theorem 7.39. *PCP with $|\Sigma| = 1$ is solvable in polynomial time.*

Proof. In this case, we only need to make sure that the top and bottom words have the same length, i.e. the same number of occurrences of our single character. This means that for each available domino, we only need to consider the difference of length between the top and bottom words, which gives us a (not necessarily positive) integer. The task then reduces to analyzing this set of integers, and selecting a subset of them (with possible repetitions) that sums up to 0.

Solving this is rather easy in polynomial time. If one of the integers is 0, then this already forms a valid sequence on its own. If not, then we need to check if there is at least one positive number $x_i > 0$ and at least one negative number $x_j < 0$ among our integers: then a sequence consisting of x_i copies of the number x_j and $|x_j|$ copies of the number x_i also sums up to 0. Otherwise, all the numbers are positive (or negative); in this case, the sum of any sequence is also positive (or negative, respectively). □

Remarks:

- For another variant, we can also restrict the size of the allowed domino sequence.

Problem 7.40 (Bounded PCP)**.** *In Bounded PCP, the input also contains an integer k, and we only accept domino sequences that have length at most k.*

Theorem 7.41. *Bounded PCP is decidable but NP-hard.*

Proof. With n dominoes, we only have n^k possible domino sequences. By enumerating and checking all these possibilities, the problem is clearly decidable in finite time.

The proof of NP-hardness can be shown through a reduction from the longest common substring problem; we do not discuss it here. □

Chapter Notes

The Turing Machine was developed by Alan Turing in 1936, long before the invention of modern day computers [15]. Turing has specifically defined the model in order to study the halting problem, and prove its incomputability. The halting problem (and its different variants) has kept its central place in the area; the majority of known incomputability results are shown through a reduction that comes either directly or indirectly from this problem.

A very similar line of thought and proof technique to the halting problem's incomputability has also appeared in the work of Kurt Gödel, who was studying incompleteness theorems and the axiomatization of natural numbers at about the same time [13]. The general message of these two results has caused a large surprise (even shock) in the scientific community, where the general belief (based on Hilbert's conjectures) was that, intuitively speaking, every well-defined question can be answered. The results have shown that this is not the case, which has far-reaching philosophical consequences.

The Turing Machine has also remained the fundamental model to study computations ever since. The closely related concepts (e.g. Universal Turing Machine, Turing-completeness, or different formulations of the Church-Turing thesis) have been gradually developed and refined in the following decades. This rapidly developing area was studied by some of the most important mathematicians of the 20th century, including John von Neumann, Alonzo Church or Stephen Kleene.

The tiling problem was first discussed by Wang in 1961 [16]. However, in his first work, Wang conjectured that whenever a tiling exists, a periodic tiling also exists. A few years later his student Berger showed that some tile sets only allow an aperiodic tiling, and that the problem is undecidable due to its connection to the halting problem [3].

Game of Life was devised by John Conway in 1970 [8], and has been analyzed in numerous papers and books since then [2]. There are many simulators online where you can create different patterns and follow their development through the rounds [1].

As for Cellular Automata in general, there is an immense literature discussing different aspects of the topic. Different variants of automata have been used in a very wide range of applications, e.g. generating pseudo-random numbers in computer science [14], modeling the crystallization of snowflakes [4], modeling the geometric patterns on seashells [6] or modeling the flow of traffic on the freeway [10].

The PCP problem was introduced by Emil Post in 1946 [12]. A long line of works have followed that tried to reduce the number of dominoes required for undecidability, going down to set of 7 dominoes in 1996 [9], and then finally 5 tiles in the work of Neary in 2015 [11]. The decidability for 2 dominoes was proven by Ehrenfeucht, Karhumäki and Rozenberg [7], while the NP-hardness of Bounded PCP was first discussed in [5].

This chapter was written in collaboration with Pál András Papp.

Bibliography

[1] John Conway's Game of Life Online. https://playgameoflife.com.

[2] Andrew Adamatzky. *Game of life cellular automata*, volume 1. Springer, 2010.

[3] Robert Berger. *The undecidability of the domino problem*. Number 66. American Mathematical Soc., 1966.

[4] Charles D Brummitt, Hannah Delventhal, and Michael Retzlaff. Packard snowflakes on the von neumann neighborhood. *Journal of Cellular Automata*, 3(1), 2008.

[5] Robert L Constable, Harry B Hunt III, and Sartaj Sahni. On the computational complexity of scheme equivalence. Technical report, Cornell University, 1974.

[6] Stephen Coombes. The geometry and pigmentation of seashells. *Nottingham: Department of Mathematical Sciences, University of Nottingham*, 2009.

[7] Andrzej Ehrenfeucht, Juhani Karhumäki, and Grzegorz Rozenberg. The (generalized) post correspondence problem with lists consisting of two words is decidable. *Theoretical Computer Science*, 21(2):119–144, 1982.

[8] Martin Gardner. Mathematical games: The fantastic combinations of john conway's new solitaire game "life". *Scientific American*, 223(4):120–123, 1970.

[9] Yuri Matiyasevich and Geraud Senizergues. Decision problems for semi-thue systems with a few rules. In *Proceedings 11th Annual IEEE Symposium on Logic in Computer Science*, pages 523–531. IEEE, 1996.

[10] Kai Nagel and Michael Schreckenberg. A cellular automaton model for freeway traffic. *Journal de physique I*, 2(12):2221–2229, 1992.

[11] Turlough Neary. Undecidability in Binary Tag Systems and the Post Correspondence Problem for Five Pairs of Words. In *32nd International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 30 of *LIPIcs*, pages 649–661, Dagstuhl, Germany, 2015.

[12] Emil L Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52(4):264–268, 1946.

[13] Panu Raatikainen. Gödel's Incompleteness Theorems. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2020 edition, 2020.

[14] Marco Tomassini, Moshe Sipper, and Mathieu Perrenoud. On the generation of high-quality random numbers by two-dimensional cellular automata. *IEEE Transactions on computers*, 49(10):1146–1151, 2000.

[15] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London mathematical society*, 2(1):230–265, 1937.

[16] Hao Wang. Proving theorems by pattern recognition—ii. *Bell system technical journal*, 40(1):1–41, 1961.