Lecture 5

Ishaan Lal

February 15, 2025

1 Introduction

We'll now step away from the SAT problem for a bit. But don't fret! We will tie it back in later! Today, we will start looking at "high-level" solvers.

Working with SAT was great due to its universality, but it truly felt very low level – having to specify variables and clauses. Instead, we will now choose to specify constraints as a problem in something closer to mathematical language.

2 Linear Programming

Consider the following problem that we will return to throughout the section:

Problem

You are going to be attending a potluck. You want to take a meal with ≥ 5000 calories and ≤ 200 mg of sodium. Additionally, you want to spend as little money as possible.

You can bring any non-negative **real-value** (not necessarily integer) amount of the following items, so long as they abide to the above constraints:

Item	Price / kg	Calories / kg	Sodium / kg
Rice	1.25	750	15
Pasta	1.65	1200	35
Couscous	1.35	1000	60

For example, I could choose to bring 0kg of Rice, 4kg of Pasta, and 0.75kg of Couscous, which would yield $0.750 + 4.1200 + 0.75 \cdot 1000 = 5550$ calories and $0.15 + 4.35 + 0.75 \cdot 60 = 185$ kg of sodium. This solution abides to the constraints posed, and its total price is $0.1.25 + 4.1.65 + 0.75 \cdot 1.35 \approx 7.61$

For a bit of vocabulary, a **linear function** is a function of the form:

$$f(x_1, ..., x_n) = a_0 + a_1 \cdot x_1 + a_2 \cdot x_2 + \ldots + a_n x_n$$

This can be thought of as an equation with only the variables, coefficients, and a constant. In other words, there are no "powers" on any of the variables. For the linear algebra inclined folks, a linear function can be expressed as the product of a row-vector of coefficients/constants with a column vector of variables.

$$f(x_1, \dots, x_n) = A \cdot X \quad \text{with} \quad A = \begin{bmatrix} a_0 & a_1 & \dots & a_n \end{bmatrix}, \quad X = \begin{bmatrix} 1 \\ x_1 \\ \vdots \\ x_n \end{bmatrix}$$

A linear inequality has the form:

 $a_1x_1 + a_2x_2 + \dots + a_nx_n \ge b$ or $a_1x_1 + a_2x_2 + \dots + a_nx_n \le b$

A linear program (LP) is a special class of optimization problems with:

- The **GOAL** being to optimize a linear function. We will either *minimize* or *maximize* a function
- Existence of constraints given as linear inequalities

Example				
As an example, an instance of a linear program may look like:				
minimize	$3x_1 + 2x_2 - 4x_3 + 5$			
subject to	$x_1 \ge 2$			
	$x_2 + 2x_3 \le 10$			
	$x_1, x_2, x_3 \ge 0$			
Solution: We can actually argue a solution by hand. First, note that we must have $x_1 = 2$, because as x_1 increases, so does our <i>objective function</i> (the function we want to minimize). Additionally, we can note that the larger x_3 is, the smaller our objective function becomes. And we want x_2 to be as small as possible. This can be obtained with $x_2 = 0$ and $x_3 = 5$.				

It turns out that this is the optimal solution.

However, not all functions are easy to work with. The more variables we have, the more complex the problem becomes from a human perspective. Same goes for constraints.

Consider the following linear program:

maximize	$2x_1 + 5x_2$
subject to	$0 \le x_1, x_2 \le 3$
	$-2x_1 + 2x_2 \le 5$
	$x_1 + 2x_2 \le 7$
	$2x_1 + x_2 \le 7$

This is a challenging problem to approach by hand. But a good first step may be to find the "feasible region" – that is, the region of (x_1, x_2) values that satisfy all of the constraints. Once we have this region, we can then find the point within the region that maximizes the function.

Since our function only has two variables, consider the Cartesian plane, where the x-axis represents the

value of x_1 and the y-axis represents the value of x_2 . Then, we can plot our constraints as lines:



Here, the blue area represents the *feasible region*. However, how do we know which point in the interior maximizes our function?

Brilliant Observation

A Brilliant observation is that the maximum **MUST** occur at one of the corners of the region!

What? How? The mathematical proof can be found, but we will sit content with a verbal, hand-wavy proof.

Instead of viewing our feasible region as just a set of points for us to test out, think of the feasible region as a mountain – where each point within the region has its own *elevation*, where the elevation is given by the functions value. That is,

elevation = $2x_1 + 5x_2$

Then, we want to find the point with highest elevation. Critically, our function is **linear**, meaning that in the context of our mountain, it is just a simple slope. Thus, consider an arrow pointing in the direction of the slope, and from an interior point within the region, trek the mountain as far as you can in the direction of the arrow until you cannot anymore. If you are at a corner, then great. If you are at an edge, you must be able to traverse the edge to get to a corner that has elevation *at least as much* as the point you reached on the edge.

This observation gives way to some major results.

- First, it was found that **any LP** is polytime solvable.
- In practice, the algorithm used is called the Simplex algorithm, which is something that you can even do by hand! It is very reminiscent of matrix row reduction.
 - The Simplex algorithm was developed by George Dantzig in 1947.
 - In the worst case, it runs in exponential time
 - But for most instances/problems, it is practically fast.

The **Simplex Algorithm** essentially performs the algorithm we came up with. Plot out the feasible region, and then test out all of the corner points and choose the location that maximizes/minimizes the objective function.

NOTE: In our example, we were dealing with only 2 variables, which allowed for easy visualization. But as we increase the amount of variables, the visualization becomes finding a feasible region in hyperspace.

It wasn't until 1984 that a new algorithm, Karmarkar's algorithm, was developed that was polytime and also practically fast. It leveraged a theory called interior-point methods, which we will not go in to.

2.1 Back to The Potluck

Let's return to the potluck problem from before. We can formalize it as an instance of a linear program. In order to do this, we must identify three things:

- 1. The variables
- 2. The Objective Function (what we want to maximize/minimize)
- 3. The constraints.

Variables are the quantities that the solver will give values to. Define your variables in a granular way so when the solver gives values to your variables, you either immediately solve your problem, or can easily derive the solution to your problem.

When we consider what our variables should be, we conclude that they should be be r, p, c, indicating the amount of rice, pasta, and couscous we purchase respectively.

Determining the objective function is straightforward: we want to minimize the price. The price, as a function of our variables is:

$$1.25 \cdot r + 1.65 \cdot p + 1.35 \cdot c$$

Lastly, we consider the constraints. First, we need ≥ 5000 calories, which can be expressed as:

 $750 \cdot r + 1200 \cdot p + 1000 \cdot c \ge 5000$

and we must also have ≤ 200 mg of sodium, expressed as:

$$15 \cdot r + 35 \cdot p + 60 \cdot c \le 200$$

Note that we also have an additional unspoken constraint of $r, p, c \ge 0$

Thus, our formalized LP is:

```
 \begin{array}{ll} \text{minimize} & 1.25 \cdot r + 1.65 \cdot p + 1.35 \cdot c \\ \text{subject to} & 750 \cdot r + 1200 \cdot p + 1000 \cdot c \geq 5000 \\ & 15 \cdot r + 35 \cdot p + 60 \cdot c \leq 200 \\ & r, p, c \geq 0 \end{array}
```

3 Max Flow

Let us now consider the Maximum Flow problem. This is a graph problem where, given a graph with edge capacities, source s and sink t, we wish to find the maximum amount of flow that can be sent from s to t so as to conserve flow (i.e. the flow coming into a vertex equals the flow going out of the vertex), and so as to not overflow an edge capacity. You can imagine this as starting at vertex s, and pouring water out from s along its outgoing edges, and any water that enters a node must leave the node. We want to figure out the most amount of water we can pour at the start. For more information, check out the first 3 minutes of this video.

The following graph shows a flow network.



Note: Solutions to the Flow problem are well known. For example, the Edmonds-Karp Algorithm solves the Max-flow problem and runs in time $\mathcal{O}(VE^2)$

3.1 Max Flow as LP

We choose to describe the Max Flow problem as an instance of LP.

First, we must choose our variables. While we might think that we should have a singular variable indicating the flow leaving the source, this proves to be a poor choice, as when we encode our constraints, they would all have to be with respect to this single variable.

We thus need more granular variables. And when we take this step deeper, we choose to have variables that represent the amount of flow on each edge. Thus, our variables are:

 f_{uv} = total flow along edge (u, v)

Our objective is to maximize the amount of flow we send out from our source s. Let $N_{OUT}(s)$ denote the neighbor set of s. That is, $N_{OUT}(s) = \{x \mid (s \to x) \in E\}$ contains all vertices adjacent to s. Then, our objective is:

maximize
$$\sum_{v \in N_{\text{OUT}}(s)} f_{sv}$$

Next, we turn to our constraints. The first constraint is that the amount of flow on each edge is non-negative and does not exceed its given capacity. Let c(u, v) denote the capacity on edge (u, v). Then, we have:

$$0 \le f_{uv} \le c(u, v) \quad \forall (u, v) \in E$$

The conservation constraint is a bit more difficult to work with. We want to say for a non-s, t vertex u, the amount of flow that enters u is equal to the amount of flow that exits u. Let $N_{\text{IN}}(u)$ denote the set of vertices such that $v \in N_{\text{IN}}(u)$ if $v \to u \in E$. Then, we have:

$$\sum_{\in N_{\rm IN}(u)} f_{vu} = \sum_{v \in N_{\rm OUT}(u)} f_{uv} \quad \forall u \in V - \{s, t\}$$

Or equivalently:

$$\sum_{v \in N_{\text{IN}}(u)} f_{vu} - \sum_{v \in N_{\text{OUT}}(u)} f_{uv} = 0 \quad \forall u \in V - \{s, t\}$$

And with this, we have a fully formed LP that we can go ahead and solve.

v

v

4 Integer (Linear) Programming

In the world of linear programming, one massive assumption we made was that fractional values are acceptable. And this is okay.

Integer Linear Programs (ILPs) are linear programs with the *additional constraint* that variables must take integer values.

Visually, recall our "feasible region" from before. Whereas before, any point within our region could be a potential solution, in ILP, the only potential solutions are points within the feasible region **that also have integer coordinates**.

Why would we ever care for an ILP when we have LP? Well, in the real world, items often come in discrete units. From our potluck example, suppose the information was given with respect to "bags" of the item instead of kilograms. It doesn't make much sense to buy 3.2 bags of rice.

ILP Example

Suppose we have 45 dollars to buy pies (5 dollars each) and cakes (9 dollars each). We can resell pies for 5 dollars of profit, and cakes for 8 dollars of profit. But we can only buy at most 6 items. How can we maximize profit?

We can write this as an instance of an ILP:

S

maximize	5p + 8c
ubject to	$5p + 9c \le 45$
	$p+c \le 6$
	$p,c \geq 0$
	$p,c\in\mathbb{Z}$

4.1 Solving ILPs?

How do we solve ILPs? One naive idea is to use our solver for LP! After all, ILP and LP are very similar problems, and we already have a solution to LP.

Suppose we do this, and our solution returns the optimal values for our variables, but some of the variables are fractional when they need to be integers. What do we do? We may choose to round these values to the nearest integers that remain within the feasible region. This doesn't work.

- Bad News 1: We can construct ILPs whose rounded LP solution is *arbitrarily* far away from the optimal.
- Bad News 2: ILP is NP-complete!

5 Mixed Integer Program

Often, we will generalize ILP to **MIP** (Mixed Integer Programming). With MIP, some variables may be constrained to integers, and some may not. Still, all objective functions and constraints are linear. More on MIP next week...