GUROBI OPTIMIZATION

The Fastest Solver in the World

Mathematical Optimization: What You Need to Know

www.gurobi.com > info@gurobi.com > 1 (713) 871-9341

Table of Contents

Mathematical Optimization: What You Need to Know

GUROBI OPTIMIZATION

Introduction

Mathematical optimization (MO) is an extremely powerful AI technology that enables companies to dramatically improve their resource utilization, operational efficiency, and overall profitability. MO technologies such as linear programming (LP) and mixed-integer programming (MIP) have been applied in a variety of business areas, empowering companies across various industries to rapidly solve their complex real-world problems, make better business decisions that improve operational efficiency, and achieve significant time and cost savings and revenue growth. According to a recent study by Forrester, 37% of data science decision makers at U.S. firms use MO frequently in their work today.

Here are a few examples of companies using MO technologies:

- The National Football League (NFL) uses optimization to solve one of the most challenging scheduling problems in existence.
- Air France estimates that it is saving around 1 percent of fuel costs for its entire fleet using optimization.
- The Federal Communications Commission (FCC) used optimization and operations research to repurpose wireless spectrum, thereby generating nearly $20 billion in revenue, with $7.3 billion going to reduce the federal deficit.

In this whitepaper, we will provide a general overview of MO, share some success stories of companies that have achieved tremendous operational and bottom-line benefits using MO technologies, explain the relationship between MO and artificial intelligence (AI), and discuss the importance of combining MO with machine learning (ML).

An Overview of Mathematical Optimization

Overview

A large number of organizations have successfully implemented MO technologies to optimize their operational efficiency and realize their business objectives while taking into account their unique resource demands, business rules, and constraints.

With MO, organizations can define and embed the key features of their complex, real-world problems in MO programming models.

MO programming models capture the important decisions in business processes. For example, should I:

- Use my resources to build product A or B?
- Send this truck to San Francisco or Dallas?
- Include this stock or a different one in my investment portfolio?

MO programming models capture the resources that are required to execute these decisions. For example:

- Building this product uses resources such as machines and employee time.
- Sending a truck consumes time, fuel, and manpower.
- Capital invested in one stock can no longer be invested in another stock.

MO programming models capture the potential conflicts between these activities. For example:

- The same machine is required to build both product A and B.
- If I send a truck to San Francisco, I can't send it to Dallas.
- The maximum capital that can be invested is $500,000.

In these and other examples, MO factors in all the relevant variables, constraints, and business objectives, instantly processes and evaluates an astronomical number of possible combinations of these factors (using a math programming solver), and then suggests a plan of action that optimizes the overall efficiency of a particular business process.

In order to optimize the efficiency of a business process, MO optimally balances the trade-offs among numerous (and often conflicting) business objectives, such as:

- Maximizing revenue
- Minimizing costs
- Minimizing delays
- Maximizing customer satisfaction

No matter what your business objectives are or how complex your business processes or problems are, MO is capable of automatically and rapidly generating an optimal solution – and this output serves as the input for your business decisions.

An Overview of Mathematical Optimization

GUROBI OPTIMIZATION

MO technologies have been utilized by a wide variety of businesses across different industries. For example:

- **Supply chain optimization:** The Gurobi Optimizer was seamlessly integrated into SAP applications – including SAP Integrated Business Planning, SAP Advanced Planning and Optimization, and SAP HANA – to help solve complex optimization problems.
- **Electrical power optimization:** Managing New York's wholesale electrical power market ($7.5 billion in annual transactions), the New York Independent System Operator (NYISO) used Gurobi's Compute Server to optimize 500 power-generation units and 11,000 miles of transmission lines with consumer demand in real time.
- **Government:** The FCC used Gurobi to create the first two-sided spectrum auction, generating $20 billion in revenue.

Some of the most successful companies in the world use the Gurobi Optimizer to tackle high-impact and complex business problems. Such problems include, but are not limited to, resource allocation, scheduling optimization, routing, logistics, and supply chain issues. Below is a sample of Gurobi customers by industry:

AIRLINES & AIRPORTS
Lufthansa
AIRFRANCE

AUTOMOTIVE
DAIMLER
Ferrari
YUTONG

SPORTS SCHEDULING
NFL

COSMETICS
ESTĒE LAUDER

METALS MATERIALS & MINING
bhpbilliton

ENERGY
ISO
ProCom
WM WASTE MANAGEMENT

FINANCIAL SERVICES
Bank of America
TransUnion
Prudential
BM&F BOVESPA The New Exchange

CONSUMER ELECTRONICS
HUAWEI

ELECTRICAL EQUIPMENT
ABB

FOOD AND BEVERAGE
STARBUCKS COFFEE

GOVERNMENT
DEPARTMENT OF DEFENSE UNITED STATES OF AMERICA

CONGLOMERATE
SIEMENS

TRANSPORTATION
SCHNEIDER LOGISTICS
LEANLOGISTICS
CSX HOW TOMORROW MOVES

MEDIA
nielsen

RETAIL
Walmart Save money. Live better.

SEMICONDUCTORS
AMD

HOSPITALITY
Marriott

TECHNOLOGY
SAP
UBER
Google
ORACLE
Microsoft

TELECOM
FCC Federal Communications Commission
Nokia Siemens Networks
AT&T

Mathematical Optimization Applications

Success Stories

In this section, we share three success stories of organizations utilizing MO to realize significant business benefits. These stories are based on four case studies, which can be found <u>here</u> on the Gurobi website.

Air France: Tail Assignment Optimization

Air France-KLM is one of the world's largest air carriers in terms of revenue and passengers transported. The company's operations research department – which strives to optimize all the carrier's operations by transforming data into smart decisions – was facing a "Tail Assignment Optimization" problem, which consists of assigning flights to aircraft while respecting operational constraints. This means building a schedule for each aircraft by setting the sequence of flights that an individual aircraft will perform.

To solve this "Tail Assignment Optimization" problem and build the most efficient schedule for its entire fleet, its operations research department decided to apply MO in order to save on fuel and operational costs, while reducing delay propagation.

To build the most efficient schedule for the entire fleet, trade-offs among different criteria and the possible impact need to be considered. This makes "Tail Assignment Optimization" a multi-criteria mixed-integer programming (MIP) problem.

The operations research team built a decision support tool, powered by the Gurobi Optimizer, that recommends a solution that internal users can use to fit their specific needs. The tool provides an optimal solution with respect to the given constraints and the criteria, but the user still has the flexibility to make decisions based on experience, if needed.

Solene Richard, Air France-KLM's Data Science and Operational Research Team Leader, commented: "With this new decision support tool powered by the Gurobi Solver, we estimate that we are saving around 1 percent of the fuel costs with huge volumes for the entire fleet. We also estimate that we save on delay propagation and on operational costs. Using mathematical optimization to solve the tail assignment problem brings large savings on a yearly basis."

Mathematical Optimization Applications

GUROBI OPTIMIZATION

National Football League (NFL): Football Season Scheduling

The NFL – one of the major sports leagues in the U.S. – uses MO to tackle its highly complicated scheduling challenges. At first glance, the NFL's scheduling problem seems simple: four people have 10 weeks to schedule 256 games over the course of a 17-week season. That might seem like plenty of planning time for seemingly few decisions, but – when you actually work it out – the number of possible schedules is well into the trillions. Making the problem particularly hard is the necessary inclusion of numerous additional constraints, including not scheduling teams from shared markets – such as the New York Jets and New York Giants – to play on the same day and time, and increasing the importance of late-season games by scheduling as many divisional match-ups as possible in the final weeks.

Until the last decade, a wooden 6-foot-square board was used to map out who would play whom, when, and where. The original goal was simply to find a feasible schedule, and accomplishing that often consumed up to 90 percent of the available planning time, with the remainder spent tweaking that feasible schedule to try to improve it. With the switch to computers and the increasing power of MO technologies, the focus turned from schedule creation toward schedule analysis.

This shift is important, since building a great schedule can be as much art as science. For example, one goal of the scheduling process is to get the best match-ups in the TV time slots that have the largest possible audience. Achieving this is an immense challenge for the NFL's planners because it requires a subjective evaluation of the match-up while also ensuring the resulting schedule doesn't violate any constraint, such as the maximum number of away games in a row for a given team. With MO, the NFL's planners can now evaluate any given schedule, identify what they want to change, conduct scenario analysis to determine the potential impact of these changes on the overall schedule, and decide if they want to implement the changes.

The complexity of the NFL's scheduling problem has increased exponentially over the past few years, necessitating a shift from a linear approach to a parallel approach to optimization. The NFL initially started with a single 24-core box, but a few years ago moved to utilizing a roomful of 16-core servers. A key advantage of the parallelization approach is that a single problem is broken down into a finite number of smaller sub-problems, which are deployed to a pool of solving resources that share information – thus creating a much more coordinated search effort. In 2017, the NFL scheduling problem was solved across a network of more than 960 cores. Additionally, the NFL has moved to a cloud environment, which provides a tremendous amount of computing power while minimizing the need for capital investment.

Since adopting Gurobi in 2013, NFL planners are able to generate and compare more than 10,000 feasible schedules despite adding more constraints to the process every year.

Mathematical Optimization Applications

GUROBI OPTIMIZATION

blueyonder

Blue Yonder GmbH: Automated Daily Retail Pricing Decisions

Blue Yonder – a leading provider of cloud-based AI solutions for the retail industry – specializes in using ML algorithms to facilitate decisions that increase sales, reduce write-off rates, and boost profits. Blue Yonder helps retail companies automate pricing decisions that are influenced by historical data and other critical factors, such as weather, public holidays, and competitor information.

Determining these dynamic price adjustments, however, is a complicated endeavor – requiring the efficient interplay of algorithmically demanding disciplines. Another challenge is automatically generating reliable sales forecasts in an ever-changing market environment, which requires the deployment of powerful ML models. Furthermore, when automating these pricing decisions, various considerations – such as the price trend in a product's life cycle, price developments in the competitive environment, as well as local factors such as time period and cost of price changes – must be taken into account.

Blue Yonder Price Optimization is a state-of-the-art automated pricing decision solution that optimizes prices for every channel and every product according to consumer demand, brand loyalty, and competitive advantage. The solution "learns" the relationship between price changes and demand while incorporating the retailer's business strategy. Blue Yonder Price Optimization rapidly senses vital demand signals from changing market conditions and data such as sales, promotions, weather, and other events. It serves the retailer's pricing strategy along the product's life cycle and measurably impacts revenues and return on investment.

With Blue Yonder Price Optimization, retailers can calculate new sales forecasts and generate up-to-date price strategies every night using the most recent data so that they have the results ready by the next morning.

To create the optimization models, product quantities must be divided into suitable clusters. About 50,000 variables are generated for each cluster for the pricing decisions and numerous business rules are converted into approximately 1 million constraints. For specific optimization processes over the product's entire life cycle, MIP models with more than 1.2 million variables are created. Depending on the uncertainties in the input data, the termination criteria of the optimization also change. A trade-off between available time and the desired solution quality is thus defined for each customer.

With Gurobi, a stable solution time and quality can be achieved consistently even when changing the input data. Usability was also a point in favor of the Gurobi Optimizer during the solver selection process. Using the Gurobi Python API, a prototype was developed in a very short time. In addition, by using Gurobi Compute Server, the optimization functionality could be seamlessly integrated into Blue Yonder's service infrastructure.

With the use of Gurobi and Blue Yonder Price Optimization, Blue Yonder's customers have been able to increase product sales by 5 percent and boost overall profitability. These companies' price adjustments – which are based on the most current data, and consider company objectives and the cost of changing prices – now take place in real time. Additionally, price reduction optimization has enabled these companies to decrease their inventories by 20 percent.

History

MO technologies – whose origins stretch back to 1947 with the invention of linear programming by George Dantzig – have evolved tremendously over the years.

On a fundamental level, MO is a methodology that consists of three steps:

- The formulation of a real problem as a mathematical model
- The development of algorithms to solve these mathematical models
- The use of software and hardware to run these algorithms and to develop mathematical programming applications

Although this essential paradigm has never changed, we have seen tremendous development and improvements in terms of the speed, flexibility, and power of MO technologies in the past few decades.

> "An optimization business problem that can be solved today in one second would have taken 55 years in 1991."

For example, an optimization business problem that can be solved today in one second would have taken 55 years in 1991, marking an overall improvement in performance speed of computers and mathematical programming algorithms of 1.75 billion times.

Beyond speed, MO technologies have also improved in their ability to handle business problems of increasing scale, complexity, and diversity. The use of MO started in the 1960s when oil companies implemented LP and MIP approaches (and saved millions of dollars). It spread in the 1970s with the availability of commercial codes running on mainframe computers and became widely adopted in the 1990s by companies in a wide range of industries including manufacturing, electrical power, and supply chain software. Today, MO has become an essential technology for companies around the world and across industries for a variety of use cases such as operational applications, business processes, logistics, pricing, and planning – and its popularity (as well as its speed and power) will undoubtedly continue to increase in the years to come.

DDA Approach

MO technologies utilize the data-driven analytics (DDA) problem-solving approach. So, before examining the details of MO technologies, we will describe the DDA approach and position MO within that context.

When attacking a business problem, we must have a deep understanding of the problem's root causes, be able to clearly define the problem being addressed, be able to describe why solving this problem is important, and be able to deploy a tool that can both be used by the company and deliver business value.

The DDA approach – which aims to achieve these objectives – involves the following steps:

1. Business understanding
2. Data understanding
3. Data preparation
4. Modeling
5. Evaluation
6. Deployment

The DDA approach is illustrated in the diagram below.

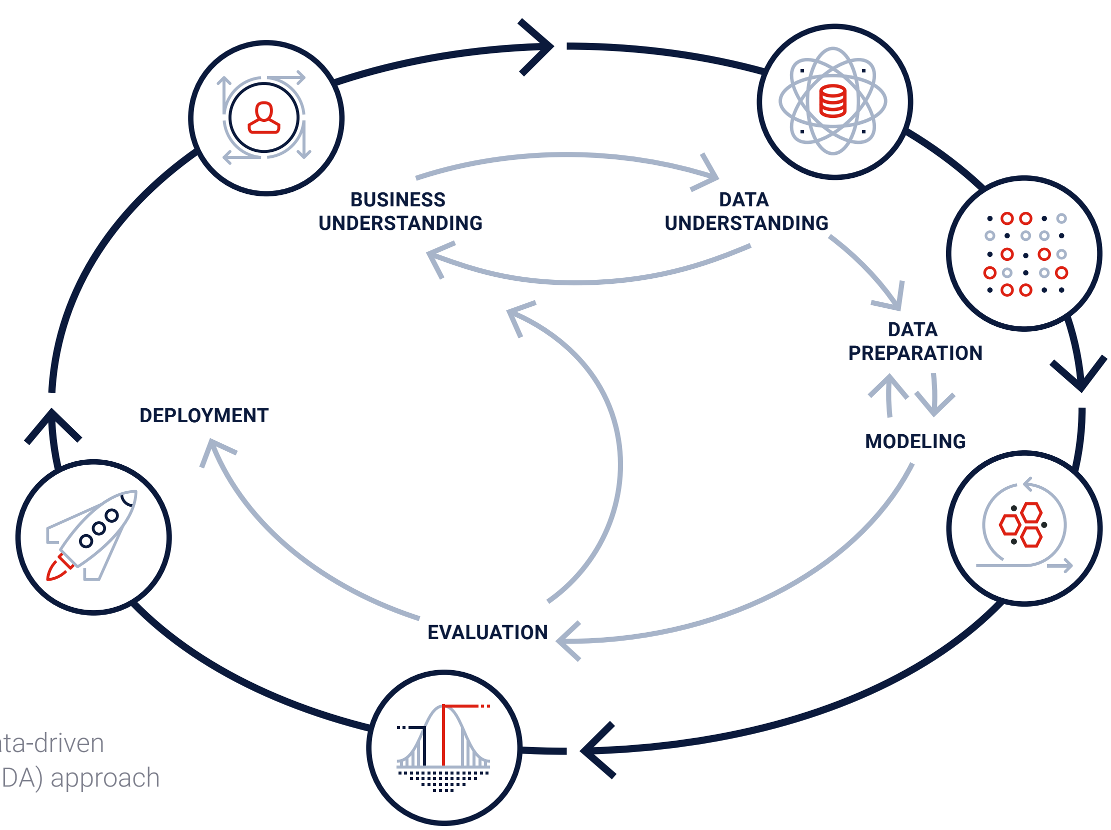

Figure 1: Data-driven analytics (DDA) approach

Data-Driven Analytics Approach

In the following sections, we take a closer look at each step of the DDA approach.

Business Understanding

In order to understand the root causes of the problem under consideration, we need to define:

- The business processes around the problem
- Business objectives and associated key performance indicator (KPI) metrics
- The decision makers

Data Understanding

Next, we must be able to describe the sources of data in detail and determine the quality of the available data in terms of:

- Availability
- Accuracy
- Completeness
- Timeliness

Data Preparation

In this step, available data is cleaned and, if the quality is low, new sources of data are identified in an effort to improve quality.

Modeling

During modeling – the most critical step in the DDA approach (which will be discussed in detail in the next section) – we identify the modeling techniques that will enable us to optimize the trade-offs among various business objectives. For each model considered, we describe the underlying assumptions, input parameters, output variables, model configuration parameters, and model characteristics, including constraints and objective functions in the case of an MO model. We also describe data preprocessing steps to create input data for the model and model output postprocessing steps to prepare data for reports and visualization.

Evaluation

During this step, the models' output reports are evaluated by the decision makers and business analysts who will use the tool. Based on their feedback, the previous steps may be revised to better satisfy business needs and expectations.

Deployment

After the decision makers and business analysts have evaluated the tool and determined that it satisfies all their business needs and expectations, a support and maintenance plan is created and the tool is moved to production.

MO technologies utilize this DDA problem-solving methodology. To successfully deploy an MO tool, it is imperative to follow all the steps in the DDA approach highlighted above.

By adhering to this DDA approach, companies can capture all the key features of a complex real-world business problem as an MO model and build decision support tools that enable them to optimize their operational efficiency.

MO Models

In this section, we describe the overall structure of an MO model in detail. The following figure illustrates this structure.

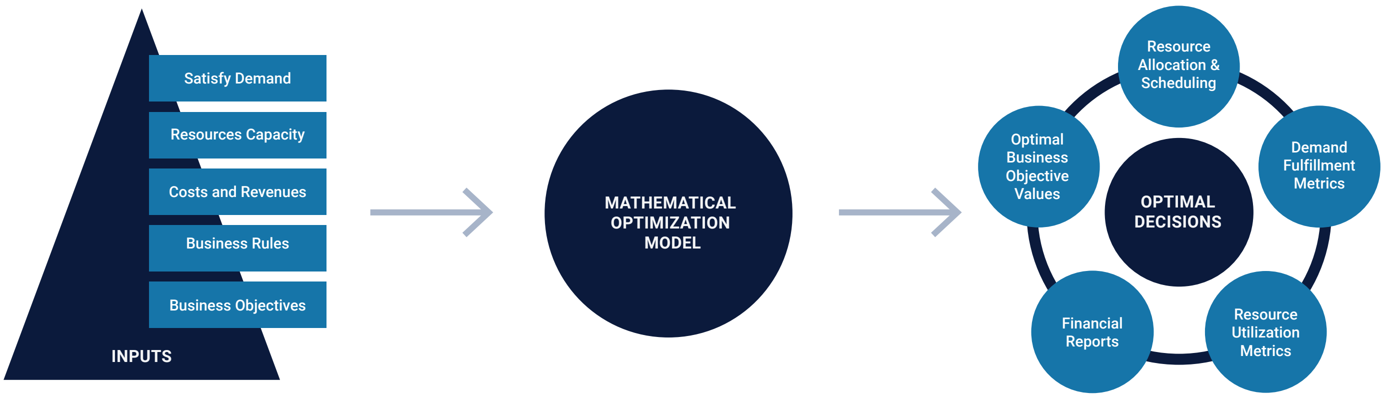

Figure 2: Overall structure of an MO model

An MO model has a set of deterministic input parameters that describe an instance of the problem we want to solve. For example, a manufacturing company has estimated the demand curve (quantity, price) of each product that a retail company can advertise and sell. Total revenue can be calculated from the expected quantity to be sold and associated unit price. In addition, consider that there is a budget to buy the raw materials required to build the finished products and that there is a buying cost of each raw material. Also, consider that the ability to build and sell a final product is limited by available labor hours and machine capacity. In a manufacturing setting, a company typically has initial inventories of raw materials and finished goods. One objective the company may have is to maximize profits. Lastly, there are business rules that the company must adhere to. For example, certain products are complementary, so they need to be built together.

The decision variables that define the possible courses of action are another component of an MO model. In our example, the decision variables are the number of finished products to build, the number of raw materials to buy, and the price at which a certain quantity of finished products must be sold.

In addition, the MO model has constraints that establish relationships between the decision variables and input parameters. In our example, there are constraints that limit the production of finished products based on the availability of labor and machines as well as constraints that enforce the complementary rule of finished goods.

There are also balance equations that relate initial inventories of raw materials and finished goods with their associated ending inventories.

Finally, the MO model has an objective function (defined in terms of decision variables and input parameters) that needs to be maximized or minimized. In our example, the objective function is to maximize total profits generated from selling finished products. This total profit is the difference between total revenues from selling finished products and total costs from buying raw materials.

The MO model generates a production plan that shows how much of each finished product to build, a sales plan that sets out how much of each finished good is expected to be sold at which price, and a procurement plan that establishes how much raw material must be obtained to support the production plan.

By combining input parameters with decision variables, we can determine demand fulfillment metrics such as the percentage of forecasted demand satisfied by the sales plan. Also, we can determine resource utilization metrics, such as what percentage of available resources (labor and machines) must be allocated to the sales plan.

Financial reports can be created that define the costs associated with the procurement plan as well as the ending inventory holding costs of raw materials and finished products. Also, reports can be created that define ROI, profits, and revenues associated with the sales plan.

Types of Mathematical Optimization Models

There are various types of MO models, each of which can be defined based on the different characteristics of its decision variables, constraints, and objective function.

Linear Programming (LP) Models

For LP models, we have continuous non-negative decision variables, linear constraints, and a linear objective function. An example of non-negative continuous decision variables are gallons of gasoline produced at a refinery.

The objective function can be maximized or minimized. For example, the model can maximize total revenue or minimize total costs.

There are several types of linear constraints:

- **Less-than-or-equal constraints** are typically considered for capacity constraints where we don't want to exceed available capacity.
- **Greater-than-or-equal constraints** are used to model demand requirements where we want to ensure that at least a certain level of demand is satisfied.
- **Equality constraints** are used when we want to match certain activities exactly with a given requirement. For example, a job position can be filled with only one resource and you have a set of possible qualified resources to assign to the job.

Mathematical Optimization Models

GUROBI OPTIMIZATION

Integer Programming (IP) Models

IP models are like LP models except that the decision variables take integer values only. For example, in the automotive industry, the production of cars requires decision variables to be non-negative integer numbers.

Binary Programming (BIP) Models

BIP models are a particular case of IP models that require that the decision variables take binary values only – that is, 0 or 1. BIP models are used to formulate combinatorial optimization problems. Let's look at a classic combinatorial optimization problem called the Traveling Salesman Problem (TSP). For this problem, consider a salesman who needs to visit *n* cities. This person wants to visit all the cities only once and wants to minimize the total distance of going from one city to another. The decision variable $x_{i,j}$ is equal to 1 if city *j* is visited immediately after visiting city *i* and it is 0 otherwise, so the TSP's decision variables are binary.

Mixed-Integer Programming (MIP) Models

MIP models are a combination of IP/BIP and LP models, where the set of decision variables can be divided into two subsets: one subset of decision variables that are continuous and non-negative, another subset of decision variables that are non-negative integers or binary. For example, let's consider a production process that has a variable linear cost of production and a onetime cost of setting up the production process. The amount to be produced is a non-negative continuous decision variable, and the cost of the production process setup is a binary decision variable.

Nonlinear Programming (NLP) Models

In nonlinear programming models, some of the constraints or the objective function (or both) are nonlinear functions. For example, let's consider a set of stocks. Each stock has a price and a return on investment. In addition, there is a forecast of the stocks' covariance matrix reflecting the risk associated with investment in this set of stocks. The decision variables are the amount of each stock to buy; the objective function is to minimize the investment risk while satisfying a minimum overall return on investment. The objective function is a quadratic function defined by the stocks' covariance matrix with a linear constraint reflecting the minimum overall return on investment.

AI Technology

In many of the world's leading companies, AI is perceived as a force of good. AI has the potential to make the world safer, healthcare more accessible, education more personal, and manufacturing more efficient. In the coming years, AI technologies are expected to have a profound (and primarily positive) impact on virtually every aspect of society and revolutionize the way businesses operate.

Forrester reports that as of 2019, more than 50% of global enterprises have already implemented AI technologies and that around 80% of global enterprises will be using AI in five years.

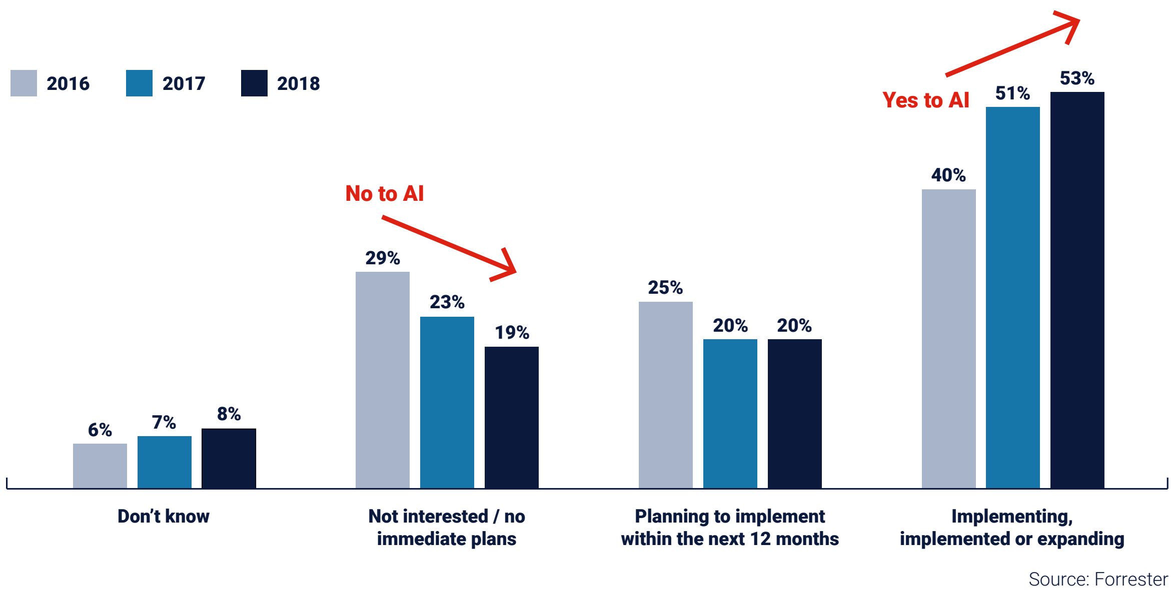

Figure 3: What are your firm's plans to use artificial intelligence?

Mathematical Optimization as an AI Technology

The business world has been transformed over the past 50 years by the introduction of computers and information technologies and continues to change due to:

- The continued growth in available data containing a wealth of valuable information
- Techniques such as analytics, ML, and MO, which make it possible to extract significant value from data
- Computing resources that can handle real-word business problems of any size or complexity

AI Technology

When considering mathematical optimization within the context of AI technologies, the key question is:

- Is mathematical optimization AI?

To answer this question, we must explore the following questions:

- How do you extract value from data?
  – By drawing conclusions.
- How do you draw conclusions?
  – With inductive reasoning, deductive reasoning, or both.

Inductive reasoning means deriving general principles from specific observations and deductive reasoning means deriving conclusions from a set of factual premises. We may say that AI is a technology that attempts to replace human decision-making. Like ML, *inductive* AI derives conclusions and makes recommendations based on patterns discovered in the data. Like MO, *deductive* AI derives conclusions and makes recommendations based on known relationships (mathematical models) and the available data.

Summarized simplistically, ML is an inductive approach that considers very large data sets and attempts to discover patterns that provide insights on the data. One critical characteristic of the ML approach is that data and algorithms are seen as a single entity where special effort is made to keep the data accurate, current, and complete. For this purpose, the data scientist identifies many sources of data.

The types of decisions handled by ML are narrow in scope – yes/no, on/off types of decisions. ML is particularly good at applying similarity scores to measure how similar two entities are. Facial recognition is a great example of a technology where ML can be applied.

However, one big issue in the future of decision technology is the combinatorial explosion in possible courses of action that can be taken – and MO (as a powerful deductive AI technology) excels in handling problems of immense complexity and size. For example, if you consider the previously mentioned Traveling Salesman Problem for 48 U.S. cities, there are 10^{61} solutions. That's actually larger than the number of atoms on earth (10^{50}). MO technology finds a provable global optimal solution in a fraction of a second.

AI Technology

MO is a declarative approach where the modeler formulates a mathematical model that captures the key features of a complex decision problem. The MO model formulation is then solved with MO algorithms, such as LP and MIP algorithms. The Gurobi Optimizer utilizes state-of-the-art algorithms to tackle MO models, while a decision support system can be built based on the MO model formulation and the Gurobi Optimizer.

> "With MO, changes in business conditions are easily reflected by new decision variables and constraints without the need to change the underlying mathematical programming algorithms."

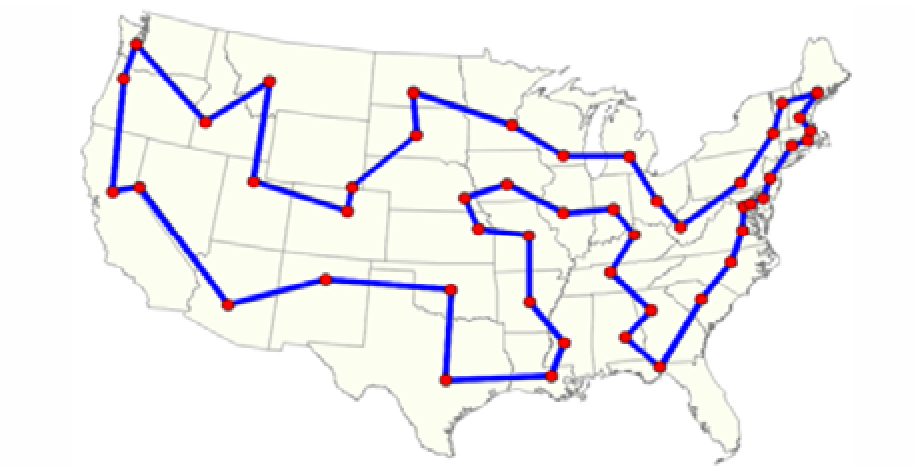

Figure 4: The global optimal solution to the Traveling Salesman Problem

Mathematical Optimization and Heuristics

In problem solving, a heuristic technique is any approach that employs a practical method not guaranteed to produce a global optimal solution. The idea of heuristic methods in AI is based on cognitive science, the study of how humans think. Indeed, humans use heuristics all the time to make decisions and solve problems. Likewise, heuristic algorithms are often used in AI to get a computer to find an approximate solution instead of an exact solution.

In AI, a metaheuristic is a higher-level procedure designed to find, generate, or select a partial search algorithm that may provide a sufficiently good solution to an optimization problem. Metaheuristics sample a set of solutions too large to be completely sampled. Metaheuristics may make a few assumptions about the optimization problem being solved and different metaheuristics are inspired by phenomena occurring in nature.

For example, a genetic algorithm is a metaheuristic based on natural selection, the process that drives biological evolution. Simulated annealing is a metaheuristic inspired by a thermodynamics process of heating and then slowly cooling metals. Compared to optimization algorithms, metaheuristics are fast but do not guarantee a global optimal solution.

As noted, the MO approach only requires a modeler to formulate a mathematical model that captures the key features of a complex decision problem. The modeler does not need to worry about how to solve the optimization problem at hand, as this is done automatically, behind the scenes. The modeler only needs to have an efficient MO model that captures the main characteristics of the optimization problem and the required data for the model. MO is easier to maintain for business applications – as business conditions change frequently. These changes are difficult to incorporate with heuristics, whereas with MO these changes are easily reflected by new decision variables and constraints without the need to change the underlying mathematical programming algorithms.

Other benefits of using a mathematical programming solver with respect to ad-hoc heuristics or metaheuristics include the following:

- An MO solver provides a systematic solution with bounds on how close that solution comes to the optimal one
- An MO solver can be used heuristically to find a provable good solution in a reasonable time
- Heuristics can be used to find feasible solutions to determine bounds and enhance the pruning of the solution search tree of an MIP approach.

“The next generation of AI applications to solve high-impact and high-value business problems will combine ML and MO techniques.”

Mathematical Optimization and Machine Learning

Machine Learning algorithms have been very successful at making predictions. For example, ML can:

- Predict supply chain issues while there is still time to remediate
- Predict who will launch a cyberattack before it happens
- Avoid wasting time by predicting experiments that are more likely to prove the hypothesis
- Predict imminent machine failure
- Predict benefits eligibility fraud
- Predict future infrastructure maintenance needs
- Predict price movements to find investment opportunities before the market does
- Predict customer propensity to buy more with targeted offers

AI Technology

Understand that ML is not without challenges. For example, ML models require a powerful infrastructure, and even though they can be formidable and very profitable, these models are about probabilities, not absolutes. Consequently, accurate models may not exist for every question.

ML models are based on correlation and possibility, but they are not causative. In addition, ML models must learn from experience and their performance can decay over time, so they must be retrained on newer data to stay fresh.

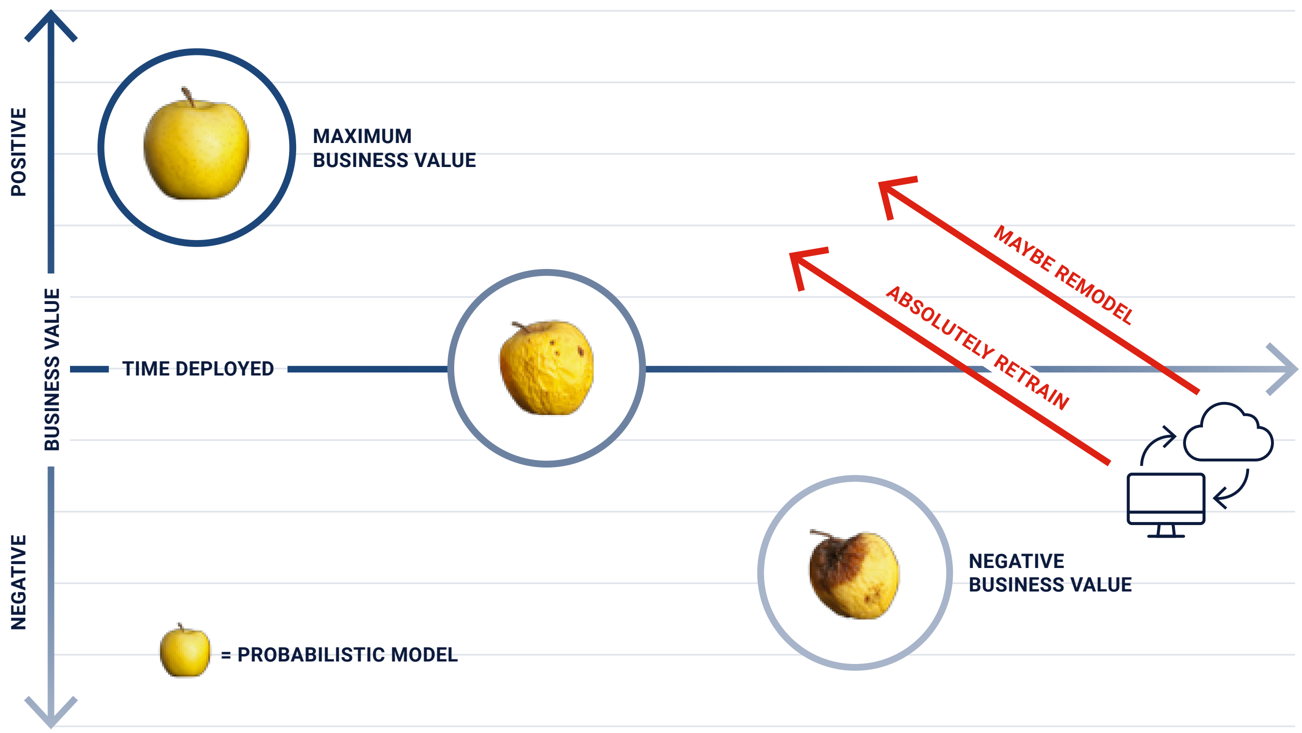

Figure 5: ML models over time

AI Technology

In a business environment, being able to make predictions is clearly not enough. Business executives need to have the capability to use these predictions to help them decide on the best course of action. As we have seen, MO determines the best possible decisions based on real-world constraints.

To illustrate, let's describe the types of decisions that can be made based on the predictions from the ML success stories highlighted above:

- ML can predict supply chain issues while there is still time to remediate.
  – MO can decide the least costly way to reroute shipments.
- ML can predict who will launch a cyberattack before it happens.
  – MO can decide which investigators to assign to potential cyber threats based on investigator skill and potential damage.
- ML can avoid wasting time by predicting experiments that are more likely to prove the hypothesis.
  – MO can decide what experiments to pursue based on talent, cost, and time.
- ML can predict imminent machine failure.
  – MO can decide when to shut the production line down to perform maintenance so as to minimize cost and maximize customer satisfaction.
- ML can predict benefits eligibility fraud.
  – MO can decide how to assign case workers to maximize recovery.
- ML can predict future infrastructure maintenance needs.
  – MO can decide how to assign maintenance teams based on cost and skill.
- ML can predict price movements to find investment opportunities before the market does.
  – MO can decide how to allocate cash across all investment vehicles.
- ML can predict customer propensity to buy more with targeted offers.
  – MO can decide how many discount coupons to offer for maximizing revenue or profit.

Yet MO, too, has its challenges. For instance, business problems must be translated into decision variables, mathematically expressed constraints, and an objective function – and not every organization possesses the technical expertise to do this. MO can also be slow without performant solver software.

In summary, we believe that the next generation of AI applications to solve high-impact and high-value business problems will combine ML and MO techniques. ML predictions can determine the need to make MO decisions and can be used as MO decision constraints. MO decisions can be used as ML model training features and as ML model scoring and inferencing inputs.

Conclusion

Maximize Efficiency

In this document, we established that MO is a mature problem-solving technology that is capable of maximizing an organization's operational efficiency while satisfying resource constraints and business rules. MO allows companies to capture the key features of a complex real-world business problem as an optimization model and enables them to build decision support tools that drive greater efficiency, revenue growth, and cost reductions.

We examined the tremendous impact that operations research and MO have on society. For example, the cumulative benefits from the projects of Franz Edelman competition finalists have topped the $292 billion mark. Leading global companies are using the Gurobi Optimizer to tackle complex problems such as resource allocation and scheduling, routing, logistics, and supply chain management.

Next, we presented three case studies that illustrate how various companies have used the Gurobi Solver to achieve significant business benefits.

Contact us

Do you want to see how much faster Gurobi is than your current solver?

We are happy to review your model and help you get the performance you need.

info@gurobi.com

1 (713) 871-9341

www.gurobi.com

After that, we examined the various steps in the DDA approach, which is used when devising MO models and deploying MO solutions.

We presented the key components of MO (data, decision variables, constraints, and objective functions) and discussed the associated output reports to help decision makers choose the best course of action. In addition, based on the characteristics of the decision variables, constraints, and objective function, we defined types of MO models.

Finally, we discussed AI technology's relationship with MO. We looked at two AI technologies, metaheuristics and ML. Most important, we discussed how ML and MO can be combined to tackle complex and critically important business challenges.

www.gurobi.com > info@gurobi.com > 1 (713) 871-9341