2019-01-01

Towards Analytical Techniques for Systems Engineering Applications

Griselda Valdepeñas Acosta
University of Texas at El Paso

TOWARDS ANALYTICAL TECHNIQUES FOR SYSTEMS ENGINEERING APPLICATIONS

GRISELDA VALDEPEÑAS ACOSTA

Doctoral Program in Electrical and Computer Engineering

APPROVED:

Eric D. Smith, Ph.D., Chair

Vladik Kreinovich, Ph.D., Co-Chair

Deidra Hodges, Ph.D.

Bill Tseng, Ph.D.

Stephen Crites, Ph.D.
Dean of the Graduate School

TOWARDS ANALYTICAL TECHNIQUES FOR SYSTEMS ENGINEERING APPLICATIONS

by

GRISELDA VALDEPEÑAS ACOSTA, M.Sc.

DISSERTATION

Presented to the Dissertation Committee

The University of Texas at El Paso

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

Doctoral Program in Electrical and Computer Engineering

THE UNIVERSITY OF TEXAS AT EL PASO

December 2019

Acknowledgements

First, and foremost, I would like to express my profound gratitude to my parents Griselda and Felipe H. Acosta and to my siblings Rolando, Felipe, and Adriana. I also wish to thank my Committee Members, Dr. Eric D. Smith, Dr. Vladik Kreinovich, Dr. Deidra Hodges, and Dr. Bill Tseng for their unwavering help, support, and encouragement.

I would like to thank all the faculty, staff, and students of the Department of Electrical and Computer Engineering and of the Systems Engineering Program at The University of Texas at El Paso for all their hard work, dedication, and help.

NOTE: This dissertation was submitted to my Supervising Committee in Fall 2019

Abstract

One of the main objectives of systems engineering is to design, maintain, and analyze systems that help the users. To design an appropriate system for an application domain, we need to know: what are the users' desires and preferences (so that we know in what direction we should aim to change this domain), what is the current state and what is the dynamics of this application domain, and how to use all this information to select the best alternatives for the system design and maintenance.

Designing a system includes selecting numerical values for many of the parameters describing the corresponding system and its subsystems. At present, in many cases, this selection is made by consulting experts and/or by following semi-heuristic recommendations (recommendations based partly on the past experience of system design and monitoring). Experience shows that such heuristic imprecise recommendations often lead to less-than-perfect results. It is therefore desirable to come up with analytical techniques for system design, techniques that would be based on valid numerical analysis and on the solution of the corresponding optimization problems.

System engineering is a very broad discipline, with many different application domains. Each domain has its own specifics and requires its own analysis and, probably, it own analytical techniques.

In this dissertation, we formulate and analyze *general* problems corresponding to different stages of system design, implementation, testing, and monitoring, and show, on appropriate examples, how the corresponding analytical techniques can be applied to different application domains. Examples of our applications range from biological and biomedical systems (ranging from cows to humans) to social-related systems (such as recommender systems) to physical systems (for which we provide a new system-based explanation for the minimum entropy principle) to engineering systems (for which we describe how to find the optimal proportion of testing on different levels of system design).

Table of Contents

Page

Acknowledgements v

Abstract vi

Table of Contents vii

Chapter

1 Formulation of the Problem 1

2 Analytical Techniques for Describing User Preferences: Justification for (and Extension of) the Matrix Factorization Technique 7

2.1 Formulation of the Problem 7

2.2 Why Matrix Factorization: Our Explanation 9

3 Analytical Techniques for Describing User Preferences: 80/20 Rule Partially Explains 7 Plus Minus 2 Law:
General System-Based Analysis 13

3.1 Formulation of the Problem 14

3.2 Our Explanation 14

4 Analytical Techniques for Analyzing Probability Distributions: How to Explain That Changes in Elderlies Depression Level Are Uniformly Distributed 17

4.1 Formulation of the Problem 18

4.2 Our Explanation 19

5 Analytical Techniques for Analyzing How Systems Change with Time: A Natural Explanation for the Minimum Entropy Production Principles 22

5.1 Formulation of the Problem 23

5.2 How Complex Problems Are Solved: Reminder and Related Analysis . . . 24

5.3 How This Analysis Helps Explain The Minimum Entropy Production Principle 28

6 Analytical Techniques for Gauging Accuracy of Expert Knowledge: A Simple System-Based Explanation of the Dunning-Kruger Effect 31
6.1 Formulation of the Problem 32
6.2 Our Explanation 33
7 Analytical Techniques Help Enhance the Results of Data Mining: Case Study of Cow Insemination 38
7.1 Formulation of the Problem 39
7.2 Our Explanation 40
8 Analytical Techniques Help Enhance the Results of Data Mining: Why Pink Noise Is Best for Enhancing Sleep and Memory 43
8.1 Formulation of the Problem 43
8.2 Our Explanation 44
9 Analytical Techniques Help Enhance the Results of Data Mining: Why Filtering Out Higher Harmonics Makes It Easier to Carry a Tune 47
9.1 Formulation of the Problem 47
9.2 Our Explanation 48
10 Case When Analytical Techniques Invalidate the Conclusions of Data Mining: Reversed Flynn Effect of Decreasing IQ Test Scores 53
10.1 Formulation of the Problem 54
10.2 Systems-Based Analysis of the Problem and the Resulting Explanation of the Reversed Flynn Effect 55
10.3 Maybe the Same Idea Can Explain the Observed Increase in Universe's Expansion Rate: A Speculative Observation 57
11 Analytical Techniques in Hypothesis Testing: Why Area Under the Curve? . . . 59
11.1 Formulation of the Problem 59
11.2 Our Explanation 62
12 It Is Important to Revisit the Selection of the Best Model When New Data Appear: Why Confirmation Bias Is a Faulty Strategy 64

12.1 Formulation of the Problem 64
12.2 Analysis of the Problem 65
13 Need for a Careful Comparison between Hypotheses: Case Study of Epicycles . 68
13.1 Epicycles: Bad Science or Genius Idea 68
13.2 Analysis of the Problem and the Resulting Conclusions 70
14 Analytical Techniques Help in Emulating Biological Systems: An Explanation of Why High-Level Attention Constantly Oscillates 73
14.1 Formulation of the Problem 73
14.2 Analysis of the Problem 75
14.3 Oscillations Are Better: Proofs 77
15 Analytical Techniques for Taking into Account Several Aspects of a Designed Systems: Case Study of Computation-Communication Tradeoff 82
15.1 Formulation of the Problem 83
15.2 Description of a Model 84
15.3 Analysis of the Model: What Is the Optimal Tradeoff Between Computation and Communication 88
15.4 Formulation of the Problem 90
15.5 Description of a Model 91
15.6 Analysis of the Model: What Is the Optimal Tradeoff Between Computation and Communication 95
16 Users Do Not Always Follow Expert Recommendations: Analytical Technique Explains Empirical Data 98
16.1 Formulation of the Problem 98
16.2 Our Explanation 99
17 Analytical Techniques for Making Recommendations More Acceptable to Users: Status Quo Bias Actually Helps Decision Makers to Take Nonlinearity into Account 102
17.1 Formulation of the Problem 103
17.2 Analysis of the Problem and the Resulting Explanation 104

18 Analytical Techniques for Testing: Optimal Distribution of Testing Resources Between Different System Levels 107
18.1 Formulation of the Problem 107
18.2 Analysis of the Problem 109
18.3 Formulation of the Problem in Precise Terms 111
18.4 Solving the Resulting Optimization Problem 112
References 115
Curriculum Vitae 130

Chapter 1

Formulation of the Problem

Main objectives of systems engineering: a brief reminder. One of the main objectives of systems engineering is to design, maintain, and analyze systems that help the users. To design an appropriate system for an application domain, we need to know:

- what are the users' desires and preferences, so that we know in what direction we should aim to change this domain, and
- what is the current state and what is the dynamics of this application domain, and
- how to use all this information to select the best alternatives for the system design and maintenance.

Need for analytical techniques. Designing a system includes selecting numerical values for many of the parameters describing the corresponding system and its subsystems. At present, in many cases, this selection is made by consulting experts and/or by following semi-heuristic recommendations (recommendations based partly on the past experience of system design and monitoring). Experience shows that such heuristic imprecise recommendations often lead to less-than-perfect results.

It is therefore desirable to come up with analytical techniques for system design, techniques that would be based on valid numerical analysis and on the solution of the corresponding optimization problems.

What we do in this dissertation: general idea. System engineering is a very broad discipline, with many different application domains. Each domain has its own specifics and requires its own analysis and, probably, it own analytical techniques.

What we do in this dissertation is we formulate and analyze *general* problems corresponding to different stages of system design, implementation, testing, and monitoring, and show, on appropriate examples, how the corresponding analytical techniques can be applied to different application domains.

What we do in this dissertation: a detailed description. We start with analytical techniques for describing the users' preferences. In the ideal world, we should be able to ask each user's opinion about each of the alternatives, but for large systems, with many possible alternatives, this is not realistic. Therefore, we need to extrapolate the user's preferences based on partial information that we can elicit from the user. There are analytical techniques for such extrapolation – e.g., the widely used matrix factorization technique. However, this technique is purely empirical – and thus, not very reliable. In Chapter 2, we provide a theoretical explanation for this techniques – and the existence of such an explanation makes it more reliable.

In analyzing user preferences, we need to take into account that these preferences are usually not very detailed – and thus, because of their approximate nature, we should not waste time trying to fit them optimally. This approximate nature is usually captured by the empirical 7 plus minus 2 law, according to which, in the first approximation, instead of sorting all the alternatives, a user usually divides them into 7 plus minus 2 groups. This law is purely empirical – and thus, its use is not as reliable as we would like it to be. To make this law more reliable, in Chapter 3, we provide a partial theoretical explanation of this law.

In addition to knowing the user preferences, we also need to know what is the current state and what is the dynamics of this application domain. This information comes from two main sources: from measurements and from expert estimates.

In analyzing this information, it is important to take into account that many real-world processes are probabilistic. In many cases, the corresponding probability distributions are Gaussian (normal) – which makes perfect sense, since such processes are affected by many independent factors and it is known that in such cases, the distributions should be close to

normal. However, there are cases when the corresponding distribution is different – e.g., uniform. In Chapter 4, on the example of a practical cases study, we explain why such distributions appear.

In general, systems change with time, and the corresponding probability distributions change. There are some general rules about such changes, some of them well-explained, some more of an empirical nature. For example, it is a well-known (and reasonable well-explained) fact that the entropy of a closed system increases with time – this is known as the Second Law of Thermodynamics. Interestingly, there is another empirical observation – which is not as well justified – that while the entropy increases, its rate of increase is often the smallest possible. The corresponding minimum entropy production principle was first formulated and explained by a future Nobelist Ilya Prigogine. Since then, many possible explanations of this principle appeared, but all of them are very technical, based on complex analysis of differential equations describing the system's dynamics. Since this phenomenon is ubiquitous for many systems, it is desirable to look for a general system-based explanation, explanation that would not depend on the specific technical details. Such an explanation is presented in Chapter 5. Our explanation is related to the well-known fact that in general, it is important to keep as many solution options open as possible: in decision making, one of the main errors is to focus too quickly and to become blind to alternatives.

Dealing with expert estimates bring additional challenges. For example, while measurement results come with guaranteed bounds on the corresponding measurement inaccuracy, the only estimates of the inaccuracy of expert estimates come from the experts themselves. It turns out that experts often misjudge the inaccuracy of their estimates. This phenomenon is known as the Dunning-Kruger effect, after the two psychologists who discovered it. Which this phenomenon has been confirmed by many follow-up experiments, it remains largely unexplained. In Chapter 6, we present an analytical model that provides a simple system-based explanation for the Dunning-Kruger effect.

Once we have the information about the system, information coming from measurements

and from expert estimates, we use this information to come up with a model describing the system. The usual way to come up with such a model is to formulate several different hypotheses and to select the one that best fits the data. Techniques for formulating hypotheses based on the available information are known as *data mining techniques*. When the amount of data is not sufficient to make statistically justified conclusions, the dependencies produced by data mining techniques are often caused by accidental coincidences and do not reflect the actual behavior of the corresponding system. To separate such accidental coincidences from true dependencies, it is important to look for possible theoretical explanation for these empirical dependencies:

- if such an explanation is possible and natural, this means that this dependence is in line with our knowledge about the system and it is, thus, highly probable that this dependence is real;
- on the other hand, if no such natural explanation is possible, this means that this "dependence" is probably an accidental coincidence.

In this dissertation, we illustrate this general approach on four examples, all four biology-related. In the first three examples, we have found a natural explanation for the observed phenomenon which confirmed the conclusions of data mining:

- the first example is a surprising observation that was made from the analysis of records of cow insemination; this example is described in Chapter 7;
- the second example is an empirical fact that pink noise enhances sleep and memory in humans, see Chapter 8;
- the third example is that filtering out higher frequencies makes it easier for a human to carry a tune; see Chapter 9.

In the fourth example – related to an observed decline in IQ scores – vice versa, a natural explanation invalidates the conclusion of data mining; see Chapter 10.

Once we have come with several reasonable models, we need to select the one that best fits the data. There are many statistical techniques for selecting the model, most of them well-justified but some more heuristic – and thus, less reliable. One of such techniques is a widely used area-under-the-curve method. In Chapter 11, we use analytical techniques to provide a theoretical explanation for this method – and thus, make it more reliable. In Chapter 12, we use analytical techniques to explain why, upon getting new data, it is desirable to revisit the selection of the best model, and why a usual practice of sticking to the original model is faulty.

In Chapter 13, we illustrate the need for a careful comparison between different hypotheses on one of the most well-known historical examples – epicycles versus more modern techniques in celestial mechanics. Our conclusion is that, contrary to what one may read in modern astronomy and physics textbooks, epicycles were actually a very efficient tool, in some aspects foreseeing modern techniques such as Fourier series – while not exactly as efficient as Fourier series.

Once we have an adequate description of the users' preferences and of the corresponding application domain, we need to come up with a system design which the most appropriate for this setting. One of the ways to come up with a good design is to use the experience of successful similar systems – engineering and even biological. Examples of such systems are plentiful: many situations in engineering and in life require constant monitoring. At first glance, this would necessitate the need for the system to maintain the same alert level. However, interestingly, recent experiments have shown that in many situations like driving, the driver's attention level constantly oscillates. In Chapter 14, we show that such an oscillation is indeed helpful – and thus, it is necessary to emulate such an oscillation when designing automatic systems, e.g., for driving.

One of the challenges in searching for an optimal system design is that we need to take into account many different aspects of the resulting system. In many practical situations, for each aspect, we have well-defined optimal design strategies, but there is no analytical techniques for taking all the aspects into account. In Chapter 15, we show how several

aspects can be taken into account on the example of a tradeoff between computation and communication needs.

When making recommendations, we need to take into account that people do not necessarily follow the expert advice. In Chapter 16, we provide that an analytical model that explains the observed non-compliance, and in Chapter 17, we use analytical techniques to explain how to make recommendations more acceptable.

On all design stages, we need to test the designed system. This testing has to be done on all levels, from the original big-picture design draft to the level of final detailed implementation. On each level, there are numerous known techniques and methods for testing. The problem is that our resources are limited, so we need to optimally distribute these testing resources between different levels. In Chapter 18, we use analytical techniques to come up with recommendation on how to optimally distribute testing resources between different system levels.

Chapter 2

Analytical Techniques for Describing User Preferences: Justification for (and Extension of) the Matrix Factorization Technique

In the ideal world, we should be able to ask each user's opinion about each of the alternatives, but for large systems, with many possible alternatives, this is not realistic. Therefore, we need to extrapolate the user's preferences based on partial information that we can elicit from the user. There are analytical techniques for such extrapolation – e.g., the widely used matrix factorization technique. However, this technique is purely empirical – and thus, not very reliable. In this chapter, we provide a theoretical explanation for this techniques – and the existence of such an explanation makes it more reliable.

Comment. The results presented in this chapter first appeared in [2].

2.1 Formulation of the Problem

Recommender systems. Many computer-based services aim at making the customers happier. For example, platforms like amazon.com that help us buy things not only allow us to buy what we want, they also advise us about we may be interested in looking at.

The system comes up with this advice based on our previous pattern of purchases and on how satisfied we were with these purchases. For example, platforms like Netflix not only

allow you to watch movies, they also use our previous selections to help the customer by providing advice on what other movies this particular customer will want to see.

To make such recommendations, for each customer i, the system uses the ratings r_{ij} that different customers made for different objects j. Based on the available values r_{ij} corresponding to different customers and different objects, the system estimates the customer's future ratings of different possible objects – and, based on these ratings, recommends, to each customer i, the objects j for which the estimated ratings r_{ij} are the largest. Such systems that, based on our past selections and our previous rating, try to predict our future preferences are known as *recommender systems*.

Matrix factorization. One of the most successful techniques in designing recommender systems is *matrix factorization*; see, e.g., [107, 135] and references therein. This method is based on the assumption that we can find parameters $c_{i1}, \ldots, c_{in}$ characterizing the i-th customer and parameters $o_{j1}, \ldots, o_{jn}$ characterizing the j-th object so that the rating r_{ij} of i-th customer on the j-th object has the form

$$r_{ij} = \sum_{k=1}^{n} c_{ik} \cdot o_{jk}. \tag{2.1}$$

Challenge. While the matrix factorization methods works well, it is not clear why a person's recommendations can be described in this way.

What we do in this chapter. In this chapter, we provide a general systems-based explanation for the matrix factorization method.

To be more precise, our theoretical analysis leads to a somewhat more general techniques. We hope that the use of this more general techniques will lead to a better representation of user preferences.

2.2 Why Matrix Factorization: Our Explanation

Formulation of the problem in precise terms. Let $p_1, \ldots, p_n$ be parameters describing a customer, and let $q_1, \ldots, q_n$ be parameters describing the object. Based on the values p_i describing the customer and on the values $q_1, \ldots, q_n$ describing the object, we need to estimate the customer's rating of the object. Let us denote the algorithm providing such an estimation by

$$f(p_1, \ldots, p_n, q_1, \ldots, q_n).$$

We want to explain why the formula (2.1) is a good model for such a dependence.

Linearization. In general, in the first approximation, we can always expand each dependence in Taylor series and keep only linear terms in the corresponding expansion. This *linearization* procedure is a general systems idea widely (and successfully) used in physics, in engineering, and in many other applications; see, e.g., [35].

In line with this general idea, let us expand the function f in Taylor series in terms of the values $p_1, \ldots, p_n$ and keep only linear terms in this expansion. We perform this procedure for each possible combination of values $q_1, \ldots, q_n$. As a result, for each possible combination of values $q_1, \ldots, q_n$, we get an expression which is linear in p_i:

$$f(p_1, \ldots, p_n, q_1, \ldots, q_n) = a_0(q_1, \ldots, q_n) + \sum_{k=1}^{n} a_k(q_1, \ldots, q_n) \cdot p_k. \tag{2.2}$$

In this expression, in general, for different combinations of values q_k, the corresponding coefficients a_k are different.

We can then apply the same linearization procedure to each of the dependencies $a_k(q_1, \ldots, q_n)$:

$$a_k(q_1, \ldots, q_n) = a_{k0} + \sum_{\ell=1}^{n} a_{k\ell} \cdot q_\ell. \tag{2.3}$$

Substituting the expressions (2.3) into the formula (2.2), we conclude that

$$f(p_1, \ldots, p_n, q_1, \ldots, q_n) = a_{00} + \sum_{k=1}^{n} a_{k0} \cdot p_k + \sum_{\ell=1}^{n} a_{0\ell} \cdot q_\ell + \sum_{k=1}^{n} \sum_{\ell=1}^{n} a_{k\ell} \cdot p_k \cdot q_\ell. \tag{2.4}$$

Singular Value Decomposition. By selecting appropriate linear combinations of p_i and q_j, we can represent the matrix $a_{k\ell}$ in the diagonal form; this is known as the *Singular Value Decomposition* of the matrix $a_{k\ell}$. In other words, if:

- instead of the original variables $p_1, \ldots, p_n$, we use their appropriate linear combinations p'_k, and
- instead of the original variables $q_1, \ldots, q_n$, we use their appropriate linear combinations q'_ℓ,

then the expression $\sum\limits_{k=1}^{n}\sum\limits_{\ell=1}^{n} a_{k\ell}\cdot p_k \cdot q_\ell$ takes a diagonalized form $\sum\limits_{k=1}^{n} \lambda_k \cdot p'_k \cdot q'_k$ for some values λ_k.

This expression can be further simplified if instead of the variables p'_k, we use variables $p''_k \stackrel{\text{def}}{=} \lambda_k \cdot p'_k$. Then, the diagonalized form takes the following simpler form: $\sum\limits_{k=1}^{n} p''_k \cdot q'_k$.

Since the new variables p''_k are linear combinations of the original variables $p_1, \ldots, p_n$, vice versa, the original variables p_k are linear combinations of the new variables $p''_1, \ldots, p''_n$. If we substitute these linear combinations into the formula $\sum\limits_{k=1}^{n} a_{k0} \cdot p_k$, we get a linear combination of the new variables p''_k, i.e., an expression of the type $\sum\limits_{k=1}^{n} a''_{k0} \cdot p''_k$, for appropriate coefficients a''_{k0}.

Similarly, since the new variables q_ℓ are linear combinations of the original variables $q_1, \ldots, q_n$, vice versa, the original variables q_ℓ are linear combinations of the new variables $q'_1, \ldots, q'_n$. If we substitute these linear combinations into the formula $\sum\limits_{\ell=1}^{n} a_{0\ell} \cdot q_\ell$, we get a linear combination of the new variables q'_ℓ, i.e., an expression of the type $\sum\limits_{\ell=1}^{n} a'_{0\ell} \cdot q'_\ell$, for appropriate coefficients $a'_{0\ell}$.

Thus, in terms of the new variables p''_i and q'_j, the expression (2.4) takes the form

$$f(p''_1, \ldots, p''_n, q'_1, \ldots, q'_n) = a_{00} + \sum_{k=1}^{n} a''_{k0} \cdot p''_k + \sum_{\ell=1}^{n} a'_{0\ell} \cdot q'_\ell + \sum_{k=1}^{n} p''_k \cdot q'_k. \tag{2.5}$$

This expression can be further simplified. The above expression can be further simplified if we introduce new variables $p'''_k = p''_k + a'_{0k}$ and $q''_\ell = q'_\ell + a''_{\ell 0}$ for which $p''_k = p'''_k - a'_{0k}$ and $q'_\ell = q''_\ell - a''_{\ell 0}$. Substituting these expressions for p''_k and q'_ℓ into the formula (2.5), we get

$$f(p'''_1, \ldots, \ldots, p'''_n, q''_1, \ldots, q''_n) =$$

$$a_{00} + \sum_{k=1}^{n} a''_{k0} \cdot (p'''_k - a'_{0k}) + \sum_{\ell=1}^{n} a'_{0\ell} \cdot (q''_\ell - a''_{\ell 0}) + \sum_{k=1}^{n} (p'''_k - a'_{0k}) \cdot (q''_k - a''_{k0}) =$$

$$a_{00} + \sum_{k=1}^{n} a''_{k0} \cdot p'''_k - \sum_{k=1}^{n} a''_{k0} \cdot a'_{0k} + \sum_{k=1}^{n} a'_{0k} \cdot q'_k - \sum_{k=1}^{n} a'_{0k} \cdot a''_{k0} +$$

$$\sum_{k=1}^{n} p'''_k \cdot q''_k - \sum_{k=1}^{n} a'_{k0} \cdot p'''_k - \sum_{k=1}^{n} a'_{0k} \cdot q''_k + \sum_{k=1}^{n} a''_{k0} \cdot a'_{0k} =$$

$$a'_{00} + \sum_{k=1}^{n} p'''_k \cdot q''_k,$$

where

$$a'_{00} = a_{00} - \sum_{k=1}^{n} a''_{k0} \cdot a'_{0k}.$$

Conclusion. Thus, in the new variables p'''_k and q''_k, the function

$$f(p'''_1, \ldots, \ldots, p'''_n, q''_1, \ldots, q''_n)$$

that estimates the customer's ratings takes the form

$$f(p'''_1, \ldots, \ldots, p'''_n, q''_1, \ldots, q''_n) = a'_{00} + \sum_{k=1}^{n} p'''_k \cdot q''_k. \tag{2.6}$$

This is – almost – the original expression (2.1), with:

- the variables $p'''_1, \ldots, p'''_n$ describing the customer and
- the variables $q''_1, \ldots, q''_n$ describing different objects.

The only difference from the formula (2.1) is that in our general formula (2.6) we have an additional constant term a'_{00}.

This term can be deleted if we appropriately re-scale the ratings, i.e., consider the ratings $r'_{ij} = r_{ij} - a'_{00}$ instead of the original ratings r_{ij}. Indeed, if, e.g., ratings on a scale from 0 to 10 satisfy the formula (2.1), this formula will no longer be valid if we use a difference numerical scale for the same ratings – e.g., a scale from -5 to 5. In this sense, the formula (2.1) already pre-assumes that we are using an appropriate scale – and thus, our general formula (2.6) indeed provides an explanation for why the empirical formula (2.1) is so ubiquitous.

Chapter 3

Analytical Techniques for Describing User Preferences: 80/20 Rule Partially Explains 7 Plus Minus 2 Law: General System-Based Analysis

In analyzing user preferences, we need to take into account that these preferences are usually not very detailed – and thus, because of their approximate nature, we should not waste time trying to fit them optimally. This approximate nature is usually captured by the empirical 7 plus minus 2 law, according to which, in the first approximation, instead of sorting all the alternatives, a user usually divides them into 7 plus minus 2 groups. This law is purely empirical – and thus, its use is not as reliable as we would like it to be. To make this law more reliable, in this chapter, we provide a partial theoretical explanation of this law.

In this explanation, we use another difficult-to-explain empirical law: the 80/20 rule, according to which, in each activity, 20% of the people contribute to the 80% of the results. In this chapter, we show that, because of the 80/20 rule, the number of classes cannot be smaller than 5. Thus, the 80/20 rule explains the lower bound (5) on the 7 ± 2 law.

Comment. The results presented in this chapter first appeared in [4].

3.1 Formulation of the Problem

Difficult-to-explain empirical facts. There are several difficult-to-explain empirical facts.

- For example, there is a ubiquitous 80/20 rule, according to which, in each human activity, 80% of the results come from 20% of the participants. For example, 20% of the people own 80% of all the wealth, 20% of researchers publish 80% of all the papers, etc.; see, e.g., [45, 63] and references therein.
- There is a known phenomenon in psychology called a 7 ± 2 law (see, e.g., [83, 105]), according to which each person usually classifies everything into a certain number of classes C; depending on the person, this number ranges from $7 - 2 = 5$ to $7 + 2 = 9$ classes.

We cannot explain these facts, but we can at least find the relation between them. There have been many attempts to explain these two facts; see, e.g., [65, 127]. However, in general, we are still far from fully understanding them.

Meanwhile, maybe we can have at least some relation between the two facts: e.g., maybe we can show that one of them explains another one – at last partially. This is what we do in this chapter: we show that the 80/20 rule partially explains the 7±2 law.

3.2 Our Explanation

Consequences of division into C classes. If we, in the first approximation, divide everything into C classes, this means that any proportion which is smaller than $1/C$ will be, in this approximation, simply ignored. For example:

- If $C = 9$, this means that any proportion smaller than $1/9 \approx 11\%$ will be ignored.
- If $C = 5$, this means that any proportion smaller than $1/5 = 20\%$ will be safely ignored, etc.

What happens if $C < 5$. If $C < 5$, i.e., if $C \leq 4$, then any proportion smaller than $1/4 = 25\%$ will be, in the first approximation, ignored.

How this is related to the 80/20 rule: wealth example. Let us see how this is related to the 80/20 rule. As we have mentioned, in general, 20% of the people own 80% of all the property, so the property owned by the remaining 80% of the people amounts to 20% of the world's wealth.

When $C \geq 5$, we can still see that: in the division into at least 5 categories, at least one of the categories is the wealth owned by the majority of the people – exactly one category out of 5 if we have $C = 5$, but still at least one such category.

If $C \leq 4$, this means that this proportion will be ignored and people will get an impression that they own nothing – that everything is owned by a few rich folks. This impression is not a recipe for social stability – it is a recipe for a violent revolution.

How this is related to the 80/20 rule: case of research productivity. In a less violent consequence, 20% of researchers publish 80% of the papers. Thus, the remaining 80% of researchers publish the remaining 20% of the papers.

When $C \geq 5$, we can still see this proportion and thus, conclude that even the least productive scientists have a chance to contribute to the world's body of knowledge.

However, if we had $C \leq 4$, then, in the first approximation, we would simply not see any possibility for anyone who is not a top researcher to publish – and this would clearly very much discourage the scientists' activity.

Another example: 20% of the letters from a text carry practically all information. An even more extreme example come from Claude Shannon's estimate that the redundancy rate of the English text is about 80%: crudely speaking, only one in five letters carries any information; see, e.g., [115], p. 152.

With $C \geq 5$, we can still notice this informative part.

However, if we had $C \leq 4$, then, in the first approximation, we would not notice any meaningful information at all – and we would thus be able to erroneously conclude that all

communications are non-informative.

Conclusion. Based on these examples, we can make the following general conclusion:

Due to the 80/20 rule, the number C of clusters on which we divide objects must be at least 5.

This explains the lower bound for the seven plus minus two law.

Chapter 4

Analytical Techniques for Analyzing Probability Distributions: How to Explain That Changes in Elderlies Depression Level Are Uniformly Distributed

In analyzing information about the application domain, it is important to take into account that many real-world processes are probabilistic. In many cases, the corresponding probability distributions are Gaussian (normal) – which makes perfect sense, since such processes are affected by many independent factors and it is known that in such cases, the distributions should be close to normal. However, there are cases when the corresponding distribution is different – e.g., uniform. This happens, e.g., when we analyze the changes in the elderlies depression level. In this chapter, on the example of this practical cases study, we explain why such distributions appear.

Comment. The results presented in this chapter first appeared in [7].

4.1 Formulation of the Problem

Elderly depression is a serious problem. Many elderly people suffer from loneliness. In general, loneliness increases the chances of depression, and depression negatively affects the person's health. As a result, many elderly people suffer from depression; it affects every seventh elderly person – a much larger proportion than in the population in general; see, e.g., [20, 131].

It is therefore desirable to monitor the changes in depression level of elderly people, especially elderly people in an at-risk category. Such monitoring has indeed been undertaken; see, e.g., [93].

Changes in depression level are uniformly distributed. Depression level is usually gauged by a number on the Geriatric Depression Scale (GDS); see, e.g., [93, 131]. The changes usually range from -8 to $+8$ units.

Interestingly, for any two moments of time, the corresponding changes are, in effect, uniformly distributed on the interval $[-d_0, d_0]$, where $d_0 \approx 8$.

Why this is interesting. The uniform distribution rarely occurs in nature. The most typical probability distribution is a normal one. Its ubiquity comes from the fact that most real-life phenomena result from the joint effect of many independent small factors. In probability theory, it is known that, under reasonable conditions, the probability distribution of the sum of a large number of small independent random variables is close to Gaussian – the corresponding result is known as the Central Limit Theorem; see, e.g., [111]. Normal distribution is what we almost always encounter – and so the emergence of the uniform distribution is highly unusual.

The fact that we have the same uniform distribution for shorter- and longer-time periods is also unusual. Indeed, as mentioned in [93], it is difficult to predict the change in one short-term period based on the observed change in another such period. With this in mind, it seems reasonable to conclude that the changes corresponding to different short-term periods are independent. The long-term difference can be represented as the sum of several such

short-term differences. If all these short-term differences are uniformly distributed, the distribution of long-term differences should correspond to the sum of several independent uniform distributions. However, it is well known that the distribution of such a sum is *not* uniform: e.g., the distribution of the sum of two identical independent uniform distribution has a triangular probability density function. In contrast, the probability distribution of long-term differences *is* uniform.

How can we explain all this?

Why? In this chapter, we provide a possible explanation for the emergence of this unexpected uniform distribution. This explanation is of very general nature, so it can be applied to other situations as well.

4.2 Our Explanation

Main idea. The fact that we cannot predict the change in one period based on the change in another period implies that the changes corresponding to adjacent time periods are kind of independent. On the other hand, they cannot be fully independent: if they were, then there would be a possibility that by combining almost d_0 unit changes in both periods, we would get a $2d_0$ units change in the long-term period – and in the long-term period, we only observe changes from $-d_0$ to d_0 units.

To describe the corresponding probability distributions, we therefore need to take into account two facts:

- that these distributions are almost independent, but
- that the distribution of the sum of these two random variables is bounded by the same interval $[-d_0, d_0]$ as each of the short-term changes.

Resulting formalization. The only limitation to independence is the observed range of the values of the sum of the two random variables. It is therefore reasonable to describe

the situation as maximally independent – with the restriction on the sum as the only available restriction.

In precise terms, we conclude that the probability distributions corresponding to two adjacent time intervals are not fully independent – the distribution of their sum corresponds to the distribution of the sum of two independent random variables, but limited to the range $[-d_0, d_0]$. In other words, the distribution for the sum can be obtained if we:

- first consider the distribution of the sum of two independent random variables each of which is distributed on the interval $[-d_0, d_0]$ and
- then we restrict this distribution to the interval $[-d_0, d_0]$, i.e., consider the conditional distribution – under the condition that the sum is located in the interval $[-d_0, d_0]$.

This formalization explains the seeming contradiction. The above formalization explains the above-mentioned seeming contradiction between:

- the fact that the long-term difference is the sum of practically independent uniformly distributed short-term differences, and
- the fact that the observed probability distribution of the long-term differences is very different from the distribution of the sum of several independent uniformly distributed random variables.

This formalization also explains the emergence of the uniform distribution. It turns out that the same formalization can explain why the distributions are uniform in the first place.

Indeed, each difference $d(t+T)-d(t)$ between depression levels at moments t and $t+T$ is the sum of the large number of very-short-term differences:

$$d(t+T)-d(t) = (d(t+\Delta t)-d(t)) + (d(t+2\Delta t)-d(t+\Delta t)) + \ldots +$$

$$d(t+(k+1)\cdot\Delta t)-d(t+k\cdot\Delta t)) + \ldots + (d(t+T)-d(t+T-\Delta t)).$$

According to our formalization, the distribution of this sum should be obtained by first taking the distribution of the sum of several independent random variables

$$d(t + (k+1) \cdot \Delta t) - d(t + k \cdot \Delta t)$$

and then limiting this distribution to the original interval $[-d_0, d_0]$.

According to the Central Limit Theorem, the distribution of the sum is close to Gaussian – and since the variance of the sum of several independent random variables is equal to the sum of the variances, this variance σ^2 grows with the number of variables in the sum. Thus, in our case, we select restrict to the interval $[-d_0, d_0]$ a Gaussian distribution corresponding to a very large value of σ^2 – and, thus, to the very large value of the standard deviation σ. On this interval, the probability density function changes from $\dfrac{1}{\sqrt{2\pi} \cdot \sigma}$ to

$$\frac{1}{\sqrt{2\pi} \cdot \sigma} \cdot \exp\left(-\frac{d_0^2}{2\sigma^2}\right).$$

The ratio of these two values – which is preserved when we consider conditional distributions – is thus equal to $\exp\left(-\dfrac{d_0^2}{2\sigma^2}\right)$ which is, for large σ, very close to 1.

Thus, under our assumption of "almost independence", the resulting probability distribution is very close to the uniform one – which is exactly what we observe.

Chapter 5

Analytical Techniques for Analyzing How Systems Change with Time: A Natural Explanation for the Minimum Entropy Production Principles

In the previous chapter, we analyzed possible probability distributions describing real-life phenomena. In general, systems change with time, and the corresponding probability distributions change. There are some general rules about such changes, some of them well-explained, some more of an empirical nature. For example, it is a well-known (and reasonable well-explained) fact that the entropy of a closed system increases with time – this is known as the Second Law of Thermodynamics. Interestingly, there is another empirical observation – which is not as well justified – that while the entropy increases, its rate of increase is often the smallest possible. The corresponding minimum entropy production principle was first formulated and explained by a future Nobelist Ilya Prigogine. Since then, many possible explanations of this principle appeared, but all of them are very technical, based on complex analysis of differential equations describing the system's dynamics. Since this phenomenon is ubiquitous for many systems, it is desirable to look for a general system-based explanation, explanation that would not depend on the specific technical details. Such an explanation is presented in this chapter. Our explanation is

related to the well-known fact that in general, it is important to keep as many solution options open as possible: in decision making, one of the main errors is to focus too quickly and to become blind to alternatives.

Comment. The results presented in this chapter first appeared in [9].

5.1 Formulation of the Problem

Minimum entropy production principle. It is well known that, according to the second law of thermodynamics, the entropy of any closed system – including the Universe as a whole – cannot decrease, it can only either increase or stay the same; see, e.g., [35, 124].

It is somewhat less well known that in many situation, this entropy increase is the smallest possible; this fact is known as the *minimum entropy production principle*. This principle was first formulated in 1945 by a future Nobelist Ilya Prigogine [101]; see also [44, 55, 62, 76, 78, 81, 102].

In contrast to the second law of thermodynamics – which is always true – the minimum entropy production principle is not always valid (see, e.g., [48]), but it is still valid in many practical situations. In particular, it explains why usually, a symmetric state, when perturbed, does not immediately turn into a state with no symmetries at all; usually, some symmetries are preserved – and the more symmetries are preserved, the more frequent are such transitions. For example, when heated, a highly symmetric solid-body state usually does not immediately turn into a completely symmetry-less gas state, it first transitions into a liquid state in which some symmetries are preserved. Sometimes, a solid state does turn directly into gas: e.g., dry ice used to keep ice cream cold goes directly into a gas state without becoming a liquid. However, usually, symmetries are broken sequentially, not all at once. This seemingly simple idea explains many physical phenomena: e.g., it explains the observable shapes of celestial bodies, relative frequency of different shapes, and how shapes change with time; see, e.g., [37, 38, 75].

Challenge: to provide a simple explanation for the minimum entropy produc-

tion principle. While the principle itself sounds reasonable, all its available derivations are very technical and non-intuitive. Usually, in physics, no matter how complex the corresponding equations, there is a reasonably simple explanation – at least a qualitative one – of the observed phenomena [35, 124]. However, for the minimum entropy production principle, such an explanation has been lacking.

What we do in this chapter. In this chapter, we provide a general system-based explanation for the ubiquity of the minimum entropy production principle, an explanation which – unlike the existing ones – uses only simple easy-to-understand math.

In this explanation, we will first start with the analysis how complex problems are solved, and then we will explain how this analysis helps explain the minimum entropy production principle.

5.2 How Complex Problems Are Solved: Reminder and Related Analysis

NP-complete problems: a brief reminder. As we have mentioned, our explanation for the minimum entropy production principle starts not with physics, but with the known fact that in real life, we need to solve complex problems:

- we may need to find a path that leads from point A to point B,
- a mechanic needs to find a way to repair a broken car,
- a medical doctor needs to cure the patients.

In most such problems, it may be difficult to come up with a solution, but once we have a candidate for a solution, we can relatively easily check whether this is indeed a solution. For example, if may be difficult to find a way to repair a car, but if we follow some sequence of actions and the car starts running, we clearly have a solution – otherwise, if the car does not start running, the sequence is not a solution.

The class of all such problems, i.e., problems in which we can, in reasonable ("feasible") time check whether a given candidate for a solution is indeed a solution, is known as the class NP. Within this class, there is a subclass of all the problems that can be *solved* in reasonable time. This subclass is usually denoted by P; see, e.g., [68, 94] for details.

Most computer scientists believe that there are problems that cannot be solved in reasonable time, i.e., that P is different from NP; however, this has never been proven, it is still an open problem. What is known is that in the class NP, there are problems which are as hard as possible – in the sense that all other problems can be reduced to this one. Such problems are known as *NP-complete*.

Historically the first NP-complete problem was the following propositional satisfiability problem for 3-SAT formulas.

- We start with *Boolean* (*propositional*) variables $x_1, \ldots, x_n$, i.e., variables that can take only two values: true (1) and false (0).
- A *literal* is either a variable x_i, or its negation $\neg x_i$.
- A *clause* (*disjunction*) is an expression of the type $a \vee b$ or $a \vee b \vee c$, where a, b, and c are literals.
- Finally, a *3-SAT formula* is an expression of the type $C_1 \,\&\, C_2 \,\&\, \ldots, \&\, C_m$, where C_j are clauses.

An example is a 3-clauses formula

$$(x_1 \vee x_2) \,\&\, (\neg x_1 \vee x_2 \vee x_3) \,\&\, (x_1 \vee \neg x_2 \vee \neg x_3).$$

The general problem is:

- given a 3-SAT formula,
- check whether this formula is *satisfiable*, i.e., whether there exist values of the variables that make it true.

How NP-complete problems are solved now. If P≠NP, this means, in particular, that no feasible algorithm is possible that would solve all the instance of the general 3-SAT problem. So, in practice, when only feasible algorithms are possible, we have to use *heuristic* algorithms, i.e., algorithms which do not always lead to a solution.

Many such algorithms start by selecting a literal – i.e., equivalently, by selecting one of the Boolean variables x_i and selecting its truth value. Then, when we substitute this value into the original formula, we get a new propositional formula with one fewer variable. If the original formula was satisfiable and we selected the literal correctly, then the new formula is also satisfiable – and so, by repeating this procedure again and again, we will confirm that the formula is satisfiable (and also find the values of the variables x_i that make the formula true).

Which literal should we select? In general, a satisfying 3-SAT formula has several satisfying vectors. For example, by trying all 8 possible combinations of truth values, we can check that the above sample 3-SAT formula has four different solutions: (101), (110), (111), and (010).

By selecting a literal, we restrict the number of solutions, from the original number N to a new – usually smaller – number $N' \le N$. A priori we do not know which vector of Boolean values are solutions, all 2^n such vectors are equally probable to be a solution. Thus, the more vectors remain, the higher the probability that by this restriction we do not miss a solution. It is therefore reasonable to select a literal for which the estimated number of satisfying vectors is the largest possible; see, e.g., [30, 31, 66, 67] and references therein.

For a general 3-SAT formula, the expected number of solutions can be estimated, e.g., as follows:

- a formula $a \vee b$ is satisfied by 3 out of 4 combinations of the values (a, b) (the only combination which does not make this formula true is $a = b =$ false); thus, the probability that this clause will be satisfied by a random Boolean vector is 3/4;

- a formula $a \vee b \vee c$ is satisfied by 7 out of 8 combinations of the values (a, b, c) (the only combination which does not make this formula true is $a = b = c =$ false); ; thus, the probability that this clause will be satisfied by a random Boolean vector is 7/8.

It is difficult to take into account correlation between the clauses, so, in the first approximation, we can simply assume that the clauses are independent, and thus, the probability that a random vector satisfies the formula is equal to the product of the corresponding probabilities – and the number of satisfying vectors can be estimated if we multiply the overall number 2^n of Boolean vectors of length n by this probability.

For example, for the above 3-SAT formula, the corresponding probability is $(3/4)\cdot(7/8)\cdot(7/8)$, and the estimates number of satisfying Boolean vectors is $(3/4)\cdot(7/8)\cdot(7/8)\cdot 2^3 \approx 4.6$. In this formula, we have three variables, so we have six possible literals. Which one should we select?

- if we select x_1 to be true, then the first and the third clauses are always satisfied, and the formula becomes $\neg x_2 \vee \neg x_3$; here, the estimated number of solutions is $(3/4) \cdot 2^2 = 3$;

- if we select a literal $\neg x_1$, i.e., we select x_1 to be false, then the second clause is satisfied, and the formula becomes $x_2 \,\&\, (\neg x_2 \vee \neg x_3)$; here, the estimated number of solutions is $(1/2) \cdot (3/4) \cdot 2^2 = 1.5$;

- if we select a literal x_2, then the formula becomes $x_1 \vee \neg x_3$; here, the estimated number of solutions is $(3/4) \cdot 2^2 = 3$;

- if we select a literal $\neg x_2$, then the formula becomes $x_1 \,\&\, (\neg x_1 \vee x_3)$; here, the estimated number of solutions is $(1/2) \cdot (3/4) \cdot 2^2 = 1.5$;

- if we select a literal x_3, then the formula becomes $(x_1 \vee x_2) \,\&\, (x_1 \vee \neg x_2)$; here, the estimated number of solutions is $(3/4) \cdot (3/4) \cdot 2^2 = 2.25$;

- finally, if we select a literal $\neg x_3$, then the formula becomes

$$(x_1 \vee x_2) \,\&\, (\neg x_1 \vee x_2);$$

  here, the estimated number of solutions is $(3/4) \cdot (3/4) \cdot 2^2 = 2.25$.

The largest estimate of remaining Boolean vectors is when we select x_1 or x_2. So, on the first step, we should select either the literal x_1 or the literal x_2. One can check that in both cases, we do not miss a solution (and in each of these cases, we actually get 3 solutions, exactly the number that we estimated).

General case. The same idea is known to be efficient for many other complex problems; see, e.g., [66]. For example, a similar algorithm has been successfully used to solve another NP-complete problem: a discrete optimization *knapsack problem*, where:

- given the resources $r_1, \ldots, r_n$ needed for each of n projects, the overall amount r of available resources, and the expected gain $g_1, \ldots, g_n$ from each of the projects,
- we need to select a set of projects $S \subseteq \{1, \ldots, n\}$ which has the largest expected gain $\sum\limits_{i \in S} g_i$ among all the sets that we can afford, i.e., among all the sets S for which $\sum\limits_{i \in S} r_i \le r$.

The corresponding algorithms are described, e.g., in [70, 112].

In general, it is important to keep as many solution options open as possible. In decision making, one of the main errors is to focus too quickly and to become blind to alternatives. This is a general problem-solving principle which the above SAT example illustrates very well.

5.3 How This Analysis Helps Explain The Minimum Entropy Production Principle

How is all this related to entropy. From the physical viewpoint, entropy is proportional to the logarithm of the number of micro-states forming a given macro-state; see, e.g.,

[35, 124]. In the case of the SAT problems, micro-states are satisfying vectors, so the number of micro-states is the number of such vectors. Similarly, in other complex problems, solution options are micro-states, and the number of micro-states is the number of such options.

As we solve each problem, the number of states decreases – but decreases as slowly as possible. Thus, the entropy – which is the logarithm of the number of states – also decreases, but decreases as slowly as possible, at the minimal possible rate.

So, if we consider the dependence of entropy on time, then, in the backward-time direction (i.e., in the direction in which entropy *in*creases), this increase is the smallest possible.

How is all this related to physics. At first glance, the above text may be more relevant for human and computer problem solving than for physics, since at first glance, nature does not solve problems.

However, in some reasonable sense it does; let us explain this. Traditionally, physical theories – starting from Newton's mechanics – have been formulated in terms of differential equations. In this formulation, there is no problem to solve: once we know the state at a given moment of time, we can compute the rate at which each variable describing the state changes with time. This computation may be tedious, may require a lot of computation time on a high-performance computer, but it does not constitute a challenging NP-complete problem.

At present, however, the most typical way to describe a physical theory is in the form of a variational principle, i.e., in the form of an objective function whose optimization corresponds to the actual behavior of the physical systems; see, e.g., [35, 69, 124]. This formulation is especially important if we take quantum effects into account:

- while in non-quantum physics, optimization is exact and is just another equivalent form of describing the corresponding differential equations,
- in quantum physics, optimization is approximate: a quantum system tries to optimize,

but its result is close to (but not exactly equal to) the corresponding optimum.

In this formulation, what nature does *is* solving the complex optimization problem: namely, trying to optimize the value of the corresponding functional.

We therefore expect to see the same pattern of entropy changes as in general problem solving: in the direction in which entropy is increasing, this increase is the smallest possible.

Increasing entropy is exactly how we determine the direction of physical time. For example:

- if we see a movie in which a cup falls down and break, we understand that this is exactly the time direction, while
- if we see the same movie played backward, when the pieces of a broken cup mysteriously come together to form a whole cup, we realize that we saw this movie in reverse.

From this viewpoint, the above statement means that in the forward-time direction – i.e., in the direction in which entropy increases – the rate of the entropy increase is the smallest possible.

We thus have a natural systems-based explanation for the minimum entropy production principle.

Chapter 6

Analytical Techniques for Gauging Accuracy of Expert Knowledge: A Simple System-Based Explanation of the Dunning-Kruger Effect

To properly design a system, we need to know what is the current state and what is the dynamics of this application domain. Often, an important part of this information comes from expert estimates. Dealing with expert estimates is challenging: while measurement results come with guaranteed bounds on the corresponding measurement inaccuracy, the only estimates of the inaccuracy of expert estimates come from the experts themselves. It turns out that experts often misjudge the inaccuracy of their estimates. Specifically, experienced experts not only provide better estimates of different situations than novice experts, but they also provide a better estimates of the accuracy of their estimates. This phenomenon is known as the Dunning-Kruger effect, after the two psychologists who discovered it. Which this phenomenon has been confirmed by many follow-up experiments, it remains largely unexplained. In this chapter, we present an analytical model that provides a simple system-based explanation for the Dunning-Kruger effect.

Comment. The results presented in this chapter first appeared in [6].

6.1 Formulation of the Problem

Dunning-Kruger effect: a brief reminder. In their 1999 paper [72], Justin Kruger and David Dunning from Cornell University showed that:

- not only experts have a better knowledge and produce more accurate estimates than novices,
- experts also assess their own accuracy with better accuracy than novices;

see also [32, 110]. This phenomenon became known as the *Dunning-Kruger effect*.

Comment. It should be mentioned that in their original paper [72], due to faulty statistical analysis, the above claim was more specific:

- that experts usually underestimate their abilities (in particular, overestimate the inaccuracy of their estimates),
- while novices usually overestimate their ability (in particular, underestimate the inaccuracy of their estimates).

A more accurate statistical analysis has shown that this specific claim is *not* supported by the evidence. However, the above claim – that experts estimate the accuracy of their own estimates better than novices – is definitely statistically valid; see, e.g., [89, 90].

Challenge. The effect is there, but how can we explain it?

What we do in this chapter. In this chapter, we provide a simple system-based qualitative explanation for the Dunning-Kruger effect.

Remaining open question and future work. We hope that future research will help transform our qualitative explanation into a more quantitative one.

6.2 Our Explanation

Towards formulating the problem in precise terms. What we want from experts is an estimate of the state of the corresponding system:

- this system is a patient if this expert is a medical doctor,
- this system is a complex machine if the expert is an engineer,
- this system is a mineral deposit if the expert is a geophysicist, etc.

Most information is usually described in terms of real numbers. In this sense, what we want from an expert is to provide some estimates of the numbers describing the corresponding system.

For example, based on his/her expertise (and on other available information), a geophysicist can provide us with:

- an estimate of the amount of oil in a given oil field,
- an estimate of the depth at which this oil will most probably be found, etc.

In all these cases, we want to predict the value of the desired quantity y based on the available values $x_1, \ldots, x_n$ of related quantities. In other words, we need to have some prediction algorithm

$$y = f(x_1, \ldots, x_n).$$

Depending on the previously available data, we can have different prediction algorithms. For example, if we want to predict tomorrow's weather, then:

- in some areas (e.g., areas shielded by mountain ranges), it is sufficient to take into account today's weather patterns only in the nearby areas, while
- in other places (e.g., on the plains where cyclones and other weather phenomena can travel large distances fast), we need to take into account today's weather in a much wider area.

In general, instead of a *single* algorithm

$$y = f(x_1, \ldots, x_n),$$

we have a *family* of algorithms

$$y = f(x_1, \ldots, x_n, c_1, \ldots, c_m)$$

depending on some parameters $c_1, \ldots, c_m$. These parameters $c_1, \ldots, c_m$ need to be determined based on the available data.

- In statistics [111] and in machine learning (see, e.g., [23, 47]) the corresponding family of algorithms is usually given explicitly.
- In expert decision making, a lot of the reasoning is happening at the subconscious level. So we do not have explicit expressions for the corresponding prediction algorithm – but, in effect, experts make predictions – and thus, subconsciously use some algorithms to make these predictions.

Resulting formulation of the problem. We have a model

$$y = f(x_1, \ldots, x_n, c_1, \ldots, c_m)$$

that describes the dependence of the desired quantity y on observable quantities $x_1, \ldots, x_n$. To make a prediction, we need to estimate the values of the parameters $c_1, \ldots, c_m$ that best describe the current situation.

In reality, measurements are approximate, and models are approximate. So, in the actual dependence of y on $x_1, \ldots, x_n$, we do not have the exact equality, there is also some noise N:

$$y = f(x_1, \ldots, x_n, c_1, \ldots, c_m) + N.$$

To estimate the values of these parameters, we can use the results of previous observations. Let K denote the number of such observations. For each observation $k = 1, \ldots, K$,

we know the values $x_1^{(k)}, \ldots, x_n^{(k)}$ and $y^{(k)}$ of the corresponding quantities. We want to find the values of the parameters $c_1, \ldots, c_m$ for which

$$y^{(k)} \approx f\left(x_1^{(k)}, \ldots, x_n^{(k)}, c_1, \ldots, c_m\right)$$

for all k.

Comment. This problem is ubiquitous in applications of statistics, it is known as *regression*.

- The most widely used and the most well-known is the *linear* regression, when the dependence $f(x_1, \ldots)$ on $x_1, \ldots, x_n$ is linear.
- However, nonlinear regression is also actively used; see, e.g., [111].

How the estimate's accuracy depends on the number of observations. The main difference between an expert and a novice is that an expert is aware of a much larger number of previous observations: the value K corresponding to the expert is much larger than the number of observations corresponding to the novice. Thus, to understand the difference between estimates by experts and estimates by novices, we need to analyze how the accuracy of an estimation depends on the available number of observations K.

According to statistics, if we have K observations with standard deviation σ, then, in general, the accuracy with which we can estimate the values of the corresponding parameters is proportional to $\frac{\sigma}{\sqrt{K}}$; see, e.g., [111].

This formula is easy to derive in the simplest situation, when we simply observe the desired quantity y several times, i.e., if we have K results $y^{(1)}, \ldots, y^{(K)}$ of measuring this same quantity. In this case, a reasonable way to combine these results into a single more accurate estimate is to take the average

$$\overline{y} = \frac{y^{(1)} + \ldots + y^{(K)}}{K}. \tag{6.1}$$

The inaccuracy of each measurement is described by the difference $\Delta y^{(k)} \stackrel{\text{def}}{=} y^{(k)} - y$ between the measurement result $y^{(k)}$ and the actual (unknown) value y of the corresponding quantity.

From the formula (6.1), we conclude that the inaccuracy $\Delta y \stackrel{\text{def}}{=} \overline{y} - y$ of the arithmetic average is equal to

$$\Delta y = \frac{\Delta y^{(1)} + \ldots + \Delta y^{(K)}}{K}.$$

Inaccuracies $\Delta y^{(k)}$ corresponding to different measurements are usually independent. For independent random variables, the variance of the sum is equal to the sum of the variances. Since each of the inaccuracies $\Delta y^{(k)}$ has the variance σ^2, the variance of the sum

$$\Delta y^{(1)} + \ldots + \Delta y^{(K)}$$

is thus equal to $K \cdot \sigma^2$. Hence, the standard deviation of the sum is equal to the square root of this variance, i.e., to $\sqrt{K} \cdot \sigma$.

When we divide a random variable by a positive constant, its standard deviation divides by the same constant. Thus, the standard deviation of the inaccuracy Δy – corresponding to the use of K observations – is equal to

$$\frac{\sqrt{K} \cdot \sigma}{K} = \frac{\sigma}{\sqrt{K}}.$$

This derivation is only valid for the simplest case, but a similar asymptotics $\frac{\sigma}{\sqrt{K}}$ holds in the general situation as well [111]. This confirms the intuitive idea that the more experience the expert, the more accurate are this expert's estimates.

Towards understanding the Dunning-Kruger effect: how accurately can we estimate the accuracy of our own estimates? Accuracy estimates presented in the previous subsection depend on the standard deviation σ of the measurement error – i.e., on the standard deviation of the differences $\Delta y^{(k)}$. Usually, we do not know this value exactly, but we can estimate it based on the results of previous observations.

For example, if we simply observe the desired quantity y several times, i.e., if we have K results $y^{(1)}, \ldots, y^{(K)}$ of measuring this same quantity, then we can estimate this standard deviation σ as

$$\sigma \approx \overline{\sigma} = \sqrt{\frac{1}{K-1} \cdot \sum_{k=1}^{K} \left(y^{(k)} - \overline{y}\right)^2}.$$

Similar formulas can be used in the general case: once we have estimates $\overline{c}_1, \ldots, \overline{c}_m$ for the parameters $c_1, \ldots, c_m$, we can then estimate σ as

$$\sigma \approx \overline{\sigma} \sqrt{\frac{1}{K-1} \cdot \sum_{k=1}^{K} \left(y^{(k)} - f\left(x_1^{(k)}, \ldots, x_n^{(k)}, \overline{c}_1, \ldots, \overline{c}_m \right) \right)^2}.$$

It is known that the relative accuracy $\dfrac{\sigma - \overline{\sigma}}{\sigma}$ of this estimate decreases with K as $\dfrac{1}{\sqrt{K}}$ – at the same rate as the accuracy itself. Thus, the more experienced the expert – i.e., the larger the corresponding K – the more accurately this expert can estimate the accuracy of his/her estimates.

This is exactly what the Dunning-Kruger effect is about. Thus, we have indeed found a simple explanation for this effect.

Chapter 7

Analytical Techniques Help Enhance the Results of Data Mining: Case Study of Cow Insemination

Once we have the information about the system, information coming from measurements and from expert estimates, we use this information to come up with a model describing the system. The usual way to come up with such a model is to formulate several different hypotheses and to select the one that best fits the data. Techniques for formulating hypotheses based on the available information are known as *data mining techniques*. When the amount of data is not sufficient to make statistically justified conclusions, the dependencies produced by data mining techniques are often caused by accidental coincidences and do not reflect the actual behavior of the corresponding system. To separate such accidental coincidences from true dependencies, it is important to look for possible theoretical explanation for these empirical dependencies:

- if such an explanation is possible and natural, this means that this dependence is in line with our knowledge about the system and it is, thus, highly probable that this dependence is real;
- on the other hand, if no such natural explanation is possible, this means that this "dependence" is probably an accidental coincidence.

In this chapter and in the three following chapters, we illustrate this general approach on four examples, all four biology-related. In the current chapter and in the two following

chapters, we consider the cases for which we have found a natural explanation for the observed phenomenon which confirmed the conclusions of data mining. The fourth chapter provides an opposite example: when a natural explanation invalidates the conclusion of data mining.

In this particular chapter, we consider the case study of cow insemination. To improve the efficiency of artificial insemination, farmers equip cows with sensors, based on which a computer system determines the cow's insemination window. Analysis of the resulting calves showed an unexpected dependence of the calf's gender on the insemination time: cows inseminated earlier in their window mostly gave birth to female calves, while cows inseminated later in their window mostly gave birth to males. In this chapter, we provide a general system-based explanation for this phenomenon.

Comment. The results presented in this chapter first appeared in [12].

7.1 Formulation of the Problem

Unexpected empirical fact. Computer-based systems are ubiquitous. For example, farmers use sensors to identify cows in heat – they then apply artificial insemination to these particular cows, thus guaranteeing that all the inseminated cows will become pregnant.

When analyzing the results of applying this technique, researchers found out an unexpected phenomenon – that the gender of the resulting calves depends on the insemination time:

- for cows inseminated at the very beginning of their sixteen-hour insemination window, most resulting calves are female (to be more precise, about 70%); while
- for cows inseminated during the later part of their sixteen-hour insemination window, most resulting calves are male;

see, e.g., [43] and references therein.

What we do in this chapter. In this chapter, we provide a natural system-based explanation for this newly discovered phenomenon.

7.2 Our Explanation

What would be a perfect mix of cow genders. Artificial insemination is a reasonably recent process. Until then, cows were always inseminated by the bulls. From this viewpoint, to understand why cows sometimes give birth to male calves and sometimes to female calves, we need to understand what proportion of cows and bulls would be perfect for such a mixed cows-and-bulls herd.

From the biological viewpoint, each species aims to reproduce as much as possible – as much as the food and other resources allow.

- From this viewpoint, if there are too few bulls in the herd, many cows will not be inseminated and thus, the herd will not achieve its reproductive potential.
- On the other hand, if there are too many bulls in the herd, much more than needed to inseminate all the cows, the herd reproductive potential will also be wasted – the same herd would reproduce more if instead of the extra useless bulls, we would have cows.

Thus, the ideal cow-herd situation is when there are exactly as many bulls as needed to inseminate all the cows – not less and not more.

What if the mix is imperfect: how to balance the situation. In real life, the cow-bull proportion may be not ideal. In this case, from the biological viewpoint, a reasonable idea is to produce calves of different genders so as to bring the cow-bull proportion closer to the ideal one:

- if there are at present too many bulls, it is desirable to balance the situation, by making sure that the majority of newborn calves are female;

- vice versa, if there are at present too few bulls, it is desirable to balance the situation, by making sure that the majority of newborn calves are male.

How can an individual cow know that the balance is imperfect. The gender of a calf is determined by the biological processes in the cow's body. How does the cow's body know when there are too many bulls or too few bulls?

At first glance, it may seem that the cow does not have this information. However, a detailed analysis of the situation shows that a cow *can* get this information.

Indeed, as we have mentioned, cows have a sixteen-hour period during which they can be inseminated. If there are sufficiently many bulls to inseminate all the cows in heat in a shorter period of time, this means that we have an excess of bulls – with fewer bulls, we would still be able to inseminate all the cows by using the remaining unused time. For example, if all the cows are inseminated during the first eight hours of their insemination period, this means that we could use half as many bulls.

In the ideal cow-bull mix, all sixteen hours of the cow's insemination period can be used. Thus, for an individual cow, the time Δt from the moment it got into heat to the moment when it is inseminated can be any value between 0 and 16 hours. The average value of Δt is thus about 8 hours.

If there are too many bulls, this means that, in general, all the cows in heat will be inseminated earlier than that – thus, Δt will be, in general, smaller.

On the other hand, if there are too few bulls, this means that many cows will be not inseminated at all, and those "lucky" ones to be inseminated will be inseminated closer to the end of their insemination window. In this case, the average value of the time Δt will be larger.

So, we arrive at the following conclusion:

- if there are too many bulls in the herd, most cows will be inseminated in the earlier part of their insemination window; while

- if there are too few bulls in the herd, most cows will be inseminated in the later part of their insemination window.

Now, we are ready to explain the above phenomenon.

Our explanation.

- When a cow gets inseminated during the earlier part of their insemination window, to the cow's organism, this is an indication that there may be too many bulls in the heard. Thus, as we have mentioned, a natural biological reaction is to decrease this dis-balance by producing mostly female calves.

- On the other hand, when the cow gets inseminated during the later part of their insemination window, to the cow's organism, this is an indication that there may be too few bulls in the heard. Thus, as we have mentioned, a natural biological reaction is to decrease this dis-balance by producing mostly male calves.

And this is exactly the phenomenon that has been observed – which we thus explained from the general system viewpoint.

Chapter 8

Analytical Techniques Help Enhance the Results of Data Mining: Why Pink Noise Is Best for Enhancing Sleep and Memory

This chapter provides a second example of how analytical techniques can help enhance the results of data mining.

This example if related to the fact that, as several researchers found out, acoustic stimulation during sleep enhances sleep and enhances memory. An interesting – and somewhat mysterious – part of this phenomenon is that out of all possible types of noise, the pink noise leads to the most efficient stimulation. In this chapter, we use general system-based ideas to explain why in this phenomenon, pink noise works best.

Comment. The results presented in this chapter first appeared in [17].

8.1 Formulation of the Problem

Acoustic stimulation helps sleep and memory: a brief description of the observed phenomenon. Several researchers found out that acoustic stimulation during sleep help patients to sleep better and enhanced their memory; see, e.g., [87, 92, 95, 125].

Qualitative explanation of the phenomenon. While the level of enhancement was

much higher than most researchers expected, the very fact that exercising some organ is good should be expected: it helps to exercise muscles, it helps to exercise brain activities, it helps to exercise visual activities, etc.

It is also understandable that a noise helps better than a signal emitted at a single frequency: just like exercising different muscles is better for a person's overall health than focusing on a single group of muscles, just like practicing different types of mental activities is better for a person's mental abilities than repeatedly performing tasks of the same type, it is reasonable to expect that processing components of different frequencies will work better than processing only one frequency.

Pink noise: why? The best stimulation results were obtained when the researchers applied *pink noise*, i.e., the noise in which the power spectral density $S(f)$ (i.e., energy per unit frequency) is inverse proportional to the frequency itself: $S(f) = \dfrac{c}{f}$ for some constant c.

From the biological viewpoint, this may not be that surprising, since the pink noise is the most common signal in biological systems; see, e.g., [118] (see also [57] and references therein). In particular, pink noise is a good description of signals corresponding to mental activities; see, e.g., [128]. Pink noise is also ubiquitous in nature in general: e.g., in describes the statistical structure of many natural images; see, e.g., [36].

However, from the scientific viewpoint, the efficiency of pink noise is still somewhat a mystery, since it is not clear why the use of the most biological signal would lead to a better enhancement than any other possible signals.

What we do in this chapter. In this chapter, we provide a system-based explanation of why pink noise is the most efficient one.

8.2 Our Explanation

Main idea. As we have mentioned earlier, the best results are usually achieved when different organs, different parts of the organs are all involved. With respect to frequencies,

this means that the biological acoustic sensors corresponding to all the frequencies should be involved.

Of course, if on one of the frequencies, we have a very weak signal, this means that the sensor corresponding to this frequency is practically not exercised ar all. So, to make sure that the exercise leads to the largest possible effect, it is reasonable to require that each of these sensors is exercised similarly, i.e., in precise terms, that the energy of the part of the signal affecting each sensor will be the same for all sensors.

To describe this idea in precise terms, we need to recall which part of the signal is affecting each of these sensors.

Acoustic perception: a particular case of the general Weber's law. In general, our perception – be it visual or acoustic or any other – follows the *Weber's law*, according to which, for each perceived quantity x, the just noticeable difference Δ_x is proportional to the actual value of this quantity, i.e., $\Delta_x = \delta \cdot x$ for some $\delta > 0$; see, e.g., [60].

In particular, for the frequency, this means that for each frequency f, the just noticeable difference in frequency $\Delta_f(f)$ should be proportional to the frequency itself, i.e., we should have $\Delta_f(f) = \delta \cdot f$, for some value $\delta > 0$. In other words, each biological acoustic sensor corresponding to a certain frequency f actually takes in all the frequencies from f to $f + \Delta_f(f) = f + \delta \cdot f$.

What will happen is we have a signal with power spectral density $S(f)$? By definition, the power spectral density is the energy per unit frequency. Thus, to get the overall energy $E(f)$ affecting this sensor, we need to multiply the power spectral density $S(f)$ by the width $\Delta_f(f)$ of the corresponding frequency interval $[f, f + \Delta_f(f)]$. As a result, we get the value

$$E(f) = S(f) \cdot \Delta_f(f) = S(f) \cdot \delta \cdot f.$$

As we have mentioned, the best effect is expected when each sensor gets the exact same amount of energy, i.e., when $E(f) = \text{const}$. For the above expression $E(f) = S(f) \cdot \delta \cdot f$, the requirement $E(f) = \text{const}$ means that, to achieve the best effect, we should use the

power spectral density $S(f) = \dfrac{\text{const}}{\delta \cdot f}$. This is exactly the pink noise, with $c = \dfrac{\text{const}}{\delta}$. Thus, we have indeed explained why, out of all possible types of acoustic noises, the pink noise leads to the most efficient stimulation of sleep and memory.

Chapter 9

Analytical Techniques Help Enhance the Results of Data Mining: Why Filtering Out Higher Harmonics Makes It Easier to Carry a Tune

This chapter provides a third (and final) example of how analytical techniques can help enhance the results of data mining.

This example if related to the fact that, as a recent patent shows, filtering out higher harmonics helps people sing in-tune. In this chapter, we use the general signal processing ideas to explain this empirical phenomenon. We also show that filtering out higher harmonics is the optimal way of increasing the signal-to-noise ratio – and thus, of making it easier for people to recognize when they are signing out of tune.

Comment. The results presented in this chapter first appeared in [1].

9.1 Formulation of the Problem

A helpful invention. According to the patent description, the patent [42] "greatly improves the singing abilities of both novice and experienced singers by amplifying the fundamental frequency of one's voice to correct tone deafness."

It works but why? The device has been successfully tested, it clearly works, but why?

Why does amplifying the fundamental frequency – or, equivalent;y, filtering out the higher harmonics – helps a person carry a tune?

This is a question to which we plan to find an answer in this chapter. This answer will be based on the general engineering signal processing ideas; see, e.g., [74, 77]. Moreover, not only we explain why this works, we also show that such filtering is the optimal way to make it easier for a person to carry a tune.

9.2 Our Explanation

Higher harmonics: a brief reminder. Each note corresponds to a certain fundamental frequency f_0. The resulting signal is periodic with the same frequency. Thus, if we perform the Fourier transform – i.e., if we represent the signal as a linear combination of sinusoids of different frequencies, then we only get components corresponding to multiples of the fundamental frequency f_0. Components corresponding to frequencies $2f_0$, $3f_0$, etc., are known as *higher harmonics*.

The fact that the frequency f_0 is fundamental means that the component corresponding to this frequency has the largest energy $S(f_0)$; the energies of the higher harmonics are smaller:

$$S(2f_0) < S(f_0); \quad S(3f_0) < S(f_0), \ldots \tag{9.1}$$

Why is it not always easy to carry a tune: signal-processing analysis. In general, in signal processing, the quality of signal detection depends on the signal-to-noise ratio. Thus, if a singing person does not understand that he/she is singing out of tune, this means that for the sound produced by this singing person, the signal-to-noise ratio is too low to detect this.

The overall energy S of the signal can be computed by adding the energies corresponding

to the fundamental frequency and to the higher harmonics:

$$S = S(f_0) + S(2f_0) + S(3f_0) + \ldots \tag{9.2}$$

Similarly, the overall energy N of the noise can be computed by adding the energies of the noise on all these frequencies, i.e.:

- the energy $N(f_0)$ of the noise on the fundamental frequency f_0,
- the energy $N(2f_0)$ of the noise on the double frequency $2f_0$, etc.:

$$N = N(f_0) + N(2f_0) + N(3f_0) + \ldots \tag{9.3}$$

In contrast to signal whose energy changes drastically from one frequency to another, the energy of the noise is usually changing very little from one frequency to another. In the first approximation, we can therefore simply assume that this energy is the same for all the involved frequencies:

$$N(f_0) = N(2f_0) = N(3f_0) = \ldots \tag{9.4}$$

Thus, the formula (9.3) has the form

$$N = k \cdot N(f_0),$$

where k is the overall number of harmonics.

The corresponding signal-to-noise ratio of the original singing signal is thus equal to

$$\frac{S}{N} = \frac{S(f_0) + S(2f_0) + S(3f_0) + \ldots + S(k \cdot f_0)}{k \cdot N(f_0)}. \tag{9.5}$$

The fact that a person has difficulty correctly carrying a tune means that this signal-to-noise ratio is too small. We need to increase it.

Let us apply filtering. In signal processing, a usual way to increase the signal-to-noise ratio is to perform some filtering. Filtering means that we either amplify or decrease certain frequencies. This amplification or damping is applied to the combination of signal and noise,

so it equally affects both. In both cases of amplification or damping, the energy of both the signal component and of the noise component is multiplied by the same coefficient $c(f) \geq 0$ depending on the frequency f:

- for the fundamental frequency f_0, the energy of the signal changes from $S(f_0)$ to $c(f_0) \cdot S(f_0)$ and the energy of the noise changes from $N(f_0)$ to $c(f_0) \cdot N(f_0)$;
- for the frequency $2f_0$, the energy of the signal changes from $S(2f_0)$ to $c(2f_0) \cdot S(2f_0)$ and the energy of the noise changes from $N(2f_0) = N(f_0)$ to $c(2f_0) \cdot N(f_0)$;
- for the frequency $3f_0$, the energy of the signal changes from $S(3f_0)$ to $c(3f_0) \cdot S(2f_0)$ and the energy of the noise changes from $N(3f_0) = N(f_0)$ to $c(3f_0) \cdot N(f_0)$; etc.

After the filtering, the overall energy of the signal is equal to

$$S' = c(f_0) \cdot S(f_0) + c(2f_0) \cdot S(2f_0) + \ldots + c(k \cdot f_0) \cdot S(k \cdot f_0), \tag{9.6}$$

the overall energy of the noise is equal to

$$N' = c(f_0) \cdot N(f_0) + c(2f_0) \cdot B(f_0) + \ldots + c(k \cdot f_0) \cdot N(f_0) =$$

$$(c(f_0) + c(2f_0) + \ldots + c(k \cdot f_0)) \cdot N(f_0), \tag{9.7}$$

and thus, the new signal-to-noise ratio is equal to

$$\frac{S'}{N'} = \frac{c(f_0) \cdot S(f_0) + c(2f_0) \cdot S(2f_0) + \ldots + c(k \cdot f_0) \cdot S(k \cdot f_0)}{(c(f_0) + c(2f_0) + \ldots + c(k \cdot f_0)) \cdot N(f_0)}. \tag{9.8}$$

Which filter is optimal: formulation of the problem. We want to find the coefficients $c(f_0)$, $c(2f_0)$, $\ldots$, $c(k \cdot f_0)$ for which the signal-to-noise ratio (9.8) attains the largest possible value – this will lead to the best possible chance of a person recognizing inaccuiracies in his/her own signing.

The optimal filter is exactly filtering out higher harmonics. Let us prove that the optimal filter is exactly the filter used in [42] – the filter that filters out all higher harmonics, i.e., the filter for which

$$c(f_0) > 0 \text{ and } c(2f_0) = \ldots = c(k \cdot f_0) = 0. \tag{9.9}$$

Indeed, for this filter, the signal-to-noise ratio is equal to

$$\frac{S'}{N'} = \frac{c(f_0) \cdot S(f_0)}{c(f_0) \cdot N(f_0)} = \frac{S(f_0)}{N(f_0)}. \tag{9.10}$$

What happens if at least one of the higher harmonics is not completely filtered out, i.e., we have $c(i \cdot f_0) > 0$ for some i? In this case, for all such harmonics i, we have, due to the inequalities (9.1), $S(i \cdot f_0) < S(f_0)$ hence

$$c(i \cdot f_0) \cdot S(i \cdot f_0) < c(i \cdot f_0) \cdot S(f_0).$$

By adding up these inequalities, we conclude that

$$S' = c(f_0) \cdot S(f_0) + c(2f_0) \cdot S(2f_0) + \ldots + c(k \cdot f_0) \cdot S(k \cdot f_0) <$$

$$c(f_0) \cdot S(f_0) + c(2f_0) \cdot S(f_0) + \ldots + c(k \cdot f_0) \cdot S(f_0) =$$

$$(c(f_0) + c(2f_0) + \ldots_c (k \cdot f_0)) \cdot S(f_0). \tag{9.11}$$

Dividing both sides of this inequality by the expression (9.7) for the noise N', we conclude that

$$\frac{S'}{N'} < \frac{(c(f_0) + c(2f_0) + \ldots + c(k \cdot f_0)) \cdot S(f_0)}{(c(f_0) + c(2f_0) + \ldots + c(k \cdot f_0)) \cdot N(f_0)}. \tag{9.12}$$

Dividing both numerator and denominator of the right-hand side by the same sum $c(f_0) + c(2f_0) + \ldots + c(k \cdot f_0)$, we thus conclude that

$$\frac{S'}{N'} < \frac{S(f_0)}{N(f_0)}. \tag{9.13}$$

Thus, if at least one of the higher harmonics is not fully filtered out, the resulting signal-to-noise ratio is smaller than the value (9.10) that we have when all these harmonics are filtered out.

In particular, by taking the values

$$c(f_0) = c(2f_0) = \ldots = c(k \cdot f_0) = 1$$

corresponding to taking the original signal as is, we conclude that the signal-to-noise ration of the optimally filtered signal is indeed larger than the signal-to-noise ratio of the original signal.

Conclusion. Thus, we have shown the following:

- We showed that filtering out higher harmonics increases signal-to-noise ratio. Thus, we explain why after this filtering, it is easier for a person to detect when he or she is signing out of tune.
- We also showed that filtering out higher harmonics is indeed the optimal approach – in the sense that it leads to the largest possible increase in the signal-to-noise ratio (and thus, to the best chance of detecting out-of-tune deviations).

Chapter 10

Case When Analytical Techniques Invalidate the Conclusions of Data Mining: Reversed Flynn Effect of Decreasing IQ Test Scores

In the previous three chapters, we provided examples when analytical techniques helped to enhance the conclusions of data mining. In this chapter, we provide an opposite example: when analytical techniques help to invalidate the conclusions of data mining.

This example is related to the fact that researchers who monitor the average intelligence of human population have reasonably recently made an unexpected observation: that after many decades in which this level was constantly growing (this is known as the Flynn effect), at present, this level has started decreasing again. In this chapter, we show that this reversed Flynn effect can be, in principle, explained in general system-based terms: namely, it is similar to the fact that a control system usually overshoots before stabilizing at the desired level. A similar idea may explain another unexpected observation – that the Universe's expansion rate, which was supposed to be decreasing, is actually increasing.

Comment. The results presented in this chapter first appeared in [15].

10.1 Formulation of the Problem

IQ tests: a brief reminder. For many decades, researchers have been using standardized test to measure Intelligent Quotient (IQ, for short), a numerical values that describes how smarter is a person that an average population:

- the IQ value of 100 means that this person has average intelligence,
- values above 100 means that this person's intelligence is above average, and
- values below 100 means that this person's intelligence is below average.

Of course, this is a rough estimation. Researchers have known that there are different types of intelligence, and that it is therefore not possible to adequately characterize one person's intelligence by using a single number. However, the IQ test score remains a reasonable overall (approximate) measure both of the individual intelligence and of the relative intelligence of different population groups. For example, a recent study showed that non-violent criminals are, on average, smarter than violent ones; this makes sense, since it takes some intelligence (ill-used but still intelligence) to steal without using violence.

Average IQ scores grow: Flynn's effect. Since the IQ scores describe the relation of a tested person's intelligence to an average intelligence at the given moment of time, researchers periodically estimate this average level of intelligence.

Somewhat unexpectedly, it turned out that for almost 100 years, the average level of intelligence has been growing; see, e.g., [22, 26, 28, 40, 41, 86, 104, 120, 126]. Specifically:

- if we give average current folks the test from the 1930s, they will, on average, score way above 100, and
- vice versa, if we measure the intelligence of the 1930s folks in a current scale, their average intelligence will be way below 100, at about the 80–90 level.

This steady increase in intelligence is known as the *Flynn effect*, after a scientists who actively promoted this idea.

Why IQ scores grow: possible explanation. There are many explanations for the growth in intelligence. One of the natural ones is that, in contrast the old days, when in many professions, physical force was all that is needed to earn a living, nowadays intelligence is very important – non-intelligent jobs have been mostly taken up by machines. No one needs a galley slave to row a boat, no one needs a strong man to lift heavy things, etc. It is therefore reasonable that modern life requires more intelligent activities, and this increase in solving intelligent problems naturally leads to an increased intelligence – just like exercising the muscles leads to an improved physique.

Reverse Flynn effect. While the intelligence scores have been steadily rising for several decades, lately, a reverse phenomenon has been observed, when the average scores no longer grow; instead, they decline. This decline is not as big as to wipe out the results of the previous decades of growth, but it is big enough to be statistically significant; see, e.g., [19, 33, 39, 49, 99, 100, 106, 117, 121, 122].

How can we explain the reverse Flynn effect? There are many different explanations for the reverse Flynn effect: that it has been caused by pollution, that it has been caused by declining education standards, etc.

In this chapter, we analyze this phenomenon from the general systems viewpoint, and conclude that, from the system's viewpoint, a current small decline is natural – and that we therefore do not need to be unnecessarily alarmed by this decline. In other words, in spite of this decline, it is still reasonable to remain optimistic.

10.2 Systems-Based Analysis of the Problem and the Resulting Explanation of the Reversed Flynn Effect

Current explanation of the Flynn's effect reformulated in general terms. The current explanation of the Flynn's effect is that the increase in intelligence is motivated by

the fact that nowadays, more and more important real-world activities require intelligence.

In other words, the previous level of intelligence – which worked optimally in the past – is no longer optimal for adequate functioning in the modern world. Thus, it is necessary to raise the average intelligence to a new higher level, a level that would guarantee effective functioning in this world.

Why general systems approach is necessary. Changing intelligence is not something we directly know how to do. It is a complex process that, probably, involves many different related quantities. The corresponding change in the values of these quantities $x_1, \ldots, x_n$ can be described by an appropriate system of differential equations

$$\frac{dx_i(t)}{dt} = f_i(x_1(t), \ldots, x_n(t)). \tag{10.1}$$

These changes are slow: they are statistically significant and impressive when we compare 1930s with 1990s, but not that noticeable year after year. Suffice it to say that the reversed Flynn effect was not noticed until a decade or so passed when, as it turned out, the intelligence scores were declining. The fact that these changes are slow means that with the passage of time, the values x_i of the corresponding quantities change very little. Let us pick some moment of time t_0. Then, the corresponding differences $\Delta x_i(t) \stackrel{\text{def}}{=} x_i(t) - x_i(t_0)$ are small. Thus, we substitute the expressions $x_i(t) = x_i(t_0) + \Delta x_i(t)$ into the right-hand side of the formula (10.1), expand this right-hand side in Taylor series and keep only linear terms in this expansion. Thus, for the new variables $\Delta x_i(t)$ for which, by the way,

$$\frac{d\Delta x_i(t)}{dt} = \frac{dx_i(t)}{dt},$$

we get a system of linear equations with constant coefficients:

$$\frac{d\Delta x_i(t)}{dt} = c_i + \sum_j c_{ij} \cdot \Delta x_j(t),$$

for appropriate coefficients c_i and c_{ij}.

The general solution to such systems of equations is well known, it depends on the eigenvalues $\lambda = a + b \cdot \mathrm{i}$ of the corresponding matrix c_{ij}, and, in general, contains not only

exponential decrease of the difference between the current and the limit state, but also oscillations (corresponding to $b \neq 0$).

This is known phenomenon in control: in an answer to a perturbation, a stable system usually not just monotonically returns to the original state, it often goes through kind of oscillations: first, it overshoots the original state, then the value goes down and get an undershoot – a smaller one than the original overshoot – then we may get one more overshoot, etc.

How this explains the reversed Flynn effect. In general, when a dynamical system tries to reach a certain level, it usually does not reach this level monotonically. It first overshoots, then undershoots, then may overshoot again, etc. In each such cycle, the deviation between the current and desired values decreases – and eventually, the system stabilizes at this new level.

This is exactly what we observe with the dynamics of average intelligence scores: first, we have a large increase, then a slight decreases. From this viewpoint, we can say that the current slight decrease does not necessarily mean that the population is becoming dumber. There is no need to be pessimistic about the future of mankind. This decline simply means that the natural dynamic phenomena that led to the original increase overshot (as is natural for dynamical systems). Our prediction is thus that this decline will continue to be small, and the resulting average intelligence level will still be higher as in the distance past. After that, we may see another – even smaller – increase, then maybe again decrease, etc.

10.3 Maybe the Same Idea Can Explain the Observed Increase in Universe's Expansion Rate: A Speculative Observation

Phenomenon. It is known, according to modern physics, the Universe expands; see, e.g., [124]. Until the late 1990s, it was assumed that – in accordance with simply physical

models – this expansion occurs at a decreasing rate. However, later observations showed that while this rate may have been indeed decreasing in the past, it is, at present, somewhat increasing; see, e.g., [96, 129]. This phenomenon even won the Nobel Prize in Physics.

Possible system-based explanation. There are many different physical explanation for this phenomenon, e.g., many explanations involving dark matter – to be more precise, using different differential equations describing the dynamics of the mysterious dark matter.

In this case, while in the cosmological time of billions of years, changes are great, year-by-year (and even million years by million years) changes are very small in comparison. Thus, similar to the IQ case, we can use linearization to analyze this phenomenon.

Our above analysis shows that there may be a general system-based explanation for this phenomenon. Namely, in general, on top of the systematic change, we usually have oscillations. Because of these oscillations, even when in the systematic component, accelerations decrease, added oscillation may make it increase or decrease all the time – and this may be a general system-based explanation for the observed phenomenon.

Chapter 11

Analytical Techniques in Hypothesis Testing: Why Area Under the Curve?

Once we have the information about the system, information coming from measurements and from expert estimates, we use this information to come up with a model describing the system. The usual way to come up with such a model is to formulate several different hypotheses and to select the one that best fits the data. There are many statistical techniques for selecting the model, most of them well-justified but some more heuristic – and thus, less reliable. One of such techniques is a widely used area-under-the-curve method. Specifically, to compare two different hypothesis testing techniques, researchers use the following heuristic idea: for each technique, they form a curve describing how the probabilities of type I and type II errors are related for this technique, and then compare areas under the resulting curves. In this chapter, we provide a justification for this heuristic idea.

Comment. The results presented in this chapter first appeared in [13].

11.1 Formulation of the Problem

Type I and type II errors. There are many different techniques for hypothesis testing, i.e., for deciding, based on the observation, whether the original (*null*) hypothesis is valid or whether this hypothesis has to be rejected (and the alternative hypothesis has to be considered true); see, e.g., [111]. In hypothesis testing, we can have two different types of errors:

- a type I error (also known as False Negative) is when the correct null hypothesis is

erroneously rejected, while

- a type II error (also known as False Positive) is when the false null hypothesis is erroneously accepted.

The probability of the type I error is usually denoted by α and the probability of the type II error is usually denoted by β.

In different situations, we have different requirements on the allowed probabilities of these two errors. For example, in early cancer diagnostics, when the null hypothesis means no cancer, type I error is not that critical – it simply means that a healthy patient has to go through an extra testing to make sure that he/she is healthy. On the other hand, a type II error means missing a potentially dangerous disease – which can lead to grave consequences. In such situations, it is desirable to minimize the probability of type II errors as much as possible – even when this leads to a larger type I error.

On the other hand, in law enforcement, we do not want to have too high a probability of type I errors – that would mean SWAT teams breaking into the houses of innocent people in the middle of the night, that would mean massively arresting people who have not done anything wrong.

Depending on the situation, we can adjust the given technique – by changing some appropriate parameters – to increase or decrease α and β. In the ideal world, we should have both errors as small as possible, but this is not possible if all we have is a finite sample. Thus:

- if we decrease α, the probability β increases, and,
- vice versa, if we decrease β, the probability α decreases.

In particular, based on the finite sample, the only way to make sure that we do not have any type I errors is to never reject the null-hypothesis. In this case, however, every time the null hypothesis is false, it will still be accepted. In other words, when the probability α of the type I error is 0, then the probability β of the type II error will be 1.

Vice versa, the only way to get $\beta = 0$ is to never accept the null hypothesis – but in this case, we will have $\alpha = 1$.

How can we compare two hypothesis testing techniques? To get a full description of the quality of a given hypothesis testing technique, we need to indicate, for each $\alpha > 0$, what probability β we can achieve with this technique, and, vice versa, for each β, what probability α we can achieve with this technique. In other words, the perfect description of this quality is a curve that describes how β depends on α – and vice versa. We have a curve $\beta = f(\alpha)$ that describes the dependence of the smallest possible β on the given value α.

For $\alpha = 0$, as we have mentioned earlier, we have $\beta = f(0) = 1$. The larger α we allow, the smaller β can be – so the dependence $f(\alpha)$ is decreasing, and it reaches the value $f(1) = 0$ for $\alpha = 1$.

How do we compare the two hypothesis testing techniques? If the required value α is given, we select the technique for which the corresponding value β is the smallest – and, vice versa, if the value β is given, we select the technique for which the corresponding value α is the smallest.

This requires that we implement all possible hypothesis testing techniques, and every time select a technique depending on the specifics of a situation. In reality, however, there are dozens and dozens of different hypothesis testing techniques. In practice, it is often not realistically possible to implement all of them on the available computational device. In such situations, we select one of the techniques and use it – with varying parameters – for all possible practical situations. (Alternatively, we can select two or more techniques – and for each given values of α or β, select the best of these. This is equivalent to selecting a *hybrid* technique: e.g., technique A for values of α which are smaller than some threshold α_0 and a technique B for all other values α.)

In all such cases, we select one of the techniques (either one of the original ones or one of the hybrid ones). The question is: how do we select? One technique may be better for

some α, another may be better for another α. The usual way to select one of the available techniques is to select the one for which the Area Under the Curve (AUC) $\int\limits_0^1 f(\alpha)\,d\alpha$ is the smallest possible.

Comment. Instead of the dependence of β on α, we can plot the dependence of $1-\beta$ on α: $1-\beta = g(\alpha)$. For this new function, the area under its curve is equal to 1 minus the area under the f-curve:

$$\int_0^1 g(\alpha)\,d\alpha = \int_0^1 (1-f(\alpha))\,d\alpha = 1 - \int_0^1 f(\alpha)\,d\alpha.$$

Thus, for these functions, minimizing the area under the f-curve is equivalent to maximizing the area under the g-curve.

Why? In practice, the AUC criterion seems to lead to reasonable results. A natural question is: why? Alternatively, we could, e.g., take different values $f(\alpha)$ with different weights $w(\alpha)$ and compare the weighted values $\int w(\alpha)\cdot f(\alpha)\,d\alpha$; so why AUC?

In this chapter, we provide a possible explanation for the empirical efficiency of the Area Under the Curve criterion.

11.2 Our Explanation

Analysis of the situation. In practice, we usually have bound on both types of error, i.e., we have bounds α_0 and β_0 for which we would like to have $\alpha \le \alpha_0$ and $\beta \le \beta_0$.

Can we achieve this requirement by using a hypothesis testing technique with a given curve $f(\alpha)$? If we cannot achieve the desired values β_0 for some $\alpha < \alpha_0$, not all is lost: we may still be able to get the desired probability of the type II error if we allow higher type I errors. Thus, to test whether the given requirements can be achieved, we should take the largest allowed value α_0 of the type I error and check whether for this value, we can get $\beta \le \beta_0$, i.e., whether we have $f(\alpha_0) \le \beta_0$.

This inequality corresponds to the point (α_0, β_0) being above the curve $\beta = f(\alpha)$. If the point (α_0, β_0) is below this curve, this means that for this hypothesis testing technique,

the corresponding requirement cannot be satisfied.

For each technique, some requirements can be satisfied, some cannot. A natural measure of the technique's quality is the frequency with which this technique succeeds – i.e., in more precise terms, the probability that this technique will succeed.

To formalize this idea, we need to select a probability distribution on the set of all pairs (α_0, β_0). To estimate the probability that a given pair of probabilities (α_0, β_0) can be achieved, we need to select some probability distribution on the set of all such pairs.

On the unit square $[0, 1] \times [0, 1]$, there are many possible probability distributions. Such situations are ubiquitous in applications of probabilistic methods. In such situations of uncertainty, a reasonable idea is not to pretend that we have less uncertainty than we do – and thus, put of all probability distributions consistent with our knowledge, to select the probability distribution with the largest uncertainty. A natural measure of the distribution's uncertainty is its entropy [56, 88]. Thus, the idea is to select the probability distribution for which the entropy $S = -\int \rho(x) \cdot \ln(\rho(x))\, dx$ is the largest possible; see, e.g., [56].

This indeed explains the AUC. It is known that among all possible probability distributions located on the unit square, the uniform distribution has the largest entropy. For the uniform distribution, the probability that the randomly selected requirements can be implemented by this technique – i.e., that randomly selected pair (α_0, β_0) will be under the curve $\beta = f(\alpha)$ – is equal to the area under this curve.

Thus, when comparing two techniques, we should indeed select the one for which the area under the f-curve is the smallest possible – or, equivalently, the technique for which the area under the g-curve is the largest possible.

Chapter 12

It Is Important to Revisit the Selection of the Best Model When New Data Appear: Why Confirmation Bias Is a Faulty Strategy

One of the biases potentially affecting systems engineers is the confirmation bias, when instead of selecting the best hypothesis based on the data, people stick to the previously-selected hypothesis until it is disproved. In this chapter, on a simple example, we show how important is to take care of this bias: namely, that because of this bias, we need twice as many experiments to switch to a better hypothesis.

Comment. The results presented in this chapter first appeared in [5].

12.1 Formulation of the Problem

Confirmation bias. It is known that our intuitive reasoning shows a lot of unexpected biases; see, e.g., [61]. One of such biases is a *confirmation bias*, when, instead of selecting the best hypothesis based on the data, people stick to the previously-selected hypothesis until it is disproved. This bias is ubiquitous in systems engineering; see, e.g., [24, 85, 98, 114].

How important is it to take the confirmation bias into account? Taking care of the confirmation bias requires some extra effort; see, e.g., [64, 113, 114, 130] and references therein. A natural question is: is the resulting improvement worth this extra effort? How better the result will we get?

In this chapter, on a simple example, we show that the result is drastically better: namely, that if we properly take this bias into account, then we will need half as many experiments to switch to a more adequate hypothesis.

12.2 Analysis of the Problem

Description of the simple example. Let us consider the simplest possible case when we have a parameter a that may be 0 and may be non-zero, and we directly observe this parameter. We will also make the usual assumption that the observation inaccuracy is normally distributed, with 0 mean and known standard deviation σ.

In this case, what we observe are the values $x_1, \ldots, x_n$ which are related to the actual (unknown) value a by a relation $x_i = a + \varepsilon_i$ $(i = 1, \ldots, n)$, where ε_i are independent normally distributed random variables with 0 mean and standard deviation σ.

Two approaches. In the ideal approach, we select one of the two models – the null-hypothesis $a = 0$ or the alternative hypothesis $a \neq 0$ – by using the usual Akaike Information Criterion (AIC); see, e.g., [111].

In the confirmation-bias approach, we estimate the value a based on the observations $x_1, \ldots, x_n$, and we select the alternative hypothesis only if the resulting estimate is statistically significantly different from 0 – i.e., e.g., that the 95% confidence interval for the value a does not contain 0.

What if we use AIC. In the AIC, we select a model for which the difference $\mathrm{AIC} \stackrel{\mathrm{def}}{=} 2k - 2\ln\left(\widehat{L}\right)$ is the smallest, where k is the number of parameters in a model and $\widehat{L}$ is the largest value of the likelihood function L corresponding to this model.

The null-model $a = 0$ has no parameters at all, so for this model, we have $k = 0$. For n independent measurement results, the likelihood function is equal to the product of the values

$$\frac{1}{\sqrt{2\pi}\cdot\sigma}\cdot\exp\left(-\frac{x_i^2}{2\sigma^2}\right)$$

of the Gaussian probability density function corresponding to these measurement results x_i. Thus,

$$L=\prod_{i=1}^{n}\frac{1}{\sqrt{2\pi}\cdot\sigma}\cdot\exp\left(-\frac{x_i^2}{2\sigma^2}\right)$$

and so, for this model,

$$\mathrm{AIC}_0=-2\ln(L)=2n\cdot\ln\left(\sqrt{2\pi}\cdot\sigma\right)+\frac{1}{\sigma^2}\cdot\sum_{i=1}^{n}x_i^2.$$

We assume that $x_i=a+\varepsilon_i$, where the mean value of ε_i is 0 and the standard deviation is σ. Thus, the expected value of x_i^2 is equal to $a^2+\sigma^2$. For large values n, due to the Law of Large Numbers (see, e.g., [111]), the average $\dfrac{1}{n}\cdot\sum_{i=1}^{n}x_i^2$ is approximately equal to the expected value $E[x_i^2]=a^2+\sigma^2$. Thus, $\sum_{i=1}^{n}x_i^2\approx n\cdot(a^2+\sigma^2)$ and hence,

$$\mathrm{AIC}_0=2n\cdot\ln\left(\sqrt{2\pi}\cdot\sigma\right)+\frac{1}{\sigma^2}\cdot n\cdot(a^2+\sigma^2). \tag{12.1}$$

The alternative model $a\neq 0$ has one parameter a, so here $k=1$. The corresponding likelihood function is then equal to

$$L=\prod_{i=1}^{n}\frac{1}{\sqrt{2\pi}\cdot\sigma}\cdot\exp\left(-\frac{(x_i-\widehat{a})^2}{2\sigma^2}\right).$$

We select the parameter a that maximizes the value of this likelihood function. Maximal likelihood is the usual way of estimating the parameters, which in this case leads to $\widehat{a}=\dfrac{1}{n}\cdot\sum_{i=1}^{n}x_i$. For large n, this estimate is close to the actual value a, so we have

$$\widehat{L}=\prod_{i=1}^{n}\frac{1}{\sqrt{2\pi}\cdot\sigma}\cdot\exp\left(-\frac{(x_i-a)^2}{2\sigma^2}\right).$$

For this model, $x_i - a = \varepsilon_i$, thus,

$$\mathrm{AIC}_1 = 2 - 2\ln\left(\widehat{L}\right) = 2 + 2n \cdot \ln\left(\sqrt{2\pi}\cdot\sigma\right) + \frac{1}{\sigma^2}\cdot\sum_{i=1}^{n}\varepsilon_i^2.$$

For large n, we have $\sum\limits_{i=1}^{n}\varepsilon_i^2 \approx n\cdot\sigma^2$, hence

$$\mathrm{AIC}_1 = 2 + 2n \cdot \ln\left(\sqrt{2\pi}\cdot\sigma\right) + \frac{1}{\sigma^2}\cdot n\cdot\sigma^2. \tag{12.2}$$

The second model is preferable if $\mathrm{AIC}_1 < \mathrm{AIC}_0$. By deleting common terms in these two values AIC_i, we conclude that the desired inequality reduces to $2 < \dfrac{n\cdot a^2}{\sigma^2}$, i.e., equivalently, to

$$n > \frac{2\sigma^2}{a^2}. \tag{12.3}$$

What if we use a confirmation-bias approach. In the confirmation-bias approach, we estimate a – and we have already mentioned that the optimal estimate is $a = \dfrac{1}{n}\cdot\sum\limits_{i=1}^{n} x_i$. It is known (see, e.g., [111]) that the standard deviation of this estimate is equal to $\sigma_e = \dfrac{\sigma}{\sqrt{n}}$. Thus, the corresponding 95% confidence interval has the form $[a - 2\sigma_e, a + 2\sigma_e]$. The condition that this interval does not contain 0 is equivalent to $|a| > 2\sigma_e$, i.e., equivalently, to $a^2 > 4\sigma_e^2$. Substituting the above expression for σ_e into this inequality, we conclude that $a^2 > 4\cdot\dfrac{\sigma^2}{n}$, i.e., equivalently, that

$$n > \frac{4\sigma^2}{a^2}. \tag{12.4}$$

Conclusion. By comparing the expressions (3) and (4) corresponding to the two approaches, we can indeed see that the confirmation-bias approach requires twice as many measurements than the approach in which we select the best model based on the data.

Thus indeed, avoiding confirmation bias can lead to a drastic improvement in our estimates and thus, in our decisions.

Chapter 13

Need for a Careful Comparison between Hypotheses: Case Study of Epicycles

In this chapter, we illustrate the need for a careful comparison between different hypotheses on one of the most well-known historical examples – epicycles versus more modern techniques in celestial mechanics. Our conclusion is that, contrary to what one may read in modern astronomy and physics textbooks, epicycles were actually a very efficient tool, in some aspects foreseeing modern techniques such as Fourier series – while not exactly as efficient as Fourier series.

Comment. The results presented in this chapter first appeared in [3].

13.1 Epicycles: Bad Science or Genius Idea

Epicycles: what they are. For purposes of navigation, since the ancient times, astronomers have studied the visible motion of the planets and the stars. In the first crude approximation, their trajectories form a circle.

To provide a more accurate description, astronomers proposed the following idea – called *epicycles*. First step is to assume that while a point corresponding to a planet follows a circular motion around the Earth, the planet itself performs a circular motion around this point. To get an even more accurate description, we assume that it is not the planet itself, but rather the second point associated with the planet that performs a

circular motion around the first point. The planet itself performs a circular motion around the second point.

An even more accurate description is that for each planet, there is a third point that performs a circular motion around the second point, and the planet itself moves in a circle around this third point, etc.

This idea was originally proposed by Apollonius of Perga (late 3rd — early 2nd centuries BCE), developed by several others, and finalized by Claudius Ptolemy (≈100–≈170).

Epicycles: traditional textbook-based negative coverage. Traditional textbooks on history of science treated epicycles as bad science, a bad idea that was overcome by the genius of Nicolaus Copernicus (1473–1543).

Epicycles as Fourier (trigonometric) series. From the mathematical viewpoint, a circular rotation around the origin of the coordinate system can be described in simple trigonometric terms:

$$x(t) = r \cdot \cos(\omega \cdot t + \varphi), \quad y(t) = r \cdot \sin(\omega \cdot t + \varphi),$$

i.e., equivalently, as

$$x(t) = x_0 \cdot \cos(\omega \cdot t) - y_0 \cdot \sin(\omega \cdot t);$$

$$y(t) = x_0 \cdot \sin(\omega \cdot t) + y_0 \cdot \cos(\omega \cdot t).$$

In these terms, Ptolemy's description means that we represent the motion of a planet as a sum of such motions – i.e., as a linear combination of sines and cosines, a combination that we now call *trigonometric series* or *Fourier series*, after Joseph Fourier (1768–1830), a researcher who, in effect, reinvented them for modern times.

Fourier series is exactly how we now describe the visible motion of the planets (see, e.g., [73]). Taking into account the ubiquity of Fourier series and their importance for science and engineering (see, e.g., [91]), the epicycles idea sounds more like a stroke of genius than the bad science as described by traditional textbooks.

What is better: Fourier series or epicycles? Epicycles are, in effect, Fourier series. So, any trajectory that can be represented by an epicycle can also be represented by the corresponding Fourier series.

Natural question are:

- are these two representations truly equivalent, i.e., can we represent each Fourier series motion in terms of epicycles?
- if yes, which representation is better? is there any computational advantage in using Fourier series in contrast to epicycles?

These are the questions that we deal with in this chapter.

13.2 Analysis of the Problem and the Resulting Conclusions

Can any Fourier series be represented in epicycle terms? In the above text, we showed that every epicycle-based motion can be represented in terms of Fourier series. A natural question is: is the opposite true? Can any Fourier-series motion be represented in epicycle terms?

When the dependence of both coordinates x and y on time is described by Fourier series, the question is whether we can separately represent the x-motion and the y-motion by epicycles; if we can, then by adding these two representations, we will be able to represent any Fourier motion this way. Thus, the question is: can epicycles represent the purely x-motion in which $y = 0$ or the purely y-motion in which $x = 0$?

The answer to this natural question was provided, for the first time, by the Arabic astronomer Nasir al-Din at-Tusi (1201–1274); see, e.g., [108]. One of the solutions that he proposed to represent an x-only motion is to have a circle of half-radius rotate inside a circle of the original radius. In this case, a point on a small circle moves only along the

x-axis. An even simpler solution would be to represent an x-only motion

$$x(t) = r \cdot \cos(\omega \cdot t + \varphi), \quad y(t) = 0$$

as the sum

$$(x(t), y(t)) = (x_1(t), y_1(t)) + (x_2(t), y_2(t)),$$

where

$$x_1(t) = \frac{r}{2} \cdot \cos(\omega \cdot t + \varphi), \quad y_1(t) = \frac{r}{2} \cdot \sin(\omega \cdot t + \varphi)$$

and

$$x_2(t) = \frac{r}{2} \cdot \cos((-\omega) \cdot t + \varphi) = \frac{r}{2} \cdot \cos(\omega \cdot t + \varphi);$$
$$y_2(t) = \frac{r}{2} \cdot \sin((-\omega) \cdot t + \varphi) = -\frac{r}{2} \cdot \sin(\omega \cdot t + \varphi).$$

Similarly, we can represent an y-only motion.

Comment. It is worth mentioning that, as shown in [108], Copernicus used al-Tusi results in his analysis of celestial mechanics – to the extent that he even borrowed some illustrative pictures from a translation of al-Tusi's book.

So which representation is computationally more efficient? In both cases – of using epicycles and of using trigonometric series – we approximate the observed motion $x(t), y(t))$ by a linear combination of several standard motions $(x_i(t), y_i(t))$:

$$x(t) \approx \sum_i a_i \cdot x_i(t), \quad y(t) \approx \sum_i a_i \cdot y_i(t);$$

see, e.g., [111].

In general, there is a small advantage in using trigonometric series. This advantage is related to the fact that in many cases, the motion is close to linear. In this case, e.g., for the x-only motion:

- in trigonometric series, we need two terms for each frequency ω, namely, we need a linear combination of the terms

  $$(x(t), y(t)) = (\cos(\omega \cdot t), 0) \text{ and } (x(t), y(t)) = (\sin(\omega \cdot t), 0),$$

  while

- to represent this particular frequency component via epicycles, we need two at-Tusi-type pairs – i.e., we need four terms and thus, four (twice as many) coefficients.

This advantage is not that big for the regular planetary motions, since such motions are indeed close to circular, but it is important for other motions – and for higher-frequency (i.e., higher-order) terms describing this motion.

Chapter 14

Analytical Techniques Help in Emulating Biological Systems: An Explanation of Why High-Level Attention Constantly Oscillates

Once we have an adequate description of the users' preferences and of the corresponding application domain, we need to come up with a system design which the most appropriate for this setting. One of the ways to come up with a good design is to use the experience of successful similar systems – engineering and even biological. Examples of such systems are plentiful: many situations in engineering and in life require constant monitoring. At first glance, this would necessitate the need for the system to maintain the same alert level. However, interestingly, recent experiments has shown that in many situations like driving, the driver's attention level constantly oscillates. In Chapter 14, we show that such an oscillation is indeed helpful – and thus, it is necessary to emulate such an oscillation when designing automatic systems, e.g., for driving.

Comment. The results presented in this chapter first appeared in [14].

14.1 Formulation of the Problem

In many real-life situations, high level of attention is crucial. In many practical situations, we concentrate on a certain task. For example, when a person drives a car,

he/she needs to keep a close attention to the road, to make sure that if a problem appears, the driver will react as soon as possible – and thus, avoid a possible accident.

What researchers assumed. In critical situations, when the maximum attention is needed, psychologists assumed that the attention is consistently kept at the maximum possible level – of course, until the person becomes too tired to maintain this level of attention.

This assumption makes perfect sense: when a lot is at stake, including the person's own life, it makes sense to concentrate all the energy on avoiding possible catastrophic situations.

A recent surprising observation. Surprisingly, recent experiments showed that while the attention indeed remains high, the attention level – as measured, e.g., by the reaction time – constantly oscillates; see, e.g., [51, 54]. This level remains high, but the reaction time still oscillates between the smallest possible value and a much larger value. This larger value of reaction time is still good, but not as perfect as the smallest value.

The problem. It is not clear what is the reason for this observed phenomenon. Are they somehow needed for survival? Or are they due to an imperfection of human physiology?

This is not just an interesting theoretical question, it has practical applications:

- If the oscillations indeed improve the system's performance, then we should add similar oscillations to the self-driving cars and other automated vehicles and systems.
- On the other hand, if the oscillations are caused by imperfections of human physiology, then we should not emulate human drives in this; we should instead keep the computer's attention level constant.

What we do in this chapter. In this chapter, we show that oscillations do make the system more efficient – and thus, appropriate oscillations should be implemented in automatic control systems.

14.2 Analysis of the Problem

Need for a numerical model. To analyze the problem, to see whether constant attention of oscillating attention are more productive, we need to formulate this problem in precise numerical terms. Let us therefore describe a simple simplified model of this phenomenon.

Towards a simplified model. Let T denote the duration of the period during which we need to maintain high attention level. Without losing generality, we can start counting time from the beginning of this period. In this case, the corresponding time interval takes the form $[0, T]$.

There are natural limitations on how many observations we can process, whether in a computer or in our brains. For a high-performance computer, these limitations are higher than for a simple laptop, but they are still there. These limitations are real: e.g., when a conference speaker makes a presentation remotely (e.g., by skype), the system often does not catch up when the speaker's movements are too fast.

Let us assume that, because of these limitations, during a certain period of time T, we can process at most N observations. In crucial situations requiring high attention, it is important that the person concentrates on the corresponding task as much as possible – and thus, that this person processes as much information as possible. This means that in such situations, we should process the maximum possible number of observations: namely, we should process exactly N observations during the time T.

These observations correspond to, in general, different moments of time. Let us sort these moments of time in chronological order. For each i from 1 to N, let us denote the time of the i-th observation by t_i. Then, we have

$$0 \le t_1 \le t_2 \le \ldots \le t_N \le T.$$

We want to detect possible obstacles as early as possible, at the time when the corresponding signals are still weak. For weak signals, a single observation is not sufficient for reliable detection, since there is always some noise level: we are not sure that the observed

signal is real or just a noise. Swerving every time when a speck appears which may be a car or a pedestrian is also a sure recipe for disaster: this means that a car would follow an unpredictable waving trajectory, like when the driver in drunk. We need to perform correcting actions only when we are reasonably sure that there is indeed a problem on the road.

The more observations confirm that there is a problem, the higher our level of confidence that this problem is real. Let m denote the smallest number of observations that make us confident. Then, if a problem appears at time $t \in [0, T]$, we will detect it when m observations pass after this time t. Let $i(t)$ denote the first index i for which $t_i \geq t$. The problem can then we observed in observations made at times $t_{i(t)}$, $t_{i(t)+1}$, $t_{i(t)+2}$, etc. The problem will be detected after m such observations, i.e., at the moment $t_{i(t)+m-1}$. The difference $\Delta(t) \stackrel{\text{def}}{=} t_{i(t)+m-1} - t$ between the time when we detect the problem and the original time t is the main component of the reaction time.

For problems appearing at the end of the time period $[0, T]$, namely for problems corresponding to times $t > t_{N-m}$, there are not enough remaining observations to observe this problem.

Definition 14.1.

- *By an* high-attention situation*, we mean a tuple (T, N, m), where $T > 0$ is a real number, and m and N are integers for which $m < N$.*
- *For each high-attention situation, by a* strategy*, we mean an increasing sequence of real numbers $t_1, \ldots, t_N$ for which $0 \leq t_1 \leq t_2 \ldots \leq t_N$.*
- *For a given strategy and for each moment $t \in [0, T]$, by the* reaction time $\Delta(t)$*, we mean the difference $t_{i(t)+m-1} - t$.*

Comment. As we have mentioned earlier, the reaction time is defined only for moments $t \leq t_{N-m}$.

Which strategy should we prefer? We want to minimize reaction time. First of all, we want to make sure that no matter when the problem appears, we should be able to deal with it within a reasonable time r – and this time should be as small as possible. This means that for all the moments $t \le T - r$, we should have $\Delta(t) \le r$. This guaranteed reaction-time r should be as small as possible.

There may be several different strategies with the same worst-case reaction time. To select between then, it is reasonable to choose the strategy with the smallest possible *average reaction time*: the average value of $\Delta(t)$ over all the moments $t \in T - r$. Thus, we arrive at the following definition.

Definition 14.2.

- *For each strategy t_i, by its* worst-case reaction time $r_w(t_i)$*, we mean the smallest positive real number r for which* $\max\limits_{0 \le t \le T-r} \Delta(t) \le r$.
- *For a strategy t_i with worst-case reaction time r, by its* average reaction time $r_a(t_i)$*, we mean the value*

$$r_a(t_i) \stackrel{\text{def}}{=} \frac{1}{T-r} \cdot \int_0^{T-r} \Delta(t)\, dt.$$

14.3 Oscillations Are Better: Proofs

Discussion. Let us use the above model to check which strategy is better: the strategy is constant or the strategy in which attention is oscillating. Let us describe these strategies in precise terms.

Constant level of attention: how to formalize. Constant level of attention means that we have the exact same difference $\delta = t_{i+1} - t_i$ between the two consecutive observations, i.e., that $t_2 = t_1 + \delta$, $t_3 = t_2 + \delta = t_1 + 2\delta$, etc., all the way to $t_N = T$.

In this case, the worst-case reaction time is $r = m \cdot \delta$ that occurs if the problem appears right after each observation, at time $t = t_i + \varepsilon$ for some small positive $\varepsilon \ll \delta$. To maintain the same reaction time for $t = 0$, it is sufficient to take $t_1 = \delta$, thus, $t_i = i \cdot \delta$. So, $\delta = T/N$.

Since we ignore moments $t > T-r$, we can as well place all the moments t_i corresponding to these times at $T - r$.

Definition 14.3. *By a* uniform strategy*, we mean the strategy in which* $t_i = i \cdot (T/N)$ *for* $i < N - m$ *and* $t_i = (N - m) \cdot (T/N)$ *for* $i \geq N - m$.

Proposition 14.1. *For the uniform strategy, the worst-case reaction time is* $r_w = m \cdot (T/N)$*, and the average reaction time is*

$$r_a = \left(m - 0.5 - \frac{m \cdot (m-1)}{2(N-m)} \right) \cdot (T/N).$$

Proof. For the worst-case reaction time, the result is straightforward.

For the average reaction time, the interval $[0, T - r]$ is divided into $N - m$ intervals $[t_{i-1}, t_i]$ of equal width δ. Thus, to compute the average reaction time over the whole interval $[0, T - r]$, it is sufficient to compute the average reaction time over each of these small intervals, and then compute the arithmetic average of these averages.

For the first $N - 2m + 1$ intervals $[t_{i-1}, t_i]$, the reaction time changes between the maximal value $m \cdot \delta$ (attained close to t_{i-1}) and the smallest value $(m-1) \cdot \delta$ (attained at t_i), so the average over this interval is $(m - 0.5) \cdot \delta$.

For $i = N - m + 2$, we have reaction time changing from $(m-1) \cdot \delta$ to $(m-2) \cdot \delta$, with an average $((m-1) - 0.5) \cdot \delta$. For the next interval, we have $((m-2) - 0.5) \cdot \delta$, etc., all the way to $0.5 \cdot \delta$ for the last interval.

In general, the average over each interval has the form $(j - 0.5) \cdot \delta$, where in $N - 2m + 1$ cases, we have $j = m$, and then we have $m - 1$ values $j = m - 1$, $j = m - 2$, $\ldots$, $j = 1$. So, the average reaction time is equal to $r_a = (E[j] - 0.5) \cdot \delta$, where $E[j]$ is the average value of j. Here,

$$E[j] = \frac{m \cdot (N - 2m + 1) + (m-1) + \ldots + 1}{N - m} = \frac{m \cdot (N - 2m + 1) + \dfrac{(m-1) \cdot m}{2}}{N - m}.$$

Since $N - 2m + 1 = (N - m) - (m - 1)$, we get

$$E[j] = \frac{m \cdot (N - m) - m \cdot (m - 1) + \dfrac{(m - 1) \cdot m}{2}}{N - m} =$$

$$\frac{m \cdot (N - m) - \dfrac{(m - 1) \cdot m}{2}}{N - m} = m - \frac{(m - 1) \cdot m}{2(N - m)}.$$

Substituting this expression for $E[j]$ into the formula $r_a = (E[j] - 0.5) \cdot \delta$, we get the desired result.

The proposition is proven.

Oscillations: how to formalize. Let us consider the extreme case of oscillations, where instead of having observations at uniformly distributed times, we bring observations in groups of m: no observations, then m of them in a row, then again no observations, then N of them in a row, etc., until we reach the last m values, i.e., the values starting with $k \cdot m + 1$, where $k = \lfloor N/m \rfloor$:

$$t_1 = t_2 = \ldots = t_m = r, \quad t_{m+1} = \ldots = t_{2m} = 2r, \ldots,$$

$$t_{k \cdot m+1} = \ldots = t_N = k \cdot r. \tag{14.1}$$

Definition 14.4. *By a* maximally oscillating strategy*, we mean the sequence (14.1), where* $k = \lfloor N/m \rfloor$ *and* $r = T/k$.

Proposition 14.2. *For the maximally oscillating strategy, the worst-case reaction time is* $r_w = T/k$*, and the average reaction time is*

$$r_a = T/(2k).$$

Discussion. For the case when N is divisible by m, we get $k = N/m$. In this case, the worst-case reaction time $r_w = T/k = m \cdot (T/N)$ is the same as for the uniform strategy.

However, the average reaction time is almost twice smaller. Thus, the oscillations indeed make the strategy more efficient.

Proof of Proposition 14.2. For the worst-case reaction time, the proof is straightforward. On each interval of width r, the reaction time changes from 0 to r. For each value t from 0 to r, the reaction time is $r - t$. Thus the average reaction time is

$$\frac{1}{r} \cdot \int_0^r (r - t)\, dt = \frac{1}{r} \cdot \left(r \cdot t - \frac{t^2}{2} \right) \Bigg|_0^r = \frac{1}{r} \cdot \left(r^2 - \frac{r^2}{2} \right) = \frac{1}{r} \cdot \frac{r^2}{2} = \frac{r}{2}.$$

The proposition is proven.

Discussion. It is possible to show that not only the maximally oscillating strategy is better than the uniform strategy, it is actually the best possible.

Definition 14.5. *Let an high-attention situation (T, N, m) be given. We say that a strategy t_i is* optimal *of for every other strategy t'_i, we have:*

- *either $r_w(t_i) < r_w(t'_i)$,*
- *or $r_w(t_i) = t_w(t'_i)$ or $r_a(t_i) \le r_a(t'_i)$.*

Proposition 14.3. *For each high-attention situation, the maximally oscillating strategy is optimal.*

Proof. Let us assume that we have an optimal strategy, and that its worst-case reaction time is equal to $r = r_w(t_i)$. For the maximally oscillating strategy, we have $r_a(t_i) = 0.5 \cdot r$. Let us show that, vice versa, we cannot have $r_a < 0.5 \cdot r$, and that if $r_a = 0.5 \cdot r$, then the corresponding strategy is maximally oscillating. This will prove that the maximally oscillating strategy is indeed optimal.

Indeed, the fact that the worst-case reaction time is equal to r means there exists a moment t_0 for which $\Delta(t)$ is as close to r as possible. This, in turn, means that between the moments t_0 and $t_0 + r - \varepsilon$, there are m values t_i, namely the values $t_{i(t_0)}$, $t_{i(t_0)+1}$, $\ldots$, $t_{i(t_0)+m-1}$. If all these m values are equal to each other, then for each moment t between

t_0 and the common value of t_i, we get $\Delta(t) = t_{i(t_0)} - t = t_0 + r - t$, and thus, the average value of $\Delta(t)$ over the corresponding interval is equal to $0.5 \cdot r$.

In general, the next m values t_i cannot be earlier that t_0+r, thus we have $\Delta(t) \geq t_0+r-t$. If for some t, we get strict equality, then the average reaction time over the corresponding interval is $> 0.5 \cdot r$. The only possibility to have this part of t_a equal to $0.5 \cdot r$ is when for all t, we have $\Delta(t) = t_0 + r - t$.

Let us show that in this case, we have at least m values t_i equal to $t_0 + r$. Indeed, let j be the last value for which $t_j < t_0 + r$. Then, any t between t_j and $t_0 + r$, the fact that we have $\Delta(t) = t_0 + r - t$ means that the next m values t_i must be $\leq t_0 + r$. Since the only value t_i between t and $t_0 + r$ is the value $t_0 + r$, this means that we have at least m values equal to $t_0 + r$. Thus, for the optimal solution, we have a group of at least m equal values, then another group of at least m equal values, etc.

If we group t_i into groups of size $> m$, then we would be divide the interval $[0, T]$ into fewer pieces than in the case when each group has exactly m values t_i. So, in this case, the distance between two consecutive groups will be larger than in the case when we have the division into groups of m; thus, this arrangement cannot be optimal. Hence, in the optimal arrangement, we should have m indices in each group of equal consecutive values t_i. This is exactly the oscillating arrangement. The proposition is proven.

Chapter 15

Analytical Techniques for Taking into Account Several Aspects of a Designed Systems: Case Study of Computation-Communication Tradeoff

Once we have an adequate description of the users' preferences and of the corresponding application domain, we need to come up with a system design which the most appropriate for this setting. One of the challenges in searching for such a design is that we need to take into account many different aspects of the resulting system. In many practical situations, for each aspect, we have well-defined optimal design strategies, but there is no analytical techniques for taking all the aspects into account. In Chapter 15, we show how several aspects can be taken into account on the example of a tradeoff between computation and communication needs.

The need for such a tradeoff can be illustrated by a recent study of chimpanzees. This study has shown that on the individual basis, the chimpanzees are, surprisingly, much better than humans in simple tasks requiring intelligence and memory. A usual explanation – called cognitive tradeoff – is that a human brain has sacrificed some of its data processing (computation) abilities in favor of enhancing the ability to communicate; as a result, while individual humans may not be as smart as possible, jointly, we can solve complex problems.

A similar cognitive tradeoff phenomenon can be observed in computer clusters: the most efficient computer clusters are not formed from the fastest, most efficient computers, they are formed from not-so-fast computers which are, however, better in their communication abilities than the fastest ones. In this chapter, we propose a simple analytical model that explains the cognitive tradeoff phenomenon.

Comment. The results presented in this chapter first appeared in [10].

15.1 Formulation of the Problem

Interesting empirical phenomenon. A recent study of chimpanzees [27, 80, 82] showed, somewhat surprisingly, that on the individual basis, they are much better than human in many tasks requiring intelligence. For example, they can remember more objects in images, and in conflict situations their behavior is much closer to the optimal behavior (as recommended by game theory) than the behavior of humans.

Cognitive tradeoff: an explanation for this phenomenon. A current explanation for this phenomenon is based on what is called *cognitive tradeoff*: humans have better communication abilities, and so, human brain has to sacrifice some individual intellectual abilities to leave space for parts needed for efficient communication.

The need for such a tradeoff is not limited to humans. A similar tradeoff phenomenon can be observed not only in humans, but in computers as well. The world's fastest computations are performed on so-called high performance computers. Each of them is, in effect, a large number of processors constantly communicating with each other.

In principle, there exist processors which are very fast and efficient, but modern supercomputers are not formed from these processors: they are formed from simpler processors – similar to the ones we use in not-very-expensive home computers. One of the reasons for this choice is that these simple processors communicate well, as opposed to more efficient processors; these more efficient processors individually perform better but which take much

longer time to communicate (another reason is that simple processors are usually much cheaper, which allows the designers to combine many more such processors within the same budget).

The ubiquity of cognitive tradeoff motivates the desired to have a universal quantitative model. The fact that cognitive tradeoff occurs in many situations, from human to computer communications, shows that there must be a simple quantitative explanation for this phenomenon.

In this chapter, we provide a simple quantitative model that explains the main ideas behind this phenomenon. We hope that this simple model can be used as a basis for more complex – and more realistic – models that would not only qualitatively explain this phenomenon, but that would also lead to quantitative predictions.

15.2 Description of a Model

Main idea behind the model. We have a computing device – be it a computer or a brain – that is involved in communication with other computing devices so that together, they can solve a certain important problem.

The main difficulty with communication is that we cannot just send the internal signals out. It does not work for humans: we sometimes do not even understand each other's gestures or words, we need to translate our knowledge from our internal knowledge-representation language to a more universal one. Similarly, computers cannot just send out signals describing 0s and 1s that serve as internal representations of the corresponding knowledge: even if the two computers use the same way of representing, e.g., arrays of real numbers, the actual representation includes the information on where exactly this arrays is stored in the computer memory – the information that is useless for the computer that receives this information.

So, in general, to communicate, computing devices need to translate their internal signals into a different, more universal communication language. For this translation, we

need a dictionary stored in the computing device.

In computing devices, usually, there are several levels of information storage. There is an operating memory where access to information is fast but the size of this memory is limited. There is usually a much larger second-tier memory that can store a much larger amount of information but where access takes much longer. There are usually several more layers, but in this chapter, for simplicity, we will simply assume that we have two memory layers.

Details. Let a denote the overall computational ability related to the top (fastest-to-access) memory level. Some part of this level memory is taken by the most frequent "words" in the dictionary – so that translation of these words and thus, sending a message would go faster. Let a_0 denote the part of this level memory that is focused on this translation; then, we have $a - a_0$ ability remaining for general computations.

Let us denote by t_0 the part of the memory that is needed, on average, to store a translation of one word. Then, in the part a_0, we can store the translations of $w \stackrel{\text{def}}{=} \dfrac{a_0}{t_0}$ words.

Let us assume that we need:

- to perform some fast computations – whose overall running time will be denoted by C – and
- to send several (M) messages (of average length of ℓ words per message); this means that overall, in addition to computations, we need to communicate $W = M \cdot \ell$ words.

Let d be the size of the dictionary, i.e., the overall number of words that can be used for communication.

In this arrangement, what is the best division of top layer memory a into parts a_0 and $a - a_0$ under which both computation and communication tasks will be performed as fast as possible?

Zipf's law. In our analysis, we will rely on the known law that describes how frequently different words appear in a message. According to this law – known as Zipf's law – if we

sort all the words from a dictionary in the decreasing order of their frequency, then the frequency f_i with which the i-word appears is equal to $f_i \approx \frac{c}{i}$, for some constant c; see, e.g., [79].

The constant c can be determined from the condition that the sum of all the frequencies $f_1, \ldots, f_d$ should be equal to 1. Thus, we get

$$\frac{c}{1} + \frac{c}{2} + \ldots + \frac{c}{d} = 1,$$

i.e., equivalently,

$$c \cdot \left(\frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{d} \right) = 1.$$

The sum in parentheses is an integral sum for the integral

$$\int_1^d \frac{1}{x}\, dx = \ln(x)|_1^d = \ln(d) - \ln(1) = \ln(d),$$

thus

$$\frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{d} \approx \ln(d),$$

hence $c \cdot \ln(d) = 1$, so $c = \frac{1}{c} = \frac{1}{\ln(d)}$ and

$$f_i = \frac{1}{\ln(d)} \cdot \frac{1}{i}.$$

Towards formulas for computation and communication times. We have $a - a_0$ elementary computational devices to perform the overall amount C of needed computations. So, if we distribute these computation tasks between these $a - a_0$ devices, then we need the time

$$\frac{C}{a - a_0}$$

to perform all these computations.

Let us now estimate the amount of computations needed to send all M needed messages. In the fast memory layer, we can store w words. To speed up computations, it is reasonable to store, in the fast memory, translations to w most frequent words. If a message contains

other words, we need to spend some time either computing its translation, or, alternatively, bringing this translation from the slower memory layer. Let us denote the average time needed to translate a not-stored-in-fast-memory word by t.

Among all $W = M \cdot \ell$ words that we need to communicate, we need the translate for all the words except for the w most frequent ones, i.e., for all the words whose frequencies are $f_{w+1}, \ldots, f_d$. The overall frequency f of all such words can be obtained by adding up all these frequencies; so, we get

$$f = f_{w+1} + \ldots + f_d = \frac{c}{w+1} + \ldots + \frac{c}{d} = c \cdot \left(\frac{1}{w+1} + \ldots + \frac{1}{d} \right).$$

The sum in the last expression is also an integral sum, this time for the integral

$$\int_{w+1}^{d} \frac{1}{x}\, dx = \ln(x)|_{w+1}^{d} \approx \ln(d) - \ln(w).$$

Thus, the frequency f is approximately equal to

$$f = c \cdot (ln(d) - \ln(w)) = \frac{\ln(d) - \ln(w)}{\ln(d)}.$$

Among all W words, we thus need to spend time on $f \cdot W$ words. Translating each word requires time t, so overall, we need to spend time $f \cdot W \cdot t$ on this translation.

Substituting the above expression for f and the formula $W = M \cdot w_0$ into this formula, we conclude that the overall time for sending M messages is equal to

$$\frac{\ln(d) - \ln(w)}{\ln(d)} \cdot M \cdot w_0 \cdot t,$$

i.e., taking into account that $w = a_0/t_0$ and thus, $\ln(w) = \ln(a_0) - \ln(t_0)$, we get

$$\frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t.$$

By adding the computation and communication time, we get the following formula for the overall time.

Resulting formula for overall computation and communication time. The overall time T needed for computation and communication is equal to

$$\frac{C}{a - a_0} + \frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t. \tag{15.1}$$

15.3 Analysis of the Model: What Is the Optimal Tradeoff Between Computation and Communication

Main idea. The desired tradeoff is described by the parameter a_0. We want to find the value of this parameter for which the overall time T needed to perform all the tasks (including both computation and communication) is the smallest possible. In other words, the expression (15.1) for this time T is our objective function.

Towards an explicit expression for the optimal value a_0. To find the optimal value a_0, let us differentiate the objective function (15.1) with respect to a_0 and equate the derivative to 0. As a result, we get the following formula:

$$\frac{C}{(a-a_0)^2} - \frac{M \cdot w_0 \cdot t}{\ln(d)} \cdot \frac{1}{a_0} = 0.$$

Multiplying both sides of this equality by $(a-a_0)^2 \cdot a_0$, we get a quadratic equation:

$$C \cdot a_0 - \frac{M \cdot w_0 \cdot t}{\ln(d)} \cdot (a-a_0)^2 = 0.$$

Dividing both sides by the coefficient at $(a-a_0)^2$ and changing the sign of both sides, we get

$$(a-a_0)^2 - k \cdot a_0 = a_0^2 - (k-2) \cdot a \cdot a_0 + a^2 = 0,$$

where we denoted

$$k \stackrel{\text{def}}{=} \frac{C \cdot \ln(d)}{M \cdot w_0 \cdot t}.$$

Dividing both sides by a^2, we get the following quadratic equation to the fraction $r_0 \stackrel{\text{def}}{=} \dfrac{a_0}{a}$ of the top-level memory allocated for communications:

$$r_0^2 - (k-2) \cdot r_0 + 1 = 0.$$

The solution of this quadratic equation is

$$r_0 = \frac{k-2}{2} \pm \sqrt{\left(\frac{k-2}{2}\right)^2 - 1},$$

and $a_0 = a \cdot r_0$.

Analysis of the problem. When there are practically no communications, i.e., when the number of messages M is very small, the second term in the expression (15.1) for the objective function is negligible, so the objective function is approximately equal to its first term:

$$T \approx \frac{C}{a - a_0}.$$

This expression is the smallest when the difference $a - a_0$ is the largest, i.e., when the value a_0 is the smallest possible – and the smallest possible value of a_0 is 0.

Thus, in situations when we do not need to perform many communications, it makes sense not to allocate any top-level memory for communications, and use it all (or almost all) for computations.

On the other hand, if the number of messages is large, then, vice versa, we can ignore the first term in the expression (15.1) for the objective function and conclude that the objective function is approximately equal to its second term:

$$T \approx \frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t.$$

In this case, the larger a_0, the larger is $\ln(a_0)$ and thus, the smaller is the above expression. So, for this expression to be as small as possible, we need to select the value a_0 which is as large as possible. The largest possible value of the communication-related portion a_0 of the top-level memory is the whole amount a of this memory: $a_0 = a$.

Thus, in situations when we need to perform a large number of communications, it makes sense to allocate practically all top-level memory for communications, and leave only the bare minimum for computations.

These are the two extreme cases, but they show that the more communications we need, the larger portion of the top-level memory should be allocated for communication purposes (and the above explicit formula for the optimal value of a_0 confirms this conclusion).

This is exactly what we observe, both in chimps and in computer networks, in terms of a tradeoff between communication and computation. Thus, our simple model indeed

captures – at least on the qualitative level – the cognitive tradeoff phenomenon.

A recent study of chimpanzees has shown that on the individual basis, they are, surprisingly, much better than humans in simple tasks requiring intelligence and memory. A usual explanation – called cognitive tradeoff – is that a human brain has sacrificed some of its data processing (computation) abilities in favor of enhancing the ability to communicate; as a result, while individual humans may not be as smart as possible, jointly, we can solve complex problems. A similar cognitive tradeoff phenomenon can be observed in computer clusters: the most efficient computer clusters are not formed from the fastest, most efficient computers, they are formed from not-so-fast computers which are, however, better in their communication abilities than the fastest ones. In this chapter, we propose a simple model that explains the cognitive tradeoff phenomenon.

15.4 Formulation of the Problem

Interesting empirical phenomenon. A recent study of chimpanzees [27, 80, 82] showed, somewhat surprisingly, that on the individual basis, they are much better than human in many tasks requiring intelligence. For example, they can remember more objects in images, and in conflict situations their behavior is much closer to the optimal behavior (as recommended by game theory) than the behavior of humans.

Cognitive tradeoff: an explanation for this phenomenon. A current explanation for this phenomenon is based on what is called *cognitive tradeoff*: humans have better communication abilities, and so, human brain has to sacrifice some individual intellectual abilities to leave space for parts needed for efficient communication.

The need for such a tradeoff is not limited to humans. A similar tradeoff phenomenon can be observed not only in humans, but in computers as well. The world's fastest computations are performed on so-called high performance computers. Each of them is, in effect, a large number of processors constantly communicating with each other.

In principle, there exist processors which are very fast and efficient, but modern supercomputers are not formed from these processors: they are formed from simpler processors – similar to the ones we use in not-very-expensive home computers. One of the reasons for this choice is that these simple processors communicate well, as opposed to more efficient processors; these more efficient processors individually perform better but which take much longer time to communicate (another reason is that simple processors are usually much cheaper, which allows the designers to combine many more such processors within the same budget).

The ubiquity of cognitive tradeoff motivates the desired to have a universal quantitative model. The fact that cognitive tradeoff occurs in many situations, from human to computer communications, shows that there must be a simple quantitative explanation for this phenomenon.

In this chapter, we provide a simple quantitative model that explains the main ideas behind this phenomenon. We hope that this simple model can be used as a basis for more complex – and more realistic – models that would not only qualitatively explain this phenomenon, but that would also lead to quantitative predictions.

15.5 Description of a Model

Main idea behind the model. We have a computing device – be it a computer or a brain – that is involved in communication with other computing devices so that together, they can solve a certain important problem.

The main difficulty with communication is that we cannot just send the internal signals out. It does not work for humans: we sometimes do not even understand each other's gestures or words, we need to translate our knowledge from our internal knowledge-representation language to a more universal one. Similarly, computers cannot just send out signals describing 0s and 1s that serve as internal representations of the corresponding knowledge: even if the two computers use the same way of representing, e.g., arrays of real

numbers, the actual representation includes the information on where exactly this arrays is stored in the computer memory – the information that is useless for the computer that receives this information.

So, in general, to communicate, computing devices need to translate their internal signals into a different, more universal communication language. For this translation, we need a dictionary stored in the computing device.

In computing devices, usually, there are several levels of information storage. There is an operating memory where access to information is fast but the size of this memory is limited. There is usually a much larger second-tier memory that can store a much larger amount of information but where access takes much longer. There are usually several more layers, but in this chapter, for simplicity, we will simply assume that we have two memory layers.

Details. Let a denote the overall computational ability related to the top (fastest-to-access) memory level. Some part of this level memory is taken by the most frequent "words" in the dictionary – so that translation of these words and thus, sending a message would go faster. Let a_0 denote the part of this level memory that is focused on this translation; then, we have $a - a_0$ ability remaining for general computations.

Let us denote by t_0 the part of the memory that is needed, on average, to store a translation of one word. Then, in the part a_0, we can store the translations of $w \stackrel{\text{def}}{=} \dfrac{a_0}{t_0}$ words.

Let us assume that we need:

- to perform some fast computations – whose overall running time will be denoted by C – and
- to send several (M) messages (of average length of ℓ words per message); this means that overall, in addition to computations, we need to communicate $W = M \cdot \ell$ words.

Let d be the size of the dictionary, i.e., the overall number of words that can be used for communication.

In this arrangement, what is the best division of top layer memory a into parts a_0 and $a - a_0$ under which both computation and communication tasks will be performed as fast as possible?

Zipf's law. In our analysis, we will rely on the known law that describes how frequently different words appear in a message. According to this law – known as Zipf's law – if we sort all the words from a dictionary in the decreasing order of their frequency, then the frequency f_i with which the i-word appears is equal to $f_i \approx \frac{c}{i}$, for some constant c; see, e.g., [79].

The constant c can be determined from the condition that the sum of all the frequencies $f_1, \ldots, f_d$ should be equal to 1. Thus, we get

$$\frac{c}{1} + \frac{c}{2} + \ldots + \frac{c}{d} = 1,$$

i.e., equivalently,

$$c \cdot \left(\frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{d} \right) = 1.$$

The sum in parentheses is an integral sum for the integral

$$\int_1^d \frac{1}{x}\, dx = \ln(x)|_1^d = \ln(d) - \ln(1) = \ln(d),$$

thus

$$\frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{d} \approx \ln(d),$$

hence $c \cdot \ln(d) = 1$, so $c = \frac{1}{c} = \frac{1}{\ln(d)}$ and

$$f_i = \frac{1}{\ln(d)} \cdot \frac{1}{i}.$$

Towards formulas for computation and communication times. We have $a - a_0$ elementary computational devices to perform the overall amount C of needed computations. So, if we distribute these computation tasks between these $a - a_0$ devices, then we need the time

$$\frac{C}{a - a_0}$$

to perform all these computations.

Let us now estimate the amount of computations needed to send all M needed messages. In the fast memory layer, we can store w words. To speed up computations, it is reasonable to store, in the fast memory, translations to w most frequent words. If a message contains other words, we need to spend some time either computing its translation, or, alternatively, bringing this translation from the slower memory layer. Let us denote the average time needed to translate a not-stored-in-fast-memory word by t.

Among all $W = M \cdot \ell$ words that we need to communicate, we need the translate for all the words except for the w most frequent ones, i.e., for all the words whose frequencies are $f_{w+1}, \ldots, f_d$. The overall frequency f of all such words can be obtained by adding up all these frequencies; so, we get

$$f = f_{w+1} + \ldots + f_d = \frac{c}{w+1} + \ldots + \frac{c}{d} = c \cdot \left(\frac{1}{w+1} + \ldots + \frac{1}{d} \right).$$

The sum in the last expression is also an integral sum, this time for the integral

$$\int_{w+1}^{d} \frac{1}{x}\, dx = \ln(x)|_{w+1}^{d} \approx \ln(d) - \ln(w).$$

Thus, the frequency f is approximately equal to

$$f = c \cdot (ln(d) - \ln(w)) = \frac{\ln(d) - \ln(w)}{\ln(d)}.$$

Among all W words, we thus need to spend time on $f \cdot W$ words. Translating each word requires time t, so overall, we need to spend time $f \cdot W \cdot t$ on this translation.

Substituting the above expression for f and the formula $W = M \cdot w_0$ into this formula, we conclude that the overall time for sending M messages is equal to

$$\frac{\ln(d) - \ln(w)}{\ln(d)} \cdot M \cdot w_0 \cdot t,$$

i.e., taking into account that $w = a_0/t_0$ and thus, $\ln(w) = \ln(a_0) - \ln(t_0)$, we get

$$\frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t.$$

By adding the computation and communication time, we get the following formula for the overall time.

Resulting formula for overall computation and communication time. The overall time T needed for computation and communication is equal to

$$\frac{C}{a - a_0} + \frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t. \tag{15.1}$$

15.6 Analysis of the Model: What Is the Optimal Tradeoff Between Computation and Communication

Main idea. The desired tradeoff is described by the parameter a_0. We want to find the value of this parameter for which the overall time T needed to perform all the tasks (including both computation and communication) is the smallest possible. In other words, the expression (15.1) for this time T is our objective function.

Towards an explicit expression for the optimal value a_0. To find the optimal value a_0, let us differentiate the objective function (15.1) with respect to a_0 and equate the derivative to 0. As a result, we get the following formula:

$$\frac{C}{(a - a_0)^2} - \frac{M \cdot w_0 \cdot t}{\ln(d)} \cdot \frac{1}{a_0} = 0.$$

Multiplying both sides of this equality by $(a - a_0)^2 \cdot a_0$, we get a quadratic equation:

$$C \cdot a_0 - \frac{M \cdot w_0 \cdot t}{\ln(d)} \cdot (a - a_0)^2 = 0.$$

Dividing both sides by the coefficient at $(a - a_0)^2$ and changing the sign of both sides, we get

$$(a - a_0)^2 - k \cdot a_0 = a_0^2 - (k - 2) \cdot a \cdot a_0 + a^2 = 0,$$

where we denoted

$$k \stackrel{\text{def}}{=} \frac{C \cdot \ln(d)}{M \cdot w_0 \cdot t}.$$

Dividing both sides by a^2, we get the following quadratic equation to the fraction $r_0 \stackrel{\text{def}}{=} \dfrac{a_0}{a}$ of the top-level memory allocated for communications:

$$r_0^2 - (k-2) \cdot r_0 + 1 = 0.$$

The solution of this quadratic equation is

$$r_0 = \frac{k-2}{2} \pm \sqrt{\left(\frac{k-2}{2}\right)^2 - 1},$$

and $a_0 = a \cdot r_0$.

Analysis of the problem. When there are practically no communications, i.e., when the number of messages M is very small, the second term in the expression (15.1) for the objective function is negligible, so the objective function is approximately equal to its first term:

$$T \approx \frac{C}{a - a_0}.$$

This expression is the smallest when the difference $a - a_0$ is the largest, i.e., when the value a_0 is the smallest possible – and the smallest possible value of a_0 is 0.

Thus, in situations when we do not need to perform many communications, it makes sense not to allocate any top-level memory for communications, and use it all (or almost all) for computations.

On the other hand, if the number of messages is large, then, vice versa, we can ignore the first term in the expression (15.1) for the objective function and conclude that the objective function is approximately equal to its second term:

$$T \approx \frac{\ln(d) + \ln(t_0) - \ln(a_0)}{\ln(d)} \cdot M \cdot w_0 \cdot t.$$

In this case, the larger a_0, the larger is $\ln(a_0)$ and thus, the smaller is the above expression. So, for this expression to be as small as possible, we need to select the value a_0 which is as large as possible. The largest possible value of the communication-related portion a_0 of the top-level memory is the whole amount a of this memory: $a_0 = a$.

Thus, in situations when we need to perform a large number of communications, it makes sense to allocate practically all top-level memory for communications, and leave only the bare minimum for computations.

These are the two extreme cases, but they show that the more communications we need, the larger portion of the top-level memory should be allocated for communication purposes (and the above explicit formula for the optimal value of a_0 confirms this conclusion).

This is exactly what we observe, both in chimps and in computer networks, in terms of a tradeoff between communication and computation. Thus, our simple model indeed captures – at least on the qualitative level – the cognitive tradeoff phenomenon.

Chapter 16

Users Do Not Always Follow Expert Recommendations: Analytical Technique Explains Empirical Data

Empirical studies show that users do not always follow expert recommendations. For example, when a medical doctor prescribes a medicine, only two third of the patients fill the prescription, and of this filling-prescription group, only half follow the doctor's instructions when taking the medicine. In this chapter, we show that a general systems approach – namely, abstracting from the specifics of this situation – helps explain these empirical observations. We also mention that systems approach can not only explains this problem, it can also help solve it – i.e., it can help increase the patients' adherence to the doctors' recommendations – and users' adherence to expert recommendations in general.

Comment. The results presented in this chapter first appeared in [16].

16.1 Formulation of the Problem

Empirical observation. A recent study [97, 119] shows that among all the patients who get a prescription from a medical doctor:

- approximately one third of the patients do not fill their prescription at all,
- one third of the patients fill the prescription, but do not exactly follow the doctor's instructions about dosage and times, and

- only the remaining one third of the patients does exactly what the doctor suggested.

Researchers tried to explain this somewhat unexpected observation; see, e.g., [50, 84, 119]. However, all these explanations have been mostly qualitative. So far, to the best of our knowledge, there has been no convincing quantitative explanation.

What we do in this chapter. In this chapter, we show that the above empirical observation can be quantitatively explained if we abstract away from medical and social details and reformulate this problem on the general systems level.

16.2 Our Explanation

Possible reactions to a doctor's recommendation. When a patient gets some recommendation from the doctor, the patient does not necessarily believe that these recommendations are correct – otherwise, he/she would obediently follow this recommendation. In general, there are three possible options:

- the first option is that the patient is confident that the doctor's recommendation is correct,
- the second option is that the patient is confident that the doctor's recommendation is not correct, and
- the third option is that the patient is not sure whether the doctor's recommendation is correct or not.

How these reactions affect the patient's behavior. If a patient is confident that the doctor's recommendation is correct, then this patient will obediently follow this recommendation. In particular, this patient will fill the doctor's prescription and take the corresponding medicine at the prescribed times and at the prescribed dosage.

If the patient is confident that the doctor's recommendation is not correct, this patient will not even bother to fill the doctor's prescription.

In the third case, when the patient is not sure whether the doctor's recommendation is correct, a reasonable idea – unless the patient goes to a second doctor for a second opinion – is to follow some part of the doctor's recommendation. In this case, the patient fills the prescription but follows only part of the doctor's instructions about taking this medicine, e.g.:

- takes only half-dose each time or
- takes it only once a day when the doctor recommended to take it twice a day, etc.

The resulting frequencies with which patients exhibit different reactions to the doctor's recommendation. The above analysis shows that, depending on the patient's reaction, we will observe exactly the three types of behavior that the above empirical study observed. Thus, the frequencies with which we observe three different behaviors are exactly the frequencies with which the patients show one of the above three reactions to a doctor's recommendation:

- approximately one third of the patients are confident that the doctor's recommendation is correct;
- about one third of the patients are confident that the doctor's recommendation is not correct, and
- the remaining one third of the patients are not sure whether the doctor's recommendation is correct or not.

Hence, to explain the above empirical observation, we need to explain the frequencies of different patient's reactions to the doctor's recommendation.

Frequencies explained. The vast majority of patients are not trained in medicine. Thus, they have no reasonable way to decide whether they should trust the doctor's recommendation or not. As a result, they face the problem of selecting one of the three alternatives

– i.e., one of the three possible reactions to a doctor's recommendation – without having any information that would help them select one of these three alternatives.

Such a situation of selection under complete uncertainty is well known. A usual way to solve it – known as *Laplace Indeterminacy Principle* – is to conclude that since we have no reason to assign different probabilities to different alternatives, a reasonable idea is to assign the exact same probability to all these alternatives; see, e.g., [56] and references therein.

In our case, there are three alternatives, so we assign the exact same probability p to each of these three alternatives. Since the probabilities of selecting one of the alternatives must add up to 1, we thus conclude that $p + p + p = 1$, i.e., that $p = 1/3$. So, we predict that the patient will select each of three alternative reactions – and thus, the corresponding behavior – in exactly 1/3 of the cases. This is exactly what the empirical study observed.

Thus, we have indeed explained the observed frequencies of different medicine adherence behavior.

Need to go beyond explanations. Explanations are nice, but it is desirable not just to explain the fact that patients do not follow the doctors' recommendation, but to come up with ways to increase their adherence to these recommendations. In this, systems approach can also help; see, e.g., [29]. Specifically, systems approach helps in developing tools that help the patients understand the reasoning behind the doctors' recommendations and thus, increase the percentage of patients who follow these recommendations. One of such approaches is described in the next chapter.

Chapter 17

Analytical Techniques for Making Recommendations More Acceptable to Users: Status Quo Bias Actually Helps Decision Makers to Take Nonlinearity into Account

One of the main motivations for designing computer models of complex systems is to come up with recommendations on how to best control these systems. Many complex real-life systems are so complicated that it is not computationally possible to use realistic nonlinear models to find the corresponding optimal control. Instead, researchers make recommendations based on simplified – e.g., linearized – models. The recommendations based on these simplified models are often not realistic but, interestingly, they can be made more realistic if we "tone them down" – i.e., consider predictions and recommendations which are close to the current status quo state. In this chapter, we analyze this situation from the viewpoint of general system analysis. This analysis explain the above empirical phenomenon – namely, we show that this "status quo bias" indeed helps decision makers to take nonlinearity into account.

Comment. The results presented in this chapter first appeared in [11].

17.1 Formulation of the Problem

Real-life problems. In his presentation [58] at the 2019 World Congress of the International Fuzzy Systems Association (IFSA), Professor Kacprzyk recalled his experience of optimizing large systems – like an economic region – at the International Institute for Applied Systems Analysis (IIASA) in Laxenburg, Austria.

Linearization is needed. Of course, many dependencies in complex real-life systems are non-linear. However, even with modern computers, optimizing a complex system under nonlinear constraints would require an unrealistic amount of computation time. As a result, in the optimization, most processes in a real-life system were approximated by linear models.

Comment. To be more precise, when analyzing the effect of one specific strategy on a system, we can afford to take non-linearity into account. However, when we need to solve an optimization problem of selecting the optimal strategy and/or optimal parameters of such a strategy, we have to use linearization.

Recommendations based on the linearized model were often not realistic. Not surprisingly, since the optimization process involved simplification of the actual system, recommendations based on the resulting simplified model were often not realistic and could not be directly implemented.

This was not just a subjective feeling: when the researchers tested, on the nonlinear model, the effect of a strategy selected based on linearization, the results were often not so good.

Status quo bias helped. One of the reasons that people listed for being reluctant to accept the center's recommendations was that these recommendations differed too much from what they expected. This phenomenon of unwillingness to follow recommendations if they are too far away from the status quo is known as the *status quo bias*; see, e.g., [59, 109].

Interestingly, when the center's researchers "toned down" their recommendations by making them closer to the status quo, the resulting recommendations led to much better results (e.g., as tested on the nonlinear models).

In other words, the toning down – corresponding to what we understand as the status quo bias – actually improved the decisions in comparison to simply using recommendations based on the simplified linear models. Thus, the status quo bias somehow takes into account non-linearity – and is, thus, not a deviation from an optimal decision making (as the word *bias* makes you think) but rather a reasonable way to come up with a better decision.

But why? The phenomenon described above seems mysterious. Why would getting closer to the status quo lead to a better solution?

In this chapter, we analyze this phenomenon from the general system approach. Our analysis allows us to explain why the status quo bias indeed helps to take some nonlinearity into account.

17.2 Analysis of the Problem and the Resulting Explanation

A general description of a system: a brief reminder. In general, the state of a system at each moment of time t can be described by listing the values $x_1(t), \ldots, x_n(t)$ of all the quantities $x_1, \ldots, x_n$ that characterize this system. Similarly, the change in the system can be described by differential equations that explain how the value of each of these quantities change with time:

$$\frac{dx_i(t)}{dt} = f_i(x_1(t), \ldots, x_n(t)), \quad i = 1, \ldots, n.$$

Here the expressions $f_i(x_1, \ldots, x_n)$ describe how the rate of change in each of the quantities depends on the state $x = (x_1, \ldots, x_n)$; in general, the expressions $f_i(x_1, \ldots, x_n)$ are non-linear.

In particular, in the simplest case when we use the value of only quantity x_1 to describe the state, we get the equation

$$\frac{dx_1(t)}{dt} = f_1(x_1(t)).$$

What happens when we linearize. When we linearize the description of the system, we thus replace the nonlinear functions $f_i(x_1, \ldots, x_n)$ by their linear approximations

$$f_i(x_1, \ldots, x_n) \approx a_i + \sum_{j=1}^{n} a_{ij} \cdot x_j.$$

In particular, in the case when we use only quantity x_1, we get the following approximate equation

$$\frac{dx_1}{dt} = a_1 + a_{11} \cdot x_1.$$

In this case, for the auxiliary variable $y_1 \stackrel{\text{def}}{=} x_1 + \dfrac{a_1}{a_1}$, we get

$$\frac{dy_1}{dt} = a_{11} \cdot y_1.$$

The solution to this simple differential equation is well-known: it is

$$y_1(t) = y_1(0) \cdot \exp(a_{11} \cdot t)$$

and thus,

$$x_1(t) = y_1(y) - \frac{a_1}{a_{11}} = y_1(0) \cdot \exp(a_{11} \cdot t) - \frac{a_1}{a_{11}}.$$

In situations when the value of x_1 (and thus, of y_1) decreases, we have $a_{11} < 0$. In such situations, the value y_1 decreases to 0. Such things happen. However, in situations when we want to describe the growth, i.e., when $a_{11} > 0$, we get an exponential growth.

Exponential growth may be a good approximation for some period of time, but eventually it starts growing too fast to be realistic. For example, in good times, economies grow – but we do not expect, e.g., production of meat to grow thousands times. Similarly, pollution grows, or, on a somewhat less negative side, populations grow, but we do not expect thousands-times increases predicted by the exponential models.

This phenomenon of models-growing-too-fast is not limited to the case when the system if described by only one variable. In general, a solution to a system of linear differential equations with constant coefficients is a linear combination of oscillatory terms and exponential terms – so, if we describe growth, the models will make it unrealistically exponential.

How to make conclusions more realistic. Linear models are not realistic – the deviations from the current values and what these models predict become too large to be realistic.

Thus, a natural way to make the models more realistic is to take this phenomenon into account – i.e., instead of the results of the linearized models, consider states which are closer to the original state

$$x(0) = (x_1(0), \ldots, x_n(0)).$$

Conclusion. This idea of considering the states which are closer to the original state than the model suggests is exactly what the status quo bias is about. Thus, indeed, the status quo bias helps take make models more realistic.

The unrealistic character of the linearized model's recommendation is caused by the fact that this model is only approximate – it ignores nonlinear terms. So, by making recommendations more realistic, the status quo bias, in effect, helps us to take nonlinearity into account.

Chapter 18

Analytical Techniques for Testing: Optimal Distribution of Testing Resources Between Different System Levels

When designing a system, we need to perform testing and checking on all levels of the system hierarchy, from the most general system level to the most detailed level. For each level, there are techniques for testing on this level. The problem is that our resources are limited, so we need to find the best way to allocate these resources, i.e., we need to decide how much efforts to use of each of the levels. In this chapter, we formulate this problem in precise terms, and provide a solution to the resulting optimization problem.

Comment. The results presented in this chapter first appeared in [8].

18.1 Formulation of the Problem

Need for system design. Sometimes, engineers and scientists concentrate on designing a specific device or a specific software. However, no device and no software works on its own, whatever we design will be a part of a system. For example, when we design a new industrial plant, we need to take into account how its functioning will affect the natural ecosystem, how the increased transportation will affect the city infrastructure, how the new people brought to this plant will change the demographic system, etc. Similarly, in

science, when we design a new radiotelescope (or even software to process signals from the radiotelescope), we need to take into account that this telescope will be mostly used as a part of a system of radiotelescopes and other astronomical instruments to observe different celestial objects.

System design is hierarchical. To properly design a system, it is important to first have a clear general structure. After that, once it becomes clear what are the system components and how they supposed to interact, we can move to designing these individual components – taking into account the need for these components to efficiently work together. These components usually also are subsystems, so we need to come up with their own components, etc. At the end, once all the tasks have been clarified, we proceed to the most detailed level, where we design individual machines and instruments and write the corresponding software.

Of course, the above sequence is an idealized representation of the actual design process: sometimes, after we go to a more detailed level of design, we realize the need to make some changes in the previously decided higher-level design structure. However, most of the time, the system design follows the above hierarchical pattern.

Need for testing and checking. On each design level, we need to check for possible problems and flaws. Flaws can occur on different levels.

For example, on the highest general-system level, we may forget an important aspect of the system – e.g., when designing a plant, we may not think about its ecological impact – and as a result, once the design is done, it may have to be redone completely. To avoid such situations, it is important to check the design on each level before starting a more detailed design level.

Flaws may also occur on the very lowest most-detailed level: e.g., we can have a software that does not always provide the correct control for the plant.

Need to allocate testing resources between different levels. We need to perform testing and checking on different levels of the system hierarchy. However, our testing re-

sources are limited. It is therefore important to efficiently distributed the available resources between different levels; see, e.g., [21, 29, 34, 46, 52, 53, 116, 132, 133, 134].

What we do in this chapter. In this chapter, we describe the problem of allocating resources in precise terms, and we provide a solution to the resulting optimization problem.

18.2 Analysis of the Problem

The cost of errors on different levels. Errors can occur on all the levels:

- we can make an error on the highest level, by deciding on a faulty overall design;
- we can also make an error on the most detailed level, e.g., making an error when manufacturing one of the system's components.

An error on a higher level is very costly: if there was indeed an error in the overall design, we have to redo the overall design and thus, redo all the details – i.e., largely, start "from scratch". On the other hand, errors on the lower levels are not that costly: if we erred in designing one small component, then all we need to do is re-design this small component.

Let us number the levels from the most general one – which will be Level 1, via the next-detailed Level 2, then even-more-detailed Level 3, etc., all the way to the most detailed Level. Let us denote the overall number of levels by n. Then, the most detailed level is Level n.

In general, an error on each level i leads to the need of redoing several details on the next-detailed level $i+1$. Let us denote the average number of details that need to be redone by q. Then, an error on Level i necessitates redoing q details on the next-detailed Level $i+1$. Each of these re-doings requires redoing q details on the next level $i+2$; thus, an error on Level i requires re-doing q^2 details on Level $i+2$. Similarly, we conclude that it requires re-doing q^3 details on Level $i+3$, and, in general, q^k details on level $i+k$. In particular, for $k = n-i$, we conclude that an error on Level i requires redoing q^{n-i} details on the most detailed Level n.

Let c denote the average cost of redoing a single detail on the most-detailed Level n. Then, the overall cost of an error on Level i can be obtained by multiplying this per-error cost c by the total number of details q^{n-i} that need to be corrected, and is, thus, equal to $c \cdot q^{n-i}$.

The cost of discovering errors. How does the number N of remaining errors depend on the effort – i.e., equivalently, on the time t spent to find these errors. We would like to find a general formula $N(t)$ for describing this dependence.

It is important to take into account that there are different way to count errors. For example, when we talk about software errors, we can count the number of modules that do not perform as we intended, we can count the number of lines of code where we made a mistake, or we can count the number of erroneous operations on each line of code. All three (and other) ways of counting errors make sense – but they differ by a factor. For example, to go from the number of erroneous moduli to the number of erroneous lines of code, we need to multiply the number of erroneous moduli by the average number of erroneous lines of code in an erroneous modulus. Thus, if we change the way we count errors, we go from the original number $N(t)$ to the new number $C \cdot N(t)$, where C is the corresponding factor.

Both the original function $N(t)$ and the new function $C \cdot N(t)$ make sense. Thus, instead of a single function $N(t)$ for describing how the number of remaining errors depends on time t, we should consider the whole family of functions $\{C \cdot N(t)\}_C$ corresponding to all possible value $C > 0$.

The time t is the time from the moment when we started testing. This may sound well-defined, but in practice, it changes from one person to another. Some programmers try to run the very first version of the program that they wrote – and thus, start debugging the code right away. Other programmers first try some on-paper tests and only start running when they are reasonably sure that they eliminate the most obvious bugs. While the results of both programmers may be similar, the starting time for measuring t is different for the second programmer: what happened for the first programmer at time t, for the second programmer, happens at time $t - t_0$, where t_0 is the time the second programmer

spent analyzing his/her code before running it. This value t_0 may be different for different programmers. It is therefore reasonable to require that the approximating family $\{C \cdot N(t)\}_C$ should not change if we simply change the way we measure the time, i.e., if we go from t to $t - t_0$.

In other words, the family $\{C \cdot N(t - t_0)\}_C$ corresponding to the shifted time $t - t_0$ should coincide with the original family $\{C \cdot N(t)\}_C$. This means, in particular, that for every t_0, the function $N(t - t_0)$ from the shifted family should belong to the original family, i.e., it should have a form

$$N(t - t_0) = C(t_0) \cdot N(t), \tag{18.1}$$

for some value $C(t_0)$ depending on t_0.

The function $N(t)$ describing the number of remaining errors after time t is (non-strictly) decreasing: when $t < t'$, then we should have $N(t) \geq N(t')$. Thus, it is measurable, and therefore, the function $C(t_0) = N(t - t_0)/N(t)$ is also measurable, as the ratio of two measurable functions. It is known (see, e.g., [18]) that for measurable functions, the only solutions to equation (18.1) have the form $N(t) = N_0 \cdot \exp(-a \cdot t)$ for some coefficients N_0 and a; see [71].

Now, we are ready to formulate the problem in precise terms.

18.3 Formulation of the Problem in Precise Terms

Towards the formulation. We want to divide the overall people-time T that we have allocated for testing into times $t_1, \ldots, t_n$ allocated to testing on different levels:

$$t_1 + \ldots + t_n = T. \tag{18.2}$$

According to the above formulas, for each level i, after the testing, we will have $N_0 \cdot \exp(-a \cdot t_i)$ errors. The cost of each error on this level is $c \cdot q^{n-i}$, so the overall cost of all these errors is $c \cdot q^{n-i} \cdot N_0 \cdot \exp(-a \cdot t_i)$.

The overall cost E coming from all the remaining errors can be computed by adding the costs corresponding to different levels:

$$E = \sum_{i=1}^{n} c \cdot q^{n-i} \cdot N_0 \cdot \exp(-a \cdot t_i). \tag{18.3}$$

Resulting formulation. We want to select the times $t_1, \ldots, t_n$ – under the constraint (18.1) – so as to minimize the overall cost E of all the errors.

In other words, we want to minimize the expression (18.2) under the constraint (18.1).

18.4 Solving the Resulting Optimization Problem

Solving the problem. A usual way to solve a constraint optimization problem is to use Lagrange multipliers, i.e., to reduce the original problem of minimizing a function $f(x)$ under a constraint $g(x) = 0$ to the unconstrained problem of minimizing an expression $f(x) + \lambda \cdot g(x)$, where the parameter λ (known as *Lagrange multiplier*) has to be determined from the condition $g(x) = 0$.

In our case, the constraint has the form

$$\sum_{i=1}^{n} t_i - T = 0,$$

so the corresponding unconstrained optimization problem means minimizing the expression

$$\sum_{i=1}^{n} c \cdot q^{n-i} \cdot N_0 \cdot \exp(-a \cdot t_i) + \lambda \cdot \left(\sum_{i=1}^{n} t_i - T \right).$$

To find the minimum of this expression, we differentiate it with respect to each unknown t_i and equate the resulting (partial) derivative to 0. As a result, we get the following formula:

$$c \cdot q^{n-i} \cdot N_0 \cdot (-a) \cdot \exp(-a \cdot t_i) + \lambda = 0,$$

i.e.,

$$\exp(-a \cdot t_i) = \frac{\lambda}{a \cdot c \cdot N_0} \cdot q^{n-i}.$$

Taking logarithms of both sides and dividing the result by $-a$, we get

$$t_i = (n-i) \cdot \frac{|\ln(q)|}{a} + c_1,$$

where we denoted

$$c_1 \stackrel{\text{def}}{=} -\frac{1}{a} \cdot \ln\left(\frac{\lambda}{a \cdot c \cdot N_0}\right).$$

Combining terms not depending on i into a single expression, we get

$$t_i = c_2 - i \cdot \frac{|\ln(q)|}{a}, \tag{18.4}$$

where

$$c_2 \stackrel{\text{def}}{=} c_1 + n \cdot \frac{|\ln(q)|}{a}.$$

In line with the main idea of the Lagrange multiplier technique, to find the value c_2, we substitute the expression (18.4) into the constraint (18.1). As a result, we get

$$T = \sum_{i=1}^{n} t_i = n \cdot c_2 - \left(\sum_{i=1}^{n} i\right) \cdot \frac{|\ln(q)|}{a}.$$

Here,

$$\sum_{i=1}^{n} i = 1 + 2 + \ldots + n = \frac{n \cdot (n+1)}{2},$$

thus

$$T = n \cdot c_2 - \frac{n \cdot (n+1)}{2} \cdot \frac{|\ln(q)|}{a},$$

and so,

$$c_2 = \frac{T}{n} + \frac{n+1}{2} \cdot \frac{|\ln(q)|}{a}.$$

Thus, we arrive at the following formula.

Resulting solution. In situation where an error on the next level costs q times less than the error on the previous level, and the number of detected errors decreases with detection time as $\exp(-a \cdot t)$, the optimal allocation of the overall testing time T into times $t_1, \ldots, t_n$ allocated to each level has the form

$$t_i = \left(\frac{T}{n} + \frac{n+1}{2} \cdot \frac{|\ln(q)|}{a}\right) - i \cdot \frac{|\ln(q)|}{a}. \tag{18.5}$$

Discussion. In other words, the allocated time linearly decreases as we go from the most abstract level to the more and more detailed levels. The fact that we allocate most of the testing time to the highest level makes perfect sense: as we have mentioned, errors on this level are the costliest ones. That the decrease should be linear follows from the specific formulas of our model.

References

[1] G. Acosta, E. Freudenthal, E. Smith, and V. Kreinovich, "Why filtering out higher harmonics makes it easier to carry a tune", *Applied Mathematical Sciences*, 2019, Vol. 13, No. 16, pp. 775–779.

[2] G. Acosta, M. Hernandez, N. Villanueva-Rosales, E. Smith, and V. Kreinovich, "Why matrix factorization works well in recommender systems: a systems-based explanation", *Journal of Uncertain Systems*, 2019, Vol. 13, No. 3, pp. 164–167.

[3] G. Acosta, E. Smith, O. Kosheleva, and V. Kreinovich, "Epicycles are almost as good as trigonometric series: general system-based analysis", *Applied Mathematical Sciences*, 2019, Vol. 13, No. 16, pp. 769–773.

[4] G. Acosta, E. Smith, and V. Kreinovich, "80/20 Rule Partially Explains 7 Plus Minus 2 Law: General System-Based Analysis", *Internatiinal Mathematical Forum*, 2019, Vol. 14, No. 5, pp. 205–208.

[5] G. Acosta, E. Smith, and V. Kreinovich, *Confirmation Bias in Systems Engineering: A Pedagogical Example*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-89, 2019.

[6] G. Acosta, E. Smith, and V. Kreinovich, *Dunning-Kruger Effect: A Simple System-Based Explanation*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-68, 2019.

[7] G. Acosta, E. Smith, and V. Kreinovich, *How to Explain That Changes in Elderlies Depression Level Are Uniformly Distributed*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-90, 2019.

[8] G. Acosta, E. Smith, and V. Kreinovich, "Optimal distribution of testing resources between different system levels", *Journal of Uncertain Systems*, 2019, Vol. 13, No. 2, pp. 84–88.

[9] G. Acosta, E. Smith, and V. Kreinovich, *A Natural Explanation for the Minimum Entropy Production Principles*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-90, 2019.

[10] G. Acosta, E. Smith, and V. Kreinovich, "A simple quantitative model of cognitive tradeoff phenomenon", *Journal of Uncertain Systems*, 2019, Vol. 13, No. 2, pp. 89–93.

[11] G. Acosta, E. Smith, and V. Kreinovich, *Status Quo Bias Actually Helps Decision Makers to Take Nonlinearity into Account: An Explanation*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-91, 2019.

[12] G. Acosta, E. Smith, and V. Kreinovich, "Unexpected empirical dependence of calf gender on insemination time: system-based explanation", *Applied Mathematical Sciences*, 2019, Vol. 13, No. 14, pp. 681–684.

[13] G. Acosta, E. Smith, and V. Kreinovich, "Why Area Under the Curve in Hypothesis Testing?", *International Mathematical Forum*, 2019, Vol. 15, No. 4, pp. 209–213.

[14] G. Acosta, E. Smith, and V. Kreinovich, *Why high-level attention constantly oscillates: system-based explanation*, University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-19-17, 2019.

[15] G. Acosta, E. Smith, and V. Kreinovich, "Why IQ test scores are slightly decreasing: possible system-based explanation for the reversed Flynn effect", *Journal of Uncertain Systems*, 2019, Vol. 13, No. 3, pp. 176–179.

[16] Griselda Acosta, Eric Smith, and Vladik Kreinovich, "Why Patients Do Not Always Follow Doctor's Advice: Systems Approach Explains Empirical Observation", *Mathematical Structures and Modeling*, 2019, Vol. 51, pp. 105–108.

[17] G. Acosta, E. Smith, and V. Kreinovich, "Why pink noise is best for enhancing sleep and memory: system-based explanation", *Applied Mathematical Sciences*, 2019, Vol. 19, No. 14, pp. 677–680.

[18] J. Aczél and J. Dhombres, *Functional Equations in Several Variables*, Cambridge University Press, 2008.

[19] A. A. Al-Shahomee, S. E. Abdalla, and R. Lynn, "An increase of intelligence in Libya from 2008 to 2017", *Personality and Individual Differences*, 2008, Vol. 130, pp. 147–149.

[20] D. Anderson, "Treating depression in old age: the reasons to be positive", *Age and Ageing*, 2001, Vol. 30, pp. 13–17.

[21] T. Bahill, R. Botta, and J. Daniels, "The Zachman framework populated with baseball models", *Journal of Enterprise Architecture*, 2006, Vol. 2, No. 4, pp. 50–68.

[22] D. R. Baker, P. J. Eslinger, M. Benavides, E. Peters, N. F. Dieckmann, and J. Leon, "The cognitive impact of the education revolution: A possible cause of the Flynn Effect on population IQ", *Intelligence*, 2015, Vol. 49, pp. 144–158.

[23] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, New York, 2006.

[24] J. Bohlman and A. T. Bahill, "Examples of mental mistakes made by systems engineers while creating tradeoff studies", *Studies in Engineering and Technology*, 2014, Vol. 1, No. 1, pp. 22–43.

[25] B. Bratsberg and O. Rogeberg, "Flynn effect and its reversal are both environmentally caused", *Proceedings of the National Academy of Sciences*, 2018, Vol. 115, No. 26, pp. 6674–6678.

[26] D. Calhoun, *The Intelligence of a People*, Princeton University Press, Princeton, New Jersey, 1973.

[27] J. Cohen, *Almost Chimpanzee: Searching for What Makes Us Human, in Rainforests, Labs, Sanctuaries, and Zoos*, Times Books, Henry Holt and Co., New York, 2010.

[28] R. Colom, J. M. Lluis-Font, and A. Andrés-Pueyo, "The generational intelligence gains are caused by decreasing variance in the lower half of the distribution: supporting evidence for the nutrition hypothesis", *Intelligence*, 2005, Vol. 33, No. 1, pp. 83–91.

[29] B. Dantu and E. Smith, "Diagnostic modeling for medical decision making", In: T. Doolen and E. Van Aken (eds.), *Proceedings of the 2011 Industrial Engineering Research Conference*, Reno, Nevada, May 21–25, 2011.

[30] O. Dubois, "Counting the number of solutions for instances of satistiability", *Theoretical Computer Science*, 1991, Vol. 81, pp. 49–64.

[31] O. Dubois and J. Carlier, "Probabilistic approach to the satisfiability problem", *Theoretical Computer Science*, 1991, Vol. 81, pp. 65–85.

[32] D. Dunning, "The Dunning-Krueger effect: On being ignorant of one's own ignorance.", In: J. M. Olson and M. P. Zanna (Eds.), *Advances in Experimental Social Psychology*, Vol. 44, Academic Press, San Diego, California, 2011, pp. 247–296.

[33] E. Dutton and R. Lynn, "A negative Flynn effect in Finland, 1997–2009", *Intelligence*, 2013, Vol. 41, No. 6, pp. 817–820.

[34] S. Ferreira, R. Valerdi, N. Medvidovic, J. Hess, I. Deonandan, T. Mikaelian, and G. Shull, "Unmanned and autonomous systems of systems test and evaluation: challenges and opportunities", *Proceedings of the 5th IEEE Systems Conference*, San Diego, California, April 5–8, 2010.

[35] R. Feynman, R. Leighton, and M. Sands, *The Feynman Lectures on Physics*, Addison Wesley, Boston, Massachusetts, 2005.

[36] D. J. Field, "Relations between the statistics of human images and the response properties of cortical cells", *Journal of the Optical Society of America, Series A*, 1987, Vol. 4, No. 12, pp. 2379–2394.

[37] A. Finkelstein, O. Kosheleva, and V. Kreinovich, "Astrogeometry: towards mathematical foundations", *International Journal of Theoretical Physics*, 1997, Vol. 36, No. 4, pp. 1009–1020.

[38] A. Finkelstein, O. Kosheleva, and V. Kreinovich, "Astrogeometry: geometry explains shapes of celestial bodies", *Geombinatorics*, 1997, Vol. VI, No. 4, pp. 125–139.

[39] R. B. Fletcher and J. Hattie, *Intelligence and Intelligence Testing*, Taylor & Francis, Milton Park, UK, 2011.

[40] J. R. Flynn, "Requiem for nutrition as the cause of IQ gains: Raven's gains in Britain 1938–2008", *Economics and Human Biology*, 2009, Vol. 7, No. 1, pp. 18–27.

[41] J. R. Flynn, *What Is Intelligence: Beyond the Flynn Effect*, Cambridge University Press, Cambridge, UK, 2009.

[42] E. Freudenthal, B. Usevitch, and E. Hanson, *Apparatus and Method to Facilitate Singing Intended Notes*, US Patent No. 9773426, 2017.

[43] T. L. Friedman, *Thank You For Being Late: An Optimists' Guide to Thriving in the Age of Accelerations*, Picador, New York, 2016.

[44] P. Glansdorff and I. Prigogine, *Thermodynamic Theory of Structure, Stability and Fluctuations*, Wiley-Interscience, London, 1971.

[45] P. Gomperts, A. Kovner, J. Lerner, and D. Scharfstein, *Skill vs. Luck in Entrepreneurship and Venture Capital: Evidence from Serial Entrepreneurs*, US National Bureau of Economic Research (NBER), 2006, Working Paper 12592, available at http://www.nber.org/papers/w12592.

[46] R. Gona and E. Smith, "Healthcare enterprise quality assessment with a Zachman-Bayesian framework", In: T. Doolen and E. Van Aken (eds.), *Proceedings of the 2011 Industrial Engineering Research Conference*, Reno, Nevada, May 21–25, 2011.

[47] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Leaning*, MIT Press, Cambridge, Massachusetts, 2016.

[48] W. T. Grandy Jr., *Entropy and the Time Evolution of Macroscopic Systems*, Oxford University Press, Oxford, 2008.

[49] R. Gray, "British teenagers have lower IQs than their counterparts did 30 years ago", *The Telegraph*, London, February 7, 2009.

[50] J. Groopman, *How Doctors Think*, Houghton Mifflin Co., New York, 2007.

[51] R. F. Helfrich, I. C. Fiebelkom, S. Z. Szczepanski, J. J. Lin, J. Parvizi, R. T. Knight, and S. Kastner, "Neural mechanisms of sustained attention are rhythmic", *Neuron*, 2018, Vol. 99, pp. 854–865.

[52] J. Hess, G. Agarwal, K. Cowart, I. Deonandan, C. R. Kenley, T. Mikaelian, and R. Valerdi, "Normative and descriptive models for test & evaluation of unmanned and autonomous systems of systems", *Proceedings of the 20th INCOSE Symposium*, Chicago, Illinois, July 12–15, 2010, Vol. 1, pp. 644–654.

[53] J. Hess and R. Valerdi, "Test and evaluation of a SoS using a prescriptive and adaptive testing framework", *Proceedings of the 5th IEEE International Conference on Systems of Systems Engineering*, Loughborough, UK, June 22–24, 2010.

[54] P. Jaekl, "Pay attention", *The Scientist*, 2018, No. 12, pp. 15–17.

[55] E. T. Jaynes, "The minimum entropy production principle", *Annual Review of Physical Chemistry*, 1980, Vol. 31, pp. 579–601.

[56] E. T. Jaynes and G. L. Bretthorst, *Probability Theory: The Logic of Science*, Cambridge University Press, Cambridge, UK, 2003.

[57] B. D. Josephson, "A trans-human source of musics?", In: P. Pylkkännen and P. Pylkkö (eds.), *New Directions in Cognitive Science*, Finnish Artifical Intelligence Society, Helsinki, Finland, 1995, pp. 280–285.

[58] J. Kacprzyk, "Cognitive biases in choice and decision making: a potential role of fuzzy logic", *Proceedings of the Joint World Congress of International Fuzzy Systems Association and Annual Conference of the North American Fuzzy Information Processing Society IFSA/NAFIPS'2019*, Lafayette, Louisiana, June 18–21, 2019.

[59] D. Kahneman, J. L. Knetsch, and R. H. Thaler, "Anomalies: the endowment effect, loss aversion, and status quo bias", *Journal of Economic Perspectives*, 1991, Vol. 5, No. 1, pp. 193–206.

[60] E. R. Kandel, J. H. Schwartz, T. M. Jessell, S. A. Siegelbaum, and A. J. Hudspeth (eds.), *Principles of Neural Science*, McGraw Hill, New York, 2013.

[61] D. Kahneman, *Thinking, Fast and Slow*, Farrar, Straus, and Giroux, New York, 2011.

[62] M. J. Klein and P. H. E. Meijer, "Principle of minimum entropy production", *Physical Review*, 1954, Vol. 96, pp. 250–255.

[63] R. Koch, *The 80/20 Principle and 92 Other Powerful Laws of Nature: The Science of Success*, Nicholas Brealey, London, UK, 2014.

[64] R. M. Koriat, S. Lichtenstein, and B. Fischhoff, "Reasons for confidence", *Journal of Experimental Psycholology: Human Learning and Memory*, 1980, Vol. 6, pp. 107–118.

[65] O. Kosheleva and V. Kreinovich, "For fuzzy logic, Occam's principle explains the ubiquity of the golden ratio and of the 80-20 rule", *Journal of Innovative Technology and Education*, 2017, Vol. 4, No. 1, pp. 13–18.

[66] V. Kreinovich, "S. Maslov's Iterative Method: 15 Years Later (Freedom of Choice, Neural Networks, Numerical Optimization, Uncertainty Reasoning, and Chemical Computing)", In: V. Kreinovich and G. Mints (eds.), *Problems of Reducing the Exhaustive Search*, American Mathematical Society, Providence, Rhode Island, 1997, pp. 175–189.

[67] V. Kreinovich and O. Fuentes, "High-Concentration Chemical Computing Techniques for Solving Hard-To-Solve Problems, and Their Relation to Numerical Optimization, Neural Computing, Reasoning Under Uncertainty, and Freedom Of Choice", In: E. Katz (ed.), *Molecular and Supramolecular Information Processing: From Molecular Switches to Logical Systems*, Wiley-VCH, Wienheim, Germany, 2012, pp. 209–235.

[68] V. Kreinovich, A. Lakeyev, J. Rohn, and P. Kahl, *Computational Complexity and Feasibility of Data Processing and Interval Computations*, Kluwer, Dordrecht, 1998.

[69] V. Kreinovich and G. Liu, "We live in the best of possible worlds: Leibniz's insight helps to derive equations of modern physics", In: R. Pisano, M. Fichant, P. Bussotti, and A. R. E. Oliveira (eds.), *The Dialogue between Sciences, Philosophy and Engineering. New Historical and Epistemological Insights, Homage to Gottfried W. Leibnitz 1646–1716*, College Publications, London, 2017, pp. 207–226.

[70] V. Kreinovich and S. Shukeilo, "A new probabilistic approach to the knapsack problem", *Proceedings of the Third USSR All-Union School on Discrete Optimization and Computers, Tashtagol, Russia, December 2–9, 1987*, Moscow, 1987, pp. 123–124 (in Russian).

[71] V. Kreinovich, Th. Swenson, and A. Elentukh, "Interval approach to testing software", *Interval Computations*, 1994, No. 2, pp. 90–109.

[72] J. Kruger and D. Dunning, "Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessment", *Journal of Personality and Social Psychology*, 1999, Vol. 6, pp. 1121–1134.

[73] J. L. Lawrence, *Celestial Calculations: A Gentle Introduction to Computational Astronomy*, MIT Press, Cambridge, Massachusetts, 2019.

[74] S. R. Laxpati and V. Goncharoff, *Practical Signal Processing and Its Applications*, World Scientific, Singapore, 2018.

[75] S. Li, Y. Ogura, and V. Kreinovich, *Limit Theorems and Applications of Set Valued and Fuzzy Valued Random Variables*, Kluwer Academic Publishers, Dordrecht, 2002.

[76] R. Livi and P. Politi, *Non-Equilibrium Statistical Physics: A Modern Perspective*, Cambridge University Press, Cambridge, UK, 2017.

[77] R. G. Lyons, *Understanding Digital Signal Processing*, Prentice Hall, Upper Saddle River, New Jersey, 2011.

[78] C. Maes and K. Netočný, "Minimum entropy production principle from a dynamical fluctuation law", *Journal of Mathematical Physics*, 2007, Vol. 48, Paper 053306.

[79] B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman, San Francisco, California, 1983.

[80] C. F. Martin, R. Bhui, P. Bassaerts, T. Matsuzawa, and C. Camerer, "Chimpanzee choice rates in competitive games match equilibrium game theory predictions", *Scientific Reports*, 2014, Vol. 4, Paper 5182.

[81] I. N. Martyushev, A. S. Nazarova, and V. D. Seleznev, "On the problem of the minimum entropy production in the nonequilibrium stationary state", *Journal of Physics A: Mathematical and Theoretical*, 2007, Vol. 40, No. 3, pp. 371–380.

[82] T. Matsuzawa, "Evolution of the brain and social behavior in chimpanzees", *Current Opinion in Neurobiology*, 2013, Vol. 23, pp. 443–449.

[83] G. A. Miller, "The magical number seven plus or minus two: some limits on our capacity for processing information", *Psychological Review*, 1956, Vol. 63, No. 2, pp. 81–97.

[84] A. Morabia, "Last stand or burgeoning rebellion against evidence-based medicine? A book review of [50]", *Preventive Medicie*, 2007, Vol. 45, pp. 491–494.

[85] R. S. Nickerson, "Confirmation bias: A ubiquitous phenomenon in many guises", *Review of General Psychology*, 1998, Vol. 2, pp. 175–220.

[86] U. Neisser, "Rising scores on intelligence tests", *American Scientist*, 1997, Vol. 85, No. 5, pp. 440–447.

[87] H. V. Ngo, T. Martinetz, J. Born, and M. Mölle, "Audiroty closed-loop stimulation of the sleep slow oscillation enhances memory", *Neuron*, 2013, Vol. 78, pp. 545–553.

[88] H. T. Nguyen, V. Kreinovich, B. Wu, and G. Xiang, *Computing Statistics under Interval and Fuzzy Uncertainty*, Springer Verlag, Berlin, Heidelberg, 2012.

[89] E. Nuhfer, C. Cogan, S. Fleischer, E. Gaze, and K. Wirth, "Random number simulations reveal how random noise affects the measurements and graphical portrayals of self-assessed competency", *Numeracy*, 2016, Vol. 9, No. 1, Paper 4.

[90] E. Nuhfer, S. Fleischer, C. Cogan, K. Wirth, and E. Gaze, "How random noise and a graphical convention subverted behavioral scientists' explanations of self-assessment data: numeracy underlies better alternatives", *Numeracy*, 2017, Vol. 10, No. 1, Paper 4.

[91] T. Olson, *Applied Fourier Analysis: From Signal Processing to Medical Imaging*, Springer, New York, 2017.

[92] J. L. Ong, J. C. Lo, N. I. Y. N. Chee, G. Santostasi, K. A. Paller, P. C. Zee, et al., "Effects of phase-locked acoustic stimulation during a nap on EEG spectra and declarative memory consolidation", *Sleep Medicine*, 2016, Vol. 20, pp. 88–97.

[93] J. Ou, H. Liang, and H. X. Tan, "Identifying elderies at risk of becoming more depressed with Internet-of-Things", In: J. Zhou and G. Salvendy (eds.), *Proceedings of the 4th International Conference on Human Aspects of IT for the Aged Population ITAP'2018*, Las Vegas, Nevada, July 15–20, 2018, Springer Lecture Notes in Computer Science, Vol. 10927, pp. 348–361.

[94] C. H. Papadimitriou, *Computational Complexity*, Addison-Wesley, San Diego, 1994.

[95] N. A. Papalambros, G. Santostasi, R. G. Malkari, R. Braun, S. Weintraub, K. A. Paller, and P. C. Zee, "Acoustic enhancement of sleep slow oscillations and concomitant memory enhancement in older adults", *Frontiers in Human Neuroscience*, 2017, Vol. 11, Article 109.

[96] P. J. E. Peebles and B. Ratra, "The cosmological constant and dark energy", *Reviews of Modern Physics*, 2003, Vol. 75, No. 2, pp. 559–606.

[97] J. B. Peterson, *12 Rules for Life: An Antidote to Chaos*, Random House Canada, Toronto, Canada, 2018.

[98] H. Petroski, "Framing hypotheses: A cautionary tale", *American Scientist*, 2003, Vol. 91, No. 1, pp. 18–22.

[99] J. Pietschnig and G. Gittler, "A reversal of the Flynn effect for spatial perception in German-speaking countries: Evidence from a cross-temporal IRT-based meta-analysis (1977—2014)", *Intelligence*, 2015, Vol. 53, pp. 145–153.

[100] J. Pietschnig and M. Voracek, "One century of global IQ gains: a formal meta-analysis of the Flynn effect (1909–2013)", *Perspectives on Psychological Science*, 2015, Vol. 10, No. 3, pp. 282–306.

[101] I. Prigogine, "Modération et transformations irréversibles des systémes ouverts", *Bulletin de la Classe des Sciences, Académie Royale de Belgique*, 1945, Vol. 31, pp. 600—606.

[102] I. Prigogine, *Etude Thermodynamique des phénoménes irréversibles*, Desoer, Liége, France, 1947.

[103] J. G. Proakis and D. K. Manolakis, *Digital Signal Processing*, Pearson, New York, 2006.

[104] J. Raven, "The Raven's Progressive Matrices: change and stability over culture and time", *Cognitive Psychology*, 2000, Vol. 41, No. 1), pp. 1–48.

[105] S. K. Reed, *Cognition: Theories and application*, Wadsworth Cengage Learning, Belmont, California, 2010.

[106] M. Rönnlund and L. G. Nilsson, "Flynn effects on sub-factors of episodic and semantic memory: parallel gains over time and the same set of determining factors". *Neuropsychologia*, 2009, Vol. 47, No. 11, pp. 2174–2180.

[107] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization", *Advances in Neural Information Processing Systems*, 2015, pp. 1257–1264.

[108] G. Saliba, *Islamic Science and the Making of the European Renaissance*, MIT Press, Cambridge, Massachusetts, 2011.

[109] W. Samuelson and R. Zeckhauser, "Status quo bias in decision making", *Journal of Risk and Uncertainty*, 1988, Vol. 1, pp. 7–59.

[110] T. Schlosser, D. Dunning, K. L. Johnson, and J. Kruger, "How unaware are the unskilled? Empirical tests of the 'signal extraction' counterexplanation for the Dunning-Kruger effect in self-evaluation of performance", *Journal of Economic Psychology*, Vol. 39, pp. 85–100.

[111] D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*, Chapman and Hall/CRC, Boca Raton, Florida, 2011.

[112] S. Shukeilo, *A New Probabilistic Approach to the Knapsack Problem*, Leningrad Electrical Engineering Institute (LETI), Master's Thesis, 1988 (in Russian).

[113] P. Slovic and B. Fischhoff, "On the psychology of experimental surprises", *Journal of Experimental Psychology: Human Perceptions and Performance*, 1977, Vol. 3, pp. 544–551.

[114] E. D. Smith and A. T. Bahill, "Attribute substitution in systems engineering", *Systems Engineering*, 2010, Vol. 13, No. 2, pp. 130–148.

[115] J. Soni and R. Goodman,*A Mind at Play: How Claude Shannon Invented the Information Age*, Simon & Schuster, New York, 2018

[116] J. M. Stecklein, J. Dabney, B. Dick, B. Haskins, R. Lovell, and G. Moroney, "Error cost escalation through the project life cycle", *Proceedings of the 14th Annual International Symposium of the International Council on Systems Engineering (INCOSE)*, Toulouse, France, June 19–24, 2004; also available as NASA Johnson Space Center Technical Report JSC-CN-8435, https://ntrs.nasa.gov/search.jsp?R=20100036670.

[117] J. Sundet, D. Barlaug, and T. Torjussen, "The end of the Flynn effect?: A study of secular trends in mean intelligence test scores of Norwegian conscripts during half a century", *Intelligence*, 2004, Vol. 32, No. 4, pp. 349–362.

[118] P. Szendro, "Pink-noise behavior of biosystems", *European Biophysics Journal*, 2001, Vol. 30, No. 3, pp. 227–231.

[119] R. Tamblyn, E. Tewodros, A. Huang, N. Winslade, and P. Doran, "The incidence and determinants of primary nonadherence with prescribed medication in primary care: a cohort study", *Annals of Internal Medicine*, 2014, Vol. 160, pp. 441–450.

[120] T. Teasdale, "Continuing secular increases in intelligence and a stable prevalence of high intelligence levels", *Intelligence*, 1989, Vol. 13, No. 3, pp. 255–262.

[121] T. W. Teasdale and D. R. Owen, "A long-term rise and recent decline in intelligence test performance: The Flynn Effect in reverse", *Personality and Individual Differences*, 2005, Vol. 39, No. 4, pp. 837–843.

[122] T. W. Teasdale and D. R. Owen, "Secular declines in cognitive test scores: A reversal of the Flynn Effect", *Intelligence*, 2008, Vol. 36, No. 2, pp. 121–126.

[123] R. L. Thorndike, "Mr. Binet's Test 70 Years Later", *Educational Researcher*, 1975, Vol. 4, No. 5, pp. 3–7.

[124] K. S. Thorne and R. D. Blandford, *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics*, Princeton University Press, Princeton, New Jersey, 2017.

[125] G. Tononi, B. A. Riedner, B. K. Hulse, F. Ferrarelli, and S. Sarasso, "Enhancing sleep slow waves with natural stimuli", *Medicamundi*, 2010, Vol. 54, pp. 82–88.

[126] L. H. Trahan, K. K. Stuebing, J. M. Fletcher, and M. Hiscock, "The Flynn effect: a meta-analysis", *Psychological Bulletin*, 2014, Vol. 140, No. 5, pp. 1332–1360.

[127] R. Trejo, V. Kreinovich, I. R. Goodman, J. Martinez, and R. Gonzalez, "A realistic (non-associative) logic and a possible explanations of 7 ± 2 law", *International Journal of Approximate Reasoning*, 2002, Vol. 29, pp. 235–266.

[128] G. C. Van Orden, J. G. Holden, and M. T. Turvey, "Self-organization of cognitive performance", *Journal of Experimental Psychology: General*, 2003, Vol. 132, No. 3, pp. 331–350.

[129] S. Weinberg, *Cosmology*, Oxford University Press, Oxford, UK, 2008.

[130] W. A. Wickelgren, *How to Solve Problems: Elements of a Theory of Problems and Problem Solving*, Freeman, San Francisco, 1974.

[131] J. A. Yesavage, T. L. Brink, T. L. Rose, O. Lum, V. Huang, M. Adey, and V. O. Leirer, "Development and validation of a geriatric depression screening scale: a preliminary report", *Journal of Psychiatric Research*, 1983, Vol. 17, pp. 37–49.

[132] J. A. Zachman, "A framework for information systems architecture", *IBM Systems Journal*, 1999, Vol. 38, No. 2–3, pp. 454–470.

[133] *Zachman Framework*, https://en.wikipedia.org/wiki/Zachman_Framework, downloaded on January 29, 2019

[134] F. Zapata, A. Posadas, A. Akundi, and E. Smith, "Combinatorial black box testing for genetic algorithms with discrete and continuous variables and potential applications for general testing methods", *Proceedings of the Joint Workshops on "System-of-Systems: A Network Integration Evaluation (NIE) Experience" and "R&D 'Use' of Propulsion Vehicles"*, El Paso, Texas, July 16–18, 2013.

[135] C. Zhao, S. Sun, L. Han, Q. Peng, "Hydrid matrix factorization for recommender systems in social networks", *International Journal on Neural and Mass-Parallel Computing and Information Systems*, 2016, Vol. 26, No. 6, pp. 559–569.

Curriculum Vitae

Griselda Valdepeñas Acosta earned her Bachelor of Engineering degree in Mechanical Engineering from New Mexico State University in 2000. In 2004 she received her Master of Science degree in Mechanical Engineering from the University of Texas at El Paso. She earned her second Bachelor's degree in Biological Sciences from the University of Texas at El Paso in 2014 and joined the University of Texas at El Paso doctoral program in Electrical and Computer Engineering in 2015.

Her peer-refereed papers were published in several academic journals published worldwide.

Permanent address: 9565 Waverly Dr., El Paso, Texas 79924