TILBURG ◆ UNIVERSITY

Dies Natalis, November 16, 2017

Speech by Rinnooy Kan

I am genuinely touched by the honor that this university has chosen to bestow on me; it came to me as a joyful surprise. It is not easy for happy recipients on occasions like these, as surprised as they are to be praised for virtues that they were themselves largely unaware of, to formulate a proper response. In fact, many of them will recall the Prizewinner's Prayer: "Lord, forgive them for their exaggerations, and forgive me for enjoying them so much." It is a privilege and a pleasure to return to Tilburg under these rewarding circumstances. Forty years ago, I represented the Erasmus University here on the fiftieth anniversary of the university, and watched Ernst Hirsch Ballin award a honorary degree to Jo van der Hoeven. We have both come a long way since then.

In gratefully accepting this degree, I am honored to represent a scientific discipline that I was lucky enough to join 45 years ago almost to the day. Its prevailing English name: Operations Research. It dates back to the Second World War, but its core notion, the development and analysis of mathematical models to represent and solve complex planning problems, turned out have tremendous peacetime potential as well. Optimization, the computation of optimal decisions, took off as a planning paradigm in the fifties and has lost none of its relevance and practicality today.

I came to Operations Research in 1972, having lost a PhD opportunity in pure mathematics - algebraic geometry - to a nationwide academic job freeze in January of that year. In retrospect, that now strikes me as an unmitigated blessing in disguise. My new supervisor gave me a German textbook to read and introduced me to Jan Karel Lenstra, whose brief and deadly review of the book revealed his impeccable quality standards for the first time, and set the stage for a happy professional collaboration that would last for more than 20 years.

Operations Research is a branch of applied mathematics and proud to be so, but its practitioners would readily acknowledge that its history can also be viewed as an ongoing love affair with the computer.

When Jan Karel and I started our research, the computer had already firmly established itself as the indispensable servant of the discipline, whose oldest and most successful tool, the simplex algorithm to solve linear programming problems, could only flourish because of its suitability for computer implementation. Over time, new mathematical insights and technological improvements went hand in hand to expand the range of successful Operations Research applications. At the same time, it had become clear that the design and analysis of algorithms, their implementation, their properties and their limitations represented a scientific discipline in its own right. In this way, ideas from theoretical computer science enabled the introduction of the computer as a theoretical research construct rather than as just an empirical computing device. For Operations Research, this turned out to be a very fruitful expansion.

In combinatorial optimization, our subdiscipline, the basic challenge is to find an optimal solution among a huge but finite set, the standard example being the task for a travelling salesman to find the shortest route that connects all the cities he needs to visit. Complete enumeration of all possible routes would obviously do the trick, but would require a forbidding amount of computing time: already for a small number of cities, it would take billions of years on the fastest computer imagineable. Can we do significantly better than that, that is to say: can we find the optimal route through a computing effort that goes up slowly and not exponentially as the number of cities increases?

A technique first proposed in 1972 allows for a tentative negative answer to this question. It enabled the identification of a large class of problems closely linked to each other, in the sense that a fast solution method for one of them could be used to construct a similarly fast solution method for all the others as well. And since this class includes all the notorious problems - including the one of the travelling salesman - for which such a fast solution method has long been sought in vain, chances are negligible that that will ever happen.

Thanks to computer science, my earlier statement is a precise mathematical one and not a fuzzy outcome of some empirical experiment. It is all the more surprising that after 45 years of hard work the final link is still missing in that the existence of that single fast solution method for all these notoriously hard problems, as unlikely as it is, has not been rigidly ruled out yet. By now, it defines one of the most intriguing open problems in mathematics. And as frustrating as that is, it also speaks to the mysterious depth of the extended paradigm.

For Operations Research - more specifically, and with due apologies to all those who do outstanding work on the stochastic side of the discipline, for the subfield of optimization - this new perspective represented a genuine breakthrough. It came as an addition to, not a substitute for, a broad tradition of research that already filled up nine solid 900 page handbooks in the nineties and would undoubtedly fill up a multiple of that number today. In a sense, it transformed the computer from a faithful, submissive servant to a challenging, occasionally surprising partner. It was a privilege for me to participate in that transformation and witness its effects, also after I left the university to move to occupations that serious mathematicians could only frown upon. From there, I watched the incredible rise of internet, leading up to the bubble around the turn of the century and its aftermath, and realised that the building blocks for internet were already in place when I visited MIT in 1984 but had remained curiously unexploited for many more years even in that most entrepreneurial of all academic institutions. The computer manages to surprise even its most gifted adherents.

The world continues to change. Today, the theme that this university has chosen to celebrate its 90th anniversary is “The Digital Society”, a society in which the by now ubiquitous internet combines with an abundance of cheap computing power and massive storage space into what many are tempted to refer to as a great digital revolution.

Revolution or not, if any discipline would appear to well positioned to contribute to a world that is flooded by digital data waiting to be transformed into useful information, it should be Operations Research. I have no doubt that its contribution is forthcoming and will be substantial. But it is ironical that it will coincide with the next step in its ongoing love affair with the computer, in which the former servant is now rapidly transforming itself into a serious competitor. That, indeed, defines a genuine revolution.

More on that in a minute. First of all, where do all these digital data come from? Partially from the traditional sources: micro- and macroeconomic data to which all of us contribute as we consume and invest. But numerical data form an increasingly small subset of a much larger whole that feeds on all human senses: texts, images, sounds, smells - all of them digitized, of course. These data come to internet from billions of sensors, soon to be increased to tens or hundreds of billions, that are finding their way into our cities, our workplaces, our cars, our homes, our clothes and, with some help from the surgeon, our bodies. Today, the new amount of data added daily is about one exabyte, a far - very far - descendant from the kilobytes and megabytes of our childhood. The amount of information in one kilobyte is about one short story; the amount in one exabyte is about equal to all the words spoken by mankind since the dawn of civilization (including these). If this amount is now added daily, the term Big Data starts to look like an understatement.

To inspect and register all those data, let alone to store and analyze them, is a stupendous task. There is nothing wrong with the traditional tools to do this - except that the numbers are so unbelievably large. But the familiar mathematical statistics of pattern recognition, including all the tools that the social sciences have added over time, can still be used to classify and detect structure, and can still identify correlation and nominate it for true causality. And where appropriate, Operations Research can still step in and feed these data into optimization models to compute optimal or (more likely) near-optimal decisions.

Except that the numbers are so unbelievably large. And so the first challenge that Big Data poses to Operations Research is to reconsider, for this reason alone, the interaction between historical data, modelparameter estimation and model calibration. Somewhat underresearched, it seems to call for an iterative, wholly or partially automated process, whereby intermediate optimization outcomes feed into the data sampling process. I was pleased to understand that Dick den Hertog is pursuing some of these ideas here in Tilburg, and I wish him all success in doing so.

To carry out effective Operations Research under these novel circumstances requires novel competences. The rapid growth of a new dicipline Data Science, of which Operations Research is a part but which also includes other relevant specializations in mathematics and computer science, is a very welcome step. The job market for young Data Scientists alone is bound to remain oustanding, and I was impressed to read about the joint initiative by the universities of Tilburg and Eindhoven to train them together in Den Bosch. If Operations Research reappears in the new curriculum under its new name Business Analytics, so be it: as insiders will know, it is the sixth label in a row under which the discipline manifests itself in The Netherlands. What matters are the results, and there I remain really proud of our undiminished international reputation, as witnessed by the high Dutch share of prestigious Franz Edelman Awards, journal editorships and consultancy firms, including Ortec, the largest of them all, whose creation I was privileged to watch thirty years ago in Rotterdam from just a few blocks away. In Operations Research, the university of Tilburg can be proud to be the highest placed non-US institution after INSEAD in a recent global ranking, followed closely, I was happy to see, by the Erasmus University of Rotterdam.

It is within Data Science that the love affair between Operations Research and the computer that I alluded to before enters into its next phase. The astonishing recent successes, after many historical ups and downs, of Artificial Intelligence have positioned the computer as an independent competitor in the art of problem solving, way up from its original humble servant’s role. Multilayer neural networks, whose structure is modelled after the human brain, generate fully automated deep learning processes that are coping successfully with notorious human challenges such as playing chess, Go and poker and, perhaps more relevantly, driving a car, playing the market and solving other complex planning puzzles.

Fascinatingly enough, the mechanism behind the greatest of these successes is fundamentally intransparent. It is simply not clear why these computer programs work as well as they do - but they do. Hence, we are being invited to delegate the solution of many of our practical problems to a mechanized version of our own brain that has managed to escape from our full control during the design process. Little wonder that serious scientists, Stephen Hawking among them, warn about a forthcoming battle with artificially intelligent machines that we are already poorly positioned to win. As these machines draw on techniques of Operations Research as one of their many tools, could it be that the computer that set out as a reliable servant would triumphantly reemerge at the end as our uncontrollable dictator?

These are very big issues, affecting our perspectives on the core of human potential and its limitations. When a human being is viewed as little more than a set of data and a collection of algorithms, indeed as little more than as an exercise from a data science textbook, as is the case in a recent bestseller with the provocative title “Homo Deus”, then that title alone underlines the need for an academic debate that stretches way beyond mathematics and computer science into the discipline represented by my fellow honorary doctor, and beyond that into the humanities at large. We need them more than ever.

From its profile and its partnerships, this university is very well positioned to contribute to this unravelling, fundamental debate. I am happy to wish it well on its ninetieth birthday, and grateful for the opportunity to speak on this festive occasion.