



## Preface

Operations Research (OR) is increasingly used to model, solve, and analyse problems in Biology, Medicine and Healthcare. In this issue of *Algorithmic Operations Research* we present seven articles at the forefront of this interface. Collectively, these articles open a succinct window into the burgeoning interactions that are broadening OR's impact on some of today's most meaningful problems. Advances in genetics and micro-biology are hurling mankind into a state of understanding that could lead to breathtaking innovations in medicine. In turn, the already intricate field of medicine is increasingly challenged by fundamental biological insights that propel things like drug design, genetic tests for disease, and techniques to correct degenerative processes. Medicine is further complicated by increasingly sophisticated technology, making the decision of how to best treat an individual patient nearly overwhelming. Consequently, managing the entirety of a healthcare system so that new science and technology is efficiently used is the daunting task faced by administrators. This issue highlights how OR can have a significant impact at each level of this chain, which begins with the fundamental discoveries in science that then transcend medical boundaries and further complicate the conundrum of healthcare management.

This issue begins with four articles at the intersection of Biology and OR. First, Renato Bruni considers a combinatorial optimisation model to determine the sequence or composition of polymeric compounds, which constitute all proteins. Information encoded in DNA is transcribed into messenger RNA and then used to form proteins, which are the structural and functional entities of a cell. Proteins are sequences of polymers, and the problem of arranging the polymers so that the resulting sequence renders a specified protein is of crucial importance. Dr. Bruni shows how to model and enumerate feasible peptide sequences through combinatorial methods.

Richard J. Forrester and Harvey J. Greenberg address four problems in computational biology that are related to our understanding of proteins. Two of these concern the problem of deciding the similarity between sequences of nucleic or amino acids, and the others concern the folding process that is undergone by all proteins as they move from a sequence of amino acids to

the final 3D structure that defines their function. Each of these problems is well established and naturally lends itself to 0 – 1 quadratic programming. However, solution methods are often tailored to the particular application, and the question posed by Drs. Forrester and Greenberg is whether or not it is possible to use a standard relaxation technique together with a commercial solver. Four different reformulation methods are considered, and after considerable numerical work, the conclusion is that state-of-the-art commercial software together with appropriate modeling are close to providing reliable solutions on real problems. This is promising since further advances in software and hardware are likely to allow these problems to be solved with reliable and robust commercial packages instead of heuristics that require constant modification as our understanding increases.

The third paper by Markus Bauer, Gunnar W. Klau, and Knut Reinert considers the similarity between different RNA strands. Unlike many other sequence alignment techniques, their method accounts for the spatial conformation of RNA. Similarity scores are modeled as the optimal value of a 0 – 1 linear program that is built from a graphical representation of the problem. A standard Lagrangian relaxation is used to solve the problem, and the solutions are shown to compare well with dynamic programming techniques when two sequences are considered. This article exhibits much of OR's applicability to problems in micro-biology, as it develops a new model from a graphical representation, continues with an analysis of both the model and the solution technique, and concludes with computational results that validate the work.

Systems biology is the study of whole-cell processes, a study made possible due to the modern high-throughput experiments that provide massive amounts of data about the elemental relationships within a cell. The goal of systems biology is to use this information to model the entirety of a cellular system. Stefan Bleuler, Philip Zimmerman, Markus Friberg and Eckart Zitzler address a problem in systems biology that attempts to associate the genes of a gene regulatory system, which is modeled as a network in which nodes are genes and arcs are regulatory relationships. The question of grouping entities with similar function is a long standing problem that is related to the similarity problems

mentioned above. The technique used in this paper is based on a multi-objective optimisation problem that is solved with a combination of an evolutionary algorithm and a greedy heuristic. Numerical comparisons with standard techniques show that the authors' approach competes favorably.

The fifth article of this issue moves from biology to medicine, and in particular, to the problem of optimising the delivery of a radiotherapy treatment. The optimal design of a radiotherapy treatment has received significant attention over the last ten years, and today all commercial planning systems incorporate an optimisation component. One of the modern treatment modalities is called Intensity Modulated Radiotherapy (IMRT), a technique defined by the use of a multi-leaf collimator to modulate a high-energy beam to better treat a patient. Thomas Kalinowski builds on previous work that models the question of optimally sequencing the collimator's movements as a network flow problem. This article advances this approach by considering a clinical consideration that is concerned with the fact that the computed dose can differ from the delivered dose due to awkward sequencing when the design of the collimator is neglected. Dr. Kalinowski shows that the adapted model and solution method minimise exposure time as

---

Matthias Ehrgott and Al Holder

well as reduce the number of collimator updates.

The next article by Hui Jiang, Yifan Liu and Zheng Su also addresses a problem in oncology, but instead of radiotherapeutic treatments, the authors consider scheduling chemotherapy treatments. Taking into account the disease's state, the authors develop a recurrence relation that models a treatment schedule. The authors rigorously show that the recurrence relation has a unique solution, and they continue to present an algorithm that is capable of calculating an optimal schedule. The authors note that their method extends to the stochastic case where the state of the disease is a random variable.

The final article by Cristina Azcárate, Aurora Gafaro and Fermín Mallor moves into the broad area of healthcare. This article considers the problem of managing a clinic in which patient arrivals are random. Such scenarios are commonly modeled by queuing theory and analysed with simulation. The authors follow a similar strategy but embed a multi-objective optimisation problem in the simulation, an idea that allows them to consider efficient states. The technique is applied to a hospital in Pamplona, Colombia, for which it allows insight into a complicated decision process.

We have enjoyed bringing these articles together and are excited about the gamut of applications that await OR in biology, medicine and healthcare. We hope readers will feel the same. Last but not least, we thank the referees who have given valuable advice on the submitted papers.