Handling uncertainty in health care management using the cardinality-constrained approach: Advantages and remarks

Bernardetta Addis [a], Giuliana Carello [b], Andrea Grosso [c], Ettore Lanzarone [d,*], Sara Mattia [e], Elena Tànfani [f]

[a] *LORIA (UMR 7503 CNRS), Université de Lorraine, INRIA Nancy-Grand Est, France*
[b] *DEIB, Politecnico di Milano, Milan, Italy*
[c] *DEI, Università di Torino, Turin, Italy*
[d] *CNR-IMATI, Milan, Italy*
[e] *CNR-IASI, Rome, Italy*
[f] *DIEC, Unversità di Genova, Genoa, Italy*

Received 13 June 2014
Accepted 10 October 2014
Available online 8 November 2014

1. Aim and scope

The purpose of this short communication is to highlight the potential of a specific robust optimisation approach, namely the cardinality-constrained model, in handling health care management problems. In fact, although robust optimisation approaches have been widely applied in health care, the cardinality-constrained model has seldom been considered. We briefly report our experience with this approach, which in our view can be fruitfully applied to health care management problems, pointing out some advantages and considerations that can be useful for operations researchers while applying this approach to the same field.

* Corresponding author.
E-mail addresses: bernardetta.addis@loria.fr (B. Addis), giuliana.carello@polimi.it (G. Carello), grosso@di.unito.it (A. Grosso), ettore.lanzarone@cnr.it (E. Lanzarone), sara.mattia@iasi.cnr.it (S. Mattia), etanfani@economia.unige.it (E. Tànfani).

2. Uncertainty in health care management

The demand for mathematically based decision support tools to optimise the delivery of health care services has been growing significantly in recent years. In Western countries, the need for efficient resource management is constantly increasing due to the ageing population on the one hand, which is increasing the demand for health care services, and the reduction in public health funding on the other. Several studies have been conducted with the aim of improving the ratio of service quality to cost in health care systems. Among the adopted quantitative techniques, operations research methods have proved to be effective in many sectors where decision making support is necessary, such as industrial planning and logistics, and are currently widely applied to health care-related problems.

Typical features of health care systems include a high level of complexity and a large amount of available data, which have only recently begun to be stored digitally. Moreover, one of the major problems faced by people working in health care management is data uncertainty [1]. Available data need to be properly analysed and processed to obtain reliable input parameters. In fact, data reliability is a key factor in guaranteeing the feasibility and efficiency of the obtained solution when it is applied to the real system. In health care, uncertainty may arise in different contexts and due to different causes. For example, in the emergency vehicles location problems, uncertainty arises from ambulance availability [2]; in planning and scheduling operating rooms, uncertain data are the duration of surgery, length of stay and rate of patient arrivals [3]; in home care services, the demand for care is uncertain in terms of the number, duration and types of visits [4]. In all cases, uncertainty cannot be neglected, as it may have a significant impact on the solution and on the quality of service provided to patients. For this reason, robust methods that explicitly take data uncertainty into account are drawing the attention of the research community.

3. Methods and approaches for robustness

The concept of solution robustness is crucial when handling uncertain data or parameters. Generally speaking, a solution is considered robust if it is able to protect against uncertainty; i.e., the solution remains feasible even when parameters change with respect to their nominal value, and the deterioration of the solution quality is controlled or limited when it is applied to the real uncertain system instead of the nominal case. Moreover, a trade-off must be pursued between these two aspects, i.e., the level of robustness (the feasibility in different scenarios defined by possible realisations of parameters), and the efficiency and cost of the solution (how much the solution deteriorates when applied to the uncertain system). Indeed, a very conservative solution tailored to unlikely scenarios may turn out to be highly expensive for more likely scenarios or, alternatively, a solution that is not conservative enough may become infeasible even for small variations with respect to the nominal problem.

Different approaches have been proposed in the literature to address uncertain parameters in optimisation problems. They belong to three main groups: stochastic programming, distributionally robust optimisation, and robust optimisation.

In stochastic programming [5,6], uncertain parameters are modelled as random variables and the probability distribution of each parameter is assumed to be known. Hence, to approach real life problems, stochastic programming requires both a strong statistical background to manage the mathematical models and a thorough knowledge of the real problem to derive the probability distributions, which are not always easy to derive. The main advantage of stochastic programming is that it produces solutions that are usually not over-conservative as they protect against the most likely realisations. The resulting optimisation problems can be difficult to solve, however, often involving a very wide number of scenarios. In addition, if the given distributions are not reliable (i.e., they are different from the real ones), the solutions produced by a stochastic method may not prove to be robust.

On the contrary, distributionally robust optimisation [7,8] and ambiguous chance-constrained approaches [9] assume that the probability distribution is not known, but lies within a known family of distributions. The solution must protect against the worst-case realisation given by the admissible probability distributions. The problem is difficult, but computationally tractable approximations exist.

Robust optimisation approaches [10–14] are a good compromise between the above mentioned methods. Robust optimisation approaches assume that each uncertain parameter belongs to a given convex set, and no detailed knowledge of its probability distribution is required. These approaches guarantee that the solution is feasible for all of the values of the parameters within the considered uncertainty set. They are usually computationally more tractable than the stochastic programming approaches. The produced solutions may be over-conservative, but the possibility of managing the convex sets and the parameters' movement within them allow the level of robustness to be adjusted.

In the simplest definition of the uncertainty set, the uncertain parameters are assumed to belong to an interval. Hence, when addressing a (possibly mixed-integer) linear program of the form

$$\begin{aligned} \text{minimise} \quad & \sum_j c_j x_j \\ \text{s.t.} \quad & \\ & \sum_j a_{ij} x_j \le b_i \quad \forall i \\ & x_j \ge 0 \quad \forall j \end{aligned}$$

a nominal value $\bar{a}_{ij}$ and a maximum deviation δ_{ij} are given for each uncertain parameter a_{ij}, i.e., $a_{ij} \in [\bar{a}_{ij} - \delta_{ij}, \bar{a}_{ij} + \delta_{ij}]$.

The cardinality-constrained robust optimisation method [14] considers such uncertainty sets. As it is unlikely that all parameters will deviate from their nominal value at the same time, the method assumes also that only a subset of at most Γ_i parameters is allowed to change in the ith constraint. In this sense, the *cardinality* of the parameters permitted to change is *constrained*. The generated solution must be feasible for the worst possible scenario in which Γ_i parameters assume their maximum values while the others assume their nominal ones.

We remark that very little knowledge of the parameters' variability is required; furthermore, the decision maker can tune the level of robustness by setting the parameters Γ_i. It is interesting to note that, as the constraints added by the cardinality-constrained approach are linear [14], the robust counterpart of an integer linear programming problem can be solved by any commercial solver, except in very large instances. Additionally, the robust counterpart of a linear programming model is still a linear programming model, and can therefore be solved in polynomial time.

The area of robust optimisation is vast and continuously expanding. For instance, scenario-based approaches are proposed in [15], whose aim is to produce solutions that have the best worst-case objective value considering all possible scenarios. In adjustable robust optimisation methods [16,17], a multistage decision policy is used. In [18,19] robust optimisation approaches have been proposed in which the robustness is enforced by cutting planes. A computational study on cutting planes versus reformulations is also provided in [20]. For other approaches and in-depth analyses of robust optimisation, we address the reader to Ben-Tal et al. [21], Bertsimas et al. [22], Gorissen et al. [23], and the references within.

4. Applying the cardinality-constrained approach to health care management problems

Several approaches are largely applied when handling uncertainty in health care problems, such as probabilistic models [2,24] or stochastic optimisation approaches [25,26]. On the contrary, to the best of our knowledge, the cardinality-constrained approach has thus far been rarely applied to health care problems. In May 2014, a search on Scopus for papers citing Bertsimas and Sim [14] in relation to health care yielded only five results [27–31] other than our contributions in a conference book [32,33], the latter also extended in a journal [34].

However, the cardinality-constrained approach seems to be suitable for handling several health care management optimisation problems. We recently tested this hypothesis by applying the cardinality-constrained approach to two relevant health care problems: the operating room planning problem and the nurse-to-patient assignment problem in home care services. In the following, we briefly describe our experience in applying this approach to these problems, focusing on general considerations, advantages, and remarks that can help operations researchers in evaluating the approach for other applications in the health care field.

5. Our experience in two applications

The surgical case assignment problem is considered in [35] (extended work of Addis et al. [32]). For each specialty, a given number of surgical blocks is available, together with their allocable timespan. Patients must be assigned to these blocks taking into account their surgery durations, with the goal of minimising the total delay for patients. The surgery duration is uncertain, and the cardinality-constrained approach is applied to compute an assignment of surgery cases to the blocks, which shall not exceed each block's capacity, even if a subset of surgery durations per block assumes the maximum value instead of the expected one. In this way, a certain degree of robustness is obtained without the need to protect against all possible realisations of surgery duration.

The nurse-to-patient assignment problem arises in home care service management when the continuity of care is pursued [34]. Continuity of care means that only one nurse (i.e., the *reference nurse*) is assigned to each patient, and the assignment is kept over a long period. The problem consists of assigning each new patient entering the service to his/her reference nurse. Nurses have a certain amount of working hours over a defined time interval, usually a week, and must be extra paid for working overtime. Moreover, a maximum overtime is allowed for each nurse. The aim is to provide a suitable service quality, minimising overtime costs and/or balancing nurses' workloads. The amount of treatment time required for each patient in the time slot may vary with respect to its expected planned value, e.g., due to sudden changes in the patient's health conditions. Therefore, the time amount for visits required by each patient in each time interval is the uncertain parameter of the planning process. We tackled the problem by applying the cardinality-constrained approach, with the aim of guaranteeing that the obtained assignments are feasible even if a subset of patients assigned to each nurse requires the maximum treatment time.

The results of both applications are promising: computational times are reasonable, thus guaranteeing the applicability of the cardinality-constrained approach in real life cases. Moreover, when tested on a set of randomly generated scenarios and/or on historical data, the solutions display good behaviour in terms of quality performance metrics and improvement with respect to their non-robust counterparts.

Indeed, concerning the surgical case assignment problem, the robust model has been tested on a set of realistic instances and has been compared with the non-robust counterpart. The obtained solutions, both robust and non-robust, have also been tested on a set of scenarios. The number of patients operated on may decrease slightly when robustness is guaranteed; however, when the solutions are applied to scenarios, the robustness dramatically reduces – up to one third – the number of patients who cannot be operated on in the scheduled day due to surgery time variations with respect to the expected value. In this way, from the patient's point of view, the quality of the solution is highly improved. Increasing the level of robustness, the number of cancelled patients is almost zero, but on the other hand, the utilisation rate is reduced. Even when the highest level of robustness is required, however, the utilisation rate never drops below 50%.

Concerning the home care application, the proposed model has been tested on both historical data and a set of generated sample paths. The robust model guarantees a better performance in terms of overtime costs, as well as a more evenly distributed workload for the nurses, with respect to the nominal model. Briefly, applying the robust model instead of the non-robust counterpart has generated savings of up to 30% on penalty costs, considering historical data of a real home care provider over six months. Both palliative and non-palliative patients are treated by the considered provider. Advantages are mainly ascribed to non-palliative patients rather than to palliative ones because the latter have a more defined care pathway associated with lower uncertainty. Computational times needed to solve the model in practice with a weekly basis are reasonable. Compared to the other methodologies applied to this problem [24,26,36], the cardinality-constrained approach is able to produce competitive solutions. It requires reduced computational time and fewer assumptions on the variability of patients' demands. On the contrary, the stochastic programming approach [26] requires a high computational time due to the inclusion of stochasticity with the scenario generation, and the analytical approach based on stochastic ordering [24,36] requires introducing several assumptions on the shape of the density functions.

6. Analysis and general conclusions

The cardinality-constrained approach proves to behave well in the two considered applications, based on which, some of the advantages of this approach that can be useful for research into other health care management problems are highlighted below.

As already mentioned, this approach uses a very simple geometry for the convex sets: parameters are assumed to lie on an interval, and a limited number of them are assumed to take the maximum possible value instead of the nominal one. Therefore, the approach does not require the knowledge of the entire probability density functions of the uncertain parameters, but only some knowledge of them. In health care applications, enormous amounts of historical data are usually available, from which the little information required can be easily derived. In contrast, it can be difficult to find the more detailed information required for other approaches. Moreover, unlike other methods, the cardinality-constrained approach does not rely on a detailed description of the uncertainty, and for the same reason, it is less sensitive to estimation errors in the data and their probability distributions. Finally, the approach provides a robust solution with a reasonable computational effort.

Another relevant advantage is that the basic idea of the approach can be easily understood by clinicians and health care service managers without any background in statistics or operations research. The robustness of the solution can be easily tuned by the service managers themselves according to the desired level of risk protection, by imposing the cardinality of the subsets of parameters assuming their maximum values. Such cardinality has a practical meaning that can easily be interpreted, and its impact can be tested to properly select the trade-off between the level of robustness and the cost of the solution.

It is worth noting several factors with respect to the parameter setting. Other parameters must be set besides cardinalities Γ_i, and their values must be carefully evaluated. Indeed, the approach assumes that the majority of uncertain parameters take their expected values, whereas a subset of them takes their maximum ones. Thus, each nominal value $\bar{a}_{ij}$ must have a real meaning in the considered application. In the health care applications described above, the lengths of treatment (e.g., surgery and visit durations) are usually continuous parameters, and their nominal values can be taken as the expected value of a suitable statistical distribution, or from standard values given by national or regional health systems. If an integer parameter is considered, however, a fractional expected value does not occur in any realisation. Regarding the maximum values $\bar{a}_{ij} + \delta_{ij}$, if a distribution is derived for a_{ij} based on historical data, attention must be paid when setting the value of the maximum deviation δ_{ij}. If we are not willing to protect against some values as they are considered to be very unlikely, these values must be outside the given interval. To do so, one possibility is to take $\bar{a}_{ij} + \delta_{ij}$ in correspondence of a given quantile of the distribution rather than the maximum value.

In addition to the benefits, we have experienced two main drawbacks in conducting our work [34,35] that we consider noteworthy. First, the approach may become computationally expensive for large instances, and computational time may also depend on the value of parameters Γ_i. Second, characterising the uncertain parameters with only two values, $\bar{a}_{ij}$ and $\bar{a}_{ij} + \delta_{ij}$, might not be sufficient to capture the behaviour of peculiar parameters. Quite straightforward improvements in the method can overcome both drawbacks. Concerning the computational effort, cutting plane-based approaches [19] can be applied to reduce computational time. On the other hand, to provide a more detailed representation of uncertain parameters, additional side-constraints may be added to the model.

In general, we think that great advantages can be obtained by the standard formulation of the cardinality-constrained approach. Should any of these two drawbacks occur in the solution of a specific application, the alternatives we discussed are worthy of implementation.

References

[1] P. Han, W. Klein, N. Arora, Varieties of uncertainty in health care: a conceptual taxonomy, Med. Decis. Making 31 (2011) 828–838.

[2] L. Brotcorne, G. Laporte, F. Semet, Ambulance location and relocation models, European J. Oper. Res. 147 (2003) 451–463.

[3] B. Cardoen, E. Demeulemeester, J. Beliën, Operating room planning and scheduling: a literature review, European J. Oper. Res. 201 (2010) 921–932.

[4] E. Lanzarone, A. Matta, G. Scaccabarozzi, A patient stochastic model to support human resource panning in home care, Prod. Plan. Control 21 (2010) 3–25.

[5] J. Birge, F. Louveaux, Introduction to Stochastic Programming, Springer-Verlag, 1997.

[6] A. Shapiro, D. Dentcheva, A. Ruszczyński, Lectures on Stochastic Programming: Modeling and Theory, in: MPS/SIAM Series on Optimization, 2009.

[7] A. Ben-Tal, D. Bertsimas, D. Brown, A soft robust model for optimization under ambiguity, Oper. Res. 58 (2010) 1220–1234.

[8] J. Goh, M. Sim, Distributionally robust optimization and its tractable approximations, Oper. Res. 58 (2010) 902–917.

[9] E. Erdogan, G. Iyengar, Ambiguous chance constrained problems and robust optimization, Math. Program. 107 (2006) 37–61.

[10] A. Soyster, Convex programming with set-inclusive constraints and applications to inexact linear programming, Oper. Res. 21 (1973) 1154–1157.

[11] A. Ben-Tal, A. Nemirowski, Robust convex optimization, Math. Oper. Res. 23 (1998) 769–805.

[12] L. El-Ghaoui, F. Oustry, H. Lebret, Robust solutions to uncertain semidefinite programs, SIAM J. Optim. 9 (1998) 33–52.

[13] D. Bertsimas, M. Sim, Robust discrete optimization and network flows, Math. Program. B 98 (2003) 49–71.

[14] D. Bertsimas, M. Sim, The price of robustness, Oper. Res. 52 (2004) 35–53.

[15] P. Kouvelis, G. Yu, Robust Discrete Optimization and its Applications, Kluwer Academic Publisher, 1997.

[16] A. Ben-Tal, A. Goryashko, E. Guslitzer, A. Nemirovski, Adjusting robust solutions of uncertain linear programs, Math. Program. 99 (2004) 351–376.

[17] X. Chen, Y. Zhang, Uncertain linear programs: extended affinely adjustable robust counterparts, European J. Oper. Res. 57 (2009) 1469–1482.

[18] D. Bienstock, Histogram models for robust portfolio optimization, J. Comput. Finance 11 (2007) 1–64.

[19] M. Fischetti, M. Monaci, Cutting plane versus compact formulations for uncertain (integer) linear programs, Math. Program. Comput. 4 (2012) 239–273.

[20] D. Bertsimas, I. Dunning, M. Lubin, Reformulations versus cutting planes for robust optimization. a computational and machine learning perspective, Optimization (2014) Online.

[21] A. Ben-Tal, L.E. Ghaoui, A. Nemirovski, Robust Optimization, in: Princeton Series in Applied Mathematics, 2009.

[22] D. Bertsimas, D. Brown, C. Caramanis, Theory and applications of robust optimization, SIAM Rev. 53 (2011) 464–501.

[23] B. Gorissen, I. Yanikoglu, D.D. Hertog, Hints for practical robust optimization, in: Center Discussion Paper 065, 2013.

[24] E. Lanzarone, A. Matta, Robust nurse-to-patient assignment in home care services to minimize overtimes under continuity of care, Oper. Res. Health Care 3 (2014) 48–58.

[25] P. Beraldi, M. Bruni, D. Conforti, Designing robust emergency medical service via stochastic programming, European J. Oper. Res. 158 (2004) 183–193.

[26] E. Lanzarone, A. Matta, E. Sahin, Operations management applied to home care services: the problem of assigning human resources to patients, IEEE Trans. Syst. Man Cybern. A 42 (2012) 1346–1363.

[27] T. Chan, T. Bortfeld, J. Tsitsiklis, A robust approach to imrt optimization, Phys. Med. Biol. 51 (2006) 2567–2583.

[28] B. Denton, A. Miller, H. Balasubramanian, T. Huschka, Optimal allocation of surgery blocks to operating rooms under uncertainty, Oper. Res. 58 (2010) 802–816.

[29] M. Holte, C. Mannino, The implementor/adversary algorithm for the cyclic and robust scheduling problem in health-care, European J. Oper. Res. 226 (2013) 551–559.

[30] C. Mannino, E. Nilssen, T. Nordlander, A pattern based, robust approach to cyclic master surgery scheduling, J. Sched. 15 (2012) 553–563.

[31] C. Banditori, P. Cappanera, F. Visintin, A combined optimization-simulation approach to the master surgical scheduling problem, IMA J. Manag. Math. 24 (2013) 155–186.

[32] B. Addis, G. Carello, E. Tànfani, A robust optimization approach for the operating room planning problem with uncertain surgery durations, in: Springer Proceedings in Mathematics & Statistics (Proceedings of HCSE 2013), 2014, pp. 185–189.

[33] E. Lanzarone, G. Carello, Applying the cardinality-constrained approach in health care systems: the home care example, in: Springer Proceedings in Mathematics & Statistics, vol. 61 (Proceedings of HCSE 2013), 2014, pp. 61–72.

[34] G. Carello, E. Lanzarone, A cardinality-constrained robust model for the assignment problem in home care services, European J. Oper. Res. 236 (2014) 748–762.

[35] B. Addis, G. Carello, E. Tànfani, A robust optimization approach for the operating room planning problem with uncertain surgery duration—an extended analysis. HAL, Hyper Articles en ligne, 2014.

[36] E. Lanzarone, A. Matta, A cost assignment policy for home care patients, Flexible Serv. Manuf. J. 24 (2012) 465–495.