OPTIMA 104

Mathematical Optimization Society Newsletter

MOS Chair's Column

June 15, 2018. The ISMP in Bordeaux is only a few weeks away when I write this column. I am looking forward to the opening session with some appropriate tension! But, I am sure that the wine tasting on Sunday, meeting up with all colleagues, some of whom I have not seen since Pittsburgh, will make everything fall into place.

It has been a hectic winter and spring, but when I see how many people are involved in the whole process of making the symposium a success, I am humbled. The organizers and program committee have been busy over a long period. But we have also had seven juries working hard on selecting the prize winners, the Publications Committee searching for new editors, and the Executive Committee helping out with a lot of practical issues. MOS is a society ran by volunteers and I am very grateful to all of you who invest considerable time in the activities of the Society! A BIG thank you to you all!

Since the last column we have new Editors-in-Chief for MPA and MPB. Jon Lee succeeded Alexander Shapiro as E-i-C for MPA, and Sven Leyffer succeeded Jong-Shi Pang as E-i-C for MPB. Both journals are in great shape and are for many of us the first choice as scientific outlet. This OPTIMA issue is also the last issue edited by Volker, Sam, Jeff, and Christoph as a team, since Volker, Sam, and Jeff step down. You have given us a series of really interesting newsletters that are very much valued. Thanks to all our editors and their boards, and good luck to the new teams!

See you in Bordeaux!

Karen Aardal
Delft Institute of Applied Mathematics
k.i.aardal@tudelft.nl

Note from the Editors

Dear MOS members,
It has been a great pleasure to edit this newsletter over the last couple of years for at least two reasons: The fascinating contributions made by the authors and interview partners and the excellent and devoted technical production work of Christoph Eyrich – many thanks! In this issue, Sebastian Sager takes us on a journey through several applications of optimization in medicine. Enjoy reading this and all the upcoming issues of Optima!

Sam Burer, Co-Editor, Volker Kaibel, Editor,
Jeff Linderoth, Co-Editor

Contents of Issue 104 / June 2018

Sebastian Sager

Optimization and Clinical Decision Support

1 Introduction

Well known to readers of *Optima*, mathematical optimization has a long standing tradition of influential impact on many aspects of modern life. There is an abundance of application areas which have been profiting from mathematical optimization technology. It has been applied for the design, layout, operation, control, calibration, or analysis of products, networks, processes, and systems. Synergistically, almost all of these application areas have been stimulating mathematical research. The goal of this article is to illustrate how clinical decision support may profit from optimization and provide new challenges for the optimization community.

There are hence two main claims and take-away messages for the reader. First, despite the obvious complexity, diversity, and uncertainty of human bodies and what is going on inside them, surprisingly simple mathematical models combined with optimization may give accurate and extremely useful insight. It is my firm conviction that there are many cases where this goes far beyond what could be achieved, e.g., with pure data-driven machine learning approaches. Some mathematical models capture the main characteristics and are based on medical knowledge – while they are also accessible for real-time optimization and personalizable with measurements.

Second, new and interesting optimization problem classes and approaches may arise from clinical decision support problems. A particular challenge is that the determination (or personalization) of the mathematical model and the optimization of it interact and have to be performed at the same time. Surprising structures may arise, e.g., when the dimension of the optimization variables depends on the optimal solution itself.

We present three case studies from the ongoing European Research Council (ERC) Consolidator Grant project MODEST, all based on a close cooperation with medical partners. Naturally, this is only a very subjective selection and many more success stories in the intersection of optimization and clinical decisions have already been written by others. Yet, it hopefully helps to raise awareness of an interesting and important application area for optimization. All three examples are presented in a consistent manner with the clinical background, the related optimization problem, modeling issues, challenges and solution approaches, as well as results. Finally, a speculative outlook is given into the direction of a model-based training of clinical decision making.

2 Inverse Simulation for Cardiac Arrhythmia

Clinical Background and Optimization Question

There are several dozen different types of cardiac arrhythmia. While most of them can be easily differentiated by cardiologists, the discrimination between atrial fibrillation (AFib) and regular atrial arrhythmias including atrial flutter (AFlu) and focal atrial tachycardia

poses a diagnostic challenge. The usually available data, a surface electrocardiogram (ECG), looks very similar in both cases, to laymen, physicians, and computerized algorithms alike. Misinterpretation rates of up to 80 % have been reported [19], which is even more concerning as different treatments (often antiarrhythmics in AFib and ablation in AFlu) are implied by the diagnosis. Furthermore, atypical forms of AFlu are becoming increasingly important in clinical practice as a complication of left atrial ablation procedures.

The electrical pacemaker signal which eventually stimulates the mechanical pumping of the heart originates in the sine node. It propagates via the atrial chambers and the atrioventricular (AV) node towards the ventricular chambers. In AFib, already the electrical activation of the atria is chaotic and intrinsically irregular. In AFlu the activation of the atria is regular, but due to a complex blocking mechanism in the AV node also here the ventricular response is irregular. Whereas in theory it should be possible to see differences in the low voltage flutter waves that correspond to atrial polarizations, they are hardly discernible in practice. In summary, we deal with a situation where only time points of the clearly visible R waves (the beeps of a heart rate monitor) corresponding to ventricular activation are provided as input data for the diagnosis. We furthermore assume that a previous diagnosis already assured that either AFib or AFlu is the cause for the observed arrhythmia.

We use a mathematical model for the complex, but deterministic blocking behavior in the AV node and optimization to personalize this model to input data. A small optimal function value (i.e., a good match between simulation and measured data) is interpreted as a high likelihood for regular behavior (AFlu), a high optimal function value as a high likelihood for chaotic behavior which cannot be explained by the model (AFib).

Mathematical Modeling of the Dynamics and of the Optimization Problem

There are many algorithmic approaches to analyze cardiac arrhythmia, e.g., based on Fourier transformations, wavelets, clustering of RR times, machine learning, cellular automata, or nonlinear time series analysis. However, none of these approaches yields satisfactory results for our task and short ECGs [12]. Also many mathematical models have been proposed to represent electrical conductivity in the heart, most prominently the Noble adaptation [13] of differential equations, for which Hodgkin and Huxley were distinguished with the Nobel Prize in 1963. In this mathematical model the electrical potential across the membrane changes due to ion currents and is related to sodium and potassium currents. A mathematical model with four differential states and several estimated model parameters successfully predicted several so far unknown phenomena and led to many extensions; see [14] for a survey article. Investigating the first-principle model for the behavior on a cellular level one observes periodically stable solutions. Changes in model parameters do not modify the qualitative behavior of ion concentrations which can be related to so-called refractory times in which the cell is not yet able again to conduct another incoming signal. But they modify the length of these time windows. Therefore we decided to concentrate on simple mathematical models which focus directly on the length of refractory periods of larger cell compounds located in the AV node, and to neglect details about underlying cell dynamics which are of no interest for the diagnosis question.

We developed a mathematical multi-level model that is based on phenomenological observations. It builds on filters of a $n+1:n$ type, i.e., out of $n+1$ incoming signals exactly n continue their way through the AV node. The travel time can vary, e.g., increase linearly from signal to signal, until it is reset after $n+1$ incoming signals. From a first-principle point of view this corresponds to signals reaching the cell at different time points of the stable oscillations. From a modeling point of view these times are formulated as

$$x_i^l = x_j^{l-1} + \tau_{\text{const}}^l + k_{i,j}^l \tau_{\text{lin}}^l \tag{1}$$

where l denotes the blocking level, i the number of the outgoing signal, and j of the corresponding incoming signal. The times τ_{const}^l and τ_{lin}^l are model parameters and thus degrees of freedom. The counter $k_{i,j}^l$ increases n times, until it is reset to 0.

Assuming a regular (AFlu) behavior, the objective function is the deviation from simulated signaling behavior from the real measurements (R wave peak times). In an appropriate norm, one minimizes $\|x^{n_l} - \bar{x}\|$, where $\bar{x}$ are the measured R wave time points. The vector x^{n_l} of time points on the lowest level n_l is the result of a forward simulation of a regular, but unknown signal vector x^0 in the atrial chambers through n_l different blocking levels. Thus, the optimization problem can be summarized as follows.

$$\min \|x^{n_l} - \bar{x}\| \quad \text{subject to} \tag{2}$$

- bounds and integrality,
- signal x^{n_l} is the result of a forward simulation based on (1),
- the incoming signal x^0 (atrial activation) is regular.

The independent discrete variables are the number of blocking levels n_l and the value n which specifies the $n+1:n$ blocking behavior on level l. The independent continuous variables are the transition times τ_{const}^l, τ_{lin}^l and the distance between two atrial signals given by $\Delta x^0 := x_i^0 - x_{i-1}^0$. Dependent variables are the signal times x_i^l and the counters $k_{i,j}^l$.

In a more involved formulation it should be considered that the blocking behavior may change between $n+1:n$ and $n+2:n+1$ during the considered time horizon, because the blocking behavior depends also on the frequency of the incoming signals, and the atrial activation interval sometimes increases or decreases by a few milliseconds.

Challenges to the State-of-the-Art of Optimization

For fixed independent optimization variables (e.g., from an outer optimization loop) the dependent variables and hence all constraints and the objective function can be evaluated. It is however by no means trivial to write (2) in a closed compact form which would allow a solution with a standard MINLP solver. The difficulties comprehend the following.

- The dimensions of the vectors x^l depend on the optimization variable Δx^0. The smaller Δx^0, the more signals arrive at subsequent levels.
- In the objective function the jth measured R wave has to be compared to the jth signal that arrives at level n_l. Whether a signal $x_i^{n_l}$ is indeed preceded by $i-j+1$ blocked signals depends on the optimization variables, though. One modeling possibility is the introduction of a matrix $\Phi \in \{0,1\}^{m_1 \times m_2}$ which has entries $\Phi_{i,j}$ which are 1 if and only if this is the case. Then the objective function could be formulated as $\min \|\Phi x^{n_l} - \bar{x}\|$. How to efficiently formulate the logical constraints to determine Φ is unclear, especially given the dependence of m_1 on Δx^0.
- The values x_i^l depend on the counters $k_{i,j}^l \in \{0, \dots, n-1\}$, compare (1), which depend on the $n+1:n$ blocking behavior on level l. The very first signal of a block has an unknown offset, whereas all following values result from augmentation and resetting. Resetting may also occur if no input signal arrives for a certain time period. Introducing $k_{i,j}^l$ as optimization variables leads to a nonlinear right hand side in (1).
- The objective function is by no means uniquely specified. How many signals should be compared (too few can always be explained, too many usually come with variations in Δx^0), which

norm shall be used, how shall a difference in the number of simulated and measured signals be penalized?

Finding a good closed formulation is an open problem.

Solution Approach

We implemented a comprehensive approach for clinical decision support. It starts with automatic data generation with a smartphone. Our mobile app *HEAT* can extract R wave times from ECG pictures and from beep recordings of a heart rate monitor. The data is sent via a secure connection to our computation solver where the optimization problem is solved. Results are directly exported as PDF files and sent back to the app.

To solve the optimization problem we developed a tailored Branch and Bound algorithm. The effect of the currently chosen multi-level block is evaluated by forward simulation. We enumerate discrete variables and the continuous variables like Δx^0 on a ms grid. This allows to move in a particular order through the forward simulations and to prune subtrees when clinically motivated constraints or the logical implications are violated. The CPU times for three blocking levels and 22 R waves are in the range of 1–3 seconds on a standard notebook.

(Preliminary) Results and Outlook

Figure 1 shows an exemplary result for AFlu. It is one among 380 different ECGs from a benchmark database which we generated together with our clinical partners in Heidelberg. All of them come with the highest possible gold standard, expert-interpreted *intracardiac* measurements. Such measurements are only available in specific clinical circumstances, and not in everyday clinical practice. Our optimization-driven approach resulted in a sensitivity (percentage of AFlu cases diagnosed as such) of 0.81 and a specificity (percentage of AFib cases diagnosed as such) of 0.87. Looking at the main classification indicator of diagnostic tests, the Area Under the Curve value of the Receiver Operating Characteristic [4], our approach is classified as excellent with a value of 0.9. More details can be found in the paper [18] and in the PhD thesis [12].

The general approach is patented [17] and the startup company *mathe.medical GmbH* currently investigates possibilities of disseminating the technology into clinical practice.

Comment by Jeremi Mizerski, PhD, heart surgeon and senior researcher at the Interdisciplinary Center for Mathematical and Computational Modeling, Warsaw, Poland: The project covers the problem of the most common rhythm disturbances in human hearts. Atrial fibrillation and atrial flutter are benign in nature, but severe in consequences. Medical treatment of the first one is entirely different than the latter. The correct diagnosis and differentiation are hence of utmost importance due to a rising prevalence according to the aging of our societies. The power and the beauty of modern science arise from convergence and interdisciplinary approaches. This tight coupling of medical research and optimization is an excellent example for this.

3 Guiding the Quest for the Ablation Point

Clinical Background and Optimization Question

Another source of unhealthy cardiac arrhythmia are premature beats (PBs). They are common in patients with a structural heart disease, and often a catheter ablation is the method of choice. It results in non-conductivity of the ablated area on the heart surface, and hence modified spatio-temporal dynamics of the electrical activation potential. An obvious challenge is to identify quickly (to minimize side effects) and reliably the spatial site of origin of PBs.

3D electroanatomic mapping systems are increasingly applied for this task. They provide a graph representation of the heart surface, with 3D positions of the nodes. The ablation device cannot only scar tissue, but also measure activation. This is used to map the measurements of local activation times (LATs) onto nodes of the graph representation of the heart surface.

The optimization task is to choose the measurement nodes such that with a high probability only few measurements are necessary to find the node with the smallest LAT. At closer inspection, two different online optimization tasks arise. First, to estimate for the so-far available data which node is the current best guess. And second, if the associated uncertainty is still too high, which measurement node will reduce this uncertainty most.

Mathematical Modeling of the Dynamics and of the Optimization Problem

We use two major simplifications compared to the first-principle models of electrical activation mentioned above. First, we restrict our considerations to the given graph representation of the heart

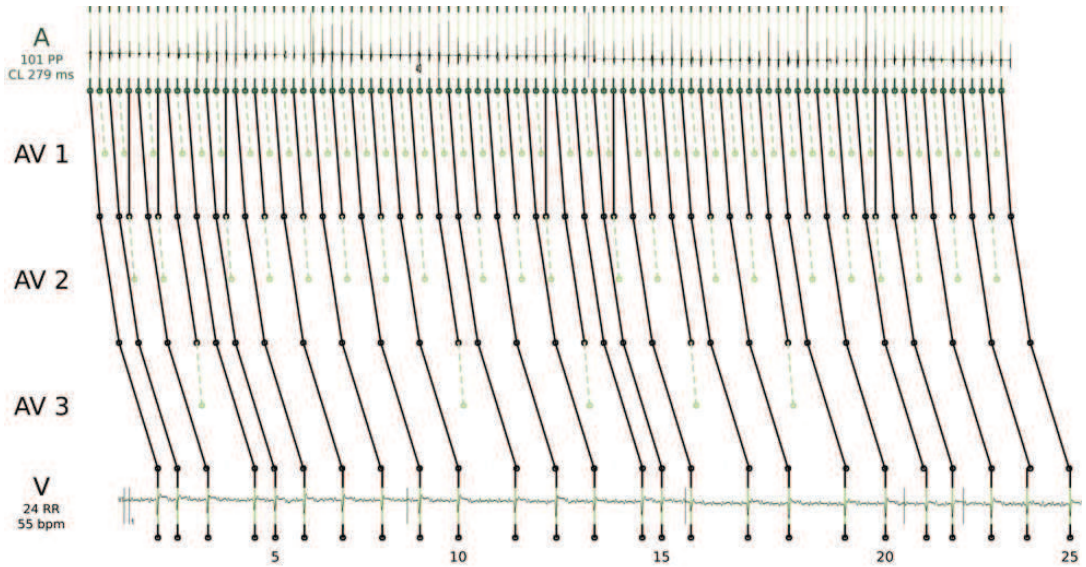

Figure 1. *Exemplary illustration of our optimization-driven solution approach from [12]. The input data, the observed ventricular (V) signals, is extracted from the surface ECG (bottom of figure). As a result of the optimization, three blocking levels with corresponding model parameters were calculated such that the calculated signal in the atrial chambers (A) is regular and the forward simulation in V is close to the measured data. The intracardiac measurements are shown for illustrative purposes (top).*

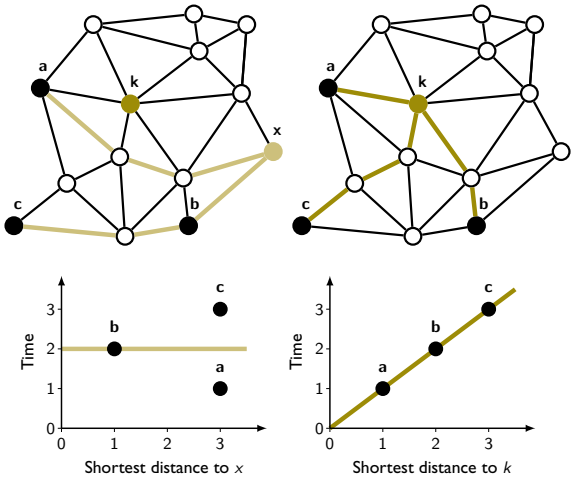

Figure 2. *(adapted from [20]). Linear regression approach for a simplified graph with equal weights 1 (distances and velocity). Left plots: The source is assumed to be node x which is different from the real source k (right plots). Comparing measured times (which correspond to the shortest distances to the real source k) with the shortest distances between x (or k) and the measurement nodes a, b, and c gives the plots in the bottom row. The linear regression on the left has a larger residual, indicating that x is not a good candidate for the unknown node of origin.*

surface. Second, we assume a constant and homogeneous velocity v of the electrical activation, i.e., the difference in LATs is proportional to the weights of the graph (Euclidean distances). The model for the LAT t_i on node i is given by

$$t_i = t_0 + \frac{d_{ik}}{v} \tag{3}$$

where t_0 is the earliest activation time (offset of the relative time measurements), d_{ik} is the shortest distance between node i and the (unknown) source node k, and v is the velocity. Figure 2 illustrates how linear regressions for different possible source nodes k can be compared.

Challenges to the State-of-the-Art of Optimization

There are two different interpretations of the optimization task. First, it can be seen as a parameter estimation problem where one "parameter", the source node $k \in V$, has a combinatorial nature. Second, it can be seen as a model discrimination problem where card(V) different models, corresponding to different assumptions on the source node, compete in explaining the measurements. The main challenge arises from the observation that the current best guess for the node of origin is not necessarily the measurement node that yields most information. Calculating this node can be interpreted as a special case of optimal experimental design.

Theoretical questions to be considered are the minimum number of LAT measurements to determine the node of origin in a deterministic setting via the metric dimension of the graph, and the connection to the bound from Carathéodory's Theorem. Algorithmic questions are related to the sequential setting (alternating sequence of measurements and optimizations) and the question when to switch from a "maximizing information gain" measurement node to a "maximizing probability of being the origin" measurement node. The clinical setting requires that solutions need to be calculated quickly, i.e., preferably in a fraction of a second.

Interestingly, a quite similar setting can be found in different source detection areas, such as contagion phenomena like drinking water pollution or influenza pandemics. The main difference in our setting is the sequential availability of only one measurement, and the possibility to verify that a node is the origin from a measurement (because of the specific shape of the LAT curve at the origin).

Comment by Prof. Dr. Eberhard Scholz, Head of Electrophysiology Lab at University Hospital Heidelberg, Germany: Imagine you are part of a rescue team searching for an avalanche victim, knowing that every elapsed minute could reduce the probability of survival. Would you dig here and there to find the victim by incidence or would you rather follow a systematic search routine using an avalanche transceiver to speed up the process? Interestingly, the localization of focal arrhythmia sources exhibits striking parallels to the situation sketched above. A small cluster of cardiac cells that is located at an unknown position of the heart muscle every now and then sends out an electric signal. However, in contrast to an experienced rescue team, most operators follow a heuristic search path to locate the arrhythmia source thereby loosing valuable time. The algorithm described in this article uses mathematical optimization to guide the operator on the shortest and hence quickest way to the desired location. This project gives an excellent example of how nicely mathematical optimization can provide decision support to physicians and might open the door to a new dimension of mapping algorithms in cardiac electrophysiology.

Solution Approach

We induce a shortest path metric on the graph and calculate all shortest distances between any two nodes a priori. Enumerating over all candidate nodes for the next measurement, this shortest path metric allows the comparison of linear regression results, as illustrated in Figure 2. The current best guess for the site of origin is the node which results in the smallest residual. The uncertainty can be estimated from the variance-covariance matrix of the parameter estimate.

In a first phase of the algorithm we compare the calculated variances (e.g., with respect to the offset t_0) and choose the node which results in the smallest value as the next measurement node. If the residuum of the current best guess is small enough, we choose in a second phase this best guess as the next measurement node.

(Preliminary) Results and Outlook

We performed a retrospective simulation study with real clinical data from 17 patients. The data was exported with the Carto 3D mapping system and consisted of the heart geometry, all measurements made by an operator, and the measured LAT. Our algorithm reduced the mean number of 42 ± 7.0 LAT measurements to 11 ± 0.89, indicating the huge potential for clinical improvement. More details can be found in the paper [20].

Future work will concentrate on less restrictive assumptions, i.e., more realistic non-constant velocities and consider also scarred tissue. A prospective study based on an integration of our algorithm into commercial mapping software is desirable.

4 Towards Optimized Blood Cancer Treatment

Clinical Background and Optimization Question

Oncologists and hematologists – and often also patients who are involved in the decision making process – need to choose between different available treatment alternatives for blood cancer. When applying chemotherapy, this choice includes combinations of different drugs and immune-boosters, scheduling of the treatments, and dosages. Additional issues like the balancing between positive (killing) impact on cancerous cells and negative (killing) impact on cells of the immune system highlight the importance to understand the body as a complex dynamical system. Obviously, simulations (what would happen if?) and optimization (what is the best treatment?) would be of tremendous value, if they were reliably applicable and personalized to the specific case. So far we have been focusing on three different kinds of decision support related to blood cancer, based on pressing practical questions and retrospective data from our clinical partners.

First, we have been investigating consolidation therapy for *acute myeloid leukemia* (AML) in adults. Acute means severe and sudden in onset (in contrast to *chronic*). Myeloid refers to where the differentiation and proliferation of blood cells takes place, here in the bone marrow. And leukemia refers to a disease affecting leukocytes, the white blood cells. The most important chemotherapeutic treatment consists of several cycles of induction therapy, followed by up to four cycles of consolidation therapy. They differ in the cytotoxic agents, dosages, and timings that are used, but also in their goals. Induction therapy tries to eradicate blasts (cancerous proliferating cells), while consolidation therapy starts when almost no more cancerous cells can be measured and tries to avoid a relapse. We have been studying the chemotherapeutical agent cytarabine, the most important component of AML treatment. Applying it, e.g., twice a day on days 1, 3, and 5 of a consolidation cycle, leads to a delayed decrease in the number of circulating (in the blood) white blood cells. This number can be seen as a surrogate for the strength of the immune system, a low number is associated with a high risk of severe side effects. The time period in which the number of white blood cells is below a certain threshold is called leukopenia. It is of high clinical relevance to be able to predict the length and depth of a leukopenia, and to avoid it, if possible, by modifying the treatment. Optimization can be applied to personalize mathematical models and predict the future dynamics of white blood cells. It can also be applied to calculate optimal experimental designs, i.e., measurement times and treatments that result in small confidence regions of parameter estimates. And, eventually, it can be used to optimize the chemotherapy schedules and dosages.

Second, we have been looking at consolidation therapy for *acute lymphoblastic leukemia* (ALL) in children. In ALL blasts occur in the lymph (the fluid circulating through the lymphatic system) and not in the bone marrow. The typical consolidation therapy for children is applied almost continuously, with orally administered drugs every two weeks over periods of several years. The clinical task consists in keeping the number of white blood cells in a target range by modifying dosage and timing of the drugs. This is usually extremely difficult, due to uncertainties, delays, and nonlinearities.

Third, we have been investigating *polycythemia vera* (PV) in adults. In PV, the bone marrow produces too many erythrocytes (red blood cells). The increase in the number of red blood cells leads to a thickening of the blood, which may be fatal if untreated due to a higher possibility of thromboembolic events. As the condition cannot (yet) be cured, treatment focuses on treating symptoms. It consists primarily of phlebotomy (blood letting), like in medieval times. Letting blood actually makes sense, as the time scale on which new red blood cells are being produced is weeks, while the blood plasma recovers within hours. As a result, the hematocrit (percentage of red blood cells in the blood) decreases after a phlebotomy. Possible side effects are well known to all blood donators and include dizziness, fatigue, or headaches. An intriguing clinical question is if personalized simulations and optimization can be used to schedule phlebotomies, e.g., to reduce the overall number, or to avoid collisions with important professional or private appointments. This could restore at least some of the patient's quality of life.

Mathematical Modeling of the Dynamics and of the Optimization Problem

There are many different levels on which cancer dynamics can be modeled [3], even for the special case of blood cancer where no spatial aspects of tumor growth need to be considered. In the specific blood cancer decision support tasks outlined above, an important ingredient is the mathematical modeling of hematopoiesis (the formation of blood cellular components by differentiation and maturation originating from hematopoietic stem cells). A rich literature exists with elaborated models with complex pathways, delays, game theoretic approaches to biological competition, or partial differential equations with cell maturation age as a spatial dimension. Such complicated models usually come with lots of model parameters, leading to identifiability issues, and are not well suited for optimization-driven decision support in real time.

We analyzed models that capture only the most important dynamics for blood cells and "agglomerate" different physiological effects into simplifying expressions. Often we also neglect dynamics of cancerous cells, as no measurable quantities are present and the focus is on the number of white or red blood cell counts, anyway. If drugs are involved, pharmacokinetics (PK; how an organism affects a drug) and pharmacodynamics (PD; how the drug affects the organism) are important. The mathematical models we have been using are compartment models. Cells of different maturation age that share common behavior, e.g., proliferating or differentiating, are clustered in compartments. In the AML project, we have been working with amounts x_1 and x_2 of the chemotherapeutic agent cytarabine in two PK compartments and with counts x_{pr} of proliferating cells, $x_{\text{tr},1}, \dots, x_{\text{tr},n_{\text{tr}}}$ of differentiating cells in n_{tr} transient compartments, and x_{ma} of mature, circulating white blood cells. The time dependent dosage of cytarabine is denoted by $u(t)$. The system of ordinary differential equations for hematopoiesis, pharmacokinetics and pharmacodynamics is based on the gold-standard model of chemotherapy–induced myelosuppression by [8],

$$\begin{aligned}
\dot{x}_1(t) &= -(k_{10} + k_{12})\, x_1(t) + k_{21}\, x_2(t) + u(t) && \text{(4a)}\\
\dot{x}_2(t) &= \quad k_{12}\, x_1(t) - k_{21}\, x_2(t) && \text{(4b)}\\
\dot{x}_{\text{pr}}(t) &= -k_{\text{tr}}\, x_{\text{pr}}(t) + F(x, k_{\text{tr}}, \gamma, B, \text{slope})\, x_{\text{pr}}(t) && \text{(4c)}\\
\dot{x}_{\text{tr},1}(t) &= \quad k_{\text{tr}}\, x_{\text{pr}}(t) - k_{\text{tr}}\, x_{\text{tr},1}(t) && \text{(4d)}\\
\dot{x}_{\text{tr},2}(t) &= \quad k_{\text{tr}}\, (x_{\text{tr},1}(t) - x_{\text{tr},2}(t)) && \text{(4e)}\\
&\dots\\
\dot{x}_{\text{tr},n_{\text{tr}}}(t) &= \quad k_{\text{tr}}\, (x_{\text{tr},n_{\text{tr}}-1}(t) - x_{\text{tr},n_{\text{tr}}}(t)) && \text{(4f)}\\
\dot{x}_{\text{ma}}(t) &= \quad k_{\text{tr}}\, x_{\text{tr},n_{\text{tr}}}(t) - k_{\text{ma}} x_{\text{ma}}(t) && \text{(4g)}
\end{aligned}$$

with a function F that models the pharmacodynamical effect of cytarabine and possible feedback mechanisms on the proliferating cells. The model parameters and initial values that can be used to personalize (4) are $p = (B, k_{\text{tr}}, \gamma, \text{slope}, x_{\text{pr}}(t_0), x_{\text{tr}}(t_0), x_{\text{ma}}(t_0))$, others are fixed to values obtained from independent in vitro studies. Figure 3 shows an exemplary result for clinical data. The model was fitted to white blood cell count measurements ($\approx x_{\text{ma}}$) from one consolidation cycle. Data from a second consolidation cycle was compared to a prediction as a cross-validation.

The models for ALL and PV are structurally similar. The predictive accuracy of (4) is quite high, and surprisingly stable against different modeling approaches of F. While this is a nice feature for simulations with a fixed treatment schedule, a model may not be appropriate to evaluate or even optimize dosage and timing of chemotherapy. The agglomerative nature of the mathematical models leads to a choice of model parameters that is not only personalized to the patient, but also to the applied schedule. It is ongoing (unpublished) work to identify mathematical models that are able to predict the outcome of different schedules well, e.g., on days 1, 2, and 3 instead of days 1, 3, 5. Obviously a validation with measurements is difficult, as identical repetitions with different treatments are not possible.

Assuming the availability of a validated personalized model, there is an abundance of optimization questions. Are high dosage or low dosage (or singular arc solutions) preferable? What impact has the timing of the chemotherapy? How important is the delay between two consolidation cycles? What is a good time to give growth stimu-

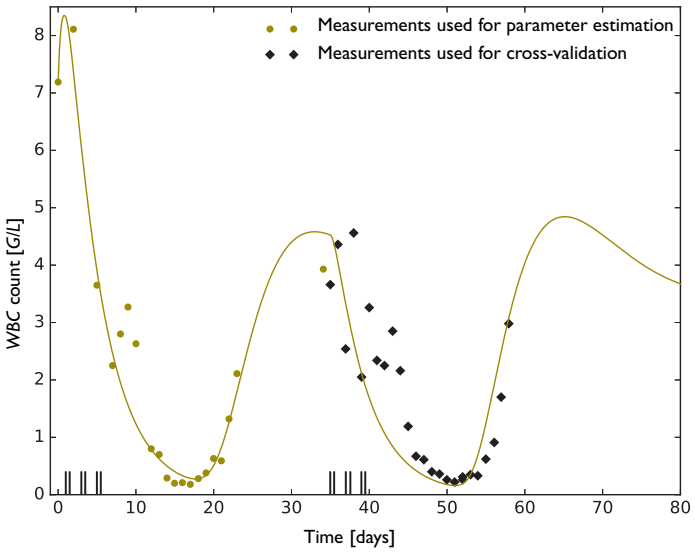

Figure 3. *Exemplary data from a patient suffering from acute myeloid leukemia from [11]. The chemotherapy schedule is indicated on the* x*-axis. The dots are white blood cell (WBC) counts. The trajectory shows the solution* $x_{ma}(\cdot)$ *of a personalized model of type* (4)*, based on the measurements from the first consolidation cycle. A cross-validation with the measurements from a second consolidation cycle illustrates the predictive accuracy of the mathematical model.*

lating factors (immune boosters)? How do all these modifications impact the relapse probability? How much do the results depend on the choice of the objective function? E.g., if x_{ma} shall be as large as possible, would you rather maximize $x_{\text{ma}}(t_f)$ or $\int_0^{t_f} x_{\text{ma}}(\tau)\,\mathrm{d}\tau$? Usually, the related objective functions are simple functions or integrals of the differential states and controls in (4).

Note that optimization can also be used in a clever way to evaluate the potential of modified treatments to check if all this is worth the trouble in the first place. In [7] we proposed to minimize *and* maximize the objective function. If the overall amount of chemotherapy is fixed in the constraints this allows to answer the question how much the objective function differs between the best and the worst treatment. If this percentage is not significantly larger than the associated level of uncertainty, the whole approach would be meaningless.

The availability of mathematical models allows another interesting experiment of thoughts. What, if some of the model parameters were control functions? In our abstracted problem formulations it is straightforward to replace constant parameters by time–dependent functions that can be optimized. Optimal solutions of these artificial control problems might give valuable hints for promising research directions for oncologists and pharmacologists, on which kinds of drugs are necessary to exploit the nonlinear dynamics.

Challenges to the State-of-the-Art of Optimization

The largest challenge originates from the assumed availability of a validated personalized model, which is obviously wrong. Personalized models with small standard deviations of parameters (and hence small uncertainty of predictions) require a certain number of measurements, which become available only during a treatment. Hence, new questions arise. What is a good treatment strategy when the uncertainty of the model is still very high? When should measurements be taken to get accurate parameter estimates? How does the chemotherapy or phlebotomy schedule influence the confidence region of subsequent state and parameter estimates? Does it make sense to give small dosages of chemotherapy when there is no direct medical benefit from it, but an indirect one from less uncertain models and more reliable predictions? The *dual control* task is to optimize the result of the treatment, which depends on a model prediction which depends also on the chosen treatment. The standard experimental design approach which minimizes a criterium of the variance-covariance matrix may lead to accurate state and parameter estimates, but in hindsight to a poor treatment.

We want to develop algorithms that suggest therapy details that lead both to an accurate real-time, patient-specific calibration of the mathematical model and to an optimal trade-off between infection risk and the therapy objectives. Although several formulations have been proposed in the literature, compare [10], there are many open theoretical and algorithmic questions in dual control.

Also of importance are the aspects of sparse controls (as patients want to and should spend as little time in hospital as possible), indicator constraints related to logical relations from law or clinical practice, and global optima.

Solution Approach

Our algorithms are based on a first-discretize, then-optimize approach (school of Georg Bock). We extended it in several research directions. On the top level we have been developing new concepts and algorithms for dual control [10] that sequentially use optimization to calculate personalized models, optimal measurement times, and optimal treatment choices. The building blocks are methods from moving horizon state and parameter estimation, robust optimal experimental design, model discrimination, nonlinear model predictive control, mixed-integer optimal control, and dual control. We developed a new sequential quadratic programming algorithm and solver for nonlinear optimization that exploits block structures as they also arise from optimal control problems, [9]. Our algorithmic developments are usually driven by clinical data, e.g., [11, 15]. To use information from other patients, we have been extending our methods to so-called mixed-effect models, which provide population parameters and their distributions, as well as individual parameters.

(Preliminary) Results and Outlook

In AML, optimally selecting 20–30 measurement times (almost daily) for one consolidation cycle by experimental design reduces the uncertainty of parameter estimates by approximately 50 % [11]. Optimal measurements from one consolidation cycle usually lead to standard deviations of parameters of approximately 10 % and thus to sufficiently accurate predictions of subsequent consolidation cycles.

These predictions can be used for optimization. In fact, they should be used for optimization. Even for quite simplistic models that do not take pharmacokinetics or effects on the immune system into account, the timing of the drug (with a fixed overall amount over the time horizon) has a large impact and may be non-intuitive, [7]. This result has been confirmed so far for the more complex processes associated with AML, ALL, and PV (all ongoing and unpublished work). We developed new mathematical models that have been personalized and cross-validated with clinical data. Using these

Comment by Prof. Dr. Thomas Fischer, Director of Department of Hematology and Oncology, University Hospital Magdeburg, Germany: Simulation and optimization may become a game changer for blood cancer treatment. For example, reliable predictions could be used for providing better care to AML patients receiving consolidation treatment. Analyzing the period of chemotherapy-induced profound leukopenia and applying optimized treatment schedules might enable prevention of severe infectious complications, sepsis, and thus delay to undergo subsequent treatment cycles. The density of chemotherapy cycles might be increased and thus deeper remissions and lower relapse rates might be achieved. This might ultimately translate into improved overall survival rates.

models we found a mathematical explanation for the superiority of dense chemotherapy treatments and a novel dynamic stratification for blood cancer patients. Our analysis indicates many cases where an optimized schedule might have avoided a leukopenia.

From my point of view, the proposed models, algorithms, and results are important steps on the way to a personalized real-time algorithm that exploits the nonlinear dynamics and combinatorial complexity of (blood) cancer treatments. Many questions and application in prospective clinical trials are still open, though.

Outlook: Optimization for Clinical Decision Training

Assuming the availability of mathematical models and tailored optimization algorithms, it is a natural question whether they can also be put to use in a training context. Simulators are already used to learn how to drive, fly, design energy-efficient houses, cars and airplanes, control large gas or water networks, determine ticket prices, or to some extent even how to construct financial regulations. While education in general will probably see many changes related to aspects like virtual reality or artificial intelligence fueling personalized trainers, also simulation and optimization can make important contributions.

Making clinical decisions can be seen as a special case of *Complex Problem Solving* (CPS), a research direction rooted in psychology. It is defined as a high-order cognitive process that involves decision making. The main intention is to understand how certain *variables* influence a solution process. In general, *personal and situational variables* are distinguished. The most typical and frequently analyzed personal variable is *intelligence*. How intelligence influences CPS is an ongoing debate. Other interesting personal variables are *working memory*, *amount of knowledge*, and *emotion regulation*. Situational variables such as the impact of *goal specificity and observation*, *feedback*, and *time constraints* have attracted less attention. Some questions on a more general level are: Why is it difficult for humans to judge and balance the consequences of decisions? What influence do cognitive representation, emotions, and feedback have?

There are two major methodological obstacles to answering these questions. How to measure performance in complex environments in an objective way? And how to systematically improve the performance by training of CPS competencies? I believe that a model-based optimization approach can help to overcome these obstacles. Although studying the interplay between optimization and

Comment by Prof. Dr. Joachim Funke, Department of Psychology, University of Heidelberg, Germany: Europe's decision makers in politics, industry, technology, and science are facing the challenge of increasing complexity. Consequently, the OECD PISA consortium defined Complex Problem Solving (CPS) competencies as an essential part of human education to help students prepare for life. For the analysis and training of CPS competencies a synergetic approach is necessary. On the one hand, the underlying cognitive processes need to be understood. On the other hand, powerful optimization methods are required to construct training environments with certain properties, to analyze the performance of participants, and to give optimization-based feedback.

For the first time in CPS research, the optimization of complex CPS tasks allows for an evaluation of a participant's performance in absolute units. This is a fundamental improvement because previous assessments since the advent of CPS tasks in the late 1970s made only statements about relative positions of participants (better or worse in terms of the analyzed population of participants). In addition, the evaluation of individual decisions in the course of multi-step decision making and new feedback and training opportunities are opening the door for future research in CPS. CPS researchers can start right away by addressing important questions such as the connection between cognitive representations, feedback, and emotion regulation with CPS performance. Results on analysis and training are expected to boost further research in this direction.

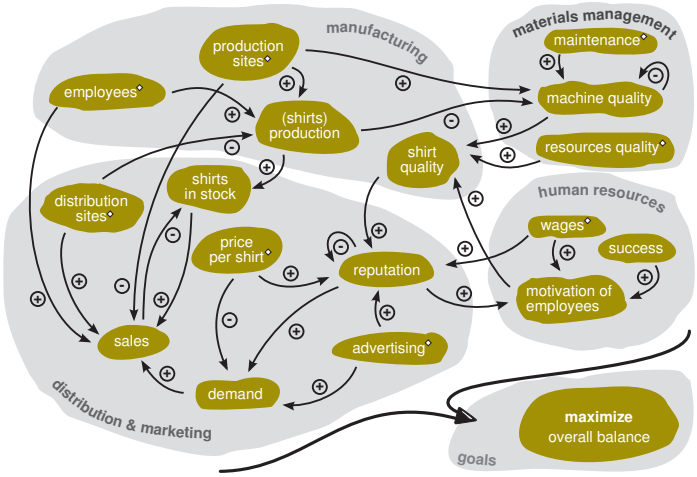

Figure 4. *Scheme of control (indicated with a diamond) and state variables in every round of the tailorshop mircoworld [5]. Arrows show dependencies. Some decisions are integer, some relations are nonlinear.*

clinical decision training is still ongoing work, we have achieved encouraging results from first studies with an economical setting. The "fruit fly" of CPS research, the *tailorshop* microworld, is illustrated in Figure 4.

In [16] we could address the first question of how to get a reliable performance indicator using a mathematical model of the *tailorshop*. An optimization is performed for every round of the participant's data, starting with exactly the same conditions as the participant. The differences between the optimal objective function values of rounds n_s and $n_s + 1$, the *How-much-is-still-possible*-function, indicate how good the performance is when compared to optimality. This measure is objective and yields a performance for each round n_s, thus taking the temporal evolution into account. This novel methodological approach has been combined with experimental studies, [1, 2, 16]. In [2] it was shown that participants who receive a *negative feedback* perform better than those who receive *positive feedback*. In [1, 16] the *ability to regulate emotion* was additionally considered. As a main result, an interaction between feedback and emotion regulation could be shown: participants with a high emotion regulation ability perform better when they get negative feedback, while those with a low ability to regulate their emotions perform badly for negative and well for positive feedback.

In [5, 6] we addressed the second question in an online study with two main hypotheses: optimization-based feedback in a training phase improves the performance in a subsequent performance phase and increases model knowledge compared to a control group. Both hypotheses could be shown to be true based on data from more than 100 participants. In particular, we used graphical visualizations of optimal decisions and of Lagrange multipliers corresponding to participants' decisions. The results indicate the large potential to improve learning using optimization, although questions like the best way to visualize results or local minima or measure the efficiency of the exploration of microworlds are still open. From an optimization point of view, the corresponding problems are challenging mixed-integer nonlinear programs with nonconvex relaxations. For the *tailorshop*, we developed a tailored decomposition and underestimation technique to get good bounds on the global optima.

Summary

We have presented several case studies that highlight the large impact that optimization may have on clinical decision support and training, and vice versa. Personalized medicine is often identified with (or defined as) stratification based on genomics. I believe in the large

IPCO XX, Ann Arbor, Michigan, USA

May 22–24, 2019. The IPCO conference is a forum for researchers and practitioners working on various aspects of integer programming and combinatorial optimization. The aim is to present recent developments in theory, computation, and applications. The scope of IPCO is viewed in a broad sense, to include algorithmic and structural results in integer programming and combinatorial optimization as well as revealing computational studies and novel applications of discrete optimization to practical problems.
Authors are invited to submit extended abstracts of their recent work by November 20, 2018 (23:59 EST); see the submission guidelines for more information. The Program Committee will select the papers to be presented on the basis of the submitted extended abstracts. Contributions are expected to be original, unpublished and not submitted to journals or conferences with proceedings before the notification date (January 31, 2019). Papers violating these requirements will not be considered by the Program Committee.
During the conference, approximately 33 papers will be presented in single-track sessions. Each lecture will be 30 minutes long. The proceedings will be published as a volume of Springer Lecture Notes in Computer Science. They will contain extended abstracts of all accepted submissions. It is expected that revised and extended versions will subsequently be submitted for publication in appropriate journals.
Program committee: ○ Amitabh Basu (Johns Hopkins U) ○ Pierre Bonami (IBM) ○ Jose Correa (U Chile) ○ Sanjeeb Dash (IBM) ○ Volker Kaibel (U Magdeburg) ○ Simge Küçükyavuz (U Washington) ○ Retsef Levi (MIT) ○ Jim Luedtke (U Wisconsin, Madison) ○ Andrea Lodi (Polytechnique Montreal, Chair) ○ Tom McCormick (UBC) ○ Ruth Misener (Imperial College) ○ Viswanath Nagarajan (U Michigan) ○ Marc Pfetsch (TU Darmstadt) ○ Martin Skutella (TU Berlin) ○ Gautier Stauffer (Kedge Business School) ○ Angelika Wiegele (Alpen-Adria U Klagenfurt)
Further Information and Submission Guidelines:
umich.edu/~ipco2019conf/

potential of a complementary *dynamic stratification*, based on fitted parameterized mathematical models.

The largest impact is related to the expected change of paradigm in clinical practice. This allows, e.g., a physician to first simulate the impact of his decisions on a computer and to consider optimized solutions. In the future, it will be the rare and unwanted exception that an important decision cannot be backed up by consultation of a model-driven decision support system or based upon a systematic model-driven training.

Sebastian Sager, Faculty of Mathematics, Otto-von-Guericke Universität Magdeburg, Germany. sager@ovgu.de

Bibliography

[1] C. Barth, "The impact of emotions on complex problem solving performance and ways of measuring this performance," Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 2010.
[2] C. Barth and J. Funke, "Negative affective environments improve complex solving performance," *Cognition and Emotion*, vol. 24, pp. 1259–1268, 2010.
[3] J. Clairambault, "Modelling physiological and pharmacological control on cell proliferation to optimise cancer treatments," *Mathematical Modelling of Natural Phenomena*, vol. 4, no. 3, pp. 12–67, 2009.
[4] I. G. Duncan, *Healthcare Risk Adjustment and Predicitve Modeling*. ACTEX Publications, 2011.
[5] M. Engelhart, "Optimization-based Analysis and Training of Human Decision Making," Ph.D. dissertation, University Heidelberg, 2015. [Online]. Available: https://mathopt.de/PUBLICATIONS/Engelhart2015.pdf
[6] M. Engelhart, J. Funke, and S. Sager, "A web-based feedback study on optimization-based training and analysis of human decision making," *Journal of Dynamic Decision Making*, vol. 3, no. 1, 2017.
[7] M. Engelhart, D. Lebiedz, and S. Sager, "Optimal control for cancer chemotherapy ODE models: Potential of optimal schedules and choice of objective function," *Mathematical Biosciences*, vol. 229, no. 1, pp. 123–134, 2011. [Online]. Available: https://mathopt.de/PUBLICATIONS/Engelhart2011.pdf
[8] L. Friberg, A. Henningsson, H. Maas, L. Nguyen, and M. Karlsson, "Model of chemotherapy-induced myelosuppression with parameter consistency across drugs," *Journal of Clinical Oncology*, vol. 20, pp. 4713–4721, 2002.
[9] D. Janka, C. Kirches, S. Sager, and A. Wächter, "An SR1/BFGS SQP algorithm for nonconvex nonlinear programs with block-diagonal Hessian matrix," *Mathematical Progamming Computation*, vol. 8, no. 4, pp. 435–459, 2016.
[10] F. Jost, S. Sager, and T. Le, "A feedback optimal control algorithm with optimal measurement time points," *Processes*, vol. 5, no. 10, pp. 1–19, 2017. [Online]. Available: http://www.mdpi.com/2227-9717/5/1/10
[11] F. Jost, K. Rinke, T. Fischer, E. Schalk, and S. Sager, "Optimum experimental design for patient specific mathematical leukopenia models," in *Proceedings of the Foundations of Systems Biology in Engineering (FOSBE) Conference*, vol. 49, no. 26, Magdeburg, Germany. Elsevier, 2016, pp. 344–349.
[12] F. Kehrle, "Inverse simulation for cardiac arrhythmia," Ph.D. dissertation, Otto-von-Guericke University Magdeburg, 2018. [Online]. Available: https://mathopt.de/PUBLICATIONS/Kehrle2018.pdf
[13] D. Noble, "A modification of the Hodgkin–Huxley equations applicable to purkinje fibre action and pacemaker potentials," *The Journal of Physiology*, vol. 160, no. 2, pp. 317–352, 1962.
[14] D. Noble, A. Garny, and P. J. Noble, "How the Hodgkin–Huxley equations inspired the cardiac physiome project," *The Journal of Physiology*, vol. 590, no. Pt 11, pp. 2613–2628, 2012.
[15] K. Rinke, F. Jost, R. Findeisen, T. Fischer, R. Bartsch, E. Schalk, and S. Sager, "Parameter estimation for leukocyte dynamics after chemotherapy," in *Proceedings of the Foundations of Systems Biology in Engineering (FOSBE) Conference*, vol. 49, no. 26, Magdeburg, Germany. Elsevier, 2016, pp. 44–49.
[16] S. Sager, C. Barth, H. Diedam, M. Engelhart, and J. Funke, "Optimization as an analysis tool for human complex problem solving," *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 936–959, 2011. [Online]. Available: https://mathopt.de/PUBLICATIONS/Sager2011c.pdf
[17] E. Scholz, S. Sager, and H. Katus, "A system and computer program product for automatically distinguishing atrial flutter from atrial fibrillation," European Patent EP2 757 940B1, 2016. [Online]. Available: https://patents.google.com/patent/EP2757940B1
[18] E. P. Scholz, F. Kehrle, S. Vossel, A. Hess, E. Zitron, H. A. Katus, and S. Sager, "Discriminating atrial flutter from atrial fibrillation using a multilevel model of atrioventricular conduction." *Heart Rhythm Journal*, vol. 11, no. 5, pp. 877–884, 2014.
[19] A. Shiyovich, A. Wolak, L. Yacobovich, A. Grosbard, and A. Katz, "Accuracy of diagnosing atrial flutter and atrial fibrillation from a surface electrocardiogram by hospital physicians: Analysis of data from internal medicine departments," *The American Journal of the Medical Sciences*, vol. 340, no. 4, pp. 271–275, 2010.
[20] T. Weber, H. Katus, S. Sager, and E. Scholz, "Novel algorithm for accelerated electroanatomical mapping and prediction of earliest activation of focal cardiac arrhythmias using mathematical optimization," *Heart Rhythm*, vol. 14, no. 6, pp. 875–882, 2017. [Online]. Available: https://mathopt.de/PUBLICATIONS/Weber2017.pdf

IMPRINT
Editor: Volker Kaibel, Institut für Mathematische Optimierung, Otto-von-Guericke Universität Magdeburg, Universitätsplatz 2, 39108 Magdeburg, Germany. kaibel@ovgu.de ■ **Co-Editors:** Samuel Burer, Department of Management Sciences, The University of Iowa, Iowa City, IA 52242-1994, USA. samuel-burer@uiowa.edu ■ Jeff Linderoth, Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706-1572, USA. linderoth@wisc.edu ■ **Founding Editor:** Donald W. Hearn ■ Published by the Mathematical Optimization Society ■ Design and typesetting by Christoph Eyrich, Berlin, Germany. optima@0x45.de ■ Printed by Oktoberdruck AG, Berlin, Germany.