



On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints

Yair Censor, Wei Chen, Patrick Louis Combettes, Ran Davidi, Gabor T. Herman

► To cite this version:

Yair Censor, Wei Chen, Patrick Louis Combettes, Ran Davidi, Gabor T. Herman. On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints. Computational Optimization and Applications, 2012, 51, pp.1065 - 1088. 10.1007/s10589-011-9401-7 . hal-01098022

HAL Id: hal-01098022

<https://hal.science/hal-01098022>

Submitted on 22 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints

Yair Censor · Wei Chen · Patrick L. Combettes ·
Ran Davidi · Gabor T. Herman

Received: 22 December 2009 / Published online: 17 March 2011
© Springer Science+Business Media, LLC 2011

Abstract The effectiveness of projection methods for solving systems of linear inequalities is investigated. It is shown that they often have a computational advantage over alternatives that have been proposed for solving the same problem and that this makes them successful in many real-world applications. This is supported by experimental evidence provided in this paper on problems of various sizes (up to tens of thousands of unknowns satisfying up to hundreds of thousands of constraints) and by a discussion of the demonstrated efficacy of projection methods in numerous scientific publications and commercial patents (dealing with problems that can have over a billion unknowns and a similar number of constraints).

Keywords Projection methods · Convex feasibility problems · Numerical evaluation · Optimization · Linear inequalities · Sparse matrices

1 Introduction

Projection methods were first used to solve systems of linear equations in Euclidean spaces in the 1930s [35, 62] and were subsequently extended to systems of linear inequalities in [1, 70, 71]. The basic step in these early algorithms consists of a projection onto an affine subspace or a half-space. Modern projection methods are much

Y. Censor (✉)

Department of Mathematics, University of Haifa, Mt. Carmel, Haifa 31905, Israel
e-mail: yair@math.haifa.ac.il

W. Chen · R. Davidi · G.T. Herman

Department of Computer Science, The Graduate Center, City University of New York,
365 Fifth Avenue, New York, NY 10016, USA

P.L. Combettes

Laboratoire Jacques-Louis Lions – UMR CNRS 7598, UPMC Université Paris 06, 75005 Paris,
France

more sophisticated [7–10, 19, 29, 38–40, 47, 48, 65] and they can solve the general *convex feasibility problem* of finding a point in the intersection of a family of closed convex sets in a Hilbert space. In such formulations, each set can be specified in various forms, e.g., as the fixed point set of a nonexpansive operator, the set of zeros of a maximal monotone operator, the set of solutions to a convex inequality, or the set of solutions to an equilibrium problem. Projection methods can have various algorithmic structures (some of which are particularly suitable for parallel computing) and they also possess desirable convergence properties and good initial behavior patterns [8, 29, 37–39, 58, 76]. The main advantage of projection methods, which makes them successful in many real-world applications, is computational. They have the ability to handle some huge-size problems of dimensions beyond which more sophisticated methods cease to be efficient or even applicable due to memory requirements. This is so because the building blocks of a projection algorithm are the projections onto the given individual sets, which are assumed to be easy to perform, and the algorithmic structure is either sequential or simultaneous, or in-between, as in the block-iterative projection methods or in the more recent string-averaging projection methods. The number of sets used simultaneously in each iteration in block-iterative methods and the number and lengths of strings used in each iteration in string-averaging methods are variable, which provides great flexibility in matching the implementation of the algorithm with the parallel architecture at hand; for block-iterative methods see, e.g., [2, 10, 18, 26, 39, 45, 49, 53, 65, 73] and for string-averaging methods see, e.g., [12, 20, 25, 27, 28, 44, 74, 77].

The convex feasibility formalism is at the core of the modeling of many problems in various areas of mathematics and the physical sciences; see [36, 37] and references therein. Over the past four decades, it has been used to model significant real-world problems in sensor networks [15], in radiation therapy treatment planning [23, 32, 59], in resolution enhancement [30], in wavelet-based denoising [34], in antenna design [56], in computerized tomography [58], in materials science [63], in watermarking [66], in data compression [68], in demosaicking [69], in magnetic resonance imaging [78], in holography [79], in color imaging [80], in optics and neural networks [81], in graph matching [82] and in adaptive filtering [84], to name but a few. In these—and numerous other—problems, projection methods have been used to solve the underlying convex feasibility problems.

We focus on the important subclass of convex feasibility problems in which finitely many sets are given and each of them is specified by a linear equality or inequality in the Euclidean space \mathbb{R}^N . For such problems, which arise in many important applications [32, 36, 58, 59], alternatives to projection methods are available (see, e.g., [3, 55] and the references therein), and it is therefore legitimate to ask whether projection methods are competitive.

In this paper we address this question and show that projection methods are indeed very competitive for some inverse problems in the environment of linear inequality constraints. In Sect. 2 we discuss their comparative performance for four different kinds of problems. In Sect. 3 we give some examples of their use in real-world applications from the research and the patent literature. Finally, we present our conclusions in Sect. 4.

2 Comparisons

2.1 Examples from traditional linear programming (LP)

In a recent paper [55] the author applies several projection methods to the Netlib LP collection of test problems and observes that the solution times of those methods are relatively long when compared to an interior point code. In this paper we point out that there are better performing projection methods than those used in [55]. In particular, for a Netlib test problem such as ISRAEL with a full dimensional feasible set, projection methods converge quite quickly. Other examples in which projection methods perform quite well will be given.

Some overarching statements in [55] about the non-effectiveness and uncompetitiveness of projection methods¹ are based on applications of some projection methods of the author's choice to simple 2-set formulations of problems only from the Netlib/CUTEr LP problem set, which is a standard test set and reflects many of the difficulties affecting real-world linear programming (LP). However, it is somewhat old and therefore the problems it contains are not huge by modern standards, but they aptly represent some of the pitfalls that are common in LP. In comparison, as indicated in the Introduction, projection methods have been used to solve highly non-linear complex problems involving a very large number of sets. Therefore, results based on the Netlib/CUTEr LP problem set are insufficient to draw general conclusions. It will also be shown in this subsection that the experiments reported in [55] use suboptimal versions of projection methods.

We start with a specific example: the problem named ISRAEL from the LP problem set Netlib/CUTEr, which has 142 variables and 316 linear inequality constraints. Gould [55] reports that none of the four variants of the projection method discussed in that paper succeeds in finding a feasible point in no more than 1,000,000 iterations, which require in total about 300 seconds of computer time. However, we now present a published projection method, called ART3+, that (starting with the same initial point as used by Gould [55]) finds a feasible point for ISRAEL in only 8,258 iterations, and requires only 0.54 millisecond to get there on an Intel Xeon 2.66 GHz processor, 16 Gbyte memory, 64 bit workstation (the time reported by Gould for his LSQP approach is 0.3 seconds).

ART3+ was introduced in [59] to solve problems of the form:

$$\text{find an } x \in \mathbb{R}^N \text{ such that } c \leq Ax \leq d, \quad (1)$$

where $c \in (\mathbb{R} \cup \{-\infty\})^M$, $d \in (\mathbb{R} \cup \{+\infty\})^M$ and, for $1 \leq i \leq M$, $c_i < d_i$ and at least one of c_i and d_i is finite. ISRAEL is a problem of this type with $M = 316$ and $N = 142$. ART3+ is an iterative procedure, in which only a single inequality-constraint-pair is treated in one iteration. Denoting by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ the Euclidean inner product and norm, respectively, the iterate $x^{(n+1)}$ is obtained from $x^{(n)}$, given

¹ Such as in the Abstract of [55]: “Unfortunately, particularly given the large literature which might make one think otherwise, numerical tests indicate that in general none of the variants [of projection methods for solving convex feasibility problems] considered are especially effective or competitive with more sophisticated alternatives.”

that the inequality-constraint-pair selected in this iteration is $c_{i(n)} \leq \langle a_{i(n)}, x \rangle \leq d_{i(n)}$, by

$$x^{(n+1)} = x^{(n)} - \begin{cases} \frac{\langle a_{i(n)}, x^{(n)} \rangle - (d_{i(n)} + c_{i(n)})/2}{\|a_{i(n)}\|^2} a_{i(n)}, \\ \quad \text{if } \langle a_{i(n)}, x^{(n)} \rangle < c_{i(n)} - (d_{i(n)} - c_{i(n)})/2, \\ 2 \frac{\langle a_{i(n)}, x^{(n)} \rangle - c_{i(n)}}{\|a_{i(n)}\|^2} a_{i(n)}, \\ \quad \text{if } c_{i(n)} - (d_{i(n)} - c_{i(n)})/2 \leq \langle a_{i(n)}, x^{(n)} \rangle < c_{i(n)}, \\ 2 \frac{\langle a_{i(n)}, x^{(n)} \rangle - d_{i(n)}}{\|a_{i(n)}\|^2} a_{i(n)}, \\ \quad \text{if } d_{i(n)} < \langle a_{i(n)}, x^{(n)} \rangle \leq d_{i(n)} + (d_{i(n)} - c_{i(n)})/2, \\ \frac{\langle a_{i(n)}, x^{(n)} \rangle - (d_{i(n)} + c_{i(n)})/2}{\|a_{i(n)}\|^2} a_{i(n)}, \\ \quad \text{if } d_{i(n)} + (d_{i(n)} - c_{i(n)})/2 < \langle a_{i(n)}, x^{(n)} \rangle. \end{cases} \quad (2)$$

If none of the conditions in (2) is satisfied, it means that $c_{i(n)} \leq \langle a_{i(n)}, x^{(n)} \rangle \leq d_{i(n)}$ and we set $x^{(n+1)} = x^{(n)}$. Note that in all cases $x^{(n+1)}$ is obtained from $x^{(n)}$ by adding a multiple of $a_{i(n)}$ to it, and $x^{(n+1)}$ is in the convex set determined by the inequality-constraint-pair used in the iteration. The control sequence $(i(n))_{n \in \mathbb{N}}$ by which the algorithm ART3+ selects the constraint-pairs is as follows. An ordered set of constraint-pairs is initialized to contain all constraint-pairs. We repeatedly cycle through this ordered set, but remove from it the currently-picked constraint-pair if it is satisfied by the current iterate $x^{(n)}$. If the ordered set becomes empty, it is filled up by the complete set of constraint-pairs, unless all the constraints are satisfied, in which case ART3+ terminates. It is proved in [59] that, provided that the set of all $x \in \mathbb{R}^N$ that satisfies all the constraints is full-dimensional (this is the case for ISRAEL), ART3+ will terminate in a finite number of iterations, irrespective of how $x^{(0)} \in \mathbb{R}^N$ is chosen. For ISRAEL, when we selected $x^{(0)}$ to be the same vector as used by Gould [55] for the initial point, we obtained a solution satisfying all the constraints over 500 times faster than the time reported by Gould using a different computer.

We point out that the Netlib/CUTEr LP problem set has very few examples for which (1) is the natural representation and of these ISRAEL is the only one for which the feasible set is full-dimensional. This by itself demonstrates the insufficiency of the Netlib/CUTEr LP problem set for benchmarking projection methods in general: problems of the kind (1) with full-dimensional feasible sets appear in many applications, of which we have some examples below. Examples of the form (1) in the Netlib/CUTEr LP problem set for which the feasible set is not full-dimensional include FIT1P and FIT2P; our experience with these examples is the same as reported in [55] in that we could not identify a projection-based approach for them that is competitive with standard interior-point algorithms. Thus, there are LP problems for which projection methods are not the methods of choice. On the other hand, as shown in Sect. 2.3 below, there do exist realistic classes of LP problems for which projection methods, properly formulated, significantly outperform standard interior-point algo-

arithms. Thus, the evidence presented in [55] should not be interpreted as meaning that projection methods are always inferior for LP problems.

The reason for the feasible set not being full dimensional for many problems in the Netlib/CUTEr set is that they include constraints of the form $Ax = b$. In the kind of problems from practical applications with which the authors of this paper work, the vector b is not known exactly; for example, because we have only physical measurements that are supposed to approximate it. In such applications, there is no point in insisting that an x satisfying $Ax = b$ be found, the mathematical task described in (1) with $c < d$, or something of similar nature, is more appropriate. This typically leads to feasibility problems with full-dimensional solution sets, a point of which can be found in a finite number of steps by a suitable projection method such as ART3+. It is of course possible that “finite” is very large. For example, as illustrated in Table 2 of [59], the execution time required to find a feasible point increases as the bound d approaches the bound c from above, thereby reducing the volume of the feasible region. However, as will be seen in the examples that follow, the actual execution times on typical problems are acceptable and compare favorably with those of alternative approaches.

The numerical experiments provided in [55] focus exclusively on the problem of solving a linear system of equations under box constraints, namely

$$\text{find an } x \in \mathbb{R}^N, \text{ such that } \begin{cases} Ax = b, \\ x \in \times_{j=1}^N [v_j, w_j], \end{cases} \quad (3)$$

where $v = (v_j)_{1 \leq j \leq N}$ and $w = (w_j)_{1 \leq j \leq N}$ are given vectors, $A \in \mathbb{R}^{M \times N}$ ($M \leq N$) has full rank, $b \in \mathbb{R}^M$, and the problem is assumed to be feasible. Problems, such as the one specified in (1), are forced into this form in [55] by the introduction of slack variables. Quite differently from the ART3+ projection method approach discussed above, it is then perceived that we have just two convex sets (the one determined by the equalities and the one determined by the box constraints) and that a projection method should consist of alternating projections onto these particular convex sets. This is indeed one way of solving LP problems by a projection method, but (as we have already demonstrated by the ISRAEL example) there may be other, more efficient projection methods available for particular classes of problems. Furthermore, we now show that, even in the restrictive setting adopted in [55], there are projection algorithms implemented with standard relaxation strategies that perform much better than those identified as projection methods in [55].

Let us denote by P_1 and P_2 the projection operators onto the closed affine subspace $S_1 = \{x \in \mathbb{R}^N \mid Ax = b\}$ and the closed convex set $S_2 = \times_{j=1}^N [v_j, w_j]$, respectively. The first operator is defined by (see Chap. 8 of [47])

$$P_1: x \mapsto x - A^\top (AA^\top)^{-1} (Ax - b), \quad (4)$$

where A^\top denotes the transpose of A . This transformation can be implemented in various fashions. For instance, in many signal and image processing problems, the matrix A is block-circulant and hence diagonalized by the discrete Fourier transform

operator, which leads to a very efficient implementation of P_1 via the Fast Fourier Transform (FFT) [4]. Here, we adopt a QR decomposition approach. Let

$$A^\top = \begin{bmatrix} Q_{11} & Q_{12} \end{bmatrix} \begin{bmatrix} R_{11} \\ 0 \end{bmatrix} \quad (5)$$

be the QR decomposition of A^\top , where R_{11} is an $M \times M$ invertible upper triangular matrix [52]. Then (4) yields

$$P_1 : x \mapsto x - Q_{11} \left(R_{11}^\top \right)^{-1} (Ax - b). \quad (6)$$

On the other hand, the projection $P_2x = (\pi_j)_{1 \leq j \leq N}$ of a vector $x = (x_j)_{1 \leq j \leq N}$ onto S_2 is obtained through a simple clipping of its components, i.e., for every $j \in \{1, \dots, N\}$, $\pi_j = \min\{\max\{x_j, v_j\}, w_j\}$ (see Sect. 8.1.1 of [16]).

Two standard projection methods to solve (3) are the alternating projection method

$$x^{(0)} \in \mathbb{R}^N \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = x^{(n)} + \lambda_n (P_1 P_2 x^{(n)} - x^{(n)}) \quad (7)$$

and the parallel projection method

$$x^{(0)} \in \mathbb{R}^N \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = x^{(n)} + \lambda_n \left(\frac{P_1 x^{(n)} + P_2 x^{(n)}}{2} - x^{(n)} \right), \quad (8)$$

where $(\lambda_n)_{n \in \mathbb{N}}$ is a sequence of strictly positive relaxation parameters. If $\lambda_n \equiv 1$ in (7), we obtain the popular *Projection Onto Convex Sets* (POCS) algorithm [36, 83]:

$$x^{(0)} \in \mathbb{R}^N \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = P_1 P_2 x^{(n)}. \quad (9)$$

The convergence of a sequence $(x^{(n)})_{n \in \mathbb{N}}$ thus constructed to a point in $S_1 \cap S_2$ was established in [17]. On the other hand, if $\lambda_n \equiv 1$ in (8), we obtain the *Parallel Projection Method* (PPM):

$$x^{(0)} \in \mathbb{R}^N \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = \frac{P_1 x^{(n)} + P_2 x^{(n)}}{2}. \quad (10)$$

The convergence of a sequence $(x^{(n)})_{n \in \mathbb{N}}$ thus constructed to a point in $S_1 \cap S_2$ was established in [5], see also [6]. In [55], (7) and (8) are used, together with variants featuring a construction of λ_n at iteration n resulting from a line search procedure and without closed-form expression. However, as the numerical results of [55] show, these relaxation schemes do not lead to significantly better convergence profiles than those obtained with the unrelaxed algorithms POCS (9) and PPM (10). In addition, nothing is said regarding the convergence of (7) and (8) with such relaxation schemes.

The potentially slow convergence of projections methods has long been recognized [42, 57, 70] and remedies have been proposed to address this problem in the form of adapted relaxation strategies that guarantee convergence. In the case of (7), it was shown in [10] that any sequence generated by the *Extrapolated Alternating Projection Method* (EAPM)

$$x^{(0)} \in S_1 \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = x^{(n)} + \rho K_n (P_1 P_2 x^{(n)} - x^{(n)}),$$

$$\text{where } 0 < \rho < 2 \text{ and } K_n = \begin{cases} \frac{\|P_2 x^{(n)} - x^{(n)}\|^2}{\|P_1 P_2 x^{(n)} - x^{(n)}\|^2}, & \text{if } x^{(n)} \notin S_2, \\ 1, & \text{if } x^{(n)} \in S_2, \end{cases} \quad (11)$$

produces a fast algorithm that converges to a solution to (3). (The word “fast” here is used in its standard intuitive sense, its use is clearly justified by the plots in Fig. 1 and numerical experiments presented in [10, 13, 39, 70, 75, 76].) This type of extrapolation scheme, which exploits the fact that S_1 is an affine subspace, actually goes back to the classical work of [57]. It has been further investigated in [11, 39] and has been extended recently to a general block-iterative scheme in [10]. Acceleration methods have also been devised for the parallel algorithm (8). Thus, the convergence of the sequence produced by the *Extrapolated Parallel Projection Method* (EPPM)

$$x^{(0)} \in \mathbb{R}^N \quad \text{and} \quad (\forall n \in \mathbb{N}) \quad x^{(n+1)} = x^{(n)} + \chi L_n \left(\frac{P_1 x^{(n)} + P_2 x^{(n)}}{2} - x^{(n)} \right),$$

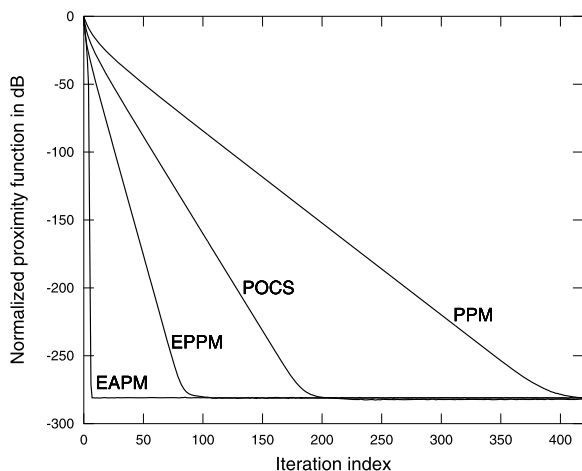
where $0 < \chi < 2$ and

$$L_n = \begin{cases} 2 \frac{\|P_1 x^{(n)} - x^{(n)}\|^2 + \|P_2 x^{(n)} - x^{(n)}\|^2}{\|P_1 x^{(n)} + P_2 x^{(n)} - 2x^{(n)}\|^2}, & \text{if } x^{(n)} \notin S_1 \cap S_2, \\ 1, & \text{if } x^{(n)} \in S_1 \cap S_2, \end{cases} \quad (12)$$

to a solution of (3) was established in [38]. This type of parallel extrapolated method goes back to [70] and [75], and it has been refined or generalized in several places [39, 65, 73].

Generally speaking, the numerical superiority of extrapolated projection methods such as EPPM and EAPM over standard projection methods such as POCS and PPM has been consistently observed in various types of problems in areas as diverse as partial differential equations, signal deconvolution, linear numerical analysis, or image restoration [10, 37, 39, 43, 53, 70, 75, 76] (see also [11, 64] for further theoretical justifications). By way of illustration, in Fig. 1, we compare the numerical performance of POCS (9), PPM (10), EAPM (11), and EPPM (12) for problems of

Fig. 1 Average performance of the algorithms when $M \times N = 600 \times 1000$



size $M \times N = 600 \times 1000$. As in [10, 37, 39], the performance of the algorithms is measured by the decibel (dB) values of the normalized proximity function, which is evaluated at the n th iterate $x^{(n)}$ by

$$10 \log_{10} \left(\frac{\|P_1 x^{(n)} - x^{(n)}\|^2 + \|P_2 x^{(n)} - x^{(n)}\|^2}{\|P_1 x^{(0)} - x^{(0)}\|^2 + \|P_2 x^{(0)} - x^{(0)}\|^2} \right). \quad (13)$$

This comparison is relevant because the computational load of each iteration resides essentially in the computation of the projection onto S_1 and it is therefore roughly the same for all four algorithms. The results are averaged over 20 runs of the algorithms initialized with $x^{(0)} = P_1 0$ and $\rho = \chi = 1.9$. In each run a matrix $A \in [-0.5, 0.5]^{M \times N}$ and a vector $x \in [0, 1]^N$ are randomly generated. The vector $b = Ax$ is then constructed so as to obtain a feasible problem using $v_j \equiv 0$ and $w_j \equiv 1$ in (3). As in [55] and many other studies, we observe that POCS is faster than PPM. However, EPPM is faster than POCS and EAPM is clearly the best method: on the average, it is about 60 times faster than PPM, 30 times faster than POCS, and it achieves full convergence in just 7 iterations. In addition, convergence to a feasible solution is guaranteed by the theory and the expression of the extrapolation parameter K_n in (11) is explicit and it requires no additional computation. It is argued in Sect. 5 of [55] that “there is a significant difference between random and real-life problems (similar observations have been made for linear equations, where random problems tend to be well-conditioned [...] and thus often easier to solve than those from applications).” Let us observe that random matrices do show up in many real-life problems, see [41, 85] and the references therein. In addition, as shown in Fig. 2, the qualitative behavior of the algorithms in the presence of poor conditioning is quite comparable to that observed in Fig. 1 (for the experiments of Fig. 2, the condition numbers vary from 3×10^4 to 3.5×10^4).

We have consistently observed this type of performance for problems of various sizes. For instance, we report in Fig. 3 on the same experiment as above on problems

Fig. 2 Average performance of the algorithms when $M \times N = 600 \times 1000$ and condition numbers vary from 3×10^4 to 3.5×10^4

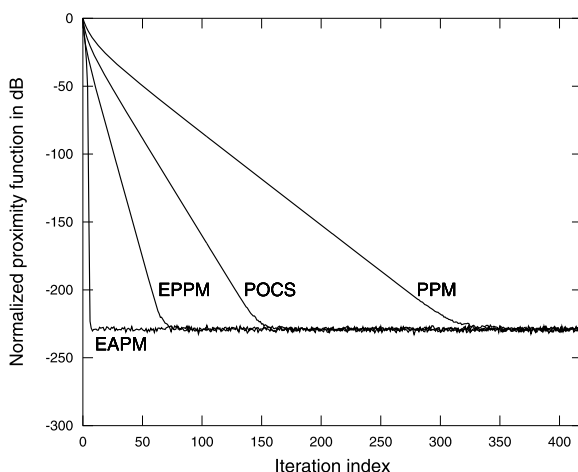
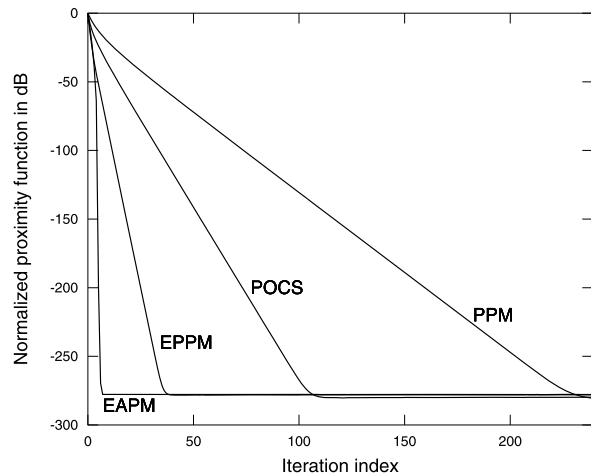


Fig. 3 Average performance of the algorithms when $M \times N = 3000 \times 7000$



of size $M \times N = 3000 \times 7000$. Here EAPM is about 45 times faster than PPM, 22 times faster than POCS, and full convergence is achieved in just 5 iterations.

These experiments, together with the above cited results of [10, 11, 37, 39, 43, 53, 64, 70, 75, 76], indicate that the results in [55] on the speed of convergence of POCS and PPM (and the variants proposed there featuring modest speed-up factors and lacking a formal convergence analysis) correspond to a suboptimal implementation of projection methods and are not representative of their performance, since significant improvements can be achieved by appropriate relaxations. Of course, this observation does not rule out the possibility that for certain LP problems there are alternatives to projection methods that have a better performance; this is pointed out above for the problems FIT1P and FIT2P of the Netlib/CUTEr set.

2.2 An example from image representation

The problem with the largest number of unknowns in the Netlib/CUTEr LP problem set used in [55] has $M \times N = 6,330 \times 22,275$ and (according to the on-line attachment to [55]), for that problem, all methods discussed in [55] need 42 seconds or more to reach the stopping tolerance on a 3.06 GHz Dell Precision 650 workstation. We found among the problems from applications that we have been investigating one that is over an order of magnitude larger and for which the projection algorithm recommended in [31] required only 25 seconds on the average on an Intel Xeon 1.7 GHz processor, 1 Gbyte memory, 32 bit workstation using the SNARK09 programming system [46]. We now give a brief description of this problem.

A $J \times J$ digitized image is one that is subdivided into J^2 square-shaped pixels within each of which the image value is uniform. Sometimes alternative representations of an image are superior. For example, in computerized tomography [58], we use the blob basis functions advocated by Lewitt [67] in some series expansion methods to reduce artifacts in the reconstruction. Such a reduction is due to the fact that blob basis functions are smoother than pixel basis functions.

The contribution to the image value at the center of any of the $M = J^2$ pixels by any of the N blob basis functions is known from the geometry of the representations.

If we are given a pixel image to start with and would like to find a good blob representation for it, the task is to find the weights x to be given to the blobs so that their combined contributions approximate the pixel values. In mathematical terms, this problem can be formulated as an instance of (1) in which the bounds c and d have to be tight to ensure a good approximation of the pixel image by the blob image. (The entries in the matrix A are the values of the various blobs at the centers of the various pixels.)

In the experiments reported in [31] $M \times N = 59,049 \times 51,152$. The algorithm that was found most efficacious among those tried is the projection method called CART3⁺⁺: the average (over 40 instances of the problem) time required by CART3⁺⁺ to find a solution to (1) was less than 25 seconds.

The algorithm CART3⁺⁺ belongs to a large family of projection methods that are usually referred to as *algebraic reconstruction techniques* (ART). These were first introduced to the tomographic image reconstruction literature in [54]; for a recent discussion, see [58, Chap. 11]. CART3⁺⁺, just like the closely related ART3+ of (2), has the property that, provided that the set of feasible vectors satisfying the inequalities in (1) is full-dimensional, it will find a feasible solution in a finite number of iterations [31]. The only difference between CART3⁺⁺ and ART3+ is the control sequence $(i(n))_{n \in \mathbb{N}}$, which provides CART3⁺⁺ with a wider range of applicability; see [31].

2.3 An example from intensity-modulated radiation therapy planning

The goal of intensity-modulated radiation therapy is to deliver sufficient doses to tumors to kill them, but without causing irreparable damage to critical organs. This requirement can be formulated as a linear feasibility problem of the kind shown in (1). The interpretation in this application is that each component of x is a to-be-determined intensity of radiation to be delivered to the patient in N separate beamlets, the components of Ax are the delivered doses at M points in the patient's body, and c and d are provided by the radiation oncologist as the desired limits on these doses. The following is based on the work performed for [32].

In the clinical case that we use as an example we have $M \times N = 302,491 \times 13,734$. The number of nonzero elements in A is 62,226,127, which is less than 1.5% of the total number of entries of A , an important consideration for the efficacy of projection methods for solving the problem. There is an additional technical consideration: since it is impossible to deliver negative radiation, each component of x has to be nonnegative, which results in an additional 13,734 inequality constraints. We use ART3+ (2) to solve this feasibility problem.

In clinical applications, it is considered desirable to find multiple feasible points, each of which is optimal according to its own criterion. A typical optimization task is “find a feasible point that results in the smallest total dose delivered to the liver.” The associated functional is linear: there is an $a \in \mathbb{R}^N$ such that $\langle a, x \rangle$ is the sum of those components of Ax that are associated with points in the liver. Recognizing the speed by which ART3+ finds a feasible point, we propose to apply it repeatedly, to solve the linear optimization problem:

$$\text{find an } x \in \mathbb{R}^N \text{ that minimizes } \langle a, x \rangle \text{ subject to } c \leq Ax \leq d. \quad (14)$$

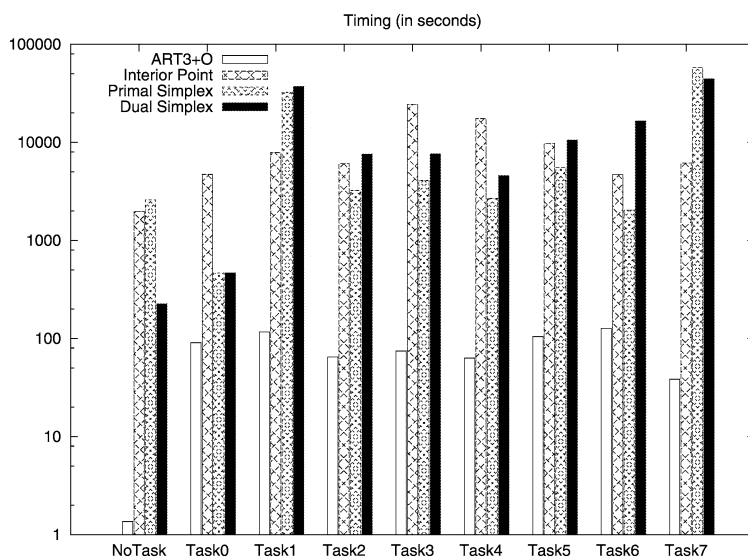


Fig. 4 Timings of the four methods for the feasibility run and the eight optimization tasks

Our method solves this problem by turning the objective function $\langle a, x \rangle$ into an additional constraint and solving

$$\text{find an } x \in \mathbb{R}^N \text{ such that } c \leq Ax \leq d \text{ and } \langle a, x \rangle \leq \rho \quad (15)$$

using ART3+. By reducing ρ using a bisection search until we obtain (within a pre-specified tolerance) the lowest value possible for it, we get a good approximation to a solution of (14). This whole process is called ART3+O [32].

The task of minimizing a linear functional subject to linear inequality constraints is the well-known LP problem and many software packages are available for solving it, see, e.g., [3]. To compare the efficiency of our proposed procedure with currently popular standard approaches, we applied them to the problem (14) for a patient with pancreatic cancer. We used all methods to find just a feasible point (No-Task) and also for eight different LP tasks representing various linear optimization criteria. The three algorithms with which we compared ART3+O were the self-dual interior point optimizer, the primal simplex optimizer and the dual simplex optimizer in the commercial software package MOSEK version 5. The results are reported in Fig. 4. Typically, for each task, ART3+O used about one to two minutes and the MOSEK algorithms needed one to several hours on an Intel Xeon 2.66 GHz processor, 16 Gbyte memory, 64 bit workstation. It is also noteworthy that the memory requirements of the MOSEK algorithms were at least twelve times as large as that of ART3+O. To enable the readers to investigate the performance of their favorite approaches to solving such problems, we posted the data for the eight LP tasks at http://dig.cs.gc.cuny.edu/~wei/web/?page_id=221.

2.4 Examples from computerized tomography

Computerized tomography is the problem of recovering an image from its measured (and hence not strictly accurate) integrals along M lines [58]. If we assume that the recovered image will be represented as a linear combination of N basis functions (see Sect. 2.2), then the task is to find the vector x the components of which are the weights to be given to the basis functions. Due to the linearity of integration and based on the knowledge of the basis functions, we can produce an $M \times N$ matrix A such that Ax is approximately the vector b of measurements. Since it is not likely that there is an x such that $Ax = b$, it is reasonable to aim instead at finding an x that minimizes

$$\sigma^2 \|b - Ax\|^2 + \|x\|^2, \quad (16)$$

where $\sigma \in \mathbb{R}$ indicates our confidence in our measurements. As explained in Section 11.3 of [58], this sought-after x is in fact the x part of the minimum norm solution of the consistent system of equations

$$[U \ \sigma A] \begin{bmatrix} u \\ x \end{bmatrix} = \sigma b, \quad (17)$$

where U is the $M \times M$ identity matrix. In the same section there is a derivation of a variant of ART that converges to the sought-after x , given by:

$$\begin{aligned} u^{(0)} &\text{ is the } M\text{-dimensional zero vector,} \\ x^{(0)} &\text{ is the } N\text{-dimensional zero vector,} \\ u^{(n+1)} &= u^{(n)} + \gamma_n e_{i(n)}, \\ x^{(n+1)} &= x^{(n)} + \sigma \gamma_n a_{i(n)}, \end{aligned} \quad (18)$$

with

$$\gamma_n = \lambda \frac{\sigma (b_{i(n)} - \langle a_{i(n)}, x^{(n)} \rangle) - u_{i(n)}^{(n)}}{1 + \sigma^2 \|a_{i(n)}\|^2}, \quad (19)$$

where, for $n \in \mathbb{N}$, $i(n) = (n \bmod M) + 1$, for $1 \leq i \leq M$, e_i is the M -dimensional vector whose i th component is 1 and whose other components are 0, a_i is the i th row of A and b_i is the i th component of b , and $0 < \lambda < 2$. Recognizing that in one iterative step only one row of the matrix is needed and that in computerized tomography most entries of each row are zero, we see that an iterative step can be carried out very rapidly, provided that we have access to the locations and the values of the nonzero entries. If the memory of the computer is large enough, this can be accommodated by storing A in a row-by-row sparse representation, otherwise the locations and values of the nonzero entries can be generated within each iterative step by some rapid mechanism, such as the *digital difference analyzer* explained, e.g., in Section 4.6 of [58].

In Sect. 5.8 of [58] there is an exact specification of the so-called *standard projection data* that are used to evaluate various reconstruction algorithms in that book, the

number of lines used in the standard projection data is $M = 223,744$. In the evaluations based on the standard projection data that are reported in [58] for reconstruction algorithms that use blob basis functions, the number of blobs used is $N = 51,152$. The first experiment on which we report in this subsection used exactly the same arrangement. (For the experiments in this subsection, the input data were created and outputs were analyzed and illustrated using SNARK09 [46].)

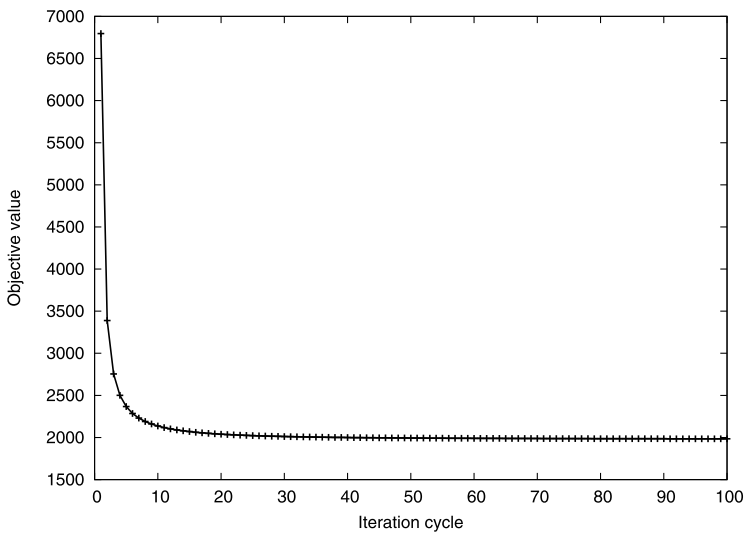
In this experiment we applied the ART algorithm of (18) and (19) with $\sigma = 5$ and $\lambda = 0.05$ to the standard projection data. In Fig. 5(a) we show the behavior of the objective function (16) as a function of iteration cycles (an *iteration cycle* is defined to be M iterations). It can be observed that the initial decrease in the objective function is very rapid.

This desirable initial behavior is even more noticeable when we evaluate the algorithm not from the purely mathematical point of view of how well the objective function is reduced, but rather from the application point of view of how good are the reconstructed images. For this purpose, we report on the *normalized mean absolute picture distance measure*, as defined in [58]. To define this measure we need a $J \times J$ digitization of the test phantom for which the data used in the reconstruction were collected; such a digitization for the phantom we used is shown in Fig. 6(a). In our definition of the measure we use $t_{h,k}$ and $s_{h,k}^{(n)}$ to denote the densities of the k th pixel of the h th row of the digitized test phantom and of the reconstruction (which is obtained from the vector $x^{(n)}$ of blob coefficients), respectively. We define the distance measure as

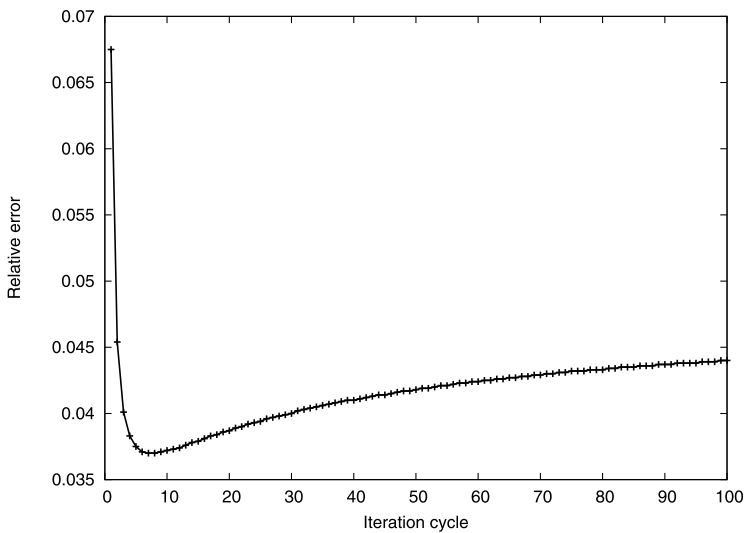
$$r^{(n)} = \frac{\sum_{h=1}^J \sum_{k=1}^J |t_{h,k} - s_{h,k}^{(n)}|}{\sum_{h=1}^J \sum_{k=1}^J |t_{h,k}|}. \quad (20)$$

In Fig. 5(b) we plot $r^{(n)}$ for this experiment. It is seen that its minimum is reached at the seventh iteration cycle, i.e., when $n = 7M$. This reflects the fact that the minimization objective (16) does not (and, in fact, it cannot in real applications where the phantom is not known to us) fully describe the application objective. For this reason it is standard practice in tomography [58] to stop the iterative process after a few iteration cycles and use the result at that time as the reconstruction. The digitization obtained from $x^{(7M)}$ produced by this experiment is shown in Fig. 6(b). The reconstruction is not perfect (as indeed it cannot possibly be since the measured data are only approximations of the line integrals assumed by the mathematics), but important features of the phantom are identifiable in the reconstruction. This ART reconstruction was carried out in 38.4 seconds on an Intel Core 1.6 GHz processor, 2 Gbyte memory, 32 bit laptop.

The plots in Fig. 5 nicely demonstrate the claim we have made in the Introduction that “Projection methods . . . possess . . . good initial behavior patterns.” ART seems to compare particularly favorably with some popular alternatives from this point of view. To demonstrate this we note that minimization of (16) can also be achieved using a *conjugate gradient* (CG) method (see, e.g., (12.48) and (12.49) in [58]). The computational cost of one iteration of the CG method is approximately the same as that of one iteration cycle by ART. However, what we get in quality of reconstruction during the early iterations is quite different: it is illustrated in Fig. 7 that the early iterates of ART significantly outperform what can be obtained at the same cost using



(a) Plot of the objective function (16).



(b) Plot of the distance measure (20).

Fig. 5 Image reconstruction by ART when $M \times N = 223,744 \times 51,152$. (One “Iteration cycle” is M ART iterative steps)

CG. This is an illustration of the well-known difference between the early behaviors of ART-type methods [58, Chap. 11] that make use of only one row of the system matrix A in any one iteration (such as the ART of (18)–(19)) and SIRT-type methods [58, Chap. 12] that make use of all rows of A simultaneously in every iteration (of which CG is an example) for image reconstruction from projections. Why this should be the case was studied for example in [72], by plotting the spectral radius per each

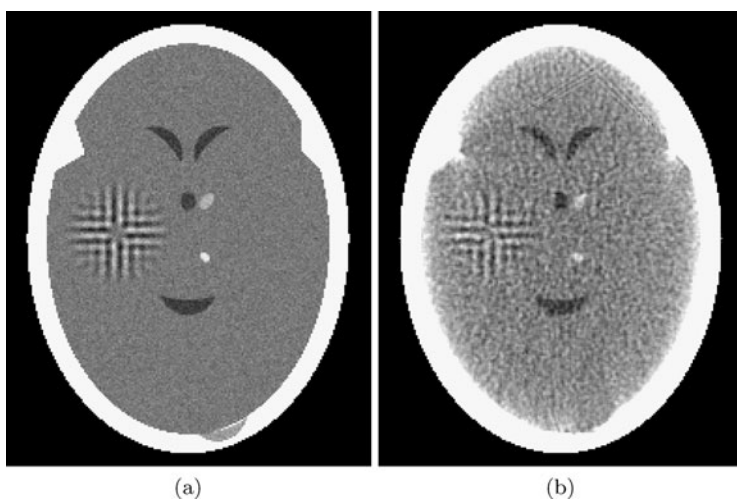


Fig. 6 Displays of a 243×243 digitized phantom (a) and of an ART reconstruction when $M \times N = 223,744 \times 51,152$ (b)

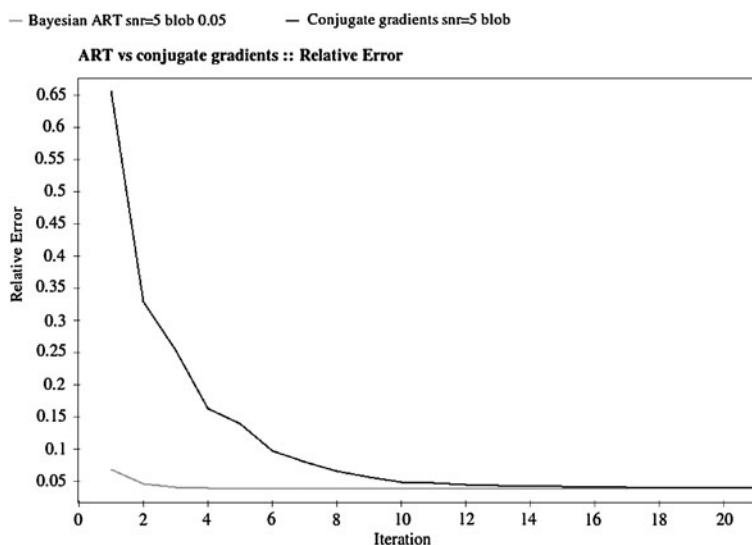


Fig. 7 Plots of the distance measure (20): image reconstruction by ART (*light*) and conjugate gradient method (*dark*) when $M \times N = 223,744 \times 51,152$. (One “Iteration” is M ART iterative steps, but it is only one iterative step by the conjugate gradient method. The computational cost of one “Iteration” is approximately the same in the two cases)

iteration cycle of the linear operator that is underlying the ART and SIRT algorithms, respectively; comparison of the spectral radii for ART (Fig. 5.7 of [72]) with those for SIRT (Fig. 5.8 of [72]) indicates why early iterates of ART produce better results than early iterates of SIRT.

We wanted to compare the time needed by ART with the time needed to solve the system (17) of consistent equations for the same data by the current implementation of the interior point method of MOSEK version 5 [3]. Unfortunately this could not be done, because the memory requirements of the MOSEK software were too large for our laptop. So we attempted to use a much more powerful Intel Xeon 2.66 GHz processor, 16 Gbyte memory, 64 bit workstation, but even the 16 Gbyte memory was too small to handle this problem using the MOSEK software. The importance of this memory requirement issue for the subject matter of this paper cannot be overemphasized: problems that routinely arise in real applications can be handled by projection methods using inexpensive laptops, while “more sophisticated alternatives” fail to produce any results even on much more powerful workstations due to their much greater demands on computer memory.

To be able to compare the efficiency of ART with that of the interior point method in MOSEK we had to reduce M and N to about a ninth of their previously-used sizes. Thus, in the second experiment on which we now report $M \times N = 24,880 \times 5,711$. For this smaller example we ran both ART and the interior point method in MOSEK (with its default parameters) on the Intel Xeon 2.66 GHz processor, 16 Gbyte memory, 64 bit workstation. In Fig. 8 we plot both the objective function and the distance measure for both algorithms as a function of time. From the point of view of the objective function, MOSEK needed over 5000 seconds to reach a value as low as ART reached in 10 seconds. The advantage of ART is more pronounced when considering the picture distance measure: the optimal value is reached by ART at 1.7 seconds (when $n = 14M$) while the interior point method never reaches a distance value that is as low as that of ART and it needs approximately 5000 seconds to reach its lowest distance measure.

Since both M and N are about a ninth of their previous sizes, we report in Fig. 9 on the 81×81 digitizations of the phantom and of the reconstruction $x^{(14M)}$. These are clearly inferior to the images in Fig. 6, demonstrating the medical necessity for the larger system of equations.

An alternative to minimizing (16) is to

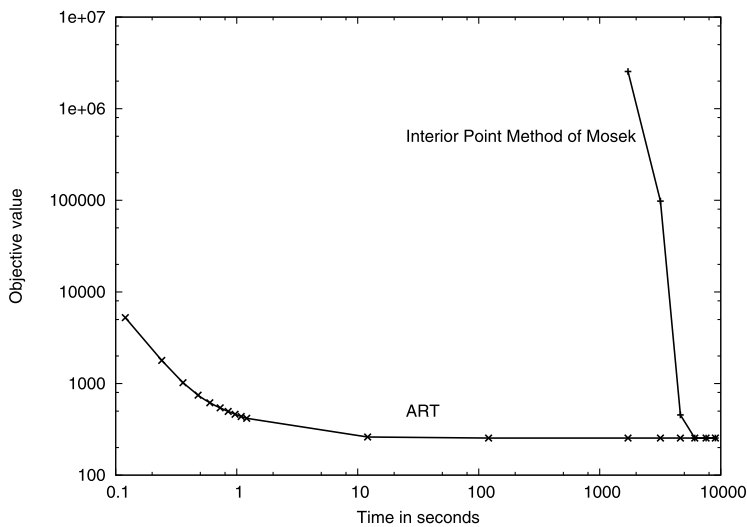
$$\text{find an } x \in \mathbb{R}^N \text{ that minimizes } \phi(x) \text{ subject to } \text{Res}(x) \leq \varepsilon, \quad (21)$$

where $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$ is a convex function, ε is a user-specified positive number and

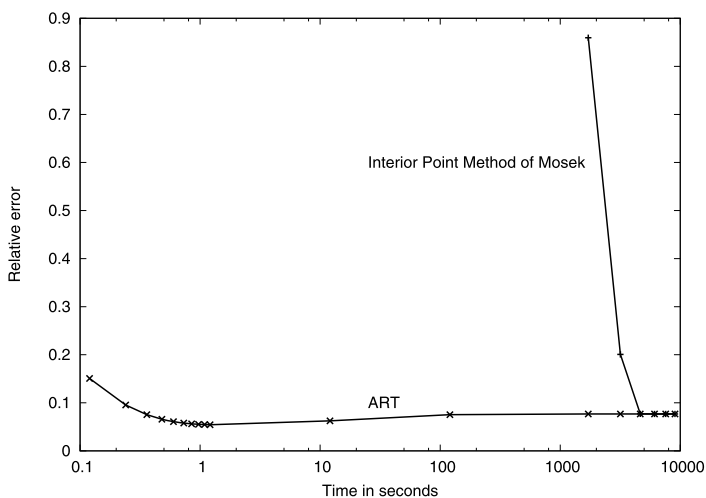
$$\text{Res}(x) = \sqrt{\sum_{i=1}^M \left(\frac{b_i - \langle a_i, x \rangle}{\|a_i\|} \right)^2}. \quad (22)$$

The intuitive intent is that $\phi(x)$ should be a measure of the undesirability of x . A choice that has become quite popular in recent years in the image processing literature is defined as follows. Let q be a $J \times J$ digitized image whose pixel values are denoted by $q_{h,k}$ ($1 \leq h \leq J$, $1 \leq k \leq J$). We define the *total variation* (TV) of q by

$$TV(q) = \sum_{h=1}^{J-1} \sum_{k=1}^{J-1} \sqrt{(q_{h+1,k} - q_{h,k})^2 + (q_{h,k+1} - q_{h,k})^2}. \quad (23)$$



(a) Plots of the objective function (16).



(b) Plots of the distance measure (20)

Fig. 8 Image reconstruction by ART and the interior point method in MOSEK when $M \times N = 24,880 \times 5,711$

By mapping q into a $(N = J^2)$ -dimensional vector x (by stacking into a single column all the columns of q), (23) gives rise to a convex function that can be used as the objective function. Indeed many algorithms have been proposed in the recent literature for TV-minimization by solving (21) with ϕ defined based on (23); for a recent representative publication see [50].

In practice, a minimization algorithm needs to be terminated at some point and so its actual output is only an approximation to a minimizer. If we happen to know the

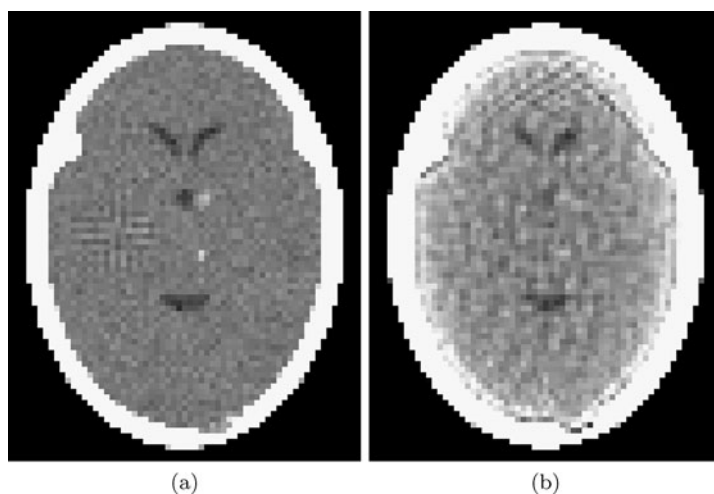


Fig. 9 Displays of an 81×81 digitized phantom (a) and of an ART reconstruction when $M \times N = 24,880 \times 5,711$ (b)

“truth” (the $J \times J$ digitization \bar{x} of the phantom for which the data b were collected), then any algorithm that solves

$$\text{find an } x \in \mathbb{R}^N \text{ such that } \phi(x) \leq \phi(\bar{x}) \text{ and } \text{Res}(x) \leq \text{Res}(\bar{x}) \quad (24)$$

should be considered satisfactory for the purpose at hand. We now report on the comparative speed on a problem of image reconstruction from projections of a state-of-the-art optimization algorithm and a projection method, both of which return a solution of (24). In brief: the projection method is more than four times faster. For the TV-minimization algorithm we used the code provided to us by Goldstein and Osher [51] in response to our request for their most efficient optimization algorithm to be used for solving (24). According to them, that code uses the algorithm they refer to as TwIST [14] with split Bregman [50] as the substep. For the projection method we used a block-iterative algorithm with superiorization [24, 45]. *Superiorization* uses perturbations to steer the iterative process of a projection method towards a minimizer of the given convex function ϕ .

For the experiment we used a 243×243 phantom presented in Fig. 10(a), with a large number of low contrast tumors introduced at various locations inside the skull. The TV value (23) of this phantom is 454.0. Data were collected based on the standard projection data from Sect. 5.8 of [58], but with 360 views making $M = 124,200$ and using pixels as the basis functions for the reconstruction algorithms with $N = 59,049$. Figure 10(b) shows the result of the reconstruction using the TV-superiorizing block-iterative projection method and Fig. 10(c) has the output image produced by the Goldstein-Osher code [51]. The TV value of these images are 451.5 and 443.9 for the image in Figs. 10(b) and 10(c), respectively. Both algorithms provide a solution to the problem (24), but the times needed to produce the images are substantially different: on an Intel Pentium Core 2 Duo 2.8 GHz (using a single CPU) with 4 Gbyte memory the image in Fig. 10(b) needed less than 7 seconds to

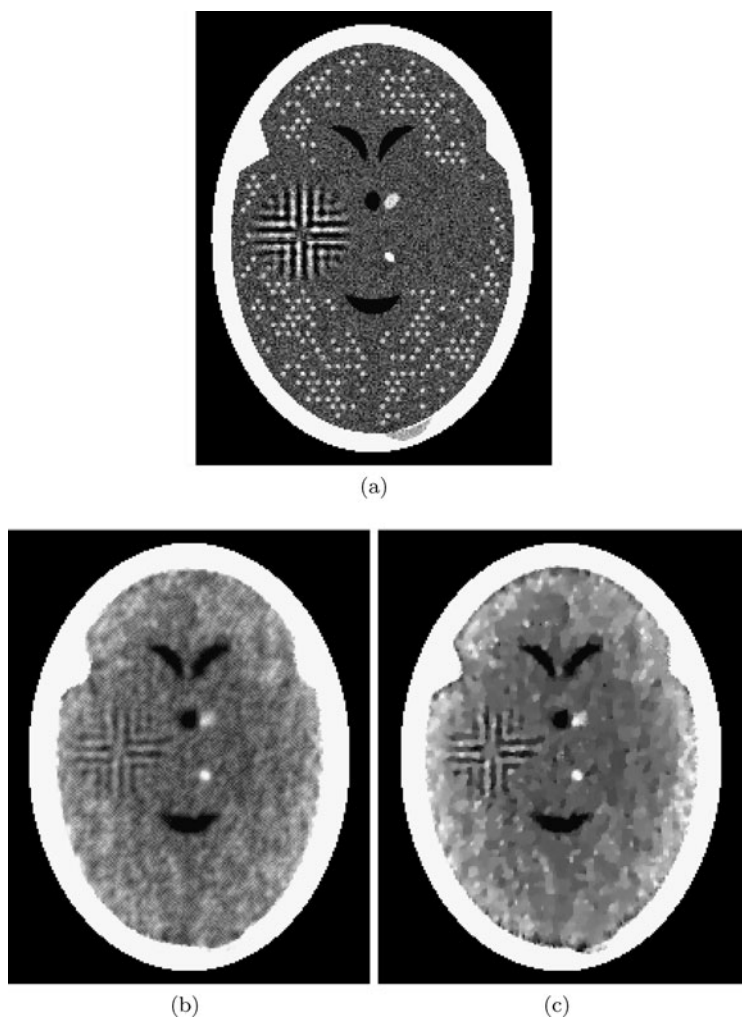


Fig. 10 (a) A head phantom for which realistically-simulated data were collected. (b) Reconstruction using a block-iterative TV-superiorizing projection algorithm (reconstruction time is less than 7 seconds). (c) Reconstruction using the Goldstein-Osher code [51] (reconstruction time is over 28 seconds)

be produced by the projection method versus slightly over 28 seconds needed by the optimization algorithm to produce the image in Fig. 10(c).

3 Published and patented results

3.1 Scientific publications

Here we give a brief glimpse into some recently published results that show the efficacy of projection methods for some large problems. In the problems discussed in

[45], the number of unknowns was 59,049. In the examples given in [59] (a paper devoted to radiation therapy planning), problems of the form (1) were considered with the number N of unknowns only 515 but the number of pairs of constraints $M = 128,688$. In four out of the six cases reported there, the projection method ART3+ (2) found a feasible point in less than three seconds, and in the remaining two cases a feasible point was found in less than 34 seconds. These times are for a standard PC, using an Intel Xeon 1.7 GHz processor and 1 Gbyte memory. The problems in [45, 59] are small compared to some of the other applications for which projection methods have been successfully used. In [21] (a paper devoted to reconstruction from electron micrographs), there are examples in which 16,777,216 unknowns are to be recovered from 4,587,520 measurements (each giving an approximate linear equality) and others in which 884,436 unknowns are to be recovered from 92,160,000 measurements. Projection methods were used in [21] to handle such large problems in a reasonable time.

In a recent paper [63] it is shown that a variant of ART can be used for crystal lattice orientation distribution function estimation from diffraction data. One of the problems discussed in [63] has 1,372,000,000 unknowns and the number of equations is potentially infinite. They are randomly generated and a projection step can be carried out as soon as a new equation is available (an ideal use of a sequential projection method of the row-action type, see [22]). The result reported in the paper for that problem is that obtained after 1,000,000,000 such projection steps.

As for all methodologies, projection methods are not necessarily the approach of choice in all applications. However, in important applications in biomedicine and image processing, projection methods work well and have been used successfully for a long time. For example, an important application of reconstruction from projections is electron microscopy and some of the leading groups in that field consider the projection method “ART with blobs” to be the method of choice, see [13]. A mathematical reason for this is that for such problems the angles between hyperplanes or half-spaces, represented by linear equalities or linear inequalities as in (1) and (3), are in general large (in the sense that the cosine of the angle between the normals of two randomly chosen hyperplanes in the system to be solved is likely to be near zero) due to the high sparsity in each of the rows of the system matrix.

3.2 Commercial patents

There is hardly better evidence for the value of projection methods than the many patents for commercial purposes that include them. Projection methods are used in commercial devices in many areas. Unfortunately, if a device is truly commercial, then the algorithm that is actually used in it is proprietary and usually not published. Many commercial emission tomography scanners use now some sort of iterative algorithms. A prime example is provided by the commercially-successful Philips Allegro scanners [33]. In x-ray computerized tomography (CT), there are reports emanating from companies that sell such scanners indicating that variants of ART are used in heart imaging; an example is presented in [61].

The first EMI (Electric & Musical Industries Ltd., London, England, UK) CT scanner, invented by G.N. Hounsfield [60], used a variant of ART. For this pioneering invention, Hounsfield shared the Nobel Prize with A.M. Cormack in 1979. Thirty years later (on September 29, 2009), a patent was issued to Philips (Koninklijke Philips Electronics N.V., Eindhoven, The Netherlands) for a “Method and device for the iterative reconstruction of cardiac images” [86]. The role of projection methods is demonstrated by the following quote from the “Summary of the Invention” included in the Patent Description:

“The iterative reconstruction applied here may particularly be based on an Algebraic Reconstruction Technique (ART) (cf. R. Gordon, R. Bender, and G.T. Herman: “Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography”, *J. Theor. Biol.*, 29:471–481, 1970) or on a Maximum Likelihood (ML) algorithm (K. Lange and J.A. Fessler: “Globally convergent algorithms for maximum a posteriori transmission tomography”, *IEEE Transactions on Image Processing*, 4(10):1430–1450, 1995), wherein each image update step uses the projections of a selected subset, i.e., projections corresponding to a similar movement phase.”

4 Conclusion

In this paper we have shown that, whether or not alternative methods are applicable, correctly implemented projection methods are very efficient for solving some convex feasibility problems with linear inequality constraints, especially for those that are large, sparse, and originate from real-life applications.

Acknowledgements The work of Y. Censor, W. Chen, R. Davidi, and G.T. Herman was supported by NIH Award Number R01HL070472 from the National Heart, Lung, and Blood Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Heart, Lung, and Blood Institute or the National Institutes of Health. The work of P.L. Combettes was supported by the Agence Nationale de la Recherche under grant ANR-08-BLAN-0294-02. The work of Y. Censor was also supported by United States-Israel Binational Science Foundation (BSF) Grant number 200912. We thank the Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School (and especially David Craft) for providing access to their computers. We are also grateful to T. Goldstein and S. Osher of UCLA for providing us with the optimization code [51] that they considered most appropriate for solving (24).

References

1. Agmon, S.: The relaxation method for linear inequalities. *Can. J. Math.* **6**, 382–392 (1954)
2. Aharoni, R., Censor, Y.: Block-iterative projection methods for parallel computation of solutions to convex feasibility problems. *Linear Algebra Appl.* **120**, 165–175 (1989)
3. Andersen, E.D., Andersen, K.D.: The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In: Frenk, H., Roos, K., Terlaky, T., Zhang, S. (eds.) *High Performance Optimization*, pp. 197–232. Kluwer, Boston (2000)
4. Andrews, H.C., Hunt, B.R.: *Digital Image Restoration*. Englewood Cliffs, Prentice-Hall (1977)
5. Auslender, A.: *Méthodes Numériques pour la Résolution des Problèmes d’Optimisation avec Contraintes*. Thèse, Faculté des Sciences, Grenoble (1969)
6. Auslender, A.: *Optimisation – Méthodes Numériques*. Masson, Paris (1976)

7. Bauschke, H.H.: The approximation of fixed points of compositions of nonexpansive mappings in Hilbert space. *J. Math. Anal. Appl.* **202**, 150–159 (1996)
8. Bauschke, H.H., Borwein, J.M.: On projection algorithms for solving convex feasibility problems. *SIAM Rev.* **38**, 367–426 (1996)
9. Bauschke, H.H., Combettes, P.L.: A weak-to-strong convergence principle for Fejér-monotone methods in Hilbert spaces. *Math. Oper. Res.* **26**, 248–264 (2001)
10. Bauschke, H.H., Combettes, P.L., Kruk, S.G.: Extrapolation algorithm for affine-convex feasibility problems. *Numer. Algorithms* **41**, 239–274 (2006)
11. Bauschke, H.H., Deutsch, F., Hundal, H., Park, S.-H.: Accelerating the convergence of the method of alternating projections. *Trans. Am. Math. Soc.* **355**, 3433–3461 (2003)
12. Bauschke, H.H., Matoušková, E., Reich, S.: Projection and proximal point methods: convergence results and counterexamples. *Nonlinear Anal.* **56**, 715–738 (2004)
13. Bilbao-Castro, J.R., Marabini, R., Sorzano, C.O.S., García, I., Carazo, J.M., Fernández, J.J.: Exploiting desktop supercomputing for three-dimensional electron microscopy reconstructions using ART with blobs. *J. Struct. Biol.* **65**, 19–26 (2009)
14. Bioucas-Dias, J., Figueiredo, M.: A new TwIST: two-step iterative shrinkage/thresholding algorithms for image restoration. *IEEE Trans. Image Process.* **16**, 2992–3004 (2007)
15. Blatt, D., Hero, A.O., III: Energy based sensor network source localization via projection onto convex sets (POCS). *IEEE Trans. Signal Process.* **54**, 3614–3619 (2006)
16. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge (2004)
17. Brègman, L.M.: The method of successive projection for finding a common point of convex sets. *Sov. Math. Dokl.* **6**, 688–692 (1965)
18. Butnariu, D., Censor, Y.: On the behavior of a block-iterative projection method for solving convex feasibility problems. *Int. J. Comput. Math.* **34**, 79–94 (1990)
19. Butnariu, D., Censor, Y., Reich, S. (eds.): *Inherently Parallel Algorithms in Feasibility and Optimization and Their Applications*. Elsevier, Amsterdam (2001)
20. Butnariu, D., Davidi, R., Herman, G.T., Kazansev, I.G.: Stable convergence behavior under summable perturbations of a class of projection methods for convex feasibility and optimization problems. *IEEE J. Sel. Top. Signal Process.* **1**, 540–547 (2007)
21. Carazo, J.M., Herman, G.T., Sorzano, C.O.S., Marabini, R.: Algorithms for three-dimensional reconstruction from imperfect projection data provided by electron microscopy. In: Frank, J. (ed.) *Electron Tomography: Methods for Three-Dimensional Visualization of Structures in the Cell*, 2nd edn., pp. 217–243. Springer, New York (2006)
22. Censor, Y.: Row-action methods for huge and sparse systems and their applications. *SIAM Rev.* **23**, 444–466 (1981)
23. Censor, Y., Altschuler, M.D., Powlis, W.D.: On the use of Cimmino's simultaneous projections method for computing a solution of the inverse problem in radiation therapy treatment planning. *Inverse Probl.* **4**, 607–623 (1988)
24. Censor, Y., Davidi, R., Herman, G.T.: Perturbation resilience and superiorization of iterative algorithms. *Inverse Probl.* **26**, 065008 (2010)
25. Censor, Y., Elfving, T., Herman, G.T.: Averaging strings of sequential iterations for convex feasibility problems. In: Butnariu, D., Censor, Y., Reich, S. (eds.) *Inherently Parallel Algorithms in Feasibility and Optimization and Their Applications*, pp. 101–114. Elsevier, Amsterdam (2001)
26. Censor, Y., Gordon, D., Gordon, R.: BICAV: A block-iterative, parallel algorithm for sparse systems with pixel-related weighting. *IEEE Trans. Med. Imaging* **20**, 1050–1060 (2001)
27. Censor, Y., Segal, A.: On the string averaging method for sparse common fixed points problems. *Int. Trans. Oper. Res.* **16**, 481–494 (2009)
28. Censor, Y., Tom, E.: Convergence of string-averaging projection schemes for inconsistent convex feasibility problems. *Optim. Methods Softw.* **18**, 543–554 (2003)
29. Censor, Y., Zenios, S.A.: *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, New York (1997)
30. Cetin, A.E., Ozaktas, H., Ozaktas, H.M.: Resolution enhancement of low resolution wavefields with POCS algorithm. *Electron. Lett.* **39**, 1808–1810 (2003)
31. Chen, W., Herman, G.T.: Efficient controls for finitely convergent sequential algorithms. *ACM Trans. Math. Softw.* **37**, Article No. 14 (2010)
32. Chen, W., Craft, D., Madden, T.M., Zhang, K., Kooy, H.M., Herman, G.T.: A fast optimization algorithm for multi-criteria intensity modulated proton therapy planning. *Med. Phys.* **7**, 4938–4945 (2010)

33. Chiang, S., Cardi, C., Matej, S., Zhuang, H., Newberg, A., Alavi, A., Karp, J.S.: Clinical validation of fully 3-D versus 2.5-D RAMLA reconstruction on the Phillips-ADAC CPET PET scanner. *Nucl. Med. Commun.* **25**, 1103–1107 (2004)
34. Choi, H., Baraniuk, R.G.: Multiple wavelet basis image denoising using Besov ball projections. *IEEE Signal Process. Lett.* **11**, 717–720 (2004)
35. Cimmino, G.: Calcolo approssimato per le soluzioni dei sistemi di equazioni lineari. *Ric. Sci. (Roma)* **1**, 326–333 (1938)
36. Combettes, P.L.: The foundations of set theoretic estimation. *Proc. IEEE* **81**, 182–208 (1993)
37. Combettes, P.L.: The convex feasibility problem in image recovery. *Adv. Imaging Electron Phys.* **95**, 155–270 (1996)
38. Combettes, P.L.: Hilbertian convex feasibility problem: Convergence of projection methods. *Appl. Math. Optim.* **35**, 311–330 (1997)
39. Combettes, P.L.: Convex set theoretic image recovery by extrapolated iterations of parallel subgradient projections. *IEEE Trans. Image Process.* **6**, 493–506 (1997)
40. Combettes, P.L., Hirstoaga, S.A.: Equilibrium programming in Hilbert spaces. *J. Nonlinear Convex Anal.* **6**, 117–136 (2005)
41. Combettes, P.L., Trussell, H.J.: Methods for digital restoration of signals degraded by a stochastic impulse response. *IEEE Trans. Acoust. Speech Signal Process.* **37**, 393–401 (1989)
42. Cottle, R.W., Pang, J.-S.: On solving linear complementarity problems as linear programs. *Math. Program. Stud.* **7**, 88–107 (1978)
43. Crombez, G.: Improving the speed of convergence in the method of projections onto convex sets. *Publ. Math. (Debr.)* **58**, 29–48 (2001)
44. Crombez, G.: Finding common fixed points of strict paracontractions by averaging strings of sequential iterations. *J. Nonlinear Convex Anal.* **3**, 345–351 (2002)
45. Davidi, R., Herman, G.T., Censor, Y.: Perturbation-resilient block-iterative projection methods with application to image reconstruction from projections. *Int. Trans. Oper. Res.* **16**, 505–524 (2009)
46. Davidi, R., Herman, G.T., Klukowska, J.: SNARK09: A programming system for the reconstruction of 2D images from 1D projections. <http://www.dig.cs.gc.cuny.edu/software/snark09/> (2011). Accessed 25 February 2011
47. Deutsch, F.: *Best Approximation in Inner Product Spaces*. Springer, New York (2001)
48. Eckstein, J., Svaiter, B.F.: General projective splitting methods for sums of maximal monotone operators. *SIAM J. Control Optim.* **48**, 787–811 (2009)
49. Eggermont, P.P.B., Herman, G.T., Lent, A.: Iterative algorithms for large partitioned linear systems, with applications to image reconstruction. *Linear Algebra Appl.* **40**, 37–67 (1981)
50. Goldstein, T., Osher, S.: The split Bregman method for L1 regularized problems. *SIAM J. Imaging Sci.* **2**, 323–343 (2009)
51. Goldstein, T., Osher, S.: Personal communication (2010)
52. Golub, G.H., van Loan, C.F.: *Matrix Computations*, 3rd edn. Johns Hopkins University Press, Baltimore (1996)
53. González-Castaño, F.J., García-Palomares, U.M., Alba-Castro, J.L., Pousada-Carballo, J.M.: Fast image recovery using dynamic load balancing in parallel architectures, by means of incomplete projections. *IEEE Trans. Image Process.* **10**, 493–499 (2001)
54. Gordon, R., Bender, R., Herman, G.T.: Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography. *J. Theor. Biol.* **29**, 471–482 (1970)
55. Gould, N.I.M.: How good are projection methods for convex feasibility problems? *Comput. Optim. Appl.* **40**, 1–12 (2008)
56. Gu, J., Stark, H., Yang, Y.: Wide-band smart antenna design using vector space projection methods. *IEEE Trans. Antennas Propag.* **52**, 3228–3236 (2004)
57. Gubin, L.G., Polyak, B.T., Raik, E.V.: The method of projections for finding the common point of convex sets. *U.S.S.R. Comput. Math. Math. Phys.* **7**, 1–24 (1967)
58. Herman, G.T.: *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*, 2nd edn. Springer, London (2009)
59. Herman, G.T., Chen, W.: A fast algorithm for solving a linear feasibility problem with application to intensity-modulated radiation therapy. *Linear Algebra Appl.* **428**, 1207–1217 (2008)
60. Hounsfield, G.N.: A method and apparatus for examination of a body by radiation such as X or gamma radiation. UK Patent No. 1283915 (1968/72)
61. Isola, A.A., Ziegler, A., Koehler, T., Niessen, W.J., Grass, M.: Motion compensated iterative cone-beam CT image reconstruction with adapted blobs as basis functions. *Phys. Med. Biol.* **53**, 6777–6797 (2008)

62. Kaczmarz, S.: Angenäherte Auflösung von Systemen linearer Gleichungen. *Bull. Acad. Sci. Pol. A* **35**, 355–357 (1937)
63. Kazantsev, I.G., Schmidt, S., Poulsen, H.F.: A discrete spherical X-ray transform of orientation distribution functions using bounding cubes. *Inverse Probl.* **25**, 105009 (2009)
64. Kiwiel, K.C.: Monotone Gram matrices and deepest surrogate inequalities in accelerated relaxation methods for convex feasibility problems. *Linear Algebra Appl.* **252**, 27–33 (1997)
65. Kiwiel, K.C., Lopuch, B.: Surrogate projection methods for finding fixed points of firmly nonexpansive mappings. *SIAM J. Optim.* **7**, 1084–1102 (1997)
66. Lee, S.-H., Kwon, K.-R.: Mesh watermarking based on projection onto two convex sets. *Multimed. Syst.* **13**, 323–330 (2008)
67. Lewitt, R.M.: Multidimensional digital image representation using generalized Kaiser-Bessel window functions. *J. Opt. Soc. Am. A* **7**, 1834–1846 (1990)
68. Liew, A.W.-C., Yan, H., Law, N.-F.: POCS-based blocking artifacts suppression using a smoothness constraint set with explicit region modeling. *IEEE Trans. Circuits Syst. Video Technol.* **15**, 795–800 (2005)
69. Lu, Y.M., Karzand, M., Vetterli, M.: Demosaicking by alternating projections: theory and fast one-step implementation. *IEEE Trans. Image Process.* **19**, 2085–2098 (2010)
70. Merzlyakov, Y.I.: On a relaxation method of solving systems of linear inequalities. *U.S.S.R. Comput. Math. Math. Phys.* **2**, 504–510 (1963)
71. Motzkin, T.S., Schoenberg, I.J.: The relaxation method for linear inequalities. *Can. J. Math.* **6**, 393–404 (1954)
72. Natterer, F., Wübbeling, F.: *Mathematical Methods in Image Reconstruction*. SIAM, Philadelphia (2001)
73. Ottavay, N.: Strong convergence of projection-like methods in Hilbert spaces. *J. Optim. Theory Appl.* **56**, 433–461 (1988)
74. Penfold, S.N., Schulte, R.W., Censor, Y., Bashkirov, V., McAllister, S., Schubert, K.E., Rosenfeld, A.B.: Block-iterative and string-averaging projection algorithms in proton computed tomography image reconstruction. In: Censor, Y., Jiang, M., Wang, G. (eds.) *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning and Inverse Problems*, pp. 347–367. Medical Physics Publishing, Madison (2010)
75. Pierra, G.: Éclatement de contraintes en parallèle pour la minimisation d’une forme quadratique. In: *Proc. 7th IFIP Conf. on Optimization Techniques: Modeling and Optimization in the Service of Man. Lecture Notes in Comput. Sci.*, vol. 41, pp. 200–218. Springer, London (1976)
76. Pierra, G.: Decomposition through formalization in a product space. *Math. Program.* **28**, 96–115 (1984)
77. Rhee, H.: An application of the string averaging method to one-sided best simultaneous approximation. *J. Korea Soc. Math. Educ. Ser. B Pure Appl. Math.* **10**, 49–56 (2003)
78. Samsonov, A.A., Kholmovski, E.G., Parker, D.L., Johnson, C.R.: POCS-SENSE: POCS-based reconstruction for sensitivity encoded magnetic resonance imaging. *Magn. Reson. Med.* **52**, 1397–1406 (2004)
79. Shaked, N.T., Rosen, J.: Multiple-viewpoint projection holograms synthesized by spatially incoherent correlation with broadband functions. *J. Opt. Soc. Am. A* **25**, 2129–2138 (2008)
80. Sharma, G.: Set theoretic estimation for problems in subtractive color. *Color Res. Appl.* **25**, 333–348 (2000)
81. Stark, H., Yang, Y.: *Vector Space Projections: A Numerical Approach to Signal and Image Processing, Neural Nets, and Optics*. Wiley-Interscience, New York (1998)
82. van Wyk, B.J., van Wyk, M.A.: A POCS-based graph matching algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**, 1526–1530 (2004)
83. Youla, D.C., Webb, H.: Image restoration by the method of convex projections: Part 1—theory. *IEEE Trans. Med. Imaging* **1**, 81–94 (1982)
84. Yukawa, M., Yamada, I.: Pairwise optimal weight realization—acceleration technique for set-theoretic adaptive parallel subgradient projection algorithm. *IEEE Trans. Signal Process.* **54**, 4557–4571 (2006)
85. Zhang, T., Hong, H.: Restoration algorithms for turbulence-degraded images based on optimized estimation of discrete values of overall point spread functions. *Opt. Eng.* **44**, 017005 (2005)
86. Ziegler, A., Grass, M., Koehler, T.: Method and device for the iterative reconstruction of cardiac images. US Patent Number 7596204 (2009)