

SOME BASIC INFORMATION ON INFORMATION-BASED COMPLEXITY THEORY

BERESFORD N. PARLETT

ABSTRACT. Numerical analysts might be expected to pay close attention to a branch of complexity theory called information-based complexity theory (IBCT), which produces an abundance of impressive results about the quest for approximate solutions to mathematical problems. Why then do most numerical analysts turn a cold shoulder to IBCT? Close analysis of two representative papers reveals a mixture of nice new observations, error bounds repackaged in new language, misdirected examples, and misleading theorems.

Some elements in the framework of IBCT, erected to support a rigorous yet flexible theory, make it difficult to judge whether a model is off-target or reasonably realistic. For instance, a sharp distinction is made between information and algorithms restricted to this information. Yet the information itself usually comes from an algorithm, so the distinction clouds the issues and can lead to true but misleading inferences. Another troublesome aspect of IBCT is a free parameter F , the class of admissible problem instances. By overlooking F 's membership fee, the theory sometimes distorts the economics of problem solving in a way reminiscent of agricultural subsidies.

The current theory's surprising results pertain only to unnatural situations, and its genuinely new insights might serve us better if expressed in the conventional modes of error analysis and approximation theory.

CONTENTS

1. Introduction
2. High level criticisms
3. Preliminaries
4. On the optimal solution of large linear systems
5. Optimal solution of large eigenpair problems

References

1. INTRODUCTION AND SUMMARY

In 1980 Traub and Wozniakowski published a monograph entitled *A general theory of optimal algorithms*, which initiated a new line of research. The subject was called analytic complexity theory, initially, but is now referred to as information-based complexity theory (IBCT hereafter). The August 1987 issue

Received by the editors November 10, 1989 and, in revised form, July 9, 1990.

1991 *Mathematics Subject Classification*. Primary 68Q25; Secondary 65F99, 65Y20.

The author gratefully acknowledges partial support from ONR Contract N00014-90-J-1372.

of Bull. Amer. Math. Soc. (N.S.) contains a summary of recent results [Pac & Wo, 1987], so acceptance of this branch of complexity theory has been swift.

One purpose of the general theory is to provide an attractive setting for the systematic study of some of the problems that have engaged numerical analysts for decades. One among the programs of IBCT is to determine the minimal cost of computing an approximate solution (of given accuracy) over *all* algorithms that one could use which restrict themselves to certain limited information about the data. It is also of interest to discover any algorithms that achieve this minimal cost or, at least, some cost close to it. Pride of place in IBCT is given to the information to which the algorithms are limited. By its choice of problems IBCT is (potentially) a branch of complexity theory that is highly relevant to numerical analysts. Whenever the minimal costs can be estimated, they provide yardsticks against which to measure actual algorithms. The program seems attractive.

The purpose of this paper is to sound a warning about IBCT, but first we point out why our task is less than straightforward and why our observations have more than local interest.

Our task would be easier if mathematics enjoyed a tradition of public debate and controversy, as exists in some other fields, from which some kind of consensus could emerge if not Truth. But too many of us treat every mathematical assertion as if it must be true or false, and if false, then mistaken, and if mistaken, then a symptom of some lapse or inaptitude. Such an oversimplification would wrongly impugn the technical competency of the founders of the IBCT and miss the point of our criticism; we intend to show that parts of IBCT are true, but mistaken. Our task would be easier if IBCT contained a logical flaw so egregious that its mere exposure rendered further discussion unnecessary. But the flaws are subtle, subtle enough that they appear only after their consequences have been analyzed in detail. That is why we must focus upon specific work. We have selected two related papers [Tr & Wo, 1984; Ku, 1986] to be examined in detail in §§4 and 5. They concern matrix computations, which is the area we understand best of all the many areas that IBCT covers. Besides, we believe these papers are typical; written in a professional manner, their analysis is not weak, and some of their observations are new and interesting in a technical way regardless of IBCT. We claim that their results, most of them, are *seriously misleading*.

Since the arguments in the papers are impeccable, the flaw must be in the framework. Yet the definitions are laid out plainly for all to see and seem to be appropriate—a puzzling situation.

One source of difficulty is the redefinition of common terms such as “eigenvalue problem” (see §2.4) or “worst case” (see §2.5) or “information” (see §4.4) or “algorithm” (see §§2.1, 3.2, 4.4, and 5.3). These slight twists to conventional meanings are subtle enough to escape notice but entail significant consequences. The results mislead because they are remembered and discussed in terms of ordinary word usage. Most readers will not even be aware of the shifts in meaning, some of which are due to the tempting but artificial distinction between information and algorithm.

Another feature of IBCT that can sometimes rob results of their relevance is the presence of a free parameter: the class F from which worst cases are drawn. The cost of testing membership in F is ignored by IBCT, so the model

loses validity whenever this cost is not negligible.

Some of our criticisms require very little knowledge of the subject matter. These criticisms are presented in the next section. After that, we get down to details, provide some background material, and then examine each paper in turn. In our summaries we try to be fair, but we encourage the interested reader to compare our effort with the original work.

A handful of reservations about IBCT have appeared in print. In a review of the second monograph [Tr, Wo & Wa, 1983], Shub [Shu, 1987] gives a couple of instances of unnatural measures of cost. In [Ba, 1987], Babuska calls on researchers in IBCT to make their models more realistic. We concur but note that the model may be so flexible as to embrace pointless investigations as readily as pertinent ones.

We make no complaint that IBCT ignores the roundoff error that afflicts implementation on digital computers. First, a good understanding of computation in exact arithmetic is a prerequisite for tackling practical issues. Second, we must acknowledge that a large part of theoretical numerical analysis confines itself to the comforts of exact arithmetic.

IBCT has already produced a large body of results, some of them surprising and consequently of potential interest. Yet each surprising result known to us, in worst-case analysis, holds only within a model sufficiently unnatural as to forfeit attention from numerical analysts. This is a pity because IBCT certainly permits realistic models, and there is plenty to do; the investigation of average case complexity of approximately solved problems is in its infancy. It would take only a few illuminating results concerning some reasonable models to restore faith in the program launched by the general theory of optimal algorithms. So, it is a pleasure to acknowledge the recent and interesting work by Traub and Wozniakowski on the average case behavior of the symmetric Lanczos algorithm [Tr & Wo, 1990]. However, we must note that the infrastructure of IBCT plays little role in this analysis.

The incentive to write this essay came from discussions held during the workshop at the Mathematical Sciences Research Institute at Berkeley in January 1986 under the title Problems Relating Computer Science to Numerical Analysis.

2. HIGH LEVEL CRITICISMS

2.1. This is not complexity theory. Numerical analysis and complexity theory are palpably different subjects. Complexity theory (CT hereafter) seeks to determine the *intrinsic* difficulty of certain tasks; whereas much of theoretical numerical analysis (NA hereafter) has been concerned with studying classes of *algorithms*, analyzing convergence and stability, developing error bounds (either a priori or a posteriori), and detecting either optimal or desirable members of a class according to various criteria. Clearly CT has more ambitious goals than does NA.

One major theme of IBCT is to find the minimal cost of achieving a certain level of approximation for the hardest case in a given problem class F , using any algorithm that confines itself to certain partial information about the case. One of the papers we examine is concerned with the solution of large systems of linear equations and the other with the matrix eigenvalue problem (see [Tr & Wo, 1984; Ku, 1986]). Now we can formulate our first complaint, one that

applies to the results in both papers.

The theorems say nothing about the intrinsic cost of computing an approximate solution to either of the problems mentioned above because the specified information is not naturally associated with the task but is acquired when a certain class of numerical methods is employed.

The class is sometimes called the Krylov subspace methods; one is *not* given a matrix A explicitly, but instead a few vectors $b, Ab, A^2b, A^3b, \dots, A^j b$, and one wishes to approximate the solution of a system of linear equations $Ax = b$ or some specified eigenpair of A . More details are given in §§3.2 and 3.3. So the invitation to minimize cost over *all* algorithms subject to the given information turns out, in these cases, to amount to the quest for the best that can be achieved at *each step* of a Krylov subspace method. This is exactly the sort of work that numerical analysts do.

We do not wish to belabor the obvious, but our suggestion that, in these cases, IBCT has the appearance of CT without the substance is important for the following reason. It might be claimed that we interpret IBCT results as though they were results about Krylov subspace methods (i.e., NA) when, in fact, they are CT results concerning Krylov *information*. In other words, perhaps we are guilty of looking at the world through NA spectacles and missing that subtle difference of emphasis characteristic of CT. This possibility needs to be considered, but the stubborn fact remains that restricting information to Krylov information is not part of the linear equations problem nor of the eigenvalue problem.

2.2. Free information in the problem class. The ingredient of IBCT that allows it to generate irrelevant results is the problem class F (see the second paragraph in §2.1). F did not appear in our brief description of the theory in the second paragraph of §1 because it is not a logically essential ingredient but rather a parameter within IBCT. Let us describe the role it plays. There is a task T to be accomplished; there is an information sequence $N = (N_1, N_2, N_3, \dots)$ coupled with a measure of cost (N_j costs j units); and there is F . For each N_j the only algorithms admitted for consideration are those that restrict themselves to use N_j and standard auxiliary computations. For a worst-case complexity analysis, the main technical goal is to determine the minimal cost, over admissible algorithms, required to achieve T for the most difficult problem *within* F that is consistent with N . This minimal cost may be called $C(F, N, T)$.

Suppose now that $F_1 \subset F_2$ and that $C(F_1, N, T) < C(F_2, N, T)$. Such a result is of little relevance to the achievement of T unless one can determine that a problem lies within F_1 rather than F_2 . To put the matter in other words, we might say that knowledge of membership in F is information and should have a cost attached to it. Whenever F is very large (for example, the class of continuous functions or the class of invertible matrices), it is realistic to assign no cost to it. On the other hand, there are examples (see §4.4) where it may be as expensive to ascertain membership in F as to achieve T , given N , over a larger class of problems. In such cases $C(F, N, T)$ bounds one part of the expense while ignoring the other. Let $C(N, T)$ denote $C(F, N, T)$ when F is as large as possible.

We may reformulate the minimax quantity $C(N, T)$ with the aid of a useful piece of notation. To each N_j there is a set \widehat{V}_j of problems (matrices in our case) that are indistinguishable by N_j . The set \widehat{V}_j , $j = 1, 2, \dots, n$, are nested and eventually reduce to a singleton. Associated with any approximation z is the set $R_j(z)$ of indistinguishable residuals (e.g., $R_j(z) = b - \widehat{A}z$, $\widehat{A} \in \widehat{V}_j$, for the linear equations problem $Ax = b$). The goal is to find the smallest natural number k such that there is a z for which $R_k(z)$ lies in the target ball (e.g., $B(0, \varepsilon\|b\|)$, the ball in R^n centered at the origin with radius $\varepsilon\|b\|$). This is $C(N, T)$.

This formulation reveals several things. First, the admissible algorithms cited in the minimax formulation of $C(N, T)$ are not really needed; what matters is the size of $R_k(z)$ for various z . Second, one reason why there is very little in the NA literature on the problem of finding the minimal k is that for most interesting tasks $k = n$, the sets $R_k(z)$ are just too big, so the problem is not interesting.

One way to reduce the indistinguishable sets is to introduce a subclass F and to use $V_j = \widehat{V}_j \cap F$ in place of \widehat{V}_j . This was discussed above. For approximation theory there is no objection to the introduction of unknown quantities that might characterize F . However, as mentioned above, IBCT seems to use F as a tuning parameter designed to keep $k < n$.

2.3. Spurious challenges. The optimality properties of various Krylov subspace methods are well known (see [Sti, 1958]). IBCT's claim to have something new to add is based on the suggestion that its theory considers any algorithm (confined to Krylov information) including those that give approximations z from *outside* the Krylov subspace. See the quotation in §4.1. The trouble with this apparent novelty is that it is not possible to evaluate the residual norm $\|b - Az\|$ for these external z because there is no known matrix A (only Krylov information). So how can an algorithm that produces z verify whether it has achieved its goal of making $\|b - Az\| < \varepsilon\|b\|$? Perhaps that is why no such new algorithm is actually exhibited? IBCT's suggestion that it goes beyond the well-known polynomial class of algorithms is more apparent than real.

2.4. A new eigenvalue problem. The task of computing some or all the eigenvalues of a matrix is acknowledged to be of practical importance. When only a few eigenvalues of a large order matrix are wanted, one seeks either the smallest eigenvalues, or the largest, or all in a given region. Unfortunately, [Ku, 1986] makes a subtle change in the problem. The redefined goal asks for *any* approximate *eigenpair* (value λ and vector x) without reference to where in the spectrum the approximate eigenvalue may lie. Of course a theorist is entitled to investigate any problem he or she chooses. However, we have yet to hear of *any* use for such output. Our complaint is that *no indication* is given that the goal is an unusual one. Very few readers would realize that the familiar relevant eigenvalue problem has been bypassed. Indeed, we missed the point ourselves until a friend pointed it out.

It is standard practice to use the size of the residual norm ($\|Ax - x\lambda\|$) as the means by which to decide whether a *specified* approximate eigenvalue is accurate enough. IBCT twists things around and makes a small residual norm into the goal. It is the old philosophical error of mistaking the means for the end.

2.5. A confusion of worst cases. An important feature of Krylov information $\{b, Ab, A^2b, \dots\}$ (see the third paragraph in §2.1) is the so-called starting vector b which is, of course, quite independent of the goal of computing eigenvalues. There are two different factors that can increase the cost of using this information to approximate eigenvalues of A . One is an intrinsic difficulty; some matrices in the given class may have unfortunate eigenvalue distributions. The other is that b may be poorly chosen. Instead of separating the effects of these factors, the eigenvalue paper combines them and ends up analyzing Krylov information with a worst possible starting vector even though satisfactory starting vectors are easy to obtain. The fact that b is treated as prescribed data is quite difficult to spot. This situation is in stark contrast to the linear equations problem where b is part of the problem of solving $Ax = b$.

The study of worst choices for b is not without interest (see [Sc, 1979], for example). Such studies are relevant to the inverse eigenvalue problem but not to the complexity of approximating eigenvalues via Krylov subspaces.

Section 5 discusses the issue in more detail, but the conclusion is that the wrong placing of b twists the model away from its original target. It is only this distorted model that permits the proof of the surprising results in [Ku, 1986, Abstract].

3. PRELIMINARIES

3.1. A word on matrix computations. The subject begins with three basic tasks.

- (i) Solve systems of linear algebraic equations; written as $Ax = b$ with x , b column vectors and A a matrix.
- (ii) Compute eigenvalues and eigenvectors of A .
- (iii) Solve least squares problems, i.e. minimize $\|b - Ax\|_2$ over all x .

There are very satisfactory programs for accomplishing these tasks when the matrices are small. An $n \times n$ matrix is said to be *small* if a couple of $n \times n$ arrays can be comfortably stored in the fast memory of the computer. These days a 50×50 matrix would be considered small on most computer systems. The reader may consult [Pa, 1984] for more information. The methods in use make explicit transformations on the given matrix. There are one or two open problems concerning convergence of some methods, but by and large the small matrix problem is in satisfactory condition with respect to conventional one-calculation-at-a-time (sequential) computers.

One reason for introducing this slice of history into the discussion is to bring out the fact that computation with large order matrices (say, 5000×5000) is a somewhat different game from computation with small ones. Sometimes the very problem itself changes. For tasks (i) and (iii) the goal does remain the same but often the product $A\nu$, for any vector ν , can be formed cheaply, so one seeks methods that exploit this feature and do not factor A . For the eigenvalue problem, there are many applications where only a few eigenvalues are wanted, perhaps 30 of 5000, and it is desirable to avoid changing A at all. Thus the task has changed; there is no desire to diagonalize A .

For all three problems it often happens that a sequence of similar cases must be treated as parameters are changed. This leads to an updating problem and ends our brief historical digressions.

As usual R denotes the real numbers, C the complex numbers, and R^n is the vector space of real n -tuples.

3.2. A word on information-based complexity. We describe a simple version of IBCT that is used in the two papers to be examined. It does not use the full panoply of concepts in the monograph or its sequel [Tr, Wo & Wa, 1983].

There are a few essential ingredients that are best seen in a specific context.

- (1) A class F , e.g., $F = \{B: B \in R^{n \times n} \text{ symmetric, positive definite}\}$;
- (2) A task, e.g., given $b \neq 0$ in R^n , and $\varepsilon > 0$, find x in R^n s.t. $\|Ax - b\| < \varepsilon\|b\|$, $A \in F$. Here $\|\cdot\|$ is some given norm.

For the eigenvalue problem, the task is stated as: find $x \in C^n$ and $\rho \in C$ such that $\|\tilde{A}x = x\rho\| \leq \varepsilon$ for all \tilde{A} in F that are indistinguishable from A . The defects in this definition were mentioned in §2.5.

- (3) information $N = (N_0, N_1, \dots)$, e.g., $N_j(A, b) = \{b, Ab, \dots, A^j b\}$ for natural numbers $j \leq n$, $A \in F$;
- (4) A measure of cost, e.g., j units for N_j . In this model the forming of linear combinations of vectors is free.

Items (2) and (3) do not make it clear that A is not known explicitly. There is more discussion of this point in §4.3.

To use an algorithm, restricted to the information N , in order to solve a problem in F will entail cost that may vary with the problem. The primary goal of worst-case IBCT is to minimize, over all such algorithms, the maximum, over all problems in F , of that cost. Determining the minimum, even roughly, is worthwhile even if an algorithm that achieves the minimum is not known. There is an analogous average-case theory.

This certainly has an appeal.

Please note that, in our example, (3) puts this theory firmly within numerical analysis. This is because the information in this example, and it is typical, is *not* part of the task. The information N_j will only be available if a certain type of method is invoked. Consequently the theory is *not* addressing the intrinsic cost, or difficulty, of solving linear systems but is confined to seeking the number of needed steps within a chosen class of methods. This is what numerical analysts do, and have done, from the beginning.

In the 1970s the word “complexity” was reserved for the intrinsic difficulty of a task and the word “cost” was used in connection with a particular algorithm. For example, see [Bo & Mu, 1975] and [Wi, 1980]. However, it is now common to talk about the complexity of an algorithm as a synonym for cost. This extension of the term complexity does no great harm. What is misleading is that the notion of information *appears* to be independent of any algorithm. This allows the theory to talk about the set of all algorithms that confine themselves to the given information. As indicated in the previous paragraph, this way of talking may sometimes be a reformulation of the standard practice of optimizing over a specified family of methods.

For $j < n$ the information $N_j(A, B)$ is partial; there are many matrices that are indistinguishable from A in the sense that each of them generates the same set of $j + 1$ vectors. The basic technical concept in the theory, the *matrix index* $k(\Phi, A)$ of an algorithm Φ , is the minimal cost using Φ to guarantee achievement of the task for all matrices in F that are indistinguishable from

A. There is more discussion in §4.2 and §2.2.

The theory seeks $\min_{\Phi} k(\Phi, A)$ and other related quantities while Φ is restricted to information N . This minimum is the complexity of the task.

IBCT is careful to distinguish itself from the older speciality now called arithmetic complexity, which is concerned with discrete problems such as the minimal number of basic arithmetic operations required to form the product of two $n \times n$ matrices (see, for example, [Ra, 1972; Str, 1969; Sch & Str, 1971; Wi, 1970]). Another way to model some aspects of scientific computing was introduced in 1989 by Blum, Shub, and Smale, [Bl, Sh & Sm, 1989]. Their algebraic complexity has the cost of a basic arithmetic operation independent of the operands just as it does in scientific computation. They analyze computation over rings other than Z and so include R and C .

3.3. Krylov subspaces. Here we sketch the conventional wisdom on this topic. These subspaces of R^n are defined by

$$K^j = K^j(A, b) = \text{span}(b, Ab, \dots, A^{j-1}b).$$

There is no loss in assuming that $\dim K^j = j$. Information N_j permits computation of any vector v in K^{j+1} , not just K^j , at no further cost provided that the coefficient γ_i in $v = \sum_{i=0}^j \gamma_i(A^i b)$ are known.

On the practical side the dominant question for the experts has been how to obtain a computationally satisfactory basis for K^j . Round off error destroys the expected linear independence of the computed vectors. Some researchers maintain that it is more efficient to use a highly redundant spanning set rather than a basis. Others recommend the additional expense of computing an orthonormal basis. In any case, it is the computation of the basis or spanning set, along with multiplication of vectors by A , that is the expensive part of the computation. The model of cost used in IBCT theory reflects this quite well. It is the number of *steps* that matters. We think of a step as adding one more vector to the Krylov sequence $\{b, Ab, A^2b, \dots\}$.

One of the founders of Krylov space methods, C. Lanczos [La, 1952], proposed a useful basis for K^j with the property that the projection of symmetric A onto K^j is a symmetric tridiagonal $j \times j$ matrix T_j . Tridiagonal matrices are easy to handle. With a basis in hand, there are diverse tasks that can be accomplished. Here is a partial list.

- (i) Compute an approximation $x^{(j)}$ to $A^{-1}b$ such that $x^{(j)} \in K^j$ and its residual $b - Ax^{(j)}$ is orthogonal to K^j . It so happens that $\|b - Ax^{(j)}\|$ may be computed without forming $x^{(j)}$ so there is no need to compute unsatisfactory $x^{(j)}$. When $A \in \text{SPD}$ (symmetric positive definite matrices) then $x^{(j)}$ coincides with the output of the conjugate gradient algorithm.
- (ii) Compute the vector $u^{(j)}$ that minimizes $\|b - Av\|$ over all $v \in K^j$ (not K^{j+1}). This is the MR (minimum residual) algorithm. The extra vector $A^j b$ is needed to ascertain the coefficients in the expansion of $u^{(j)}$.
- (iii) Compute some, or all, of the Rayleigh-Ritz approximations (θ_i, y_i) , $i = 1, \dots, j$, to eigenpairs of symmetric A . Here $\theta_i \in R$ and $\{y_1, \dots, y_j\}$ is an orthonormal basis for K^j . For each i , $Ay_i - y_i\theta_i$ is orthogonal to K^j .

Krylov subspace methods are not really iterative. All the basic tasks mentioned in §3.1 are solved exactly (in exact arithmetic) in at most n steps. However the interest in this approach is due to the fact that in many instances far fewer than n steps are required to produce acceptable approximations. In other words, to take n steps is the practical equivalent of failure. However, for each basic task there are data pairs (A, b) that do require n steps even for loose tolerances such as $\varepsilon = 10^{-3}$, so research has focussed on explanations of why so often the cost is much smaller. The gradual realization of the efficacy of changing the $Ax = b$ problem to an equivalent one via a technique called preconditioning has enhanced the use of Krylov subspace methods.

In a sense the *convergence* of all these methods is completely understood in the symmetric case and is embodied in the error bounds published by Kaniel and improved by Paige and Saad (see [Ka, 1966; Sa, 1980; Pa, 1980] for the details). The error depends on two things; the eigenvalue distribution of A and the components of the starting vector b along the eigenvectors. Of course, all this analysis supposes a given matrix A , not a set of indistinguishable matrices.

From this conventional viewpoint the thrust of these two complexity papers is to see to what extent the standard algorithms (CG, MR, Lanczos) do not make best use of the information on hand. Recall that N_j contains an extra vector not in K^j . This is a reasonable project and the results can be expressed within the usual framework. The term “complexity theory” appears to a numerical analyst like window dressing.

4. ON THE OPTIMAL SOLUTION OF LARGE LINEAR SYSTEMS

These sections offer a description of and commentary on [Tr & Wo, 1984].

4.1. **Spurious generality.** Here is a quotation from the introduction:

We contrast our approach with that which is typical in the approximate solution of large linear systems. One constructs an algorithm Φ that generates a sequence $\{x_k\}$ approximating the solution $\alpha = A^{-1}b$; the calculation of x_k requires k matrix-vector multiplication and x_k lies in the Krylov subspace spanned by $b, Ab, \dots, A^k b$. The algorithm Φ is often chosen to guarantee good approximation properties of the sequence $\{x_k\}$. In some cases Φ is defined to minimize some measure of the error in a *restrictive* class of algorithms. For instance, let this class be defined as the class of ‘polynomial’ algorithms; that is

$$\alpha - x_k = W_k(A)\alpha, \quad \text{where } W_k(0) = 1.$$

Here W_k is a polynomial of degree at most k .

...

[Some omitted sentences define the minimum residual and conjugate gradient algorithms.]

It seems to us that this procedure is unnecessarily restrictive. It is not clear, a priori, why an algorithm has to construct x_k of the form $\alpha - x_k = W_k(A)\alpha$. Indeed, we show that for orthogonally invariant classes of matrices the minimum residual

algorithm (MR) is within at most one matrix vector multiplication of the lower bound without any restriction on the class of algorithms. However, if the class is not orthogonally invariant, the optimality of MR may disappear.

Our first point was made earlier. The information N does not come with the linear equations problem. The brief answer to the quoted rhetorical question (why must an algorithm construct x_k of the given form?) that serves to justify the whole paper is the following. To any vector x not in the Krylov subspace K^k , there is an admissible matrix A , such that the residual norm is as large as you please. This holds even when A is required to be symmetric and positive definite. An admissible matrix A is one that is consistent with the Krylov information (more on this below).

4.2. Definitions and optimality. In this section we put down the definitions made in [Tr & Wo, 1984]. Our comments are reserved for the next section.

- (i) Let F be a subclass of the class $GL(n, R)$ of $n \times n$ nonsingular real matrices.
- (ii) Let $b \in R^n$ with $\|b\| = (b, b)^{1/2} = 1$ be given. For $0 \leq \varepsilon < 1$, find $x \in R^n$ such that $\|b - Ax\| \leq \varepsilon$ (it would have been kinder to add $A \in F$).
- (iii) Krylov information: $N_j(A, b) = \{b, Ab, \dots, A^j b\}$, $j = 0, 1, \dots$.
- (iv) Measure of cost: N_j costs j units.
- (v) An algorithm $\Phi = \{\phi_j\}$ is a sequence of mappings $\phi_j: N_j(F, R^n) \rightarrow R^n$.
- (vi) The set of indistinguishable matrices for given $N_j(A, b)$:

$$V(N_j(A, b)) = \{\tilde{A} : \tilde{A} \in F : N_j(\tilde{A}, b) = N_j(A, b)\}.$$

- (vii) The matrix index of an algorithm Φ :

$$k(\Phi, A) = \min \left\{ j : \max_{\tilde{A} \in V(N_j)} \|b - \tilde{A}x_j\| \leq \varepsilon \right\}, \quad N_j = N_j(A, b),$$

where $x_j = \Phi_j(N_j(A, b))$. If the set of j values is empty, then $k(\Phi, A) = \infty$.

- (viii) The class index of an algorithm

$$\Phi: k(\Phi, F) = \max_{B \in F} k(\Phi, B).$$

- (ix) The optimal matrix index: $k(A) = \min_{\Phi} k(\Phi, A)$ over Φ restricted to N .
- (x) The optimal class index: $k(F) = \max_{B \in F} k(B)$.
- (xi) Strong optimality: Φ is strongly optimal iff $k(\Phi, B) = k(B)$, for each $B \in F$.
- (xii) Optimality: Φ is optimal iff $k(\Phi, F) = k(F)$.
- (xiii) Almost strong optimality: Φ is almost strongly optimal iff $k(\Phi, B) \leq k(B) + c$, for every $B \in F$, for some small integer c .

Remark 1. Since $A^i b = A(A^{i-1} b)$, it follows that Krylov information $N_j(A, b)$ requires j applications of the operator A . That is why the cost is given as j units. In practice, one uses better bases for the Krylov subspace K^j than is

provided by $N_j(A, b)$; but for theoretical purposes, this modification may be ignored.

Remark 2. It can happen that $k(A) \ll k(F)$. For this reason, it is of interest to find algorithms with small matrix index.

Remark 3. For simplicity, the dependence of all concepts on n , N_j , b , and ε is suppressed. The idea is to compute $k(A)$ and $k(F)$ for interesting classes F and to find strongly optimal or optimal algorithms if possible.

4.3. Discussion of the basic concepts. In §2 we pointed out how misleading it can be to compute complexity for restricted classes F , which are difficult to discern in practice. Here we wish to point out that F is introduced into the basic definitions, such as V in (vi), and there is no need for it.

To add to any confusion, the basic definitions do not make clear the role of A . In the context of numerical analysis there is a particular matrix A on hand and this permits one to test the residual $r = b - Av$ for any vector v . However, in the context of IBCT, that is not quite the case. In this game we consider a specific A , but it is not available explicitly. That odd situation certainly warrants some discussion, and it faithfully reflects the state of affairs at a typical step in a Krylov subspace method. The matrix is hidden inside a subprogram, and the user has only a basis for the Krylov subspace corresponding to Krylov information $N_j(A, b) = \{b, Ab, \dots, A^j b\}$. Associated with $N_j(A, b)$ is

$$\widehat{V}(N_j(A, b)) = \{\tilde{A} : N_j(\tilde{A}, b) = N_j(A, b)\},$$

the set of matrices indistinguishable from A by $N_j(A, b)$. Contrast \widehat{V} with V in §4.2 (vi).

With \widehat{V} defined above, the natural definition of the matrix index of an algorithm Φ is

$$\hat{k}(\Phi, A) = \min \left\{ j : \max_{\tilde{A} \in \widehat{V}(N_j)} \|b - \tilde{A}x_j\| \leq \varepsilon \right\}, \quad N_j = N_j(A, b),$$

where

$$x_j = \Phi_j(N_j(A, b)).$$

If the set of j values is empty then $\hat{k}(\Phi, A) = \infty$. Please note that in contrast to (vii) in §4.2 there are no hidden parameters in \hat{k} . It is the first step in the process at which the task is accomplished by Φ for all matrices indistinguishable from A by N_1 through N_k . Then the optimal index is

$$\hat{k}(A) = \min_{\Phi} \hat{k}(\Phi, A)$$

over all Φ such that Φ_j uses only $N_j(A, b)$ and standard arithmetic operations.

There is no logical need for F . However, given a class F , one may define

$$\hat{k}(\Phi, F) = \max_{B \in F} \hat{k}(\Phi, B); \quad \hat{k}(F) = \min_{\Phi} \hat{k}(\Phi, F).$$

Why did IBCT not follow this simple approach? Why does IBCT use $V = \widehat{V} \cap F$ to define the matrix index $k(\Phi, A)$ and thus suppress the role of F ? The reason, we suspect, is that with these natural definitions the “polynomial”

algorithms, deemed to be restrictive in the introduction to [Tr & Wo, 1984] are mandatory; and, consequently, IBCT has nothing new to offer. Here is a result of ours that shows why the nonpolynomial algorithms are of no interest in worst case complexity, i.e., $\hat{k}(\Phi, A) = \infty$.

Theorem. Assume $A = A^t \in R^{n \times n}$. Let $K^j = \text{span}(b, Ab, \dots, A^{j-1}b)$ have dimension j ($< n$) and $\|b\| = 1$. To each $v \notin K^j$ there exists $\tilde{A} \in \hat{V}(N_j(a, b))$ such that $\|b - \tilde{A}v\| > 1$.

Sketch of proof. (1) Should b happen to be orthogonal to some eigenvectors of A , it is always possible to choose an $\bar{A} \in \hat{V}(N_j(A, b))$ such that b is not orthogonal to any of \bar{A} 's eigenvectors. If necessary replace A by \bar{A} .

- (2) There is a distinguished orthonormal basis for K^j that can be extended to a basis for R^n and in which A is represented by the matrix

$$\begin{bmatrix} T & E^t \\ E & U \end{bmatrix}$$

where $T = T^t \in R^{j \times j}$, E is null except for a $\beta \neq 0$ in the top right entry, U is unknown. Moreover T and β are determined by $N_j(A, b)$.

- (3) In this distinguished basis b is represented by $e_1 = (1, 0, 0, \dots, 0)^t$ and let any $v \in K^j$ be represented in partitioned form (f, g) where $f \in R^j$, $g \in R^{n-j}$. By hypothesis $g \neq 0$, since $v \notin K^j$, and

$$\|b - \tilde{A}v\|^2 = \|e_1 - Tf - E^t g\|^2 + \|Ef + \tilde{U}g\|^2,$$

where \tilde{U} is the $(2, 2)$ block in the representation of \tilde{A} and is unconstrained.

- (4) To each $\tilde{U} \in R^{(n-j) \times (n-j)}$ there is an $\tilde{A} \in \hat{V}(N_j(A, b))$ and for any $g \neq 0$, there exists \tilde{U} such that

$$\|Ef + \tilde{U}g\|^2 = \|e_1 \beta f(j) + \tilde{U}g\|^2 > 1.$$

In particular, it is possible to select \tilde{U} to be symmetric, positive definite, and diagonal. Here ends the sketch of the proof.

Symmetry is not needed in the result given above, If A is not symmetric there is still a distinguished orthonormal basis for $K^j(A, b)$ and R^n such that b is represented by e_1 , and A is represented by

$$\begin{bmatrix} H & J \\ E & L \end{bmatrix}.$$

Now H , E , and the first column of J are determined by $N_j(A, b)$. Moreover, $Je_1 \neq 0$ and for \tilde{A} indistinguishable from A we have

$$\|b - \tilde{A}v\|^2 = \|e_1 - Hf - Jg\|^2 + \|e_1 \beta f(j) + \tilde{L}g\|^2.$$

This can be made as large as desired. In the language of IBCT, $\hat{k}(\Phi, A) = \infty$ for any Φ such that $\Phi(N_j(A, b))$ takes values outside $K^j(A, b)$.

Only two choices were left to IBCT. Either turn away to unsolved problems or cut down the number of indistinguishable matrices by using

$$V = \widehat{V} \cap F$$

instead of \widehat{V} .

Here is the quandary for IBCT. If F is chosen too small, the model loses realism. If F is allowed to be large, then the standard “polynomial” regime is optimal.

4.4. Discussion of results. The main result of the paper concerns the minimal residual algorithm MR: this polynomial algorithm’s output for the information $N_j = \{b, Ab, \dots, A^j b\}$ is the vector in the Krylov subspace $K^j = \text{span}(b, Ab, \dots, A^{j-1}b)$ that minimizes the residual norm, so MR is optimal in K^j . Given N_j MR needs $k(\text{MR}, A)$ steps to achieve an ε -approximation in the worst case. However, IBCT says

Theorem 3.1 [Tr & Wo, 1984]. *If F is orthogonally invariant then*

$$k(\text{MR}, A) \geq k(A) \geq k(\text{MR}, A) - 1 \text{ for any } A \in F.$$

Furthermore, both the upper and lower bounds can be achieved.

Recall that $k(A)$ is the minimal number of steps over all admissible algorithms.

The fact that MR is *not* always strongly optimal for the given information appears to give substance to the theory. It will astonish the numerical analyst, so let us look at the example that purports to show that MR is not always optimal for the given information (Example 3.2 in the paper). This class is

$$\widetilde{F}_\rho = \{A : A = I - B, B = B^t, \|B\| \leq \rho < 1\}.$$

When ε , ρ , and n are specially related so that

$$q(\varepsilon) = \left\lceil \frac{\ln((1 + (1 - \varepsilon^2)^{1/2})/\varepsilon)}{\ln((1 + (1 - \rho^2)^{1/2})/\rho)} \right\rceil < n,$$

then MR is just beaten by another polynomial algorithm, called the Chebyshev algorithm, because

$$k(\text{Cheb}, A) = q(\varepsilon), k(\text{MR}, A) = q(\varepsilon) + 1.$$

A word of explanation is in order. Recall that $A^j b \in N_j(A, b)$, but $A^j b \notin K^j$. The MR algorithm needs $A^j b$ to compute the coefficients γ_i in

$$\text{MR}(N_j(A, b)) = \sum_{i=0}^j \gamma_i(A^i b).$$

This always beats the Cheb output from K^j . However, Cheb can use the well-known three-term recurrence, based on ρ , to obtain its equioscillation approximation from K^{j+1} , not just K^j . With the right relation between ε , ρ , and n , one has

$$\|b - A \text{MR}(N_{q(\varepsilon)})\| > \|b - A \text{Cheb}(N_{q(\varepsilon)})\| > \|b - A \text{MR}(N_{q(\varepsilon)+1})\|.$$

Is it fair to compare them? The theory claims to compare algorithms restricted solely to information N_j . So how could the Cheb algorithm obtain the

crucial parameter ρ ? The answer is that ρ is found in the definition of the problem class \tilde{F}_ρ ! In other words, knowledge that Cheb can use is passed along through the problem class, not the information.

The important point we wish to make is not just that comparisons may not be fair but that the results of IBCT tell us as much about the idiosyncrasies of its framework as they do about the difficulty of various approximation problems. With a realistic class such as SPD (sym. pos. def.), MR is optimal (strongly) as it was designed to be and as is well known.

In more recent work [Wo, 1985; Tr & Wo, 1988], the flaw mentioned above appears to be corrected and the parameter ρ is put into the information explicitly. Again Cheb wins by 1 because it uses ρ while MR does not. However, this new clarity comes at the expense of realism; the Krylov information is scrupulously priced while ρ comes free. Yet membership in \tilde{F}_ρ may be more difficult to ascertain than the approximate solution.

Although the IBCT paper does not mention the possibility, Krylov information may be used to obtain lower bounds on ρ that get increasingly accurate as the dimension of the Krylov subspace increases. Algorithms that exploit knowledge of the spectrum will have good average behavior but there is little room for improvement in the worst case.

The simple facts are well known: Chebyshev techniques are powerful and users are willing to do considerable preliminary work to estimate parameters such as ρ . It is not clear, and depends strongly on the data, when it is preferable to use a weaker method such as MR that does not need extra parameters. The result that MR is only almost strongly optimal is a striking example of obfuscation. The framework of IBCT permits unnatural comparison of algorithms.

Embracing the conjugate gradient algorithm. In §4 of their paper the authors generalize the framework to cover other known methods such as the conjugate gradient algorithm. All that is necessary (for IBCT) is to introduce a parameter p into the basic task. Now an ε -approximation to $A^{-1}b$ is redefined as any $x \in R^n$ that satisfies

$$\|A^p(x - A^{-1}b)\| < \varepsilon\|A^{p-1}b\|.$$

The cases $p = 0, 1/2, 1$ are the most important, and when p is not an integer, it is appropriate to restrict attention to the symmetric, positive definite subset (SPD) of $R^{n \times n}$. When $p = 1$ we recover MR and the results of §3. To generalize MR to $p = 0$, it is necessary to use the normal equations of a given system. The new feature, slipped in without mention, is that with $p < 1$ the right-hand side of the new definitions is *not directly computable*. So how does an algorithm know when to terminate? Please note that approximation theory can present results that are not computable without a blush because it merely exhibits relationships. In computation, however, there is no virtue in attaining the desired accuracy if this fact cannot be detected. Since IBCT defines an algorithm as a sequence of mappings, it dodges the difficult task of knowing when to stop.

The fact that the Conjugate Gradient algorithm ($p = 1/2$) minimizes $\|b - Ax\|_A$ at each step is well known (see [Da, 1967]). Nevertheless, in practice, the algorithm is usually terminated by the less desirable condition $\|b - Ax\| < \varepsilon\|b\|$ because the desirable A norm is not available (see [Go & Va, 1984]).

For the reasons given in the two previous paragraphs, Theorem 4.2 in [Tr & Wo, 1984], which states that under certain technical conditions the generalized MR algorithm ($0 \leq p \leq 1$) is almost strongly optimal, is a true theorem about improper algorithms.

4.5. An interesting result. Recall that

$$\begin{aligned} N_j(A, b) &= \{b, Ab, \dots, A^j b\}, \\ K^{j+1} &= \text{span}\{b, Ab, \dots, A^j b\}, \\ \widehat{V}(N_j(A, b)) &= \{\widetilde{A} : N_j(\widetilde{A}, b) = N_j(A, b)\}. \end{aligned}$$

Theorem (reformulated by us from Theorem 3.1 in [Tr & Wo, 1984]). *If $y \in R^n$ yields an ε residual norm, $\|b - \widehat{A}y\| < \varepsilon\|b\|$, for all $\widetilde{A} \in \widehat{V}(N_j(A, b))$ then so does its orthogonal projection z onto K^{j+1} .*

We offer a simplified version of the argument in [Tr & Wo, 1984].

Proof. There are five steps.

(i) Either $y \in K^{j+1}$, and there is nothing more to prove, or $y = z + w$, with $z \in K^{j+1}$, and $0 \neq w$ orthogonal to K^{j+1} .

(ii) For the vector w defined in (i) there is a unique symmetric orthogonal matrix $H = H(w)$, called the *reflector that reverses w* . In particular: (a) $Hx = x$, for x orthogonal to w , (b) $Hw = -w$.

Define an auxiliary matrix \widehat{A} by

(iii) $\widehat{A} = HAH \in V(N_j(A, b))$ since, by use of (ii)(a), $\widehat{A}^i b = HA^i Hb = HA^i b = A^i b$, $i = 1, \dots, j$.

Note that

$$\begin{aligned} \widehat{A}y - b &= HAHy - b \\ &= HA(z - w - b), \quad \text{using (i) and (ii)(b).} \\ &= H(Az - b - Aw), \quad \text{using } Hb = b. \end{aligned}$$

This shows the crucial relationship

(iv) $\|\widehat{A}y - b\| = \|Az - b - Aw\|$, since H preserves norms.

Hence

(v)

$$\begin{aligned} \|Az - b\| &\leq \frac{1}{2}\{\|Az - b - Aw\| + \|Az - b + Aw\|\}, \\ &\quad \text{by the triangle inequality,} \\ &= \frac{1}{2}\{\|\widehat{A}y - b\| + \|Ay - b\|\}, \quad \text{by (iv), and (i),} \\ &\leq \varepsilon\|b\|, \quad \text{using the hypothesis.} \end{aligned}$$

Recall that z is y 's projection onto K^{j+1} . Hence $\widetilde{A}z = Az$ for all $\widetilde{A} \in \widehat{V}(N_j(A, b))$ and so

$$\|\widetilde{A}z - b\| = \|Az - b\| \leq \varepsilon\|b\|, \quad \text{by (v). Q.E.D.}$$

This theorem explains why MR cannot lag more than one step behind any algorithm that produces an ε residual norm for *all* matrices indistinguishable

from A . For, *by definition*, MR produces from $N_{j+1}(A, b)$ (note the increased subscript) the unique vector in K^{j+1} that gives the smallest residual and is at least as good as the vectors y and z defined in the proof above. But y could be the output of a rival algorithm.

Our formulation of the lemma omits any mention of the class F . Now it is clear why the hypothesis that F should be orthogonally invariant appears in most of the theorems. Recall that $\widehat{V}(N_j(A, b))$ is *too big*. To cut down the number of *indistinguishable* matrices, the theory uses $\widehat{V} \cap F = V$. To make use of the theorem, it is necessary to have $H A H \in F$ and this will be true provided that F is orthogonally invariant.

4.5. Summary. One hidden defect in the framework for discussing the MR algorithms is the far reaching feature that allows the family F to convey what most people would call free information behind the back of the information operator N .

More disturbing than the previous defect is that we cannot see how any algorithm other than the well-studied polynomial algorithms could know when it had achieved an ε -approximation if it is restricted to the given information. This gives rise to a feeling that [Tr & Wo, 1984] managed to create an artificial problem where no real puzzle exists. The mentioned theorems (3.1 and 4.2) reflect only the propensity of their general theory of optimal algorithms for creating such situations.

5. OPTIMAL SOLUTION OF LARGE EIGENPAIR PROBLEMS

This section offers a description of and commentary on [Ku, 1986]. The paper demonstrates cleverness and clean exposition but, nevertheless, suffers from design flaws; it equates different versions of a given algorithm and it redefines a standard task. From the abstract:

The problem of approximation of an eigenpair of a large $n \times n$ matrix A is considered. We study algorithms which approximate an eigenpair of A using the partial information on A given by $b, Ab, A^2b, \dots, A^j b, j \leq n$, i.e., by Krylov subspaces. A new algorithm called the generalized minimal residual algorithm (GMR) is analyzed. Its optimality for some classes of matrices is proved. We compare the GMR algorithm with the widely used Lanczos algorithm for symmetric matrices. The GMR and Lanczos algorithms cost essentially the same per step and they have the same stability characteristics. Since the GMR algorithm never requires more steps than the Lanczos algorithm, and sometimes uses substantially fewer steps, the GMR algorithm seems preferable.

... The Fortran subroutine is also available via ...

This last phrase shows that the subject matter is firmly within the field of numerical analysis. Implementation issues concerning GMR are described in [Ku, 1985].

5.1. A subtle change of goal. Here are the first five lines of the paper. "Suppose we wish to find an approximation to an eigenpair of a very large matrix A . That is, we wish to compute (x, ρ) , where x is an $n \times 1$ normalized vector,

$\|x\| = 1$, and ρ is a complex number s.t.

$$(1.1) \quad \|Ax - x\rho\| < \varepsilon$$

for a given positive ε . Here $\|\cdot\|$ denotes the 2-norm."

It is all too easy to assent to this statement of the problem and pass on to the rest of the article. However it is *not* the normal eigenvalue problem. We are not aware of any demand at all for the accomplishment of this particular task. The users of eigenvalue programs (engineers, theoretical chemists, theoretical physicists) want eigenvalues in *specified parts of the spectrum*; occasionally, they want the whole spectrum. The main concern of this article is with symmetric matrices; and because their eigenvalues are real, the usual demands are for the leftmost p eigenvalues (for some $p \leq n$) or the rightmost p eigenvalues or for all eigenvalues in a given interval. Eigenvectors may or may not be wanted. There is nothing inherently wrong with restricting attention to the rather special case $p = 1$ and a few articles (not cited by Kuczynski) have been devoted to it (see [O'L, Ste & Va, 1979; Pa, Si & Str, 1982] for the details).

To support our description of user's demands we refer to three publications from different fields, [Cu & Wi, 1985, Introduction; Je, 1977, Ch. 7; Sh, 1977, §6].

The consequences of leaving out a vital aspect of the usual task are most serious, precisely when one seeks optimal performance.

One reason why it is so easy to overlook the omission in the problem statement is that, for symmetric matrices, almost everyone does use the residual norm $\|Ax - x\rho\|$ to judge the accuracy of an approximate eigenpair (x, ρ) . However, it is not very interesting to minimize the residual norm if that might yield a ρ in the unwanted part of the spectrum. Now a pure mathematician is free to define his goal at will. What is regrettable is that no hint is given to the reader that the goal is not standard.

We say more about the ε appearing in (1.1) in §5.5.

We mention one other fact which may be news to readers who are not much concerned with eigenvalue problems. It suggests why the direction taken by Kuczynski has not appeared in the literature before. If we are given a symmetric matrix A and seek a single eigenvalue (with or without its eigenspace) then the wanted eigenvalue is almost certain to be either the leftmost or the rightmost. Recall that the Rayleigh quotient of a column vector $x \neq 0$ in R^n is the number $x^t Ax / x^t x$. The extreme eigenvalues of A are the extreme values of the Rayleigh quotient over all possible vectors x in R^n . It happens that, for the given Krylov information N_j , the Lanczos algorithm is optimal for this task in the strong sense that it yields the leftmost and rightmost values of the Rayleigh quotient over all vectors in the *available* space K^j . The last vector $A^j b$ in N_j is needed to ascertain the extreme values over K^j . Thus the problem is settled. It is a pity that this well-known fact was not mentioned.

5.2. Choosing a bad starting vector. The particular aspect of Information-Based Complexity Theory adopted in the paper under review is called worst-case complexity. It seeks bounds on the cost of computing ε -approximations over all matrices in certain classes F and over *all* starting directions b . Theorems 3.1, 3.2, 4.1, 5.1 (there is no theorem 1.1 or 2.1) in [Ku, 1986] are examples. In particular, the theory must cover what can happen with the *worst possible* starting vector. Theorem 3.1 is quoted in full in §5.4.

There is nothing wrong with studying the worst case. Indeed it has already been done. [Sc, 1979] is a paper with the clear title *How to make the Lanczos algorithm converge slowly* in which the author gives formulae for a starting vector that prevents any Rayleigh Ritz approximation from converging until the final step! Scott's paper appeared in *Mathematics of Computation*, the American Mathematical Society's principal outlet for numerical analysis, but it is not referenced in [Ku, 1986]. The fact that *some*, Krylov subspaces can be very badly aligned with A 's eigenvectors does prevent worst-case analysis from shedding much light on how Krylov subspaces approach certain eigenvectors in the usual case of a random starting vector. That study, of course, comes under average-case analysis and is ripe for attention.

Please note that this comment is quite independent of comparisons of GMR and Lanczos. The point is this: the starting vector b is a *free parameter* in the eigenvalue problem (in contrast to the linear equations problem $Ax = b$). It is not given and may be chosen to *improve* performance. In the absence of extra information, it is the almost universal habit to pick b with the aid of a random number generator. Recent theoretical work on Lanczos has been concerned to explain why this choice is so powerful (see [Sa, 1980; Pa, 1980]). Note that two quite different situations have been pushed together under the label 'worst case'. It is quite normal to consider the most difficult matrices A because they are part of the problem. On the other hand, a bad b is a self-inflicted handicap rather than a genuine difficulty. It is the confounding of these cases that is unfortunate, not their study.

Returning to the eigenvalue problem, we can rephrase our complaint as follows: Kuczynski's focus, perhaps unwittingly, is on *Krylov-subspaces-with-worst-possible-starting-vectors*. What a pity that this was not emphasized! The numerical examples given in the paper are not without interest. The starting vector there, though not perhaps worst possible, is very bad. Both methods, GMR and Lanczos converge very slowly. The chosen matrices are extensions of the one used by Scott to illustrate how bad Rayleigh Ritz approximations can be (see [Pa, 1980, p. 218]).

We ran these examples with our local Lanczos program. It uses a random starting vector, and convergence was quite satisfactory. The results are given in §5.5.

There is a different context in which the focus on worst starting values is much more relevant. The GMR algorithm presented by Kuczynski is a generalization of the MR (minimum residual) algorithm used to compute approximations to $A^{-1}b$. There one seeks vectors x in R^n s.t. $\|Ax - b\| < \varepsilon\|b\|$. A well-chosen subspace may be used to generate approximate solutions at low cost. It is advisable to ensure that the right-hand side is in the chosen subspace, and this consideration leads one to choose the subspace K^j . In this context b is part of the data (A, b) and is not at our disposal. The study of bad b 's is relevant to a study of the complexity of Krylov space methods for *linear equations*. However, it has been appreciated from the beginning that for reasonable ε and unfortunate b then n steps will be required unless A is close to the identity matrix (see [Ka, 1966; Thr. 4.3; and Me, 1963]).

To burden the Lanczos algorithm (or GMR) with *unnecessarily* bad starting vectors for the eigenvalue problem is like studying the performance of Olympic

athletes only when they suffer from some rare affliction like poison ivy.

5.3. Redefining the Lanczos algorithm. The new algorithm GMR is contrasted with the well-known Lanczos algorithm. Here is Kuczynski's definition of the Lanczos algorithm, from p. 142. The subspace A_j is our Krylov subspace K^j .

Perform the following steps:

- (1) Find an orthonormal basis q_1, q_2, \dots, q_j of the subspace A_j ; let $Q_j = (q_1, \dots, q_j)$ be the $n \times j$ matrix.
- (2) Form the $j \times j$ matrix $H_j = Q_j^t A Q_j$; compute eigenpairs of H_j ; $H_j g_i = h_i g_i, (g_i, g_m) = \delta_{im}, i, m = 1, \dots, j$.
- (3) Compute the Ritz vectors $z_i = Q_j g_i$ and the residual

$$r_j^L = \min_{1 \leq i \leq j} \|Az_i - \theta_i z_i\| \quad \text{for } 1 \leq i \leq j.$$

- (4) Define $Z_j = \{z_i, \theta_i\}, i = 1, 2, \dots, j: \|Az_i - z_i \theta_i\| = r_j^L$

The j th step of the L algorithm is defined by

$$\Phi_j^L(N_j(A, b)) = (x_k, \rho_k),$$

where (x_k, ρ_k) is an arbitrary element from Z_j .

The trouble is that steps 3 and 4 have been changed from the usual ones to conform with the idiosyncratic goal discussed in §5.1. However, no mention of this fact is made.

Here is the conventional description wherein it is supposed that p eigenvalues are to be approximated. It is from [Pa, 1980, p. 214].

- (1) Is the same as above.
- (2) Form the $j \times j$ matrix $H_j = Q_j^t A Q_j$; compute the $p(\leq j)$ eigenpairs of H_j that are of interest, say

$$H_j g_i = g_i \theta_i, \quad i = 1, \dots, p$$

The θ_i are Ritz values, $\theta_1 < \theta_2 < \dots < \theta_j$. Equality is not possible.

- (3) If desired, compute the p Ritz vectors $z_i = Q_j g_i, i = 1, \dots, p$. The full set $\{(\theta_i, z_i), i = 1, \dots, j\}$ is the best set of j approximations to eigenpairs of A that can be derived from A_j alone.
- (4) Residual error bounds. Form the p residual vectors $r_i = Az_i - z_i \theta_i$. Each interval $[\theta_i - \|r_i\|, \theta_i + \|r_i\|]$ contains an eigenvalue of A . If some intervals overlap then a bit more work is required to guarantee approximations to p eigenvalues. See [Pa, 1980, §11-5].
- (5) If satisfies then stop.

In the context of Kuczynski's investigations his modification is entirely reasonable; i.e., he selects at each step one Ritz pair with a minimal residual norm. However, it is most misleading to those not familiar with the Lanczos algorithm to suggest that its purpose is simply to product this Ritz pair. In fact, as indicated above, the Lanczos algorithm produces an approximation, albeit crude, to the whole spectrum, namely $\theta_1, \dots, \theta_j$ and the user is free to select from this set to suit the specific goal. Thus, to approximate the right-most eigenvalue, one concentrates on θ_j and continues until its error bound $\|r_j\|$ is satisfactory. In practice, more refined error bounds can be made but that is not germane here (see [Pa & No, 1985]).

It would have been preferable to state the Lanczos algorithm conventionally and then specify the modifications appropriate for the purpose in hand. This

action would make clear that the Lanczos algorithm is not trying to minimize one residual norm. That is why it is inferior to GMR for that purpose.

It is worth pointing out here, that in the model of arithmetic used in these studies, the cost of all the Rayleigh-Ritz approximations and of finding the minimal residual norm over the subspace K^j is taken as nil. It might occur to the reader that in this context it would cost no more per step to compute *all* the Rayleigh-Ritz approximations and use whatever approximations one desires. Thus setting up GMR and Lanczos as competing algorithms is artificial. Moreover, in practice, it is much more expensive to compute the minimal residual than to compute the Rayleigh-Ritz residuals. Kuczynski has devoted a whole report to the task.

5.4. Theoretical results. From p. 138 of [Ku, 1986]. We are ready to formulate the main theory of the paper.

Theorem 3.1. *If F is unitarily (orthogonally) invariant, then the GMR is almost strongly optimal in F , i.e., $k(\Phi^{gmr}, A, b) = \min_{\Phi} k(\Phi, A, b) + a$, for any $(A, b) \in F \times S_n$, where $a \in \{0, 1, 2\}$.*

Here $k(\Phi, A, b)$ is the minimal number of steps j required to guarantee an ε -residual with algorithm Φ over all matrices \tilde{A} that are indistinguishable from A with the given information $N = N_j = [b, Ab, \dots, A^j b]$. The algorithm Φ returns a pair ρ, x whose residual norm $\|\tilde{A}x - x\rho\|$ is to be compared with ε .

Recall that with information N_j , the GMR algorithm, by definition, picks out a unit vector $x \in K^j$ and a $\rho \in C$ that produce the minimal residual norm.

How could any other algorithm possibly do better? Well, there might be special circumstances in which one could deduce a suitable additional component of that last vector $A^j b$ that is not used by GMR in forming x , although $A^j b$ is used in calculating the coefficients of GMR's approximation from K^j . The proof studies this possibility and concludes that GMR would make up any discrepancy in at most two more steps.

The argument is very nice. In many important cases, when A is Hermitian for example, then the constant a in Theorem 3.1 is actually 0.

There are other clever results. Theorem 4.2 shows that for symmetric matrices the residual norm of GMR must be strictly decreasing at least at every other step. Theorem 5.1 yields a beautiful but esoteric fact about Krylov subspaces generated by Hermitian A . For the worst starting vector there is a unit vector v in K^j such that

$$\frac{\|A\|}{2j} \leq \|Av - v\rho\| \leq \frac{\|A\|}{j},$$

for $j < n$.

As indicated above, these nice results from approximation theory do not add up to a case for replacing Rayleigh-Ritz with some rival algorithm. We sometimes abbreviate Rayleigh-Ritz by R-R.

5.5. Numerical examples. All the numerical results reported in [Ku, 1986] concern symmetric tridiagonal matrices with starting vector e_1 (the first column of the identity matrix I). This starting vector ensures that the Lanczos algorithm reproduces the original matrix. At this point we should recall that the original goal of the Lanczos algorithm was to reduce a symmetric matrix to tridiagonal

form. So, the numerical results to be seen below do not relate to the Lanczos recurrence itself but merely indicate alternative rules for stopping. With the goal of tridiagonalization, the algorithm always stopped at step n . Later, it was realized that excellent approximations to a few eigenvectors were usually obtained early in the process. There is no single stopping criterion for current Lanczos algorithms; termination depends on what the user wants (see [Cu & Wi, 1985; Pa, 1980; and Go & Va, 1984]).

Kuczynski provides the Lanczos algorithm with a stopping criterion to suit his purposes, but his algorithm GMR could have been called (with more justice) the Lanczos algorithm with a new stopping criterion. It uses the minimum residual in the whole Krylov subspace instead of the usual (cheap) Rayleigh-Ritz approximations. So, the numerical results simply indicate the effect of these different termination criteria.

The first batch of results concern tridiagonals with nonzero elements chosen at random from $[-\frac{1}{3}, \frac{1}{3}]$. The most striking feature is that the GMR residual and the smallest Rayleigh-Ritz residual slip below the given ε at the same step in the vast majority of cases, particularly for $\varepsilon < 10^{-3}$. In Table 8.1, with $\varepsilon = 10^{-6}$, the step was the same in 18 out of 20 cases. In the other two, the Lanczos algorithm (i.e. the minimal R - R norm) took one more step (17 as against 16).

Some weight is given to the fact that the smallest R - R residual norm is rarely monotone decreasing from one step to another whereas GMR does enjoy this property. However, if the approximate eigenvalue associated with the minimum residual happens to change position in the spectrum from step to step, then this monotonicity of GMR is not associated with the convergence to a specific eigenvalue of the original matrix. No indication is given in the results of how the approximate eigenvalue implicitly chosen by GMR jumps around the spectrum. In practice, the interesting thing to know is *how many* Ritz values have *converged*, and to what accuracy, when the algorithm is terminated. Unfortunately, this information is excluded from the GMR viewpoint and is not reported.

The next results, Examples 8.1 and 8.2 in [Ku, 1986], exhibit the dramatic *failure* of the Lanczos algorithm. On a tridiagonal matrix of order 201 and norm near 1 the minimal R - R residual remained at its initial value 0.035 for all steps except the last (at which it must be 0). In contrast, the GMR residual declined slowly from the initial 0.035 to 0.0039 at step 200. If $\varepsilon = 0.034$, then GMR takes 2 steps while Lanczos takes 201! However, with $\varepsilon \leq 10^{-3}$, both algorithms need 201 steps. We repeat, once again, that GMR will not know which eigenvalue it has approximated.

Unfortunately, no attempt is made to put this example in context. It illustrates the phenomenon explored in some detail in [Sc, 1979], namely that for every symmetric matrix with distinct eigenvalues there is a set (with positive Lebesgue measure on the sphere) of starting vectors such that no Rayleigh-Ritz approximation is any good until the last step. We must repeat that the Lanczos algorithm is not obliged to use a poor initial vector. We ran *our* Lanczos code on this matrix. Our code starts with a normalized version of Ar , where A is the given matrix (or operator) and r 's elements are chosen at random from a uniform random distribution. The reason for starting with Ar is compelling

when an operator A has unwanted infinite eigenvalues. The results are given in the Table 1.

TABLE 1. Convergence of Ritz pairs on T_{201} .

STEP	ε	Number of good Ritz values
20	10^{-2}	1
30	10^{-3}	1
40	10^{-3}	4
50	10^{-3}	5
60	$10^{-4}, 10^{-6}$	7, 4
70	$10^{-5}, 10^{-7}$	10, 5
80	$10^{-5}, 10^{-7}$	14, 10
90	$10^{-5}, 10^{-7}$	18, 14
100	$10^{-5}, 10^{-7}$	23, 18
110	$10^{-5}, 10^{-7}$	28, 24
120	$10^{-5}, 10^{-7}$	35, 30
130	$10^{-5}, 10^{-7}$	44, 37
140	$10^{-5}, 10^{-7}$	52, 44
150	$10^{-5}, 10^{-7}$	60, 54
160	$10^{-5}, 10^{-7}$	70, 61
170	$10^{-5}, 10^{-7}$	83, 75
180	$10^{-5}, 10^{-7}$	99, 89
190	$10^{-5}, 10^{-7}$	118, 109

The accepted eigenvalues (104 of them at step 190) agreed with those computed by EISPACK to all of the fifteen decimals printed out. The efficiency is not at all bad considering that this is a difficult eigenvalue distribution for Krylov space methods.

Example 8.2, a tridiagonal of order 501 with null diagonal and monotonely increasing off diagonal elements, caused the minimal R - R residual norm to *increase* from 0.001 initially in 0.011 at steps 499 and 500. In contrast GMR residual norms declined to 0.00036 at steps 499 and 500. Thus, with $\varepsilon = .00099$, GMR terminates at step 2 whereas Lanczos terminates at step 501! However, with $\varepsilon \leq 10^{-4}$, both take 501 steps.

As with Example 8.1 e_1 is a bad starting vector yielding a poor Krylov subspace. We ran our Lanczos program and found the results given in Table 2.

TABLE 2. Convergence of Ritz pairs on T_{501} .

STEP	ε	Number of good Ritz values
20	10^{-2}	1
30	10^{-3}	1
40	10^{-3}	2
50	10^{-4}	3
60	$10^{-4}, 10^{-6}$	5, 2
70	$10^{-5}, 10^{-7}$	7, 3
80	$10^{-5}, 10^{-7}$	9, 6
90	$10^{-5}, 10^{-7}$	13, 10
100	$10^{-5}, 10^{-7}$	16, 13
110	$10^{-5}, 10^{-7}$	20, 15
120	$10^{-5}, 10^{-7}$	23, 20
130	$10^{-5}, 10^{-7}$	29, 23
140	$10^{-5}, 10^{-7}$	34, 29
150	$10^{-5}, 10^{-7}$	39, 34

We quote the final paragraph of the article.

From all the tests we have performed we conclude that the GMR algorithm is essentially superior to the Lanczos Algorithm on matrices with constant or increasing codiagonal elements. For random matrices or matrices with decreasing codiagonal elements, both algorithms produce nearly the same residuals.

The revealing word here is “codiagonal.” The author has worked exclusively with tridiagonal matrices and has forgotten that the goal of the Lanczos recurrence is to produce a tridiagonal matrix! Given such a matrix one has *no need of either Lanczos or GMR*. As our results indicate, a random starting vector permits the Lanczos algorithm to perform satisfactorily even on such craftily-designed matrices. The quotation reveals just how far a mathematical theory can stray from relevance.

5.6. Summary. Here is an attempt to formulate the numerical analyst’s version of Complexity Theory for Krylov subspaces and eigenvalues.

For each symmetric $n \times n$ matrix there are initial vectors that yield an eigenvalue in one step, and initial vectors that yield an eigenvalue only at the n th step. The nontrivial result contained

in the Kaniel-Paige-Saad error bounds (see [Pa, 1980, Chap. 12]) is that with most starting vectors the extreme eigenvalues can be found in a modest number of steps that depends on the distribution of the spectrum and is nearly independent of n .

We summarize our criticism of [Ku, 1986] but wish to note that the paper is essentially the author's Ph.D. dissertation, and it would not be charitable to hold him responsible for the vagaries to which IBCT is subject.

Section 1 exposes a serious flaw in the model, namely the goal.

Section 2 exposes a subtle way in which features of a method are pushed into the problem statement; the starting vector.

Section 3 shows how standard terms can be redefined; the Lanczos algorithm.

Section 4 contains some clever and interesting results on the approximating power of Krylov subspaces.

Section 5 shows how very misleading numerical results can be in the absence of proper interpretation.

ACKNOWLEDGMENT

The author has been assisted by many friends in composing this essay but would like to single out Stan Eisenstat, G. W. Stewart, and W. Kahan for their helpful and detailed comments.

REFERENCES

- [Ba, 1987] I. Babuska, *Information-based numerical practice*, J. Complexity 3 (1987), 331–346.
- [Bl, Shu, & Sm, 1989] L. Blum, M. Shub, and S. Smale, *On a theory of computation and complexity over the real numbers; NP completeness, recursive functions, and Turing machines*, Bull. Amer. Math. Soc. (N.S.) 21 (1989), 1–46.
- [Bo & Mu, 1975] A. Borodin and I. Munro, *Computational complexity of algebraic and numeric problems*, Amer. Elsevier, New York, 1975.
- [Cu & Wi, 1985] J. K. Cullum and R. A. Willoughby, *Lanczos algorithms for large symmetric eigenvalue computations, Vol. I: Theory*, Progr. Comput. Sci., Birkhauser, Basel, 1985.
- [Da, 1967] J. W. Daniel, *The conjugate gradient method for linear and nonlinear operator equations*, SIAM J. Numer. Anal. 4 (1967), 10–26.
- [Go & Va, 1984] G. H. Golub and C. V. Van Loan, *Advanced matrix computations*, Johns Hopkins Univ. Press, Maryland, 1984.
- [Je, 1977] A. Jennings, *Matrix computation for engineers and scientists*, John Wiley & Sons, Chichester, 1977.
- [Ka, 1966] S. Kaniel, *Estimates for some computational techniques in linear algebra*, Math. Comp. 20 (1966), 369–378.
- [Ku, 1986] J. Kuczynski, *A generalized minimal residual algorithm for finding an eigenpair of a large matrix*, J. Complexity 2 (1986), 131–162.
- [Ku, 1985] ———, *Implementation of the GMR algorithm for large symmetric eigenproblems*, Tech. Report, Comp. Sci. Dept., Columbia University, New York, 1985.
- [Ku & Wo, 1992] J. Kuczynski and H. Woźniakowski, *Average case complexity of the symmetric Lanczos algorithm*, SIAM J. Matrix Anal. Appl., 1992 (to appear).
- [La, 1952] C. Lanczos, *Solution of systems of linear equations by minimized iterations*, J. Res. Nat. Bur. Standards 49 (1952), 33–51.
- [Me, 1963] G. Meinardus, *Über eine Verallgemeinerung einer Ungleichung von L. V. Kantorowitsch*, Numer. Math. 5 (1963), 14–23.

- [O'L, Ste & Va, 1979] D. O'Leary, G. W. Stewart and J. S. Vandergraft, *Estimating the largest eigenvalue of a positive definite matrix*, Math. Comp. **33** (1979), 1289–1292.
- [Pa, 1980] B. N. Parlett, *The symmetric eigenvalue problem*, Prentice-Hall, New Jersey, 1980.
- [Pa, 1982] —, *Two monitoring schemes for the Lanczos algorithm*, Computing Methods in Applied Sciences and Engineering (V. R. Glowinski and J. L. Lions, eds.), North-Holland Press, 1982.
- [Pa, 1984] —, *The software scene in the extraction of eigenvalues from sparse matrices*, SIAM J. Sci. Statist. Comput. **5** (1984), 590–603.
- [Pa & No, 1985] B. N. Parlett and B. Nour-Omid, *The use of refined error bounds when updating eigenvalues of tridiagonals*, Linear Algebra Appl. **68** (1985), 179–219.
- [Pa, Si & Str, 1982] B. N. Parlett, H. A. Simon, and L. M. Stringer, *On estimating the largest eigenvalue with the Lanczos algorithm*, Math. Comp. **38** (1982), 153–165.
- [Pac, 1986] E. Packel, Review of *A general theory of optimal algorithms*, by J. F. Traub and H. Woźniakowski (Academic Press, New York, 1980), SIAM Rev. **28** (1986), 435–437.
- [Pac & Wo, 1987] E. Packel and H. Woźniakowski, *Recent developments on information-based complexity*, Bull. Amer. Math. Soc. (N.S.) **17** (1987), 9–26.
- [Ra, 1972] M. O. Rabin, *Solving linear equations by means of scalar products*, Complexity of Computer Computations (R. E. Miller and J. W. Thatcher, eds.), Plenum Press, New York, 1972, pp. 11–20.
- [Sa, 1980] Y. Saad, *On the rates of convergence of the Lanczos and the block-Lanczos methods*, SIAM J. Numer. Anal. **17** (1980), 687–706.
- [Sc, 1979] D. S. Scott, *How to make the Lanczos algorithm converge slowly*, Math. Comp. **33** (1979), 239–247.
- [Sch & Str, 1971] A. Schoenhage and V. Strassen, *Schnelle Multiplikation Grosser Zahlen*, Computing **7** (1971), 281–292.
- [Sh, 1977] I. Shavitt, *The method of configuration interaction*, Modern Theoretical Chemistry, vol. 3 (H. P. Schaefer, ed.), Plenum Press, 1977.
- [Shu, 1987] M. Shub, Review of *Information, uncertainty, complexity* by J. F. Traub, H. Woźniakowski et al. (Addison-Wesley, Reading, MA 1983), SIAM Rev. **29** (1987), 495–496.
- [Sti, 1958] E. Stiefel, *Kernel polynomials in linear algebra and their numerical applications*, NBS Appl. Math. **43** (1958), 1–22.
- [Str, 1969] V. Strassen, *Gaussian elimination is not optimal*, Numer. Math. **13** (1969), 354–356.
- [Tr & Wo, 1980] J. F. Traub and H. Woźniakowski, *A general theory of optimal algorithms*, Academic Press, New York, 1980.
- [Tr & Wo, 1984] —, *On the optimal solution of large linear systems*, J. Assoc. Comput. Mach., **31** (1984), 545–559.
- [Tr & Wo, 1988] —, *Information-based complexity*, Academic Press, New York, 1988.
- [Tr, Wo & Wa, 1983] J. F. Traub, H. Woźniakowski, and G. Wasilkowski, *Information, uncertainty, complexity*, Addison-Wesley, Reading, 1983.
- [Wi, 1970] S. Winograd, *On the number of multiplication necessary to compute certain functions*, Comm. Pure. Appl. Math. **23** (1970), 165–179.
- [Wi, 1980] —, *Arithmetic complexity of computations*, CBMS-NSF Regional Conf. Ser. In Appl. Math., vol. 33, SIAM, Philadelphia, 1980
- [Wo, 1985] H. Woźniakowski, *A survey of information-based complexity*, J. Complexity **1** (1985), 11–44.