A PROBLEM-CENTRED APPROACH TO CANONICAL MATRIX FORMS

J. Sylvestre
Department of Science
University of Alberta, Augustana Campus
Camrose, Alberta T4V 2R3, CANADA
jeremy.sylvestre@ualberta.ca

A PROBLEM-CENTRED APPROACH TO CANONICAL MATRIX FORMS

Abstract: This article outlines a problem-centred approach to the topic of canonical matrix forms in a second linear algebra course. In this approach, abstract theory, including such topics as eigenvalues, generalized eigenspaces, invariant subspaces, independent subspaces, nilpotency, and cyclic spaces, is developed in response to the patterns discovered in studying similarity classes of square matrices, rather than as a disconnected prerequisite to this topic. Furthermore, the subtopics involved afford an opportunity to highlight many common mathematical problem-solving techniques and philosophies.

Keywords: problem-centred learning, linear algebra, Jordan canonical form

1 INTRODUCTION

In [1], the authors advocate a problem-centred approach to the teaching of undergraduate linear algebra. In such an approach, abstract theory is developed to use as a tool in solving specific problems, which may be of an applied or purely mathematical nature. The purpose of this approach is to give students an immediate *raison d'être* for the theory, as opposed to the "much-delayed gratification" ([1]) of the traditional *theory-then-applications* order of topics, and to allow the students to see theory as arising naturally out of the study of problems. This method could be viewed as a form of *inquiry-based learning*, and the outline of the study of canonical matrix forms given in this article could be adapted for use in such a learning environment in a straightforward manner.

The relation of similarity is an equivalence relation on the set of $n \times n$ matrices, and the members of a given equivalence class share all of the most important properties of a matrix: rank, nullity, invertibility, characteristic polynomial, eigenvalues, dimensions of eigenspaces, determinant, and trace. Most of elementary linear algebra reduces to the study of homogeneous systems of linear equations, and two similar matrices represent the same system, up to a linear change of variables. The central problem in the topic of canonical matrix forms could be phrased as: what is the "simplest" representative member of any given similarity class? Attempting to answer this question, and discovering that the answer can take on many forms, provides an enriching theme of inquiry for one unit of a second linear algebra course. In particular, a step-by-step investigation of the topic, through examples and special cases, allows an instructor to model for students many common mathematical problem-solving and theory-building techniques, such as

- considering special cases as prelude to more complicated situations,
- generalizing concepts,
- applying and adapting old methods to new situations,
- breaking objects into smaller constituent parts, and analyzing each piece individually,
- special analysis of "cusp" cases, on the boundary between one type of behaviour and another,
- using indirect methods of gathering information, and
- collecting results into a general theory encompassing all cases,

all in one self-contained unit. These techniques, and the pedagogical choice to explicitly highlight their use in teaching the topic of canonical matrix forms, could be considered to be in the spirit of the modern problem solving heuristics of Pólya (see [3]).

2 FIRST STEPS

2.1 Diagonalizable Matrices

Philosophy: Start with a simple case.

A good place to start is with a discussion of what "*simplest* member of a similarity class" should even mean. Students will usually offer the zero matrix and the identity matrix as the simplest matrices they know, but each of these is alone in its respective similarity class. The observation can be made that analyzing *too* simple a case is often not valuable, if it provides no clues as to how to proceed in other cases. (One could also pause here to investigate: what other matrices are the sole member of their similarity class?)

Eventually diagonal matrices will come up. Students will

$$D = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & -1 \end{bmatrix} \qquad A = \begin{bmatrix} 7 & -12 & -4 \\ 4 & -9 & -4 \\ -4 & 12 & 7 \end{bmatrix} \qquad P = \begin{bmatrix} -1 & 2 & -1 \\ 0 & 1 & -1 \\ -1 & -1 & 1 \end{bmatrix}$$

Example 1. How is it that $P^{-1}AP = D$?

probably have seen eigenvalues and eigenvectors in their first course in linear algebra, and they might remember that the answer to the question posed in Example 1 is that the columns of P must be eigenvectors of A, but they probably do not remember why. It is easy to convince the students that the requirement $AP = PD$ is much nicer to work with than $P^{-1}AP = D$, as it avoids a messy inverse matrix calculation. From here, we can borrow an idea from the concept of matrix-as-linear-transformation, and make the connection to the vector space $\mathbf{C}^n$ (where $\mathbf{C}$ is the field of complex numbers) — it can only be true that AP and PD are the same matrix if they act the same way on $\mathbf{C}^n$. That is, we must have $AP\mathbf{e}_j = PD\mathbf{e}_j$ for each standard basis vector $\mathbf{e}_j$ of $\mathbf{C}^n$. Now, multiplying a standard basis vector by a matrix picks off the corresponding column of the matrix, and since the columns of D are just scalar multiples of these basis vectors, we immediately get that $A\mathbf{p}_j = \lambda_j \mathbf{p}_j$, where $\mathbf{p}_j$ is the j^{th} column of P, and λ_j is the j^{th} diagonal entry of D. Of course, this says that $\lambda_j, \mathbf{p}_j$ must be an eigenvalue-eigenvector pair of A, and we are led directly to the attendant theory of such objects.

2.2 Similarity Revisited

Philosophy: Find the essence of what made the analysis of the simple case work.

Looking back at the diagonalizable case, it is easy for students to focus on the fact that eigenvalues and eigenvectors made an appearance, and to get lost in details like algebraic and geometric multiplicities of eigenvalues. While these objects and details will certainly be important in further cases, the analysis that *led* to the consideration of eigenvalues and eigenvectors is much more important for the purpose of developing a framework for investigating *all* possible cases.

If A and B are similar matrices, and P is an invertible matrix that realizes the similarity relation between them, then the idea of analyzing the equality $AP\mathbf{e}_j = PB\mathbf{e}_j$ (similar to the analysis of the diagonalizable case), leads to

$$\begin{aligned} A\mathbf{p}_j &= P\mathbf{b}_j \\ &= P(b_{1j}\mathbf{e}_1 + b_{2j}\mathbf{e}_2 + \cdots + b_{nj}\mathbf{e}_n) \\ &= b_{1j}\mathbf{p}_1 + b_{2j}\mathbf{p}_2 + \cdots + b_{nj}\mathbf{p}_n. \end{aligned}$$

This observation provides the guiding principle by which to analyze similarity in general: $P^{-1}AP = B$ *is true if and only if the columns of* B *encode the action of* A *on the basis of* $\mathbf{C}^n$ *formed by the columns of* P.

2.3 Block Diagonal Form

Philosophy: Generalize.

If a matrix is not diagonalizable, when is it similar to something vaguely diagonal in form? Block diagonal form is a simple example of taking old ideas in a new direction to produce a generalization.

$$B = \begin{bmatrix} 3 & 1 & & \\ 1 & 1 & & \\ & & 2 & 1 \\ & & 1 & 2 \end{bmatrix} \qquad A = \begin{bmatrix} 6 & 4 & -1 & -3 \\ 7 & 10 & -2 & -6 \\ 0 & 6 & 0 & -3 \\ 14 & 12 & -3 & -8 \end{bmatrix} \qquad P = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 1 \\ 0 & 0 & 1 & -2 \\ 2 & 4 & 2 & 3 \end{bmatrix}$$

Example 2. How is it that $P^{-1}AP = B$?

Consider the matrices in Example 2. Using the fact that the form matrix B encodes the action of A on the columns of P, we see that A sends each of $\mathbf{p}_1$ and $\mathbf{p}_2$ to a linear combination of the two, and similarly for $\mathbf{p}_3$ and $\mathbf{p}_4$. This example leads to the theory of *invariant subspaces* and of *collections of independent subspaces.* (I prefer to avoid the concept of direct sum at this level as unnecessarily complicated. The equivalent concept of independent subspaces is more natural as a generalization of independent vectors.)

Of course, for a given matrix, it is not so easy to come up with a collection of independent, invariant subspaces. But geometric examples can be used here, such as rotations in $\mathbf{R}^3$. And, of course, the prototypical example of the collection of eigenspaces of a diagonalizable matrix should be discussed as a prelude to a future topic.

3 SQUEEZING MORE OUT OF EIGENVALUES

3.1 Scalar-Triangular Form

Philosophy: Try an incrementally more difficult case.

What can we do with a matrix that is not diagonalizable? The students will likely identify triangular matrices as the appropriate next step after diagonal matrices. Here it is useful to use the special case of an upper triangular matrix with a single repeated eigenvalue (I call this a *scalar-triangular matrix*) as an incremental step.

$$T = \begin{bmatrix} 3 & 1 & 2 \\ 0 & 3 & -1 \\ 0 & 0 & 3 \end{bmatrix} \qquad A = \begin{bmatrix} -2 & -14 & 5 \\ 1 & 6 & -1 \\ -2 & -5 & 5 \end{bmatrix} \qquad P = \begin{bmatrix} 1 & -2 & 1 \\ 0 & 1 & -1 \\ 1 & 1 & -1 \end{bmatrix}$$

Example 3. How is it that $P^{-1}AP = T$?

Consider the matrices in Example 3. Once again, using the point of view that the form matrix T encodes the action of A on the columns of P, we immediately see that the first column of P must be an eigenvector of A. But the second column is a puzzle — it is *almost* an eigenvector, except for a pesky extra multiple of $\mathbf{p}_1$. With a little algebraic manipulation, we find that the vector $\mathbf{p}_2$ does not solve the homogeneous system $(\lambda I - A)\mathbf{x} = \mathbf{0}$ (as it should if it were an eigenvector), but instead solves the nonhomogeneous system $(\lambda I - A)\mathbf{x} = -t_{12}\mathbf{p}_1$. If possible, it would be preferable to replace this relationship between $\mathbf{p}_1$ and $\mathbf{p}_2$ with a homogeneous condition involving just $\mathbf{p}_2$. Toward that end, we can use the already obtained knowledge that $\mathbf{p}_1$ is an eigenvector to turn $(\lambda I - A)\mathbf{p}_2 = -t_{12}\mathbf{p}_1$ into $(\lambda I - A)^2\mathbf{p}_2 = \mathbf{0}$. From here, the pattern becomes evident: we need $(\lambda I - A)^j\mathbf{p}_j = \mathbf{0}$ to be true for each column index j, and so are led naturally to the theory of *generalized eigenvectors.*

3.2 Triangular-Block Form

Philosophy: Boldly forge ahead using whatever tools worked before.

If A has more than one eigenvalue, can it still be similar to a triangular matrix? Generalized eigenvectors had something to do with the scalar-triangular case, so it stands to reason that they could play a role in this more general setting. Working with some examples, we quickly find that generalized eigenspaces of a matrix A are invariant under multiplication by A, and the collection of generalized eigenspaces of A form an independent collection. So the theory attached to block diagonal form can be applied to put any matrix into a block form where each block is in scalar-triangular form.

4 ATTACKING THE BLOCKS

4.1 Nilpotent Matrices

Philosophy: Break the problem apart.

Can each of the blocks of a matrix in triangular-block form be "simplified" any further? It is not immediately obvious how to choose bases for each generalized eigenspace to do this. The scalar diagonal of each block is as simple as possible already, so it is natural to focus on the messy upper triangular part. Breaking each block into the sum $\lambda I + N$, and recognizing the simple algebraic identity $P^{-1}(\lambda I + N)P = \lambda I + P^{-1}NP$, leads naturally to the consideration of the special case of nilpotent matrices. And just as the study of matrix forms is an attempt to understand the partition of the set of all $n \times n$ matrices into similarity classes, nilpotent matrices are one cell in a partition of the set of singular $n \times n$ matrices with respect to the multiplicity of the eigenvalue 0. The set of nilpotent matrices can be further subdivided with respect to the degree of their nilpotency — that is, the size of the exponent in the first power of the matrix to become zero. (See Figure 1.)

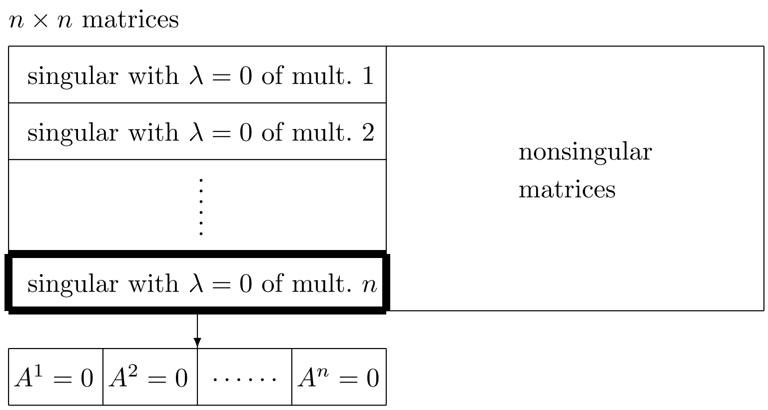

Figure 1. Cells of singular matrices.

4.2 Elementary Nilpotent Form

Philosophy: Analyze the "cusp" case.

Of the cells of nilpotent $n \times n$ matrices in the partition illustrated in Figure 1, the cell containing those matrices A that satisfy $A^{n-1} \neq 0$ is the furthest from the cell containing just the zero matrix, and could be considered to be on the "cusp" between the larger cells of nilpotent and nonnilpotent singular matrices. For this reason, the matrices in this cell are a good initial special case to consider. It turns out that such matrices are *all* similar to a particularly simple triangular form (and, therefore, to each other). The prototypical example of a nilpotent matrix in this class is one that is zero in every entry, except for an unbroken line of ones along the first subdiagonal (I call this *elementary nilpotent form*). Here, we switch to a lower triangular form so that the theory will work out a little more cleanly.

$$N = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \qquad A = \begin{bmatrix} 1 & 2 & 0 & -1 \\ 3 & 7 & -1 & -4 \\ -1 & 9 & -2 & -4 \\ 6 & 10 & -1 & -6 \end{bmatrix} \qquad P = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 1 \\ 0 & 0 & 1 & -2 \\ 2 & 4 & 2 & 3 \end{bmatrix}$$

Example 4. How is it that $P^{-1}AP = N$?

In Example 4, we again consider the action of A on the columns of P, and we see that the columns of P form an A-cyclic basis for $\mathbf{C}^n$, with $\mathbf{p}_j = A^j\mathbf{p}_1$ for $2 \leq j \leq 4$. In general, if A is an $n \times n$ nilpotent matrix with $A^{n-1} \neq 0$, we can always take P to be the matrix with columns

$\mathbf{p}_j = A^{j-1}\mathbf{e}_k$, where k is any fixed index such that the k^{th} column of A^{n-1} is nonzero, and $\mathbf{e}_k$ is the corresponding standard basis vector. In this way, we are led to the theory of *cyclic subspaces*.

4.3 Triangular-Block Form For Nilpotent Matrices

Philosophy: Gather information indirectly.

From here, depending on the level of the course and the time available, some "hand waving" may be called for. A full journey to the Cyclic Decomposition Theorem is not really necessary, and, based on the results so far, students should have little trouble accepting that a nilpotent matrix that becomes zero *before* the n^{th} power is similar to one made up of smaller blocks in elementary nilpotent form. Rather than focus on algorithm development, I prefer here to set students to some detective tasks, to indirectly determine the form of a nilpotent matrix, without explicitly computing P.

Using the knowledge that nilpotent A is similar to a block diagonal matrix, with each block in elementary nilpotent form, how can we determine the number and sizes of the blocks? This information completely determines the form matrix, and, by similarity, the ranks of the powers of A are enough to work out the form. Each block in the form has rank one less than full, so the number of blocks is equal to $n - \text{rank}(A) = \text{nullity}(A)$. If k is the first positive exponent such that $A^k = 0$, then the largest block is $k \times k$, and there are $\text{rank}(A^{k-1})$ such blocks. From here, the term-to-term jumps in the sequence $\text{rank}(A^{k-1}), \text{rank}(A^{k-2}), \ldots, \text{rank}(A)$ tell the number and sizes of the smaller blocks.

5 PUTTING IT ALL TOGETHER

5.1 Jordan Canonical Form

Philosophy: Develop a general theory encompassing all cases.

Finally, the results of the analysis of the nilpotent case can be applied back to the triangular-block form case, to arrive at the Jordan canonical form. Again, if procedural proficiency is not a main objective of the course, it may be more appropriate to tackle this form from the same philosophy as the triangular-block form for nilpotent matrices: the number and sizes of elementary Jordan blocks corresponding to a given eigenvalue λ_j can be deduced from $\text{nullity}(A_j)$ and the sequence $\text{rank}(A_j^{m_j}), \text{rank}(A_j^{m_j-1}), \ldots, \text{rank}(A_j)$, where $A_j = \lambda_j I - A$, and m_j is the algebraic multiplicity of λ_j.

6 ALGORITHM DEVELOPMENT TO REINFORCE THEORY

In [2], it is suggested that "learning may be improved by helping students *construct* knowledge in their own minds in a context that is designed to aid, or even stimulate, that construction," and that one way to structure such a learning experience is "to have the student program mathematical constructions in a computer language designed so that the act of programming parallels the construction of the underlying mathematical processes."

Rather than hand the students ready-made algorithms for computing, for a given matrix A, an invertible matrix P so that $P^{-1}AP$ is in a desired form, the instructor can ask the students to develop the algorithms on their own, forcing them to think more deeply about how the theory leads to a solution of the problem than they would in just running through the (quite tedious) calculations in exercises. The algorithms could take the form of a sequence of commands to be entered into a computer algebra system such as Maple or Maxima, or could simply be in a pseudo-code language of the instructor's or students' own devising.

```
BEGINFUNCTION (INPUT matrix A)
    COMPUTE eigenvalues of A
    IF A has more than one distinct eigenvalue THEN
        RETURN ERROR
    ENDIF
    SET e := eigenvalue of A
    SET I := identity matrix of same dimension as A
    COMPUTE basis B for nullspace of (e*I − A)
    SET k := 2
    WHILE B contains fewer vectors than dimension of A DO
        COMPUTE basis B2 for nullspace of (e*I − A)^k
        WHILE B contains fewer vectors than B2 DO
            APPEND to B a vector in B2 that is not in span of B
        DONE
        INCREMENT k
    DONE
    RETURN matrix P obtained by concatenating column vectors in B
ENDFUNCTION
```

Listing 1. Pseudo-code for an algorithm to compute from input matrix A an output matrix P so that $P^{-1}AP$ is in scalar-triangular form.

In Listing 1, a sample pseudo-code algorithm for putting a matrix in scalar-triangular form is given. As can be seen in this sample listing, developing such algorithms also forces students to revisit fundamental topics in elementary linear algebra, such as basis and span, and how to extend a basis for a subspace to one for a larger space. These abstract concepts, with which students often struggle, become more fixed or concrete in the mind of the student when they consider them as related to objects (matrices and vectors) which are supposed by the student to exist, if only in the memory of a (possibly fictitious) computer, but are still abstract in that they are not specific matrices and vectors (see [2, p. 235]).

7 CONCLUSION

In this article, we described how the topic of canonical matrix forms can be used as the context and motivation for the development of several important general topics in elementary linear algebra. By situating these general topics in a problem-solving context, rather than presenting them as disconnected prerequisites, it is hoped to provide a more engaging learning experience for the student, as well as an opportunity for the student to reflect on various approaches to problem solving.

REFERENCES

[1] Alm, J. F. and D. A. Andrews 2012. Solutions in search of problems. *Notices of the American Mathematical Society.* 59(7): 963–964.

[2] Dubinsky, E. and D. Tall 1991. Advanced mathematical thinking and the computer. In Tall, D. (Ed.), *Advanced Mathematical Thinking* (pp. 231–243). Dordrecht; Boston : Kluwer Academic Publishers.

[3] Pólya, G. 1945. *How to Solve It.* Princeton, New Jersey: Princeton University Press.

BIOGRAPHICAL SKETCH

J. Sylvestre received his Ph.D. from the University of Toronto in 2008. Since then, he has enjoyed developing his teaching as an Assistant Professor of mathematics at Augustana Campus, a small, liberal arts, undergraduate focused teaching campus of the University of Alberta.