The Abdus Salam International Centre for Theoretical Physics

United Nations Educational, Scientific and Cultural Organization

IAEA International Atomic Energy Agency

2044-6

Summer School and Advanced Workshop on Trends and Developments in Linear Algebra

22 June - 10 July, 2009

Geometrical techniques in linear algebra-lecture notes

Peter Semrl

Faculty of Mathematics and Physics
University of Ljubljana, Ljubljana, Slovenia

peter.semrl@fmf.uni-lj.si

Geometrical techniques in linear algebra-lecture notes, ICTP, Trieste 2009

Peter Šemrl
Faculty of Mathematics and Physics
University of Ljubljana
Jadranska 19
SI-1000 Ljubljana
Slovenia
peter.semrl@fmf.uni-lj.si

June 21, 2009

1 Introduction

Linear algebra is a powerful tool in almost all branches of mathematics. This is a well-known fact. There are a lot of examples where we have the opposite situation. We start with a problem in linear algebra that is easy to understand. However, to solve it we need a lot of deep mathematical ideas coming from analysis, algebra, geometry, topology, etc. Two spectacular recent cases are Boyle-Handelman partial solution of the nonnegative inverse eigenvalue problem (M. Boyle and D. Handelman, The spectra of nonnegative matrices via symbolic dynamics, Ann. Math. 133 (1991), 249-316) and the complete solution of Alfred Horn's conjecture on eigenvalues of sums of Hermitian matrices obtained by Klyachko, and Knutson and Tao (see the paper R. Bhatia, Linear Algebra to quantum cohomology: the story of Alfred Horn's inequalities, Amer. Math. Monthly 108 (2001), 289-318).

Here we are going to deal with less spectacular results. Working as a linear algebraists I have realized that geometrical techniques can be quite useful when solving certain linear algebra problems. My plan is to present three exmples illustrating this fact. In these lecture notes the emphasis will be on main ideas. So, no detailed proofs will be given. However, we will give references where an interested reader can find complete proofs of the statements presented in these notes.

In the first section we will give an example showing that affine and projective geometry arise naturally in some linear algebra problmes. We will continue with differential geometry and the last section will be devoted to an application of algebraic geometry.

Each section will be organized in the following way. We will start with some basic definitions and facts from geometry. All the definitions and facts will be at the elementary level. Then we will present a linear algebra problem. And finally, we will try to explain how geometrical methods can be applied to solve the problem.

2 Affine and projective geometry

2.1 Basic concepts

Let V be a vector space over a field $\mathbb{F}$. A map $\phi : V \to V$ is called a semilinear map if there exists an automorphism τ of the field $\mathbb{F}$ such that

$$\phi(x+y) = \phi(x) + \phi(y), \quad x, y \in V,$$

and

$$\phi(\lambda x) = \tau(\lambda)\phi(x), \quad \lambda \in \mathbb{F}, \ x \in V.$$

It is well-known that the identity is the only automorphism of the real field. Hence, semilinear maps on real vector spaces are linear maps. Examples of semilinear maps on complex vector spaces V include linear maps and conjugate-linear maps, that is, maps $\phi : V \to V$ satisfying

$$\phi(x+y) = \phi(x) + \phi(y), \quad x, y \in V,$$

and

$$\phi(\lambda x) = \overline{\lambda}\phi(x), \quad \lambda \in \mathbb{C}, \ x \in V.$$

Besides the identity and the complex conjugation there are other automorphisms of the complex field. Consequently, there are semilinear maps on complex space V that are neither linear, nor conjugate-linear.

Let again V be a vector space over an arbitrary field $\mathbb{F}$ and $x, y \in V$ arbitrary vectors with $y \neq 0$. The set

$$\{x + \lambda y \, : \, \lambda \in \mathbb{F}\}$$

is called a line. Suppose that $\phi : V \to V$ is a bijective semilinear map. Then $\phi(L)$ is a line for every line $L \subset V$.

Theorem 1 *(Fundamental theorem of the affine geometry) Let $\mathbb{F}$ be a field, $\mathbb{F} \neq \mathbb{F}_2$, and V a vector space over $\mathbb{F}$, $\dim V \geq 2$. Let $\psi : V \to V$ be a bijective map with the property that $\psi(L)$ is a line for every line $L \subset V$. Then there exist a bijective semilinear map $\phi : V \to V$ and a vector $x_0 \in V$ such that*

$$\psi(x) = \phi(x) + x_0, \quad x \in V.$$

The proof of this statement can be found in the book: Z.-X. Wan, Geometry of matrices, World Scientific, 1996. This book contains proofs of most of the theorems that will be mentioned in this section.

For a non-zero vector $x \in V$ we denote by $[x]$ the one-dimensional subspace of V spanned by x. The set of all one-dimensional subspaces of V is called the projective space over V denoted by $\mathbb{P}V$,

$$\mathbb{P}V = \{[x] \, : \, x \in V \setminus \{0\}\}.$$

Theorem 2 *(Fundamental theorem of projective geometry) Let $\mathbb{F}$ be any field and V a vector space over $\mathbb{F}$, $\dim V \geq 3$. Let $\psi : \mathbb{P}V \to \mathbb{P}V$ be a bijective map such that for every $x, y, z \in V \setminus \{0\}$ we have*

$$[x] \subset [y] + [z] \iff \psi([x]) \subset \psi([y]) + \psi([z]).$$

Then there exists a bijective semilinear map $\phi : V \to V$ such that

$$\psi([x]) = [\phi(x)], \quad x \in V \setminus \{0\}.$$

2.2 Fundamental theorem of the geometry of matrices

We denote by $M_{m\times n}$ the algebra of all $m \times n$ matrices over a field $\mathbb{F}$. In the case when $m = n$ we simply write $M_n = M_{n\times n}$. For $A, B \in M_{m\times n}$ we define the arithmetic distance $d(A, B)$ between A and B to be

$$d(A, B) = \text{rank}\,(B - A).$$

The rank is subadditive (that is, $\text{rank}\,(A + B) \leq \text{rank}\, A + \text{rank}\, B$), and it follows easily that $(M_{m\times n}, d)$ is a metric space. Then one can ask what are all isometries of $M_{m\times n}$ with respect to the distance d. We can ask an even more general question, that is, we would like to know the general form of bijective distance one preserving maps on $M_{m\times n}$. We say that $A, B \in M_{m\times n}$ are adjacent if $d(A, B) = 1$. In order to formulate the fundamental theorem of the geometry of matrices we need one more notation. Let $A = [a_{ij}] \in M_{m\times n}$ be any matrix and τ an automorphism of the field $\mathbb{F}$. Then A_τ denotes the matrix obtained from A by applying the automorphism τ entrywise, $[a_{ij}]_\tau = [\tau(a_{ij})]$.

Theorem 3 *Let $\mathbb{F}$ be any field, m, n positive integers ≥ 2, and $\phi : M_{m\times n} \to M_{m\times n}$ a bijective map. Assume that for every pair $A, B \in M_{m\times n}$ we have*

$$A \text{ and } B \text{ are adjacent} \iff \phi(A) \text{ and } \phi(B) \text{ are adjacent}.$$

Then there exist invertible matrices $P \in M_m$ and $Q \in M_n$, an arbitrary matrix $T \in M_{m\times n}$, and an automorphism $\tau : \mathbb{F} \to \mathbb{F}$ such that

$$\phi(A) = PA_\tau Q + T, \quad A \in M_{m\times n}.$$

In the case when $m = n$, we have an additional possibility that

$$\phi(A) = PA_\tau^T Q + T, \quad A \in M_n.$$

In fact, an analogous statement holds true also for bijective adjacency preserving maps on matrices over an arbitrary division ring. If $\phi : M_{m\times n} \to M_{m\times n}$ is an adjacency preserving map, then the map $\psi : M_{m\times n} \to M_{m\times n}$ defined by $\psi(A) = \phi(A) - \phi(0)$ preserves adjacency as well. So, when studying adjacency preserving maps ϕ there is no loss of generality in assuming that $\phi(0) = 0$ (then, of course, we have $T = 0$ in the conclusion of our theorem). From now on, we will always assume that $\phi(0) = 0$ when speaking of an adjacency preserving map ϕ. It is a remarkable fact that after a harmless normalization $\phi(0) = 0$ the semilinear character of ϕ is not an assumption but a conclusion.

This beautiful result due to Hua has many applications. Let us mention here two of them. Let $G(n, k)$ be the Grassmann space consisting of all k-dimensional subspaces of an n-dimensional vector space V. We recall that two elements of $G(n, k)$ are adjacent if their intersection has dimension $k - 1$. In 1949 Chow (W.-L. Chow, On the geometry of algebraic homogeneous spaces, Ann. Math. 50 (1949), 32-67) determined all bijections $\phi : G(n, k) \to G(n, k)$ that preserve adjacency in both directions. Except in the trivial cases $k = 1$ and $k = n - 1$ every such map is induced by a bijective semilinear transformation $\psi : V \to V$. In the case when $n = 2k$ we have the additional possibility that such a map is induced by a semilinear transformation from V onto its dual V^*.

Let m, n be positive integers, $m, n \geq 2$. Then to each point in $G(m+n, m)$, that is, to each m-dimensional subspace U of $\mathbb{F}^{m+n}$, we can associate an $m \times (m + n)$ matrix whose rows are coordinates of the vectors that form a basis of U. Each $m \times (m + n)$ matrix will be written in the block form $[X\ Y]$. Here, X is an $m \times n$ matrix and Y is an $m \times m$ matrix. It is an elementary exercise in linear algebra to prove that two matrices $[X\ Y]$ and $[X'\ Y']$ are associated to the same subspace U (their rows represent two bases of U) if and only if $[X\ Y] = P[X'\ Y']$ for some invertible $m \times m$ matrix P. If this is the case, then either both Y and Y' are invertible, or both Y and Y' are singular. So, to each point in a Grassmann space we have associated a (not uniquely determined) matrix $[X\ Y]$. If Y is singular, then the corresponding point in the Grassmann space is called a point at infinity. Otherwise, it is called a finite point. Observe that a finite point $[X\ Y]$ can be represented also with the matrix $[Y^{-1}X \quad I]$. The matrix $Y^{-1}X$ in such a representation is uniquely determined by the point in the Grassmann space. So, if U and V are two m-dimensional subspaces that are finite points in the Grassmann space, then they can be represented with two uniquely determined $m \times n$ matrices T and S, respectively, and it is easy to see that the subspaces U and V are adjacent if and only if the matrices T and S are adjacent. Using this connection it is possible to deduce Chow's description of bijective maps on the Grassmann space preserving adjacency in both directions from Theorem 3.

Theorem 3 is of fundemantal importance in the theory of linear preservers. A linear map $\phi : M_n \to M_n$ is called a linear preserver if it preserves a certain property or a certain subset or a certain relation. Let us mention here three examples. We say that ϕ preserves rank one if $\phi(A)$ is a matrix of rank one for every $A \in M_n$ of rank one. Further, ϕ preserves invertibility if $\phi(A)$ is invertible whenever $A \in M_n$ is invertible. And finally, ϕ preserves commutativity if $\phi(A)$ and $\phi(B)$ commute whenever A and B commute.We want to describe the general form of such maps. The problem of characterizing linear invertibility preserving maps was formulated by Kaplansky and has been extensively studied not only on matrix algebras, but also on operator algebras and even more general semi-simple Banach algebras. The algebra M_n becomes a Lie algebra if we introduce the Lie product $[A, B] = AB - BA$. Hence, linear maps preserving commutativity are linear maps preserving zero Lie products. Thus, describing the general form of linear commutativity preserving maps we get as a corollary the structural result for Lie homomorphisms. It is therefore not surprising that one can find in the literature a lot of papers on linear preservers of commutativity on various algebras. The most frequently used method in the theory of linear preservers is the reduction of a problem of characterizing certain linear preservers to the problem of characterizing linear maps preserving rank one matrices. Let us illustrate this with an example. Suppose we want to describe the general form of bijective linear maps $\phi : M_n \to M_n$ preserving invertibility in both directions, that is, for every $A \in M_n$ we have

$$A \text{ is invertible} \iff \phi(A) \text{ is invertible.} \tag{1}$$

We assume that the underlying field $\mathbb{F}$ is not of characteristic 2. The first step in the proof is the following characterization of rank one matrices. For a nonzero $A \in M_n$ the following two conditions are equivalent:

- $\operatorname{rank} A = 1$,
- for every invertible matrix $B \in M_n$ the matrix $B + \lambda A$ is invertible for all but at most one scalar λ.

Assume for a moment that we have already proved that these two statements are equivalent. Then it follows from bijectivity of ϕ, (1), and the above characterization that for every $A \in M_n$ we have

$$\operatorname{rank} A = 1 \iff \operatorname{rank} \phi(A) = 1.$$

Thus, two matrices A and B are adjacent if and only if $\operatorname{rank}(A - B) = 1$, which is equivalent to $\operatorname{rank} \phi(A - B) = \operatorname{rank}(\phi(A) - \phi(B)) = 1$. In other words, A and B are adjacent if and only if $\phi(A)$ and $\phi(B)$ are adjacent. So, we can apply Theorem 3 and linearity to conclude that either

$$\phi(A) = PAQ, \quad A \in M_n,$$

or

$$\phi(A) = PA^T Q, \quad A \in M_n$$

for some invertible matrices $P, Q \in M_n$. Most of linear preserver problems can be solved using this pattern: we want to characterize linear maps on M_n having a certain preserving property and then we first show that such maps preserve rank one operators which further yields that they preserve adjacency. And then we are in a position where we can apply Theorem 3.

There is an important observation here. Following the above approach we start with a linear preserver problem and we reduce it to the problem of characterizing linear bijective maps on matrices preserving adjacency. But Theorem 3 characterizes bijective adjacency preserving maps without assuming linearity. This suggests that Theorem 3 can be applied to study not only linear, but also general (non-linear) preservers. And indeed, recently the first results on non-linear preservers have been proved either using Theorem 3, or proof techniques similar to those used in the proof of Theorem 3.

We will conclude this section by proving the equivalence of the above two conditions. Assume first that $\operatorname{rank} A = 1$ and $B \in M_n$ is invertible. Then

$$B + \lambda A = B(I + \lambda R)$$

(here, $R = B^{-1}A$ is a matrix of rank one) is invertible if and only if $I + \lambda R$ is invertible. It's a rather easy exercise to show that for any rank one matrix R the matrix $I + \lambda R$ is either invertible for all scalars λ, or is singular for exactly one scalar λ_0 and invertible for all other scalars λ.

To prove the converse assume that $\operatorname{rank} A \geq 2$. Identifying matrices with linear operators we can find two vectors x_1, x_2 such that $Ax_1 = y_1$ and $Ax_2 = y_2$ are linearly independent. There exists an invertible matrix B such that $Bx_1 = y_1$ and $Bx_2 = -y_2$. It follows that both matrices $B + A$ and $B - A$ are singular.

2.3 Geometric ideas used in the proof of Theorem 3

We can assume with no loss of generality that $\phi(0) = 0$. We identify vectors in $\mathbb{F}^m$ with $m \times 1$ matrices. So, each matrix $A \in M_{m\times n}$ of rank at most one can be written as

$$A = xy^T,$$

where $x \in \mathbb{F}^m$ and $y \in \mathbb{F}^n$ (of course, A is zero if and only if $x = 0$ or $y = 0$, otherwise A is of rank one).

We fix nonzero $x \in \mathbb{F}^m$ and $y \in \mathbb{F}^n$ and denote

$$L_x = \{xu^T \,:\, u \in \mathbb{F}^n\}$$

and

$$R_y = \{wy^T \,:\, w \in \mathbb{F}^m\}.$$

Clearly, $L_x \subset M_{m\times n}$ has the following two properties:

- $0 \in L_x$,
- $A, B \in L_x$ and $A \neq B \Rightarrow A$ and B are adjacent.

The same is true for R_y. Moreover, if $S \subset M_{m\times n}$ is any subset of pairwise adjacent matrices containing the zero matrix, then S is contained in either some L_x or some R_y. As $\phi(0) = 0$ and ϕ preserves adjacency, it maps maximal subsets of pairwise adjacent matrices containing the zero matrix onto the subsets of the same type. Thus, every L_x is mapped onto some L_w, or onto some R_z. It is not too difficult to show that either every L_x is mapped onto L_w for some $w \in \mathbb{F}^m$, or every L_x is mapped onto R_z for some $z \in \mathbb{F}^n$.

Let us consider only the first case. So for each nonzero $x \in \mathbb{F}^m$ there exists $w \in \mathbb{F}^m \setminus \{0\}$ such that

$$\phi(L_x) = L_w. \tag{2}$$

Clearly, $L_{x_1} = L_{x_2}$ if and only if the nonzero vectors x_1 and x_2 are linearly dependent. Thus, applying (2) we see that ϕ induces a map ξ on $\mathbb{PF}^m$ such that

$$\xi([x]) = [w] \iff \phi(L_x) = L_w, \quad x, w \in \mathbb{F}^m \setminus \{0\}.$$

In the next step we show that

$$[x] \subset [y] + [z] \iff \xi([x]) \subset \xi([y]) + \xi([z]), \quad x, y, z \in \mathbb{F}^m \setminus \{0\}.$$

It is enough to prove the implication in one direction only. It is also enough to consider the case when y and z are linearly independent (then, as it is easy to see, $\xi([y]) \neq \xi([z])$). Note that ϕ maps rank one matrices to rank one matrices (these are precisely the matrices that are adjacent to 0) and it maps matrices of rank two into matrices of rank two. Indeed, if A is of rank two, then it is adjacent to some rank one matrix. It follows that $\phi(A)$ is adjacent to some rank one matrix and is therefore of rank at most two. But it cannot be of rank one (because $\phi(A)$ is not adjacent to 0) and it cannot be 0 because of the bijectivity and the fact that $\phi(0) = 0$. It is also rather easy to see that if B is a rank one matrix adjacent to a rank two matrix A, then $\operatorname{Im} B \subset \operatorname{Im} A$ (here we have identified matrices with linear operators).

Now, if $[x] \subset [y] + [z]$, then we can find nonzero vectors u_1, u_2, u_3 and a rank two matrix B such that xu_1^T, yu_2^T, zu_3^T are all adjacent to B. Thus, $\phi(xu_1^T), \phi(yu_2^T), \phi(zu_3^T)$ are all adjacent to $\phi(B)$, and consequently, $\xi([x]) = \operatorname{Im}\phi(xu_1^T) \subset \operatorname{Im}\phi(B)$, $\xi([y]) = \operatorname{Im}\phi(yu_2^T) \subset \operatorname{Im}\phi(B)$, and $\xi([z]) = \operatorname{Im}\phi(zu_3^T) \subset \operatorname{Im}\phi(B)$. The image of $\phi(B)$ is two-dimensional, and since $\xi([y]) \neq \xi([z])$, we have $\xi([x]) \subset \xi([y]) + \xi([z])$, as desired.

So, we have used the fundamental theorem of projective geometry to get some information on the behaviour of the map ϕ on rank one matrices. We can get additional information using the fundamental theorem of affine geometry. To explain the main idea we consider L_{e_1}, where e_1 is the first vector in the standard basis of $\mathbb{F}^m$. We know that $\phi(L_{e_1}) = L_u$ for some nonzero $u \in \mathbb{F}^m$. Both sets, L_{e_1} and L_u are isomorphic to $\mathbb{F}^n$ and we want to show that ϕ restricted to L_{e_1} considered as a map from L_{e_1} onto L_u behaves nicely. In particular we want to show that this is a semilinear map. By the fundamental theorem of affine geometry it is enough to show that this restriction of ϕ maps lines to lines. All we know is that ϕ preserves adjacency in both directions, and thus, we have to characterize lines in L_{e_1} and L_u using the adjacency relation. We do not want to go into all details. Our aim is to present the main idea only. So, let us just say here that such a characterization is possible and to explain the main idea we observe that the set of all elements of

$$L_{e_1} = \left\{ \begin{bmatrix} * & * & \dots & * \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \right\}$$

that are adjacent to the rank two matrix

$$\begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}$$

is exactly the line in L_{e_1} consisting of all matrices of the form

$$\begin{bmatrix} 1 & * & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}.$$

With the above two ideas we can completely describe the behaviour of ϕ on rank one matrices. This enables us to get the complete description of ϕ on the set of all matrices of rank two. Namely, every rank two matrix A is uniquely determined by the set of all rank one matrices that are adjacent to A. We continue with rank three matrices, and then with rank four matrices,... All these steps are far from being trivial. Nevertheless, the main ideas of the proof of Theorem 3 is to consider the sets L_x and R_y and to apply fundamental theorems of affine and/or projective geometry. All other parts of the proof require just technical skills.

3 Differential geometry

3.1 Basic concepts

Throughout this section we denote by M_n the algebra of all $n \times n$ complex matrices equipped with the usual operator norm. Let $N_n \subset M_n$ be the subset of all normal matrices. A continuous map $\gamma : [a, b] \to N_n$ is called a normal path or a normal curve. The matrices $\gamma(a)$ and $\gamma(b)$ are called the endpoints of this curve. The length of γ is defined to be

$$l_\gamma = \sup \sum_{k=0}^{m-1} \|\gamma(t_{k+1}) - \gamma(t_k)\|.$$

Here, the supremum is taken over all the partitions:

$$a = t_0 < t_1 < \ldots < t_m = b.$$

Of course, the supremum may be finite or infinite. In the first case we say that γ is rectifiable.

Assume that γ is a piecewise C^1 function. Then

$$l_\gamma = \int_a^b \|\gamma'(t)\| dt.$$

3.2 Spectral variation

Let A and B be $n \times n$ Hermitian matrices. It seems natural to expect that if B is a small perturbation of A, then the spectrum of B is close to the spectrum of A. This is indeed the case. A precise formulation will be given below. It is clear that this kind of statements are important in pure and applied mathematics.

Let

$$t_1 \geq t_2 \geq \ldots \geq t_n$$

and

$$s_1 \geq s_2 \geq \ldots \geq s_n$$

be the eigenvalues of A and B, respectively. Assume for a moment that A and B commute. Then they are simultaneously unitarily similar to diagonal matrices, that is, there exists a unitary matrix U such that

$$UAU^* = \begin{bmatrix} t_1 & 0 & \ldots & 0 \\ 0 & t_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & t_n \end{bmatrix} \quad \text{and} \quad UBU^* = \begin{bmatrix} s_{\tau(1)} & 0 & \ldots & 0 \\ 0 & s_{\tau(2)} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & s_{\tau(n)} \end{bmatrix},$$

where τ is a permutation on n symbols. It follows that

$$\|A - B\| = \|UAU^* - UBU^*\| = \max_j |t_j - s_{\tau(j)}|.$$

It is trivial to see that the minimal value of the right-hand side of this equality is achieved when $\tau = id$. Thus,

$$\max_j |t_j - s_j| \leq \|A - B\|.$$

We have seen that the above inequality holds true for all commuting pairs of Hermitian matrices A and B. Moreover, it is sharp.

Weyl's perturbation theorem says that the above statement holds true without the commutativity assumption.

Theorem 4 *Let A and B be any $n \times n$ Hermitian matrices with spectra $t_1 \geq t_2 \geq \ldots \geq t_n$ and $s_1 \geq s_2 \geq \ldots \geq s_n$, respectively. Then*

$$\max_j |t_j - s_j| \leq \|A - B\|.$$

The proof of this and other statements in this section can be found in the book: R. Bhatia, Matrix Analysis, Springer, New York, 1996.

We continue with normal matrices. Once again we first assume that normal matrices $A, B \in M_n$ commute. Let $\lambda_1, \ldots, \lambda_n$ and $\mu_1, \ldots, \mu_n$ be the eigenvalues of A and B, respectively. Then we can see in exactly the same way as above that

$$\min_\tau \max_j |\lambda_j - \mu_{\tau(j)}| \leq \|A - B\|.$$

Here, of course, j runs from 1 to n and τ runs over all permutations on n symbols. The expression on the left-hand side of this inequality is called the optimal matching distance between the unordered n-tuples $\{\lambda_1, \ldots, \lambda_n\}$ and $\{\mu_1, \ldots, \mu_n\}$.

Once again we can ask if the above inequality holds true also for non-commuting normal matrices A and B. This question was open for quite a long time and then finally answered in the negative by Holbrook. However, Bhatia, Davis, and McIntosh proved the following result.

Theorem 5 *There exists a constant c, $1 < c < 3$, such that the optimal matching distance $d(\sigma(A), \sigma(B))$ between the eigenvalues of any two normal matrices A and B is bounded as*

$$d(\sigma(A), \sigma(B)) \leq c\|A - B\|.$$

The optimal value of c is not known. An interesting question is under what additional conditions on the pair of normal matrices A and B we get the above inequality with $c = 1$. Weyl's theorem says that this happens if we assume that A and B are Hermitian. The same is true for unitary matrices.

Theorem 6 *Let A and B be unitary matrices. Then*

$$d(\sigma(A), \sigma(B)) \leq \|A - B\|.$$

In the next section we will try to explain the main geometrical ideas that lead to this result. This approach to the spectral variation problem was initiated by Bhatia and then further developed by Bhatia and Holbrook.

3.3 Differential geometry and spectral variation

We start with the following statement.

Lemma 7 *Let $A, B \in M_n$ and assume that A is normal. Then each eigenvalue of B is within a distance $\|A - B\|$ of some eigenvalue of A.*

Proof. Assume that the statement does not hold true. Let β be an eigenvalue of B whose distance to each eigenvalue of A is strictly larger than $\|A-B\|$. With no loss of generality we can assume that $\beta = 0$. It follows that A is invertible and all its eigenvalues are located outside the closed disk $D(0, \|A - B\|)$. As A is normal, we have

$$\|A^{-1}\| < \frac{1}{\|A - B\|},$$

and consequently, $\|A^{-1}(B - A)\| < 1$, which further yields that

$$B = A(I + A^{-1}(B - A))$$

is invertible, contradicting the fact that $\beta = 0 \in \sigma(B)$.

□

Corollary 8 *Let $A, B \in M_n$ with A being normal and assume that $\|A - B\|$ is smaller than half the distance between any two distinct eigenvalues of A. Then*

$$d(\sigma(A), \sigma(B)) \leq \|A - B\|.$$

Instead of giving a proof we will try to explain just the main idea. So, assume that A has three distinct eigenvalues with multiplicities 2, 2, and 3. Thus, the eigenvalues of A (counting thier multiplicities) are $\lambda_1, \lambda_1, \lambda_2, \lambda_2, \lambda_3, \lambda_3, \lambda_3$. The three disks with centers at $\lambda_1, \lambda_2, \lambda_3$ and radius $\|A - B\|$ are disjoint. All we have to show is that two eigenvalues of B belong to $\overline{D}(\lambda_1, \|A - B\|)$, two eigenvalues of B belong to $\overline{D}(\lambda_2, \|A - B\|)$, and three eigenvalues of B belong to $\overline{D}(\lambda_3, \|A - B\|)$. By Lemma 7, all eigenvalues of B belong to the disjoint union of these three closed disks. So, all we have to do is to show that they are distributed in the way we want (two in the first disk, two in the second disk, and three in the last one centered at λ_3). This can be done using a continuity argument.

We are now ready to present the main idea of this section. Assume that A and B are normal matrices and $\gamma(t)$, $t \in [a, b]$, a normal curve with the endpoints A and B. For every $t \in [a, b]$ the spectral variation inequality of Corollary 8 holds in a small neighbourhood of $\gamma(t)$. It follows (the details can be found in already mentioned Bhatia's book) that the total spectral variation between A and B is bounded by the length of the curve γ. More precisely, we have the following statement.

Theorem 9 *Let A and B be normal matrices and let γ be a rectifiable normal path joining them. Then*

$$d(\sigma(A), \sigma(B)) \leq l_\gamma.$$

Of course, $l_\gamma \geq \|A-B\|$. It is now clear that one possibility to find a "good" upper bound for $d(\sigma(A), \sigma(B))$ is to find the length of the normal geodesic between A and B. The problem of finding normal geodesics between two normal matrices is not easy as the geometry of the set N_n is not well-understood.

Here, we will consider two consequences of Theorem 9. The first one is an extension of Weyl's theorem.

Theorem 10 *If A, B are normal matrices such that $A-B$ is also normal, then*

$$d(\sigma(A), \sigma(B)) \leq \|A - B\|.$$

Proof. All we have to do is to show that if $A, B, A-B$ are all normal, then each matrix $A + t(B-A)$, $0 \leq t \leq 1$, is normal. The length of the line segment joining A and B is equal to $\|B - A\|$. So, the statement follows directly from Theorem 9.

□

It is not so straightforward to show that Theorem 6 also follows from Theorem 9. Namely, the straight line between two unitary matrices does not belong to N_n in general. Nevertheless, it is possible to show that the set $\mathbb{C}U_n$ of all scalar multiples of $n \times n$ unitary matrices is metrically flat, that is, any two matrices pA and qB, where $p, q \in \mathbb{R}$ and $A, B \in U_n$, can be joined with a path of length $\|pA - qB\|$ that lies entirely in $\mathbb{C}U_n$. Once we show that $\mathbb{C}U_n$ is metrically flat, Theorem 6 becomes a direct consequence of Theorem 9 (in fact, we conlude that Theorem 6 holds true under the weaker assumption that A and B are scalar multiples of unitary matrices).

To see that $\mathbb{C}U_n$ is metrically flat we choose matrices $pA, qB \in \mathbb{C}U_n$. As BA^{-1} is unitary, we may assume with no loss of generality that BA^{-1} is diagonal, $BA^{-1} = \operatorname{diag}(e^{i\theta_1}, \ldots, e^{i\theta_n})$ with $|\theta_n| \leq \ldots \leq |\theta_1| \leq \pi$. We will consider only the case when $|\theta_1| < \pi$. We have

$$\|pA - qB\| = \|p - qBA^{-1}\| = |p - qe^{i\theta_1}|.$$

Parametrise the straight line between p ans $qe^{i\theta_1}$ in the complex plane as $r(t)e^{it\theta_1}$, $0 \leq t \leq 1$. Set

$$A(t) = r(t)\operatorname{diag}(e^{it\theta_1}, \ldots, e^{it\theta_n})A, \quad 0 \leq t \leq 1.$$

One can verify that this is a smooth curve in $\mathbb{C}U_n$ with endpoints pA and qB and length $\|qB - pA\|$, as desired.

In particular, we have seen here that the geometry with respect to the operator norm is essentially different from the standard Riemannian geometry. In fact, this geometry is quite complicated with geodesics that are not straight lines. An interested reader can find more information on differential geometry related to spectral variation and geometric means in two articles that appeared in Math. Intelligencer (R. Bhatia, Spectral variation, normal matrices, and Finsler geometry, R. Bhatia and J. Holbrook, Noncommutative geometric means). The pdf files of both articles are available on the web page together with these lecture notes.

4 Algebraic geometry

4.1 Basic concepts

We denote by $\mathbb{C}[X_1, \ldots, X_n]$ the ring of all complex polynomials in n indeterminates. Let S be a subset of $\mathbb{C}[X_1, \ldots, X_n]$ and denote by I the ideal in $\mathbb{C}[X_1, \ldots, X_n]$ generated by S. It is not difficult to verify that I consists of all polynomials of the fom

$$p_1 q_1 + \ldots + p_m q_m,$$

where m is a positive integer, $p_1, \ldots, p_m \in S$, and $q_1, \ldots, q_m$ are arbitrary polynomials in $\mathbb{C}[X_1, \ldots, X_n]$.

Theorem 11 *Each ideal in $\mathbb{C}[X_1, \ldots, X_n]$ is finitely generated.*

For an arbitrary ideal $I \subset \mathbb{C}[X_1, \ldots, X_n]$ we define

$$\mathrm{rad}\,(I) = \{p \in \mathbb{C}[X_1, \ldots, X_n] \, : \, p^m \in I \text{ for some } m \in \mathbb{N}\}.$$

It tuns out that rad (I) is an ideal.

An algebraic subset $V(S)$ of $\mathbb{C}^n$ is the set of common zeroes of some set $S \subset \mathbb{C}[X_1, \ldots, X_n]$:

$$V(S) = \{(a_1, \ldots, a_n) \, : \, p(a_1, \ldots, a_n) = 0 \text{ for all } p \in S\}.$$

It is rather easy to check that $V(S) = V(I)$, where I is the ideal gnerated by S (note that I is finitely generated, and consequently, we may always assume that S is a finite set).

Let $p = p(X, Y)$ be a polynomial in 2 indeterminates. The set $\mathcal{C}$ of all points $(a, b) \in \mathbb{C}^2$ satisfying the equation

$$p(a, b) = 0$$

is called a curve in $\mathbb{C}^2$. The equation of the tangent in the point $(a, b) \in \mathcal{C}$ is

$$\frac{\partial p}{\partial X}(a, b)\,(X - a) + \frac{\partial p}{\partial Y}(a, b)\,(Y - b) = 0.$$

If $\frac{\partial p}{\partial X}(a,b) \neq 0$ or $\frac{\partial p}{\partial Y}(a,b) \neq 0$, then the tangent is a line and we say that (a,b) is a smooth point. Otherwise the tangent is $\mathbb{C}^2$ and we say that (a,b) is a singular point.

In the case that (a,b) is a smooth point the equation of the tangent at (a,b) can be writen as

$$(X,Y) = (a,b) + \lambda(u,v)$$

where (u,v) is the vector satisfying the equation

$$\frac{dp}{d\lambda}((a,b) + \lambda(u,v))_{|\lambda=0} = 0.$$

We say that the tangent space at the point (a,b) is a one-dimensional subspace of $\mathbb{C}^2$ spanned by (u,v).

More generally, let $A \subset \mathbb{C}^n$ be an algebraic set. Let $A = V(I)$ for some ideal I and let $a \in A$. The tangent space at the point a is defined as

$$\tan(I,a) = \{b \in \mathbb{C}^n \ : \ \frac{dp}{d\lambda}(a + \lambda b)_{|\lambda=0} = 0 \ \text{ for all } \ p \in I\}.$$

Similarly, we set

$$\tan(\mathrm{rad}\,(I),a) = \{b \in \mathbb{C}^n \ : \ \frac{dp}{d\lambda}(a + \lambda b)_{|\lambda=0} = 0 \ \text{ for all } \ p \in \mathrm{rad}\,(I)\}.$$

Both are vector spaces and $\tan(\mathrm{rad}\,(I),a) \subset \tan(I,a)$.

4.2 Linear preservers of nilpotents

In the second section we have mentioned the problem of characterizing linear maps preserving invertibility. Let $\phi : M_n \to M_n$ be a bijective linear invertibility preserving map on the algebra of all $n \times n$ complex matrices. When studying such maps there is no loss of generality in assuming that $\phi(I) = I$. Indeed, if this is not the case we can replace ϕ by the map $A \mapsto \phi(I)^{-1}\phi(A)$, $A \in M_n$. This map is linear, bijective and unital. Clearly, it preserves invertibility.

So, we may and we will assume that ϕ is unital, that is, $\phi(I) = I$. Then we have

$$\sigma(\phi(A)) \subset \sigma(A)$$

for every $A \in M_n$. Indeed,

$$\lambda \notin \sigma(A) \Rightarrow \lambda I - A \ \text{ is invertible}$$

$$\Rightarrow \lambda I - \phi(A) \ \text{ is invertible} \Rightarrow \lambda \notin \sigma(\phi(A)).$$

In particular, $\phi(\mathcal{N}) \subset \mathcal{N}$, where $\mathcal{N} = \{N \in M_n \ : \ N^n = 0\}$ is the set of all nilpotents.

So, instead of trying to characterize linear maps on M_n preserving invertibility one can think of an even more general problem of characterizing linear maps on M_n preserving nilpotents. We denote by $sl_n \subset M_n$ the subspace of all trace zero matrices. Clearly, all nilpotents are trace zero matrices. It is not difficult to see that $\operatorname{span} \mathcal{N} = sl_n$. Hence, when considering linear preservers of nilpotent matrices we may as well assume that these linear maps are defined on sl_n. The following theorem was proved by Botta, Pierce and Watkins (Pacific J. Math. 104 (1983), 39- 46.

Theorem 12 *Let $n \geq 3$ and let $\phi : sl_n \to sl_n$ be a bijective linear map such that $\phi(\mathcal{N}) \subset \mathcal{N}$. Then there exist an invertible matrix P and a nonzero scalar c such that either*

$$\phi(A) = cPAP^{-1}, \quad A \in sl_n,$$

or

$$\phi(A) = cPA^T P^{-1}, \quad A \in sl_n.$$

The main idea is the same as explained in the second section. We want to show that ϕ preserves rank one matrices. As we are dealing with trace zero matrices only, we want to show that every nilpotent of rank one is mapped into a nilpotent of rank one. And then the standard methods (a slight modification of methods from the second section) yield our result. The set of all nilpotent matrices is an algebraic set. Thus, it is not surprising that algebraic geometry plays an important role in proving the fact that $\phi(N)$ is a nilpotent of rank one for every nilpotent N of rank one. The main ideas of the proof of this statement will be explained in the next section.

4.3 Algebraic geometry and linear nilpotent preserving maps

We start with some preliminaries. Then we will sketch the proof of Theorem 12. All the details can be found in the paper: P. Botta, S. Pierce, and W. Watkins, Linear transformations that preserve the nilpotent matrices, Pacific J. Math. 104 (1983), 39-46.

Clearly, $\mathcal{N} \subset M_n$ is an algebraic set. Namely, $A \in \mathcal{N}$ if and only if $A^n = 0$. That means that each entry of A^n must be zero. Thus, the set $\mathcal{N}$ is determined by n^2 algebraic equations.

A more efficient way (with less equations) is the following one. We have $A \in \mathcal{N}$ if and only if the characteristic polynomial of A is $p(X) = X^n$. This means that all coefficients of the characteristic polynomial but the leading one are zero. Hence, $\mathcal{N} = V(I)$, where I is the ideal of polynomials in n^2 indeterminates generated by $E_1, \ldots, E_n$. Here, $E_r(A)$ is the r-th elementary symmetric function of the matrix A, that is, $E_r(A)$ is the sum of all principal $r \times r$ subdeterminants. Thus, the set $\mathcal{N}$ can be completely determined by only n algebraic equations.

The matrix $A \in \mathcal{N}$ is said to be of nilindex n if $A^{n-1} \neq 0$. Clearly, $A \in \mathcal{N}$ is of nilindex n if and only if it is similar to

$$N = \begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & 0 & \dots & 0 \end{bmatrix}.$$

It is straightforward to verify that $B \in M_n$ is a nilpotent of nilindex n that commutes with N if and only if B is a strictly upper triangular Toeplitz matrix with a nonzero first super diagonal, that is

$$B = \begin{bmatrix} 0 & b_1 & b_2 & b_3 & \dots & b_{n-1} \\ 0 & 0 & b_1 & b_2 & \dots & b_{n-2} \\ 0 & 0 & 0 & b_1 & \dots & b_{n-3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & b_1 \\ 0 & 0 & 0 & 0 & \dots & 0 \end{bmatrix}.$$

with $b_1 \neq 0$. One can verify that

$$\tan(\mathrm{rad}(I), N) = \tan(\mathrm{rad}(I), B).$$

It turns out that the above property is characteristic for commuting pairs of nilpotents of maximal nilindex.

Lemma 13 *Let $A, B \in \mathcal{N}$ be both of nilindex n. Then $AB = BA$ if and only if*

$$\tan(\mathrm{rad}(I), A) = \tan(\mathrm{rad}(I), B).$$

We denote by $\mathcal{T}$ the algebra of all strictly upper triangular matrices,

$$\mathcal{T} = \left\{ \begin{bmatrix} 0 & * & * & * & \dots & * \\ 0 & 0 & * & * & \dots & * \\ 0 & 0 & 0 & * & \dots & * \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & * \\ 0 & 0 & 0 & 0 & \dots & 0 \end{bmatrix} \right\}$$

Clearly $\mathcal{T}$ is a subalgebra of M_n consisting of nilpotent mastrices only and $\dim \mathcal{T} = \frac{n(n-1)}{2}$. Gerstenhaber (M. Gerstenhaber, On nilalgebras and linear varieties of nilpotent matrices I, Amer. J. Math. 80 (1958), 614-622) proved the following theorem.

Theorem 14 *Let $\mathcal{S} \subset \mathcal{N}$ be a linear subspace. Then*

$$\dim \mathcal{S} \leq \frac{n(n-1)}{2}.$$

If

$$\dim \mathcal{S} = \frac{n(n-1)}{2}$$

then there exists an invertible matrix P such that $P\mathcal{S}P^{-1} = \mathcal{T}$.

Let us make two more observations. It is easy to see that

$$Z(\mathcal{T}) = \{M \in \mathcal{T} : MK = KM \text{ for every } K \in \mathcal{T}\} = \text{span}\,\{E_{1n}\}.$$

Here, E_{1n} denotes the matrix whose all entries are zero but the $(1,n)$-entry which is equal to 1.

Any subalgebra similar to $\mathcal{T}$ will be called a maximal triangularizable nilpotent subalgebra. One can verify that each nilpotent matrix of nilindex n is contained in exactly one maximal triangularizable nilpotent subalgebra. And it is also easy to see that any nilpotent M which is not of maximal nilindex (that is, it satisfies $M^{n-1} = 0$) belongs to at least two maximal triangularizable nilpotent subalgebras.

We are now ready to sketch the proof of Theorem 12. So, we have a bijective linear map $\phi : sl_n \to sl_n$ with $\phi(\mathcal{N}) \subset \mathcal{N}$. We want to show that it maps rank one nilpotents into rank one nilpotents. In the paper: J. Dixon, Rigid embeddings of simple groups in the general linear group, Canad. J. Math. 29 (1977), 384-391, one can find a proof (rather short, but tricky) that if a bijective linear map on a finite-dimensional vector space maps a certain algebraic set into itself, then it actually maps it onto itself. Thus, we have $\phi(\mathcal{N}) = \mathcal{N}$. It follows from Gerstenhaber's result that ϕ maps every maximal triangularizable nilpotent subalgebra onto some maximal triangularizable nilpotent subalgebra. By the remark in the previous paragraph we have for every $A \in sl_n$:

$$A \text{ is nilpotent of nilindex } n \iff \phi(A) \text{ is nilpotent of nilindex } n.$$

Let M be a nilpotent of rank one. We may assume with no loss of generality that $M = E_{1n}$. Moreover, we know that $\phi(\mathcal{T})$ is similar to $\mathcal{T}$. So, after composing ϕ with a similarity transformation we may assume that

$$\phi(\mathcal{T}) = \mathcal{T}.$$

The set of all nilpotents $A \in \mathcal{T}$ of nilindex n is mapped onto itself. For every such A, the matrix $A + E_{1n}$ is again a nilpotent of nilindex n commuting with A. We have

$$\phi(\tan(\text{rad}\,(I), A)) = \tan(\text{rad}\,(I), \phi(A)).$$

Hence, $\phi(A) + \phi(E_{1n})$ is a nilpotent of nilindex n commuting with $\phi(A)$. It follows that for every $B \in \mathcal{T}$ of nilindex n the matrix $B + \phi(E_{1n})$ commutes with B. The set of nilpotents of nilindex n is dense in $\mathcal{T}$. Hence, $\phi(E_{1n})$ commutes with every $B \in \mathcal{T}_n$. Consequently, $\phi(E_{1n}) \in \operatorname{span}\{E_{1n}\}$. We have shown that rank one nilpotents are mapped into rank one nilpotents, as desired.

Let us conclude these lecture notes with one remark on geometrical techniques and elementary proofs. As $\mathcal{N}$ is an algebraic set it is natural to expect that algebraic geometry could be an important tool when proving the characterization of linear nilpotent preserving maps. The same is true for the above mentioned Gerstenhaber's result. However, once we have a result in linear algebra that was proved using deep ideas from other parts of mathematics (analysis, topology, algebra, geometry,...) we still have an open question. Is it necessary to use difficult techniques to solve the linear algebra problem? Or can we prove the result using just elementary linear algebra techniques? Let us mention here that an elementary proof of Gerstenhaber's result has been found by Mathes, Omladič, and Radjavi (Linear spaces of nilpotent matrices, Linear Algebra Appl. 149 (1991), 215-225) and a very short elementary proof of Theorem 12 has been found by the author of these lecture notes (Characterization of matrices having rank k, Linear and Multilinear Algebra 42 (1997), 233-238).