Linear Algebra 2
Minimal polynomial, generalized eigenvectors and the JNF, inner products and generalizations
by Ivan Aidun

These are notes interspersed with exercises. The purpose of these notes is to be more fleshed out than Evan Dummit's old notes, but shorter and more focused than a textbook treatment. Halfway between Dummit and Dummit and Foote, so to speak. I hope these are helpful to you!

Even more eigen-stuff!

This section talks about what I think are the most important linear algebra constructions to know for the Qual: the minimal polynomial and Jordan Normal (or sometimes Canonical) Form.

Minimal polynomial

In the previous notes we discussed the Cayley-Hamilton Theorem, which says that A satisfies its own characteristic polynomial. Since A satisfies some polynomial, there must be a polynomial A satisfies of least degree. Such a polynomial is called the **minimal polynomial** of A.

Exercise 1. (Basic properties of the minimal polynomial) Let A be a linear operator on a finite-dimensional complex vector space V. Let $p(x)$ be the minimal polynomial of A.

(a) Prove that if f is a polynomial A satisfies, then $p \mid f$.

(b) Prove that only eigenvalues of A can be roots of p.

(c) Prove that if λ is an eigenvalue of A, then λ is in fact a zero of p. (Hint: let v be an eigenvector, and look at $p(A)v$.)

(d) Prove that two similar matrices have the same minimal polynomial.

(e) Prove that a matrix A is diagonalizable if and only if its minimal polynomial has no repeated roots. (Hint: think about $\dim\ker A - \lambda I$ as λ ranges over the roots of p.)

(f) Let $W \subset V$ be a subspace that is A-stable, that is, so that for every $w \in W$, $Aw \in W$. Then we can restrict the operator $A\colon V \to V$ to an operator $A|_W\colon W \to W$. Let $q(x)$ be the minimal polynomial of the restriction $A|_W$. What is the relation between $q(x)$ and $p(x)$?

How can the minimal polynomial differ from the characteristic polynomial? The next exercise offers one example.

Exercise 2. Take I_n the $n \times n$ identity matrix.

(a) What is its characteristic polynomial?

(b) What is its minimal polynomial?

(c) What is the characteristic/minimal polynomial of the shearing transformation

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}?$$

The above exercise suggests that the minimal polynomial has low degree when A has few eigenvalues and many eigenvectors. Indeed, in the next section we will see the exact sense in which this is true.

Generalized eigenspaces and Jordan Normal Form

As mentioned in the previous set of notes, v being an eigenvector for A is equivalent to v being a nontrivial vector in the kernel of $\lambda I - A$ for some λ. We saw two ways that a matrix could fail to be diagonalizable: if the field in question wasn't algebraically closed, the eigenvalues (and hence eigenvectors) might not be in the base field, or even over an algebraically closed field A could fail to have n linearly independent eigenvectors. Examples of the latter include

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

For both of these matrices, the obstacle preventing them from being diagonalizable is the same: they have only one eigenvalue, and the matrix $\lambda I - A$ is nilpotent, but not zero. We say a vector v is a **generalized eigenvector** of eigenvalue λ if v is in the kernel of $(\lambda I - A)^k$ for some natural number k. The smallest integer k such that $(\lambda I - A)^k v = 0$ I will call the **degree** of the generalized eigenvector v. (Wikipedia calls it the "rank", but I feel like that's a misleading term in this context because it is kinda the opposite of the rank of a matrix. I prefer "degree" since it aligns with the terminology "nilpotence degree".) Thus, an ordinary eigenvector is an eigenvector of degree 1.

The following exercise shows that the generalized eigenvectors are exactly the vectors we are "missing" if we want to find a good matrix similar to A.

Exercise 3. Let V be a finite-dimensional complex vector space, and $A\colon V \to V$ a linear transformation.

(a) For each eigenvalue λ of A, let $V_{\lambda,k} = \{v \in V : v \text{ a } \lambda\text{-generalized-eigenvector of degree } \leq k\}$. Prove that $V_{\lambda,k}$ is a subspace of V, and that $AV_{\lambda,k} \subset V_{\lambda,k}$. The subspace $V_{\lambda,k}$ is the **degree k λ-generalized-eigenspace** of A.

(b) Suppose v_k is a λ-generalized-eigenvector of degree k. Show that $v_{k-1} = (\lambda I - A)v_k$ is a λ-generalized-eigenvector of degree $k-1$. By induction, this shows that $(\lambda I - A)^\ell v_k$ is a generalized eigenvector of degree $k - \ell$. The sequence of vectors $v_k, v_{k-1}, \ldots, v_1$ is called a **Jordan chain**.

(c) Let $V_{\lambda,\max}$ be the maximal λ-generalized-eigenspace for A, that is, $V_{\lambda,\max} = V_{\lambda,k}$ for the maximal degree k of a λ-generalized-eigenvector of A. Let $n = \dim V_{\lambda,1}$. Show that $V_{\lambda,\max}$ has a basis consisting of n disjoint Jordan chains.

(d) Let $\mathcal{B}$ be a basis as in the previous part. What does the matrix of A with respect to $\mathcal{B}$ look like?

The form of matrix you found in part (d) above is called a **Jordan block**. The above exercise implies that for any matrix A, we can find a basis for V so that the matrix for A is almost diagonal: it is block-diagonal, with each block a Jordan block. This form of matrix is called the **Jordan**

normal form of A. Moreover, it is easy to see that up to rearranging the blocks, this form is unique. This is worth stating as a theorem.

Theorem. Two matrices A, A' are similar if and only if they have the same Jordan normal form (up to rearranging the blocks).

Here's a more concrete exercise if you're having trouble connecting the discussion in Exercise 3 to the matrix form you expect to be seeing.

Exercise 4. (A concrete JNF computation) Consider the matrix

$$A = \begin{pmatrix} 4 & 1 & 0 \\ -11 & -3 & -1 \\ 5 & 2 & 2 \end{pmatrix}.$$

I will tell you right now that the only eigenvalue of A is 1.

(a) Given the fact that the only eigenvalue of A is 1, find a basis for the 1-eigenspace of A. (This is a kernel computation, which you know how to do.)

(b) Based on your computation from part (a), how many Jordan chains will you need to find in order to span $\mathbb{R}^3$?

(c) (Optional) Write down a basis for the degree 2 generalized eigenspace of A.

(d) Choose a generalized eigenvector for A of maximum degree. Write down its Jordan chain.

(e) Find a basis for $\mathbb{R}^3$ consisting of Jordan chains. Let B be the matrix with those vectors as its columns. Compute B^{-1}.

(f) Write down the matrix of A with respect to this basis by computing $B^{-1}AB$. Notice that this is the JNF.

(g) If you were writing an exercise for someone else to compute a JNF, how would you come up with a matrix that makes it so the problem isn't trivial, but also so that it isn't too hard? Asking for a friend...

Exercise 5.

(a) Given a matrix in Jordan normal form, how can you tell what its characteristic polynomial is?

(b) If the matrix A has only one Jordan block, what is its minimal polynomial?

(c) If the matrix A has two Jordan blocks (possibly for different eigenvalues) what is its minimal polynomial?

(d) In general, given a matrix in Jordan normal form, how can you tell what its minimal polynomial is?

Exercise 6.

(a) Give an example of two linear transformations with the same characteristic polynomial but different minimal polynomials.

(b) Give an example of two linear transformations on the same vector space with the same minimal polynomials which are not similar.

Commuting matrices

There's a standard trick involving commuting matrices which comes up often enough that I want to mention it.

Exercise 7. (August 2014 Problem 3 (d)) Suppose that T and U are two linear transformations on a finite-dimensional vector space V over $\mathbb{C}$ which satisfy $TU = UT$. Prove that there is some $v \in V$ which is an eigenvector for both T and U.

There are lots of variations on this theme, but the moral is that whenever you see two matrices which commute (or almost do so) you should be looking for a trick along these lines. One particular application of the above fact is given in the following exercise.

Exercise 8. Suppose that T and U are as in Exercise 7.

(a) Suppose T and U are diagonalizable. Then prove that in fact, there is a basis such that the matrices of T and U are *both* diagonal with respect to that basis. This is called **simultaneous diagonalization**, and it appears in various parts of math, but I learned of it because in number theory you simultaneously diagonalize Hecke operators to learn things about modular forms/elliptic curves.

(b) One might suspect that without the assumption of diagonalizability, we might be able to simultaneously *Jordanize* T and U (that is, put them into simultaneous Jordan form). Find a counterexample showing that this is false.

Inner products and bilinear forms

Inner products

Something that usually comes up in an introductory linear algebra class but which has so far not come up at all in these notes are inner products. The vector space $\mathbb{R}^n$ comes with an important function usually called the **dot product**, or sometimes the "standard inner product". It is defined by $u \cdot v = u_1v_1 + u_2v_2 + \cdots + u_nv_n$, where the u_i and v_i are the coordinates of the vectors u and v. This comes up all the time in physical applications.

The essential properties of the dot product are encoded in the definition of an **inner product**. An inner product on a real vector space V is a function $V^2 \to \mathbb{R}$, often denoted with angle brackets $(u, v) \mapsto \langle u, v\rangle$, with the following properties.

(i) (bilinearity) $\langle au_1 + bu_2, v\rangle = a\langle u_1, v\rangle + b\langle u_2, v\rangle$, and similarly on the other side.

(ii) (symmetry) $\langle u, v\rangle = \langle v, u\rangle$

(iii) (positive definite) We have $\langle v, v\rangle \geq 0$, with equality if and only if $v = 0$.

Inner products can be used to define the angle between two vectors in an arbitrary vector space. We won't need this fact, except to mention that two vectors are **orthogonal** if $\langle v, u\rangle = 0$. They also can be used to define the length of a vector via $\|v\| = \sqrt{\langle v, v\rangle}$. This does in fact induce a metric, which is a consequence of the Cauchy-Schwarz theorem.

Given a subspace U of a vector space V, the **orthogonal complement** $U^\perp$ is the set of all vectors which are orthogonal to everything in U: $U^\perp = \{v \in V : \forall u \in U,\ \langle v, u\rangle = 0\}$. The orthogonal complement to a subspace is "the best" complement to that subspace, because by design the vectors from the original subspace essentially do not interact at all with those in the orthogonal space.

Exercise 9.

(a) Give an example of an inner product on $\mathbb{R}^2$ that is not the dot product.

(b) Let $\mathcal{C}([0, 1])$ be the set of real-valued continuous functions on the interval $[0, 1]$. Define a pairing by

$$\langle f, g\rangle = \int_0^1 f(x)g(x)\ \mathrm{d}x.$$

Prove that this is an inner product.

Exercise 10. (August 2020 Problem 2) Let M be an $n \times n$ matrix with real entries. Let M^T be its transpose. Each of these two matrices defines a linear transformation $\mathbb{R}^n \to \mathbb{R}^n$; let A denote the transformation corresponding to M, and let B denote that corresponding to M^T.

(a) Prove that any vector in $\ker(A)$ is orthogonal to any vector in $\mathrm{im}(B)$ with respect to the standard inner product (dot product).

(b) Prove that $\mathbb{R}^n$ is the direct sum of $\ker(A)$ and $\mathrm{im}(B)$. That is to say, show that for any vector $u \in \mathbb{R}^n$, there exist unique vectors $v \in \ker(A)$ and $w \in \mathrm{im}(B)$ so that $u = v + w$.

Using inner products, one can create especially nice bases: orthonormal bases. In an **orthonormal basis**, each basis vector has length 1 and is orthogonal to all the other basis vectors. If $\{e_1, \ldots, e_n\}$ is an orthonormal basis, it is particularly easy to determine the coefficient of e_i in the expansion of an arbitrary vector v: it is just $\langle e_i, v\rangle$.

Exercise 11. Consider the functions $\{\sqrt{2}\sin(2n\pi x), \sqrt{2}\cos(2n\pi x)\} \subset \mathcal{C}([0, 1])$. Prove that these functions are orthonormal with respect to the inner product from Exercise 9. (Hint: Look up your double angle formulas, and your product-to-sum formulas.)

(The functions $\{\sqrt{2}\sin(2n\pi x), \sqrt{2}\cos(2n\pi x)\}$ don't quite form a basis in the traditional sense, because this is a countable set and $\mathcal{C}([0, 1])$ has dimension equal to the cardinality of the continuum. But, the span of these functions is dense, which in the context is a more natural condition to work with. Analysts call the normal definition of a basis an "algebraic basis" or "Hamel basis", and they will often call a system like this a "Schauder basis" or just a "basis" if there isn't an algebraist in earshot.)

Exercise 12. Let O be an $n \times n$ real matrix. Show that the following are equivalent.

(i) The columns of O form an orthonormal basis.

(ii) $O^{-1} = O^T$.

If O satisfies either of the conditions in the preceding problem, it is called an **orthogonal matrix**. The set of $n \times n$ orthogonal matrices is a group, called the **orthogonal group**, and denoted $\mathrm{O}(n)$. The subgroup of orthogonal matrices with determinant 1 is called the **special orthogonal group**, and is denoted $\mathrm{SO}(n)$.

Exercise 13. (Spectral theorem) Let V be a finite-dimensional inner product space. A linear transformation $A\colon V \to V$ is called **symmetric** if $\langle Au, v\rangle = \langle u, Av\rangle$.

(a) Let $\{e_1, \ldots, e_n\}$ be an orthonormal basis for V. Show that A is symmetric if and only if the matrix of A is symmetric with respect to the chosen basis.

(b) Let A be a symmetric linear operator. Show that there is some orthonormal basis such that the matrix of A is diagonal.

Other products

The term "inner product space" is also applied to complex vector spaces, altho there it has a somewhat different definition. An inner product on a complex vector space V is a function $(V)^2 \to \mathbb{C}$ satisfying the following.

(i) (linearity in the second term) $\langle u, av_1 + bv_2\rangle = a\langle u, v_1\rangle + b\langle u, v_2\rangle$.

(ii) (conjugate-linearity in the first term) $\langle au_1 + bu_2, v\rangle = \bar{a}\langle u_1, v\rangle + \bar{b}\langle u_2, v\rangle$

(iii) (conjugate symmetry) $\langle u, v\rangle = \overline{\langle v, u\rangle}$

(iv) (positive definite) $\langle v, v\rangle \geq 0$, with equality if and only if $v = 0$. (Note that conjugate symmetry forces $\langle v, v\rangle \in \mathbb{R}$, so it makes sense to talk about whether or not it is greater than 0.)

Of course, it is a matter of convention whether or not you take an inner product to be conjugate linear in the first variable and linear in the second or visa versa. I think the convention above is the most common one, on account of the fact that it makes $u^T v$ into the standard inner product. A function satisfying only (i)-(iii) above is called a **Hermitian product**, and a function satisfying only (i)-(ii) is sometimes called a **sesquilinear product**. Over complex inner product spaces, the analog of an orthogonal matrix is called a **unitary matrix**.

You might wonder (as I have) why when we extend the definition of an inner product in a way that treats the two variables differently, rather than demanding bilinearity which seems more straightforward and natural. One reason is that without conjugate-linearity, we cannot guarantee that $\langle v, v\rangle \in \mathbb{R}$, which we want in order to talk about whether or not it is positive-definite. A related reason is that if an inner product were required to be symmetric and bilinear instead of Hermitian, you could get strange behavior extending an inner product from a real vector space contained inside a complex vector space. For example, take $\mathbb{R}^2$ with the dot product contained inside $\mathbb{C}^2$, and let $v_1, v_2 \in \mathbb{R}^2$ be orthonormal. If we extended the dot product to $\mathbb{C}^2$ bilinearly instead of sesquilinearly, we would get that $\langle v_1 + iv_2, v_1 + iv_2\rangle = \langle v_1, v_1\rangle + i^2\langle v_2, v_2\rangle = 1 - 1 = 0$, so $v_1 + iv_2$ would be a nonzero vector with "length 0".

So far our discussion has been limited to just vector spaces over $\mathbb{R}$ and $\mathbb{C}$, but it would be pretty sad if there were no analogs for vector spaces over other fields. Luckily, there are. For a vector space V over a field k, a bilinear function $(V)^2 \to k$ is called a **bilinear form**, or sometimes just a **pairing**. (Altho sometimes you can have pairings between different vector spaces.) Often people think about symmetric bilinear forms in imitation of the inner product case, but the condition of positive-definiteness cannot be straightforwardly ported to an arbitrary field. In its place, the next best thing is to require a bilinear form to be **nondegenerate:** if $\langle u, v\rangle = 0$ for all $u \in V$, then $v = 0$. Symmetric bilinear forms are closely related to what are called "quadratic forms": polynomials in n variables each of whose terms has total degree 2.

Exercise 14. Show that the set of bilinear forms on V is a vector space.

There are other kinds of products considered in various parts of mathematics. For example, a **symplectic product** on a vector space is a nondegenerate bilinear form that is required to be *anti*-symmetric, rather than symmetric: $\langle u, v\rangle = -\langle v, u\rangle$. These don't appear on the qual, so I will say no more of them other than that they exist and some people care about them.

The dual vector space

Discussing bilinear forms is a great time to mention a construction that is good to know: the dual vector space. Given a vector space V over a field k, the dual vector space is the set of all linear maps $V \to k$. This does in fact form a vector space, variously denoted $V^\vee$, V^*, V', $\overline{V}$, $\widetilde{V}$, and V^T. I will use $V^\vee$. The elements in $V^\vee$ are variously called **covectors** or **functionals** depending on context.

Exercise 15. Let V be a vector space over a field k.

(a) Suppose that $\dim V = n$ and that $\{v_1, \ldots, v_n\}$ is a basis for V. Define $\phi_i \colon V \to k$ to be the unique map such that $\phi_i(v_j) = 1$ if $i = j$ and 0 otherwise. Show that the ϕ_i form a basis for $V^\vee$ (and, in particular, that $\dim V = \dim V^\vee$). This basis is called the **dual basis** to $\{v_1, \ldots, v_n\}$.

(b) If you are a set theory lover, show that $\dim V$ need not equal $\dim V^\vee$ if $\dim V$ is not finite. (This is pretty similar to showing that the power set of $\mathbb{N}$ is uncountable.)

(c) Suppose $A \colon V \to V$ is an invertible linear transformation. Then $\{Av_1, \ldots, Av_n\}$ is also a basis for V. What is the relationship between the dual basis to $\{Av_1, \ldots, Av_n\}$ and the dual basis to $\{v_1, \ldots, v_n\}$?

Exercise 16. Let V, W be vector spaces over a field k. Denote by $\mathrm{Hom}_k(V, W)$ the set of all linear maps $V \to W$.

(a) Show that $\mathrm{Hom}_k(V, W)$ is a vector space.

(b) If $\dim V = m$ and $\dim W = n$, what is $\dim \mathrm{Hom}_k(V, W)$?

Exercise 17. I'm thinking of a linear isomorphism between the vector space of bilinear forms on V and $\mathrm{Hom}_k(V, V^\vee)$. This isomorphism can be written out explicitly, and doesn't require me to tell you anything about V or about k (in particular, it doesn't require me to tell you a basis for V). Can you find the isomorphism?

Given a linear map $A\colon V \to W$, there is an associated linear map $A^\vee\colon W^\vee \to V^\vee$ given by "pulling back functionals": $A^\vee(f)(v) = f(Av)$. **Exercise 18.** Check that if $V = \mathbb{R}^m$ and $W = \mathbb{R}^n$ with the standard bases, then the matrix of $A^\vee$ with respect to the dual bases is A^T.

"Generalized products" and tensor products

The great utility of inner products, and bilinear forms more generally, induced people to want more general kinds of products, which could be between two different vector spaces, and were required to satisfy nothing other than bilinearity. Given vector spaces V and W, we can introduce the symbol $v \otimes w$ to stand in for the value this "most general possible product" would take on the pair $(v, w) \in V \times W$. Then, the assumption of bilinearity forces upon us certain relations among these symbols:

- $(v_1 + v_2) \otimes w = v_1 \otimes w + v_2 \otimes w$, and conversely on the other side.
- $(av) \otimes w = a(v \otimes w) = v \otimes (aw)$.

This, in fact, has a name, given to it by Einstein in the 1910s: these symbols subject to these relations are called **tensors**. The vector space of all tensors from V and W is called the **tensor product** of V and W, denoted $V \otimes_k W$. A tensor can be a linear combination of symbols, like $v_1 \otimes w_1 + v_2 \otimes w_2 + \dots$. The tensors that can be written as $v \otimes w$ are called **simple tensors**.

Exercise 19. Suppose V, W are vector spaces over a field k, and that $\{v_1, \dots, v_m\}$ is a basis for V and $\{w_1, \dots, w_n\}$ is a basis for W. Show that the set $v_i \otimes w_j$ is a basis for $V \otimes_k W$.

Exercise 20. (January 2022 Problem 4) Let V and W be vector spaces over a field k, not necessarily finite dimensional. As usual, we denote by $\mathrm{Hom}_k(V, W)$ the space of linear maps from V to W, and by $V^\vee$ the dual of V, so that $V^\vee = \mathrm{Hom}_k(V, k)$.

(a) I am thinking of an injective linear map

$$\Phi\colon V^\vee \otimes_k W \to \mathrm{Hom}(V, W).$$

This map can be written down explicitly, and doesn't require me to tell you anything further about V, W or about k. Can you find the map?

(b) Prove that the operator Φ is an isomorphism if V and W are finite-dimensional.

(c) Prove a necessary and sufficient condition that a linear map $A \in \mathrm{Hom}(V, W)$ is in the image of Φ. (The assumption that V and W are finite-dimensional does not apply to this part.)

If you have a vector space V defined over a field K, and a field L which contains K, then (as we'll discuss when it comes to field theory) L is itself also a K-vector space. For example, $\mathbb{C}$ can be viewed as a two-dimensional real vector space. Then the tensor product $V \otimes_K L$ has a particularly nice interpretation: it is just the vector space V extended to be a vector space over L instead of over K. For example, $\mathbb{R}^3 \otimes_{\mathbb{R}} \mathbb{C} \cong \mathbb{C}^3$ is just what we get when we take a basis for $\mathbb{R}^3$, and declare that we will form a complex vector space with exactly that basis. This process is called **extension of scalars** and it is very good to know about. The next exercise shows how we can use extension of scalars to understand a little bit about the structure of orthogonal transformations.

Exercise 21. Let V be a finite-dimensional real inner product space. Let $O\colon V \to V$ be an orthogonal transformation. Let $V_{\mathbb{C}} = V \otimes_{\mathbb{R}} \mathbb{C}$ be the **complexification** of V. The transformation O extends to a transformation $V_{\mathbb{C}} \to V_{\mathbb{C}}$, which by abuse of notation I will also call O.

(a) Show that all the eigenvalues of O in $\mathbb{C}$ lie on the unit circle.

(b) Show that O is diagonalizable. (Hint: $\langle v, w\rangle = \langle Ov, Ow\rangle$. Try looking at the orthogonal complement to an eigenvector.)

(c) If $v \in V_{\mathbb{C}}$ is an eigenvector of O of eigenvalue $\lambda \in \mathbb{C} \setminus \mathbb{R}$, then $\bar{v}$ is also an eigenvector of eigenvalue $\bar{\lambda}$. Consider the subspace W of the original V given by $V \cap \operatorname{span}(v, \bar{v})$. Show that $OW = W$.

(d) Let the argument (= angle in the complex plane) of λ be θ. Prove that O acts on W by rotation by $\pm\theta$. (To be honest, I didn't think about this one very long. Idk if it's easy or not.)

A bit about fields of definition

Discussing extension of scalars brings me to one more somewhat conceptually tricky topic. The above discussion of Jordan normal form was entirely over $\mathbb{C}$ for the purpose of making sure that a matrix always has an eigenvector. However, the theory works over any field, as long as the eigenvalues of A actually lie in that field.

I bring this up because of a trick that shows up occasionally. Sometimes, you want to show that two matrices A, B over $\mathbb{Q}$ or $\mathbb{R}$ are similar, but for some reason this is hard. However, after passing to a field where all the eigenvalues are defined (e.g. $\mathbb{C}$) it becomes easy to show that the two matrices have the same JNF, and so are similar. A priori, the matrix C such that $A = CBC^{-1}$ could have complex entries (or entries in whatever field you passed to). Is it the case that we can choose C with entries in the field we started with?

The answer is yes, here's one way to think about it. Given A, B matrices over, say, $\mathbb{R}$, we can define a map $L\colon M_n(\mathbb{R}) \to M_n(\mathbb{R})$ via $C \mapsto AC - CB$. This is in fact a linear map on the real vector space $M_n(\mathbb{R})$, and a matrix C so that $A = CBC^{-1}$ would be the same thing as an element in $\ker L \cap \mathrm{GL}_n(\mathbb{R})$. If we know that A and B are similar over $\mathbb{C}$, that's the same as saying that once we extend scalars to consider L as a linear transformation $M_n(\mathbb{C}) \to M_n(\mathbb{C})$, we know that $\ker L \cap \mathrm{GL}_n(\mathbb{C})$ is nonempty. The next exercise asks you to show that at least this implies that $\ker L$ is nonempty over $\mathbb{R}$, and hopefully you can convince yourself that moreover $\ker L \cap \mathrm{GL}_n(\mathbb{R})$ is nonempty.

Exercise 22. Suppose V is a real vector space, and $L\colon V \to V$ is a linear transformation. Let $V_{\mathbb{C}} = V \otimes_{\mathbb{R}} \mathbb{C}$ be the complexification of V, and extend L to a linear map $L_{\mathbb{C}}\colon V_{\mathbb{C}} \to V_{\mathbb{C}}$. Prove that if $\ker L_{\mathbb{C}}$ is nonempty, so too is $\ker L$. (Try to do so without invoking complex conjugation. This fact applies to any extension of fields $K \subset L$ whatsoever, so to prove it in maximum generality you should avoid using anything special about the relationship between $\mathbb{R}$ and $\mathbb{C}$.)

Exercise 23. Let $V, V_{\mathbb{C}}, L, L_{\mathbb{C}}$ be as in the previous problem. Suppose $L_{\mathbb{C}}$ has an eigenvector $v \in V_{\mathbb{C}}$ of eigenvalue $\lambda \in \mathbb{R}$. Prove that L also has an eigenvector in V of eigenvalue λ.

While not a consequence of the above, the next exercise has similarities to the reasoning.

Exercise 24. (January 2019 Problem 5) Let F be a field and n be a positive integer. Fix an $n \times n$ matrix S over F that is invertible and symmetric. Writing A^t for the transpose of a matrix, we let

$$V \coloneqq \{n \times n \text{ matrices } A \text{ over } F : A^t = SAS^{-1}\}.$$

Note that V is a vector space (you do not need to prove this). Find the dimension of V in terms of n.