CONGRUENCE; SYLVESTER'S LAW OF INERTIA

1. Changing basis for quadratic forms

We've seen before that a quadratic form $q : \mathbb{R}^n \to \mathbb{R}$ can be represented by a symmetric matrix A:
$$q(\vec{x}) = \vec{x}^T A \vec{x}.$$
Since we haven't specified which basis of $\mathbb{R}^n$ we're defining A relative to, it's safe to assume that everything in this previous equation — the coordinates used to express $\vec{x}$ and the matrix A itself — is phrased relative to the standard basis. What if we wanted to represent this quadratic form in another coordinate system? What matrix would we use then?

Suppose that $\mathcal{B} = \{\vec{v}_1, \cdots, \vec{v}_n\}$ is another basis for $\mathbb{R}^n$, and let S be the matrix whose columns are the vectos $\vec{v}_i$. Now the value of a quadratic form q at a vector $\vec{x}$ — expressed in standard coordinates — is $[\vec{x}]_{\text{std}}^T A[\vec{x}]_{\text{std}}$ (I've included the "std" notation to emphasize that this is happening relative to standard coordinates). But the standard coordinate representation of $\vec{x}$ and the $\mathcal{B}$-coordinate representation of $\vec{x}$ are linked:
$$S[\vec{x}]_{\mathcal{B}} = [\vec{x}]_{\text{std}}.$$
Hence we get
$$[\vec{x}]_{\text{std}}^T A[\vec{x}]_{\text{std}} = (S[\vec{x}]_{\mathcal{B}})^T A(S[\vec{x}]_{\mathcal{B}}) = [\vec{x}]_{\mathcal{B}}^T S^T AS[\vec{x}]_{\mathcal{B}}.$$
Hence if we want to define our quadratic form relative to $\mathcal{B}$-coordinates, the matrix we would choose instead of A is $S^T AS$. This leads to the following

Definition 1.1. Two symmetric matrices A and B are called congruent if there exists an invertible matrix S such that
$$A = S^T BS.$$

It's important to recognize that the notion of similarity is different from the notion of congruence. Whereas similar matrices can be interpretted as matrix representations for the same *linear operator* on two bases for a space, congruence matrices are different matrix representations for the same *quadratic form* on two bases for a space.

Having defined what it means for matrices to be congruence, it's natural to ask how one can go about determining whether two matrices A and B are congruent. The natural approach — trying to find an invertible matrix S so that $A = S^T BS$ — can be quite difficult to do in practice. In today's class we're going to cover a theorem which allows us to quickly determine whther two matrices are congruent, without ever having to compute the matrix S in the definition of congruence!

To start, let's give a necessary condition that matrices A and B must satisfy in order to be congruent.

Theorem 1.1. *If A and B are congruent, then the rank of A is the same as the rank of B.*

To prove this theorem, we'll first need some results which we could have covered when we first discussed dimensions of images and kernels.

Lemma 1.2. *Suppose that C and D are matrices. Then*

- $\ker(CD) \supseteq \ker(D)$, *with equality if C is invertible; and*
- *im*$(CD) \subseteq$ *im*(C), *with equality if C is invertible.*

Proof. We'll prove the first result. Let $\overrightarrow{x} \in \ker(D)$ be given. Then we have
$$(CD)\overrightarrow{x} = C(D\overrightarrow{x}) = C\overrightarrow{0} = \overrightarrow{0}.$$
Hence $\overrightarrow{x} \in \ker(CD)$. Now if C is invertible, assume that $\overrightarrow{x} \in \ker(CD)$. Then we have $CD\overrightarrow{x} = \overrightarrow{0}$; multiplying both sides by C^{-1} gives
$$D\overrightarrow{x} = C^{-1}CD\overrightarrow{x} = C^{-1}\overrightarrow{0} = \overrightarrow{0}.$$
Hence $\overrightarrow{x} \in \ker(D)$, as desired.

The proof of the second result is similar. □

Now we're ready to prove our theorem

Proof that congruent matrices have the same rank. Suppose that A and B are congruent, so that $A = S^TBS$. Notice that since S is invertible, so too is S^T. If we apply the first part of the previous lemma to the matrix BS, then since S is invertible we know that $\mathrm{im}(B) = \mathrm{im}(BS)$. Hence we also have the rank of B and the rank of BS are the same. Likewise, since S^T is invertible, we know that $\ker(S^T(BS)) = \ker(BS)$. Applying the rank nullity theorem and the equality between images we've already shown, we therefore have
$$\mathrm{rk}(A) = \mathrm{rk}(S^TBS) = n - \dim(\ker(S^TBS)) = n - \dim(\ker(BS)) = n - \dim(\ker(B)) = \mathrm{rk}(B).$$
□

2. Sufficient conditions

We now see that in order for A and B to be congruent, it is necessary that their ranks are identical. Is this enough? No!

Example. Consider the matrices $A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} -1 & 0 \\ 0 & 0 \end{pmatrix}$. The quadratic form defined by A is takes on positive values, while the quadratic form defined by B does not. Since A and B take different values, they cannot be congruent.

Clearly there is more to the picture than just rank. The obstruction to congruence in the previous problem had to do with positivity and negativity, so perhaps this is enough to determine congruence. Inspired by, we give the following

Definition 2.1. Suppose that D is a diagonal matrix. We say that the index of D is the number of positive entries in D, and the signature of D is the the number of positive entries of D minus the number of negative entries in D.

If we let r denote the rank of a diagonal matrix D (i.e., the total number of nonzero entries in D), p the index and s the signature, then it's not hard to see
$$s = r - p.$$
Likewise we can solve for any of the two quantities from $\{r, p, s\}$ in terms of the other two. Surprisingly, these invariants are enough to determine when matrices are congruent!

Theorem 2.1 (Sylvester's Law of Inertia)**.** *Two symmetric $n \times n$ matrices B and C are congruent if and only if the diagonal representations for B and C have the same rank, index and signature.*

This theorem is extremely powerful: it tells us we can determine whether two symmetric matrices are congruent just by computing their eigenvalues!

Example. Determine which of the following matrices are congruent to each other:

$$A = \begin{pmatrix} 1 & -1 & 3 \\ -1 & 2 & 1 \\ 3 & 1 & 1 \end{pmatrix}, B = \begin{pmatrix} 1 & 2 & 1 \\ 2 & 3 & 2 \\ 1 & 2 & 1 \end{pmatrix}, C = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 2 \\ 1 & 2 & 1 \end{pmatrix}.$$

(Note: in class we had a different (non-symmetric) matrix for A. This is the matrix I should have written down.)

To answer this question, we know we need to compute the rank, index and signature for these matrices. If one sits down and does the calculations, then one finds the following:

	Eigenvalues	**rank**	**index**	**signature**
A	$4, \pm\sqrt{6}$	3	2	1
B	$0, \frac{5\pm\sqrt{33}}{2}$	2	1	0
C	$1, 1\pm\sqrt{5}$	3	2	1

Since A and C have the same rank and index (and signature), they are congruent. Since B has a different rank from either A or C, they aren't congruent.

We'll spend the balance of class proving Sylvester's Law of Inertia.

3. Proof of Sylvester's Law of Inertia

Proof. Suppose that B and C are congruent, so they represent the same quadratic form (in different bases). Now let $\{\vec{v}_1, \cdots, \vec{v}_n\}$ be the eigenbasis which diagonalizes B, and let a_i be the eigenvalue which corresponds to $\vec{v}_i$. Likewise, let $\{\vec{w}_1, \cdots, \vec{w}_n\}$ be the eigenbasis diagonalizing C, with e_i the eigenvalue corresponding to $\vec{w}_i$. We'll assume that the bases are arranged so that all positive eigenvalues come first, all negative eigenvalues second, and all 0 eigenvalues last.

We already know that B and C have the same rank, so to reach a contradition we'll assume they have different index. Let p be the index for B, and q the index for C. Just to recap, this means that the diagonalizing basis for B has

$$\{\underbrace{\vec{v}_1, \cdots, \vec{v}_p}_{\text{pos e-values}}, \underbrace{\vec{v}_{p+1}, \cdots, \vec{v}_r}_{\text{neg e-values}}, \underbrace{\vec{v}_{r+1}, \cdots, \vec{v}_n}_{\text{zero e-values}}\}$$

and, likewise, that

$$\{\underbrace{\vec{w}_1, \cdots, \vec{w}_q}_{\text{pos e-values}}, \underbrace{\vec{w}_{q+1}, \cdots, \vec{w}_r}_{\text{neg e-values}}, \underbrace{\vec{w}_{r+1}, \cdots, \vec{w}_n}_{\text{zero e-values}}\}.$$

With the bookkeeping out of the way, we'll now define a linear operator on $\mathbb{R}^n$ as follows:

$$L(\overrightarrow{x}) = \begin{pmatrix} \overrightarrow{v}_1^T B\overrightarrow{x} \\ \overrightarrow{v}_2^T B\overrightarrow{x} \\ \vdots \\ \overrightarrow{v}_p^T B\overrightarrow{x} \\ \overrightarrow{w}_{q+1}^T C\overrightarrow{x} \\ \vdots \\ \overrightarrow{w}_r^T C\overrightarrow{x} \end{pmatrix}.$$

Notice that this is a map from $\mathbb{R}^n \to \mathbb{R}^{p+(r-q)}$, and hence rank nullity says that

$$\dim(\ker(L)) = n - \mathrm{rk}(L) \geq n - (p + r - q) > n - r.$$

This means that there exists some vector $\overrightarrow{v}_0$ in the kernel of L so that

$$\overrightarrow{v}_0 \notin \mathrm{span}\{\overrightarrow{v}_{r+1}, \cdots, \overrightarrow{v}_n\} \text{ and } \overrightarrow{v}_0 \notin \mathrm{span}\{\overrightarrow{w}_{r+1}, \cdots, \overrightarrow{w}_n\}.$$

Let's see what else we can say about this magical vector $\overrightarrow{v}_0$. Certainly $\overrightarrow{v}_0$ can be expressed in terms of the two bases we have in hand:

$$\overrightarrow{v}_0 = \sum_{i=1}^{n} c_i \overrightarrow{v}_i = \sum_{i=1}^{n} d_i \overrightarrow{w}_i.$$

We just argued that at least one of the coefficients c_i and d_i must be nonzero for $i \leq r$. Can we say more? Notice that we have

$$\overrightarrow{v}_k^T B\overrightarrow{v}_0 = \overrightarrow{v}_k^T B(\sum_{i=1}^{n} c_i \overrightarrow{v}_i) = \sum_{i=1}^{n} c_i \overrightarrow{v}_k^T B\overrightarrow{v}_i.$$

But since $\overrightarrow{v}_i$ is an eigenvector for B with eigenvalue a_i, this simplifies to

$$\sum_{i=1}^{n} c_i a_i \overrightarrow{v}_k^T \overrightarrow{v}_i.$$

Because the $\overrightarrow{v}_i$ form an orthonormal basis, we know this dot product is 0, unless $k = i$ when it is 1. Hence we get

$$\overrightarrow{v}_k^T B\overrightarrow{v}_0 = \sum_{i=1}^{n} c_i a_i \overrightarrow{v}_k^T \overrightarrow{v}_i = c_k a_k.$$

If we choose k to be between 1 and p, then since $\overrightarrow{v}_k^T B\overrightarrow{v}_0$ represents the kth coordinate of the output vector $L\overrightarrow{v}_0$ — and since $L(\overrightarrow{v}_0) = \overrightarrow{0}$ since $\overrightarrow{v}_0 \in \ker(L)$ — we in fact have

$$\overrightarrow{v}_k^T V\overrightarrow{v}_0 = c_k a_k = 0.$$

Since $a_k > 0$ for these values of k, the only way this equality can hold is if $c_k = 0$. Hence we have that

$$\overrightarrow{v}_0 = \sum_{i=p+1}^{n} c_i \overrightarrow{v}_i.$$

One can run an analogous argument to show that

$$\overrightarrow{v}_0 = \sum_{\substack{1 \leq i \leq q \\ r+1 \leq i \leq n}} d_i \overrightarrow{w}_i.$$

What does all this buy for us? We'll use the expressions for $\overrightarrow{v}_0$ in the different coordinate systems to compute the value of the underlying quadratic form at $\overrightarrow{v}_0$. First, we compute it's value relative to the diagonalizing eigenbasis for B:

$$\overrightarrow{v}_0^T B \overrightarrow{v}_0 = \left(\sum_{i=p+1}^{n} c_i \overrightarrow{v}_i\right)^T B \left(\sum_{j=p+1}^{n} c_j \overrightarrow{v}_j\right) = \sum_{p+1 \leq i,j \leq n} c_i c_j \overrightarrow{v}_i^T B \overrightarrow{v}_j.$$

As before, we know that the term $\overrightarrow{v}_i^T B \overrightarrow{v}_j$ is 0 unless $i = j$, in which case the value is instead a_i. Hence we get

$$\overrightarrow{v}_0^T B \overrightarrow{v}_0 = \sum_{p+1 \leq i \leq n} c_i^2 a_i.$$

Now the a_i under consideration are all non-positive, and at least one of the $c_i \neq 0$ for $p+1 \leq i \leq r$ (we worked to show this when we first defined $\overrightarrow{v}_0$). Hence we know that the value of the quadratic form at $\overrightarrow{v}_0$ is negative.

On the other hand, we can compute the value of the quadratic form at $\overrightarrow{v}_0$ relative to the diagonalizing eigenbasis for C:

$$\overrightarrow{v}_0^T B \overrightarrow{v}_0 = \left(\sum_{\substack{1 \leq i \leq q \\ r+1 \leq i \leq n}} d_i \overrightarrow{w}_i\right)^T C \left(\sum_{\substack{1 \leq j \leq q \\ r+1 \leq j \leq n}} d_j \overrightarrow{w}_j\right) = \sum_{\substack{1 \leq i,j \leq q \\ r+1 \leq i,j \leq n}} d_i d_j \overrightarrow{w}_i^T C \overrightarrow{w}_j.$$

As before, we know that the term $\overrightarrow{w}_i^T C \overrightarrow{w}_j$ is 0 unless $i = j$, in which case the value is instead e_i. Hence we get

$$\overrightarrow{v}_0^T C \overrightarrow{v}_0 = \sum_{\substack{1 \leq i \leq q \\ r+1 \leq i \leq n}} c_i^2 e_i.$$

Now the e_i under consideration are all non-negative, and at least one of the $d_i \neq 0$ for $1 \leq i \leq q$ (we worked to show this when we first defined $\overrightarrow{v}_0$). Hence we know that the value of the quadratic form at $\overrightarrow{v}_0$ is positive.

This should set off alarm bells – the quadratic form can't be both positive and negative at $\overrightarrow{v}_0$. Hence we have a contradiction, and our assumption that the index of B and C are different must be false. □

Since rank and index are sufficient to describe when two matrices are congruent, it's natural to pick one "canonical" matrix for each congruence class. That is to say, to find a set of matrices so that any given symmetric matrix is congruent to exactly *one* of the matrices in our set. For us, these "canonical forms" will be given by matrices we'll call $J_{r,p}$. Specifically, if an $n \times n$ diagonal matrix is going to have rank r and index p, then we could choose the positive entries of the matrix to just be 1. The $r - p$ negative entries could then be selected as -1, and we'd have the matrix

$$J_{r,p} = \begin{pmatrix} I_p & & & & \\ & -I_{r-p} & & & \\ & & 0 & & \\ & & & \ddots & \\ & & & & 0 \end{pmatrix}.$$