Linear Algebra (part ∞) : Bilinear and Quadratic Forms (by Evan Dummit, 2019, v. 0.01)

Contents

0.1 Bilinear and Quadratic Forms

- We discuss bilinear and quadratic forms.
- Let V be a vector space over the field F.

0.1.1 Bilinear Forms

- Definition: A function $\Phi : V \times V \to F$ is a bilinear form on V if it is linear in each variable when the other variable is fixed. Explicitly, this means $\Phi(\mathbf{v}_1 + \alpha\mathbf{v}_2, y) = \Phi(\mathbf{v}_1, \mathbf{w}) + \alpha\Phi(\mathbf{v}_2, \mathbf{w})$ and $\Phi(\mathbf{v}, \mathbf{w}_1 + \alpha\mathbf{w}_2) = \Phi(\mathbf{v}, \mathbf{w}_1) + \alpha\Phi(\mathbf{v}, \mathbf{w}_2)$ for arbitrary $\mathbf{v}_i, \mathbf{w}_i \in V$ and $\alpha \in F$.
  - It is easy to see that the set of all bilinear forms on V forms a vector space (under componentwise addition and scalar multiplication).
  - An inner product on a real vector space is a bilinear form, but an inner product on a complex vector space is not, since it is conjugate-linear in the second component.
  - For example, if $V = L^2[a, b]$ is the space of (real-valued) square-integrable functions on $[a, b]$, then $\Phi(f, g) = \int_a^b f(x)g(x)\,dx$ is a bilinear form on V.
  - If there exists a nonzero vector $\mathbf{x} \in V$ such that $\Phi(\mathbf{x}, \mathbf{v}) = 0$ for all $\mathbf{v} \in V$, we say Φ is degenerate, and otherwise (if there is no such $\mathbf{x}$) we say Φ is nondegenerate.
- A large class of examples of bilinear forms arise as follows: if $V = F^n$, then for any matrix $A \in M_{n\times n}(F)$, the map $\Phi_A(\mathbf{v}, \mathbf{w}) = \mathbf{v}^T A\mathbf{w}$ is a bilinear form on V.
  - Example: The matrix $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ yields the bilinear form $\Phi_A\left(\begin{bmatrix} x_1 \\ y_1 \end{bmatrix}, \begin{bmatrix} x_2 \\ y_2 \end{bmatrix}\right) = [x_1\ y_1]\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}\begin{bmatrix} x_2 \\ y_2 \end{bmatrix} = x_1x_2 + 2x_1y_2 + 3x_2y_1 + 4y_1y_2$.
- Indeed, if V is finite-dimensional, then by choosing a basis of V we can see that every bilinear form arises in this manner, as follows:
- Definition: If V is a finite-dimensional vector space, $\beta = \{\beta_1, \dots, \beta_n\}$ is a basis of V, and Φ is a bilinear form on V, the associated matrix of Φ with respect to β is the matrix $[\Phi]_\beta \in M_{n\times n}(F)$ whose (i, j)-entry is the value $\Phi(\beta_i, \beta_j)$.
- Proposition (Associated Matrices): Suppose that V is a finite-dimensional vector space, $\beta = \{\beta_1, \dots, \beta_n\}$ is a basis of V, and Φ is a bilinear form on V.

1. If $\mathbf{v}$ and $\mathbf{w}$ are any vectors in V, then $\Phi(\mathbf{v},\mathbf{w}) = [\mathbf{v}]_\beta^T [\Phi]_\beta [\mathbf{w}]_\beta$.[1]
   - Proof: If $\mathbf{v} = \beta_i$ and $\mathbf{w} = \beta_j$ then the result follows immediately from the definition of matrix multiplication and the matrix $[\Phi]_\beta$. The result for arbitrary $\mathbf{v}$ and $\mathbf{w}$ then follows by linearity.
2. The association $\Phi \mapsto [\Phi]_\beta$ of a bilinear form with its matrix representation yields an isomorphism of the space $\mathcal{B}(V)$ of bilinear forms on V with $M_{n\times n}(F)$. In particular, $\dim_F \mathcal{B}(V) = n^2$.
   - Proof: The inverse map is defined as follows: given a matrix $A \in M_{n\times n}(F)$, define a bilinear form Φ_A via $\Phi_A(\mathbf{v},\mathbf{w}) = [\mathbf{v}]_\beta^T A [\mathbf{w}]_\beta$.
   - It is easy to verify that this map is a well-defined linear transformation and that it is inverse to the map given above. The dimension calculation is immediate.
3. If Φ^T is the "reverse form" of Φ defined via $\Phi^T(\mathbf{v},\mathbf{w}) = \Phi(\mathbf{w},\mathbf{v})$, then $[\Phi^T]_\beta = [\Phi]_\beta^T$.
   - Proof: By definition we have $\Phi^T(\mathbf{v},\mathbf{w}) = [\mathbf{w}]_\beta^T [\Phi]_\beta [\mathbf{v}]_\beta$. Since the matrix product on the right is a scalar, it is equal to its transpose, which is $[\mathbf{v}]_\beta^T [\Phi]_\beta^T [\mathbf{w}]_\beta$.
   - This means $[\Phi^T]_\beta$ and $[\Phi]_\beta^T$ agree, as bilinear forms, on all pairs of vectors $[\mathbf{v}]_\beta$ and $[\mathbf{w}]_\beta$ in F^n, so they are equal.
4. If γ is another basis of V and $Q = [I]_\beta^\gamma$ is the change-of-basis matrix from β to γ, then $[\Phi]_\gamma = Q^T [\Phi]_\beta Q$.
   - Proof: By definition, $[\mathbf{v}]_\gamma = Q[\mathbf{v}]_\beta$. Hence $[\mathbf{v}]_\beta^T Q^T [\Phi]_\beta Q [\mathbf{w}]_\beta = [\mathbf{v}]_\gamma^T [\Phi]_\beta [\mathbf{w}]_\gamma$.
   - This means that $Q^T [\Phi]_\beta Q$ and $[\Phi]_\gamma$ agree, as bilinear forms, on all pairs of vectors $[\mathbf{v}]_\beta$ and $[\mathbf{w}]_\beta$ in F^n, so they are equal.

- The last result of the proposition above tells us how bilinear forms behave under change of basis: rather than the more familiar conjugation relation $B = QAQ^{-1}$, we instead have a slightly different relation $B = Q^T AQ$.

0.1.2 Symmetric Bilinear Forms and Diagonalization

- In the same way that we classified the linear operators on a vector space that can be diagonalized, we would also like to classify the diagonalizable bilinear forms.
- Definition: If V is finite-dimensional, a bilinear form Φ on V is diagonalizable if there exists a basis β of V such that $[\Phi]_\beta$ is a diagonal matrix.
  - The matrix formulation of this question is as follows: we say that matrices B and C are congruent if there exists an invertible matrix Q such that $C = Q^T BQ$.
  - Then the matrices B and C are congruent if and only if they represent the same bilinear form in different bases (the translation being $B = [\Phi]_\beta$ and $C = [\Phi]_\gamma$, with $Q = [I]_\beta^\gamma$ being the corresponding change-of-basis matrix).
- It turns out that when $\text{char}(F) \neq 2$, there is an easy criterion for diagonalizability.
- Definition: A bilinear form Φ on V is symmetric if $\Phi(\mathbf{v},\mathbf{w}) = \Phi(\mathbf{w},\mathbf{v})$ for all $\mathbf{v},\mathbf{w} \in V$.
  - Notice that Φ is symmetric if and only if it equals its reverse form Φ^T.
  - By taking associated matrices, we see immediately that if V is finite-dimensional with basis β, then Φ is a symmetric bilinear form if and only if $[\Phi]_\beta$ is equal to its transpose, which is to say, when it is a symmetric matrix.
  - Now observe that if Φ is diagonalizable, then $[\Phi]_\beta$ is a diagonal matrix hence symmetric, and thus Φ must be symmetric.
- When the characteristic of F is not equal to 2, the converse holds also:
- Theorem (Diagonalization of Bilinear Forms): Let V be a finite-dimensional vector space over a field F of characteristic not equal to 2. Then a bilinear form on V is diagonalizable if and only if it is symmetric.

[1] Recall that $[\mathbf{v}]_\beta$ is the coefficient vector in F^n for $\mathbf{v}$ when it is expressed in terms of the basis β.

- Proof: The forward direction was established above. For the reverse, we show the result by induction on $n = \dim_F V$. The base case $n = 1$ is trivial, so suppose the result holds for all spaces of dimension less than n, and let Φ be symmetric on V.
- If Φ is the zero form, then clearly Φ is diagonalizable. Otherwise, suppose Φ is not identically zero: we claim there exists a vector $\mathbf{x}$ with $\Phi(\mathbf{x}, \mathbf{x}) \neq 0$.
- By hypothesis, Φ is not identically zero so suppose that $\Phi(\mathbf{v}, \mathbf{w}) \neq 0$. If $\Phi(\mathbf{v}, \mathbf{v}) \neq 0$ or $\Phi(\mathbf{w}, \mathbf{w}) \neq 0$ we may take $\mathbf{x} = \mathbf{v}$ or $\mathbf{x} = \mathbf{w}$. Otherwise, we have $\Phi(\mathbf{v} + \mathbf{w}, \mathbf{v} + \mathbf{w}) = \Phi(\mathbf{v}, \mathbf{v}) + 2\Phi(\mathbf{v}, \mathbf{w}) + \Phi(\mathbf{w}, \mathbf{w}) = 2\Phi(\mathbf{v}, \mathbf{w}) \neq 0$ by the assumption that $\Phi(\mathbf{v}, \mathbf{w}) \neq 0$ and $2 \neq 0$ in F (here is where we require the characteristic not to equal 2), and so we may take $\mathbf{x} = \mathbf{v} + \mathbf{w}$.
- Now consider the linear functional $T : V \to F$ given by $T(\mathbf{v}) = \Phi(\mathbf{x}, \mathbf{v})$. Since $T(\mathbf{x}) = \Phi(\mathbf{x}, \mathbf{x}) \neq 0$, we see that $\text{im}(T) = F$, so $\dim_F \ker(T) = n - 1$.
- Then the restriction of Φ to $\ker(T)$ is clearly a symmetric bilinear form on $\ker(T)$, so by induction, there exists a basis $\{\beta_1, \ldots, \beta_{n-1}\}$ of $\ker(T)$ such that the restriction of Φ is diagonalized by this basis, which is to say, $\Phi(\beta_i, \beta_j) = 0$ for $i \neq j$.
- Now set $\beta_n = \mathbf{x}$ and observe that since $\mathbf{x} \notin \ker(T)$, the set $\beta = \{\beta_1, \ldots, \beta_{n-1}, \beta_n\}$ is a basis of V. Since $\Phi(\mathbf{x}, \beta_i) = \Phi(\beta_i, \mathbf{x}) = 0$ for all $i < n$ by definition of T, we conclude that β diagonalizes Φ, as required.

- We will note that the assumption that $\text{char}(F) \neq 2$ in the theorem above cannot be removed.
  - Explicitly, if $F = \mathbb{F}_2$ is the field with 2 elements, then if Φ is the bilinear form on F^2 with associated matrix $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, then Φ is symmetric but cannot be diagonalized.
  - Explicitly, suppose $Q = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$: then $Q^T A Q = \begin{bmatrix} a & c \\ b & d \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 0 & ad + bc \\ ad + bc & 0 \end{bmatrix}$, so the only possible diagonalization of Φ would be the zero matrix, but that is impossible because Φ is not the zero form.
  - In this example we can see that[2] $\Phi(\mathbf{x}, \mathbf{x}) = 0$ for all $\mathbf{x} \in F^2$, which causes the inductive argument to fail.

- As an immediate corollary, we see that every symmetric matrix is congruent to a symmetric matrix in characteristic $\neq 2$:
- Corollary: If $\text{char}(F) \neq 2$, then every symmetric matrix over F is congruent to a diagonal matrix.
  - Proof: The result follows immediately by diagonalizing the corresponding bilinear form.
- We can give an explicit procedure for writing a symmetric matrix S in the form $D = Q^T S Q$ that is similar to the algorithm for computing the inverse of a matrix.
  - Recall that if E is an elementary row matrix (obtained by performing an elementary row operation on the identity matrix), then EA is the matrix obtained by performing that elementary row operation on A.
  - Likewise, if C is an elementary column matrix, then AC is the matrix obtained by performing that elementary column operation on A.
  - Hence if E is an elementary row matrix, then EAE^T is the matrix obtained by performing the elementary row operation on A (given by E) and then the corresponding elementary column operation (given by E^T).
  - Since the invertible matrix Q is a product $E_1 \cdots E_d$ of elementary row matrices, we see that $Q^T S Q = E_d^T \cdots E_1^T S E_1 \cdots E_d$ is obtained from S by performing a sequence of these paired row-column operations.
  - Our result on diagonalization above ensures that there is a sequence of these operations that will yield a diagonal matrix.

[2] This follows by noting that $\Phi(\beta_i, \beta_i) = 0$ for each basis element β_i. Then if $\mathbf{v} = a_1\beta_1 + \cdots + a_n\beta_n$, expanding $\Phi(\mathbf{v}, \mathbf{v})$ linearly and applying symmetry shows that every term $a_i a_j \Phi(\beta_i, \beta_j)$ for $i \neq j$ has a coefficient of 2, so $\Phi(\mathbf{v}, \mathbf{v}) = 0$ for all $\mathbf{v}$.

- We may find the proper sequence of operations by performing these "paired" operations using a method similar to row-reduction: using the (1,1)-entry, we apply row operations to clear out all the entries in the first column below it. (If this entry is zero, we add an appropriate multiple of another row to the top row to make it nonzero.)
- This will also clear out the column entries to the right of the (1,1)-entry, yielding a matrix whose first row and column are now diagonalized. We then restrict attention to the smaller $(n-1)\times(n-1)$ matrix excluding the first row and column, and repeat the procedure recursively until the matrix is diagonalized.
- Then we may obtain the matrix $Q^T = E_d^T \cdots E_1^T I$ by applying all of the elementary row operations (in the same order) starting with the identity matrix.

- We may keep track of these operations using a "double matrix" as in the algorithm for computing the inverse of a matrix: on the left we start with the symmetric matrix S, and on the right we start with the identity matrix I.
  - At each step, we select a row operation and perform it, and its corresponding column operation, on the left matrix. We also perform the row operation (but only the row operation) on the right matrix.
  - When we are finished, we will have transformed the double-matrix $[S|I]$ into the double-matrix $[D|Q^T]$, and we will have $Q^T SQ = D$.

- Example: For $S = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 0 \\ 3 & 0 & 2 \end{bmatrix}$, find an invertible matrix Q and diagonal matrix D such that $Q^T SQ = D$.
  - We set up the double matrix and perform row/column operations as listed (to emphasize again, we perform the row and then the corresponding column operation on the left side, but only the row operation on the right side):

$$\left[\begin{array}{ccc|ccc} 1 & 2 & 3 & 1 & 0 & 0 \\ 2 & 1 & 0 & 0 & 1 & 0 \\ 3 & 0 & 2 & 0 & 0 & 1 \end{array}\right] \underset{C_2-2C_1}{\overset{R_2-2R_1}{\longrightarrow}} \left[\begin{array}{ccc|ccc} 1 & 0 & 3 & 1 & 0 & 0 \\ 0 & -3 & -6 & -2 & 1 & 0 \\ 3 & -6 & 2 & 0 & 0 & 1 \end{array}\right]$$

$$\underset{C_3-3C_1}{\overset{R_3-3R_1}{\longrightarrow}} \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & -3 & -6 & -2 & 1 & 0 \\ 0 & -6 & -7 & -3 & 0 & 1 \end{array}\right] \underset{C_3-2C_2}{\overset{R_3-2R_2}{\longrightarrow}} \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & -3 & 0 & -2 & 1 & 0 \\ 0 & 0 & 5 & 1 & -2 & 1 \end{array}\right]$$

  - The matrix on the left is now diagonal.
  - Thus, we may take $D = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -3 & 0 \\ 0 & 0 & 5 \end{bmatrix}$ with $Q^T = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 1 & -2 & 1 \end{bmatrix}$ and thus $Q = \begin{bmatrix} 1 & -2 & 1 \\ 0 & 1 & -2 \\ 0 & 0 & 1 \end{bmatrix}$. Indeed, one may double-check that $Q^T SQ = D$, as claimed.

0.1.3 Quadratic Forms

- In the proof that symmetric forms are diagonalizable, the existence of a vector $\mathbf{x} \in V$ such that $\Phi(\mathbf{x}, \mathbf{x}) \neq 0$ played a central role. We now examine this (non-linear!) function $\Phi(\mathbf{x}, \mathbf{x})$ more closely.
- Definition: If Φ is a symmetric bilinear form on V, the function $Q : V \to F$ given by $Q(\mathbf{v}) = \Phi(\mathbf{v}, \mathbf{v})$ is called the quadratic form associated to Φ.
  - Example: If Φ is the symmetric bilinear form with matrix $A = \begin{bmatrix} 1 & 3 \\ 3 & 4 \end{bmatrix}$ over F^2, then the corresponding quadratic form has $Q(\begin{bmatrix} x \\ y \end{bmatrix}) = x^2 + 6xy + 4y^2$. (The fact that this is a homogeneous quadratic function of the entries of the input vector is the reason for the name "quadratic form".)
- Clearly, Q is uniquely determined by Φ. When $\text{char}(F) \neq 2$, the reverse holds as well.

- Explicitly, since $Q(\mathbf{v}+\mathbf{w}) = \Phi(\mathbf{v}+\mathbf{w},\mathbf{v}+\mathbf{w}) = Q(\mathbf{v}) + 2\Phi(\mathbf{v},\mathbf{w}) + Q(\mathbf{w})$, we can write $\Phi(\mathbf{v},\mathbf{w}) = \frac{1}{2}\left[Q(\mathbf{v}+\mathbf{w}) - Q(\mathbf{v}) - Q(\mathbf{w})\right]$, and so we may recover Φ from Q.
- Also, observe that for any scalar $\alpha \in F$, we have $Q(\alpha\mathbf{v}) = \Phi(\alpha\mathbf{v},\alpha\mathbf{v}) = \alpha^2\Phi(\mathbf{v},\mathbf{v}) = \alpha^2 Q(\mathbf{v})$.
- This last two relations provide us a way to define a quadratic form without explicit reference to the underlying symmetric bilinear form.

- Definition: If V is a vector space, a quadratic form is a function $Q : V \to F$ such that $Q(\alpha\mathbf{v}) = \alpha^2 Q(\mathbf{v})$ for all $\alpha \in F$, and the function $Q(\mathbf{v}+\mathbf{w}) - Q(\mathbf{v}) - Q(\mathbf{w})$ is linear in $\mathbf{v}$ and $\mathbf{w}$.
  - By setting $\alpha = 0$ we see $Q(\mathbf{0}) = 0$, and by setting $\alpha = -1$ we see $Q(-\mathbf{v}) = Q(\mathbf{v})$.
  - Like with bilinear forms, the set of all quadratic forms on V forms a vector space.
  - Example (again): The function $Q(\begin{bmatrix} x \\ y \end{bmatrix}) = x^2 + 6xy + 4y^2$ is a quadratic form on F^2.
  - Example: If $V = L^2[a,b]$ is the space of square-integrable (real-valued) functions on the interval $[a,b]$, then the function $Q(f) = \int_a^b f(x)^2\,dx$ is a quadratic form on V.

- If $\mathrm{char}(F) \neq 2$, then the function $\frac{1}{2}[Q(\mathbf{v}+\mathbf{w}) - Q(\mathbf{v}) - Q(\mathbf{w})]$ is the bilinear pairing associated to Q. It is easy to see that we obtain a correspondence between quadratic forms and bilinear pairings in this case.
  - In particular, any homogeneous quadratic function on F^n (i.e., any polynomial function all of whose terms have total degree 2) is a quadratic form on F^n: for variables $x_1, \dots, x_n$, such a function has the general form $\sum_{1 \le i \le j \le n} a_{i,j} x_i x_j$. [3]
  - Then we can see that the associated matrix A for the corresponding bilinear form has entries $a_{i,j} = a_{j,i} = \begin{cases} a_{i,i} & \text{for } i = j \\ a_{i,j}/2 & \text{for } i \neq j \end{cases}$; this of course requires $\mathrm{char}(F) \neq 2$ in order to be able to divide by 2.
  - Example: The function $Q(x_1, x_2, x_3) = x_1^2 + 2x_1x_3 - 3x_2x_3 + 4x_3^2$ is a quadratic form on F^3. When $\mathrm{char}(F) \neq 2$ the matrix for the associated symmetric bilinear form is $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & -\frac{3}{2} \\ 1 & -\frac{3}{2} & 4 \end{bmatrix}$.
  - Example: The function $Q(x_1, \dots, x_n) = x_1^2 + 2x_2^2 + 3x_3^2 + \cdots + nx_n^2$ is a quadratic form on F^n. Its associated matrix is the diagonal matrix with entries $1, 2, \dots, n$.

0.1.4 Quadratic Forms Over $\mathbb{R}^n$

- In the event that V is a finite-dimensional vector space over $F = \mathbb{R}$, quadratic forms are particularly pleasant. By choosing a basis we may assume that $V = \mathbb{R}^n$ for concreteness.
  - Then, per the real spectral theorem, any real symmetric matrix is orthogonally diagonalizable, meaning that if S is any real symmetric matrix, then there exists an orthogonal matrix Q (with $Q^T = Q^{-1}$) such that $QSQ^{-1} = D$ is diagonal.
  - But since $Q^T = Q^{-1}$, if we take $R = Q^T$ then this condition is the same as saying $R^T S R = D$ is diagonal. This is precisely the condition we require in order to diagonalize a symmetric bilinear form.
  - Hence: we may diagonalize a symmetric bilinear form over $\mathbb{R}$ by computing the (regular) diagonalization of the corresponding matrix: this is quite efficient as it only requires finding the eigenvalues and eigenvectors.
  - The corresponding diagonalization represents "completing the square" in the quadratic form via a change of variables that is orthogonal (i.e., arises from an orthonormal basis), which corresponds geometrically to a rotation of the standard coordinate axes.

[3] It can also be verified directly from the definition that this is a quadratic form via some mild calculations; this also shows the statement is true even when $\mathrm{char}(F) = 2$.

- Example: Find an orthogonal change of basis that diagonalizes the quadratic form $Q(x, y, z) = 5x^2 + 4xy + 6y^2 + 4yz + 7z^2$ over $\mathbb{R}^3$.
  - We simply diagonalize the matrix for the corresponding bilinear form, which is $A = \begin{bmatrix} 5 & 2 & 0 \\ 2 & 6 & 2 \\ 0 & 2 & 7 \end{bmatrix}$.
  - A short calculation produces the eigenvalues $\lambda = 3, 6, 9$ with corresponding eigenvectors $\begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}$, $\begin{bmatrix} -2 \\ -1 \\ 2 \end{bmatrix}$, $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$.
  - Hence we may take $Q = \frac{1}{3}\begin{bmatrix} 2 & -2 & 1 \\ -2 & -1 & 2 \\ 1 & 2 & 2 \end{bmatrix}$, and then we have $Q^T = Q^{-1}$ and $QAQ^{-1} = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & 9 \end{bmatrix} = D$.
  - Therefore the desired change of basis is $x' = \frac{1}{3}(2x - 2y + z)$, $y' = \frac{1}{3}(-2x - y + 2z)$, $z' = \frac{1}{3}(x + 2y + 2z)$, and with this change of basis it is not hard to verify that, indeed, $Q(x, y, z) = 3(x')^2 + 6(y')^2 + 9(z')^2$.
- One application of the existence of such a diagonalization is to classify the conics in $\mathbb{R}^2$ and the quadric surfaces in $\mathbb{R}^3$.
  - For conics in $\mathbb{R}^2$, the general equation is $Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0$. By diagonalizing, we may eliminate the xy term, and so the quadratic term can be put into the form $Ax^2 + Cy^2$. We then have various cases depending on the signs of A and C.
  - If A and C are both zero then the conic degenerates to a line. If one is zero and the other is not, then by rescaling and swapping variables we may assume $A = 1$ and $C = 0$, in which case the equation $x^2 + Dx + Ey + F = 0$ yields a parabola upon solving for y.
  - If both A, C are nonzero, then we may complete the square to eliminate the linear terms, and then rescale so that $F = -1$. The resulting equation then has the form $A'x^2 + C'y^2 = 1$. If A', C' have the same sign, then the curve is an ellipse, while if A', C' have the opposite sign, the curve is a hyperbola.
  - For quadric surfaces in $\mathbb{R}^3$ we may likewise eliminate cross-terms by diagonalizing, which yields a reduced equation $Ax^2 + By^2 + Cz^2 + Dx + Ey + Fz + G = 0$.
  - We can then perform a similar analysis (based on how many of A, B, C are zero and the relative signs of the nonzero coefficients and the linear terms) to obtain all of the possible quadric surfaces in $\mathbb{R}^3$.
  - In addition to the "degenerate" surfaces (e.g., a point, a plane, two planes), after rescaling the variables, one obtains 9 different quadric surfaces: the ellipsoid (e.g., $x^2 + y^2 + z^2 = 1$), the elliptic, parabolic, and hyperbolic cylinders (e.g., $x^2 + y^2 = 1$, $y = x^2$, and $x^2 - y^2 = 1$), the hyperboloid of one sheet (e.g., $z^2 - x^2 - y^2 = 1$), the elliptical cone (e.g., $z^2 = x^2 + y^2$), the hyperboloid of two sheets (e.g., $x^2 + y^2 - z^2 = 1$), the elliptic paraboloid (e.g., $z = x^2 + y^2$), and the hyperbolic paraboloid (e.g., $z = x^2 - y^2$).
  - Seven of the quadric surfaces are plotted in Figure 1 (the parabolic and hyperbolic cylinders are omitted).
- We can also use quadratic forms to prove the famous "second derivatives test" from multivariable calculus:
- Theorem (Second Derivatives Test in $\mathbb{R}^n$): Suppose f is a function of n variables $x_1, \dots, x_n$ that is twice-differentiable and P is a critical point of f, so that $f_{x_i}(P) = 0$ for each i. Let H be the Hessian matrix, whose (i, j)-entry is the second-order partial derivative $f_{x_i x_j}(P)$. If all eigenvalues of H are positive then f has a local minimum at P, if all eigenvalues of H are negative then f has a local maximum at P, if H has at least one eigenvalue of each sign then f has a saddle point at P, and in all other cases (where H has at least one zero eigenvalue and does not have one of each sign) the test is inconclusive.
  - Proof (outline): By translating appropriately, assume for simplicity that P is at the origin.

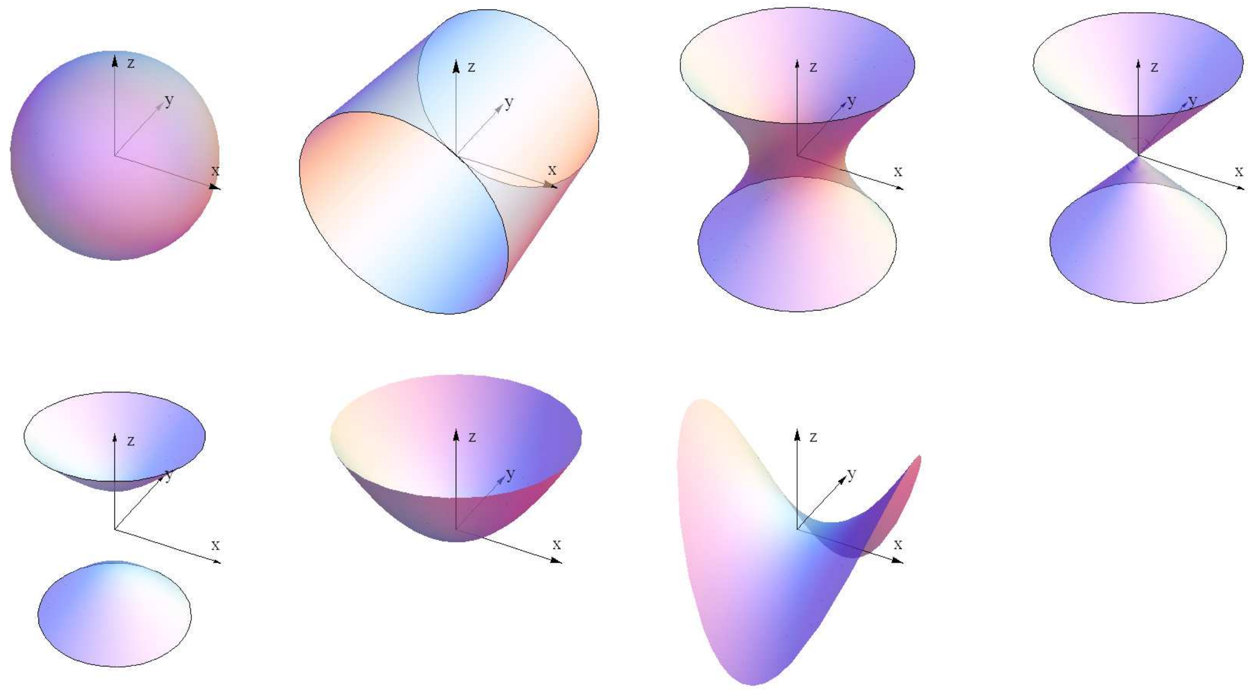

Figure 1: (Top row) Ellipsoid, Circular Cylinder, Hyperboloid of One Sheet, Cone,
(Bottom row) Hyperboloid of Two Sheets, Elliptic Paraboloid, Hyperbolic Paraboloid

- Then by the multivariable version of Taylor's theorem in $\mathbb{R}^2$, the function $f(x_1, \ldots, x_n) - f(P)$ will be closely approximated by its degree-2 Taylor polynomial T, which has the form $T = \sum_{1 \le i \le j \le n} a_{i,j} x_i x_i$, where $a_{i,j} = \begin{cases} f_{x_i,x_i}(P)/2 & \text{for } i = j \\ f_{x_i,x_j}(P) & \text{for } i \neq j \end{cases}$.
- Specifically, Taylor's theorem says that $\lim_{(x_1,\ldots,x_n)\to P} \dfrac{f(x_1,\ldots,x_n) - T - f(P)}{x_1^2 + \cdots + x_n^2} = 0$, which we can write more compactly as $f(x_1, \ldots, x_n) - f(P) = T + O(x_1^2 + \cdots + x_n^2)$.
- Now observe T is a quadratic form whose associated bilinear form has matrix $H/2$. By using an orthonormal change of basis, we may diagonalize this quadratic form, and the entries on the diagonal of the diagonalization are the eigenvalues of $H/2$.
- If $x_1', \ldots, x_n'$ is the new coordinate system, this means $f(x_1, \ldots, x_n) - f(P) = \frac{1}{2}\lambda_1 (x_1')^2 + \cdots + \frac{1}{2}\lambda_n (x_n')^2 + O[(x_1')^2 + \cdots + (x_n')^2]$.
- If all of the λ_i are positive then the error term is smaller than the remaining terms, and so we see that $f(x_1, \ldots, x_n) - f(P) > 0$ sufficiently close to P, meaning that P is a local minimum.
- Likewise, if all of the λ_i are negative then the error term is smaller than the remaining terms, and so we see that $f(x_1, \ldots, x_n) - f(P) < 0$ sufficiently close to P, meaning that P is a local maximum.
- If there is at least one positive eigenvalue λ_i and one negative eigenvalue λ_j, then approaching P along the direction x_i' yields values of f less than P, while approaching P along the direction x_j' yields values of f greater than P, so P is a saddle point.
- The other cases are inconclusive[4] because we can take (for example) the functions $f = x_1^2 + x_2^4$ and $g = x_1^2 - x_2^4$: then H has a single nonzero eigenvalue (corresponding to x_1), but f has a local minimum while g has a saddle point.

[4] Ultimately, the issue is that if some eigenvalues are zero, say, λ_1 then if we approach P along the direction x_1', the quadratic form T is constant, and so the sign of $f(x_1, \ldots, x_n) - f(P)$ along that path will be determined by the error term. Except in the case where the function is known to take both positive and negative values (guaranteeing a saddle point), any of the other possible behaviors not ruled out by the existence of positive or negative eigenvalues could occur.

- A fundamental component of the classification in the second derivatives test was the behavior of the quadratic form (and in particular, whether it was "always positive" or "always negative" for nonzero inputs). This behavior is quite important and we will record it:

- Definition: A quadratic form on a real vector space is positive definite if $Q(\mathbf{v}) > 0$ for every nonzero vector $\mathbf{v} \in V$, and it is negative definite if $Q(\mathbf{v}) < 0$ for every nonzero vector $\mathbf{v} \in V$.

  - If V is a real inner product space, then the square of the norm function $||\mathbf{v}||^2 = \langle \mathbf{v}, \mathbf{v} \rangle$ is a positive-definite quadratic form on V.
  - Indeed, it is not hard to see that the definition of an inner product on a real vector space is equivalent to saying that $\langle \mathbf{v}, \mathbf{v} \rangle$ is a quadratic form on V. Thus, all of the results for norms on inner product spaces also hold for positive-definite quadratic forms: for example, positive-definite quadratic forms obey the Cauchy-Schwarz inequality.
  - There are also a moderately useful weaker versions of these conditions: we say Q is positive semidefinite if $Q(\mathbf{v}) \geq 0$ for all $\mathbf{v} \in V$ and negative semidefinite if $Q(\mathbf{v}) \leq 0$ for all $\mathbf{v} \in V$.
  - As noted in the proof of the second derivatives test, if a real quadratic form is positive definite, then all the diagonal entries in its diagonalization are positive. Likewise, if a real quadratic form is negative definite, then all the diagonal entries in its diagonalization are negative. (The statements for semidefinite forms are similar, upon replacing "positive" with "nonnegative" and "negative" with "nonpositive".)

0.1.5 Sylvester's Law of Inertia

- We now discuss another fundamental result (which was, in fact, somewhat implicit in our earlier discussion of conics) regarding the possible diagonal entries for the diagonalization of a real quadratic form.

  - By making different choices for the matrix P (e.g., by rescaling it or selecting different row operations), we may obtain different diagonalizations of a given real quadratic form.
  - For example, with the quadratic form $Q(x, y) = x^2 + 2y^2$, which is already diagonal, if we change basis to $x' = x/2$, $y' = y/3$, then we obtain $Q(x, y) = 4(x')^2 + 18(y')^2$.
  - Indeed, it is clear that given any diagonalization, if we scale the ith row of the diagonalizing matrix by the scalar α, then the coefficient of the ith variable will be scaled by α^2.
  - Hence, by rescaling, we may change any positive coefficient to an arbitrary positive value and any negative coefficient to an arbitrary negative value.
  - It turns out that this is essentially the only possible change we may make to the diagonalization over $\mathbb{R}$.

- Theorem (Sylvester's Law of Inertia): Suppose V is a finite-dimensional real vector space and Q is a quadratic form on V. Then the numbers of positive diagonal entries, zero diagonal entries, and negative diagonal entries in any diagonalization of Q is independent of the choice of diagonalization.

  - The idea of this result is that we may decompose V as a direct sum of three spaces: one on which Q acts as a positive-definite quadratic form (corresponding to the positive diagonal entries), one on which Q acts as the zero map (corresponding to the zero entries), and one on which Q acts as a negative-definite quadratic form (corresponding to the negative diagonal entries).
  - Since this decomposition of V depends only on Q, these three spaces (and thus their dimensions) are independent of the choice of diagonalizing basis, and so the number of positive, zero, and negative diagonal entries in any diagonalization is necessarily fixed.
  - Proof: Since we are over a field of characteristic not 2, we may equivalently work with the symmetric bilinear form Φ associated to Q.
  - Let V_0 be the subspace of V given by $V_0 = \{\mathbf{v}_0 \in V : \Phi(\mathbf{v}_0, \mathbf{v}) = 0 \text{ for all } \mathbf{v} \in V\}$: then Φ acts as the zero map on V_0. Now write $V = V_0 \oplus V_1$: we claim that Φ is nondegenerate on V_1.
  - To see this, suppose $\mathbf{y} \in V_1$ has $\Phi(\mathbf{y}, \mathbf{v}_1) = 0$ for all $\mathbf{v}_1 \in V$: then for any $\mathbf{w} \in V$ we may write $\mathbf{w} = \mathbf{v}_0 + \mathbf{v}_1$ for $\mathbf{v}_i \in V_i$, in which case $\Phi(\mathbf{y}, \mathbf{w}) = \Phi(\mathbf{y}, \mathbf{v}_0) + \Phi(\mathbf{y}, \mathbf{v}_1) = 0$. But this would imply $\mathbf{y} \in V_0$ whence $\mathbf{y} = \mathbf{0}$.

- Now we will show that if Φ is nondegenerate on V_1, then V_1 decomposes as a direct sum $V_1 = V_+ \oplus V_-$, where Φ is positive-definite on V_+ and negative-definite on V_-.
- Let V_+ be the maximal subspace of V_1 on which Φ is positive-definite (since the condition is defined only on individual vectors, this subspace is well-defined), and define $V_- = \{\mathbf{w} \in V : \Phi(\mathbf{v}_+, \mathbf{w}) = 0 \text{ for all } \mathbf{v}_+ \in V_+\}$. Then by an application of Gram-Schmidt[5] (via Φ, rather than an inner product), we see that $V_1 = V_+ \oplus V_-$.
- It remains to show that Φ is negative-definite on V_-, so let $\mathbf{z} \in V_-$ be nonzero. Then by assumption, Φ is not positive-definite on $V_+ \oplus \langle \mathbf{z} \rangle$, so there exists some nonzero $\mathbf{v} = \mathbf{v}_+ + \alpha \mathbf{z}$ with $\mathbf{v}_+ \in V_+$ and $\alpha \in \mathbb{R}$ such that $\Phi(\mathbf{v}, \mathbf{v}) \leq 0$.
- We cannot have $\alpha = 0$ since then positive-definiteness would imply $\mathbf{v}_+ = 0$. Since $\Phi(\mathbf{v}, \mathbf{v}) = \Phi(\mathbf{v}_+, \mathbf{v}_+) + 2\alpha\Phi(\mathbf{v}_+, \mathbf{z}) + \alpha^2\Phi(\mathbf{z}, \mathbf{z}) = \Phi(\mathbf{v}_+, \mathbf{v}_+) + \alpha^2\Phi(\mathbf{z}, \mathbf{z})$, we have $\Phi(\mathbf{z}, \mathbf{z}) = \dfrac{1}{\alpha^2}[\Phi(\mathbf{v}, \mathbf{v}) - \Phi(\mathbf{v}_+, \mathbf{v}_+)]$.
- Then both terms are less than or equal to zero, and both cannot be zero. Hence $\Phi(\mathbf{z}, \mathbf{z}) < 0$ for all nonzero $\mathbf{z} \in V_-$ and so Φ is negative-definite on V_-.
- The desired result then follows from the direct sum decomposition $V = V_0 \oplus V_+ \oplus V_-$: if we select any diagonalization, then the restriction to the subspace generated by the basis vectors with diagonal entries 0, positive, negative (respectively) is trivial, positive-definite, negative-definite (respectively), and thus the number of diagonal elements is at least $\dim(V_0)$, $\dim(V_+)$, $\dim(V_-)$ (respectively). But since the total number of diagonal elements is $\dim(V) = \dim(V_0) + \dim(V_+) + \dim(V_-)$, we must have equality everywhere.
- Hence the numbers of positive diagonal entries, zero diagonal entries, and negative diagonal entries in any diagonalization of Q is independent of the choice of diagonalization, as claimed.

- We will also mention that there is some classical terminology associated with this result: the index of a quadratic form is the number of positive diagonal entries (in any diagonalization) and the signature is the difference between the number of positive and negative diagonal entries.
  - Some authors instead refer to the triple $(\dim V_+, \dim V_-, \dim V_0)$, or some appropriate permutation, as the signature of the quadratic form. These three values themselves are called the invariants of the form, and the value of any two of them (along with the dimension of the ambient space V) is sufficient to find the value of the other one.
  - For nondegenerate forms, where there are no 0 entries (so $\dim V_0 = 0$), the dimension of the space along with the value of $\dim V_+ - \dim V_-$ is sufficient to recover the two values.
  - Example: The quadratic form $Q(x, y, z) = x^2 - y^2 - z^2$ over $\mathbb{R}^3$ has index 1 and signature -1.
  - Example: The quadratic form $Q(x, y, z) = x^2 - z^2$ over $\mathbb{R}^3$ has index 1 and signature 0.
  - Example: The quadratic form $Q(x, y, z) = 5x^2 + 4xy + 6y^2 + 4yz + 7z^2$ over $\mathbb{R}^3$ has index 3 and signature 3, since we computed its diagonalization to have diagonal entries $3, 6, 9$.

Well, you're at the end of my handout. Hope it was helpful.

[5] The argument here is the same as for showing that $\dim(W) + \dim(W^\perp) = \dim(V)$ for an inner product. The Gram-Schmidt algorithm does not use the positive-definiteness of the inner product (it requires only linearity and symmetry), so the same argument also works for any bilinear form.