Summary of week 7 (lectures 19, 20 and 21)

Recall that if A is an $n \times p$ matrix over the field $\mathbb{R}$ then the column space of A, which by definition is the subspace of $\mathbb{R}^n$ consisting of all linear combinations of the columns of A, is given by

$$\mathrm{CS}(A) = \{\, Au \mid u \in \mathbb{R}^p \,\}.$$

We address the following problem: given $b \in \mathbb{R}^n$, find an element $u \in \mathbb{R}^p$ such that Au is as close as possible to b. In other words, we wish to find $u \in \mathbb{R}^p$ such that $Au = P(b)$, where P is the orthogonal projection from $\mathbb{R}^n$ to $\mathrm{CS}(A)$. We remark that although $P(b)$ is uniquely determined by b and A, and the equation $Au = P(b)$ is guaranteed to have a solution u, the solution need not be unique. However, if the columns of A are linearly independent then they form a basis of $\mathrm{CS}(A)$; in this case each element of $\mathrm{CS}(A)$ is uniquely expressible as a linear combination of the columns of A, and so the u such that $Au = P(b)$ will be unique.

From the discussion of orthogonal projections in last week's lectures, we know that $P(b)$ is the unique element of $\mathrm{CS}(A)$ such that $b - P(b)$ is orthogonal to all elements of $\mathrm{CS}(A)$. Thus we require $u \in \mathbb{R}^p$ to satisfy

$$(Av) \cdot (b - Au) = 0 \tag{1}$$

for all $v \in \mathbb{R}^p$. Note that if x and y are column vectors then ${}^{\mathrm{t}}x$ is a row vector, and the dot product $x \cdot y$ is the same as the matrix product $({}^{\mathrm{t}}x)y$. (Note that 1×1 matrices are identified with scalars; similarly, row and column vectors are identified with matrices having a single row or column.) Since ${}^{\mathrm{t}}(Av) = ({}^{\mathrm{t}}v)({}^{\mathrm{t}}A)$, Eq. (1) gives

$$({}^{\mathrm{t}}v)({}^{\mathrm{t}}A)(b - Au) = 0.$$

So, writing $y = ({}^{\mathrm{t}}A)(b - Au)$, we see that y is a column vector (in $\mathbb{R}^p$) with the property that $({}^{\mathrm{t}}v)y = 0$ for all $v \in \mathbb{R}^p$. That is, $v \cdot y = 0$ for all $v \in \mathbb{R}^p$. This forces y to be the zero vector, as no nonzero vector can be orthogonal to all vectors v. (Indeed, if $v \cdot y = 0$ for all v then, in particular, $y \cdot y = 0$, and by the positive definiteness of the dot product this forces $y = 0$.)

We have shown that if $Au = P(b)$ then $({}^{\mathrm{t}}A)(b - Au) = 0$. It is trivially checked that, conversely, if $({}^{\mathrm{t}}A)(b - Au) = 0$ then $(Av) \cdot (b - Au) = 0$ for all $v \in \mathbb{R}^p$, and so $Au = P(b)$. Now $({}^{\mathrm{t}}A)(b - Au) = ({}^{\mathrm{t}}A)b - ({}^{\mathrm{t}}A)Au$, and so $Au = P(b)$ if and only if

$$({}^{\mathrm{t}}A)Au = ({}^{\mathrm{t}}A)b. \tag{2}$$

Given b, this is a system of p linear equations for the p unknowns that are the entries of u. The equations have a unique solution U if and only if the coefficient matrix $({}^{\mathrm{t}}A)A$ is invertible; from the discussion above we know that this occurs if the columns of A are linearly independent.

For example, suppose that

$$A = \begin{pmatrix} 1 & a_1 \\ 1 & a_2 \\ \vdots & \vdots \\ 1 & a_n \end{pmatrix}, \qquad b = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}.$$

Then we find that

$$({}^{\mathrm{t}}A)b = \begin{pmatrix} 1 & 1 & \dots & 1 \\ a_1 & a_2 & \dots & a_n \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = \begin{pmatrix} \sum b_i \\ \sum a_i b_i \end{pmatrix}$$

$$({}^{\mathrm{t}}A)A = \begin{pmatrix} 1 & 1 & \dots & 1 \\ a_1 & a_2 & \dots & a_n \end{pmatrix} \begin{pmatrix} 1 & a_1 \\ 1 & a_2 \\ \vdots & \vdots \\ 1 & a_n \end{pmatrix} = \begin{pmatrix} n & \sum a_i \\ \sum a_i & \sum a_i^2 \end{pmatrix}.$$

If the numbers a_1, a_2, $\dots$, a_n are not all equal then the unique solution to Eq. (2) is given by

$$u = \begin{pmatrix} n & \sum a_i \\ \sum a_i & \sum a_i^2 \end{pmatrix}^{-1} \begin{pmatrix} \sum b_i \\ \sum a_i b_i \end{pmatrix}.$$

We now give an important practical application of the above theory. Suppose that we are given n pairs of real numbers, (a_1, b_1), (a_n, b_2), $\dots$, (a_n, b_n), and we plot these as n points in the Cartesian plane (using the a_i as the x-cordinates and the b_i as the y-coordinates). We wish to find the "line of best fit": the straight line that is, in some sense, as close as possible to all the points.

Of course, we need a way of deciding how close a line is to a set of points. However, it is helpful to reformulate the problem before attempting to say how closeness should be measured.

If the line has equation $y = c + mx$ then the n points (a_i, b_i) all lie on the line if the n equations $c + ma_i = b_i$ are all satisfied. These n equations can be written as a single vector equation

$$c \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} + m \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}, \tag{3}$$

or, equivalently, as a matrix equation

$$\begin{pmatrix} 1 & a_1 \\ 1 & a_2 \\ \vdots & \vdots \\ 1 & a_n \end{pmatrix} \begin{pmatrix} c \\ m \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}. \tag{4}$$

The numbers a_i and b_i are given, and the task is to find c and m. If there is a straight line passing through all the points then c and m can be chosen so that Eq. (4) holds. If there is no straight line passing through all the points then we want to choose c and m so that the left hand side of Eq. (4) is as close as possible to the right hand side.

Formulated in this way, the problem becomes the same as the problem we solved above. In effect, we are saying that the distance from the line $y = c + mx$ to the set of points $\{(a_1, b_1), (a_2, b_2), \ldots, (a_n, b_n)\}$ can be identified with

$$\left\| \begin{pmatrix} c + ma_1 \\ c + ma_2 \\ \vdots \\ c + ma_n \end{pmatrix} - \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} \right\|,$$

the distance between two vectors in $\mathbb{R}^n$, and the line of best fit is to be found by minimizing this quantity. Now

$$\left\| \begin{pmatrix} c + ma_1 \\ c + ma_2 \\ \vdots \\ c + ma_n \end{pmatrix} - \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} \right\|^2 = \sum_{i=1}^{n} (c + ma_i) - b_i)^2,$$

and $((c+ma_i) - b_i)^2$ is the square of the distance in $\mathbb{R}^2$ between the points (a_i, b_i) and $(a_i, c + ma_i)$. The point $(a_i, c + ma_i)$ is the point on the line $y = c + mx$ that has the same x-coordinate as (a_i, b_i). Thus the line of best fit can be characterized as the line that minimizes S, the sum of the squares† of the vertical distances from

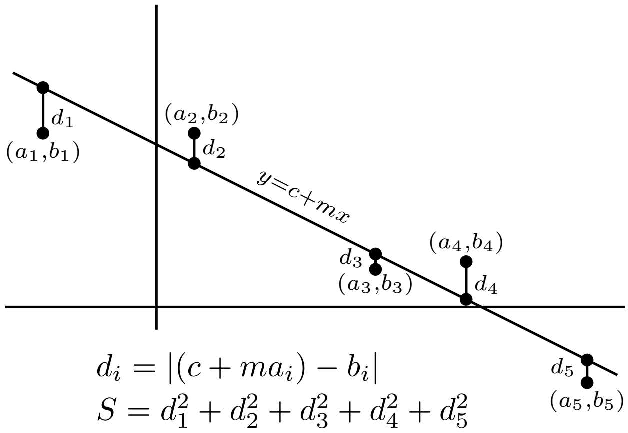

$$d_i = |(c + ma_i) - b_i|$$
$$S = d_1^2 + d_2^2 + d_3^2 + d_4^2 + d_5^2$$

the data points to the line. As we saw above, the coefficients c and m are given by the formula

$$\begin{pmatrix} c \\ m \end{pmatrix} = \begin{pmatrix} n & \sum a_i \\ \sum a_i & \sum a_i^2 \end{pmatrix}^{-1} \begin{pmatrix} \sum b_i \\ \sum a_i b_i \end{pmatrix}.$$

† For this reason, it is commonly called the *least squares* line of best fit.

For example, let us find the line of best fit for the points $(0,1),(1,2),(1,3)$ and $(2,5)$. We have the following table of values

i	1	2	3	3
a_i	0	1	1	3
b_i	1	2	3	5
a_ib_i	0	2	3	10
a_i^2	0	1	1	4

and we conclude that the line of best fit is $y = c + mx$, where

$$\begin{aligned}\begin{pmatrix} c \\ m \end{pmatrix} &= \begin{pmatrix} n & \sum a_i \\ \sum a_i & \sum a_i^2 \end{pmatrix}^{-1} \begin{pmatrix} \sum b_i \\ \sum a_i b_i \end{pmatrix} \\ &= \begin{pmatrix} 4 & 4 \\ 4 & 6 \end{pmatrix}^{-1} \begin{pmatrix} 11 \\ 15 \end{pmatrix} \\ &= \frac{1}{8} \begin{pmatrix} 6 & -4 \\ -4 & 4 \end{pmatrix}^{-1} \begin{pmatrix} 11 \\ 15 \end{pmatrix} = \begin{pmatrix} 3/4 \\ 2 \end{pmatrix}.\end{aligned}$$

The line of best fit is $y = (3/4) + 2x$.

A similar analysis can be applied to the problem of finding the least squares best fitting parabola $y = c + dx + ex^2$ for a given set of data points. Indeed, the same procedure can be applied for polynomials of any given degree, and other families of curves as well. For the case $y = c + dx + ex^2$, the task is to find c, d, e so that

$$c \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} + d \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} + e \begin{pmatrix} a_1^2 \\ a_2^2 \\ \vdots \\ a_n^2 \end{pmatrix} \approx \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix},$$

where '$\approx$' means 'as close as possible to'. The required c, d, e have the property that

$$\begin{pmatrix} 1 & a_1 & a_1^2 \\ 1 & a_2 & a_2^2 \\ \vdots & \vdots & \vdots \\ 1 & a_n & a_n^2 \end{pmatrix} \begin{pmatrix} c \\ d \\ e \end{pmatrix} = P \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}$$

where P is the projection of $\mathbb{R}^n$ onto the column space of the matrix

$$A = \begin{pmatrix} 1 & a_1 & a_1^2 \\ 1 & a_2 & a_2^2 \\ \vdots & \vdots & \vdots \\ 1 & a_n & a_n^2 \end{pmatrix}.$$

By calculations similar to those given above, we find that

$$
{}^{\mathrm{t}}AA\begin{pmatrix} c \\ d \\ e \end{pmatrix} = {}^{\mathrm{t}}A\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix},
$$

and if the columns of A are linearly independent then the unique solution is

$$
\begin{pmatrix} c \\ d \\ e \end{pmatrix} = \begin{pmatrix} n & \sum a_i & \sum a_i^2 \\ \sum a_i & \sum a_i^2 & \sum a_i^3 \\ \sum a_i^2 & \sum a_i^3 & \sum a_i^4 \end{pmatrix}^{-1} \begin{pmatrix} \sum b_i \\ \sum a_i b_i \\ \sum a_i^2 b_i \end{pmatrix}.
$$

It can be shown that the columns of A are linearly independent provided there are at least two distinct numbers in the set $\{a_1, a_2, \ldots, a_n\}$.

A large part of Lecture 20 was devoted to a discussion of Fourier series, essentially as in #9 on p. 113 of [VST]. We consider the set $\mathscr{C}'$ consisting of all continuous functions $f\colon \mathbb{R} \to \mathbb{R}$ such that $f(x) = f(x + 2\pi)$ for all $x \in \mathbb{R}$. It is clear that $\mathscr{C}'$ is a subspace of the space of all functions $\mathbb{R} \to \mathbb{R}$, and since any f such that $f(x) = f(x + 2\pi)$ for all x is uniquely determined by its restriction to $[-\pi, \pi]$, the space $\mathscr{C}'$ can, in effect, be identified with the space of all continuous functions on $[-\pi, \pi]$ that take the same value at π as at $-\pi$. We can make $\mathscr{C}'$ into an inner product space by defining $\langle f, g\rangle = \int_{-\pi}^{\pi} f(x)g(x)\,dx$ for all f, $g \in \mathscr{C}'$.

You should all remember, from secondary school mathematics, the formulas for $\sin(\alpha + \beta)$ and $\cos(\alpha + \beta)$, and be able to use these to derive (in a few seconds) the formulas for the products $2\sin(nx)\cos(mx)$, $2\cos(nx)\cos(mx)$ and $2\sin(nx)\sin(mx)$ given on p. 113 of [VST]. Using these formulas it can be shown (see [VST]) that the functions c_0, s_1, c_1, s_2, c_2, …, s_k, c_k form an orthogonal basis of a subspace of $\mathscr{C}'$. Let us call this space $\mathscr{C}_k$, and let P be the orthogonal projection from $\mathscr{C}'$ to $\mathscr{C}_k$. By Lemma 5.5 of [VST], if f is an arbitrary element of $\mathscr{C}'$ then $P(f)$ is given by

$$
P(f) = \frac{\langle c_0, f\rangle}{\langle c_0, c_0\rangle} c_0 + \frac{\langle s_1, f\rangle}{\langle s_1, s_1\rangle} s_1 + \frac{\langle c_1, f\rangle}{\langle c_1, c_1\rangle} c_1 + \cdots + \frac{\langle c_k, f\rangle}{\langle c_k, c_k\rangle} c_k,
$$

and $P(f)$ is the best approximation to f by an element of $\mathscr{C}_k$. Here the "best approximation to f" means the function $g \in \mathscr{C}_k$ such that $\int_{-\pi}^{\pi}(f(x) - g(x)^2\,dx$ is minimal.

If $\alpha \in \mathbb{C}$ then $\overline{\alpha}\alpha = |\alpha|^2$; so if V is a complex inner product space and $v \in V$ then for all $\alpha \in \mathbb{C}$,

$$
\|\alpha v\| = \langle \alpha v, \alpha v\rangle = \overline{\alpha}\langle v, \alpha v\rangle = \overline{\alpha}\alpha\langle v, v\rangle = |\alpha|^2\|v\|
$$

(by the fact that the inner product is semilinear in the first variable and linear in the second). In particular, if we put $\alpha = \frac{1}{\|v\|}$ then we obtain $\|\alpha v\| = 1$. It follows that an orthogonal sequence of vectors $v_1, v_2, \ldots, v_n$ can be converted into an orthonormal sequence of vectors $u_1, u_2, \ldots, u_n$ by dividing each vector in the sequence by its own length. That is, we define $u_i = \frac{1}{\|v_i\|}v_i$. This process is known as *normalizing* the vectors.

If $u_1, u_2, \ldots, u_n$ is an orthonormal basis of V then

$$\begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_n \end{pmatrix} \mapsto \lambda_1 u_1 + \lambda_2 u_2 + \cdots + \lambda_n u_n$$

is a bijective linear map—that is, a vector space isomorphism—from $\mathbb{C}^n$ to V. Writing f for this map, we find that for all $\lambda_i, \mu_i \in \mathbb{C}$,

$$\begin{aligned} \left\langle f\begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_n \end{pmatrix}, f\begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{pmatrix} \right\rangle &= \langle \sum_{i=1}^{n} \lambda_i u_i, \sum_{j=1}^{n} \mu_j u_j \rangle \\ &= \sum_{i=1}^{n} \sum_{j=1}^{n} \overline{\lambda}_i \mu_j \langle u_i, u_j \rangle \\ &= \sum_{i=1}^{n} \overline{\lambda}_i \mu_j \qquad \text{since } \langle u_i, u_j \rangle = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases} \\ &= \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_n \end{pmatrix} \cdot \begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{pmatrix}. \end{aligned}$$

We can express this property in words by saying that the vector space isomorphism f *preserves inner products*, and illustrate it, in a way, with the diagram below.

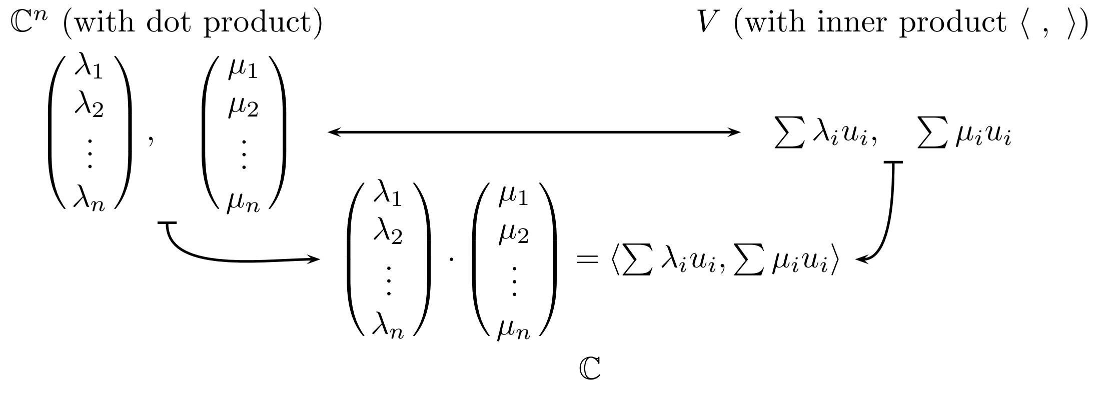

The upshot is that V with the inner product $\langle\ ,\ \rangle$ is essentially just a copy of $\mathbb{C}^n$ with the dot product.

If u_1, u_2, $\ldots$, u_n is an orthonormal basis of V and $v \in V$ then there exist scalars λ_i such that

$$v = \lambda_1 u_1 + \lambda_2 u_2 + \cdots + \lambda_n u_n.$$

It is easy to compute the λ_i, since

$$\langle u_i, v\rangle = \langle u_i, \sum_j \lambda_j u_j\rangle = \sum_j \lambda_j \langle u_i, u_j\rangle = \lambda_i$$

(since $\langle u_i, u_j\rangle = 0$ for all $j \neq i$, and $\langle u_i, u_j\rangle = 1$). Thus

$$v = \langle u_1, v\rangle u_1 + \langle u_2, v\rangle u_2 + \cdots + \langle u_n, v\rangle u_n. \tag{5}$$

This formula works for all $v \in V$ and all orthonormal bases u_1, u_2, $\ldots$, u_n. Observe, consequently, that for all v, $v' \in V$ we have that

$$\langle v, v'\rangle = \langle \sum_i \langle u_i, v\rangle u_i, \sum_j \langle u_j, v'\rangle u_j\rangle = \sum_i \overline{\langle u_i, v\rangle} \langle u_j, v'\rangle u_j.$$

Of course, this is the same as the dot product

$$\begin{pmatrix} \langle u_1, v\rangle \\ \langle u_2, v\rangle \\ \vdots \\ \langle u_n, v\rangle \end{pmatrix} \cdot \begin{pmatrix} \langle u_1, v\rangle \\ \langle u_2, v\rangle \\ \vdots \\ \langle u_n, v\rangle \end{pmatrix}$$

since, as explained above, inner products are preserved by the vector space isomorphism $\mathbb{C}^n \to V$ that the orthonormal basis u_1, u_2, $\ldots$, u_n provides.

For a basis v_1, v_2, $\ldots$, v_n that is merely orthogonal, rather than orthonormal, the formulas we have derived need to be modified slightly. For example, Eq. (5) above becomes

$$v = \frac{\langle v_1, v\rangle}{\langle v_1, v_1\rangle} v_1 + \frac{\langle v_2, v\rangle}{\langle v_2, v_2\rangle} v_2 + \cdots + \frac{\langle v_n, v\rangle}{\langle v_n, v_n\rangle} v_n.$$

See Proposition 5.10 of [VST].

Refer now to the the proof of 5.10 (i) given in [VST], and observe that the inequality $\langle v, v\rangle \geq \langle P(v), P(v)\rangle$ will be strict if $x \neq 0$, since $\langle x, x\rangle > 0$ in this case. Furthermore, $x = 0$ if and only if $P(v) = v$ (since $x = v - P(v)$), and since $P(v)$ is by definition the point of U closest to v it is clear that $P(v) = v$ if and only if $v \in U$. Thus we obtain the following strengthened version of Proposition 5.10 (i).

Proposition. *If* $(u_1, u_2, \ldots, u_n)$ *is an orthogonal basis for a subspace* U *of an inner product space* V*, then for all* $v \in V$ *we have*

$$\|v\|^2 \geq \sum_{i=1}^{n} \frac{|\langle u_i, v\rangle|^2}{\|u_i\|^2}$$

with equality if and only if v *is in the subspace* U*.*

This permits us also to strengthen Proposition 5.11 (i) of [VST]. Suppose that v, $w \in V$ with $w \neq 0$ and v not in the 1-dimensional space spanned by w. Following the proof given in [VST] and using the our strengthened form of 5.10 (i), we find that $|\langle v, w\rangle| < \|v\|\,\|w\|$. It is easily verified that if v is a scalar multiple of w (or if w is a scalar multiple of v) then in fact $|\langle v, w\rangle| = \|v\|\,\|w\|$.

Proposition. *If* V *is an inner product space and* v, $w \in V$*, then*

$$|\langle v, w\rangle| \leq \|v\|\,\|w\|$$

with equality if and only if v, w *are linearly dependent.*

Proposition 5.11 (ii) of [VST] was also proved in lectures, although we did not investigate the conditions under which equality holds. Interested students might work this out for themselves.

A start was also made on §5c of [VST]. We defined what it means to say that a linear transformation from one inner product space to another preserves inner products or preserves lengths. Proposition 5.12 was also proved in lectures, but I fear that the proof was poorly explained; so students are strongly encouraged to read the proof in [VST]. (Students are always strongly encouraged to read the proofs in [VST], but even more so in this case.) The crucial point of the proof is that it is possible to give a formula for the inner product $\langle u, v\rangle$, where u, $v \in V$ are arbitrary, in terms of lengths of vectors. The formula is proved (though not explicitly stated) in [VST]:

$$\langle u, v\rangle = \frac{1}{2}(\|u+v\|^2 - \|u\|^2 - \|v\|^2) + \frac{i}{2}(\|iu+v\|^2 - \|iu\|^2 - \|v\|^2)$$

where here $i = \sqrt{-1} \in \mathbb{C}$. If $T\colon V \to W$ is a linear transformation that preserves lengths then $\|T(u)\| = \|u\|$ and $\|T(v)\| = \|v\|$ (by the definition); moreover

$$\|T(u) + T(v)\| = \|T(u+v)\| = \|u+v\|,$$
$$\|iT(u)\| = \|T(iu)\| = \|iu\|,$$
$$\|iT(u) + T(v)\| = \|T(iu+v)\| = \|iu+v\|.$$

Consequently it follows that

$$\frac{1}{2}(\|T(u)+T(v)\|^2 - \|T(u)\|^2 - \|T(v)\|^2) + \frac{i}{2}(\|iT(u)+T(v)\|^2 - \|iT(u)\|^2 - \|T(v)\|^2)$$

is equal to

$$\frac{1}{2}(\|u+v\|^2 - \|u\|^2 - \|v\|^2) + \frac{i}{2}(\|iu+v\|^2 - \|iu\|^2 - \|v\|^2),$$

showing that $\langle T(u), T(v)\rangle = \langle u, v\rangle$.