Math 396. Orientations

In the theory of manifolds there will be a notion called "orientability" that is required to integrate top-degree differential forms. Orientations arise from certain notions in linear algebra, applied to tangent and cotangent spaces of manifolds. The aim of this handout is to develop these relevant foundations in linear algebra, and the globalization on manifolds will be given later in the course.

All vector spaces are assumed to be finite-dimensional in what follows.

1. Definitions and examples

Let V be a nonzero $\mathbf{R}$-vector space with dimension d, and let $\mathbf{v} = \{v_1, \dots, v_d\}$ be an ordered basis of V. This gives rise to a nonzero vector

$$\wedge(\mathbf{v}) \stackrel{\text{def}}{=} v_1 \wedge \cdots \wedge v_d \in \wedge^d(V)$$

in the line $\wedge^d(V)$. If $\mathbf{v}'$ is a second ordered basis, then $\wedge(\mathbf{v}')$ is another nonzero vector in the same line $\wedge^d(V)$, so $\wedge(\mathbf{v}') = c\wedge(v)$ for a unique $c \in \mathbf{R}^\times$. Concretely, if $T_{\mathbf{v},\mathbf{v}'} : V \simeq V$ is the unique linear automorphism satisfying $v_i' = T(v_i)$ for all i (it is the "change of basis matrix" from $\mathbf{v}'$-coordinates to $\mathbf{v}$-coordinates), then $c = \det T_{\mathbf{v},\mathbf{v}'}$ and $1/c = \det T^{-1}_{\mathbf{v},\mathbf{v}'} = \det T_{\mathbf{v}',\mathbf{v}}$. Hence, $c > 0$ (or equivalently, $1/c > 0$) if and only if $\wedge(\mathbf{v})$ and $\wedge(\mathbf{v}')$ lie in the same connected component of $\wedge^d(V) - \{0\}$.

Definition 1.1. An *orientation* μ of V is a choice of connected component of $\wedge^d(V) - \{0\}$, called the *positive* component with respect to μ. An *oriented vector space* is a nonzero vector space V equipped with a choice of orientation μ.

Clearly there are exactly two orientations on V, and if μ is one then we write $-\mu$ to denote the other (so $-(-\mu) = \mu$). The notion of orientation rests on the fact that the multiplicative group $\mathbf{R}^\times$ has exactly two connected components, and that the component containing the identity is a subgroup; since $\mathbf{C}^\times$ is connected, there is no analogous "$\mathbf{C}$-linear" concept of orientation (but in §3 we will see that the theory of orientations is especially pleasant on the $\mathbf{R}$-vector spaces underlying $\mathbf{C}$-vector spaces).

Here is an equivalent, and perhaps more concrete, formulation of the notion of orientation. Declare two ordered bases $\mathbf{v}$ and $\mathbf{v}'$ to be *similarly oriented* precisely when the linear automorphism of V that relates them (in either direction!) has positive determinant, and otherwise we say that $\mathbf{v}$ and $\mathbf{v}'$ are *oppositely-oriented*. The property of being similarly oriented is an equivalence relation on the set of ordered bases of V (check!), and the assignment $\mathbf{v} \mapsto \wedge(\mathbf{v}) \in \wedge^d(V)$ sends two ordered bases into the same connected component if and only if they are similarly oriented. Thus, we see that there are exactly two equivalence classes and that they correspond to orientations on V (the correspondence being that if we are given some equivalence class of $\mathbf{v}$'s then the associated orientation is the common connected component containing all such $\wedge(\mathbf{v})$'s). The ordered bases $\mathbf{v}$ of V such that $\wedge(\mathbf{v})$ lies in the positive component for μ are called *positive* bases with respect to μ (or *μ-oriented bases*).

Example 1.2. If $V = \mathbf{R}^d$ and $\{e_1, \dots, e_d\}$ is the standard ordered basis, the corresponding orientation class (i.e., connected component of $\wedge^d(V)$ containing $e_1 \wedge \cdots \wedge e_d$) is called the *standard orientation* of $\mathbf{R}^d$.

Remark 1.3. In order to include 0-dimensional manifolds in the general theory later on, it is convenient to use the convention $\wedge^0(V) = \mathbf{R}$ for all V, including $V = 0$. This is consistent with the result that $\wedge^d(V)$ is 1-dimensional for $d = \dim V$ when V is nonzero. However, when working with $\wedge^d(V)$ for $d = \dim V = 0$ we cannot interpret orientations of this line in terms of the language of bases of V. Largely for this reason, we will ignore the 0 space in this handout; anyway, including it in various results later on is always either meaningless or a triviality.

Definition 1.4. Let (V, μ) and (W, ν) be oriented vector spaces with common dimension $d > 0$. A linear isomorphism $T : V \simeq W$ is *orientation-preserving* if $\wedge^d(T) : \wedge^d(V) \simeq \wedge^d(W)$ carries the positive component to the positive component. Otherwise T is *orientation-reversing*.

Note that if $W = V$ and $\nu = \mu$ then $T : V \simeq V$ is orientation-preserving if and only if $\det T > 0$ because $\wedge^d(T) : \wedge^d(V) \simeq \wedge^d(V)$ is multiplication by $\det T \neq 0$ and such a multiplication map preserves a connected component of $\wedge^d(V) - \{0\}$ if and only if $\det T > 0$ (in which case both components are preserved). It is clear that $T : V \simeq W$ is orientation-preserving if and only if $T^{-1} : W \simeq V$ is, and that the property of being orientation-preserving is stable under composition. In particular, negation on V is orientation preserving if and only if $d = \dim V$ is even (as $\det_V(-1) = (-1)^d$), and so in general when T is orientation-preserving the same is true for $-T$ if and only if d is even. As a special case, if $\mathbf{e} = \{e_1, \ldots, e_d\}$ is μ-positive ordered basis of V, then $\{-e_1, \ldots, -e_d\}$ is in the orientation class $(-1)^d\mu$ on V; in particular, the opposite orientation class $-\mu$ on V is *not* generally represented by negating all vectors in a μ-positive basis (except when d is odd). If we simply negate one of the basis vectors e_i, then we get an ordered basis in the orientation class $-\mu$. Also, if $d \geq 2$ and we swap two of the e_i's then the new ordered basis $\mathbf{e}'$ has associated wedge product $\wedge(\mathbf{e}') = - \wedge (\mathbf{e})$, and so $\mathbf{e}'$ lies in the class $-\mu$; iterating this, for any permutation $\sigma \in S_d$, the ordered basis $\{e_{\sigma(1)}, \ldots, e_{\sigma(d)}\}$ lies in the orientation class $\mathrm{sign}(\sigma)\mu$.

A linear automorphism $T : V \simeq V$ is orientation-preserving with respect to one orientation μ (on both copies of V) if and only if it is so using $-\mu$ (on both copies of V), as each condition is equivalent to the property "$\det T > 0$" that has nothing to do with the choice of μ. Hence, linear automorphisms $T : V \simeq V$ are called *orientation-preserving* when $\det T > 0$, as this is equivalent to T being orientation-preserving in the sense of the general definition above when both copies of V are endowed with a common orientation on V.

Example 1.5. If L is a 1-dimensional vector space, then an orientation of L is a choice of connected component of $L - \{0\}$ (or more concretely, it is a choice of "half-line"). A vector in the chosen component will be called *positive*. (We can say in general that orienting a nonzero V of dimension d is the same thing as orienting the line $\wedge^d(V)$.)

If L' is another 1-dimensional space, and if any two of the three lines L, L', or $L \otimes L'$ are oriented, there is a unique orientation on the third such that for nonzero $e \in L$ and $e' \in L'$ (so $e \otimes e' \in L \otimes L'$ is nonzero) if any two of three vectors e, e', or $e \otimes e'$ is positive in their respective line then so is the third. Such a triple of orientations on L, L', and $L \otimes L'$ is called *compatible*. Fix an orientation on L. If we orient L' then for compatibility to hold we must orient $L \otimes L'$ by declaring the positive component of $(L \otimes L') - \{0\}$ to be the common one containing $e \otimes e'$ for all positive $e \in L$ and $e' \in L'$ (why do all such lie in a common component of $(L \otimes L') - \{0\}$?), and if we orient $L \otimes L'$ then we must orient L' by declaring the positive component of $L' - \{0\}$ to consist of those nonzero $e' \in L'$ such that $e \otimes e'$ is positive in $L \otimes L'$ whenever e is positive in L. In the same manner, if $L_1, \ldots, L_N$ is an ordered collection of lines and each is oriented, then we get a natural orientation on $L_1 \otimes \cdots \otimes L_N$ by declaring the tensor product of positive vectors to be positive. These simple constructions in the 1-dimensional case have many pleasing "applications" in the higher-dimensional case, as we shall now see.

Example 1.6. Let (V, μ) and (V', μ') be oriented vector spaces with respective dimensions d and d'. Select an ordering among V and V', say V first and V' second. There is a natural isomorphism of lines

$$\wedge^d(V) \otimes \wedge^{d'}(V') \simeq \wedge^{d+d'}(V \oplus V') \tag{1}$$

by Theorem 2.4 from the handout on tensor algebras. Since the lines $\wedge^d(V)$ and $\wedge^{d'}(V')$ are oriented, we thereby get an orientation on $\wedge^{d+d'}(V \oplus V')$ via Example 1.5 and hence we get an orientation on $V \oplus V'$. More concretely, the positive component of $(\wedge^{d+d'}(V \oplus V')) - \{0\}$ is the one containing $e \wedge e'$ where $e \in \wedge^d(V)$ and $e' \in \wedge^{d'}(V')$ are any vectors in the positive components of the nonzero loci on these lines. In terms of ordered bases, if $\{v_i\}$ is a μ-oriented basis of V and $\{v'_i\}$ is a μ'-oriented basis of V' then $\{v_1, \dots, v_d, v'_1, \dots, v'_{d'}\}$ is a positive basis of $V \oplus V'$; this rests on how we *defined* the isomorphism (1). This orientation of $V \oplus V'$ is denoted $\mu \oplus \mu'$ and is called the *direct sum* orientation (with respect to the chosen ordering among V and V').

Note that if we swap the roles of V and V' then we are led to use the isomorphism

$$\wedge^{d'}(V') \otimes \wedge^d(V) \simeq \wedge^{d+d'}(V' \oplus V) \simeq \wedge^{d+d'}(V \oplus V')$$

that multiplies the orientation by a sign of $(-1)^{dd'}$. Hence, when at least one of V or V' is even-dimensional over $\mathbf{R}$ (such as $\mathbf{C}$-vector spaces!) then the orientation on $V \oplus V'$ is intrinsic, but otherwise it depends on the choice of ordering among the two factor spaces. Clearly if $T : V \simeq W$ and $T' : V' \simeq W'$ are orientation-preserving isomorphisms between oriented $\mathbf{R}$-vector spaces and we pick compatible orderings (with V and W first and V' and W' second, or the other way around) then $T \oplus T'$ is orientation-preserving for the direct sum orientations; when one of these vector spaces is even-dimensional then $T \oplus T'$ is orientation-preserving regardless of how we order the pairs $\{V, V'\}$ and $\{W, W'\}$.

Example 1.7. Let (V, μ) be an oriented vector space with dimension d and let $W \subseteq V$ be a proper subspace with dimension $d_0 \geq 0$. By Theorem 2.5 in the handout on tensor algebras, there is a natural isomorphism

$$\wedge^{d_0}(W) \otimes \wedge^{d-d_0}(V/W) \simeq \wedge^d(V) \tag{2}$$

(defined in the evident manner in case $d_0 = 0$) and the line $\wedge^d(V)$ is oriented, so the tensor product on the left is oriented. Hence, by Example 1.5, orienting one of W or V/W immediately determines a preferred orientation on the other via this isomorphism relating the top exterior-power lines. By inspecting the definition of this isomorphism, we get the following translation in terms of oriented bases when $d_0 > 0$ (i.e., $W \neq 0$). Let $\{w_j\}$ be an ordered basis of W and let $\{\overline{v}_i\}$ be an ordered basis of V/W. The property of whether or not the ordered basis

$$\{w_1, \dots, w_{d_0}, v_1, \dots, v_{d-d_0}\}$$

of V is positive (where $v_i \in V$ represents $\overline{v}_i \in V/W$) is independent of the v_i's because under (2) the wedge product

$$w_1 \wedge \cdots \wedge w_{d_0} \wedge v_1 \wedge \cdots \wedge v_{d-d_0} \in \wedge^d(V)$$

corresponds to

$$(w_1 \wedge \cdots \wedge w_{d_0}) \otimes (\overline{v}_1 \wedge \cdots \wedge \overline{v}_{d-d_0}) \in \wedge^{d_0}(W) \otimes \wedge^{d-d_0}(V/W)$$

and hence is *independent* of the choices of $v_i \in V$ lifting $\overline{v}_i \in V/W$.

To make this more concrete, the condition relating orientations on W and V/W (given the initial fixed orientation on V) for $W \neq 0$ is that upon picking an orientation on one of W or V/W there is a unique orientation on the other such that when both $\{w_i\}$ and $\{\overline{v}_i\}$ are positive bases then $\{w_1, \dots, w_{d_0}, v_1, \dots, v_{d-d_0}\}$ is a positive basis for V. We call a triple of orientations μ on V, ν on W, and $\overline{\mu}$ on V/W *compatible* when they satisfy the relation just stated. If we first pick ν then the uniquely determined compatible $\overline{\mu}$ is called the *quotient* orientation (with respect to μ and ν), whereas if we first pick $\overline{\mu}$ then the uniquely determined compatible ν is called the *subspace* orientation (with respect to μ and $\overline{\mu}$). Likewise, in case $W = 0$ the data of an orientation ν on W is just the specification of a sign ± 1 and for a compatible triple $(\mu, \overline{\mu}, \nu)$ the sign is 1 (resp. -1) if and only if the orientations μ and $\overline{\mu}$ on V and $V/W = V$ are the same (resp. opposite).

Example 1.8. Let (V, μ) be an oriented vector space and let $W_1, \dots, W_N$ be an *ordered set* of mutually transverse proper nonzero subspaces. Let $W' = \cap W_j$. Pick orientations ν_j on each W_j, so by the preceding example we get quotient orientations $\overline{\mu}_j$ on each V/W_j. By Theorem 2.6 in the handout on tensor algebras, there is a canonical isomorphism

$$\wedge^{c_1}(V/W_1) \otimes \cdots \otimes \wedge^{c_N}(V/W_N) \simeq \wedge^{c}(V/W') \tag{3}$$

with $c_j = \operatorname{codim}(W_j) = \dim(V/W_j)$ and $c = \operatorname{codim}(W') = \dim(V/W')$. The orientation $\overline{\mu}_j$ on V/W_j puts an orientation on the line $\wedge^{c_j}(V/W_j)$ for all j, so via the above isomorphism we get an orientation on $\wedge^c(V/W')$. Hence, we get an orientation on V/W'. Since V is oriented, by the preceding example we thereby obtain a natural orientation on W'!

In terms of ordered bases, if $W' \neq 0$ (resp. if $W' = 0$) when is a basis of W' positive (resp. when is the orientation sign on $W' = 0$ equal to 1 rather than -1)? The transversality condition ensures (for dimension reasons) that the injective map

$$V/W' \to (V/W_1) \oplus \cdots \oplus (V/W_N) \tag{4}$$

is an isomorphism, and so if we pick positive bases of each W_j then the positive bases of V/W_j are those whose lifts combine with the positive bases of W_j to be a positive bases of V (using the ordering that puts the basis vectors from W_j first). More specifically, choose an ordered set of vectors in V lifting a positive basis of V/W_1, followed by a lift of a positive basis of V/W_2, and so on, so we arrive at an ordered set of c independent vectors $\{v_1, \dots, v_c\}$ in V whose first c_1 members lift a positive basis of V/W_1, whose next c_2 members lift a positive basis of V/W_2, and so on. This ordered list of c vectors lifts a basis of V/W', and if $W' \neq 0$ then a basis $\{w'_1, \dots, w'_{d-c}\}$ of W' is positive (resp. if $W' = 0$ then its orientation sign is 1 rather than -1) precisely when the ordered basis

$$\{w'_1, \dots, w'_{d-c}, v_1, \dots, v_c\}$$

of V is positive (with respect to the initial orientation given on V). The reason that (4) is related to (3) for the purposes of computing orientations is because both rest on the same ordering of the set $\{V/W_1, \dots, V/W_N\}$ (namely, first V/W_1, then V/W_2, and so on).

Example 1.9. We wish to explain how choosing an orientation μ on an $\mathbf{R}$-vector space V determines an orientation on the $\mathbf{R}$-linear dual $V^\vee$. We first focus on the 1-dimensional case. If L is 1-dimensional then the evaluation pairing

$$L \otimes L^\vee \to \mathbf{R}$$

is an *isomorphism*, and so upon orienting L and using the canonical orientation of $\mathbf{R}$ (with the usual positive half-line as the preferred component of $\mathbf{R} - \{0\}$) Example 1.5 provides a preferred orientation on $L^\vee$. Concretely, if $e \in L$ is a basis and $e^* \in L^\vee$ is the dual basis (so $e \otimes e^*$ maps to $e^*(e) = 1 > 0$ under the evaluation pairing) then e^* is positive in $L^\vee - \{0\}$ if and only if e is positive in $L - \{0\}$.

Now we pass to the higher-dimensional case (as the case $V = 0$ is handled in a similar trivial manner, left to the reader). Let V be a nonzero $\mathbf{R}$-vector space with dimension $d > 0$, and let μ be an orientation on V. In particular, by definition of "orientation" the line $\wedge^d(V)$ is thereby oriented, and so its dual $(\wedge^d(V))^\vee$ has an induced orientation by the preceding paragraph. By

Corollary 3.2 in the handout on tensor algebras, there is a natural identification between $\wedge^d(V^\vee)$ and $(\wedge^d(V))^\vee$, and so we get a preferred orientation on $\wedge^d(V^\vee)$. Equivalently, we have a preferred orientation on $V^\vee$, denoted $\mu^\vee$ and called the *dual orientation* (recovering the procedure in the preceding paragraph when $d = 1$). Concretely, if $\{e_i\}$ is a basis of V with dual basis $\{e_i^*\}$ then since $e_1 \wedge \cdots \wedge e_d$ is dual to $e_1^* \wedge \cdots \wedge e_d^*$ under the duality pairing between the lines $\wedge^d(V)$ and $\wedge^d(V^\vee)$ it follows that $\{e_i\}$ is positive if and only if $\{e_i^*\}$ is positive. (In the 19th century, the dual orientation would be *defined* by this criterion, and it would be checked to be independent of the choice of positive basis of V as follows: when changing from one positive basis of V to another, the resulting change of basis matrix between dual bases of $V^\vee$ is the transpose of the change of basis matrix between the initial positive bases of V and so has positive determinant. The modern language of tensor algebra permits us to give definitions that suppress such explicit use of bases.) The basis criterion shows that $-(\mu^\vee) = (-\mu)^\vee$ and that the isomorphism $V \simeq V^{\vee\vee}$ we have $\mu = \mu^{\vee\vee}$.

Remark 1.10. Since duality interchanges subspaces and quotients, it is natural to ask how the formation of the dual orientation interacts with subspaces and quotient orientations (as in Example 1.7). The relation is a little tricky. Let W be a nonzero proper subspace of V, and choose compatible orientations μ, ν, and $\overline{\mu}$ of V, W, and V/W respectively. We get a triple of respective dual orientations $\mu^\vee$, $\overline{\mu}^\vee$, and $\nu^\vee$ on $V^\vee$, its nonzero proper subspace $(V/W)^\vee$, and the associated nonzero quotient $W^\vee$. Is this latter triple a compatible one? If V has dimension d and W has dimension d_0, then we claim that this triple of dual orientations is compatible up to a sign of $(-1)^{d_0(d-d_0)}$ (i.e., it is compatible when $d_0(d-d_0)$ is even and is not when $d_0(d-d_0)$ is odd).

To see what is going on, it is simplest to look at bases. Let $\{w_1, \dots, w_{d_0}\}$ be a ν-positive basis of W and $\{v_1, \dots, v_{d-d_0}\}$ a lift of a $\overline{\mu}$-positive basis of V/W, so $\{w_1, \dots, w_{d_0}, v_1, \dots, v_{d-d_0}\}$ is a μ-positive basis of V due to the hypothesis of compatibility among μ, ν, and $\overline{\mu}$. The ordered dual basis

$$\{w_1^*, \dots, w_{d_0}^*, v_1^*, \dots, v_{d-d_0}^*\}$$

of $V^\vee$ is therefore $\mu^\vee$-positive. Re-ordering this basis of $V^\vee$ by moving each w_j^* pass all $d - d_0$ of the v_i^*'s introduces $d_0(d-d_0)$ minus signs on the orientation class, so the ordered basis

$$\{v_1^*, \dots, v_{d-d_0}^*, w_1^*, \dots, w_{d_0}^*\} \tag{5}$$

of $V^\vee$ is in the orientation class of $\varepsilon\mu^\vee$ with $\varepsilon = (-1)^{d_0(d-d_0)}$. However, $W^\vee$ is a quotient of $V^\vee$ and the w_j^*'s in $V^\vee$ lift the $\nu^\vee$-positive ordered basis of $W^\vee$ dual to the ν-positive $\{w_j\}$, and likewise $(V/W)^\vee$ is a subspace of $V^\vee$ that contains the v_i^*'s as an ordered basis dual to the $\overline{\mu}$-positive ordered basis of $\overline{v}_i$'s in V/W, so $\{v_i^*\}$ in $(V/W)^\vee$ is $\overline{\mu}^\vee$-positive. (Note that the v_i^*'s as elements of $(V/W)^\vee$ really do not depend on the choices of liftings v_i because the v_i^*'s kill the w_j's and hence kill W, thereby inducing linear functionals on V/W that one checks are indeed dual to the $\overline{v}_i$'s.) To summarize, the orientation $\varepsilon\mu^\vee$ on $V^\vee$ determined by the ordered basis in (5) satisfies the condition to be compatible with $\overline{\mu}^\vee$ and $\nu^\vee$, whence the triple of dual orientations $\mu^\vee$, $\nu^\vee$, and $\overline{\mu}^\vee$ is compatible up to the sign $\varepsilon = (-1)^{d_0(d-d_0)}$. The preceding discussion has a trivial analogue in case $W = 0$ that we leave to the interested reader.

2. Symmetric bilinear pairings

We now study symmetric bilinear pairings $B : V \times V \to F$ for a d-dimensional vector space V over a field F, and we will then specialize to the case $F = \mathbf{R}$. When working with $F = \mathbf{R}$ (which is not a field of characteristic 2), we will also use the equivalent language of non-degenerate quadratic spaces. In the algebraic generality, we have seen in §3 in the handout on tensor algebras that if

B is a perfect pairing, then the induced bilinear pairings $B^{\otimes n}$ and $\wedge^n(B)$ on $V^{\otimes n}$ and $\wedge^n(V)$ are perfect, and so is $\mathrm{Sym}^n(B)$ when $n! \neq 0$ in F. These pairings are also symmetric:

Lemma 2.1. *If $B : V \times V \to F$ is symmetric, then so are $B^{\otimes n}$, $\mathrm{Sym}^n(B)$, and $\wedge^n(B)$ as bilinear forms on $V^{\otimes n}$, $\mathrm{Sym}^n(B)$, and $\wedge^n(V)$. Moreover, if $F = \mathbf{R}$ and B is positive-definite then these pairings are positive-definite.*

In general, if $F = \mathbf{R}$ and B has mixed signature then the signatures of the induced pairings on the nth tensor, symmetric, and exterior powers of V are given by some combinatorial formulas that we leave to the interested reader to sort out. (We will not require them.)

Proof. If $\beta : W \times W \to F$ is a bilinear form and $\{w_i\}$ is a spanning set of W then to verify the symmetry identity $\beta(w, w') = \beta(w', w)$ it suffices to consider w and w' in the spanning set $\{w_i\}$. Thus, the condition of symmetry may be checked in each case by working with elementary products (of tensor, symmetric, or exterior type), and in these cases we have explicit formulas for the values of the pairings, namely $\prod_{i=1}^n B(v_i, v_i')$, $\sum_{\sigma \in S_n} \prod_{i=1}^n B(v_i, v'_{\sigma(i)})$, and $\det(B(v_i, v_j'))$. It must be checked that swapping v_i and v_i' for *all* i does not change the value. This is clear in the first case because $B(v_i, v_i') = B(v_i', v_i)$. In the last case, since $B(v_i, v_j') = B(v_j', v_i)$ we get the desired equality because the determinant of a matrix is unaffected by passing to the transpose. In the case of symmetric powers we use the identity

$$\prod_{i=1}^n B(v_i, v'_{\sigma(i)}) = \prod_{i=1}^n B(v_{\sigma^{-1}(i)}, v_i') = \prod_{i=1}^n B(v_i', v_{\sigma^{-1}(i)})$$

to get the desired equality after summing over all $\sigma \in S_n$ (as $\sigma \mapsto \sigma^{-1}$ is a permutation of the elements of S_n).

Now suppose that $F = \mathbf{R}$ and B is positive-definite. We certainly *cannot* check positive-definiteness by merely working with spanning sets; for example, $B((x, y), (x', y')) = xx' - 10xy' + yy'$ is not positive-definite ($B((1, 1), (1, 1)) = -8$) but for $v = (1, 0)$ or $v = (0, 1)$ we have $B(v, v) = 1$. Hence, we have to use another viewpoint. Recall from the handout on quadratic spaces that for any non-degenerate symmetric bilinear form over any field, there exists a basis $\{e_i\}$ of V such that $B(e_i, e_j) = 0$ for all $i \neq j$ and $B(e_i, e_i) \neq 0$ for all i. The positive-definteness aspect over $\mathbf{R}$ is encoded in precisely the fact that $B(e_i, e_i) > 0$ for all i. We fix such a basis on V, and aim to use it to find analogous such bases (satisfying the positivity condition as well) in the tensor, symmetric, and exterior powers of V. Direct calculation shows that the respective habitual bases of $V^{\otimes n}$ and $\wedge^n(V)$ in terms of the e_i's are a set of pairwise orthogonal vectors for $B^{\otimes n}$ and $\wedge^n(V)$ whose self-pairings are 1. We now work out the pairings in the case of symmetric powers, and the reader will see that it is the same argument as in the discussion following Corollary 3.2 in the handout on tensor algebras. We suppose $d = \dim V$ is positive and we fix a monotone sequence $I = \{i_1, \dots, i_n\}$ of integers between 1 and d; define $e_I = e_{i_1} \cdot \dots \cdot e_{i_n}$. For two monotonically increasing sequences of n indices I and I' we see that $\prod_{r=1}^n \langle e_{i_r'}, e_{i_{\sigma(r)}} \rangle$ vanishes unless $i_{\sigma(r)} = i_r'$ for all r (in which case it equals 1), and the monotonicity property of the i_r's and the i_r''s implies that the non-vanishing holds if and only if $I' = I$ and for each $1 \leq j \leq d$ the permutation σ individually permutes the set of $m_j(I)$ indices r such that $i_r = j$. There are $m_j(I)!$ such permutations of the set of r's with $i_r = j$ when $m_j(I) > 0$, and so there are $\prod_{j=1}^d m_j(I)!$ such permutations in S_n in total. This shows that the e_I's are pairwise perpendicular and that the self-pairing of e_I is $\prod_{j=1}^d m_j(I)! > 0$. ∎

Remark 2.2. For later purposes, we remind the reader of the conclusion of the study of duality pairings on symmetric powers in the handout on tensor algebras: over any field F, the duality morphism $\mathrm{Sym}^n(V^\vee) \to (\mathrm{Sym}^n(V))^\vee$ (an isomorphism when $n! \neq 0$ in F) satisfies

$$e_{i_1}^* \cdot \cdots \cdot e_{i_n}^* \mapsto \prod_{j=1}^{d} m_j(I)! \cdot (e_{i_1} \cdot \cdots \cdot e_{i_n})^*.$$

Since a positive-definite symmetric bilinear form over $F = \mathbf{R}$ is just an inner product by another name, we conclude that if V is an inner product space then $V^{\otimes n}$, $\mathrm{Sym}^n(V)$, and $\wedge^n(V)$ have natural structures of inner product space. For applications to general relativity we wish to avoid positive-definiteness restrictions (since the quadratic form of interest in spacetime is $x^2 + y^2 + z^2 - c^2t^2$, and this has signature $(3,1)$). Hence, we shall focus on general non-degenerate quadratic spaces over $\mathbf{R}$, which is to say V endowed with a non-degenerate symmetric bilinear form $\langle \cdot, \cdot \rangle$ that may have mixed signature. Thus, the tensor, symmetric, and exterior powers of V are likewise non-degenerate quadratic spaces over $\mathbf{R}$ with possibly mixed signature. A vector $v \in V$ is a *unit vector* if $\langle v, v \rangle = \pm 1$, and a basis $\{e_i\}$ of V is *orthonormal* if $\langle e_i, e_j \rangle = 0$ for $i \neq j$ and $\langle e_i, e_i \rangle = \varepsilon_i = \pm 1$ for each i; the Gram-Schmidt process produces such bases. Let us record an important and useful observation that was essentially shown in the preceding proof (up to the tracking of signs):

Corollary 2.3. *Let V be a quadratic space of dimension $d > 0$ over $\mathbf{R}$, with $\langle \cdot, \cdot \rangle$ the corresponding non-degenerate symmetric bilinear form. Let $\{e_i\}$ be an orthonormal basis and $\varepsilon_i = \langle e_i, e_i \rangle = \pm 1$. The bases*

$$\{e_{i_1} \otimes \cdots \otimes e_{i_n}\}_{1 \le i_j \le d}, \ \ \{e_{i_1} \wedge \cdots \wedge e_{i_n}\}_{1 \le i_1 < \cdots < i_n \le d}$$

of $V^{\otimes n}$ and $\wedge^n(V)$ are orthonormal bases for the induced quadratic structures on these spaces, with $e_{i_1} \otimes \cdots \otimes e_{i_n}$ and $e_{i_1} \wedge \cdots \wedge e_{i_n}$ each having self-pairing $\prod_{j=1}^n \varepsilon_{i_j} = \pm 1$. An orthonormal basis of $\mathrm{Sym}^n(V)$ is given by

$$\left\{ \frac{e_{i_1} \cdot \cdots \cdot e_{i_n}}{\sqrt{\prod_{j=1}^d m_j(I)!}} \right\}_{1 \le i_1 \le \cdots \le i_n \le d}$$

for $1 \le i_1 \le \cdots \le i_n \le d$ where for $I = \{i_1, \dots, i_n\}$ we define $m_j(I)$ to be the number of $1 \le r \le n$ such that $i_r = j$; the self-pairing of $(e_{i_1} \cdot \cdots \cdot e_{i_n})/\sqrt{\prod_j m_j(I)!}$ is $\prod_{j=1}^n \varepsilon_{i_j} = \pm 1$.

Now we focus on the especially interesting case of the exterior powers $\wedge^r(V)$ when $(V, \langle \cdot, \cdot \rangle)$ is a nonzero non-degenerate quadratic space over $\mathbf{R}$ and $1 \le r \le \dim V$. Let $\langle \cdot, \cdot \rangle_r$ denote the induced symmetric bilinear form on $\wedge^r(V)$, so for $v_1, \dots, v_r, v'_1, \dots, v'_r \in V$ we have

$$\langle v_1 \wedge \cdots \wedge v_r, v'_1 \wedge \cdots \wedge v'_r \rangle_r = \det(\langle v_i, v'_j \rangle).$$

In particular, if $v'_j = v_j$ for all j and our quadratic space is positive-definite then this determinant must be non-negative (and is positive whenever $v_1 \wedge \cdots \wedge v_r \neq 0$, which is to say that the v_i's are linearly independent); such a positivity result for $\det(\langle v_i, v_j \rangle)$ can be seen directly without going through any of the preceding considerations (see the end of the proof of Theorem 2.6). If $\|\omega\|_r = \sqrt{|\langle \omega, \omega \rangle_r|}$ denotes the associated "length" on $\wedge^r V$ (which satisfies $\|c\omega\|_r = |c| \|\omega\|_r$ but otherwise violates the norm axioms except when $\langle \cdot, \cdot \rangle_r$ is positive-definite or negative-definite) then we have a simple formula for the "length" of an elementary wedge product:

$$\|v_1 \wedge \cdots \wedge v_r\|_r = \sqrt{|\det(\langle v_i, v_j \rangle)|}.$$

Passing to the case $r = d = \dim V$, for any ordered basis $\{v_1, \dots, v_d\}$ of V the nonzero $v_1 \wedge \cdots \wedge v_d$ in the 1-dimensional non-degenerate quadratic space $\wedge^d(V)$ has length $\sqrt{|\det(\langle v_i, v_j \rangle)|}$. This has a very interesting interpretation in the positive-definite case: it is the *volume* of the parallelotope

$$P = \{\sum t_i v_i \,|\, 0 \le t_i \le 1\} \subseteq V$$

"spanned" by $v_1, \dots, v_d$. For example, if V is a positive-definite plane then for independent $v, w \in V$ we claim that

$$\|v \wedge w\|_2 = \sqrt{\|v\|\|w\| - \langle v, w\rangle^2}$$

is the area of the parallelogram "spanned" by v and w.

In order to make sense of this claim concerning volumes, we need to clarify the meaning of volume in the inner product space V that has not been coordinatized. If $\mathbf{e} = \{e_i\}$ is an ordered orthonormal basis of V then such a basis defines a linear isomorphism $\phi_{\mathbf{e}} : V \simeq \mathbf{R}^d$ (sending e_i to the ith standard basis vector) and $\phi_{\mathbf{e}}$ carries the given inner product on V over to the standard one on $\mathbf{R}^d$. If $\mathbf{e}'$ is a second ordered orthonormal basis of V then $\phi_{\mathbf{e}'} : V \simeq \mathbf{R}^d$ is another linear isomorphism carrying the given inner product over to the standard one. Hence, the linear automorphism $T = \phi_{\mathbf{e}'} \circ \phi_{\mathbf{e}}^{-1}$ of $\mathbf{R}^d$ *preserves* the standard inner product and so is represented by an orthogonal matrix; in particular, $|\det T| = 1$. Since $\phi_{\mathbf{e}'} = T \circ \phi_{\mathbf{e}}$, it follows that for any subset $S \subseteq V$ we have $\phi_{\mathbf{e}'}(S) = T(\phi_{\mathbf{e}}(S))$. By the linear case of the Change of Variables formula, we conclude that the property of $\phi_{\mathbf{e}}(S)$ being rectifiable and (in such cases) the value of its volume (perhaps infinite) is *independent* of the choice of $\mathbf{e}$. In this sense we can speak of volume intrinsically for subsets of an inner product space without choosing linear coordinates (or rather, we use an arbitrary system of linear orthonormal coordinates, the choice of which does not matter). Now we can prove:

Theorem 2.4. *If $\{v_1, \dots, v_d\}$ is an ordered basis of an inner product space $(V, \langle\cdot,\cdot\rangle)$ then the parallelotope P spanned by the v_j's has volume $\|v_1 \wedge \dots \wedge v_d\|_d = \sqrt{\det(\langle v_i, v_j\rangle)}$.*

Proof. Let $\{e_i\}$ be an orthonormal basis, so if $T : V \simeq V$ is the linear map carrying e_i to v_i then

$$v_1 \wedge \dots \wedge v_d = T(e_1) \wedge \dots \wedge T(e_d) = \wedge^d(T)(e_1 \wedge \dots \wedge e_d) = \det(T) \cdot (e_1 \wedge \dots \wedge e_d),$$

and hence

$$\|v_1 \wedge \dots \wedge v_d\|_d = |\det T| \|e_1 \wedge \dots \wedge e_d\|_d = |\det T|$$

since $\det(\langle e_i, e_j\rangle) = \det(\mathrm{id}_d) = 1$. The parallelotope P is exactly $T(C)$ where C is the parallelotope spanned by the e_i's, so by using $\{e_i\}$ to identify V with $\mathbf{R}^d$ (respecting inner products) we carry C over to $[0,1]^d$ and so the Change of Variables formula tells us that the volume of $T(C)$ is $|\det T|$ times the volume of C, which is to say that $P = T(C)$ has volume $|\det T|$. ∎

We next consider the dual situation, as this will be essential in pseudo-Riemannian geometry. The bilinear form on V induces a natural "dual bilinear form" on $V^\vee$: we have an isomorphism $V \simeq V^\vee$ via $v \mapsto \langle v, \cdot\rangle = \langle\cdot, v\rangle$, and so for $\ell, \ell' \in V^\vee$ we *define*

$$\langle \ell, \ell'\rangle^\vee \stackrel{\text{def}}{=} \langle v, v'\rangle$$

with $\ell = \langle v, \cdot\rangle$ and $\ell' = \langle v', \cdot\rangle$ for *unique* $v, v' \in V$; the reader should check that $\langle\cdot,\cdot\rangle^\vee$ is indeed bilinear, symmetric, and non-degenerate on $V^\vee$, and when repeating the process to make a bilinear form on $V^{\vee\vee}$ we recover the initial bilinear form on V via the natural isomorphism $V \simeq V^{\vee\vee}$. To make this concrete, note that if $\{e_i\}$ is an orthonormal basis of V then its dual basis $\{e_i^*\}$ is an orthonormal basis of V because orthonormality of $\{e_i\}$ implies that

$$e_i^* = \langle e_i, e_i\rangle \cdot \langle e_i, \cdot\rangle = \langle \varepsilon_i e_i, \cdot\rangle$$

with $\varepsilon_i = \langle e_i, e_i\rangle = \pm 1$ (so moreover $\varepsilon_i^* = \langle e_i^*, e_i^*\rangle^\vee$ is equal to $\varepsilon_i^2\langle e_i, e_i\rangle = \varepsilon_i^3 = \varepsilon_i$, and hence the quadratic space $V^\vee$ has the same signature as V). By iterating the procedure from $V^\vee$ to $V^{\vee\vee} \simeq V$, we likewise see the converse: if $\{v_i\}$ is a basis of V whose dual basis is an orthonormal basis of $V^\vee$

then $\{v_i\}$ is an orthonormal basis of V. In words, a basis of a non-degenerate quadratic space is orthonormal if and only if its dual basis is orthonormal with respect to the dual symmetric bilinear form on the dual space.

The dual bilinear form on $V^\vee$ induces non-degenerate symmetric bilinear forms on $(V^\vee)^{\otimes n}$, $\mathrm{Sym}^n(V^\vee)$, and $\wedge^n(V^\vee)$ for all n, by Lemma 2.1 applied to $V^\vee$ and $B = \langle \cdot, \cdot \rangle^\vee$. There now arises a very important compatibility question, due to the fact that there is an entirely different natural method to get quadratic structures on these spaces: we have non-degenerate symmetric bilinear forms on $V^{\otimes n}$, $\mathrm{Sym}^n(V)$, and $\wedge^n(V)$ via Lemma 2.1 for V and $B = \langle \cdot, \cdot \rangle$, and so the recipe just given for dualizing general non-degenerate quadratic spaces may be applied to each of these spaces to endow their duals $(V^{\otimes n})^\vee$, $(\mathrm{Sym}^n(V))^\vee$, and $(\wedge^n(V))^\vee$ with non-degenerate symmetric bilinear forms — but in §3 in the handout on tensor algebras we saw that these duals are naturally identified with $(V^\vee)^{\otimes n}$, $\mathrm{Sym}^n(V^\vee)$, and $\wedge^n(V^\vee)$ respectively! Consequently, we obtain a second method for constructing quadratic structures on these latter spaces. Does it agree with what we get by passing to tensor, symmetric, and exterior powers on the dual bilinear form made on $V^\vee$? We now briefly digress to solve this problem in the affirmative before going further.

Theorem 2.5. *The two methods for putting non-degenerate symmetric bilinear forms on tensor, symmetric, and exterior powers of $V^\vee$ coincide. That is, the natural isomorphisms*

$$(V^\vee)^{\otimes n} \simeq (V^{\otimes n})^\vee, \;\; \mathrm{Sym}^n(V^\vee) \simeq (\mathrm{Sym}^n(V))^\vee, \;\; \wedge^n(V^\vee) \simeq (\wedge^n(V))^\vee$$

carry the tensor, symmetric, and exterior powers of the symmetric bilinear form on $V^\vee$ over to the dual of the tensor, symmetric, and exterior powers of the symmetric bilinear form on V.

Proof. This is a problem of chasing orthonormal bases, and the case of symmetric powers will require a little care due to the annoying factors $m_j(I)!$. Take $\{e_i\}$ to be an orthonormal basis of V, so the tensor and wedge products of the e_i's give orthonormal bases of tensor and exterior powers of V. The dual basis $\{e_i^*\}$ is an orthonormal basis of $V^\vee$ with $\varepsilon_i^* = \langle e_i^*, e_i^* \rangle^\vee$ equal to $\varepsilon_i = \langle e_i, e_i \rangle$ for all i, so the tensor and wedge products of the e_i^*'s give orthonormal bases of $(V^\vee)^{\otimes n}$ and $\wedge^n(V^\vee)$ and under the natural isomorphisms with $(V^{\otimes n})^\vee$ and $(\wedge^n(V))^\vee$ these bases are dual to the tensor and wedge products of the e_i's (by Corollary 2.3). Since a basis of a non-degenerate quadratic space is orthonormal if and only if its dual basis is orthonormal with respect to the dual symmetric bilinear form on the dual space, and moreover $\varepsilon_i^* = \varepsilon_i$ for all i, we obtain the desired agreement in the case of tensor and exterior powers (as two symmetric bilinear forms agree if they share a common orthogonal basis with the same self-pairings on the basis vectors).

For the case of nth symmetric powers, we let $I = \{i_1, \dots, i_n\}$ be a monotone sequence of integers between 1 and d, define $m(I) = \prod_{j=1}^d m_j(I)!$, and define

$$e_I = e_{i_1} \cdot \dots \cdot e_{i_n}, \;\; e_{I^*} = e_{i_1}^* \cdot \dots \cdot e_{i_n}^*.$$

We write $\{e_I^*\}$ for the dual basis to $\{e_I\}$ in $(\mathrm{Sym}^n(V))^\vee$. The isomorphism $\mathrm{Sym}^n(V^\vee) \simeq (\mathrm{Sym}^n(V))^\vee$ carries e_{I^*} to $m(I)e_I^*$, by Remark 2.2. We also know that the e_{I^*}'s are pairwise orthogonal (as the e_i^*'s are so in $V^\vee$) and the e_I^*'s are pairwise orthogonal (as they are dual to the pairwise orthogonal vectors e_I in $\mathrm{Sym}^n(V)$). Thus, the problem is to compare lengths and signs of self-pairings: we want the length of e_{I^*} to be $m(I)$ times the length of e_I^* and the sign of the self-pairings of e_{I^*} and e_I^* to coincide. The latter aspect follows from Corollary 2.3, so the problem is one of lengths. By Corollary 2.3, e_I has length $\sqrt{m(I)}$ (as $e_I/\sqrt{m(I)}$ is a unit vector), and since $\{e_i^*\}$ is an orthonormal basis of $V^\vee$ we likewise conclude that e_{I^*} also has length $\sqrt{m(I)}$. Thus, we want the dual-basis vector e_I^* to have length $1/\sqrt{m(I)}$, or equivalently to have absolute self-pairing $1/m(I)$. Since e_I has absolute self-pairing $m(I)$ and the e_I's are pairwise orthogonal, our problem is a special case of the general claim that if W is a vector space endowed with a non-degenerate symmetric bilinear form and $\{w_j\}$ is an orthogonal basis then the dual basis $\{w_j^*\}$ of $W^\vee$ is orthogonal with

$$\langle w_j^*, w_j^* \rangle^\vee = \frac{1}{\langle w_j, w_j \rangle}.$$

If we let $c_j = \langle w_j, w_j \rangle \neq 0$ then orthogonality of the basis of W implies $w_j^* = \langle w_j/c_j, \cdot \rangle$. Hence, by definition of $\langle \cdot, \cdot \rangle^\vee$ we have

$$\langle w_j^*, w_k^* \rangle^\vee = \langle w_j/c_j, w_k/c_k \rangle = \frac{\langle w_j, w_k \rangle}{c_j c_k},$$

and this vanishes if $j \neq k$ but equals $1/c_j$ if $j = k$. ■

Now that we are assured that there is no ambiguity concerning the natural quadratic structure on tensor, symmetric, and exterior powers of $V^\vee$, we turn to an item that will be of much importance in the theory of volume on pseudo-Riemannian manifolds, and that is the study of $\wedge^d(V^\vee)$ with $d = \dim V > 0$.

Theorem 2.6. *Let $(V, \langle \cdot, \cdot \rangle)$ be a nonzero non-degenerate quadratic space and let $\mathbf{v} = \{v_i\}$ be an ordered basis of V with corresponding dual basis $\{v_i^*\}$ in $V^\vee$. The vector*

$$\omega_{\mathbf{v}} = \sqrt{|\det(\langle v_i, v_j \rangle)|}\, v_1^* \wedge \cdots \wedge v_d^* \in \wedge^d(V^\vee)$$

is a unit vector for the natural quadratic structure on the line $\wedge^d(V^\vee)$, and if $\mathbf{v}'$ is a second ordered basis of V then $\omega_{\mathbf{v}'} = \varepsilon \omega_{\mathbf{v}}$ with $\varepsilon = \pm 1$ equal to the sign of $\det T_{\mathbf{v},\mathbf{v}'}$ where $T_{\mathbf{v},\mathbf{v}'} : V \simeq V$ is the map satisfying $v_i \mapsto v_i'$.

Concretely, this theorem associates a unit vector in the line $\wedge^d(V^\vee)$ to any ordered basis of V, and two ordered bases get assigned to the same unit vector if and only if the linear automorphism of V (or "change of basis matrix") relating them (in either direction!) has positive determinant. The element $\omega_{\mathbf{v}}$ is generally called the *volume form* associated to $\mathbf{v}$, for reasons that will become clear in the applications to differential geometry, and we see that if the quadratic space V is *oriented* then the orientation μ picks out a preferred volume form, namely the unique one in the positive half-line for the dual orientation (i.e., it is $\omega_{\mathbf{v}}$ for any μ-positive ordered basis $\mathbf{v}$).

Proof. Let us first prove the unit-vector claim by pure thought, rather than by reducing it to a calculation in the case of an orthonormal basis. We know that $\sqrt{|\det(\langle v_i, v_j \rangle)|}$ is the length of the vector $v_1 \wedge \cdots \wedge v_d$ in $\wedge^d(V)$, and so the unit-vector assertion is the statement that the vectors

$$v_1 \wedge \cdots \wedge v_d \in \wedge^d(V), \quad v_1^* \wedge \cdots \wedge v_d^* \in \wedge^d(V^\vee)$$

have reciprocal lengths. We have already seen that the natural isomorphism $\wedge^d(V^\vee) \simeq (\wedge^d(V))^\vee$ carries the exterior power of $\langle \cdot, \cdot \rangle^\vee$ on $V^\vee$ over to the dual of the exterior power of $\langle \cdot, \cdot \rangle$ on V. This isomorphism also carries $v_1^* \wedge \cdots \wedge v_d^*$ to the linear functional on $\wedge^d(V)$ dual to the basis $v_1 \wedge \cdots \wedge v_d$ of the line $\wedge^d(V)$. Hence, our problem is in the theory of 1-dimensional quadratic spaces over $\mathbf{R}$: if (L, q) is a 1-dimensional non-degenerate quadratic space with basis vector $v \in L$ then with respect to the dual quadratic structure $q^\vee$ on $V^\vee$ (corresponding to the dual symmetric bilinear form) the length $\sqrt{|q^\vee(v^*)|}$ of the dual vector $v^* \in L^\vee$ determined by the condition $v^*(v) = 1$ is reciprocal to the length $\sqrt{|q(v)|}$ of v. Even better, we have a reciprocal relationship without absolute values and square roots: $q(v) q^\vee(v^*) = 1$. This was seen at the end of the proof of Theorem 2.5.

It remains to work out the dependence of $\omega_{\mathbf{v}}$ on $\mathbf{v}$. If $\mathbf{v}' = \{v'_i\}$ is another ordered basis and $v'_i = T(v_i)$ for all i then the dual bases are related by $v'^*_i = (T^\vee)^{-1}(v^*_i) = (T^{-1})^\vee(v^*_i)$ because

$$((T^{-1})^\vee(v^*_i))(v'_j) = v^*_i(T^{-1}(v'_j)) = v^*_i(v'_j)$$

vanishes when $i \neq j$ and equals 1 when $i = j$. Thus,

$$v_1'^* \wedge \cdots \wedge v_d'^* = \det(T^\vee)^{-1} v_1^* \wedge \cdots \wedge v_d^* = \det(T)^{-1} v_1^* \wedge \cdots \wedge v_d^*,\ \ (\langle v'_i, v'_j \rangle) = [T]^{\mathrm{t}}(\langle v_i, v_j \rangle)[T]$$

with $[T] = {}_{\mathbf{v}}[T]_{\mathbf{v}'}$ the change of basis matrix for T (from $\mathbf{v}'$-coordinates to $\mathbf{v}$-coordinates). By passing to determinants and using that transpose does not affect determinant, we get

$$\omega_{\mathbf{v}'} = \sqrt{|\det(\langle v'_i, v'_j \rangle)|} v_1'^* \wedge \cdots \wedge v_d'^* = \frac{\sqrt{|\det(T)|^2}}{\det(T)} \omega_{\mathbf{v}} = \varepsilon \omega_{\mathbf{v}}$$

with ε equal to the sign of $\det(T)$. ■

3. Orientations of C-vector spaces

We now turn to an example that is very important in the theory of complex manifolds, and contains a few subtle points. Considering $\mathbf{C}$ as an $\mathbf{R}$-vector space, we can identify the two orientations with the two choices of square root of -1 in $\mathbf{C}$. Indeed, if i denotes such a square root (so $-i$ is the other one) then the two $\mathbf{R}$-bases $\{1, i\}$ and $\{1, -i\}$ of $\mathbf{C}$ are related by an automorphism whose determinant is -1 (concretely, $1 \wedge (-i) = -(1 \wedge i)$ in $\wedge^2_{\mathbf{R}}(\mathbf{C})$), so these represent the two equivalence classes of $\mathbf{R}$-bases. We write μ_i to denote the orientation of $\mathbf{C}$ (as an $\mathbf{R}$-vector space) determined by the ordered $\mathbf{R}$-basis $\{1, i\}$ (i.e., it is the connected component of $\wedge^2_{\mathbf{R}}(\mathbf{C}) - \{0\}$ containing $1 \wedge i$). Of course, neither of these two orientations $\mu_{\pm i}$ is "better" than the other since there is no natural square root of -1 in $\mathbf{C}$.

Another way to think about the situation is as follows. The field $\mathbf{C}$ over $\mathbf{R}$ has a unique nontrivial field automorphism $\sigma : z \mapsto \overline{z}$ fixing $\mathbf{R}$. (To see the uniqueness, observe that the set of solutions to $X^2 + 1 = 0$ must get permuted by such an automorphism, yet knowing the action of such an automorphism on this set of roots determines it uniquely because of the description of $\mathbf{C}$ from high school.) By viewing this field automorphism as an $\mathbf{R}$-linear automorphism of $\mathbf{C}$, we see that the points fixed by σ form a line $L^+ \subseteq \mathbf{C}$ (namely, $L^+ = \mathbf{R}$) and the points negated by σ form a line $L^- \subseteq \mathbf{C}$ (namely, the "imaginary axis"). We have $\mathbf{C} = L^+ \oplus L^-$ as $\mathbf{R}$-vector spaces, and $L^+ = \mathbf{R}$ is canonically oriented using the $\mathbf{R}$-basis 1. However, L^- is not canonically oriented, and we may distinguish the two components of $L^- - \{0\}$ by the unique square root of -1 in each. Thus, choosing $\sqrt{-1} \in \mathbf{C}$ gives an orientation to L^- (by declaring the positive component of $L^- - \{0\}$ to be the one containing the choice of $\sqrt{-1}$) and $\mu_{\sqrt{-1}}$ is the resulting direct sum orientation on $\mathbf{C}$ when the lines L^+ and L^- are considered as an ordered pair via the rule "L^+ first, L^- second".

Lemma 3.1. *Let $T : V \simeq V$ be a $\mathbf{C}$-linear automorphism of a nonzero $\mathbf{C}$-vector space, and let $T_{\mathbf{R}}$ be the $\mathbf{R}$-linear automorphism of the underlying $\mathbf{R}$-vector space $V_{\mathbf{R}}$. Then* $\det T_{\mathbf{R}} = |\det T|^2 > 0$.

Proof. We wish to induct on $\dim_{\mathbf{C}} V$. First suppose $\dim_{\mathbf{C}} V = 1$, so T is multiplication by some $\lambda \in \mathbf{C}^\times$. We want $\det T_{\mathbf{R}} = |\lambda|^2 = \lambda\overline{\lambda}$. This is clear if $\lambda \in \mathbf{R}^\times$, so suppose otherwise. Hence, $\{1, \lambda\}$ is an $\mathbf{R}$-basis of $\mathbf{C}$ and λ is a root of

$$(X - \lambda)(X - \overline{\lambda}) = X^2 - aX + b \in \mathbf{R}[X]$$

with $a = \lambda + \overline{\lambda}$ and $b = \lambda\overline{\lambda}$. Using the $\mathbf{R}$-basis $\mathbf{e} = \{1, \lambda\}$ of $\mathbf{C}$, $T_{\mathbf{R}}$ has matrix

$${}_{\mathbf{e}}[T_{\mathbf{R}}]_{\mathbf{e}} = \begin{pmatrix} 0 & -b \\ 1 & -a \end{pmatrix},$$

so $\det T_{\mathbf{R}} = b = \lambda\overline{\lambda}$. (Alternatively, since $\det T_{\mathbf{R}}$ is the scaling effect of $\wedge^2(T_{\mathbf{R}})$ on $\wedge^2_{\mathbf{R}}(\mathbf{C})$ and $1 \wedge \lambda$ is nonzero in this line, we compute the multiplier via the equalities

$$\wedge^2(T_{\mathbf{R}})(1 \wedge \lambda) = T_{\mathbf{R}}(1) \wedge T_{\mathbf{R}}(\lambda) = \lambda \wedge \lambda^2 = \lambda \wedge (-a\lambda - b) = -a\lambda \wedge \lambda - b\lambda \wedge 1 = b(1 \wedge \lambda),$$

so $\det T_{\mathbf{R}} = b = \lambda\overline{\lambda}$.)

Now assume $\dim_{\mathbf{C}} V > 1$ and that the result is known in smaller dimensions. Suppose that W is a nonzero proper subspace of V that is T-stable, and let $T_W : W \to W$ and $\overline{T} : V/W \to V/W$ be induced by T; one such W is a line spanned by an eigenvector of T (which exists, by the Fundamental Theorem of Algebra). Since $\dim_{\mathbf{C}} W, \dim_{\mathbf{C}} V/W < \dim_{\mathbf{C}} V$, by induction the result is known for $\mathbf{C}$-linear self-maps of W and V/W. We have $\det T = \det T_W \det \overline{T}$. Clearly $T_{\mathbf{R}}$ preserves $W_{\mathbf{R}}$ and it restricts to $(T_W)_{\mathbf{R}}$ on $W_{\mathbf{R}}$ and induces $\overline{T}_{\mathbf{R}}$ on $(V/W)_{\mathbf{R}}$ via the obvious identification of $(V/W)_{\mathbf{R}}$ with $V_{\mathbf{R}}/W_{\mathbf{R}}$. Thus, $\det T_{\mathbf{R}} = \det(T_W)_{\mathbf{R}} \det \overline{T}_{\mathbf{R}} = |\det T_W|^2 |\det \overline{T}|^2 = |\det T|^2$, as desired. ∎

Let us now exploit this positivity result. Fix an orientation $\mu = \mu_{\sqrt{-1}}$ on $\mathbf{C}$ once and for all. By Example 1.6, for all $n \geq 1$ we get an n-fold direct sum orientation $\mu^{\oplus n} = \mu \oplus \cdots \oplus \mu$ on $(\mathbf{C}^n)_{\mathbf{R}} \simeq (\mathbf{C}_{\mathbf{R}})^n$. For any nonzero complex vector space V, there exists a $\mathbf{C}$-linear isomorphism $T : V \simeq \mathbf{C}^n$, and so by using the underlying $\mathbf{R}$-linear isomorphism $T_{\mathbf{R}}$ we may transfer the orientation $\mu^{\oplus n}$ on $(\mathbf{C}^n)_{\mathbf{R}}$ to an orientation μ_T on $V_{\mathbf{R}}$. The miracle is:

Theorem 3.2. *The orientation μ_T on $V_{\mathbf{R}}$ is* independent *of the choice of $\mathbf{C}$-linear isomorphism $T : V \simeq \mathbf{C}^n$. Writing μ_V to denote the resulting orientation on $V_{\mathbf{R}}$ (that depends only on $\mu = \mu_{\sqrt{-1}}$ and the given $\mathbf{C}$-structure on $V_{\mathbf{R}}$ provided by V), if $L : V \simeq V'$ is a $\mathbf{C}$-linear isomorphism then $L_{\mathbf{R}}$ is orientation-preserving with respect to the orientations μ_V and $\mu_{V'}$ on $V_{\mathbf{R}}$ and $V'_{\mathbf{R}}$ respectively.*

We call the orientation μ_V on $V_{\mathbf{R}}$ the *canonical orientation* (arising from the initial choice of μ on $\mathbf{C}$ and the given $\mathbf{C}$-linear structure on $V_{\mathbf{R}}$ provided by V).

Proof. Any $\mathbf{C}$-linear isomorphisms $T : V \simeq \mathbf{C}^n$ and $T' : V \simeq \mathbf{C}^n$ satisfy $T' = T \circ \phi$ for a $\mathbf{C}$-linear automorphism $\phi = T^{-1} \circ T'$ of V, and by the lemma $\det \phi_{\mathbf{R}} > 0$. Hence, $\phi_{\mathbf{R}}$ is orientation-preserving on $V_{\mathbf{R}}$ with respect to any fixed orientation of $V_{\mathbf{R}}$, yet the identity $T'_{\mathbf{R}} = T_{\mathbf{R}} \circ \phi_{\mathbf{R}}$ implies that $\phi_{\mathbf{R}}$ interchanges μ_T and $\mu_{T'}$. Thus, $\mu_T = \mu_{T'}$.

If $L : V \simeq V'$ is a $\mathbf{C}$-linear isomorphism between nonzero $\mathbf{C}$-vector spaces, then for a $\mathbf{C}$-linear isomorphism $T' : V \simeq \mathbf{C}^n$ we have a $\mathbf{C}$-linear isomorphism $T' \circ L : V \simeq \mathbf{C}^n$, so by definition $\mu_{T'}$ is carried to $\mu_{T' \circ L}$ by $L_{\mathbf{R}}$. That is, $L_{\mathbf{R}}$ respects $\mu_{V'}$ and μ_V. ∎

This is a remarkable state of affairs: simply by choosing $\sqrt{-1} \in \mathbf{C}$ all nonzero $\mathbf{C}$-vector spaces are endowed with a canonical orientation (on their underlying $\mathbf{R}$-vector spaces) that is respected by all $\mathbf{C}$-linear isomorphisms (viewed as $\mathbf{R}$-linear isomorphisms). One typically sees this fact described by the phrase that "all $\mathbf{C}$-vector spaces are canonically oriented", with the dependence on the initial choice of $\sqrt{-1} \in \mathbf{C}$ suppressed. Since most mathematicians (including some Fields medalists) have the mistaken belief that there is a preferred square root of -1 in $\mathbf{C}$ (they're wrong), most believe that there is a God-given orientation on the $\mathbf{R}$-vector space underlying any $\mathbf{C}$-vector space (though depending on the $\mathbf{C}$-linear structure). This is slightly incorrect: we first have to pick $\sqrt{-1} \in \mathbf{C}$, and then everything else is canonically determined. To be completely explicit, if $\{v_1, \dots, v_d\}$ is a $\mathbf{C}$-basis of V, then μ_V is the orientation class of the $\mathbf{R}$-basis

$$\{v_1, \sqrt{-1} \cdot v_1, \dots, v_d, \sqrt{-1} \cdot v_d\}$$

of $V_{\mathbf{R}}$. (Be sure you see why this really is an $\mathbf{R}$-basis, and that this description is what comes out of the preceding considerations.)

Example 3.3. Let V and V' be nonzero $\mathbf{C}$-vector spaces, so $V_{\mathbf{R}}$ and $V'_{\mathbf{R}}$ are even-dimensional over $\mathbf{R}$. There is a canonical isomorphism $(V \oplus V')_{\mathbf{R}} \simeq V_{\mathbf{R}} \oplus V'_{\mathbf{R}}$, and it is easy to check that under this isomorphism $\mu_{V \oplus V'}$ goes over to the direct sum orientation $\mu_V \oplus \mu_{V'}$ (which does not depend on the ordering of the set $\{V_{\mathbf{R}}, V'_{\mathbf{R}}\}$ since these spaces are even-dimensional). In this sense, the canonical orientation (for a fixed choice of $\sqrt{-1} \in \mathbf{C}$ is compatible with the formation of direct sums.

Example 3.4. Let V be a nonzero $\mathbf{C}$-vector space and $W \subseteq V$ a nonzero proper subspace. The three nonzero $\mathbf{C}$-vector spaces V, W, and V/W has corresponding canonical orientations μ_V, μ_W, and $\mu_{V/W}$, and via the evident inclusion of $W_{\mathbf{R}}$ into $V_{\mathbf{R}}$ and the isomorphism $V_{\mathbf{R}}/W_{\mathbf{R}} \simeq (V/W)_{\mathbf{R}}$ it makes sense to ask if μ_V, μ_W, and $\mu_{V/W}$ are compatible (in the sense of Example 1.7). They are, and the proof is a matter of basis-chasing that we leave to the reader (keeping in mind how μ_V is defined by using a $\mathbf{C}$-basis of V and the choice of $\sqrt{-1} \in \mathbf{C}$).

Example 3.5. Let V be a nonzero $\mathbf{C}$-vector space and let $W_1, \dots, W_N \subseteq V$ be proper nonzero subspaces that are mutually transverse (either in the $\mathbf{R}$-linear or $\mathbf{C}$-linear sense – why are these equivalent concepts?), and assume $W' = \cap W_j$ is nonzero. Clearly $W'_{\mathbf{R}} = \cap (W_j)_{\mathbf{R}}$, so the canonical orientations on $V_{\mathbf{R}}$ and the $(W_j)_{\mathbf{R}}$'s determine a preferred orientation on the transverse intersection $W'_{\mathbf{R}}$ in accordance with Example 1.8. Does this agree with the canonical orientation $\mu_{W'}$? Indeed it does, and this again is left as an exercise in basis-chasing.

Let us conclude our investigation of the $\mathbf{C}$-linear case by working out the "dual" of the canonical orientation, as this presents a mild surprise. Let V be a $\mathbf{C}$-vector space, and let $V^\vee$ be its $\mathbf{C}$-linear dual; the underlying $\mathbf{R}$-vector space of $V^\vee$ is *not* literally the $\mathbf{R}$-linear dual of $V_{\mathbf{R}}$, as elements of the former are $\mathbf{C}$-linear functionals on V and elements of the latter are $\mathbf{R}$-linear functions on $V_{\mathbf{R}}$ (with values in $\mathbf{R}$, not $\mathbf{C}$). However, these two duals can be identified:

Lemma 3.6. *Let $t : \mathbf{C} \to \mathbf{R}$ be the $\mathbf{R}$-linear "trace" defined by $z \mapsto z + \overline{z}$. The $\mathbf{R}$-linear map $(V^\vee)_{\mathbf{R}} \to (V_{\mathbf{R}})^\vee$ defined by $\ell \mapsto t \circ \ell_{\mathbf{R}}$ is an isomorphism.*

Proof. The map in question is certainly $\mathbf{R}$-linear, and both sides have the same $\mathbf{R}$-dimension (namely, twice the $\mathbf{C}$-dimension of V), and hence it suffices to check injectivity. That is, if $t \circ \ell_{\mathbf{R}} = 0$ then we want $\ell = 0$. If $t \circ \ell_{\mathbf{R}} = 0$ then $\ell_{\mathbf{R}}$ takes values in the imaginary axis. However, $\ell : V \to \mathbf{C}$ is a $\mathbf{C}$-linear map and so if it is not zero then it is surjective and hence cannot have image contained in the imaginary axis. This only leaves the option $\ell = 0$. ■

Fixing an orientation $\mu = \mu_{\sqrt{-1}}$ on $\mathbf{C}$, we get canonical orientations μ_V and $\mu_{V^\vee}$ on $V_{\mathbf{R}}$ and $(V^\vee)_{\mathbf{R}}$ for any nonzero $\mathbf{C}$-vector space V. The lemma permits us to identify $(V^\vee)_{\mathbf{R}}$ with $(V_{\mathbf{R}})^\vee$, so it makes sense to ask how $\mu_{V^\vee}$ and $\mu_V^\vee$ are related. The answer turns out to be not what one might expect:

Theorem 3.7. *Let V be a finite-dimensional nonzero $\mathbf{C}$-vector space with $N = \dim_{\mathbf{C}} V > 0$. Via the natural isomorphism between $(V^\vee)_{\mathbf{R}}$ and $(V_{\mathbf{R}})^\vee$, $\mu_{V^\vee}$ goes over to $(-1)^N \mu_V^\vee$.*

The moral is that for *even-dimensional* $\mathbf{C}$-vector spaces there is no confusion concerning the meaning of "dual orientation", but in the odd-dimensional case one has to make clear whether one is $\mathbf{R}$-dualizing μ_V or one is working with the canonical orientation associated to the $\mathbf{C}$-linear dual.

Proof. One checks that the isomorphism $(V^\vee)_{\mathbf{R}} \simeq (V_{\mathbf{R}})^\vee$ is naturally compatible with direct sums in the $\mathbf{C}$-vector space V, and likewise for an ordered pair of oriented $\mathbf{R}$-vector spaces (W, ν) and (W', ν') we have $\nu^\vee \oplus \nu'^\vee = (\nu \oplus \nu')^\vee$ via the natural isomorphism $W^\vee \oplus W'^\vee \simeq (W \oplus W')^\vee$. Also, for an ordered m-tuple of oriented $\mathbf{R}$-vector spaces (W_j, ν_j), on the direct sum $\oplus W_j$ the orientation

$\oplus(-\nu_j)$ is off by a sign of $(-1)^m$ from $\oplus\nu_j$. Hence, by additivity of $\mathbf{C}$-dimension on direct sums we conclude that the problem is compatible with formation of direct sums in V and that it is unaffected by passing to a $\mathbf{C}$-linearly isomorphic space. Since V is $\mathbf{C}$-linearly isomorphic to a direct sum of copies of $\mathbf{C}$, we may therefore reduce to the case $V = \mathbf{C}$. In this case, a positive basis of $V_{\mathbf{R}}$ is $\{1, \sqrt{-1}\}$, and so the positive dual basis of $(V_{\mathbf{R}})^\vee$ is $\{x, y\}$ with the $\mathbf{R}$-linear maps $x, y : \mathbf{C} \rightrightarrows \mathbf{R}$ computing real and imaginary parts: $\lambda = x(\lambda)\cdot 1 + y(\lambda)\sqrt{-1}$ for $\lambda \in \mathbf{C}$.

Meanwhile, the $\mathbf{C}$-linear dual $\mathbf{C}^\vee$ admits as a singleton $\mathbf{C}$-basis the identity map $\mathbf{1} : \mathbf{C} \to \mathbf{C}$, and so the resulting $\mathbf{C}$-linear isomorphism $\mathbf{C}^\vee \simeq \mathbf{C}$ (associating $\lambda \in \mathbf{C}$ to $\lambda \cdot \mathbf{1} \in \mathbf{C}^\vee$) gives the canonical orientation on $(\mathbf{C}^\vee)_{\mathbf{R}}$ as $\{1\cdot\mathbf{1}, \sqrt{-1}\cdot\mathbf{1}\} = \{\mathbf{1}, \sqrt{-1}\cdot\mathbf{1}\}$. The isomorphism $(\mathbf{C}^\vee)_{\mathbf{R}} \simeq (\mathbf{C}_{\mathbf{R}})^\vee$ is composition with the trace $t : \mathbf{C} \to \mathbf{R}$, and so it carries $\{\mathbf{1}, \sqrt{-1}\cdot\mathbf{1}\}$ to $\{t, t(\sqrt{-1}(\cdot))\}$. Hence, our problem is to show that the ordered $\mathbf{R}$-bases $\{x, y\}$ and $\{t, t(\sqrt{-1}(\cdot))\}$ of $\mathrm{Hom}_{\mathbf{R}}(\mathbf{C}, \mathbf{R})$ are oppositely oriented. For $\lambda = a + b\sqrt{-1}$ with $a, b \in \mathbf{R}$ we compute $t(\lambda) = \lambda + \overline{\lambda} = 2a$ and $t(\sqrt{-1}\cdot\lambda) = t(-b + a\sqrt{-1}) = -2b$, so $t = 2x$ and $t(\sqrt{-1}(\cdot)) = -2y$. The change of basis matrix going between $\{x, y\}$ and $\{2x, -2y\}$ is the diagonal matrix with diagonal entries 2 and -2, so its determinant is $-4 < 0$. This gives the required sign discrepancy. ■

Now let V be a nonzero $\mathbf{C}$-vector space with $\dim_{\mathbf{C}} V = d$. Using our initial choice of $\sqrt{-1}$, any $\mathbf{C}$-linear functional $\ell : V \to \mathbf{C}$ can be uniquely decomposed into real and imaginary parts, which is to say that we can uniquely write $\ell = \ell_+ + \ell_-\sqrt{-1}$ with $\mathbf{R}$-linear functionals $\ell_+, \ell_- : V_{\mathbf{R}} \rightrightarrows \mathbf{R}$. (In particular, if we negate $\sqrt{-1}$ then the "imaginary part" ℓ_- of ℓ is negated but the "real part" ℓ_+ of ℓ is unchanged.) The $\mathbf{C}$-linearity condition on ℓ imposes a link between ℓ_+ and ℓ_- through the identity

$$\ell_+(\sqrt{-1}\cdot v) + \ell_-(\sqrt{-1}\cdot v)\sqrt{-1} = \ell(\sqrt{-1}\cdot v) = \sqrt{-1}\cdot\ell(v) = -\ell_-(v) + \ell_+(v)\sqrt{-1},$$

or equivalently $\ell_-(v) = -\ell_+(\sqrt{-1}\cdot v)$. Both ℓ_+ and ℓ_- may be viewed as $\mathbf{R}$-linear functionals on V, which is to say that they are elements of $(V_{\mathbf{R}})^\vee$.

Fix an ordered $\mathbf{C}$-basis $\{v_1, \dots, v_d\}$ of V, and let $\{z_1, \dots, z_d\}$ be the resulting ordered $\mathbf{C}$-dual basis of the $\mathbf{C}$-linear dual $V^\vee$. Let $z_j = x_j + y_j\sqrt{-1}$ be the decomposition of the $\mathbf{C}$-dual basis elements z_j with $x_j, y_j \in (V_{\mathbf{R}})^\vee$. The proof of Theorem 3.7 shows that the $\mathbf{R}$-dual orientation on $(V_{\mathbf{R}})^\vee$ (dual to the canonical orientation μ_V on $V_{\mathbf{R}}$ arising from the $\mathbf{C}$-structure on V and the initial choice of $\sqrt{-1}$) is represented by the ordered basis $\{x_1, y_1, \dots, x_d, y_d\}$, whereas the canonical orientation arising from the $\mathbf{C}$-structure on $V^\vee$ and the trace-isomorphism $(V^\vee)_{\mathbf{R}} \simeq (V_{\mathbf{R}})^\vee$ is represented by $\{x_1, -y_1, \dots, x_d, -y_d\}$. This explicitly exhibits the sign discrepancy of $(-1)^d$, and so for odd d one needs to be very careful to set down one's convention for "dual orientation" in the setting of $\mathbf{C}$-vector spaces.

Remark 3.8. If V is a nonzero $\mathbf{C}$-vector space with $\dim_{\mathbf{C}} V = d$, how does the canonical orientation μ_V change when we change the initial choice of $\sqrt{-1} \in \mathbf{C}^\times$? Changing this choice carries the ordered $\mathbf{R}$-basis $\{1, \sqrt{-1}\}$ of $\mathbf{C}$ to the ordered $\mathbf{R}$-basis $\{1, -\sqrt{-1}\}$ that is in the opposite orientation class, and so $\mu_{\mathbf{C}}$ is replaced with $-\mu_{\mathbf{C}}$. Since μ_V is obtained from $\mu_{\mathbf{C}}^{\oplus d}$ for any $\mathbf{C}$-linear isomorphism $V \simeq \mathbf{C}^d$, the identity $(-\mu_{\mathbf{C}})^{\oplus d} = (-1)^d \mu_{\mathbf{C}}^{\oplus d}$ shows that μ_V is changed by a sign of $(-1)^d$ when we change the initial choice of $\sqrt{-1}$. Hence, for *even-dimensional* $\mathbf{C}$-vector spaces V there truly is a canonical orientation of $V_{\mathbf{R}}$ that depends on the given $\mathbf{C}$-linear structure on V but is otherwise independent of all choices. We have likewise seen above that the notion of "dual orientation" in the $\mathbf{C}$-linear setting is especially well-behaved in the even-dimensional (over $\mathbf{C}$) case. Thus, for even-dimensional $\mathbf{C}$-vector spaces there is an particularly pleasant theory of orientations, whereas for odd-dimensional $\mathbf{C}$-vector spaces there is a rather good theory but one does need to pay attention to the choice of $\sqrt{-1}$ and any intervention of dualities.