LECTURE

Covectors as Geometric Objects

Labeled Hyperplane Foliations, Dual Coordinates, and Support Functions

Antonio Franchi

University of Twente and
Sapienza University of Rome

November 2026

Ver. 0.3.0

Guiding question

What geometric object is a covector before coordinates represent it by an array and before an inner product associates it with a vector?

Central message. A nonzero covector is a linear scalar observation. Geometrically, it is a parallel family of affine hyperplanes together with a linear labeling of the leaves. A basis records the covector through finitely many labels, and the dual basis isolates the coordinate contributions of vectors. This labeled-foliation picture also makes the support function of a set transparent: the support value is the highest covector label attained on the set, and the corresponding leaf is a supporting hyperplane. After an inner product is chosen, the same covector may additionally be represented by a normal vector.

Use, attribution, and disclaimer

Educational use and attribution

These notes may be used, copied, shared, and adapted for personal study, teaching, and other non-commercial educational purposes. No prior permission is required for such use, provided that the source is acknowledged and that modified versions are not presented as the author's original text. Suggested citation:

> Antonio Franchi, *Covectors as Geometric Objects: Labeled Hyperplane Foliations, Dual Coordinates, and Support Functions*, November 2026.

Accuracy and responsibility

These notes are educational material and may contain typographical errors, mathematical inaccuracies, incomplete arguments, or statements requiring additional hypotheses. Readers are encouraged to verify definitions, derivations, and proofs independently before relying on them in research, teaching, software, engineering applications, or other professional work.

Feedback and corrections

Questions, corrections, and suggestions are welcome. When reporting a possible error, please indicate the relevant section, equation, proposition, or page number.

Author: Antonio Franchi

Affiliations: University of Twente and Sapienza University of Rome

Document date: November 2026

Contents

Use, attribution, and disclaimer 1

Prerequisites and scope 4

Learning objectives and conceptual map 4

1 Recommended references 6

Recommended references 6

2 Methodological principle and ambient structure 7

3 Representing a map inside its domain 7
3.1 Maps as labeled partitions . . . 7
3.2 Real-valued maps and level sets . . . 7

4 From linear scalar observation to affine foliation 7
4.1 Formal definition . . . 7
4.2 The kernel has codimension one . . . 8
4.3 All fibers are parallel affine hyperplanes . . . 8

5 The complete geometric content of a covector 8
5.1 The kernel determines the leaf direction . . . 9
5.2 The linear labeling completes the foliation . . . 9
5.3 The equivalent nested-half-space representation . . . 9

6 Transverse parameterization and quotient geometry 9
6.1 A transverse vector . . . 9
6.2 Normalization without a norm . . . 10
6.3 Reconstructing a covector from geometric data . . . 10

7 A basis gives finite access to the complete covector 10
7.1 Recovering the kernel from the basis labels . . . 11

8 The dual basis as coordinate-isolating affine foliations 12
8.1 Recovering a vector as the intersection of its coordinate leaves . . . 13
8.2 Why the dual covectors form a basis of V^* . . . 14
8.3 Basis evaluations and dual coordinates are the same numbers . . . 14

9 Scalar multiplication as geometric relabeling 15
9.1 Positive scaling . . . 15
9.2 Negative scaling . . . 15
9.3 Zero scaling . . . 15

10 Addition as pointwise label addition and geometric shear 15
10.1 Pointwise label addition . . . 16
10.2 Linearly dependent covectors . . . 16
10.3 Linearly independent covectors . . . 16
10.4 Shear interpretation . . . 16
10.5 Recovering the complete geometry of the sum . . . 17

11 Metric-free audit **17**

12 Application: support functions and supporting hyperplanes **18**
12.1 Why support functions are useful 18
12.2 The supporting leaf 18
12.3 Set inclusion as support-function comparison 19
12.4 Basic consequences of the definition 20

13 Inner products and normal-vector representations **20**
13.1 The final geometric link 20
13.2 One covector, many vector representations 22

14 Final synthesis **22**

Some Exercises **23**

Prerequisites and scope

This lecture is a geometric introduction for students and researchers in engineering, robotics, control, mechanics, physics, and optimization. It assumes basic linear algebra but no differential geometry. Its purpose is to build a precise geometric picture of a covector before using the concept in optimization, mechanics, robotics, and manifold geometry.

Linear-algebra prerequisites

The reader should be familiar with:

- finite-dimensional real vector spaces and linear subspaces;
- linear maps, kernels, images, rank, and rank–nullity;
- bases, basis coefficients, and direct sums; and
- affine subspaces and parallel translates.

Dual spaces, dual bases, covector transformation, and covector-to-vector identifications are developed in the lecture rather than assumed.

Conventions and level of treatment

All vector spaces are finite-dimensional and real. The lecture distinguishes:

1. a covector and its labeled affine foliation;
2. auxiliary choices, such as a basis or transverse vector;
3. support functions obtained by maximizing covector labels over sets;
4. normal-vector representations supplied by inner products; and
5. coordinate arrays used for computation.

The arguments are intended to be rigorous at the level needed by an engineering researcher while remaining readable and geometrically motivated.

What this lecture will develop

The lecture explains:
- how a linear scalar observation becomes a labeled affine foliation;
- why the kernel alone does not determine a covector;
- how basis evaluations and the dual basis describe the same covector;
- how scaling and addition act on labeled foliations;
- how support functions select extremal labels and supporting hyperplanes;
- how an inner product represents a covector by a normal vector.

Learning objectives and conceptual map

At the end of this lecture, the reader should be able to describe a nonzero covector as a linearly labeled family of parallel affine hyperplanes; reconstruct a covector from basis evaluations; interpret the dual basis as a family of coordinate-isolating foliations; explain scaling and addition geometrically; use covector labels to interpret support functions and supporting hyperplanes; and explain how an inner product supplies a normal-vector representation.

Conceptual progression

$$\alpha : V \to \mathbb{R} \quad \longrightarrow \quad \text{labeled affine foliation} \quad \longrightarrow \quad \text{supporting leaves of sets} \quad \longrightarrow \quad \alpha^{\sharp_g}.$$

The covector supplies the labels. Optimization selects extremal labels. An inner product later represents the same labeling by a normal vector.

1 Recommended references

This lecture is intended as a compact geometric introduction to covectors for students and researchers in engineering, robotics, mechanics, physics, and applied mathematics. The algebraic foundations and the extensions toward multivariable calculus and differential geometry are developed in greater depth in the following references.

For a systematic treatment of vector spaces, linear maps, dual spaces, quotient spaces, and inner-product spaces, the principal linear-algebra reference is Axler [1]. Its intrinsic treatment of finite-dimensional linear algebra provides the natural mathematical foundation for the distinction between a vector space and its dual used throughout this lecture.

For a broader development connecting linear algebra, multivariable calculus, differential forms, and geometry, see Hubbard and Hubbard [2]. This reference is particularly useful for readers who wish to follow the present ideas toward differentials, manifolds, and applications while retaining a unified geometric and computational viewpoint.

A closely related companion lecture

The slides *Vectors and Covectors* by Albert Chern, developed for the University of California San Diego course *CSE 270: Discrete Differential Geometry*, are strongly recommended as a companion to this lecture [3, 4].
The two treatments use a closely related language and are best read together. Both emphasize that vectors and covectors are different types of geometric objects, that the dual pairing is not an inner product, and that a covector can be visualized through parallel level hyperplanes.
The emphases are complementary. Chern's slides connect these ideas rapidly to adjoints, differentials, discrete differential geometry, and geometric computation. The present lecture develops in greater detail the linear labeling of the affine leaves, the geometric meaning of the dual basis, and the distinction between a covector and the normal-vector representation supplied by an inner product. Together, the two presentations provide a broader and more complete conceptual picture.

Readers wishing to continue immediately from covectors to differentials may also consult Chern's subsequent lecture *Function Differential* [5]. It develops the differential of a scalar function as a covector and distinguishes it from the gradient obtained only after an inner product has been selected.

References

[1] S. Axler, *Linear Algebra Done Right*, 4th ed., Undergraduate Texts in Mathematics, Springer, 2024. doi:10.1007/978-3-031-41026-0.

[2] J. H. Hubbard and B. B. Hubbard, *Vector Calculus, Linear Algebra, and Differential Forms: A Unified Approach*, 5th ed., Matrix Editions, 2015. ISBN 978-0-9715766-8-1. https://matrixeditions.com/5thUnifiedApproach.html.

[3] A. Chern, *Vectors and Covectors*, lecture slides for CSE 270: Discrete Differential Geometry, University of California San Diego, Winter 2024. https://cseweb.ucsd.edu/~alchern/teaching/cse270_wi24/3-1VectorCovector.pdf.

[4] A. Chern, *CSE 270: Discrete Differential Geometry*, course materials, University of California San Diego, Winter 2024. https://cseweb.ucsd.edu/~alchern/teaching/cse270_wi24/.

[5] A. Chern, *Function Differential*, lecture slides for CSE 270: Discrete Differential Geometry, University of California San Diego, Winter 2024. https://cseweb.ucsd.edu/~alchern/teaching/cse270_wi24/3-2Differential.pdf.

2 Methodological principle and ambient structure

The guiding rule is

Takeaway

Do not assume less than needed, and do not assume more than can be justified.

Let V be a finite-dimensional real vector space of dimension n, and let

$$V^* := \mathrm{Hom}(V, \mathbb{R})$$

be its algebraic dual. We want to provide a description of covectors as geometric objects using only:

- vector addition and scalar multiplication in V;
- the affine structure canonically induced by V;
- linear maps from V to $\mathbb{R}$;
- the vector-space and order structures of $\mathbb{R}$.

We do not assume a norm, inner product, angle, orthogonality, length, distance, unit vector, or normal vector. In particular, there is no canonical identification $V \simeq V^*$.

3 Representing a map inside its domain

3.1 Maps as labeled partitions

Let $f : X \to Y$ be an arbitrary map between sets. For $y \in Y$, its fiber is

$$f^{-1}(y) = \{x \in X : f(x) = y\}.$$

The nonempty fibers partition X: each $x \in X$ belongs to exactly one fiber, namely $f^{-1}(f(x))$. If each fiber retains its label, then the labeled partition contains precisely the information of the map:

$$\boxed{f \quad \longleftrightarrow \quad \big\{\big(f^{-1}(y), y\big) : y \in f(X)\big\}.} \tag{1}$$

Indeed, $f(x)$ is recovered as the label of the unique fiber containing x.

3.2 Real-valued maps and level sets

For $f : X \to \mathbb{R}$, the fibers are the level sets for any $c \in \mathbb{R}$

$$f^{-1}(c) = H_{f,c} = \{x \in X : f(x) = c\}.$$

Thus a real-valued map can be represented inside its domain as scalar-labeled level sets. For a general nonlinear map, these sets need not possess any special geometry.

4 From linear scalar observation to affine foliation

4.1 Formal definition

A covector is a linear scalar observation

$$\boxed{\alpha \in V^* \quad \Longleftrightarrow \quad \alpha : V \to \mathbb{R} \text{ is linear.}} \tag{2}$$

The scalar $\alpha(x)$ is the label assigned by α to $x \in V$.

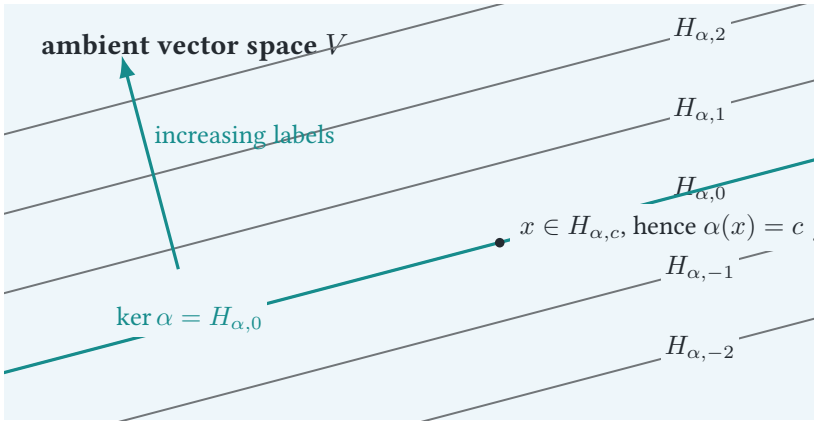

Figure 1: A covector represented inside V by its labeled fibers. The covector determines which points lie on the same leaf, which leaves are parallel, and the label and ordering of each leaf. The apparent angles and spacings belong only to the drawing.

4.2 The kernel has codimension one

Assume $\alpha \neq 0$. Then $\operatorname{im}\alpha$ is a nonzero linear subspace of the one-dimensional vector space $\mathbb{R}$, hence $\operatorname{im}\alpha = \mathbb{R}$. Rank–nullity yields

$$\boxed{\dim \ker \alpha = n - 1, \quad \operatorname{codim}(\ker \alpha) := \dim V - \dim \ker \alpha = 1.} \tag{3}$$

Therefore $\ker \alpha$ is a linear hyperplane through the origin. By contrast, $\ker 0 = V$, where 0 is the zero map in V^*.

An equivalent direct proof chooses $u \in V$ with $\alpha(u) = 1$. Every $x \in V$ admits the decomposition

$$\boxed{x = \big(x - \alpha(x)u\big) + \alpha(x)u \in \ker \alpha \oplus \operatorname{span}\{u\}.} \tag{4}$$

4.3 All fibers are parallel affine hyperplanes

For $c \in \mathbb{R}$, define

$$H_{\alpha,c} := \{x \in V : \alpha(x) = c\}.$$

Choose any x_c satisfying $\alpha(x_c) = c$. Then

$$\boxed{H_{\alpha,c} = x_c + \ker \alpha.} \tag{5}$$

Indeed, $x \in H_{\alpha,c}$ if and only if $\alpha(x - x_c) = 0$. Consequently,

$$\boxed{V = \bigsqcup_{c \in \mathbb{R}} H_{\alpha,c}.} \tag{6}$$

These affine hyperplanes are parallel because they are translates of the same direction subspace $\ker \alpha$.

Takeaway

Parallelism is an affine concept; no metric is involved..

5 The complete geometric content of a covector

5.1 The kernel determines the leaf direction

Let $\alpha \neq 0$. The kernel $\ker \alpha$ determines the common direction of all affine leaves $H_{\alpha,c}$. If the labels are erased, only the parallel foliation remains. This unlabeled foliation does not determine fully the covector: every nonzero scalar multiple $r\alpha$, with $r \in \mathbb{R}$, has the same kernel.

5.2 The linear labeling completes the foliation

The covector does more than determine which points lie on the same leaf. It assigns to every leaf the real number c for which that leaf is $H_{\alpha,c}$. The signs of these labels determine the coorientation, their order determines the ordering of the leaves, and their numerical values determine the calibration. These are not independent additions to a previously complete labeling; they are aspects of the same signed linear labeling.

The intrinsic formulation is the induced isomorphism

$$\boxed{\overline{\alpha} : V/\ker \alpha \to \mathbb{R}, \quad [x] \mapsto \alpha(x).} \tag{7}$$

It is well defined because $[x] = [y]$ means $x - y \in \ker \alpha$, hence $\alpha(x) = \alpha(y)$. It is injective because $\overline{\alpha}([x]) = 0$ implies $x \in \ker \alpha$, and it is surjective because a nonzero map from V to the one-dimensional space $\mathbb{R}$ has image $\mathbb{R}$.

Takeaway

nonzero covector = hyperplane direction and induced foliation + linear labeling of the leaves.

Coorientation is induced by the complete labeling: the positive labels distinguish one transverse side of the kernel from the other.

5.3 The equivalent nested-half-space representation

The order on $\mathbb{R}$ turns the labeled foliation into nested half-spaces:

$$\boxed{H^-_{\alpha,c} = \{x : \alpha(x) \leq c\}, \quad H^+_{\alpha,c} = \{x : \alpha(x) \geq c\}.} \tag{8}$$

Their common boundary is $H_{\alpha,c}$, and

$$c_1 \leq c_2 \quad \Longrightarrow \quad H^-_{\alpha,c_1} \subseteq H^-_{\alpha,c_2}.$$

Takeaway

nonzero covector = a linearly labeled nested family of affine half-spaces..

6 Transverse parameterization and quotient geometry

6.1 A transverse vector

Choose a transverse vector to $\ker \alpha$ so $a := \alpha(u) \neq 0$. The term transverse means only $u \notin \ker \alpha$, equivalently $\operatorname{span}\{u\} \cap \ker \alpha = \{0\}$. It does not mean orthogonal.

Then

$$V = \ker \alpha \oplus \operatorname{span}\{u\},$$

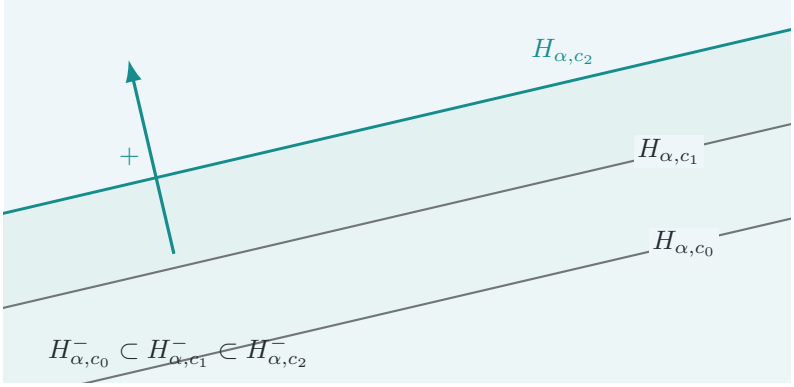

Figure 2: The ordered labels turn the affine foliation into a nested family of lower half-spaces. Inclusion follows from $c_0 < c_1 < c_2$, not from any metric distance between the boundary hyperplanes.

and every $x \in V$ is uniquely expressible as $x = h + \lambda u$, with $h \in \ker \alpha$. Linearity gives

$$\boxed{\alpha(h + \lambda u) = \lambda a.} \tag{9}$$

Therefore

$$\boxed{H_{\alpha,\lambda a} = \lambda u + \ker \alpha.} \tag{10}$$

Here a is the calibration of the chosen transverse vector; λ is its coordinate; and λa is the actual leaf label.

We have seen that a nonzero covector intrinsically identifies the one-dimensional quotient $V/\ker \alpha$ with $\mathbb{R}$. Choosing a transverse vector merely supplies concrete representatives of the quotient classes.

6.2 Normalization without a norm

Rescaling u by choosing $u' = \alpha(u)^{-1}u$ permits the algebraic normalization $\alpha(u') = 1$. This does not make u' a unit vector. It merely selects a representative whose covector label is one.

6.3 Reconstructing a covector from geometric data

Let $H \subset V$ be a hyperplane. Choose one transverse vector $u \notin H$ and prescribe one nonzero value $a \in \mathbb{R}$. Since

$$V = H \oplus \operatorname{span}\{u\},$$

every $x \in V$ has a unique expression $x = h + \lambda u$ with $h \in H$. There is then exactly one covector with kernel H and $\alpha(u) = a$, namely

$$\boxed{\alpha(h + \lambda u) = \lambda a.} \tag{11}$$

The hyperplane determines which points receive the same label. The pair (u, a) determines how the resulting one-dimensional leaf space is identified with $\mathbb{R}$.

7 A basis gives finite access to the complete covector

Let

$$\mathcal{B} = (v_1, \ldots, v_n)$$

be a basis of V. A covector α assigns one scalar label to each basis vector:

$$\alpha(v_1), \ldots, \alpha(v_n).$$

These n labels determine the complete labeled foliation associated with α. No additional values of α need to be given.

$$\boxed{\text{A covector on an } n\text{-dimensional vector space is uniquely determined by its values on any chosen basis.}} \tag{12}$$

Indeed, every $x \in V$ has a unique decomposition

$$x = x^1 v_1 + \cdots + x^n v_n.$$

Linearity then forces

$$\boxed{\alpha(x) = x^1 \alpha(v_1) + \cdots + x^n \alpha(v_n).} \tag{13}$$

Thus the basis labels determine the label of every point, and therefore every affine leaf $H_{\alpha,c}$.

Conversely, arbitrary prescribed labels $a_1, \dots, a_n \in \mathbb{R}$ define exactly one covector by

$$\boxed{\alpha\left(\sum_{i=1}^{n} x^i v_i\right) := \sum_{i=1}^{n} a_i x^i.} \tag{14}$$

This covector satisfies $\alpha(v_i) = a_i$ for every i.

The assignment

$$\boxed{\mathcal{C}_{\mathcal{B}} : V^* \longrightarrow \mathbb{R}^n, \quad \mathcal{C}_{\mathcal{B}}(\alpha) = \big(\alpha(v_1), \dots, \alpha(v_n)\big)} \tag{15}$$

is therefore bijective. It is also linear because labels add and scale pointwise:

$$\mathcal{C}_{\mathcal{B}}(a\alpha + b\beta) = a\mathcal{C}_{\mathcal{B}}(\alpha) + b\mathcal{C}_{\mathcal{B}}(\beta).$$

Hence $\mathcal{C}_{\mathcal{B}}$ is a linear isomorphism and

$$\boxed{\dim V^* = \dim V = n.} \tag{16}$$

The equality of dimensions does not canonically identify V with V^*. The representation in $\mathbb{R}^n$ depends on the selected basis.

7.1 Recovering the kernel from the basis labels

The same conclusion has a direct geometric interpretation. Let

$$a_i := \alpha(v_i).$$

If all a_i vanish, then every point has label zero and $\alpha = 0$. Otherwise choose an index p such that $a_p \neq 0$. For every $i \neq p$, define

$$\boxed{h_i := v_i - \frac{a_i}{a_p} v_p.} \tag{17}$$

Then

$$\alpha(h_i) = a_i - \frac{a_i}{a_p} a_p = 0,$$

so every h_i belongs to $\ker \alpha$. If $a_i = 0$, then $h_i = v_i$: a basis vector already lying in the zero leaf is retained unchanged.

The $n - 1$ vectors h_i, with $i \neq p$, are linearly independent. Indeed, any linear relation among them becomes a linear relation among the original basis vectors, and all its coefficients must vanish. Since

$\ker\alpha$ has dimension $n-1$, these vectors form a basis of the kernel:

$$\boxed{\ker\alpha = \operatorname{span}\left\{ v_i - \frac{\alpha(v_i)}{\alpha(v_p)} v_p : i \neq p \right\}.} \tag{18}$$

Geometrically, each v_i is displaced along the selected transverse direction v_p until its label is cancelled. The resulting vectors lie on the zero leaf and span it. The nonzero value $\alpha(v_p)$ then fixes the calibration transverse to that leaf. Thus the n basis labels recover both defining parts of the covector:

$$\boxed{\begin{array}{lll} \alpha(v_1), \dots, \alpha(v_n) & \Longrightarrow & \text{the kernel hyperplane,} \\ & & \text{the coorientation and scalar calibration.} \end{array}} \tag{19}$$

8 The dual basis as coordinate-isolating affine foliations

Let $\mathcal{B} = (v_1, \dots, v_n)$ be a basis of V. For each i, the other $n-1$ basis vectors span the complementary hyperplane

$$H_i := \operatorname{span}\{v_j : j \neq i\}.$$

Because

$$V = H_i \oplus \operatorname{span}\{v_i\},$$

there is a unique covector $\varepsilon^i \in V^*$ whose zero leaf is H_i and whose leaf through v_i has label one:

$$\boxed{\ker\varepsilon^i = \operatorname{span}\{v_j : j \neq i\}, \quad \varepsilon^i(v_i) = 1.} \tag{20}$$

Equivalently,

$$\boxed{\varepsilon^i(v_j) = \delta^i_j = \begin{cases} 1, & j = i, \\ 0, & j \neq i. \end{cases}} \tag{21}$$

The family of covectors

$$\mathcal{B}^* = (\varepsilon^1, \dots, \varepsilon^n)$$

is called the basis dual to $\mathcal{B}$.

Each dual covector has a direct geometric meaning. Every affine leaf parallel to H_i intersects $\operatorname{span}\{v_i\}$ in exactly one point λv_i. The covector ε^i assigns to the entire leaf

$$\lambda v_i + H_i$$

the label λ, namely the relative scale between the intersection point λv_i and the calibrating vector v_i.

The two components of the direct sum

$$V = H_i \oplus \operatorname{span}\{v_i\}$$

therefore have distinct geometric roles:

- H_i determines which points receive the same label;
- v_i determines the scale used to label the leaves.

Indeed, if

$$x = x^1 v_1 + \cdots + x^n v_n,$$

then, relative to the i-th direct-sum decomposition,

$$x = \underbrace{\sum_{j\neq i} x^j v_j}_{\in H_i} + \underbrace{x^i v_i}_{\in \mathrm{span}\{v_i\}}.$$

Consequently,

$$\boxed{\varepsilon^i(x) = x^i.} \tag{22}$$

Thus ε^i isolates the scalar multiplying v_i in the basis decomposition, i.e., the i-th coordinate of the vector x w.r.t. the basis $\mathcal{B}$. This is not an orthogonal projection. It uses only the direct-sum decomposition

$$V = H_i \oplus \mathrm{span}\{v_i\}.$$

No inner product, length, angle, or identification of vectors with covectors is involved.

8.1 Recovering a vector as the intersection of its coordinate leaves

The equality

$$\varepsilon^i(x) = x^i$$

has a stronger geometric interpretation than the mere extraction of a number. It identifies the unique leaf that contains x in of the foliation induced by the i-th covector of the dual basis:

$$\boxed{x \in x^i v_i + H_i = \{y \in V : \varepsilon^i(y) = x^i\}.} \tag{23}$$

The intersection of this leaf with $\mathrm{span}\{v_i\}$ is the point $x^i v_i$. Therefore, the label x^i is the relative scale between this intersection point and the calibrating vector v_i.

Considering only ε^i isolates one coordinate, but it does not identify x uniquely. Every point on the $n-1$-dimensional affine hyperplane

$$x^i v_i + H_i$$

has the same i-th coordinate. The remaining $n-1$ coordinates are free to vary within the direction subspace H_i.

The complete dual basis removes this freedom one coordinate at a time. For every $i = 1, \ldots, n$, the covector ε^i selects the affine leaf

$$x^i v_i + H_i$$

containing x. The vector x is the unique point common to all these leaves:

$$\boxed{\{x\} = \bigcap_{i=1}^{n} \big(x^i v_i + H_i\big).} \tag{24}$$

To prove this statement, first observe that x belongs to every leaf, because $\varepsilon^i(x) = x^i$. Hence

$$x \in \bigcap_{i=1}^{n} \big(x^i v_i + H_i\big).$$

Conversely, let y belong to the intersection. Then, for every i,

$$y \in x^i v_i + H_i,$$

and therefore

$$\varepsilon^i(y) = x^i = \varepsilon^i(x).$$

It follows that

$$\varepsilon^i(y - x) = 0 \quad \text{for every } i.$$

Writing

$$y - x = \sum_{j=1}^{n} z^j v_j,$$

and applying ε^i, we obtain

$$z^i = \varepsilon^i(y - x) = 0 \quad \text{for every } i.$$

Thus $y - x = 0$, so $y = x$.

Takeaway

One dual covector isolates one coordinate by selecting one affine hyperplane. The complete dual basis recovers the vector as the unique intersection of the n selected coordinate hyperplanes.

This construction uses only direct sums, affine translates, and affine intersections. No orthogonality is involved.

8.2 Why the dual covectors form a basis of V^*

The covectors $\varepsilon^1, \dots, \varepsilon^n$ are linearly independent. If

$$\sum_{i=1}^{n} c_i \varepsilon^i = 0,$$

evaluating at v_j gives $c_j = 0$. They also span V^*, because every covector α satisfies

$$\boxed{\alpha = \sum_{i=1}^{n} \alpha(v_i)\varepsilon^i.} \tag{25}$$

Both sides assign the same label $\alpha(v_j)$ to every basis vector v_j, and a covector is uniquely determined by those labels.

Takeaway

A vector basis decomposes every vector into scalar contributions. The dual basis consists of the geometric covectors that extract those contributions.

8.3 Basis evaluations and dual coordinates are the same numbers

The preceding constructions now fit together without introducing any additional structure. Given $\alpha \in V^*$, first measure the vector basis:

$$a_i := \alpha(v_i), \quad i = 1, \dots, n.$$

The same numbers have three simultaneous meanings:

1. a_i is the label assigned by α to the vector v_i;

2. the tuple $(a_1, \dots, a_n)$ reconstructs the kernel by the label-cancellation construction of (18);
3. a_i is the coordinate multiplying the dual-basis covector ε^i in (25).

Thus evaluation on the vector basis and expansion in the dual basis are not two unrelated procedures. They are inverse descriptions of the same geometric covector. Define

$$\mathcal{C}_{\mathcal{B}} : V^* \longrightarrow \mathbb{R}^n, \qquad \alpha \longmapsto \big(\alpha(v_1), \dots, \alpha(v_n)\big), \tag{26}$$

$$\mathcal{S}_{\mathcal{B}^*} : \mathbb{R}^n \longrightarrow V^*, \qquad (a_1, \dots, a_n) \longmapsto \sum_{i=1}^{n} a_i \varepsilon^i. \tag{27}$$

The covector $\mathcal{S}_{\mathcal{B}^*}(a_1, \dots, a_n)$ assigns the label a_j to v_j, because

$$\left(\sum_{i=1}^{n} a_i \varepsilon^i\right)(v_j) = \sum_{i=1}^{n} a_i \delta^i_j = a_j.$$

Consequently,

$$\boxed{\mathcal{S}_{\mathcal{B}^*} \circ \mathcal{C}_{\mathcal{B}} = \mathrm{id}_{V^*}, \quad \mathcal{C}_{\mathcal{B}} \circ \mathcal{S}_{\mathcal{B}^*} = \mathrm{id}_{\mathbb{R}^n}\,.} \tag{28}$$

Geometrically, each ε^i is a coordinate-isolating calibrated foliation that extracts the contribution along v_i relative to the complementary hyperplane spanned by the other basis vectors. A general covector is obtained by weighting these dual-basis labelings by the measured labels a_i and adding them pointwise.

Takeaway

Measuring the vector basis and combining the dual-basis foliations are inverse descriptions of the same covector.

9 Scalar multiplication as geometric relabeling

For $r \in \mathbb{R}$, define $(r\alpha)(x) = r\alpha(x)$.

9.1 Positive scaling

If $r > 0$, then

$$\ker(r\alpha) = \ker \alpha, \quad H_{r\alpha,c} = H_{\alpha,c/r}.$$

The unlabeled foliation and coorientation remain unchanged; the labels are rescaled by $c \mapsto rc$.

9.2 Negative scaling

If $r < 0$, the unlabeled foliation remains unchanged, the labels are rescaled, and the coorientation reverses.

9.3 Zero scaling

If $r = 0$, every point receives label zero. The codimension-one foliation degenerates into the single fiber V.

Takeaway

$r\alpha$: same unlabeled foliation, rescaled labels, and reversed coorientation exactly when $r < 0$.

10 Addition as pointwise label addition and geometric shear

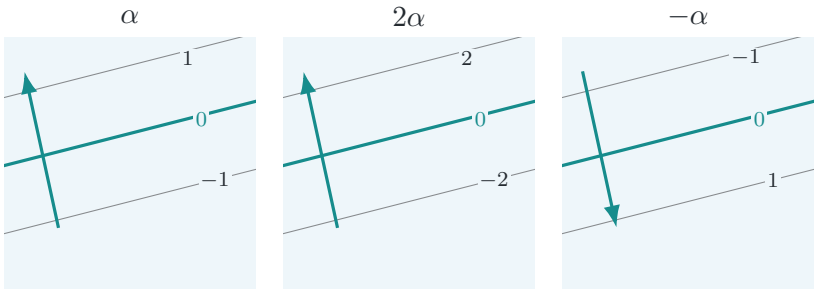

Figure 3: Scalar multiplication changes labels, not the unlabeled foliation. Positive scaling preserves the coorientation; negative scaling reverses it. No statement about metric spacing is intended, even though the leaves must be placed apart on the page.

10.1 Pointwise label addition

For $\alpha, \beta \in V^*$, define

$$\boxed{(\alpha+\beta)(x) = \alpha(x) + \beta(x).} \tag{29}$$

Geometrically, each point already carries an α-label and a β-label. The sum covector assigns their sum. The kernel consists of the points at which the two labels cancel:

$$\boxed{\ker(\alpha+\beta) = \{x : \alpha(x) + \beta(x) = 0\}.} \tag{30}$$

10.2 Linearly dependent covectors

If $\beta = r\alpha$, then $\alpha + \beta = (1+r)\alpha$. Hence

$$\boxed{\ker(\alpha+\beta) = \begin{cases} \ker\alpha, & r \neq -1, \\ V, & r = -1. \end{cases}} \tag{31}$$

10.3 Linearly independent covectors

Assume α and β are linearly independent. There exists $u_\beta \in \ker\beta$ with $\alpha(u_\beta) \neq 0$. Otherwise $\ker\beta \subseteq \ker\alpha$; since both are hyperplanes, their kernels would coincide, forcing linear dependence.

Because $u_\beta \notin \ker\alpha$, every $v \in V$ has a unique decomposition

$$v = u_\alpha + \lambda u_\beta, \quad u_\alpha \in \ker\alpha.$$

Now impose $(\alpha+\beta)(v) = 0$:

$$0 = \beta(u_\alpha) + \lambda\alpha(u_\beta),$$

so

$$\lambda = -\frac{\beta(u_\alpha)}{\alpha(u_\beta)}.$$

Therefore

$$\boxed{\ker(\alpha+\beta) = \left\{ u_\alpha - \frac{\beta(u_\alpha)}{\alpha(u_\beta)} u_\beta : \ u_\alpha \in \ker\alpha \right\}.} \tag{32}$$

10.4 Shear interpretation

Define

$$\boxed{\Phi_{\alpha,\beta}(u_\alpha) = u_\alpha - \frac{\beta(u_\alpha)}{\alpha(u_\beta)} u_\beta.} \tag{33}$$

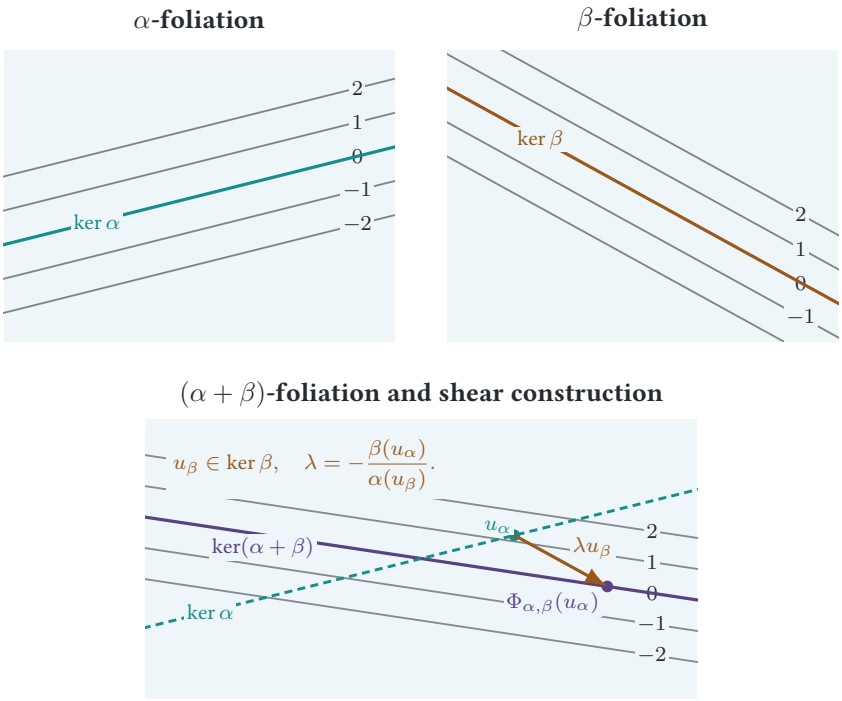

Figure 4: Addition of two covectors. The upper panels show the labeled foliations of α and β. In the lower panel, the labels are added pointwise. Starting from $u_\alpha \in \ker \alpha$, the displacement λu_β, with $u_\beta \in \ker \beta$, reaches $\Phi_{\alpha,\beta}(u_\alpha) \in \ker(\alpha + \beta)$, where $\lambda = -\beta(u_\alpha)/\alpha(u_\beta)$. The displacement is parallel to $\ker \beta$; no orthogonality is used.

This is a linear isomorphism from $\ker \alpha$ onto $\ker(\alpha + \beta)$. Geometrically:

1. start with $u_\alpha \in \ker \alpha$;
2. read its β-label $\beta(u_\alpha)$;
3. move along the fixed direction $u_\beta \in \ker \beta$;
4. choose the displacement that makes the total label zero.

The new kernel is therefore a metric-free linear shear of $\ker \alpha$. Exchanging α and β gives a symmetric construction from $\ker \beta$.

Conceptual warning

The kernels alone do not determine the kernel of the sum. Addition depends on the coorientations and relative calibrations of the two covectors. Replacing β by $r\beta$ preserves $\ker \beta$ but generally changes $\ker(\alpha + r\beta)$.

10.5 Recovering the complete geometry of the sum

Once $\ker(\alpha + \beta)$ is known, every other leaf is an affine translate:

$$H_{\alpha+\beta,c} = x_c + \ker(\alpha + \beta), \quad (\alpha + \beta)(x_c) = c.$$

The sign of $\alpha + \beta$ determines the coorientation, and the values c give the calibration.

11 Metric-free audit

The following concepts have been used without an inner product:

- kernel, image, dimension, and codimension;
- affine hyperplanes, affine translates, and parallelism;
- fibers, level sets, foliations, and nested half-spaces;
- linear labeling and its induced coorientation;
- transversality as direct-sum complementarity;
- quotient spaces, basis labels, and dual bases;
- scalar multiplication, label addition, and the shear construction.

Without additional structure, one must not claim that a covector is a normal vector, that a transverse vector is perpendicular to the kernel, that vectors have magnitudes, that leaves possess metric distances, or that V is canonically identifiable with V^*.

12 Application: support functions and supporting hyperplanes

12.1 Why support functions are useful

Optimization often asks how far a set extends in a selected linear direction.

Let $C \subseteq V$ be nonempty. The support function of C is the map

$$\boxed{\sigma_C : V^* \longrightarrow \overline{\mathbb{R}},}$$

where

$$\overline{\mathbb{R}} := \mathbb{R} \cup \{+\infty\},$$

defined by

$$\boxed{\sigma_C(\alpha) := \sup_{x \in C} \alpha(x), \quad \alpha \in V^*.} \tag{34}$$

It is used in convex optimization, separation, duality, set propagation, and geometric descriptions of feasible regions.

The formula alone can appear abstract. The labeled-foliation picture makes it immediate: move through the leaves $H_{\alpha,c}$ in the direction of increasing labels. The support value is the last label reached while the leaf still meets C.

12.2 The supporting leaf

Assume that the supremum is attained at $x_\star \in C$. Then

$$\alpha(x_\star) = \sigma_C(\alpha),$$

and every $x \in C$ satisfies

$$\alpha(x) \le \sigma_C(\alpha).$$

Consequently,

$$\boxed{C \subseteq H^-_{\alpha,\sigma_C(\alpha)}, \quad x_\star \in H_{\alpha,\sigma_C(\alpha)}.} \tag{35}$$

The leaf

$$H_{\alpha,\sigma_C(\alpha)} = \{x \in V : \alpha(x) = \sigma_C(\alpha)\}$$

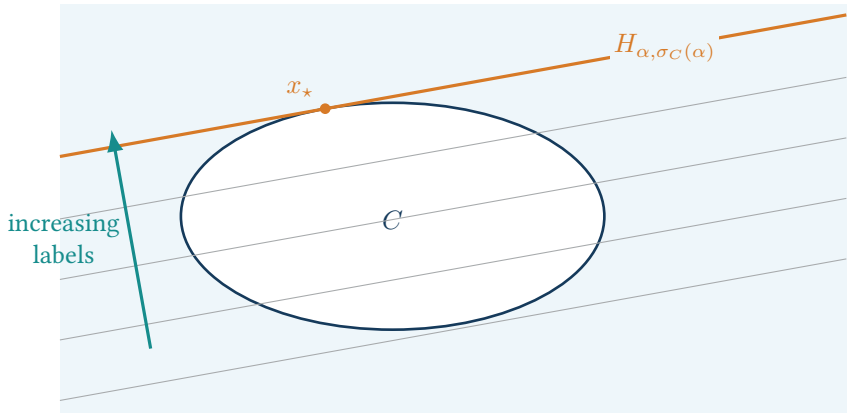

Figure 5: The covector α labels a parallel family of affine hyperplanes. The support value $\sigma_C(\alpha)$ is the largest label attained on C; the corresponding leaf supports the set at the contact point $x_\star$. The drawing uses no metric interpretation of α.

is therefore a supporting hyperplane of C. The contact set is

$$\boxed{\arg\max_{x\in C} \alpha(x) = C \cap H_{\alpha,\sigma_C(\alpha)}.}$$

For a compact set C, the maximum is attained. Convexity is not required to define σ_C, although support functions are especially informative for convex sets.

Why the covector picture helps

The support function does not first require a normal vector. Its argument is a covector: a rule that labels parallel affine hyperplanes. Maximizing the covector over C selects the extremal label and therefore the supporting leaf. An inner product may later represent α by a normal vector, but the supporting-hyperplane construction is already complete.

12.3 Set inclusion as support-function comparison

If $C \subseteq D$, then

$$\boxed{\sigma_C(\alpha) \le \sigma_D(\alpha) \quad \text{for every } \alpha \in V^*.}$$

For each covector, the last labeled leaf reached by C cannot lie beyond the last leaf reached by D.

For nonempty closed convex sets, the converse follows from separation:

$$\boxed{C \subseteq D \iff \sigma_C \le \sigma_D \quad \text{pointwise on } V^*.} \tag{36}$$

Indeed, if $x \in C \setminus D$, a separating covector satisfies $\alpha(x) > \sigma_D(\alpha)$, whereas $\sigma_C(\alpha) \ge \alpha(x)$.

Consequently, for nonempty closed convex sets,

$$C = D \iff \sigma_C = \sigma_D.$$

For an arbitrary nonempty set,

$$\boxed{\sigma_C = \sigma_{\overline{\mathrm{conv}}C},}$$

so the support function represents the closed convex hull rather than any nonconvex internal structure.

12.4 Basic consequences of the definition

For $r \geq 0$,

$$\boxed{\sigma_C(r\alpha) = r\,\sigma_C(\alpha).}$$

For $\alpha, \beta \in V^*$,

$$\boxed{\sigma_C(\alpha+\beta) \leq \sigma_C(\alpha) + \sigma_C(\beta).}$$

Thus, whenever it is finite, the support function is a positively homogeneous, subadditive function on the dual space. These properties follow directly from scaling and adding covector labels.

Takeaway

$$\boxed{\text{support function} = \text{largest covector label attained on the set.}}$$

The corresponding extremal leaf is a supporting hyperplane. For closed convex sets, pointwise comparison of support functions is equivalent to set inclusion:

$$C \subseteq D \quad \Longleftrightarrow \quad \sigma_C \leq \sigma_D.$$

Thus, the support function converts geometric comparison of closed convex sets into scalar comparison of functions on the dual space.

13 Inner products and normal-vector representations

An inner product g supplies lengths, angles, and orthogonality. It also converts each covector into a vector whose metric pairing reproduces the covector labels. This vector is normal to the labeled leaves.

Define

$$g^\flat(v) = g(v, \cdot), \quad g^\sharp = (g^\flat)^{-1}.$$

For $\alpha \in V^*$, write

$$\alpha^{\sharp_g} = g^\sharp(\alpha).$$

It is the unique vector satisfying

$$\boxed{\alpha(x) = g(\alpha^{\sharp_g}, x) \quad \forall x \in V.} \tag{37}$$

If $h \in \ker\alpha$, then $g(\alpha^{\sharp_g}, h) = \alpha(h) = 0$. Therefore

$$\alpha^{\sharp_g} \in (\ker\alpha)^{\perp_g}.$$

Because $\ker\alpha$ is a hyperplane, its orthogonal complement is one-dimensional. Orthogonality determines the transverse line; the requirement that $g(\alpha^{\sharp_g}, \cdot)$ reproduce every label of α determines the sign and scale on that line.

13.1 The final geometric link

Let $\alpha \neq 0$ be a covector. Without an inner product, choose any transverse vector

$$u \notin \ker\alpha,$$

and set $a = \alpha(u) \neq 0$. Every labeled leaf can then be written as

$$\boxed{H_{\alpha,c} = \frac{c}{a}u + \ker\alpha.} \tag{38}$$

This parameterization uses the direct sum

$$V = \ker \alpha \oplus \operatorname{span}\{u\}.$$

The covector determines the kernel and the linear labeling of $V/\ker\alpha$, but no transverse line has yet been selected. Every line complementary to $\ker\alpha$ gives a valid vector-based parameterization of the same labeled foliation.

After choosing an inner product g, the situation changes. The hyperplane $\ker\alpha$ now determines the distinguished transverse line

$$(\ker\alpha)^{\perp_g}.$$

On this line, there is a unique vector $\alpha^{\sharp_g}$ satisfying

$$\boxed{\alpha(x) = g(\alpha^{\sharp_g}, x) \quad \forall x \in V.} \tag{39}$$

The line $\operatorname{span}\{\alpha^{\sharp_g}\}$ is fixed by g-orthogonality to the kernel, while the exact vector $\alpha^{\sharp_g}$ on that line is fixed by the requirement that it reproduce the complete labeling defined by α.

The unique vector on this distinguished line that reaches the leaf labeled one is

$$\boxed{u_g = \frac{\alpha^{\sharp_g}}{g(\alpha^{\sharp_g}, \alpha^{\sharp_g})}, \quad \{u_g\} = H_{\alpha,1} \cap (\ker\alpha)^{\perp_g}.} \tag{40}$$

Indeed,

$$\alpha(u_g) = g(\alpha^{\sharp_g}, u_g) = 1.$$

Therefore, the arbitrary parameterization (38) is replaced, relative to the enriched structure (V, g), by the distinguished parameterization

$$\boxed{H_{\alpha,c} = c\,u_g + \ker\alpha = \frac{c}{g(\alpha^{\sharp_g}, \alpha^{\sharp_g})}\alpha^{\sharp_g} + \ker\alpha.} \tag{41}$$

Notice that u_g is normalized by the covector label, $\alpha(u_g) = 1$, not by its g-length.

The essential geometric consequence is that the single vector $\alpha^{\sharp_g}$ now encodes both components of the covector geometry. Its line $\operatorname{span}\{\alpha^{\sharp_g}\}$ determines the kernel through g-orthogonality, while its g-pairing with every vector determines the complete scale and sign of the leaf labeling:

$$\boxed{\ker\alpha = \big(\operatorname{span}\{\alpha^{\sharp_g}\}\big)^{\perp_g} = \{x \in V : g(\alpha^{\sharp_g}, x) = 0\}, \quad H_{\alpha,c} = \{x \in V : g(\alpha^{\sharp_g}, x) = c\}.} \tag{42}$$

Takeaway

Before an inner product is chosen, a covector determines its kernel and the linear labels of the leaves, while any complementary line may be used to parameterize them.

After choosing an inner product g, the orthogonal complement $(\ker\alpha)^{\perp_g}$ supplies a transverse line, and the complete labeling selects the unique vector $\alpha^{\sharp_g}$ that reproduces all covector values.

13.2 One covector, many vector representations

The covector α, its kernel, its affine leaves, and their labels do not change when the inner product changes. The vector representation generally does. Denoting by $\mathrm{Inn}(V)$ the set of inner products on V, every choice defines an isomorphism

$$g^\sharp : V^* \to V,$$

and hence a map

$$\mathrm{Inn}(V) \to \mathrm{Iso}(V^*, V), \quad g \mapsto g^\sharp.$$

Consequently, one covector gives rise to the family of vector representations

$$\left\{ \alpha^{\sharp_g} : g \in \mathrm{Inn}(V) \right\}.$$

Each $\alpha^{\sharp_g}$ encodes the same labeled foliation, but it does so through the orthogonality and scale supplied by its particular inner product g. Changing g changes the normal-vector representation, while the covector, its leaves, and their labels remain unchanged.

Takeaway

α is the covector, whereas $\alpha^{\sharp_g}$ is its normal-vector representation relative to g.

14 Final synthesis

A nonzero covector admits four mutually reinforcing descriptions:

$$\begin{aligned} \text{map view:} &\quad \alpha : V \to \mathbb{R} \text{ linear}, \\ \text{foliation view:} &\quad \text{parallel affine hyperplanes with linear labels}, \\ \text{quotient view:} &\quad \overline{\alpha} : V/\ker\alpha \xrightarrow{\sim} \mathbb{R}, \\ \text{basis view:} &\quad \alpha = \sum_{i=1}^{n} \alpha(v_i)\varepsilon^i. \end{aligned} \tag{43}$$

The same geometric picture organizes the principal constructions of the lecture:

Construction	Geometric meaning
$\ker\alpha$	zero-labeled leaf and common direction of the foliation
$H_{\alpha,c}$	leaf carrying label c
ε^i	coordinate-isolating labeled foliation
$r\alpha$	relabeling of the same foliation, with reversal if $r < 0$
$\alpha + \beta$	pointwise addition of labels and the induced shear of the zero leaf
$\sigma_C(\alpha)$	largest α-label attained on C
$\alpha^{\sharp_g}$	normal-vector representation supplied by the chosen inner product

Thus, a covector is first understood as a linear scalar observation and a labeled affine foliation. Support functions use this observation directly to select supporting leaves. An inner product adds the familiar normal-vector picture without replacing the underlying covector geometry.

Some Exercises

Exercise 14.1 (Equal-label leaves)**.** Let $\alpha \in (\mathbb{R}^2)^*$ be defined by

$$\alpha(x^1, x^2) = 3x^1 + 2x^2.$$

(a) Determine $\ker \alpha$.

(b) Describe $H_{\alpha,c}$ for arbitrary $c \in \mathbb{R}$.

(c) Explain which geometric information remains if all labels are erased.

(d) Compare the labeled foliations of α, 2α, and $-\alpha$.

Exercise 14.2 (Recovering a covector from basis labels)**.** Let $\mathcal{B} = (v_1, v_2, v_3)$ be a basis and suppose

$$\alpha(v_1) = 2, \quad \alpha(v_2) = -1, \quad \alpha(v_3) = 3.$$

(a) Reconstruct $\alpha(x)$ for $x = x^1 v_1 + x^2 v_2 + x^3 v_3$.

(b) Use v_3 as the transverse basis vector and construct a basis of $\ker \alpha$ by label cancellation.

(c) Write the leaf $H_{\alpha,c}$ as a translate of $\ker \alpha$.

Exercise 14.3 (Coordinate leaves of the dual basis)**.** For a basis $\mathcal{B} = (v_1, \dots, v_n)$, let $H_i = \operatorname{span}\{v_j : j \neq i\}$.

(a) Show that every affine leaf parallel to H_i intersects $\operatorname{span}\{v_i\}$ in exactly one point.

(b) Prove that ε^i labels the leaf $\lambda v_i + H_i$ by λ.

(c) Prove that $\{x\} = \bigcap_{i=1}^n (x^i v_i + H_i)$.

Exercise 14.4 (One covector, two inner products)**.** On $V = \mathbb{R}^2$, let

$$\alpha(x) = 3x^1 + 2x^2.$$

Consider the inner products represented in the standard basis by

$$G_1 = I, \quad G_2 = \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix}.$$

(a) Compute $\alpha^{\sharp_{g_1}}$ and $\alpha^{\sharp_{g_2}}$.

(b) Verify that both reproduce the same covector through their respective inner products.

(c) Verify that each is orthogonal to $\ker \alpha$ in its own inner product.

Exercise 14.5 (Support function and supporting leaves)**.** Let $C \subset \mathbb{R}^2$ be the triangle with vertices $(0, 0)$, $(2, 0)$, and $(0, 1)$, and let $\alpha(x^1, x^2) = x^1 + 2x^2$.

(a) Compute $\sigma_C(\alpha)$.

(b) Determine the contact set $C \cap H_{\alpha,\sigma_C(\alpha)}$.

(c) Describe the containing half-space $H^-_{\alpha,\sigma_C(\alpha)}$.

(d) Explain the result using only the labeled-foliation interpretation.

Exercise 14.6 (Force as a covector)**.** Let $F \in V^*$ assign virtual work $F(\delta q)$ to an infinitesimal displacement $\delta q \in V$.

(a) Interpret $\ker F$ physically.

(b) Interpret the affine leaves $H_{F,w}$.

(c) After choosing an inner product g, explain why $F^{\sharp_g}$ is normal to the equal-work leaves.

(d) Explain why a different inner product may produce a different force vector while leaving the force covector unchanged.