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

Version 0.2.0 (Geometric edition)

Guiding question

What geometric object is a covector before any metric identifies it with a vector?

PART

Front Matter

The complete route from scalar observation to geometric covector.

Lecture architecture

Geometric covector

1. maps as labeled fibers;
2. affine foliations and half-spaces;
3. basis labels and the dual basis;
4. scaling and addition.

Applications and enrichment

1. support functions;
2. supporting hyperplanes;
3. set inclusion by function comparison;
4. inner-product normal representation.

Takeaway

A covector is first a linear scalar observation and a labeled affine foliation. Optimization and inner products use that geometry in different ways.

Ambient assumptions

Let V be a finite-dimensional real vector space, $\dim V = n$, and let

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

Available structure

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

Not assumed

- norm (length) or inner product;
- angle or orthogonality;
- distance or unit vectors;
- a canonical identification $V \simeq V^*$.

Methodological principle

Takeaway

Do not assume **less** structure than **needed** and **more** than can be **justified**.

Geometric interpretation

Every construction will be audited by asking which structure it actually uses.

PART I

From Maps to Foliations

A map is represented inside its domain by the fibers carrying its values.



How can a map be represented inside its own domain?

Keep its fibers together with their labels.



Every map determines a labeled partition

Let $f : X \to Y$. For $y \in Y$,

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

The nonempty fibers partition X, and the labels recover the map:

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

Key idea

For each x, the value $f(x)$ is the label of the unique fiber containing x.

Real-valued maps produce level sets

For $f : X \to \mathbb{R}$, define

$$H_c(f) = \{x \in X : f(x) = c\}.$$

Map view

f assigns a scalar to every point.

Domain-internal view

The domain is decomposed into scalar-labeled level sets.

Conceptual warning

For a general nonlinear map, the level sets need not be affine, parallel, or even regular.

What does linearity add to the labeled-fiber picture?

Every nonzero covector produces a parallel codimension-one affine foliation.

A covector is a linear scalar observation

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

Geometric interpretation

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

Takeaway

A covector is not initially a row, a column, or a normal vector. It is a linear map to scalars.

The zero fiber is a hyperplane

Assume $\alpha \neq 0$. Then $\operatorname{im}\alpha = \mathbb{R}$, so rank–nullity gives

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

Key idea

$\ker \alpha$ is a linear hyperplane through the origin.

Conceptual warning

For the zero covector, $\ker 0 = V$; the codimension-one geometry degenerates.

A direct splitting

Choose $u \in V$ such that $\alpha(u) = 1$. Then

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

because every $x \in V$ admits the unique decomposition

$$\boxed{x = \underbrace{x - \alpha(x)u}_{\in \ker \alpha} + \underbrace{\alpha(x)u}_{\in \operatorname{span}\{u\}}.}$$

Indeed, $\alpha\big(x - \alpha(x)u\big) = 0$.

Geometric interpretation

The first term lies within the zero leaf, $\ker \alpha$.
The second selects the leaf through its scalar coefficient $\alpha(x)$.

Every fiber is an affine hyperplane

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

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

If $\alpha(x_c) = c$, then

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

Therefore

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

Takeaway

All leaves are parallel because they are translates of the same direction subspace $\ker \alpha$. Parallelism is affine, not metric.

The labeled affine foliation

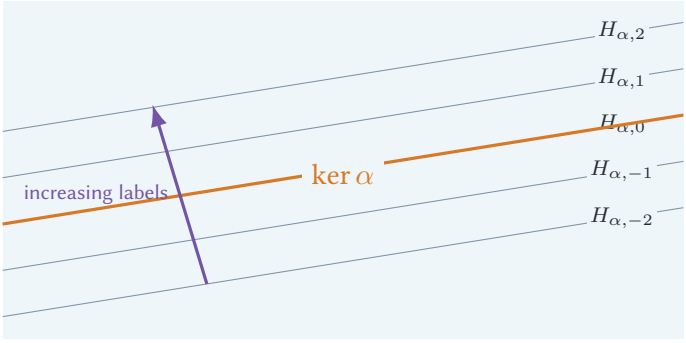

Conceptual warning

The apparent angles, spacings, and distances belong only to the drawing.
The labeled foliation defined by α specifies only the leaves, their parallelism, their labels, and the ordering induced by those labels.

PART II

II

Complete Geometric Content

The kernel determines the leaves; their linear labeling completes the covector.

Does the kernel determine a covector?

No. It leaves both sign and scale undetermined.



Direction: the common leaf direction

The kernel $\ker \alpha$ determines the common direction of all leaves $H_{\alpha,c}$.

Conceptual warning

The kernel alone does not determine:

- which side is positive;
- which scalar labels the nonzero leaves;
- the relative scale of those labels.

Coorientation: the positive transverse side

The covector distinguishes the two open half-spaces

$$\{x : \alpha(x) > 0\} \quad \text{and} \quad \{x : \alpha(x) < 0\}.$$

Geometric interpretation

This choice is a *coorientation* of the kernel hyperplane.

Conceptual warning

Coorientation is not an orientation of the hyperplane. It chooses a sign in the one-dimensional transverse quotient.

The labels determine orientation and scale

The covector α assigns to every leaf $H_{\alpha,c}$ its scalar label c.

Sign of the labels

The distinction between positive and negative labels determines the coorientation.

Magnitude of the labels

The numerical values of the labels determine the calibration of the foliation.

Quotient formulation

The induced linear isomorphism identifies the one-dimensional space of leaves with $\mathbb{R}$

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

Takeaway

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

Nested half-spaces

Define

$$H^-_{\alpha,c} = \{x : \alpha(x) \le c\}, \quad H^+_{\alpha,c} = \{x : \alpha(x) \ge c\}.$$

Then

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

Geometric interpretation

A covector is equivalently a linearly labeled nested family of affine half-spaces.

What captures the transverse leaf information?

The quotient space of leaves, not any selected transverse vector.



A transverse vector parameterizes the leaves

Choose $u \notin \ker \alpha$ and set $a = \alpha(u) \neq 0$. Then

$$V = \ker \alpha \oplus \operatorname{span}\{u\}, \quad \alpha(h + \lambda u) = \lambda a.$$

Hence

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

Conceptual warning

Transverse means complementary: $u \notin \ker \alpha$. It does not mean perpendicular.

Normalization without a norm

Rescaling a transverse vector gives a representative satisfying

$$\alpha(u) = 1.$$

Key idea

This is an algebraic normalization: the selected point lies on the leaf labeled one.

Conceptual warning

It does not make u a unit vector. No notion of length has been introduced.

The quotient is the leaf space

The quotient $V/\ker\alpha$ consists exactly of the affine cosets parallel to $\ker\alpha$. The map

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

is a well-defined linear isomorphism.

Takeaway

A nonzero covector identifies the one-dimensional space of leaves with $\mathbb{R}$. A transverse vector only chooses representatives.

Reconstructing a covector from geometric data

Let $H \subset V$ be a hyperplane. A nonzero covector with kernel H is fixed by:

1. the hyperplane H;
2. one nonzero calibration $\alpha(u) = a$ for some $u \notin H$.

Geometric interpretation

For $x = h + \lambda u$, linearity then forces $\alpha(x) = \lambda a$.

PART III

III

Basis Data and the Dual Basis

Finite labels reconstruct the complete geometric covector.

How many labels determine a covector on an n-dimensional space?

Exactly its values on any basis.



Basis labels determine every value

Let $\mathcal{B} = (v_1, \dots, v_n)$ be a basis and write $x = \sum_i x^i v_i$. Then

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

Takeaway

The n numbers $\alpha(v_1), \dots, \alpha(v_n)$ determine every point label and therefore the complete affine foliation.

Arbitrary basis labels define one covector

Given $a_1, \dots, a_n \in \mathbb{R}$, define, for every $x = \sum_{i=1}^{n} x^i v_i \in V$,

$$\boxed{\alpha(x) = \sum_{i=1}^{n} a_i x^i.}$$

In particular, $\alpha(v_i) = a_i$. Thus

$$\mathcal{C}_{\mathcal{B}} : V^* \to \mathbb{R}^n, \quad \alpha \mapsto \big(\alpha(v_1), \dots, \alpha(v_n)\big)$$

is a linear isomorphism.

Conceptual warning

$\dim V^* = \dim V = n$ does not canonically identify V with V^*.

Recovering the kernel by label cancellation

Set $a_i = \alpha(v_i)$, and choose p with $a_p \neq 0$. For $i \neq p$, define

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

Then $\alpha(h_i) = 0$, and the $n-1$ vectors h_i, $i \neq p$, form a basis of $\ker \alpha$.

Geometric interpretation

Choose a transverse vector v_p from the given basis.
Displace each other basis vector along the selected transverse direction until its label cancels.
The resulting $n-1$ vectors are a basis of the kernel of the covector.

What the basis labels recover

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

Since $v_p \notin \ker \alpha$, every $x \in V$ has a unique decomposition

$$x = h + \lambda v_p, \quad h \in \ker \alpha,$$

and therefore

$$\alpha(x) = \lambda \alpha(v_p).$$

Zero-leaf geometry

The cancellation vectors reconstruct the kernel hyperplane.

Labelling all leaves

The leaf with label c is

$$H_{\alpha,c} = \frac{c}{\alpha(v_p)} v_p + \ker \alpha.$$

What geometric objects form the basis dual to a vector basis?

Calibrated foliations that isolate one basis coefficient at a time.

Geometric definition of a dual covector

Let $\mathcal{B} = (v_1, \dots, v_n)$ be a basis, for each i, let

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

There is a unique covector $\varepsilon^i \in V^*$ satisfying

$$\boxed{\ker \varepsilon^i = H_i, \quad \varepsilon^i(v_i) = 1.}$$

Equivalently,

$$\varepsilon^i(v_j) = \delta^i_j.$$

Geometric interpretation

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 λv_i and the calibrating vector v_i.

One dual covector isolates one coordinate

From

$$V = H_i \oplus \operatorname{span}\{v_i\}, \quad x = h_i + x^i v_i,$$

the vector x belongs to the leaf

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

Geometric interpretation

The leaf is labeled by the relative scale x^i of its intersection $x^i v_i$ with the span of v_i.

Takeaway

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

The dual covectors recover the vector

Each dual basis covector selects the leaf containing x in its own affine foliation:

$$x \in x^i v_i + H_i.$$

Key idea

The n selected leaves recover x as their unique common point:

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

Conceptual warning

This uses only direct sums and affine intersections, not orthogonality.

The dual covectors form a basis

The family $\mathcal{B}^* = (\varepsilon^1, \dots, \varepsilon^n)$ is linearly independent and spans V^*. Every covector satisfies

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

Takeaway

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

Basis evaluation and dual-basis synthesis

The basis $\mathcal{B} = (v_1, \dots, v_n)$ defines the evaluation map

$$\mathcal{C}_{\mathcal{B}} : V^* \to \mathbb{R}^n, \quad \alpha \mapsto \big(\alpha(v_1), \dots, \alpha(v_n)\big),$$

while the dual basis $\mathcal{B}^* = (\varepsilon^1, \dots, \varepsilon^n)$ defines the synthesis map

$$\mathcal{S}_{\mathcal{B}^*} : \mathbb{R}^n \to V^*, \quad (a_1, \dots, a_n) \mapsto \sum_{i=1}^{n} a_i \varepsilon^i.$$

Takeaway

These maps are inverse:

$$\boxed{\mathcal{S}_{\mathcal{B}^*} = \mathcal{C}_{\mathcal{B}}^{-1}, \quad \alpha = \sum_{i=1}^{n} \alpha(v_i)\varepsilon^i.}$$

PART IV

Vector-Space Operations

Scaling relabels a foliation; addition combines labels pointwise.

What does multiplying a covector by a scalar do geometrically?

It relabels the same foliation, with a sign reversal exactly for negative scalars.

Positive scaling

For $r > 0$,

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

Key idea

The unlabeled foliation and coorientation remain fixed; every label is multiplied by r.

Negative and zero scaling

If $r < 0$, the unlabeled foliation is unchanged, the labels are rescaled, and the coorientation reverses. If $r = 0$, every point receives label zero.

Takeaway

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

Scaling changes labels, not metric spacing

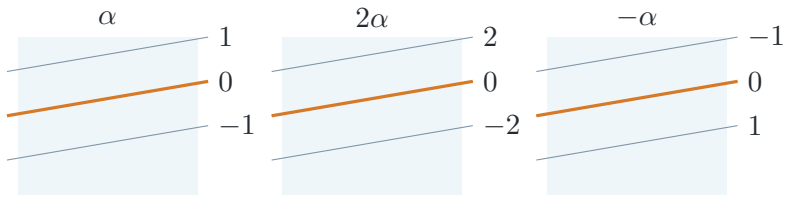

Conceptual warning

No statement about distances between leaves is intended. Only the labels change.

How are two geometric covectors added?

Add their labels at every point, then read the new calibrated foliation.



Addition is pointwise label addition

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

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

Therefore

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

Geometric interpretation

The new zero leaf contains exactly the points where the two scalar labels cancel.

Dependent covectors

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

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

Conceptual warning

The case $r = -1$ collapses the complete foliation into the single zero fiber V.

Independent covectors: select a shear direction

Assume α and β are linearly independent. Choose a vector in $\ker \beta$ and transverse to $\ker \alpha$

$$u_\beta \in \ker \beta \quad \text{with} \quad \alpha(u_\beta) \neq 0.$$

Every $v \in V$ is uniquely expressible as

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

Geometric interpretation

The direction u_β preserves the β-label while changing the α-label.

The new kernel is a shear of the old one

Imposing $(\alpha + \beta)(v) = 0$ gives

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

Hence

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

Takeaway

The resulting kernel is obtained by a metric-free linear shear of $\ker \alpha$.

The shear construction step by step

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.

Conceptual warning

The kernels alone do not determine the kernel of the sum. Addition depends on the relative calibration.

PART V

V

Geometric Audit and Support Functions

The labeled-foliation picture becomes a direct tool in optimization.

Which statements used no metric?

All constructions based on linearity, affine geometry, quotient structure, and label arithmetic.



What required no inner product

Linear and affine structure

- kernels and codimension;
- affine translates and parallelism;
- labeled foliations and half-spaces;
- quotient and transverse parameterizations.

Dual and algebraic structure

- basis labels and dual basis;
- coordinate-isolating leaves;
- scalar relabeling;
- pointwise addition and shear.

Conceptual warning

A metric interpretation of angles, lengths, distances, orthogonality, or normal vectors has not yet been used.

How far does a set extend in one covector direction?

Read the largest label attained by the set.



Support function: type and definition

Let $C \subseteq V$ be nonempty and define

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

The support function of C is the map

$$\boxed{\sigma_C : V^* \to \overline{\mathbb{R}}, \quad \sigma_C(\alpha) = \sup_{x \in C} \alpha(x).}$$

Geometric interpretation

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.

The extremal label selects a supporting leaf

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

$$\alpha(x_\star) = \sigma_C(\alpha), \quad \alpha(x) \le \sigma_C(\alpha) \quad \forall x \in C.$$

Hence

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

The contact set is

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

Takeaway

The maximization selects both the extremal label and its supporting leaf.

Support function as the last attained label

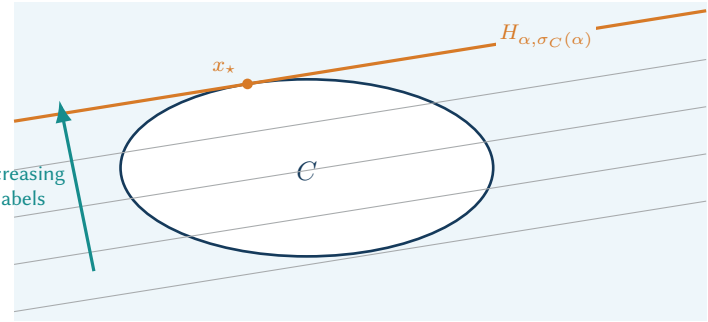

Takeaway

support function = largest covector label attained on the set.

Set inclusion becomes function comparison

If $C \subseteq D$, then

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

For nonempty closed convex sets, separation gives the converse:

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

Geometric interpretation

In every covector direction, the last leaf reached by C must not lie beyond the corresponding last leaf reached by D.

What support functions remember

For nonempty closed convex sets,

$$C = D \quad \Longleftrightarrow \quad \sigma_C = \sigma_D.$$

For an arbitrary nonempty set,

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

Key idea

The support function encodes the closed convex hull, not the internal nonconvex structure of the original set.

Support functions inherit label arithmetic

For $r \geq 0$,

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

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

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

Takeaway

Whenever finite, σ_C is positively homogeneous and subadditive on V^*. These properties come directly from scaling and adding covector labels.

PART VI

VI

Inner-Product Representation

An inner product turns the labeled foliation into a normal-vector picture.

What does an inner product add?

A normal-vector representation that reproduces every covector label.



An inner product supplies the sharp map

Let g be an inner product on V. Define

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

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

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

It is the unique vector satisfying

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

Why the representing vector is normal

If $h \in \ker \alpha$, then

$$g(\alpha^{\sharp_g}, h) = \alpha(h) = 0.$$

Therefore

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

Direction

Orthogonality selects the transverse line $(\ker \alpha)^{\perp_g}$.

Sign and scale

Reproducing all labels selects the unique vector $\alpha^{\sharp_g}$ on that line.

The normal vector encodes every leaf

Define the label-one transverse vector

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

Then

$$\boxed{H_{\alpha,c} = c\, u_g + \ker \alpha = \{x \in V : g(\alpha^{\sharp_g}, x) = c\}.}$$

Conceptual warning

u_g is normalized by its covector label, not by its g-length.

One covector, many normal-vector representations

Fixed geometric object

$$\alpha : V \to \mathbb{R}$$

$$\ker \alpha, \quad H_{\alpha,c}, \quad \text{labels}$$

Representation chosen by g

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

$$g(\alpha^{\sharp_g}, \cdot) = \alpha$$

Takeaway

Changing g changes the normal-vector representation, while the covector, its leaves, and their labels remain unchanged.

VII

PART VII

Final Synthesis

One geometric object, several compatible descriptions and uses.

What has been described in several languages?

A linear scalar observation and its labeled affine foliation.

Four views of one covector

Takeaway

map view: $\alpha : V \to \mathbb{R}$ linear,

foliation view: parallel affine hyperplanes with linear labels,

quotient view: $\overline{\alpha} : V/\ker\alpha \xrightarrow{\sim} \mathbb{R}$,

basis view: $\alpha = \sum_{i=1}^{n} \alpha(v_i)\varepsilon^i$.

The geometric picture organizes every construction

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

One continuous chain

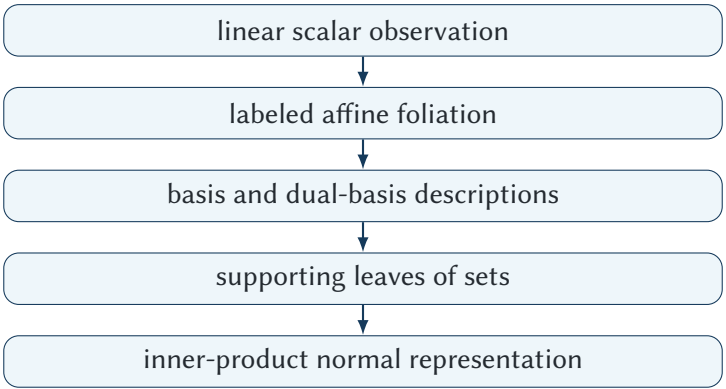

Final conclusion

Takeaway

A nonzero covector is a linear scalar observation, represented geometrically by a parallel family of affine hyperplanes with linear labels.

Optimization

The support function selects the largest label attained on a set and hence its supporting leaf.

Inner-product geometry

An inner product represents the same covector by a normal vector that reproduces every label.

Contact

Antonio Franchi

University of Twente
Sapienza University of Rome

https://r-franchi.eu/

Corrections, suggestions, and discussion

schol@r-franchi.eu