

The Metric-based Compactness Hypothesis and Method of Potential Functions for Machine Learning

Oleg Seredin

Tula State University
Tula, Russia

Vadim Mottl

Computing Centre of the Russian
Academy of Sciences
Moscow, Russia



A conceptual basis of dependency estimation: The compactness hypothesis

The set of real-world objects	$\omega \in \Omega$
The hidden characteristic of an object (goal characteristic)	$y \in \mathbb{Y}$
The sought-for decision rule	$\hat{y}(\omega): \Omega \rightarrow \mathbb{Y}$

The main idea:

A metric is to be chosen in the set of real-world objects

$$\rho(\omega', \omega'') : \Omega \times \Omega \rightarrow \mathbb{R}$$

$$\rho(\omega', \omega'') = \rho(\omega'', \omega') \geq 0, \quad \rho(\omega', \omega'') > 0 \text{ if } \omega' \neq \omega'', \quad \rho(\omega', \omega'') + \rho(\omega'', \omega''') \geq \rho(\omega', \omega''')$$

The choice is appropriate if the metric meets **the compactness hypothesis**

(Emmanuel Braverman. Ph.D. Thesis: “Experiments on learning a machine to distinguish patterns. 1961, Institute of Automation and Remote Control, Moscow):

If two objects $\omega', \omega'' \in \Omega$ are similar to each other in the sense of the chosen metric

$\rho(\omega', \omega'') \cong 0$, their values of the goal characteristic are mostly almost the same $y(\omega') \cong y(\omega'')$.

Hence, the decision rule follows:

For close objects $\rho(\omega', \omega'') \cong 0$, close decisions are to be made

$$\hat{y}(\omega') = \hat{y}(\omega'') \quad \text{in the problem of pattern recognition } \mathbb{Y} = \{y_1, \dots, y_m\}$$

$$\hat{y}(\omega') \cong \hat{y}(\omega'') \quad \text{in the problem of regression estimation } \mathbb{Y} = \mathbb{R}$$

Dipole in a metric space

The metric space of real-world objects: $\omega \in \Omega$, $\rho(\omega', \omega'')$ – the metric

A dipole in the metric space – an ordered pair: $\langle \alpha_{-1}, \alpha_1 \rangle \in \Omega \times \Omega$

A simplest instantiation of the compactness hypothesis

Membership of an arbitrary object $\omega \in \Omega$ in one of two classes

$\hat{y}(\omega) = \begin{cases} 1 & \text{if } \rho(\alpha_1, \omega) < \rho(\alpha_{-1}, \omega), \\ -1 & \text{if } \rho(\alpha_1, \omega) > \rho(\alpha_{-1}, \omega). \end{cases}$ There are too few elements in the set of objects Ω . Besides, only a finite training set of objects is accessible to the observer $\{\omega_j, j = 1, \dots, N\}$

How should the dipole be chosen?

A more “delicate” realization of the compactness hypothesis in a dense hull of the initially sparse metric space

A hypothetical dense space, in which the set of real-world objects is subset of, maybe, isolated elements: $\tilde{\Omega} \supset \Omega$. The dipole is to be chosen in this dense hull.

$\alpha_{-1}, \alpha_1 \in \tilde{\Omega}$, $\mathcal{H}(\alpha_{-1}, \alpha_1) = \{\vartheta \in \tilde{\Omega} : \rho(\alpha_{-1}, \vartheta) = \rho(\alpha_1, \vartheta)\}$ – metric “hyperplane” in $\tilde{\Omega}$


$\omega_{\mathcal{H}}(\alpha_{-1}, \alpha_1) \in \mathcal{H}(\alpha_{-1}, \alpha_1)$ – projection of $\omega \in \Omega$ onto the “hyperplane” in $\tilde{\Omega}$

Score function: sign-dependent distance of the point from the “hyperplane” in $\tilde{\Omega}$

$$d(\omega | \alpha_{-1}, \alpha_1) = \begin{cases} \rho(\omega_{\mathcal{H}}, \omega) & \text{if } \rho(\alpha_1, \omega) \leq \rho(\alpha_{-1}, \omega), \\ -\rho(\omega_{\mathcal{H}}, \omega) & \text{if } \rho(\alpha_1, \omega) > \rho(\alpha_{-1}, \omega). \end{cases}$$

This is an analog of the linear approach to dependence estimation:

All the decisions on the hidden characteristic y of a real-world object ω are to be made only from the sign-dependent distance $d(\omega | \alpha_{-1}, \alpha_1)$.



Dipole in a metric space

The metric space of real-world objects: $\omega \in \Omega$, $\rho(\omega', \omega'')$ – the metric

A dipole in the metric space – an ordered pair: $\langle \alpha_{-1}, \alpha_1 \rangle \in \Omega \times \Omega$

A simplest instantiation of the compactness hypothesis

Membership of an arbitrary object $\omega \in \Omega$ in one of two classes

$\hat{y}(\omega) = \begin{cases} 1 & \text{if } \rho(\alpha_1, \omega) < \rho(\alpha_{-1}, \omega), \\ -1 & \text{if } \rho(\alpha_1, \omega) > \rho(\alpha_{-1}, \omega). \end{cases}$ There are too few elements in the set of objects Ω . Besides, only a finite training set of objects is accessible to the observer $\{\omega_j, j = 1, \dots, N\}$

How should the dipole be chosen?

A more “delicate” realization of the compactness hypothesis in a dense hull of the initially sparse metric space

A hypothetical dense space, in which the set of real-world objects is subset of, maybe, isolated elements: $\tilde{\Omega} \supset \Omega$. The dipole is to be chosen in this dense hull.

$\alpha_{-1}, \alpha_1 \in \tilde{\Omega}$, $\mathcal{H}(\alpha_{-1}, \alpha_1) = \{\vartheta \in \tilde{\Omega} : \rho(\alpha_{-1}, \vartheta) = \rho(\alpha_1, \vartheta)\}$ – metric “hyperplane” in $\tilde{\Omega}$


$\omega_{\mathcal{H}}(\alpha_{-1}, \alpha_1) \in \mathcal{H}(\alpha_{-1}, \alpha_1)$ – projection of $\omega \in \Omega$ onto the “hyperplane” in $\tilde{\Omega}$

Score function: sign-dependent distance of the point from the “hyperplane” in $\tilde{\Omega}$

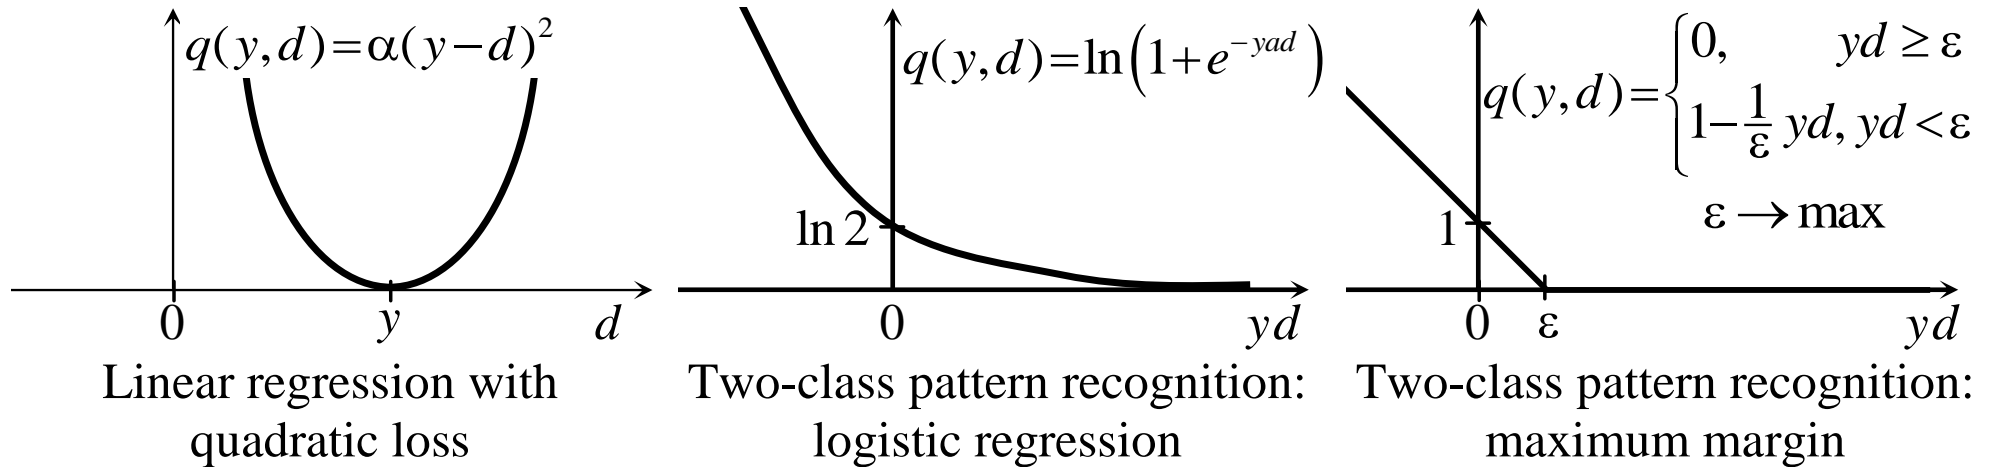
$$d(\omega | \alpha_{-1}, \alpha_1) = \begin{cases} \rho(\omega_{\mathcal{H}}, \omega) & \text{if } \rho(\alpha_1, \omega) \leq \rho(\alpha_{-1}, \omega), \\ -\rho(\omega_{\mathcal{H}}, \omega) & \text{if } \rho(\alpha_1, \omega) > \rho(\alpha_{-1}, \omega). \end{cases}$$

$$\text{Classification: } \hat{y}(\omega) = \begin{cases} 1 & \text{if } d(\omega | \alpha_{-1}, \alpha_1) + b > 0, \\ -1 & \text{if } d(\omega | \alpha_{-1}, \alpha_1) + b < 0. \end{cases}$$

$$\text{Regression: } \hat{y}(\omega) = a d(\omega | \alpha_{-1}, \alpha_1) + b$$



Hinge functions: A bridge to particular problems of dependence estimation



But first, the initially sparse metric space of real-world objects is to be embedded into a dense hull $\tilde{\Omega} \supseteq \Omega$, in which the dipole could be continuously chosen.

A way:

Embedding an arbitrary metric space into a linear space.

However, it will not be a proper metric space.

The metric will be defined in it not for all pairs of elements.

But the initial metric space will be embedded isometrically.

An instrument of embedding a metric space into a linear space: Commonality of two elements of a metric space

Metric space Ω with metric $\rho(\omega', \omega'')$:

$$\rho(\omega, \omega) = 0, \quad \rho(\omega', \omega'') > 0 \text{ if } \omega' \neq \omega''; \quad \rho(\omega', \omega'') = \rho(\omega'', \omega');$$

$$\rho(\omega', \omega'') + \rho(\omega'', \omega''') \geq \rho(\omega', \omega''') \text{ – triangle inequality}$$

Let us choose an element of the metric space $\phi \in \Omega$ as its “center”

Two-argument commonality function $K_\phi(\omega', \omega''): \Omega \times \Omega \rightarrow \mathbb{R}$

$$K_\phi(\omega', \omega'') = \frac{1}{2} \left[\rho^2(\omega', \phi) + \rho^2(\omega'', \phi) - \rho^2(\omega', \omega'') \right]$$

Theorem 1. Properties of the commonality function:

$$K_\phi(\omega', \omega'') = K_\phi(\omega'', \omega') \text{ – symmetry}$$

$$K_\phi(\omega, \omega) = \rho^2(\omega, \phi) \geq 0 \text{ – non-negativity for } \omega' = \omega''$$

$$|K_\phi(\omega', \omega'')| \leq \sqrt{K_\phi(\omega', \omega')} \sqrt{K_\phi(\omega'', \omega'')} \text{ – inequality of Cauchy-Bunyakowsky type}$$

$$K_{\tilde{\phi}}(\omega', \omega'') = K_\phi(\omega', \omega'') - K_\phi(\omega', \tilde{\phi}) - K_\phi(\omega'', \tilde{\phi}) + K_\phi(\tilde{\phi}, \tilde{\phi}) \text{ – the rule of center translation}$$

$$\rho^2(\omega', \omega'') = K_\phi(\omega', \omega') + K_\phi(\omega'', \omega'') - 2K_\phi(\omega', \omega'') \text{ – return to the metric}$$

It resembles very much the inner product!

But there are no linear operations as yet!?

Commonality matrix for a finite metric space

Let us assume, for simplicity sake, that the set of real-world objects with metric $\rho(\omega', \omega'')$ is finite $|\Omega| = M$, $\Omega = \{\omega_1, \dots, \omega_M\}$.

$\mathbf{K}_\phi = [K_\phi(\omega_i, \omega_j), i, j = 1, \dots, M]$ – the symmetric commonality matrix for some center $\phi \in \Omega$

$\xi_{\phi,1} \in \mathbb{R}, \dots, \xi_{\phi,N} \in \mathbb{R}$

eigen values are real numbers

$\underbrace{\xi_{\phi,1} \geq 0, \dots, \xi_{\phi,p_\phi} \geq 0}_{p_\phi}, \underbrace{\xi_{\phi,p_\phi+1} < 0, \dots, \xi_{\phi,M} < 0}_{q_\phi}$

for an arbitrary metric, the matrix may be not positive definite

$p_\phi + q_\phi = M$ – signature of the commonality

Theorem 2. *The signature of matrix \mathbf{K}_ϕ does not depend on the choice of the center in the metric space $\phi \in \Omega$.*

Thus, any metric on a finite metric space $|\Omega| = M$ is characterized by its signature $p + q = M$.

Definition. *Metric is said to be proto-Euclidean one, if for any subset of objects*

$\{\omega_1, \dots, \omega_N\}$ matrix $\mathbf{P}_N = [-\rho^2(\omega_j, \omega_l), j, l = 1, \dots, N]$ is conditionally positive semidefinite:

$$\mathbf{c}^T \mathbf{P}_N \mathbf{c} \geq 0, \quad \mathbf{1}^T \mathbf{c} = \sum_{j=1}^N c_j = 0.$$

Theorem 3. *The signature of a proto-Euclidean metric $p = M, q = 0$.*

In other words, the commonality matrix of a finite metric space with a proto-Euclidean metric (finite proto-Euclidean metric space) is positive semidefinite for any choice of the center.

Indefinite inner product and pseudo-Euclidean linear space

Let metric $\rho(\omega', \omega'')$ on the set of objects $|\Omega| = M$ be of signature $p + q = M$.

Let the following two-argument real-valued function be defined over the entire linear space \mathbb{R}^M :

$$K(\mathbf{x}', \mathbf{x}'') = \mathbf{x}'^T \mathbf{J}_{p,q} \mathbf{x}'' : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$$

signature depends only on the metric

$$\mathbf{J}_{p,q} = \begin{pmatrix} \mathbf{I}_{p \times p} & \mathbf{0}_{p \times q} \\ \mathbf{0}_{q \times p} & -\mathbf{I}_{q \times q} \end{pmatrix} (M \times M)$$

Properties:

Symmetry

$$K(\mathbf{x}', \mathbf{x}'') = K(\mathbf{x}'', \mathbf{x}')$$

Bilinearity

$$K(c'\mathbf{x}' + c''\mathbf{x}'', \mathbf{x}''') = c'K(\mathbf{x}', \mathbf{x}''') + c''K(\mathbf{x}'', \mathbf{x}''')$$

The property of non-negativity when the arguments coincide is absent

$$K(\mathbf{x}, \mathbf{x}) < 0 \text{ for some } \mathbf{x} \in \mathbb{R}^M$$

$$K(\mathbf{x}, \mathbf{x}) = \mathbf{x}^T \mathbf{J}_{p,q} \mathbf{x} = (\mathbf{u}^T \ \mathbf{v}^T) \begin{pmatrix} \mathbf{I}_{p \times p} & \mathbf{0}_{p \times q} \\ \mathbf{0}_{q \times p} & -\mathbf{I}_{q \times q} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} = \mathbf{u}^T \mathbf{u} - \mathbf{v}^T \mathbf{v} < 0 \text{ if } \mathbf{u}^T \mathbf{u} < \mathbf{v}^T \mathbf{v}$$

This is no inner product!

Such a function is said to be indefinite inner product, and the linear space \mathbb{R}^M – pseudo-Euclidean linear space.

A finite-dimensional linear space with indefinite inner product is called the Krein space.

Mark Krein (1907-1989),

Professor of Odessa Civil-Engineering Institute.

Vector length and distance between vectors in the pseudo-Euclidean linear space

The pseudo-Euclidean linear space is neither normed (not all elements have norm) nor metric (metric is defined not for all pairs of elements).

Nevertheless, it is possible to isometrically embed an arbitrary finite metric space of signature $p + q = M$ into a pseudo-Euclidean linear space of the same signature.

Isometric embedding of an arbitrary metric space into a pseudo-Euclidean linear space

Let $\Omega = \{\omega_1, \dots, \omega_M\}$ be a finite metric space with metric $\rho(\omega', \omega'')$, $p + q = M$ be its signature, and $\phi \in \Omega$ be an arbitrary element assigned as the center.

Theorem 4. *There exist M real-valued vectors $\mathbf{x}_{\phi,j} \in \mathbb{R}^M$, $j = 1, \dots, N$, in the pseudo-Euclidean linear space of signature $p + q = M$, for which are defined:*

<i>the norm</i>	$r(\mathbf{x}_{\phi,j}, \mathbf{0}) = \underbrace{(\mathbf{x}_{\phi,j}^T \mathbf{J}_{p,q} \mathbf{x}_{\phi,j})}_{\geq 0}^{1/2} = \rho(\omega_j, \phi)$	<i>distance to the zero point</i>
<i>the metric</i>	$r(\mathbf{x}_{\phi,j}, \mathbf{x}_{\phi,l}) = \underbrace{[(\mathbf{x}_{\phi,j} - \mathbf{x}_{\phi,l})^T \mathbf{J}_{p,q} (\mathbf{x}_{\phi,j} - \mathbf{x}_{\phi,l})]}_{\geq 0}^{1/2} = \rho(\omega_j, \omega_l)$	<i>distance between pairs of vectors</i>

Discriminant (score) functions in a pseudo-Euclidean linear space


Discriminant dipole:

An ordered pair of vectors $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$, $\alpha_{-1}, \alpha_1 \in \mathbb{R}^M$ – the nodes of the dipole. In what follows, we shall consider only dipoles, for which the metric distance between the nodes is defined, i.e., the squared distance between the nodes is positive

$$r^2(\alpha_{-1}, \alpha_1) = (\alpha_{-1} - \alpha_1)^T \mathbf{J}_p (\alpha_{-1} - \alpha_1) > 0$$

The axis defined by a dipole:

$$\{\mathbf{x}_c \in \mathbb{R}^M: \mathbf{x}_c = (1-c)\alpha_{-1} + c\alpha_1 \in \mathbb{R}^M, c \in \mathbb{R}\} \subset \mathbb{R}^M$$



Discriminant functions in a pseudo-Euclidean linear space

Discriminant dipole:


An ordered pair of vectors $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$, $\alpha_{-1}, \alpha_1 \in \mathbb{R}^M$ – the nodes of the dipole. In what follows, we shall consider only dipoles, for which the metric distance between the nodes is defined, i.e., the squared distance between the nodes is positive

$$r^2(\alpha_{-1}, \alpha_1) = (\alpha_{-1} - \alpha_1)^T \mathbf{J}_p (\alpha_{-1} - \alpha_1) > 0$$

The axis defined by a dipole:

$$\{\mathbf{x}_c \in \mathbb{R}^M: \mathbf{x}_c = (1-c)\alpha_{-1} + c\alpha_1 \in \mathbb{R}^M, c \in \mathbb{R}\} \subset \mathbb{R}^M$$

$\hat{\mathbf{x}}(\alpha_{-1}, \alpha_1)$ – projection of a vector $\mathbf{x} \in \mathbb{R}^M$ on the axis defined by the dipole



Discriminant functions in a pseudo-Euclidean linear space

Discriminant dipole:

An ordered pair of vectors $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$, $\alpha_{-1}, \alpha_1 \in \mathbb{R}^M$ – the nodes of the dipole. In what follows, we shall consider only dipoles, for which the metric distance between the nodes is defined, i.e., the squared distance between the nodes is positive

$$r^2(\alpha_{-1}, \alpha_1) = (\alpha_{-1} - \alpha_1)^T \mathbf{J}_p (\alpha_{-1} - \alpha_1) > 0$$

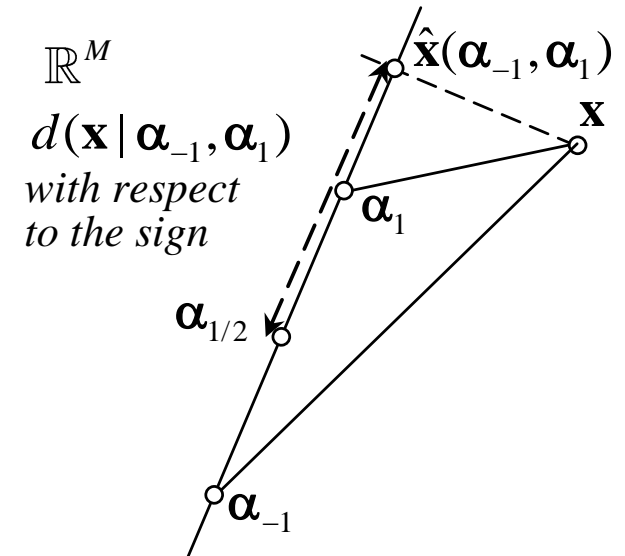
The axis defined by a dipole:

$$\{\mathbf{x}_c \in \mathbb{R}^M: \mathbf{x}_c = (1-c)\alpha_{-1} + c\alpha_1 \in \mathbb{R}^M, c \in \mathbb{R}\} \subset \mathbb{R}^M$$

$\hat{\mathbf{x}}(\alpha_{-1}, \alpha_1)$ – projection of a vector $\mathbf{x} \in \mathbb{R}^M$ on the axis defined by the dipole

Theorem 5. If $r^2(\alpha_{-1}, \alpha_1) > 0$ then $d^2(\mathbf{x} | \alpha_{-1}, \alpha_1) \geq 0$ for any vector $\mathbf{x} \in \mathbb{R}^M$, and, with respect to the sign

$$d(\mathbf{x} | \alpha_{-1}, \alpha_1) = \frac{1}{2} \left(r^2(\alpha_{-1}, \mathbf{x}) - r^2(\alpha_1, \mathbf{x}) \right) \frac{1}{r(\alpha_{-1}, \alpha_1)}.$$



Discriminant functions in a pseudo-Euclidean linear space

Discriminant dipole:

An ordered pair of vectors $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$, $\alpha_{-1}, \alpha_1 \in \mathbb{R}^M$ – the nodes of the dipole. In what follows, we shall consider only dipoles, for which the metric distance between the nodes is defined, i.e., the squared distance between the nodes is positive

$$r^2(\alpha_{-1}, \alpha_1) = (\alpha_{-1} - \alpha_1)^T \mathbf{J}_p (\alpha_{-1} - \alpha_1) > 0$$

The axis defined by a dipole:

$$\{\mathbf{x}_c \in \mathbb{R}^M: \mathbf{x}_c = (1-c)\alpha_{-1} + c\alpha_1 \in \mathbb{R}^M, c \in \mathbb{R}\} \subset \mathbb{R}^M$$

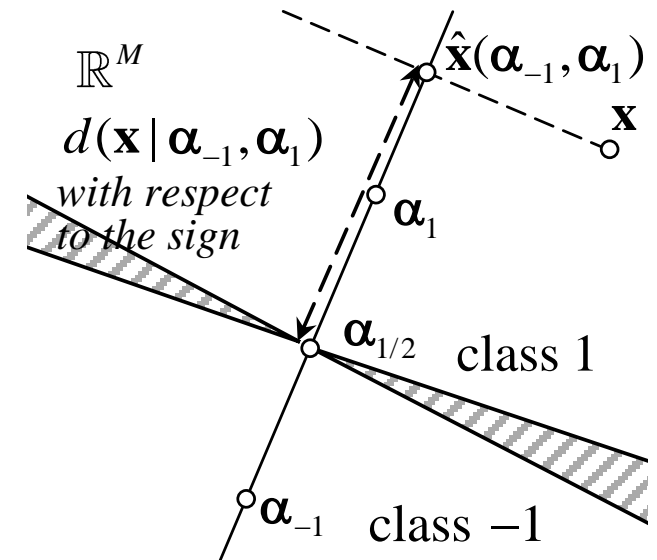
$\hat{\mathbf{x}}(\alpha_{-1}, \alpha_1)$ – projection of a vector $\mathbf{x} \in \mathbb{R}^M$ on the axis defined by the dipole

Theorem 5. If $r^2(\alpha_{-1}, \alpha_1) > 0$ then $d^2(\mathbf{x} | \alpha_{-1}, \alpha_1) \geq 0$ for any vector $\mathbf{x} \in \mathbb{R}^M$, and, with respect to the sign

$$d(\mathbf{x} | \alpha_{-1}, \alpha_1) = \frac{1}{2} \left(r^2(\alpha_{-1}, \mathbf{x}) - r^2(\alpha_1, \mathbf{x}) \right) \frac{1}{r(\alpha_{-1}, \alpha_1)}.$$

Parametric family of discriminant functions:

$$d(\mathbf{x} | \alpha_{-1}, \alpha_1) \begin{cases} > 0 \Rightarrow \text{the positive part,} \\ = 0 \Rightarrow \text{the neutral part,} \\ < 0 \Rightarrow \text{the negative part.} \end{cases}$$



Training: Choice of the discriminant dipole

A training set in the metric space of real-world objects $\omega \in \Omega$:

$\{(\omega_j, y_j), j = 1, \dots, N\}$, distances $[\rho(\omega_j, \omega_l), j, l = 1, \dots, N]$, classes $y_j = y(\omega_j)$.

The center $\phi \in \Omega$ + a real-world object $\omega \in \Omega \Rightarrow \mathbf{0} \in \mathbb{R}^M$ + vector $\mathbf{x}_{\phi, \omega} \in \mathbb{R}^M$
(in the pseudo-Euclidean space).

The training process boils down to the choice of the dipole $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$:

We shall consider the dipole's nodes as linear (affine) combinations of the vectors $\{\mathbf{x}_{\phi, \omega_1}, \dots, \mathbf{x}_{\phi, \omega_N}\}$, into which the objects of the training set are mapped:

$$\begin{aligned} \alpha_{-1} &= \sum_{j=1}^N c_{-1,j} \mathbf{x}_{\phi, \omega_j}, \quad \sum_{j=1}^N c_{-1,j} = 1, \\ \alpha_1 &= \sum_{j=1}^N c_{1,j} \mathbf{x}_{\phi, \omega_j}, \quad \sum_{j=1}^N c_{1,j} = 1, \end{aligned} \quad a_j = c_{1,j} - c_{-1,j}, \quad \sum_{j=1}^N a_j = 0.$$

Score function:
$$d(\mathbf{x} | \alpha_{-1}, \alpha_1) = \frac{1}{2} \left(r^2(\alpha_{-1}, \mathbf{x}) - r^2(\alpha_1, \mathbf{x}) \right) \frac{1}{r(\alpha_{-1}, \alpha_1)}$$

Theorem 6: For any real-world object $\omega \in \Omega$, any training set $\{\omega_j, j = 1, \dots, N\} \subset \Omega$, and any $\phi \in \Omega$:

$$\boxed{r^2(\alpha_{-1}, \mathbf{x}_{\phi, \omega}) - r^2(\alpha_1, \mathbf{x}_{\phi, \omega}) = \sum_{l=1}^N \left(-\rho^2(\omega_j, \omega_l) \right) a_j + b, \quad \boxed{r^2(\alpha_{-1}, \alpha_1) = \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N \left(-\rho^2(\omega_j, \omega_l) \right) a_j a_l.}$$

Training: Choice of the discriminant dipole

A training set in the metric space of real-world objects $\omega \in \Omega$:

$\{(\omega_j, y_j), j = 1, \dots, N\}$, distances $[\rho(\omega_j, \omega_l), j, l = 1, \dots, N]$, classes $y_j = y(\omega_j)$.

The center $\phi \in \Omega$ + a real-world object $\omega \in \Omega \Rightarrow \mathbf{0} \in \mathbb{R}^M$ + vector $\mathbf{x}_{\phi, \omega} \in \mathbb{R}^M$
(in the pseudo-Euclidean space).

The training process boils down to the choice of the dipole $\langle \alpha_{-1}, \alpha_1 \rangle \in \mathbb{R}^M \times \mathbb{R}^M$:

We shall consider the dipole's nodes as linear (affine) combinations of the vectors $\{\mathbf{x}_{\phi, \omega_1}, \dots, \mathbf{x}_{\phi, \omega_N}\}$, into which the objects of the training set are mapped:

$$\begin{aligned} \alpha_{-1} &= \sum_{j=1}^N c_{-1,j} \mathbf{x}_{\phi, \omega_j}, \quad \sum_{j=1}^N c_{-1,j} = 1, \\ \alpha_1 &= \sum_{j=1}^N c_{1,j} \mathbf{x}_{\phi, \omega_j}, \quad \sum_{j=1}^N c_{1,j} = 1, \end{aligned} \quad a_j = c_{1,j} - c_{-1,j}, \quad \sum_{j=1}^N a_j = 0.$$

Scaled score function:
$$r(\alpha_{-1}, \alpha_1) d(\mathbf{x} | \alpha_{-1}, \alpha_1) = \frac{1}{2} (r^2(\alpha_{-1}, \mathbf{x}) - r^2(\alpha_1, \mathbf{x}))$$

For any real-world object $\omega \in \Omega$, any training set $\{\omega_j, j = 1, \dots, N\} \subset \Omega$, and any $\phi \in \Omega$:

$$r^2(\alpha_{-1}, \mathbf{x}_{\phi, \omega}) - r^2(\alpha_1, \mathbf{x}_{\phi, \omega}) = \sum_{l=1}^N (-\rho^2(\omega_l, \omega)) a_l + b,$$

$$r^2(\alpha_{-1}, \alpha_1) = \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l.$$

The general training problem

Training set: $\{(\omega_j, y_j), j = 1, \dots, N\}$, $\rho(\omega_j, \omega_l), j, l = 1, \dots, N$

Unified items of the training criterion	Score function	$d(\omega a_1, \dots, a_N, b) = \sum_{j=1}^N (-\rho^2(\omega_j, \omega)) a_j + b$
	Regularization function: squared dipole length	$\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0, \sum_{l=1}^N a_l = 0$
The only problem-specific item	Hinge function	$q(y, d(\omega a_1, \dots, a_N, b)) = q(y, \omega, a_1, \dots, a_N, b)$

The unified training criterion

$$\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l + C \underbrace{\sum_{j=1}^N q(y_j, \omega_j, a_1, \dots, a_N, b)} \rightarrow \min(a_1, \dots, a_n, b), \sum_{l=1}^N a_l = 0.$$

Particular versions of the problem-specific hinge function

Regression estimation	$\sum_{j=1}^N \left\{ y - \left(\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right) \right\}^2$
Logistic regression	$\sum_{j=1}^N \ln \left\{ 1 + \exp \left[-y \left(\sum_{l=1}^N (-\rho^2(\omega_l, \omega)) a_l + b \right) \right] \right\}$
Maximum margin	$\sum_{j=1}^N \max \left\{ 0, 1 - y \left(\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right) \right\}$

All the particular hinge-function summands are convex.

The general training problem

Training set: $\{(\omega_j, y_j), j = 1, \dots, N\}$, $\rho(\omega_j, \omega_l), j, l = 1, \dots, N$

Unified items of the training criterion	Score function	$d(\omega a_1, \dots, a_N, b) = \sum_{j=1}^N (-\rho^2(\omega_j, \omega)) a_j + b$
	Regularization function: squared dipole length	$\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0, \sum_{l=1}^N a_l = 0$
The only problem-specific item	Hinge function	$q(y, d(\omega a_1, \dots, a_N, b)) = q(y, \omega, a_1, \dots, a_N, b)$

The unified training criterion

$$\underbrace{\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l}_{\text{Regularization}} + C \sum_{j=1}^N q(y_j, \omega_j, a_1, \dots, a_N, b) \rightarrow \min(a_1, \dots, a_n, b), \underbrace{\sum_{l=1}^N a_l = 0}_{\text{Constraint}}$$

Particular versions of the problem-specific hinge function

Regression estimation	$\sum_{j=1}^N \left\{ y - \left(\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right) \right\}^2$
Logistic regression	$\sum_{j=1}^N \ln \left\{ 1 + \exp \left[-y \left(\sum_{l=1}^N (-\rho^2(\omega_l, \omega)) a_l + b \right) \right] \right\}$
Maximum margin	$\sum_{j=1}^N \max \left\{ 0, 1 - y \left(\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right) \right\}$

All the particular hinge-function summands are convex.

But the unified regularization term may be nonconvex in the case of an arbitrary metric.

Nonconvex maximum margin classifier

The training criterion

$$\left\{ \begin{array}{l} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l + C \sum_{j=1}^N \xi_j \rightarrow \min(a_1, \dots, a_N, b, \delta_1, \dots, \delta_N), \\ y_j \left[\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right] \geq 1 - \xi_j, \quad \xi_j \geq 0, j = 1, \dots, N. \end{array} \right. \quad \begin{array}{l} \text{Restriction} \\ \sum_{j=1}^N a_j = 0 \\ \text{is met automatically.} \end{array}$$

In the general case, this is a nonconvex quadratic programming problem.

Function $\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l$ under equality restriction $\sum_{j=1}^N a_j = 0$ is convex only for a proto-Euclidean metric.

The necessary condition of the minimum: The dual problem

$$\left\{ \begin{array}{l} \sum_{j=1}^N \lambda_j - (1/4) \sum_{j=1}^N \sum_{l=1}^N y_j y_l (-\rho^2(\omega_j, \omega_l)) \lambda_j \lambda_l \rightarrow \max(\lambda_1, \dots, \lambda_N), \\ \sum_{j=1}^N y_j \lambda_j = 0, 0 \leq \lambda_j \leq C, j = 1, \dots, N. \end{array} \right.$$

$$a_j = (1/2) y_j \lambda_j, \quad \sum_{j=1}^N a_j = 0, \quad b = \frac{(1/2) \sum_{j:0 < \lambda_j < \bar{C}} \lambda_j \sum_{l:\lambda_l > 0} \rho^2(\omega_j, \omega_l) y_l \lambda_l - C \sum_{j:\lambda_j = \bar{C}} y_j}{\sum_{j:0 < \lambda_j < C} \lambda_j},$$

For an arbitrary metric, the solution of the non-concave dual problem may result in the negative value of the squared maximum margin:

$$1/\varepsilon^2 = (1/2) \sum_{j=1}^N \sum_{l=1}^N y_j y_l (-\rho^2(\omega_j, \omega_l)) \lambda_j \lambda_l \geq 0.$$

Nonconvex maximum margin classifier

The training criterion

$$\left\{ \begin{array}{l} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l + C \sum_{j=1}^N \xi_j \rightarrow \min(a_1, \dots, a_N, b, \delta_1, \dots, \delta_N), \\ y_j \left[\sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_l + b \right] \geq 1 - \xi_j, \quad \xi_j \geq 0, j = 1, \dots, N. \end{array} \right. \quad \begin{array}{l} \text{Restriction} \\ \sum_{j=1}^N a_j = 0 \\ \text{is met automatically.} \end{array}$$

In the general case, this is a nonconvex quadratic programming problem.

Function $\sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l$ under equality restriction $\sum_{j=1}^N a_j = 0$ is convex only for a proto-Euclidean metric.

The necessary condition of the minimum: The dual problem

$$\left\{ \begin{array}{l} \sum_{j=1}^N \lambda_j - (1/4) \sum_{j=1}^N \sum_{l=1}^N y_j y_l (-\rho^2(\omega_j, \omega_l)) \lambda_j \lambda_l \rightarrow \max(\lambda_1, \dots, \lambda_N), \\ \sum_{j=1}^N y_j \lambda_j = 0, 0 \leq \lambda_j \leq C, j = 1, \dots, N. \end{array} \right. \quad \begin{array}{l} \sum_{j=1}^N \sum_{l=1}^N y_j y_l (-\rho^2(\omega_j, \omega_l)) \lambda_j \lambda_l \geq \delta > 0 \\ \text{quadratically constrained} \\ \text{quadratic optimization} \end{array}$$

$$a_j = (1/2) y_j \lambda_j, \quad \sum_{j=1}^N a_j = 0, \quad b = \frac{(1/2) \sum_{j:0 < \lambda_j < \bar{C}} \lambda_j \sum_{l:\lambda_l > 0} \rho^2(\omega_j, \omega_l) y_l \lambda_l - C \sum_{j:\lambda_j = \bar{C}} y_j}{\sum_{j:0 < \lambda_j < C} \lambda_j},$$

In this case, the squared maximum margin remains always positive for an arbitrary metric:

$$1/\varepsilon^2 = (1/2) \sum_{j=1}^N \sum_{l=1}^N y_j y_l (-\rho^2(\omega_j, \omega_l)) \lambda_j \lambda_l > 0.$$

Decision rule (score function) in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

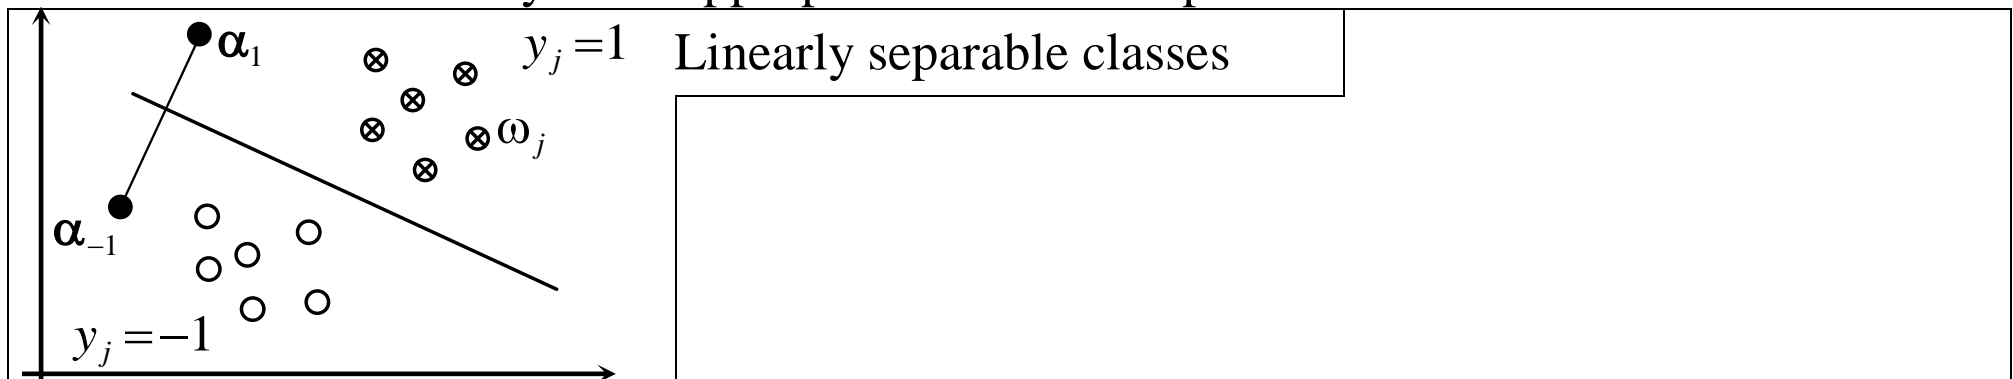
$$d(\omega \mid \underbrace{a_1, \dots, a_N, b}_{\text{Parameters to be estimated}}) = \frac{1}{2} \left[\sum_{j=1}^N a_j \underbrace{(-\rho^2(\omega_j, \omega))}_{\substack{\text{objects of the training set} \\ \swarrow \text{new object}}} + b \right] \geq 0,$$

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



Decision rule in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

$$d(\omega \mid \underbrace{a_1, \dots, a_N}_{\text{Parameters to be estimated}}, b) = \frac{1}{2} \left[\sum_{j=1}^N a_j (-\rho^2(\omega_j, \omega)) + b \right] \geq 0,$$

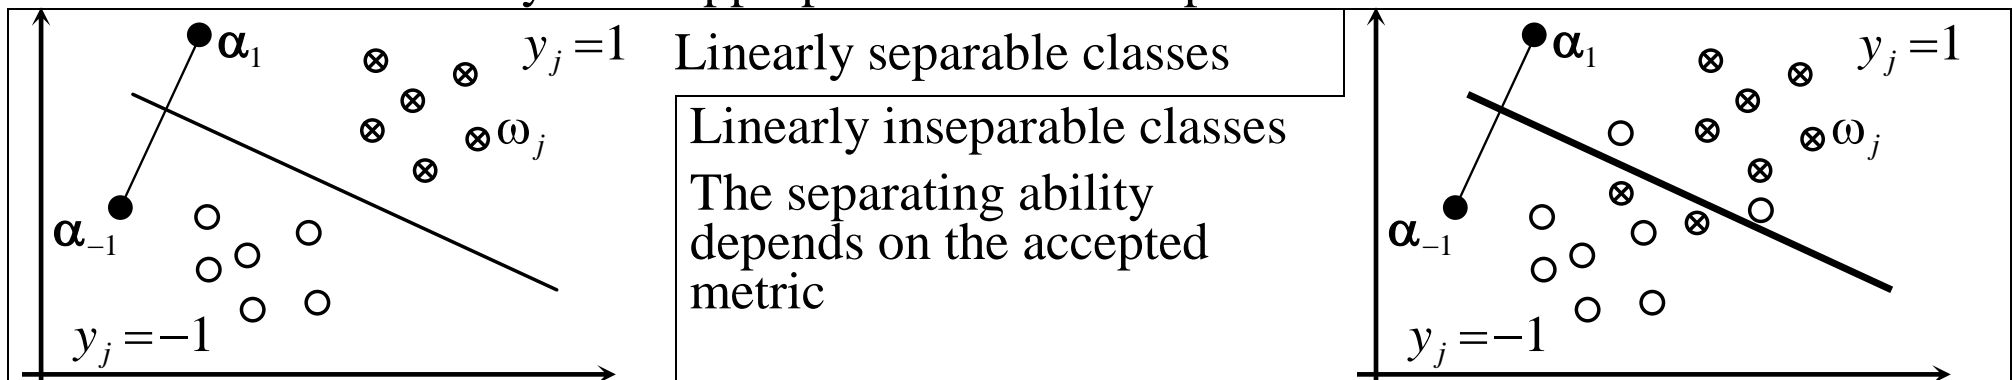
\swarrow *new object*
 \nwarrow *objects of the training set*

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



Decision rule in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

$$d(\omega \mid \underbrace{a_1, \dots, a_N}_{\text{Parameters to be estimated}}, b) = \frac{1}{2} \left[\sum_{j=1}^N a_j (-\rho^2(\omega_j, \omega)) + b \right] \geq 0,$$

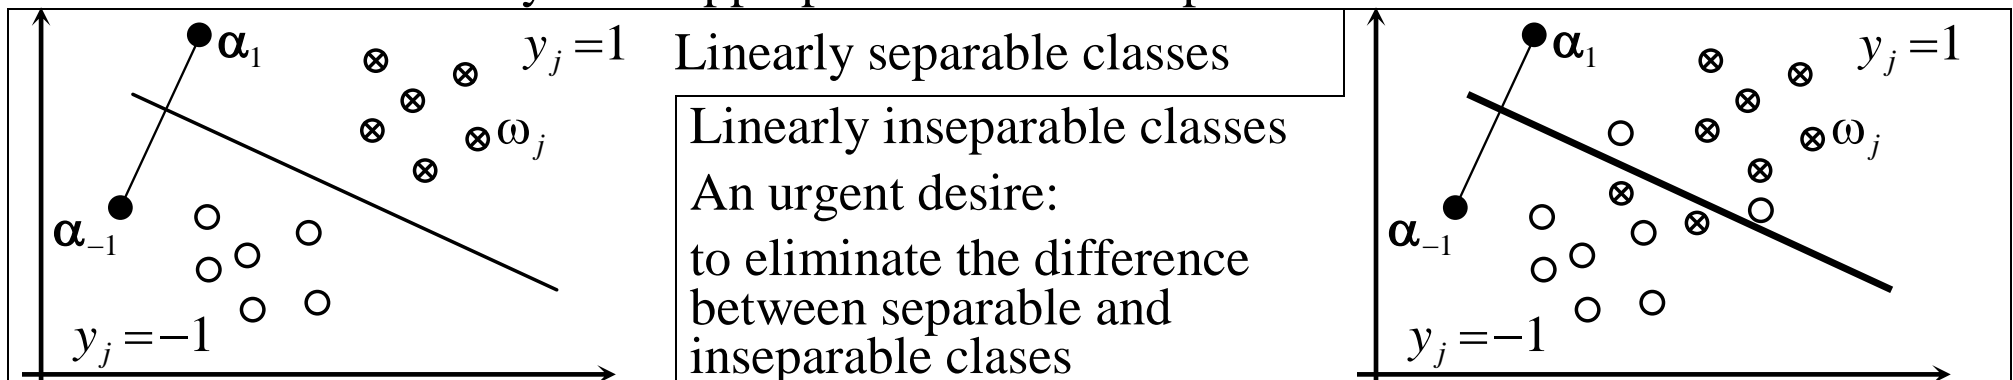
\swarrow *new object*
 \nwarrow *objects of the training set*

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



Decision rule in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

$$d(\omega \mid \underbrace{a_1, \dots, a_N, b}_{\text{Parameters to be estimated}}) = \frac{1}{2} \left[\sum_{j=1}^N a_j (-\rho^2(\omega_j, \omega)) + b \right] \geq 0,$$

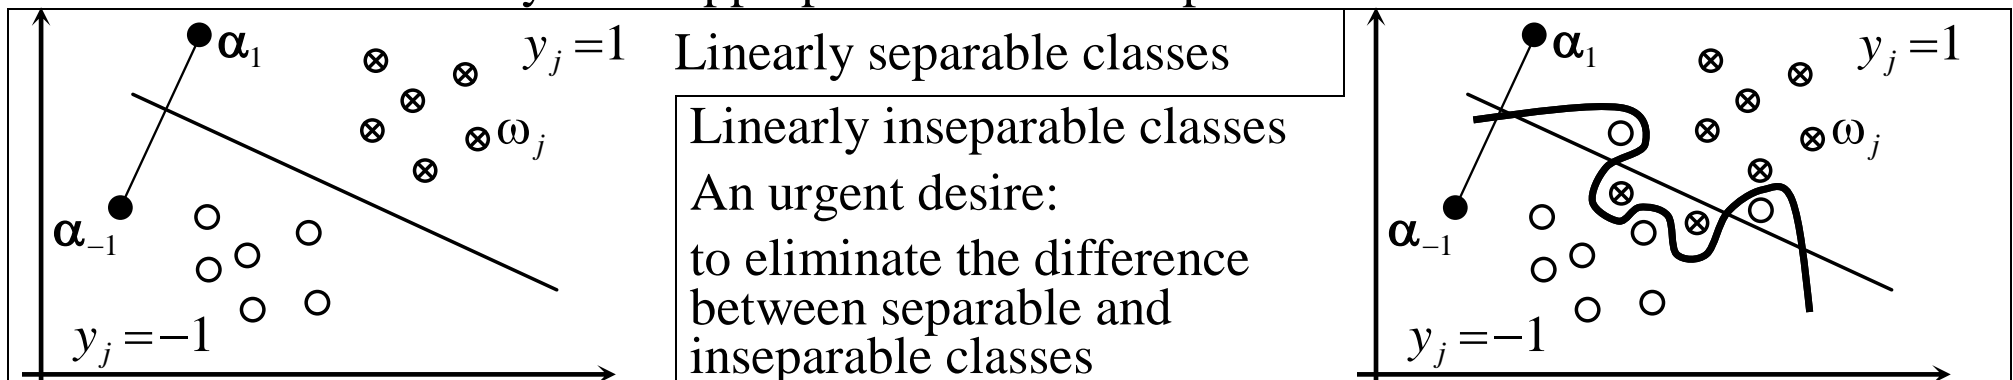
\swarrow *new object*
 \nwarrow *objects of the training set*

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



Decision rule in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

$$d(\omega \mid \underbrace{a_1, \dots, a_N}_{\text{Parameters to be estimated}}, b) = \frac{1}{2} \left[\sum_{j=1}^N a_j (-\rho^2(\omega_j, \omega)) + b \right] \geq 0,$$

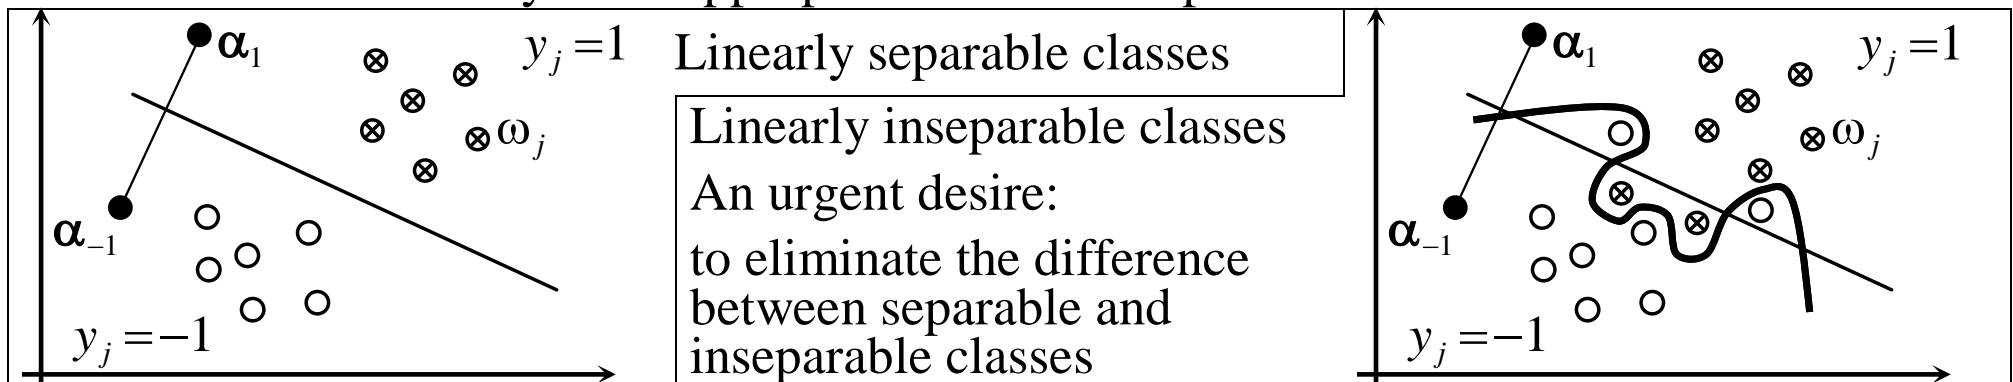
\swarrow *new object*
 \nwarrow *objects of the training set*

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



Way-out: An extension of the well-known notion of potential functions from finite-dimensional linear spaces onto metric spaces.

Decision rule in an arbitrary metric space

The equivalent form of the discriminant function, which is completely determined by $N+1$ real numbers (a_1, \dots, a_N, b) :

$$d(\omega \mid \underbrace{a_1, \dots, a_N}_{\text{Parameters to be estimated}}, b) = \frac{1}{2} \left[\sum_{j=1}^N a_j (-\rho^2(\omega_j, \omega)) + b \right] \geq 0,$$

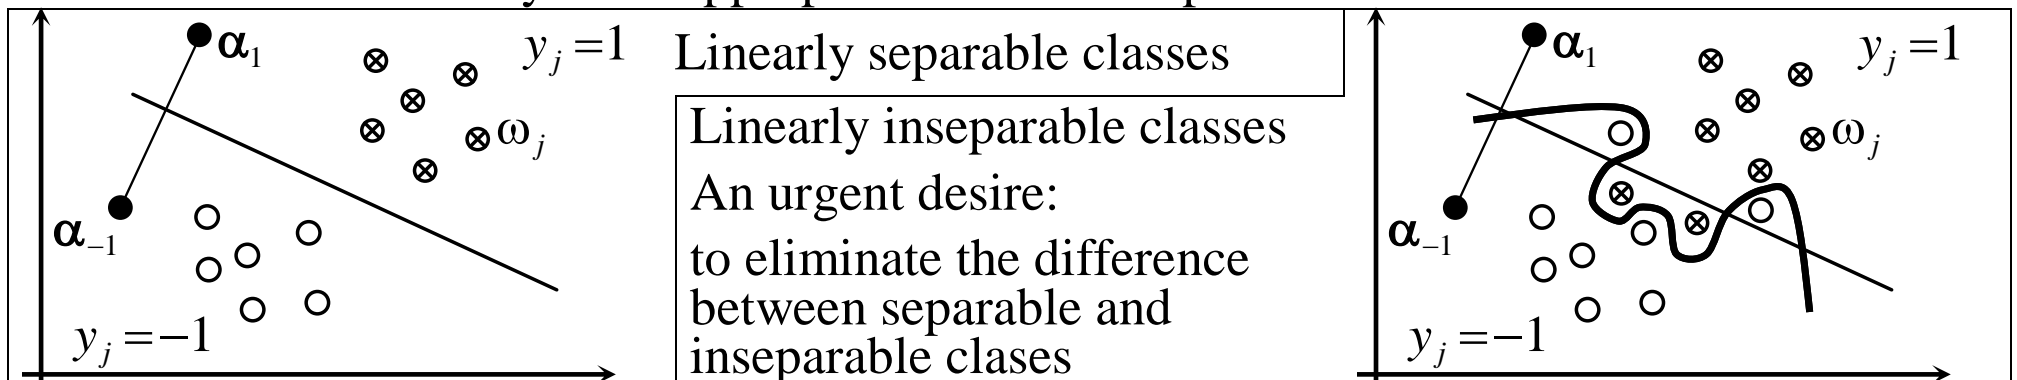
\swarrow *new object*
 \nwarrow *objects of the training set*

$$\sum_{j=1}^N a_j = 0, \quad \frac{1}{2} \sum_{j=1}^N \sum_{l=1}^N (-\rho^2(\omega_j, \omega_l)) a_j a_l > 0.$$

This is a linear decision rule in the pseudo-Euclidean space spanned over the given metric space of real-world objects

Eliminating the difference between linear and nonlinear decision rules in a pseudo-Euclidean space

Despite the fact that neither a metric space nor its pseudo-Euclidean embedding do not lend themselves to easy geometrical interpretation, we can conventionally demonstrate the essence of this boundary in an appropriate Euclidean space:



M. Aizerman, E. Braverman, L. Rozonoer. Theoretical foundations of the potential function method in pattern recognition learning. Automation and Remote Control, 1964, Vol. 25.

Potential function (kernel) in a finite-dimensional linear space

The initial inner product: $K(\mathbf{x}', \mathbf{x}'') = \mathbf{x}'^T \mathbf{x}'', \quad \mathbf{x}', \mathbf{x}'' \in \mathbb{R}^n$

The initial Euclidean metric: $\rho(\mathbf{x}', \mathbf{x}'') = [K(\mathbf{x}', \mathbf{x}') + K(\mathbf{x}'', \mathbf{x}'') - 2K(\mathbf{x}', \mathbf{x}'')]^{1/2} = [(\mathbf{x}' - \mathbf{x}'')^T (\mathbf{x}' - \mathbf{x}'')]^{1/2}$

Hilbert-space embedding –
a new inner product: $\tilde{K}(\mathbf{x}', \mathbf{x}'') = \exp(-\beta \rho^2(\mathbf{x}', \mathbf{x}''))$ – radial potential function

The new metric
in the Hilbert space: $\tilde{\rho}(\mathbf{x}', \mathbf{x}'') = [\tilde{K}(\mathbf{x}', \mathbf{x}') + \tilde{K}(\mathbf{x}'', \mathbf{x}'') - 2\tilde{K}(\mathbf{x}', \mathbf{x}'')]^{1/2} = \sqrt{2} [1 - \exp(-\beta \rho^2(\mathbf{x}', \mathbf{x}''))]^{1/2}$

However: $\tilde{\rho}(\mathbf{x}', \mathbf{x}'') \rightarrow 0$ when $\beta \rightarrow 0$

Normalization: $\tilde{\rho}(\mathbf{x}', \mathbf{x}'') = \frac{1}{\sqrt{\beta}} [1 - \exp(-\beta \rho^2(\mathbf{x}', \mathbf{x}''))]^{1/2}$

The kernel trick:

A linear function in the Hilbert space will produce a nonlinear one in the original space.

The idea: To apply such a transformation to an arbitrary metric.

Metric decision functions of growing complexity

The separating ability depends on the accepted metric $\rho(\omega', \omega'')$

$$d(\omega | a_1, \dots, a_N, b) = \sum_{j=1}^N (-\rho^2(\omega_j, \omega)) a_j + b \geq 0, \\ \sum_{j=1}^N a_j = 0$$


A parametric family of metric transformations, which guarantees improving the separating ability of the initial metric (family of saturable metrics):

$$\tilde{\rho}(\omega', \omega'' | \beta) = \frac{1}{\sqrt{\beta}} \left[1 - \exp(-\beta \rho^2(\omega', \omega'')) \right]^{1/2}$$

Theorem 7. If $\rho(\omega', \omega'')$ is a metric on Ω then $\tilde{\rho}(\omega', \omega'' | \beta)$ is a metric, too, with any value of the parameter $\beta \geq 0$, and $\tilde{\rho}(\omega', \omega'' | \beta) \rightarrow \rho(\omega', \omega'')$ when $\beta \rightarrow 0$.

Theorem 8. Any training set $\{(\omega_j, y_j = \pm 1), j=1, \dots, N\}$ with any metric $\rho(\omega', \omega'')$ is separable in the metric space $\tilde{\rho}(\omega', \omega'' | \beta)$ if the parameter $\beta > 0$ is large enough.

The choice of the metric transformation parameter β is analogous to the choice of the parameter of a radial potential function in a finite-dimensional feature space:




Metric decision functions of growing complexity. Illustration

Separable case:


Non-Separable case:

Using of metric transformations




Metric decision functions of growing complexity. Illustration

Separable case:



Non-Separable case:
Using of metric transformations



Thank you for your attention!