PRA Group

Università degli Studi di Cagliari

Exploiting Dissimilarity Representations for Person Re-Identification

Riccardo Satta, Giorgio Fumera, Fabio Roli

Pattern Recognition and Applications Group

Dept. of Electrical and Electronic Engineering

University of Cagliari, Italy

Outline

1. Introduction to the Person Re-Identification problem

2. Approaching Person Re-Identification through dissimilarity representations: motivations and advantages

3. A dissimilarity-based framework for Person Re-Identification

4. Preliminary experimental analysis

5. Further developments

Introduction

Person Re-Identification: the ability to determine if an individual has already been observed over a network of cameras.

In general, face recognition techniques are ineffective (*low resolution!*), therefore one must consider the global “appearance” of the individual

→ *short term problem*

Scenarios:

- **Online** *(i.e. tracking a person over different, non-overlapping cameras)*
- **Offline** *(i.e. retrieval of all the video sequences which contain a person of interest)*

Problem formulation

Person Re-Identification has been modelled as a classical *matching* problem

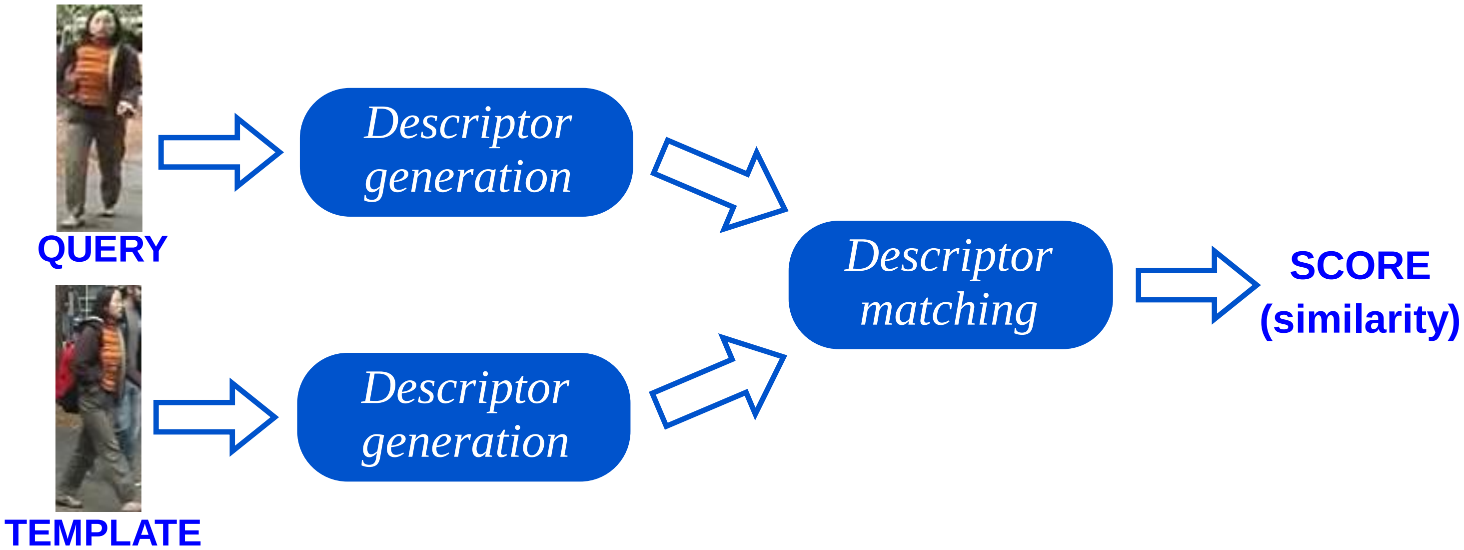

given a gallery set of templates $\mathrm{T} = \{\boldsymbol{T}_1,\ldots,\boldsymbol{T}_n\}$, *and a query* $\boldsymbol{Q}$, *find the most similar template* $\boldsymbol{T}^* \in \mathrm{T}$ *with respect to a similarity measure* $D(\cdot,\cdot)$:

$$\boldsymbol{T}^* = \arg\min_{\boldsymbol{T}_i} D(\boldsymbol{T}_i, \boldsymbol{Q})$$

An unifying view of the methods in literature

Despite the wide variety of methods in literature, most of the existing works

- adopt some part-based body subdivision (e.g. *torso* – *legs*), and/or
- represent the body or its parts via multiple components (e.g. *patches, SIFT points*).

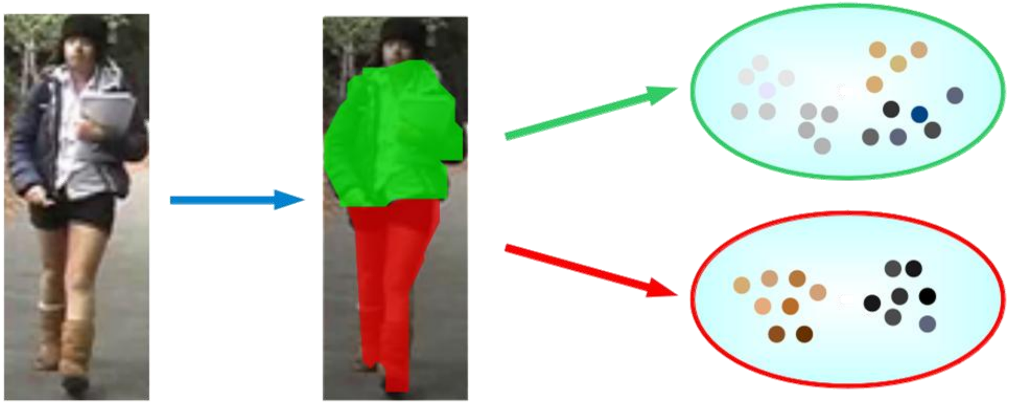

Part subdivision
in this example,
upper and lower body

Multiple components
represented here
as coloured dots

The authors exploited these similarities in a previous work to develop a general framework for person re-identification, **MCM**, able to frame most existing works:

[Satta et al., *A Multiple Component Matching Framework for Person Re-Identification*, ICIAP 2011]

Person Re-Identification and dissimilarity representations

Descriptors of different people often share similar components, related to common characteristics of their clothes.

We propose to represent each person as a vector of dissimilarity values in respect to pre-defined **visual prototypes**, drawn from their multiple component representations.

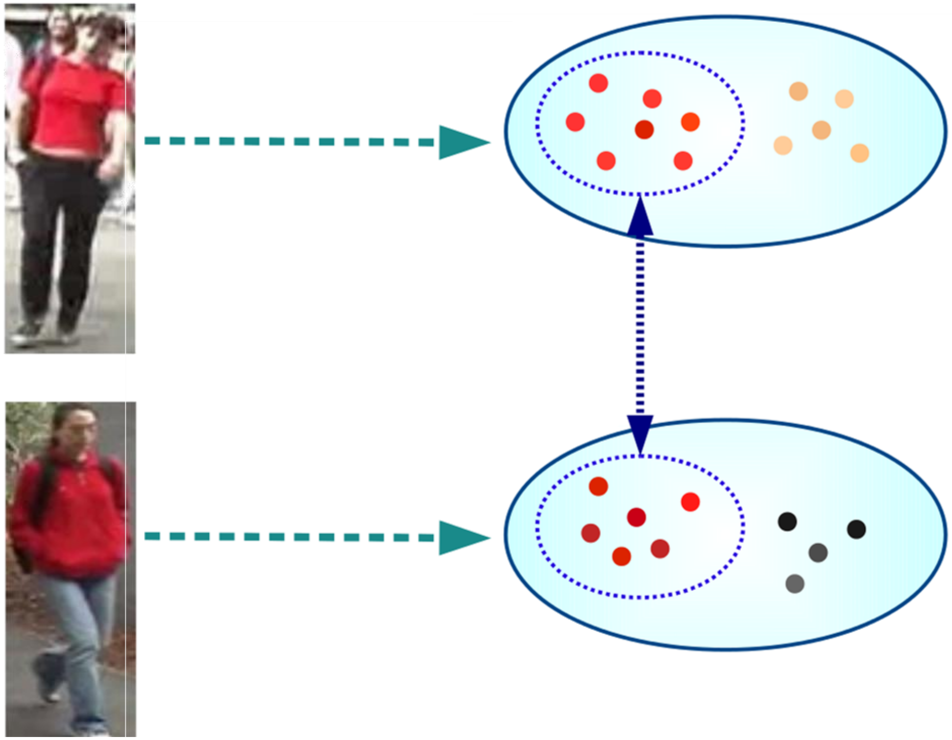

Each prototype corresponds to a common "visual characteristics".

Advantages

- **Reduced storage requirements**
  - people descriptors become as compact as a vector of real values
- **Extremely fast matching**
  - comparing two descriptors is equivalent to computing distances between vectors, an almost immediate operation with modern CPUs

A dissimilarity-based framework for Person Re-Identification

We present a new dissimilarity-based framework called **Multiple Component Dissimilarity (MCD)**, aimed at transposing a generic person re-identification method to a dissimilarity-based form.

MCD is based on the MCM framework introduced previously, and inherits the concepts of part subdivision and multiple components representation, in order to be applicable to any method that can be framed in MCM.

The MCM framework in short

- the individual is subdivided into different *body parts*
- each part is described by a *set of components*
- each component is represented by a feature vector
- components can be any kind of local features (patches, SIFT points, etc.), randomly chosen or selected through an interest operator

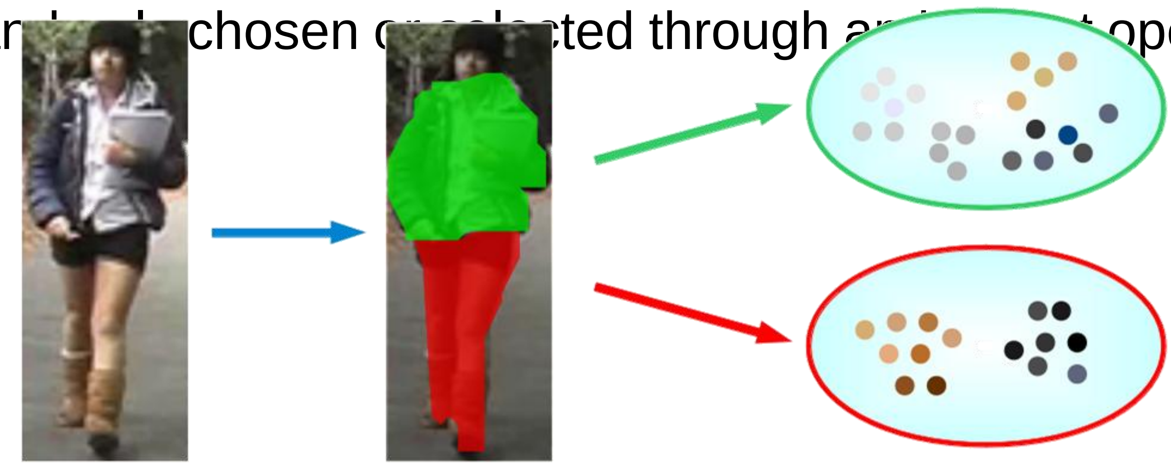

An example in which the body is subdivided in two parts.
Components are represented by coloured dots.

The Multiple Component Dissimilarity framework

Given

- a MCM-based person re-identification method
- a gallery of template individuals $\mathrm{T} = \{T_1, \ldots, T_n\}$
- a query individual Q

transposing the problem into a dissimilarity space requires four steps:

1. define a set of visual prototypes for each body part;
2. represent each element of T via dissimilarity vectors, one for each part;
3. represent Q via dissimilarity vectors, one for each part;
4. find the element of T which is most similar to Q in the dissimilarity space.

The Multiple Component Dissimilarity framework

Step 1 – a

First, all the components corresponding to each part are harvested from the template gallery and put together.

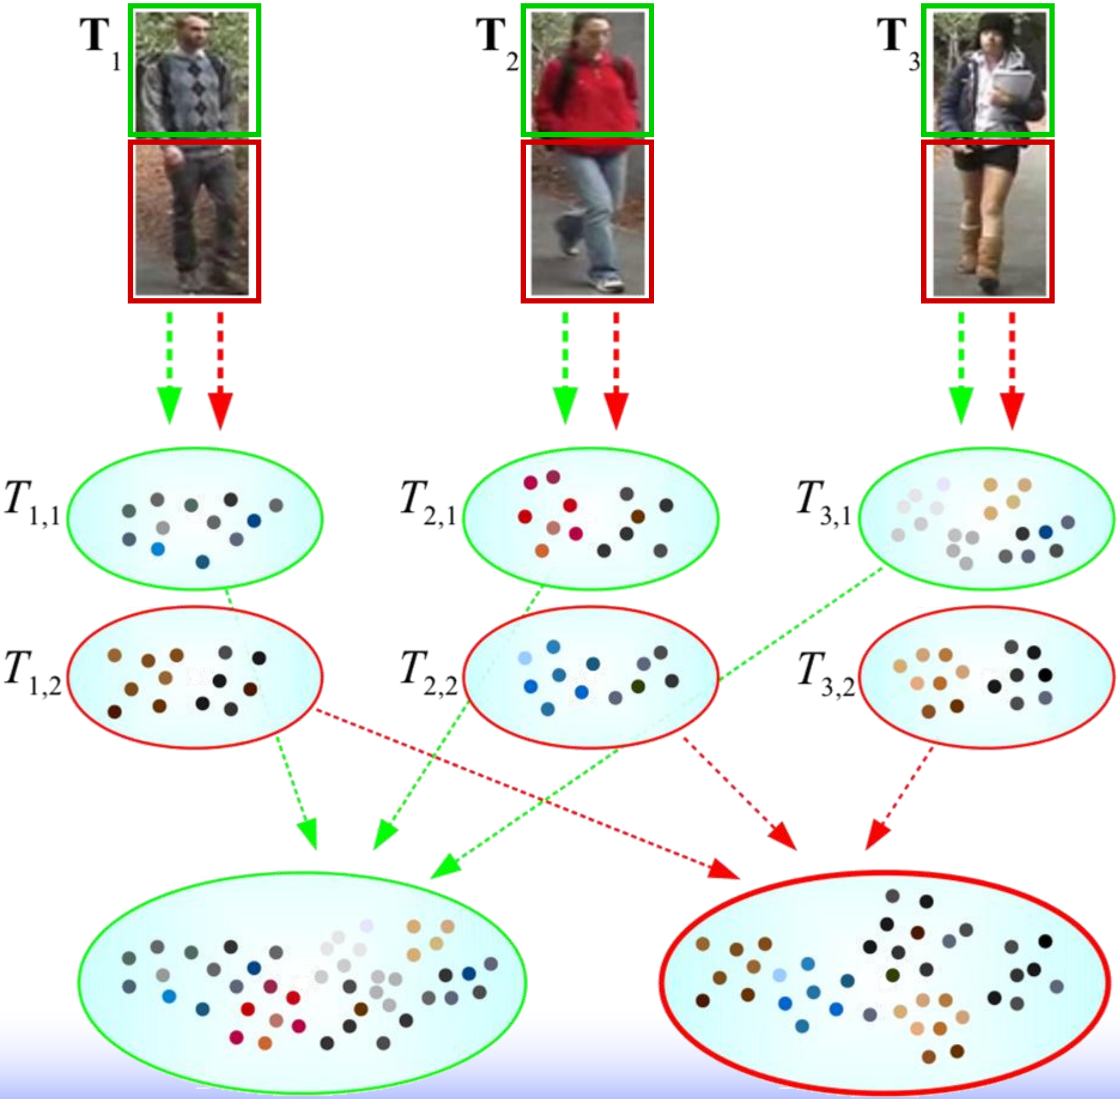

The Multiple Component Dissimilarity framework

Step 1 – b

Then, a clustering operator is applied and a number of **prototypes** is defined for each part.

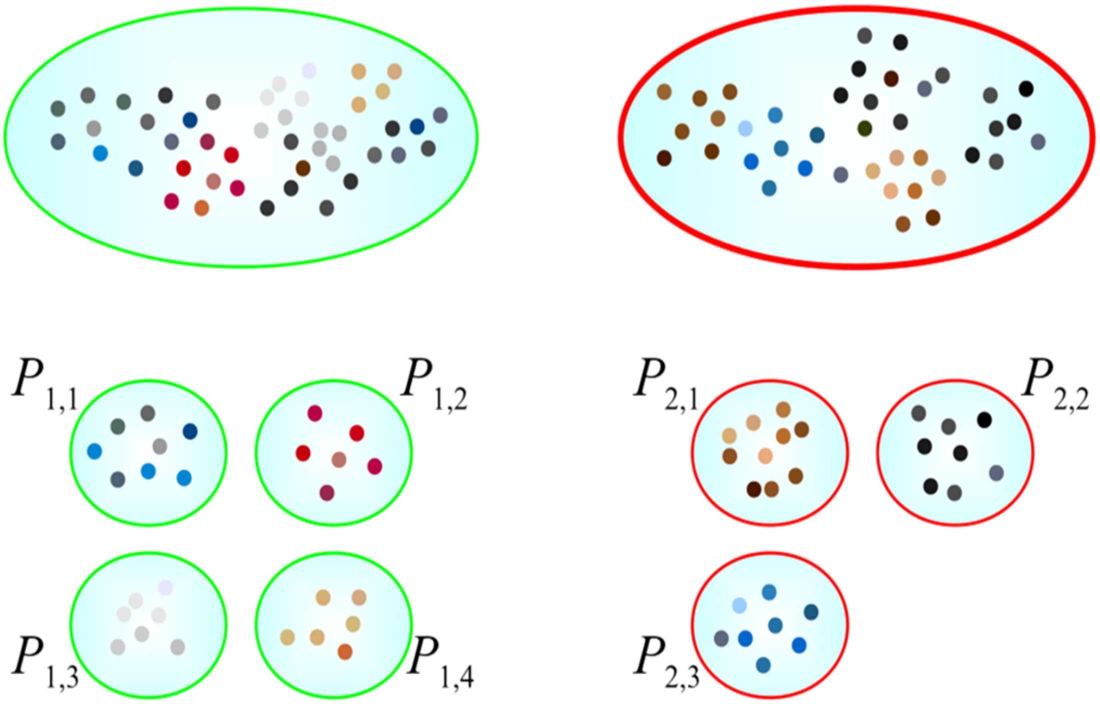

NOTE: each prototype is a set of components!

The Multiple Component Dissimilarity framework

Steps 2 and 3

Templates and the query are described with dissimilarity vectors in respect to the visual prototypes (one vector for each part).

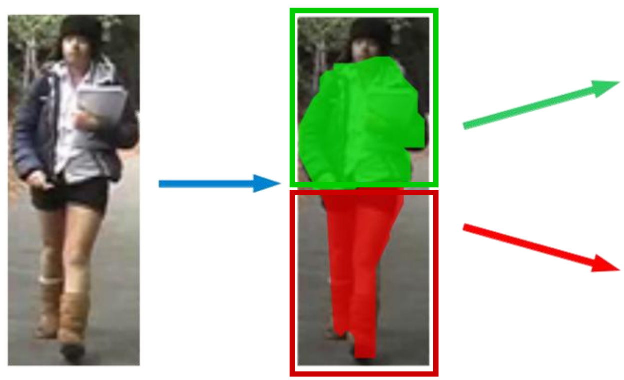

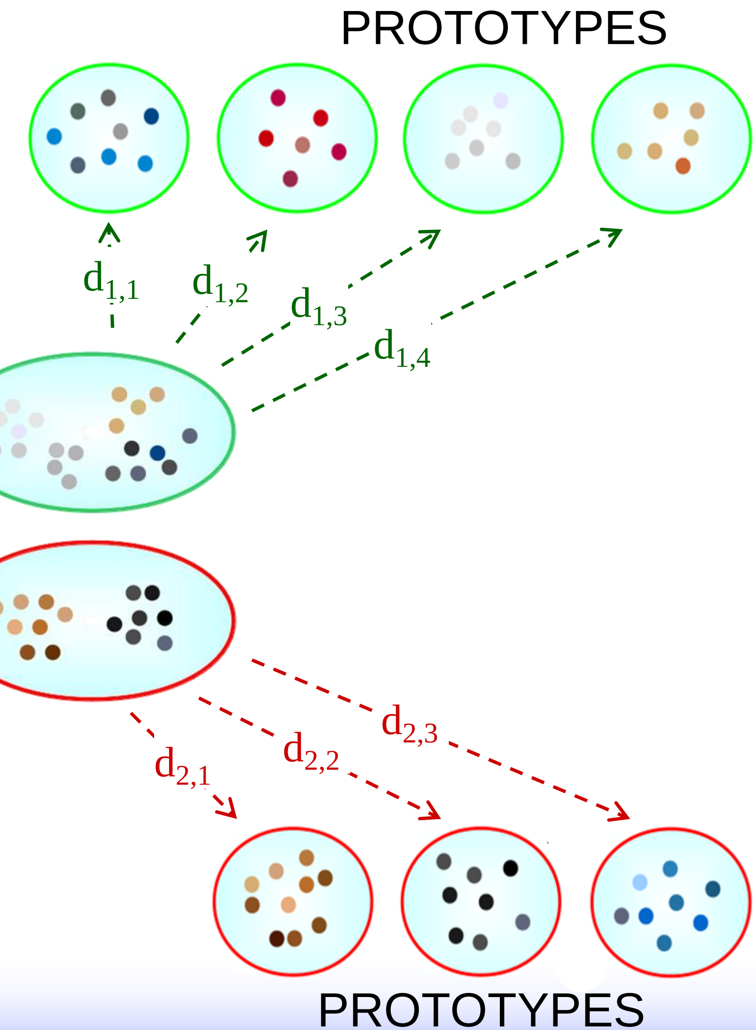

The Multiple Component Dissimilarity framework

Steps 2 and 3

Since each descriptor (either a template or the query) is made up of a set of components for each part, and each prototype is itself a set of components, dissimilarities are **distances between sets**.

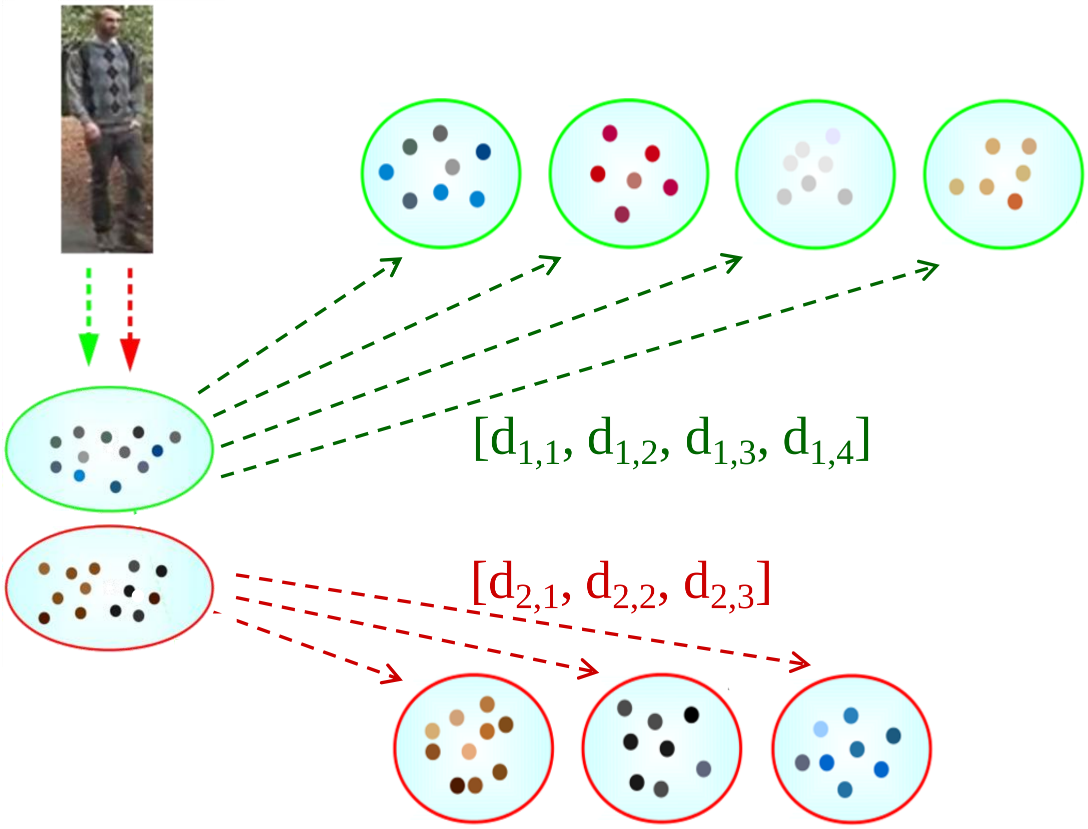

The Multiple Component Dissimilarity framework

Steps 2 and 3

TEMPLATES

$[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]$

$[d_{2,1}, d_{2,2}, d_{2,3}]$

$[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]$

$[d_{2,1}, d_{2,2}, d_{2,3}]$

$[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]$

$[d_{2,1}, d_{2,2}, d_{2,3}]$

QUERY

$[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]$

$[d_{2,1}, d_{2,2}, d_{2,3}]$

The Multiple Component Dissimilarity framework

PRA Group

Step 4

Similarities between the query and each template are evaluated by computing vector distances separately for each part, then fusing the part-level scores.

TEMPLATE

$$\underbrace{[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]}_{D_{part\ 1}} \quad \underbrace{[d_{2,1}, d_{2,2}, d_{2,3}]}_{D_{part\ 2}}$$

QUERY

$$\overbrace{[d_{1,1}, d_{1,2}, d_{1,3}, d_{1,4}]} \quad \overbrace{[d_{2,1}, d_{2,2}, d_{2,3}]}$$

$$\text{SIMILARITY} = f\left(D_{part\ 1}, D_{part\ 2} \right)$$

Experimental evaluation

Preliminary experiments were conducted by applying MCD to *MCMimpl*, a method proposed in [Satta et al., ICIAP 2011]

MCMimpl in short:

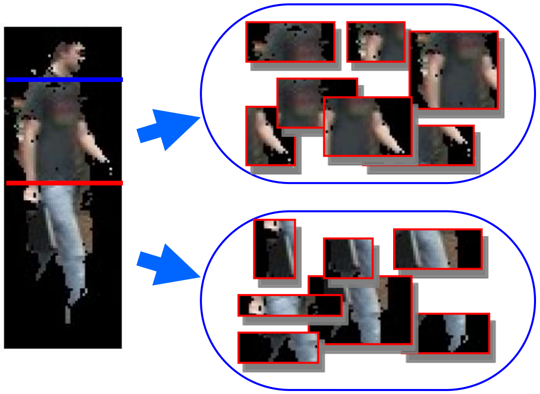

- part subdivision: torso – legs, discarding the head
- multiple component representation: for each part, randomly taken and partly overlapping patches, described by HSV histograms

Application of MCD to *MCMimpl*:

- visual prototypes generated through a two-stage clustering algorithm
- modified k-th Hausdorff distance to compute dissimilarities

Experimental evaluation

Experiments were carried out on 10 different subsets of 63 pedestrians taken from the VIPeR dataset, in which two different views per pedestrian are available.

- the first view of each pedestrian was used to build the template gallery;
- the second view of each pedestrian was used to build the probe (query) gallery.

Templates were ranked in respect to their similarity to each element of the probe gallery.
We evaluated

- computational times and storage requirements,
- recognition performance by means of the average Cumulative Matching Characteristics (CMC) curve.

Experimental evaluation

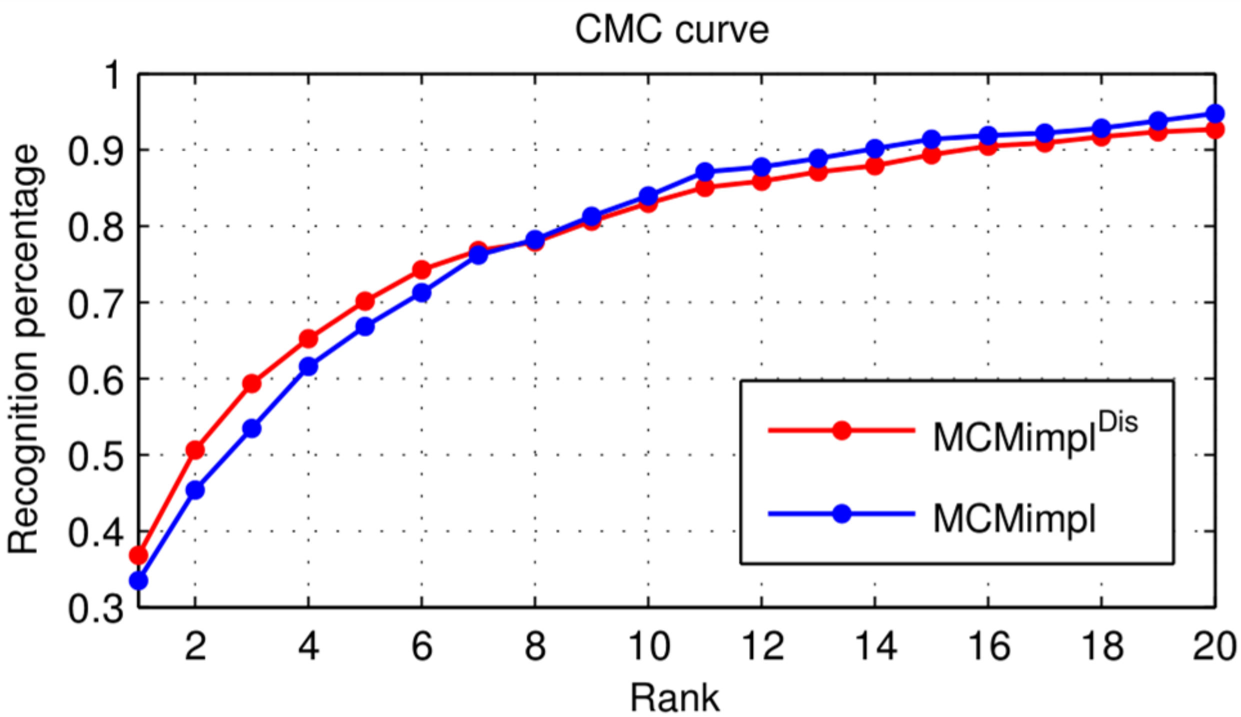

	$MCMimpl$	$MCMimpl^{Dis}$
Size of the descriptor	$96KB$	$640B$
Average matching time	$28.6ms$	$< 0.01ms$

The dissimilarity-based version of *MCMimpl* is denoted as $MCMimpl^{Dis}$.

PRA Group

Further developments

Appearance-based people grouping

Visual prototypes actually encode common visual characteristics of the clothes.

Dissimilarity representations can then be exploited to group people in respect to some visual characteristics of interest (e.g. *all people wearing a black shirt*).

Possible applications:

- reducing the number of candidates in a person re-identification task, by considering only those who share a certain amount of visual characteristics with the query
- appearance-based people retrieval, that is, retrieve people according to a given query related to some visual characteristics of interest

Thank you

Time for questions

The MCM framework in short

Descriptor matching

- set of components relative to each body part are compared separately through a similarity measure *d* between sets (e.g. the Hausdorff distance)

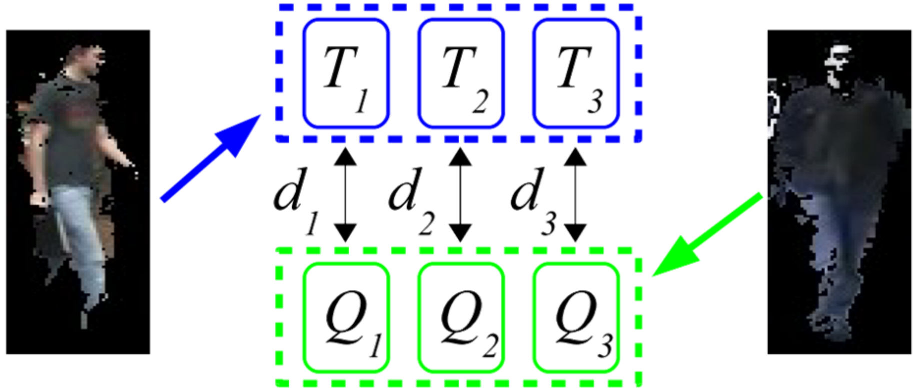

- the final matching score is a combination *D* (e.g. a linear combination) of the similarities at part level

$$D = f\left(d_1 , \ldots, d_i \right)$$