

A Transfer Learning Approach to Cross-modal Object Recognition: from Visual Observation to Robotic Haptic Exploration

Pietro Falco ^{*†} *Member IEEE*, Shuang Lu^{*}, Ciro Natale[‡], Salvatore Pirozzi[‡],
and Dongheui Lee^{*§} *Member IEEE*

^{*}Technical University of Munich, Germany

[†]ABB Corporate Research, Sweden

[‡]University of Campania Luigi Vanvitelli, Italy

[§]Institute of Robotics and Mechatronics, German Aerospace Center, Germany

Abstract—In this work, we introduce the problem of cross-modal visuo-tactile object recognition with robotic active exploration. With this term, we mean that the robot observes a set of objects with visual perception and, later on, it is able to recognize such objects only with tactile exploration, without having touched any object before. Using a machine learning terminology, in our application we have a visual training set and a tactile test set, or vice versa. To tackle this problem, we propose an approach constituted by four steps: finding a visuo-tactile common representation, defining a suitable set of features, transferring the features across the domains, and classifying the objects. We show the results of our approach using a set of 15 objects, collecting 40 visual examples and five tactile examples for each object. The proposed approach achieves an accuracy of 94.7%, which is comparable with the accuracy of the monomodal case, i.e., when using visual data both as training set and test set. Moreover, it performs well compared to the human ability, which we have roughly estimated carrying out an experiment with ten participants.

Index Terms—Cross-modal object recognition, Tactile Perception, Visual Perception, Robotic Manipulation

I. INTRODUCTION

Multi-modal perception technologies are the key enablers of robot autonomy operating in unstructured environments. On one hand, computer vision is fundamental for scene analysis and motion planning of the robot or for monitoring the robot workspace. On the other hand, vision cannot be the only solution to the perceptual need of a robot autonomously interacting with an unknown environment. In fact, effectiveness of a visual system is affected by lighting conditions, occlusions and limited field of view, especially during physical interaction with the world. Therefore, vision has to be supported by additional perceptual abilities, such as force and tactile feedback, that is extremely rich of useful information when the robot is in contact with the environment, e.g., measuring the contact force allows the robot to immediately identify constrained directions where the motion is not allowed. These considerations motivated a lot of research effort in the last decade towards advancements in the multi-modal perception technology, especially in the combined use of vision and force/tactile

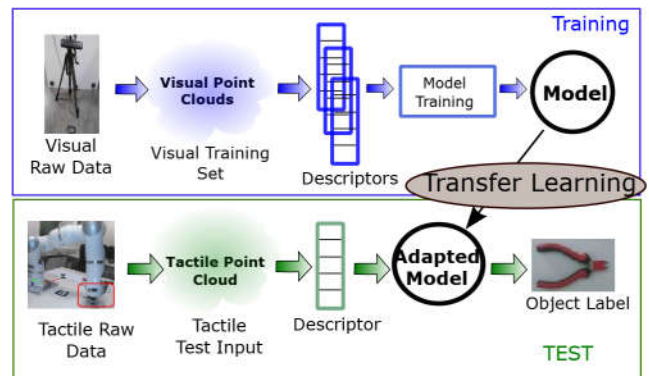


Fig. 1: Cross-modal recognition concept: training pipeline (top) and execution pipeline (bottom)

sensing. However, the field of cross-modal perception has been explored only superficially, while leveraging knowledge acquired in a perceptual domain during the execution of an action exploiting a different sensing modality could lead to a great level of autonomy. The present paper focuses on a typical application of cross-modal perception, namely the visuo-tactile object recognition, which means recognizing a previously seen object (never touched before) by simply touching it.

In neuroscience and psychology, cross-modal (or inter-modal) object recognition is defined as *the name for the ability to recognize an object, previously inspected with one modality like vision, via a second modality like touch* [1], without prior training in the second modality [2]. Cross-modal perceptual ability could enhance autonomy of a robot that is performing a task exploiting a given mono-modal perception system. The used sensor may unexpectedly become unavailable, e.g. vision due to lighting failure or occlusion, and we explored if the robot is able to perform its task using another sensor modality, e.g. a tactile sensor. To achieve this objective, the robot must be able to exploit its a-priori knowledge gained in the first sensing modality and use it at run time exploiting the second sensing modality. The concept of this approach is depicted in Fig. 1. We investigate if it is possible to recognize an object by using only tactile data and a classifier trained only with

visual data.

In order to make a cross-modal object recognition algorithm effective, two research challenges have to be faced, i.e.,

- 1) Selecting a common representation for both visual and tactile data. It should be fully interchangeable in a transparent way in both domains.
- 2) Selecting the right descriptors for the chosen visuo-tactile representation, possibly by re-using descriptors from the computer vision community.

This paper presents how we tackle these challenges by supporting our proposed choices with a large set of experiments. The preliminary version of the proposed approach is presented in [3] and extended by leveraging results from transfer learning [4] to further improve the performance. With respect to the work in [3], the novel additional research questions are

- do transfer learning approaches help to improve the performance of cross-modal visuo-tactile problems?
- are transfer learning approaches alternative to the solution proposed in [3] or can they be effectively combined?

The rest of the paper is organized as follows. Section II discusses the related work. Section III describes the proposed unified representation and the proposed descriptor. Both visual and tactile sensing setup are described in Sec. IV. Section V presents diverse experiments in order to show the performance of the proposed cross-modal classifier. Conclusions are reported in Sec. VI.

II. RELATED WORK

While the literature of the robotics community contains a number of contributions on both mono-modal and multi-modal object recognition problem, the cross-modal approach has been explored mainly in the neuroscience and psychology communities. As explained in the introduction, a mono-modal classifier is trained with one sensing modality and then it is queried with the same sensing modality to recognize the observed object. Concerning the visual sensing modality, the computer vision approach is a well-explored field and the related literature is vast. Owing to the widespread use of low-cost RGB-D cameras, the purely visual recognition approaches have been accompanied by recognition algorithms based on 3D point clouds. Specific descriptors have been proposed in the last decade exploiting visual 3D point clouds, e.g., Persistent Feature Histograms (PFH) [5], Fast PFH (FPPH) [6], Unique Signatures of Histograms (SHOT) [7], and Ensemble of Shape Functions (ESF) [8], and Spin Images (SI) [9]. A tutorial that describes and compares the most widespread descriptors is [10].

In contrast to visual approaches, a second sensing modality widely explored for object recognition is the tactile perception [11], which can be used not only for object classification but also for recognizing specific physical object features [12] such as texture or friction [13]. Zhang et al. [14] propose a descriptor to recognize objects based on data collected by a robotic hand equipped with tactile sensors. In [15], a bag-of-words approach is adopted to recognize objects from low-resolution tactile images acquired during the grasping with a

sensorized gripper. A stochastic approach based on the bag-of-features is proposed in [16] to estimate the probability distribution over object identity by object tactile exploration.

Multi-modal perception techniques are usually adopted to improve accuracy of the object classification algorithm by exploiting both visual and tactile data in the training phase. The deep learning method based on Convolutional Neural Networks (CNNs) proposed in [17] achieves very good performance in recognizing some material properties. The algorithm presented in [18] fuses visual and range data to recognize objects, while [19] combines visual features with tactile glances to refine object models, obtaining more accurate information about surfaces. In [20], visual-tactile recognition is carried out with 18 household objects. Visual and tactile data are used not only for recognition, as in [21], where 3D models of unknown objects are reconstructed based on multi-modal data acquired during object grasping. In [22], both monomodal and multimodal perception are used to detect what is inside a container using robotic grasping. The approach proposed in [23] exploits visuo-tactile multimodal perception to reduce the problem of pairing, also discussed in [24].

Cross-modal perception has been investigated in the neuroscience and psychology literature, e.g., in [25], [2], where some studies on animals have been carried out to demonstrate how cross-modality is actually exploited in nature. An entire chapter of [26] is dedicated to cross-modal object recognition. In [27] a study is reported concerning intermodal matching on infants, while [28] investigates visuo-tactile cross-modal perception in apes.

However, to the best of our knowledge, cross-modal visuo-tactile object recognition has been investigated in the robotics community for the first time in this work, together with our previous conference paper [3], which introduces the problem of cross-modal object recognition and proposes an empirical solution. The present paper extends the work in [3] and includes transfer learning techniques that improve significantly the accuracy of cross-modal recognition.

III. CROSS-MODAL OBJECT RECOGNITION

This section describes the elements of our cross-modal visuo-tactile framework, i.e., unified representation, features definition, visuo-tactile transfer learning, and learning algorithm.

A. Representation and Preprocessing

The first point we address is how to represent visual and tactile data to allow an effective cross-modal perception. RGB-D cameras allow us to represent an object O as a set of points $\mathcal{P} = \{\mathbf{p}_0, \mathbf{p}_1, \dots\}$, defined hereafter as point cloud of O . Each vector $\mathbf{p} = (p_x, p_y, p_z)$ denotes the 3D position coordinates of the point \mathbf{p} . With the symbols \mathcal{P}^v and \mathcal{P}^t we indicate that the point cloud \mathcal{P} and the point $\mathbf{p} \in \mathcal{P}$ is captured with visual perception. Note that the number of points of \mathcal{P} is always different at each acquisition. In order to derive a unified, compatible representation, *we represent tactile raw data as point clouds*. With raw data we mean the contact points between the object and the sensor. Even




Fig. 2: A representation for cross-modal perception considers common information between the modalities

though representations based on tactile point clouds were used for shape reconstruction [29] and creation of object bounding boxes [30], this choice may appear naive for object recognition applications. In fact, modern tactile sensors can provide richer information than a point cloud, such as contact forces, textures, pressure maps, and friction coefficients. However, as graphically shown in Fig. 2, in order to achieve cross-modal capabilities, a representation is required that contains information common to both visual and tactile perception. The tactile point cloud representation of the object O is denoted as $\mathcal{P}^t = \{p_0^t, p_1^t, \dots\}$, where the symbol t denotes a point cloud acquired with a tactile perception system. Tactile and visual point clouds present significant differences in point density, in partiality of data, and in the characteristics of noise that affects the measurements. To derive a more effective unified representation, we equalize \mathcal{P}^v and \mathcal{P}^t in order to reduce the difference in point density and in partiality. Data partiality consists in missing points in visual and tactile clouds. Even when the position and orientation of the objects are the same in both tactile and visual exploration, the tactile and visual point clouds have different missing points. Besides partiality, visual and tactile point clouds present also different point densities. In order to alleviate these differences, we preprocess both tactile and visual point clouds through two main steps: *equalizing partiality of the data* and *uniforming point density*.

1) *Equalizing Partiality*: The method we adopt to handle data partiality is the Moving Least Squares (MLS) surface reconstruction [31]. This step allows us to filter the measurement noise and to recreate the missing parts of the surface. The core of the MLS approach is composed by three basic steps. We assume to have a set of points \mathcal{P} . Given a query point $p \in \mathcal{P}$, the first step consists in finding a plane H that approximates locally the surface S in a region I of center p and radius r , called "search radius". The plane H is computed by using Principal Component Analysis (PCA). The points of the set I are projected onto H and upsampled with a step u_s of 0.3 mm. With this operation, we transform the set I into the

Algorithm 1 Equalization

- 1: function $\mathcal{P} = \text{equalize}(\text{PointCloud } \mathcal{P}^*)$
 - 2: $\tilde{\mathcal{P}} = \text{MLS}(\mathcal{P}^*, u_s = 0.3 \text{ mm}, r = 6 \text{ cm}, p_d = 2)$
 - 3: $\mathcal{P} = \text{voxelGridFilter}(l = 5 \text{ mm})$
 - 4: return(\mathcal{P})
-

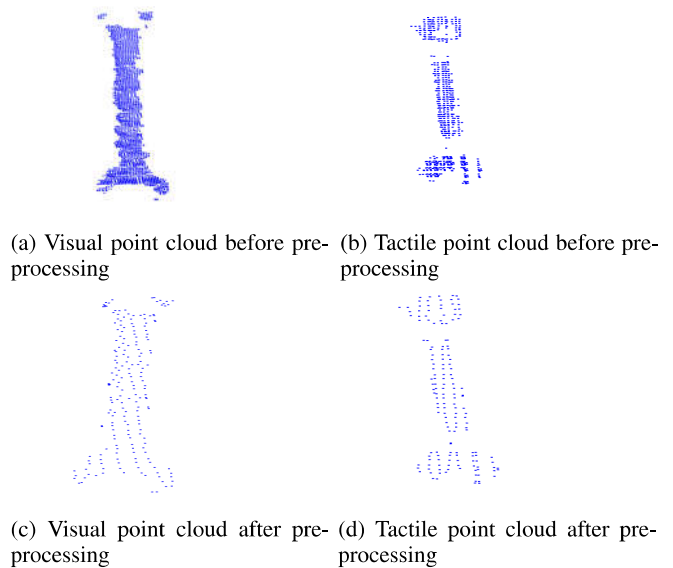


Fig. 3: Visual and tactile point clouds

set \tilde{I} . The second step consists in fitting with a polynomial of order p_d the height of the points projected on H . We choose $p_d = 2$ and $r = 6$ cm. Setting $r = 6$ cm confers rather strong filtering behavior and we can lose information in proximity of sharp edges. Typical values in monomodal visual perception are $r \in [1.5, 3]$ cm. However, in cross-modal perception, a strong filter can equalize cross-modal noise and in our case study allows achieving better performance. A more detailed and formal description of the procedure can be found in [31]. In this work, the parameters are chosen with a grid search approach, maximizing the recognition accuracy.

2) *Uniforming Density*: The second step of the equalization procedure consists in applying a voxel grid filter [32] to downsample and ensure a more uniform point density. We apply the voxel filtering approach implemented in Point Cloud Libraries (PCL) [33]. In this approach, the space is divided in 3D cubes (called voxels). All the points contained in each 3D box are substituted with their centroids. Following this procedure, the number of points will be equal to the number of filled voxels. Selecting appropriately the dimension of the voxels, the similarity of point density between tactile and visual data can be improved. In this work, we have empirically chosen cubic voxel with edge length $l = 5$ mm.

The procedure is summarized in Algorithm 1. An example of visual and tactile point clouds before and after preprocessing is shown in Fig. 3. The equalization step plays a key role in order to improve the performance (see Sec. V).

B. A Suitable Descriptor

After defining a unified representation based on point clouds, it is important to choose a suitable feature descriptor for cross-modal recognition. The choice of the feature descriptors strongly depends on the chosen representation of raw data. Since our unified representation is based on point clouds, we orient our research towards 3D point cloud descriptors.

From the results reported in the computer vision literature ([7], [8], [34]), SHOT and ESF seem promising candidates for our problem. However, our experimental data shows that even after the preprocessing step, the differences in noise, resolution and partiality of the data found in the training and test set cannot be equalized perfectly. Therefore, there is a need of a new, robust descriptor suitable for cross-modality. Following a strategy commonly adopted in communication engineering, we propose to increase the redundancy of the information associated to the descriptors. A way to increase the redundancy is finding a smart combination of different descriptors. In our case we expect benefit by combining SHOT and ESF, since they encode information with two different approaches, i.e. normal-based and normal-free respectively. In particular, SHOT encodes point clouds with histograms of normal vectors [7], while ESF encodes information based on shape functions [8]. Also, they both present the best performance according to our comparisons (Table II) and do well according to works in the literature, e.g., [34], [8]. The ESF descriptor is an ensemble of 10 concatenated histograms of shape functions consisting of angle, point distance, and area functions. Each histogram has 64 bin, for a total of 640 elements [8]. The SHOT computes a local reference frame Σ_c using the eigenvalue decomposition around an input point c , in our case c is the centroid of the point cloud \mathcal{P} . Given the frame Σ_c , a sphere S_c of center c and radius r_c is defined. S_c is then split into 32 divisions and for each division a 11-bin histogram is computed. Each histogram contains angles that describe the directions of the normal vectors to each point $p \in \mathcal{P}$ in the frame Σ_c . The descriptor concatenates the histograms into the final signature, obtaining a vector of 352 elements. We compute a single SHOT feature for each object and use it as a global feature.

Let d_{SHOT} be the SHOT descriptor associated to the point cloud \mathcal{P} and d_{ESF} the ESF descriptor associated to the same point cloud. Both descriptors are column vectors. The first possible way to improve the performance is to simply concatenate d_{SHOT} and d_{ESF} so that:

$$d_c = [d_{SHOT}^T \quad d_{ESF}^T]^T, \quad (1)$$

where d_c is the concatenated descriptor. Even if the concatenated descriptor d_c contains more information than d_{ESF} , the improvement in accuracy was not significant. This can happen because the dimension of d_c is much higher than both SHOT and ESF. As a consequence, the classification problem is affected by the curse of dimensionality. Moreover, the high increase in dimension can be a limitation also in terms of training time and classification time, especially when scaling to very large databases. The SHOT descriptor associated to a point cloud is a vector of 352 elements, while an ESF descriptor is a vector of 640 elements. Therefore, we decided to exploit a data compression method. To compress the descriptor d_c , we organize the vectors d_{ESF} and d_{SHOT} in the matrix:

$$\hat{D} = [d_{ESF} \quad \tilde{d}_{SHOT}], \quad (2)$$

where $\tilde{d}_{SHOT} = [d_{SHOT}^T \quad \bar{\mathbf{0}}]^T$ and $\bar{\mathbf{0}}$ is a 0-vector of dimension $1 \times (640 - 352)$. We want to compress the information carried from the matrix $\hat{D} \in \mathbb{R}^{640 \times 2}$ into a vector $d_r \in \mathbb{R}^{640}$.

We leverage the data compression capability of Singular Value Decomposition (SVD) [35]. First, we center the matrix \hat{D} , so that all columns are zero-mean. Let D be such a mean-centered matrix. The SVD of the matrix D is

$$D = U \Sigma V^T, \quad (3)$$

where $U = [u_1, u_2, \dots, u_{640}] \in \mathbb{R}^{640 \times 640}$, $V = [v_1, v_2] \in \mathbb{R}^{2 \times 2}$, and $\Sigma \in \mathbb{R}^{640 \times 2}$ is the matrix that contains the singular values σ_1 and σ_2 , U and V contain the left and right singular vectors, respectively. We choose the compressed SVD descriptor d_r as

$$d_r = \sigma_1 u_1. \quad (4)$$

Since we make zero-mean the columns of the matrix D , the singular value decomposition is equivalent to a decomposition based on principal component analysis. This descriptor has dimension 640 and is a linear combination of the columns of the matrix D . The best rank-1 approximation of the matrix D is given by the matrix $D_1 = \sigma_1 u_1 v_1^T$. As a consequence, $d_r v_1^T$ is the 1-rank matrix that minimizes the norm $\|D - D_1\|$.

Using d_r as a descriptor we obtain significantly better performance than using the d_c [3]. The descriptor d_r , in fact, carries more information than both ESF or SHOT, but is less affected by the curse of dimensionality than the concatenated descriptor d_c . We call the descriptor derived in Eq. (4) Cross-modal point cLoUd dEscriptor (CLUE). The CLUE descriptor consists in the basic ESF enriched with the information carried by SHOT. CLUE performs better than both ESF and SHOT in the cross-modal case, while it performs very similarly in the monomodal and multimodal cases. Hence, we guess that the robustness of combining a normal-based descriptor (SHOT) and a normal-free descriptor (ESF) is beneficial for cross-modal applications in particular.

C. Transfer Learning

Exploiting the equalization step and the CLUE descriptor, we obtain a significant improvement in the accuracy [3]. In order to further improve the accuracy, we propose to adopt techniques from transfer learning. Transfer learning approaches have been adopted effectively in computer vision and textual document recognition [4].

In Transfer Learning, and in particular domain adaptation, we define a source domain \mathcal{D}_S and a target domain \mathcal{D}_T . A generic domain \mathcal{D} is constituted by the couples $(x, p(\mathbf{X}))$, where $x \in X$ is the features vector, $p(\mathbf{X})$ is the probability distribution of the feature space. In this work, the vector x is constituted by the elements of the descriptor, e.g., CLUE shown in Sec. III-B. Beside a domain, we define also a learning task $\mathcal{T} = \{\mathbf{Y}, f(\cdot)\}$, where \mathbf{Y} is a set of labels and $f : x \in \mathcal{X} \rightarrow y \in \mathcal{Y}$ is the function which associates feature vectors to classes (or labels). The source dataset $D_S = \{x_{S_1}, \dots, x_{S_N}\}$ is constituted by the source feature vectors. Moreover, we denote with $\tilde{D}_S = \{(x_{S_1}, y_{S_1}), \dots, (x_{S_N}, y_{S_N})\}$ the source dataset with labels. The target dataset $D_T = \{(x_{T_1}), \dots, (x_{T_N})\}$ is constituted only by feature vectors but typically not by labels. Moreover, we can define a source task $\mathcal{Y}_S = \{Y_S, f_S(\cdot)\}$ and a target task \mathcal{Y}_T . Given a source domain \mathcal{D}_S , a source learning task \mathcal{T}_S , a target domain \mathcal{D}_T , and a target learning task \mathcal{T}_T ,

transfer learning aims to improve the learning of \mathcal{T}_T by using the knowledge of \mathcal{D}_S and \mathcal{T}_S . When $\mathcal{T}_T = \mathcal{T}_S$, as in our case, we have a domain adaptation problem.

We focus on using the visual data as source domain and the tactile data as target domain. The principal reason is that it is easier, in general, to collect several images and it is more demanding to collect many tactile examples. Collecting tactile data requires physical interaction with the external environments and, on the long run, devices such as sensors can be damaged or a significant amount of time or energy can be required for haptic exploration. In simpler words, the transfer learning problem can be formulated this way: given the knowledge of the labeled set \tilde{D}_S , can we estimate labels for the set D_T ?

In the literature, transfer learning approaches have been used in text classification and in image recognition to transfer knowledge from a data base to another, exploiting the same perception modality. In our case, we investigate if such techniques can be adopted to transfer knowledge across different perception modalities. To this aim, we apply to our problem two classes of approaches. The first class is based on dimensionality reduction, which includes Transfer Component Analysis (TCA) [36] and Subspace Alignment (SA) [37]. The second class is based on the geodesic flow associated to different subspaces. This class includes Geodesic Flow Kernel (GFK) [38]. We briefly describe here the GFK approach. For details on TCA and SA, please refer to [36] and [37], respectively.

Existing approaches such as TCA and SA focus on learning feature representations that are invariant across domains. The basic idea of GFK is to integrate an infinite number of subspaces that characterize changes in geometric and statistical properties from the source domain to the target domain. For our application in cross-modal object recognition, GFK shows the best performance, and combined with equalization of partiality and resolution, achieves similar performance of monomodal recognition and slightly outperforms visuo-tactile multimodal recognition. The method we used in this work is based on [38]. In order to apply GFK, we introduce the concept of *Grassmann Manifold*. Given a vector space \mathcal{D} of dimension D , a Grassmannian $G(\delta, D)$ is a manifold which includes all δ -dimensional linear subspaces of the vector space \mathcal{D} . In our case, the source domain is denoted \mathcal{D}_S and the target domain is denoted as \mathcal{D}_T . Both the spaces have dimension D . The first step of GFK transfer learning consists in reducing the dimensionality of both source and target domains with a linear operator. A typical option is principal component analysis:

$$X_S = \text{PCA}(D_S, \delta) \quad (5)$$

$$X_T = \text{PCA}(D_T, \delta), \quad (6)$$

where, similarly as in SA, δ is a parameter of the algorithm. The matrix $X_S \in \mathbb{R}^{D \times \delta}$ represents the basis for the linear subspace of \mathcal{D}_S obtained through PCA. Similarly, the matrix $X_T \in \mathbb{R}^{D \times \delta}$ represents the PCA basis for the linear subspace of \mathcal{D}_T . Since X_S and X_T can be seen as δ -dimensional subspaces of a D -dimensional space, they are points of a Grassmann Manifold $G(\delta, D)$. Let R_S and R_T be the orthog-

onal complements to X_S and X_T , respectively. The geodesic flow is represented with the parametric function

$$\Phi : t \in [0, 1] \rightarrow \Phi(t) \in G(\delta, D), \quad (7)$$

where

$$\Phi(0) = X_S \quad (8)$$

$$\Phi(1) = X_T. \quad (9)$$

The parametric function $\Phi(t)$, in brief, maps the values of the parameter t to all the subspaces that connect the source domain and the target domain. In particular, as shown in [38], the function Φ has the form:

$$\Phi(t) = P_S U_1 \Gamma(t) - R_S U_2 \Sigma(t), \quad (10)$$

where $\Gamma(t) = \text{diag}\{\sin(t\theta_1), \sin(t\theta_2), \dots, \sin(t\theta_\delta)\}$, $\Sigma(t) = \text{diag}\{\cos(t\theta_1), \cos(t\theta_2), \dots, \cos(t\theta_\delta)\}$, and $\{\theta_1, \theta_2, \dots, \theta_\delta\}$ are the principal angles between the source and target domains. It holds $\theta_1 < \theta_2 < \dots < \theta_\delta < \pi/2$. The matrices $U_1 \in \mathbb{R}^{\delta \times \delta}$ and $U_2 \in \mathbb{R}^{(D-\delta) \times \delta}$ are computed by the following singular value decompositions

$$X_S^T X_T = U_1 \Gamma V^T \quad (11)$$

$$R_S^T P_T = -U_2 \Sigma V^T. \quad (12)$$

After defining the function Φ , the GFK approach uses the infinite-dimensional feature vector z^∞ defined as

$$z^\infty = [\Phi(0), \dots, \Phi(t), \dots, \Phi(1)]^T \mathbf{x}, \quad (13)$$

where $t \in [0, 1]$, \mathbf{x} is the feature vector as derived in Sec. III-B, e.g. CLUE. Since the vector z^∞ has infinite dimension, it is not usable directly by a digital computer. However, exploiting the kernel trick [39], we can define a distance between two infinite-dimensional vectors \mathbf{z}_i^∞ and \mathbf{z}_j^∞ using the scalar product

$$\langle \mathbf{z}_i^\infty, \mathbf{z}_j^\infty \rangle = \int_0^1 (\Phi(t)^T \mathbf{x}_i)^T (\Phi(t)^T \mathbf{x}_j) dt. \quad (14)$$

It is possible to show that the integral in Eq. (14) is equal to the following quadratic form:

$$\int_0^1 (\Phi(t)^T \mathbf{x}_i)^T (\Phi(t)^T \mathbf{x}_j) dt = \mathbf{x}_i^T \mathbf{G} \mathbf{x}_j, \quad (15)$$

where the matrix $\mathbf{G} \in \mathbb{R}^{D \times D}$ is such that

$$\mathbf{G} = [X_S U_1 \quad R_S U_2] \begin{bmatrix} \Lambda_1 & \Lambda_2 \\ \Lambda_2 & \Lambda_3 \end{bmatrix} \begin{bmatrix} U_1^T X_S^T \\ U_2^T X_T^T \end{bmatrix} \quad (16)$$

with

$$\Lambda_i = \text{diag}\{\lambda_{i_1}, \lambda_{i_2}, \dots, \lambda_{i_\delta}\}, \quad i = 1, 2, 3 \quad (17)$$

and

$$\lambda_{1_j} = 1 + \frac{\sin(2\theta_j)}{2\theta_j} \quad (18)$$

$$\lambda_{2_j} = \frac{\cos(2\theta_j) - 1}{2\theta_j} \quad (19)$$

$$\lambda_{3_j} = 1 - \frac{\sin(2\theta_j)}{2\theta_j}, \quad j = 1, 2, \dots, \delta. \quad (20)$$

It is interesting to note that the GFK approach compute the similarity between two feature descriptors by leveraging the

Algorithm 2 CMR Training

```

1: function  $\mathcal{M}^v = \text{training}(\mathcal{P}_S, \text{Labels } Y_S)$ 
2:  $\bar{\mathcal{P}}_S = \text{equalize}(\mathcal{P}_S)$ 
3:  $D_S = \text{computeCLUE}(\bar{\mathcal{P}}_S)$ 
4: Model  $\mathcal{M}^v = \text{train}(D_S, Y_S)$ 
5: return( $\mathcal{M}^v$ )

```

Algorithm 3 TL-CMR Training

```

1: function  $\mathcal{M}^v = \text{training}(\mathcal{P}_S, \mathcal{P}_T, \text{Labels } Y_S)$ 
2:  $\bar{\mathcal{P}}_S = \text{equalize}(\mathcal{P}_S)$ 
3:  $\bar{\mathcal{P}}_T = \text{equalize}(\mathcal{P}_T)$ 
4:  $D_S = \text{computeCLUE}(\bar{\mathcal{P}}_S)$ 
5:  $D_T = \text{computeCLUE}(\bar{\mathcal{P}}_T)$ 
6:  $\mathbf{G} = \text{GFK}(D_S, D_T, \delta)$ 
7: Model  $\mathcal{M}^v = \text{train}(D_S, Y_S, \mathbf{G})$ 
8: return( $\mathcal{M}^v$ )

```

matrix \mathbf{G} . Such matrix depends on both the source domain and the target domain. When $\mathbf{G} = \mathbf{I}$, the similarity becomes the classical scalar product in the Euclidean space. The sole parameter of the algorithm is the subspace dimension δ . Detailed guidelines on how to tune this parameter are reported in [38]. It is important to remark that the GFK transfer learning algorithm exploits both data from the source domain and data for the target domain. However, none of such data is labelled. Data from the target domain, in fact, are used only in the adaptation phase in unsupervised fashion and not in the training phase.

D. Classification Algorithm

We compare k -Nearest Neighbor (k -NN), with different values of k and radial basis function kernel Support Vector Machines (RBF-SVM), and linear SVM. Both k -NN and SVM are simple and widely-used algorithms for classification problems. We apply such learning algorithms to several state-of-the-art visual descriptors and with the one proposed in this work.

In more detail, to deal with the cross-modal recognition problem, we perform two steps. The first step consists in building a model \mathcal{M}^v , which embeds a-priori knowledge derived from visual perception. The second step is to exploit a-priori knowledge embedded in \mathcal{M}^v with data from a different sensing modality. In this work, we build the model by following and comparing two procedures, i.e., Cross Modal Recognition (CMR) pipeline and Transfer Learning-based CMR (TL-CMR). For CMR, we need only labeled source-domain data to build the model, while for TL-CMR we need labeled source-domain data and a set of unlabeled target-domain data.

1) *Building the model using CMR*: We use visual point clouds of 15 objects and for each object we collect 40 examples. Each example i consists in a point cloud \mathcal{P}_i . For each point cloud \mathcal{P}_i , we apply the equalization procedure described in Sec. III-A and compute the CLUE descriptor \mathbf{d}_i , which is the representation of the point cloud in the feature

space. Let D_S be the set of the CLUE descriptors associated to all the examples, i.e., $D_S = (\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_n)$, with $n = 600$ in our case. Also, let Y_S be the set of all the labels associated to the elements of D_S . The labeled source dataset is denoted as $\tilde{D}_S = (D_S, Y_S)$. We derive the model \mathcal{M}^v using the set \tilde{D}_S . When using k -NN, the model \mathcal{M}^v simply consists of the elements of D_S . In case of SVM and other methods that require an explicit training step, we use D_S as training set and the model \mathcal{M}^v is the trained classifier. The procedure to build a CMR model is described in Algorithm 2. The sensing system and the procedure to collect visual data is described in Sec. IV-A. We call this process Cross Modal Recognition (CMR). In this case, we use only labeled source training examples to build our model.

2) *Building the model using TL-CMR*: Besides the CMR pipeline, we exploit the transfer learning techniques described in Sec. III-C to improve the performance. In order to build a model with the SA and GFK approaches we need not only the labeled source dataset \tilde{D}_S , but also unlabeled examples of the target domain D_T . Hence, when applying TL-CMR we assume that the robot has haptically explored the objects of the target dataset but it is not provided with labels. The procedure of building the model with TL-CMR is described in Algorithm 3. The TL-CMR procedure consists in building the source labeled dataset \tilde{D}_S by using equalization procedure and descriptor computation on the acquired point clouds. Then, the target unlabeled dataset D_T is built applying equalization and descriptor computation to the target unlabeled point clouds. The next step is to apply the transfer learning algorithm to the unlabeled source data D_S and the unlabeled target domain data D_T . The output of the transfer learning algorithm is a domain adaptation factor. In training the classifier, the adaptation factor is taken into account to reduce the differences between target and source domain. In the case of GKF, the adaptation factor is the matrix \mathbf{G} , which replaces the scalar product with a quadratic form to compute the similarity between two feature vectors [38]. In the case of SA, we have a matrix \mathbf{M} that aims at aligning source and target domains in a reduced subspace, as described in Sec. III-C. In order to apply the overall TL-CMR pipeline we need labeled source-domain data and unlabeled target-domain data. Target-domain data are in fact exploited only for the adaptation step and not for model training. If collecting a few unlabeled examples of the target domain is not possible for a specific case study, the simpler CMR pipeline

Algorithm 4 Cross-modal Recognition

```

1: function  $l = \text{recognize}(\text{PointCloud } \mathcal{P}_o^t, \text{Model } \mathcal{M}^v)$ 
2:  $\bar{\mathcal{P}}_o^t = \text{equalize}(\mathcal{P}_o^t)$ 
3:  $\mathbf{d}_{SHOT} = \text{computeSHOT}(\bar{\mathcal{P}}_o^t)$ 
4:  $\mathbf{d}_{ESF} = \text{computeESF}(\bar{\mathcal{P}}_o^t)$ 
5:  $\hat{\mathbf{D}} = [\mathbf{d}_{SHOT} \quad \mathbf{d}_{ESF}]$ 
6:  $\mathbf{D} = \text{center}(\hat{\mathbf{D}})$ 
7:  $[\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}] = \text{svd}(\mathbf{D})$ 
8:  $\mathbf{d}_{CLUE} = \mathbf{U}(:, 1)\mathbf{\Sigma}(\mathbf{1}, \mathbf{1})$ 
9:  $l = \text{classify}(\mathbf{d}_{CLUE}, \mathcal{M}^v)$ 
10: return( $l$ )

```

should be used.

3) *Exploiting the model for cross-modal recognition:* We exploit the knowledge accumulated with visual perception in order to interpret tactile data at execution time. To test the performance of cross-modal recognition, we classify the outcome of five tactile explorations per object. The tactile sensing system and the exploration procedure are described in Sec. IV-B. After the acquisition of the tactile point cloud, we derive the descriptor d with the procedure described in Sec. III-B. To recognize the object through visual a-priori knowledge, we provide d as an input to the classifier which embeds the model \mathcal{M}^v . The output of such a classifier is the estimated class of the object.

The entire process of visuo-tactile recognition that adopts the CLUE descriptor is summarized in Algorithm 4. The inputs of the algorithm are the model \mathcal{M}^v , derived by visual data a-priori known and the point cloud observed by tactile sensors \mathcal{P}_O^t at execution time. The output is the label l of the explored object O .

IV. SENSING SYSTEM

An experimental setup has been prepared in order to test the effectiveness of the cross-modal object recognition approach proposed in Sec. III. The system is constituted by a robotic arm (KUKA LWR-IV) equipped with a tactile sensor and an external visual perception system.

A. Visual Perception

Figure 4 shows the visual perception system, constituted by an Asus Xtion Pro Live RGB-D camera, which has been used to collect the visual point clouds. All relative positions between camera and objects are as shown Fig. 4. For an object O the collected point cloud is separated from the rest of the scene by using an ECE (Euclidean Cluster Extraction) algorithm, available from the PCL libraries. This algorithm removes the planes from the scene and clusters the remaining points by using a kd-tree approach. Each object has been placed in two different poses during the acquisitions to add




Fig. 4: Visual Sensing System




Fig. 5: Experimental setup for tactile perception

more variability to the data. The descriptors used in our approach, though, are invariant to position and orientation of the objects.

B. Tactile Perception

1) *Tactile Sensing Setup:* In the experimental setup, the tactile skin developed within the SAPHARI Project [40] has been fixed on the end effector of the KUKA robotic arm, as shown in Figs. 5a, 5b.

The distributed tactile sensor, originally presented in [41], consists of a PCB (Printed Circuit Board) constituted by couples (emitter/detector) of optoelectronic devices, used to detect the local deformations of the deformable layer covering the optoelectronic layer. These deformations are related to the external contact forces applied to the sensor.

The used sensor is constituted by an interconnection of a number of identical sensing modules, each capable to estimate the three components of the contact force applied to it. In particular, each sensing module is constituted by four optoelectronic couples organized in a 2×2 matrix. The whole sensor consists of a 6×6 grid of sensing modules with a total size equal to $5 \times 5 \text{ cm}^2$. Each sensing module, shown in Figs. 5c, 5d, has a unique spatial representation in the robotic base frame and provides the estimated three force components.

The i th contact point p_i is selected, when the contact force intensity $\|F_i\|$ estimated by the i th module is larger than a threshold β . For the experimental results, reported in this paper, the threshold value has been empirically chosen as $\beta = 0.8 \text{ N}$. Then, for each object, the tactile readings are represented as three-dimensional point clouds, as described in Sec. III.

2) *Exploration Strategy:* For the tactile object representation and recognition, the tactile exploration is a fundamental phase. An appropriate strategy guarantees a good quality of the tactile point clouds. For the experiments discussed in this paper, the objects are explored by pressing on them along the z axis of the robot base frame, highlighted in Fig. 5a. The base frame has been selected as the unified world reference frame, and all tactile point clouds are represented in this frame. During the exploration phase, if a module of the tactile sensor

senses a contact with the object O , the corresponding point is included in its point cloud \mathcal{P}^t . The tactile readings are constituted by six dimensional vector, encoding the pose and the force data. All exploration experiments have been carried out by using a Cartesian impedance controller for the robot, in order to allow a compliant interaction with the objects. In particular, the robot is compliant along the z axis, so that the exploration phase does not damage any object. For each object a grid as reported in Fig. 4b is defined. The end effector of the robot is moved to each vertex of this grid. Then, after the reaching of a vertex, the robot presses with the tactile sensor on the object. The i -th point between the skin and the object \mathbf{p}_i is represented by its coordinates (p_x, p_y, p_z) , with respect to the robot base frame. In this paper, the tactile frame Σ_t coincides with the robot base frame, but the feature selection is frame-independent. During the proposed experiments, the objects are fixed on the table. The points of the table are removed with the planar filter algorithm implemented in PCL. The simple exploration strategy, adopted in this paper, is particularly suitable for planar objects. The whole procedure is detailed in Algorithm 5.

Algorithm 5 Exploration Strategy

```

1:  $Tra_{j1} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$   $\triangleright$  grid vertices in Fig. 4b
2: for  $\mathbf{v}_j \in Tra_{j1}$  do
3:   moveTo( $\mathbf{v}_j$ )  $\triangleright$  it brings robot from vertex to vertex
4:   press on vertex  $\mathbf{v}_j$ 
5:   for each sensor module  $i$  do
6:     read( $\mathbf{F}_i$ )
7:     if  $\|\mathbf{F}_i\| \geq 0.8N$  then
8:        $\mathbf{p}_i \leftarrow (p_x, p_y, p_z)$ 
9:        $\mathcal{P} = \mathcal{P} \cup \{\mathbf{p}_i\}$ 
10:    end if
11:  end for
12: end for

```

V. EXPERIMENTS

A. Description of the Dataset

We selected 15 objects, depicted in Fig. 6, which are typical of domestic and industrial environments. For each object, the tactile exploration procedure, described in Algorithm 5 has been repeated 5 times. Then, 40 samples from each object have been collected with the visual system in Fig. 4. After the data acquisition procedure, we have 40 visual and 5 tactile point clouds for each object. The visual point clouds are used to build the a-priori knowledge. In this case study, a-priori knowledge is constituted by the classifier trained with visual data, denoted in our work with the symbol \mathcal{M}^v . The tactile exploration data are then classified exploiting the a priori knowledge \mathcal{M}^v . It is important to emphasize that, with the proposed approach, the robot can classify an object using the sense of touch when the object has never been touched before, but only seen by vision. In this work, we used only rigid object. The application of this strategy to deformable objects is planned as future work.



Fig. 6: Objects used for experiments, depicted near to a 1-euro coin

Feature descriptor	1-NN	3-NN	5-NN	SVM
PFH	15.20%	8.44%	8.44%	8.44%
FPFH	10.06%	10.06%	10.06%	13.34%
SI	21.33%	22.67%	21.33%	32.00%
SHOT	32.63%	32.63%	31.25%	35.79%
ESF	45.26%	45.26%	42.11%	32.63%
\mathbf{d}_c	45.26%	45.26%	42.11%	26.32%
CLUE	55.79%	57.89%	53.68%	33.68%

TABLE I: Cross-modal Recognition result without preprocessing

B. Classification results

In order to assess the performance of the framework, we evaluate, in terms of accuracy, the proposed combination of (1) unified representation, (2) unified descriptor, (3) transfer learning approach, and (4) learning algorithm. The accuracy is defined as the number of correct classification operations over the total number of classification operations. We compare different state-of-the-art descriptors with the proposed CLUE, as shown in Table I and Table II. Also, we show how the application of transfer learning techniques improves the performance in terms of recognition accuracy.

1) *Results without Transfer Learning:* First, we evaluate the proposed framework without using transfer learning. In this

Feature Descriptor	1-NN	3-NN	5-NN	SVM
PFH	5.33%	12.00%	12.00%	10.67%
FPFH	9.33%	12.00%	14.67%	16.00%
SI	22.67%	28.00%	28.00%	40.00%
SHOT	37.33%	36.00%	34.67%	32.00%
ESF	60.00%	65.33%	65.33%	49.33%
\mathbf{d}_c	62.67%	68.00%	66.67%	30.67%
CLUE	72.00%	73.33%	77.33%	40.00%

TABLE II: Cross-modal Recognition Result with preprocessing

a	1.0	0	0	0	0	0	0	0.2	0	0	0	0	0.2	0	0	0
b	0	1.0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
c	0	0	1.0	0	0.4	0	0	0	0	0	0	0	0	0	0	0
d	0	0	0	1.0	0	0	0	0	0	0	0	0	0	0	0	0
e	0	0	0	0	1.0	0	0	0	0	0	0	0	0	0	0	0
f	0	0	0	0	0	1.0	0	0	0	0	0	0	0	0.2	0	0
g	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0	0	0
h	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0	0
i	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0
j	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0
k	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0
l	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0
m	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0
n	0	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0
o	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0
	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	

TABLE III: Confusion matrix for the CMR architecture with CLUE descriptor and 5-NN classification algorithm.

a	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
b	0	1.0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
c	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
d	0	0	0	1.0	0.2	0	0	0	0	0	0	0	0	0	0	0
e	0	0	0	0	1.0	0.6	0	0	0	0	0	0	0	0	0	0
f	0	0	0	0	0	1.0	0.2	0	0	0	0	0	0	0	0	0
g	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0	0	0
h	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0	0
i	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0	0
j	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0
k	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0
l	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0
m	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0.4
n	0	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0
o	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0.6
	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	

TABLE IV: Confusion matrix for the TL-CMR architecture with CLUE descriptor, GFK transfer learning, and RBF SVM classifier.

case study, we adopt the Cross Modal Recognition process (CMR) proposed in [3]. In order to show how strongly the preprocessing step described in Sec. III-A impacts on the performance of cross-modal object recognition, we indicate in Table I and Table II the classification accuracy without and with preprocessing, respectively. In the light of the results mentioned above, the preprocessing allows an improvement of up to 20% in the recognition accuracy and the importance of this step is crucial when dealing with visuo-tactile cross-modality. As classification algorithms, we compare k -NN and radial basis function kernel SVM, since they are simple and have recognized effectiveness in classification problems. The best performance is achieved by 5-NN combined with the CLUE descriptor.

The accuracy is evaluated by classifying five tactile explorations per object using the visual knowledge embedded in the model \mathcal{M}^v , obtained through the different classification methods reported in Table II. We can observe that the descriptors using estimated normal vectors, i.e., PFH, FPFH, SI and SHOT, perform worse than ESF. CLUE is particularly suitable for cross-modal recognition, as the improvement with respect to basic ESF is, in our case study, almost 15% and the dimension remains the same. CMR achieves 77% accuracy. In Table III, the confusion matrix is reported for a more detailed analysis. We can notice the recognition performance is between 80% and 100% for all the objects except for four objects: (j) the hairpin achieves 20% accuracy, (n) the tape 40% accuracy, (e) the socket wrench 60% accuracy, and (h) the ruler 60% accuracy.

2) *Results with Transfer Learning*: Our second analysis is performed by applying the transfer learning approaches described in Sec. III-C, hence, using the Transfer Learning-

based Cross Modal Recognition process (TL-CMR). We used four transfer learning techniques: the simple PCA, Transfer Component Analysis (TCA), Subspace Alignment (SA), and Geodesic Flow Kernel (GFK). In Table V, we report the classification accuracy without preprocessing, i.e., without the equalization procedure performed in Sec. III-A. The transfer learning approaches based on dimensionality reduction such as PCA, TCA and SA show poor accuracy. The best approach is GFK combined with CLUE, which improves the accuracy of CMR, achieving 82% accuracy. Table VI reports the results of TL-CMR including the equalization steps. Also in this case, PCA, TCA, and SA do not significantly improve the performance with respect to the CMR architecture. In fact, in most transfer learning approaches the accuracy is even lower than the CMR accuracy. However, GFK combined with the CLUE descriptor achieves 94% accuracy using the RBF-SVM classification algorithm and 89% using a simple 3-NN. We can see that the equalization step allows a 13% increase in the accuracy when using the TL-CMR architecture as well. With respect to the simple CMR architecture, using the GFK transfer learning approaches increase the accuracy of almost 20%. In Table IV, the confusion matrix is reported for the best case of TL-CMR architecture. We can notice that most objects are correctly classified in all the trials, except object o (mouse) and object e (socket wrench). In particular, the mouse is confused in 40% of the cases with the small tape. The socket wrench is confused in 20% of the cases with the wrench and in 20% of cases with the spanner. For GFK, we choose the parameters $\delta = 27$, following the guideline in [38] and fine-tuning with a grid search. While GFK does well, methods based on principal component analysis perform badly and reduce the accuracy with respect to the CMR case. Our guess is that a dimensionality reduction based on the covariance matrix such as PCA is not convenient for the CLUE descriptor, which can contain important information with low variance. Even though GFK uses PCA internally, the final measurement of the distance is performed taking into account all the elements of the descriptors, as shown in Eq. (15). As a consequence, GFK does not cut off significant low-variance information. In order to give at least a rough clue on the computational time efficiency of the proposed approach, we estimated the time required for classifying the whole tactile dataset, given the knowledge acquired from the visual dataset. The classification time for the entire dataset is 0.6 s using CLUE + GFK with 3-NN implemented in MATLAB, and 0.02 s using SVM with a Python implementation. The algorithm was executed on a personal computer with an Intel i7-8550 CPU.

C. Comparison with human cross-modal object recognition

In order to have an ideal reference for assessing the performance of artificial cross-modal recognition, and because in the literature it is hard to find a cross-modal recognition algorithm, we compare the performance of our framework with a "golden standard", which is represented by the performance of humans. It is important to note that a rigorous, complete comparison with human performance is beyond the scope of the paper.

Feature Descriptor	1-NN	3-NN	Linear SVM	RBF SVM
SHOT + PCA	49.3%	48.0%	37.3%	41.3%
ESF + PCA	53.3%	58.6%	46.4%	57.3%
CLUE + PCA	42.6%	46.7%	57.3%	57.3%
SHOT + TCA	25.3%	18.6%	22.6%	44.0%
ESF + TCA	54.6%	64.0%	40.0%	60.0%
CLUE + TCA	44.0%	40.0%	54.6%	57.3%
SHOT + SA	37.3%	33.3%	24.0%	33.3%
ESF + SA	45.3%	50.6%	32.0%	52.0%
CLUE + SA	58.6%	56.0%	36.0%	40.0%
SHOT + GFK	50.6%	50.6%	29.3%	53.3%
ESF + GFK	68.0%	65.3%	54.6%	74.6%
CLUE + GFK	82.6%	82.6%	58.6%	81.3%

TABLE V: Cross-modal recognition result applying transfer learning without preprocessing. Visual data are used as a training set, tactile data are used as a test set.

Feature Descriptor	1-NN	3-NN	Linear SVM	RBF SVM
SHOT + PCA	53.3%	53.3%	48.0%	65.3%
ESF + PCA	54.8%	48.0%	56.0%	62.7%
CLUE + PCA	52.0%	46.7%	41.3%	53.3%
SHOT + TCA	41.3%	42.7%	17.3%	38.7%
ESF + TCA	41.3%	49.3%	52.0%	57.3%
CLUE + TCA	49.3%	49.3%	40.0%	52.0%
SHOT + SA	33.3%	34.7%	25.3%	24.0%
ESF + SA	42.7%	42.7%	14.6%	48.0%
CLUE + SA	42.7%	44.0%	34.7%	48.0%
SHOT + GFK	54.7%	54.7%	42.7%	66.7%
ESF + GFK	77.3%	81.3%	57.3%	77.3%
CLUE + GFK	88.0%	89.3%	76.0%	94.7%

TABLE VI: Cross-modal recognition result applying transfer learning with preprocessing. Visual data are used as a training set, tactile data are used as a test set.

The comparison with humans is intended to provide a rough indicator of reasonable cross-modal recognition performance. We arranged a simple experiment to have a first estimation of human performance in visuo-tactile cross-modal object recognition tasks. In this experiment, ten participants, ranging in age from 20 to 30 years, were invited to look at the set of objects shown in Figure 6 for 2 minutes. Afterwards, each participant was blindfolded and explored each object with one hand. Since human skin can sense also the temperature of the object, the participants wore a thin glove to maintain the tactile perception capability, but to reduce the perception of the temperature. The objects were placed on the table. The participants explored by touching each object for 10 s without seeing the objects. After that, the participants removed the blindfold and were invited to say which object was explored by looking at the objects. In more detail, this experimental protocol is summarized below:

- 1) The subject is invited to see all the objects for 2 minutes



Fig. 7: Tactile exploration performed by a human

a	0.95	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
b	0	1.0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
c	0.05	0	0.75	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
d	0	0	0.05	0.75	0.1	0	0	0	0	0	0.1	0	0	0	0	0	0	0	0
e	0	0	0	0.15	0.9	0.1	0	0	0	0	0	0	0	0	0	0	0	0	0
f	0	0	0	0	0	0.9	0	0	0	0	0	0	0	0	0	0	0	0	0
g	0	0	0	0	0	0	0.95	0	0	0	0	0	0	0	0	0	0	0	0
h	0	0	0.1	0	0	0	0.05	0.8	0.2	0	0	0	0	0	0	0	0	0	0
i	0	0	0.1	0	0	0	0	0.2	0.8	0	0	0	0	0	0	0	0	0	0
j	0	0	0	0	0	0	0	0	0	0.75	0	0	0	0	0	0	0	0	0
k	0	0	0	0.1	0	0	0	0	0	0	0.15	1.0	0	0	0	0	0	0	0
l	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0	0	0	0	0	0
m	0	0	0	0	0	0	0	0	0	0	0	0	0	0.95	0.05	0	0	0	0
n	0	0	0	0	0	0	0	0	0	0	0	0	0	0.05	0.95	0	0	0	0
o	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.0	0
	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o				

TABLE VII: Confusion matrix of human object recognition.

- 2) The subject is invited to wear a thin glove that prevents from sensing the temperature of objects
- 3) The subject is blinded with a blind fold
- 4) The objects are put in a bag, one object is picked out and put on a table
- 5) The subject explores the object with the hand for 10 s
- 6) The subject removes the blindfold
- 7) The subject sees all the objects and tells the name of the explored object
- 8) Go to step 1 until every object has been picked out

A picture of human tactile exploration is shown in Fig. 7. The average accuracy achieved is 89.7% and in Table VII the confusion matrix is reported. The accuracy of humans in this experiment is 12% better than the performance of the CMR method based on processed tactile point clouds and the CLUE descriptor. However, using the TL-CMR method the performance of the robot becomes 6% better than the human estimated performance. Comparing Tables III with VII, we can notice that for most objects the performance of CMR are close to human performance in this case study. It is important to note that the tactile skin does not measure the roughness since it only touches the object without rubbing.

D. Comparison with monomodal object recognition

The results of the cross-modal visuo-tactile object recognition framework are also compared with the monomodal visual and monomodal tactile recognition case. The results of the visual and tactile monomodal cases are reported in Table VIII. In this case study, the classifier is trained and tested with the same modality. The accuracy has been evaluated with a 10-fold cross-validation method. From Table VIII it is possible to see that both visual and tactile monomodal problems are, as expected, less challenging than the cross-modal case, since training set and test set are generated from the same perception modality. Most state-of-the-art descriptors achieve more than 90% accuracy in the monomodal case with 1-NN. We also notice that the performance of tactile classification is slightly less accurate than visual classification.

Feature Descriptor	Visual	Tactile
SHOT	97.17%	92.00%
ESF	97.33%	94.67%
CLUE	98.67%	94.67%

TABLE VIII: Monomodal Recognition Result

VI. CONCLUSION AND FUTURE WORK

In this work, we deal with robotic cross-modal visuo-tactile object recognition. We train a classifier by using visual data from an Asus Xtion Pro Live camera and we recognize objects at execution time only with tactile data, without any a-priori tactile information. The preliminary version of this work [3] showed for the first time that cross-modal visuo-tactile object recognition is feasible with respectable performance. It leverages empirical methods such as equalization of partiality and resolution, as well as a novel descriptor that performs well across different modalities. In this paper, we extend the framework proposed in [3] by combining the empirical ideas with formal transfer learning techniques. We show that combining the equalization of partiality and resolution, the CLUE descriptor, and a transfer learning techniques called geodesic flow kernel, we achieve an accuracy that is very close to the monomodal case. It is interesting to emphasize that our method reaches the peak of performance only when the transfer learning algorithms are combined with the pipeline proposed in [3]. Therefore, the mere application of TL without a preprocessing and without using CLUE cannot substitute the CMR effectively, but a smart combination significantly improves the results.

Future work will follow different directions. The first is implementing more complex exploration strategies. In this work, we used a simple strategy suitable mainly for quasi-planar rigid objects. In order to extend the TL-CMR to arbitrary objects, we have to introduce more sophisticated exploration algorithms, based for example on a nonprehensile manipulation strategy. Also, in current implementation the robot has to perform an extensive tactile exploration to build the point cloud. Novel exploration strategies can exploit the measurement of both tangential and normal contact forces to achieve a more effective exploration by manipulating the object and selecting specific sample points. A second direction is investigating novel visuo-tactile descriptors or in applying deep learning methods to train a model from a huge amount of visual data and transfer the acquired knowledge to a small amount of tactile data. Collecting a huge amount of tactile data, in fact, can be unpractical as robot and environments are in physical contact. In fact, robot movements for collecting data have a cost in terms of energy and tactile sensors can deteriorate, thus making collected data less accurate. On the other hand, visual images are less expensive to collect and, therefore, more suitable for approaches based on big data. A future work direction will be therefore to train big deep networks with visual data and reuse the knowledge with tactile data, without the need of huge tactile data collection. A starting point for our investigation will be 3D Convolutional Neural Networks (CNN) for object recognition [42], [43]. The third direction is to perform a larger number of experiments with human subjects to compare the performance of human and robot more accurately. Moreover, cross-modal perception will be applied in other fields, such as data efficient learning [44], [45] and object tracking [46].

ACKNOWLEDGMENT

This work has been partially funded by the Marie Curie Action LEACON, EU project 659265.

REFERENCES

- [1] "Psychology dictionary," <http://psychologydictionary.org/>.
- [2] S. Schumacher, T. B. de Perera, J. Thenert, and G. von der Emde, "Cross-modal object recognition and dynamic weighting of sensory inputs in a fish," *Proceedings of the National Academy of Sciences*, vol. 113, no. 27, pp. 7638–7643, 2016.
- [3] P. Falco, S. Lu, A. Cirillo, C. Natale, S. Pirozzi, and D. Lee, "Cross-modal visuo-tactile object recognition using robotic active exploration," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 5273–5280.
- [4] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [5] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning point cloud views using persistent feature histograms," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.
- [6] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *IEEE International Conference on Robotics and Automation*, 2009.
- [7] S. Salti, F. Tombari, and L. Di Stefano, "Shot: unique signatures of histograms for surface and texture description," *Computer Vision and Image Understanding*, vol. 125, pp. 251–264, 2014.
- [8] W. Wohlkinger and M. Vincze, "Ensemble of shape functions for 3d object classification," in *Robotics and Biomimetics (ROBIO), 2011 IEEE International Conference on*, 2011, pp. 2987–2992.
- [9] A. Johnson, "A representation for 3d surface matching," Ph.D. dissertation, Ph. D. thesis, Carnegie Mellon University, 1997.
- [10] A. Aldoma, Z.-C. Marton, F. Tombari, W. Wohlkinger, C. Potthast, B. Zeisl, R. B. Rusu, S. Gedikli, and M. Vincze, "Tutorial: Point cloud library: Three-dimensional object recognition and 6 dof pose estimation," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, pp. 80–91, 2012.
- [11] S. Luo, W. Mou, K. Althoefer, and H. Liu, "Novel tactile-sift descriptor for object shape recognition," *Sensors Journal, IEEE*, vol. 15, no. 9, pp. 5001–5009, 2015.
- [12] H. Liu, X. Song, J. Bimbo, L. Seneviratne, and K. Althoefer, "Surface material recognition through haptic exploration using an intelligent contact sensing finger," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 52–57.
- [13] G. De Maria, P. Falco, C. Natale, and S. Pirozzi, "Integrated force/tactile sensing: The enabling technology for slipping detection and avoidance," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 3883–3889.
- [14] M. M. Zhang, M. D. Kennedy, M. A. Hsieh, and K. Daniilidis, "A triangle histogram for object classification by tactile sensing," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 4931–4938.
- [15] A. Schneider, J. Sturm, C. Stachniss, M. Reiser, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 243–248.
- [16] Z. Pezzementi, E. Plaku, C. Reyda, and G. D. Hager, "Tactile-object recognition from appearance information," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 473–487, 2011.
- [17] Y. Gao, L. A. Hendricks, K. J. Kuchenbecker, and T. Darrell, "Deep learning for tactile understanding from visual and haptic data," *arXiv preprint arXiv:1511.06065*, 2015.
- [18] S. Gould, P. Baumstarck, M. Quigley, A. Y. Ng, and D. Koller, "Integrating visual and range data for robotic object detection," in *Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications-M2SFA2*, 2008.
- [19] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, 2013, pp. 3180–3186.
- [20] H. Liu, Y. Yu, F. Sun, and J. Gu, "Visual-tactile fusion for object recognition," *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 2, pp. 996–1008, 2017.
- [21] J. Ilonen, J. Bohg, and V. Kyrki, "Fusing visual and tactile sensing for 3-d object reconstruction while grasping," in *IEEE Int. Conference on Robotics and Automation (ICRA)*, 2013, pp. 3547–3554.

- [22] P. Güler, Y. Bekiroglu, X. Gratal, K. Pauwels, and D. Kragic, "What's in the container? classifying object contents from vision and touch," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2014*. IEEE, 2014, pp. 3961–3968.
- [23] O. Kroemer, C. H. Lampert, and J. Peters, "Learning dynamic tactile sensing with robust vision-based training," *IEEE transactions on robotics*, vol. 27, no. 3, pp. 545–557, 2011.
- [24] H. Liu, Y. Wu, F. Sun, B. Fang, and D. Guo, "Weakly paired multimodal fusion for object recognition," *IEEE Transactions on Automation Science and Engineering*, 2017.
- [25] J. M. Cloke, D. L. Jacklin, and B. D. Winters, "The neural bases of crossmodal object recognition in non-human primates and rodents: a review," *Behavioural brain research*, vol. 285, pp. 118–130, 2015.
- [26] G. Calvert, C. Spence, and B. E. Stein, *The handbook of multisensory processes*. MIT press, 2004.
- [27] A. N. Meltzoff and R. W. Borton, "Intermodal matching by human neonates," *Nature*, 1979.
- [28] R. K. Davenport, C. M. Rogers, and I. S. Russell, "Cross modal perception in apes," *Neuropsychologia*, vol. 11, no. 1, pp. 21–28, 1973.
- [29] M. Meier, M. Schöpfer, R. Haschke, and H. Ritter, "A probabilistic approach to tactile shape reconstruction," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 630–635, 2011.
- [30] K. Charusta, D. Dimitrov, A. J. Lilienthal, and B. Iliev, "Extraction of grasp-related features by human dual-hand object exploration," in *International Conference on Advanced Robotics*, 2009.
- [31] D. Levin, "Mesh-independent surface interpolation," in *Geometric modeling for scientific visualization*. Springer, 2004, pp. 37–49.
- [32] "Downsampling a pointcloud using a voxelgrid filter," http://pointclouds.org/documentation/tutorials/voxel_grid.php.
- [33] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 1–4.
- [34] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. M. Kwok, "A comprehensive performance evaluation of 3d local feature descriptors," *International Journal of Computer Vision*, vol. 116, no. 1, pp. 66–89, 2016.
- [35] D. Kalman, "A singularly valuable decomposition: the svd of a matrix," *The college mathematics journal*, vol. 27, 1996.
- [36] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang, "Domain adaptation via transfer component analysis," *IEEE Transactions on Neural Networks*, vol. 22, no. 2, pp. 199–210, 2011.
- [37] B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars, "Unsupervised visual domain adaptation using subspace alignment," in *Proceedings of the IEEE international conference on computer vision*, 2013, pp. 2960–2967.
- [38] B. Gong, Y. Shi, F. Sha, and K. Grauman, "Geodesic flow kernel for unsupervised domain adaptation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012, pp. 2066–2073.
- [39] C. M. Bishop, *Pattern recognition and machine learning*. springer, 2006.
- [40] A. Cirillo, F. Ficuciello, C. Natale, S. Pirozzi, and L. Villani, "A conformable force/tactile skin for physical human-robot interaction," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 41–48, 2016.
- [41] A. Cirillo, P. Cirillo, G. De Maria, C. Natale, and S. Pirozzi, "An artificial skin based on optoelectronic technology," *Sensors and Actuators A: Physical*, vol. 212, pp. 110–122, 2014.
- [42] D. Maturana and S. Scherer, "Voxnet: A 3d convolutional neural network for real-time object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 922–928.
- [43] J. Huang and S. You, "Point cloud labeling using 3d convolutional neural network," in *Pattern Recognition (ICPR), 2016 23rd International Conference on*. IEEE, 2016, pp. 2670–2675.
- [44] M. Saveriano, Y. Yin, P. Falco, and D. Lee, "Data-efficient control policy search using residual dynamics learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2017*.
- [45] P. Falco, A. Attawia, M. Saveriano, and D. Lee, "On policy learning robust to irreversible events: An application to robotic in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1482–1489, July 2018.
- [46] S. Koo, D. Lee, and D.-S. Kwon, "Incremental object learning and robust tracking of multiple objects from rgb-d point set data," *Journal of Visual Communication and Image Representation*, vol. 25, no. 1, pp. 108–121, 2014.



Pietro Falco received his PhD in 2012 at University of Campania Luigi Vanvitelli. From December 2010 to July 2011 he was a visiting scholar at Karlsruhe Institute of Technology (KIT). Until 2015, he worked within several EU projects such as DEXMART, ECHORD, SAPHARI and EUROC at University of Campania. In 2011, Pietro co-founded Aeromechs srl, a successful startup in the field of energy management for aeronautics and home automation. In 2015, he left the active co-leadership of Aeromechs and moved to Technical University of Munich, in order to continue research and teaching activities in robotics. He was awarded a TUM Foundation Fellowship sponsored by Rohde & Schwarz in March 2015 and a Marie Curie Individual Fellowship for experienced researchers in March 2016 with the project LEACON "LEARNING-CONTROL tight interaction: a novel approach to robust execution of mobile manipulation tasks". From January 2018, Pietro is a tenured senior scientist at ABB corporate research, Sweden. His research focuses on control theory and machine learning for robotics, human-robot interaction, and human motion interpretation. Pietro is serving as associate editor for IEEE Robotics and Automation Letters (RA-L).



Shuang Lu was born in Heilongjiang, China, in 1987. She received the B.E. degree in electrical engineering from Northeast Electric Power University, China, in 2009, and M.Sc. degree in electrical engineering and information technology from Technical University of Munich, Germany, in 2016.



Ciro Natale Ciro Natale received the Laurea degree and the Research Doctorate degree in Electronic Engineering from the University of Naples in 1995 and 2000, respectively. From 2000 to 2004 he has been Research Associate at the Department of Industrial and Information Engineering of the University of Campania "Luigi Vanvitelli", where he currently holds the position of Full Professor of Robotics and Mechatronics. From November 1998 to April 1999 he was a Visiting Scholar at the German Aerospace Center (DLR) in Oberpfaffenhofen, Germany. His research interests include modeling and control of industrial manipulators, force and visual control, cooperative robots, as well as modeling and control of flexible structures, active noise and vibration control and modeling, identification and control of smart actuators. He has published more than 120 journal and conference papers, including two monographs published by Springer. He served or is serving as Associate Editor of various international journals and he had responsibility roles in many European projects funded under the 7th and the 8th Framework Programmes.



Salvatore Pirozzi was born in Napoli, Italy, on April 1977. He received the Laurea and the Research Doctorate degree in Electronic Engineering from the Seconda Università degli Studi di Napoli, Aversa, Italy, in 2001 and 2004, respectively. From 2008 to 2016 he has been Research Associate at Seconda Università degli Studi di Napoli. He currently holds the position of Associate Professor at Università degli Studi della Campania "Luigi Vanvitelli". His research interests, in the aeronautics application sector, include modeling and control of smart actuators

for active noise and vibration control. His research activities on robotics include design and modelling of innovative sensors, in particular of tactile solutions, as well as interpretation and fusion of data acquired from the developed sensors. He published more than 70 international journal and conference papers and he is co-author of the book "Active Control of Flexible Structures", published by Springer. He currently serves as Associate Editor of the IEEE Transaction on Control Systems Technology.



Dongheui Lee is an associate professor at Department of Electrical and Computer Engineering, Technical University of Munich (TUM) and a head of human-centered assistive robotics group at DLR, Germany. Prior, she was assistant professor at TUM (2009-2017), project assistant professor at the University of Tokyo (2007-2009), and a research scientist at Korea Institute of Science and Technology (KIST) (2001-2004). She received mechanical engineering B.S. (2001) and M.S. (2003) degrees from Kyunghee University, Korea and PhD degree

(2007) from the department of Mechano-Informatics, the University of Tokyo, Japan. She was awarded for a Carl-von-Linde fellowship at TUM Institute for Advanced Study (2011) and a Helmholtz professorship prize (2015). She is Coordinator of euRobotics Topic Group on physical Human Robot Interaction and the co-coordinator of TUM Center of Competence Robotics, Autonomy and Interaction. Her research interests include human motion understanding, human robot interaction, machine learning in robotics, and assistive robotics.