



HAL
open science

MCA-NMF: Multimodal Concept Acquisition with Non-Negative Matrix Factorization

Olivier Mangin, David Filliat, Pierre-Yves Oudeyer

► **To cite this version:**

Olivier Mangin, David Filliat, Pierre-Yves Oudeyer. MCA-NMF: Multimodal Concept Acquisition with Non-Negative Matrix Factorization. 2015. hal-01137529v2

HAL Id: hal-01137529

<https://hal.science/hal-01137529v2>

Preprint submitted on 4 May 2015 (v2), last revised 29 Oct 2015 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

MCA-NMF: Multimodal Concept Acquisition with Non-Negative Matrix Factorization

Olivier Mangin, David Filliat, and Pierre-Yves Oudeyer

Abstract—In this paper we introduce MCA-NMF, a computational model of the acquisition of multimodal concepts by an agent grounded in its environment. We propose such computational models as an answer to the question of what concepts are and not only of how they can be learnt; more precisely our model finds patterns that characterizes associations across modalities. We detail why multimodality is essential to lower the ambiguity of learnt concepts as well as communicate about them. We then present a set of experiments that demonstrate the learning of such concepts from real non-symbolic data consisting of sounds, images, and motion acquisitions. Finally we consider structure in perceptual signals and demonstrate that a detailed knowledge of this structure, named *compositional understanding* can emerge from, instead of being a prerequisite of, global understanding. An open-source implementation of the MCA-NMF learner as well as scripts to reproduce the experiments are publicly available.

Keywords—multimodal concept acquisition; language acquisition; speech processing; vision; motion; non-negative matrix factorization; structure learning

I. INTRODUCTION

Whether they are seen as grounding the meaning of words or giving structure to our high level perception, concepts are an essential aspect of human cognition as well as a desirable feature for robots and other artificial cognitive systems. Yet the exact definition of what concepts are, let alone, the question of how they may be acquired by an agent in interaction with its physical environment are still mostly unanswered.

A. Multimodal concepts

Concepts may be defined as *mental representations*. What these mental representations are, would they be forged by a human or a robot, is however not firmly defined. Thus, before building models of the acquisition of concepts, and therefore indirectly of the concepts themselves, some aspects must be clarified. We propose here a perspective on learning concepts as patterns in a flow of multimodal perception, a view that is in line with the embodied cognition principle (see [Wilson, 2002](#)) and perceptual concepts as the ones from [Barsalou \(1999\)](#) (for a more detailed review see [Cangelosi, 2010](#)).

More precisely we consider the notion of *multimodal concepts* for several reasons. First, many concepts do not lie in a single perceptual modality. For example, the emergence of the concept of *dog* is not only related to the ability to recognize pictures of dogs but also to the sound of a dog barking and the touch of a dog’s fur. Furthermore, many concepts cannot be completely characterized without grounding them on several modalities: the concept *metallic* cannot be characterized without taking into account its perceptual expression on several

modalities (for example visual aspect, sound, touch, or taste), together with the recognition of the spoken or written word.

Another reason is that concepts often occur in the context of language and hence involve a linguistic modality. This aspect is emphasized in particular in the symbol grounding problem, as introduced by [Harnad \(1990\)](#) and discussed by [Glenberg & Kaschak \(2002\)](#). It points out that learning language is not only about learning the signs of communication such as words, but also requires to relate them to their semantic content, that emerges from and is grounded in the interaction with the world. This process is denoted as the *semiotic association*. From that perspective, learning a concept may involve learning and relating its semantic content to a given symbol. Natural communication channels however do not contain such thing as absolute symbols but rather manifestations of these symbols: the same word is never heard the same twice and the image of that word written is not perceived the same depending on the font it is written with, the angle it is viewed from, as well as the ambient luminosity. Thus perceiving symbols is by itself an analogous problem to perceiving their meaning. Hence the symbol grounding problem becomes a weak instance of the larger problem of learning meanings in one modality and linguistic symbols in another modality. Also, human communication is not generally reduced to one modality such as speaking or writing; instead full featured communication makes extensive use of facial expressions, physical contact, and eye gaze. A famous evidence of the multimodal nature of communication was given by [McGurk & MacDonald \(1976\)](#) and is referred as the *McGurk effect*: observing lips pronouncing ‘ga’ while hearing ‘ba’ is most often reported as perceiving the sound ‘da’ (see also [Schwartz, 2010](#)). In that perspective the multimodal character of natural language makes it very similar to the kind of concepts it may refer to; so that linguistic elements may be seen as multimodal concepts themselves. Furthermore we propose to study the learning of multimodal concepts that span several modalities, including linguistic and non-linguistic channels. In particular this approach emphasizes the co-organization of language and meaning, which is in line with growing evidences of the influence of language in learning concepts (see [Lupyan et al., 2007](#), [Waxman Sandra & Markow Dana, 1995](#)).

However trying to define what it means to learn multimodal concepts reveals that this task is prone to many ambiguities. As pointed out by [Belpaeme & Morse \(2012\)](#): “*The challenge which cross-situational learning needs to solve is not only one of mapping a word to a meaning, but of distinguishing that meaning from possible distractors.*” Indeed, Quine’s *indeterminacy of reference* ([Quine, 1960](#)) states that relating words

to meanings when learning a foreign language is intrinsically ambiguous. On the other hand, many models of learning semantic components from one modality also encounter similar ambiguity issues. For example, [Cederborg & Oudeyer \(2013\)](#) draw a parallel between Quine’s indeterminacy and ambiguity in imitation learning, that they call *the motor gavagai problem*. Another example is encountered with concepts that corresponds to categories. Indeed infants learning categories face the alternative possibilities of *thematic* and *taxonomic* associations of concepts as explained in [Markman \(1990\)](#): thematic association refers to the association of concepts that are related because they interact together, as *milk* and *cow*; taxonomic association refers to concepts that belongs to the same class, such as *cow* and *pig*. Other analogies can be drawn between this phenomenon and the ambiguity of word segmentation (see [Brent, 1999](#)), but also with multistability phenomenon in perception as described by [Blake \(1989\)](#), [Leopold & Logothetis \(1999\)](#), [Schwartz et al. \(2012\)](#), and the *cocktail party effect* (see [Cherry, 1953](#)). Many multimodal problems feature ambiguity in one or several modalities, but, in a somehow paradoxical manner, integrating information from several modalities may be efficiently used to overcome such ambiguity. In other words, considering the problem of concept learning separately in each modality suffers from the presence of ambiguities, but looking at the same problem in several modalities at the same time might help resolving that ambiguity instead of increasing it. For example the role of multimodal perceptions relatively to multistability is discussed by [Schwartz et al. \(2012\)](#). Similarly [Schwartz et al. \(2004\)](#) explore the role of vision of the lips for improving intelligibility of spoken sound and [Sodoyer et al. \(2004\)](#) present an algorithms for source separation taking advantage of audio-visual information. Finally, [Massera et al. \(2010\)](#) demonstrate, in an other experiment, that a robot can reach better performance on a motor task when a linguistic guidance is present, even if it does not initially know the meaning of the symbols composing the guidance.

B. Learning associations

A central issue toward defining the problem of learning multimodal concepts is to formalize what it means to *learn* a concept. Importantly *concepts* do not necessarily refer to an explicit representation; this notion rather targets emerging behaviors that are interpreted as the mastering of that concept. For example a child is said to master the concept *dog* not by looking into his brain for a neuron spiking each time a dog is seen but rather by its ability to relate the sight of a dog with the sound of a barking dog or the sound of the name ‘*dog*’. One way of modelling this is to specify a behavioral evaluation of the learning process. In particular this article focuses on the ability to classify stimuli in specific way.

From a machine learning perspective, unlike supervised learning, unsupervised learning, or reinforcement learning, multimodal learning is not a specific class of algorithms. Indeed, multimodal data can be treated as unimodal data on which an unsupervised learning is applied (some examples provided in this article fall under this category). It can also

be considered a supervised regression problem that consists in predicting the signal in one modality, knowing the others. The focus of this work is on learning that occurs in an unsupervised manner, that is how multimodal perception self-organizes in a way that can explain the emergence of concepts. The kind of behaviors under consideration are classification behaviors. However they do not correspond to supervised classification in machine learning, that is the association of a symbol to a given stimuli, but rather to unsupervised association between stimuli from a same semantic class: in the way a child would group a car with a truck instead of with a cow.

This article actually only focuses on one type of concepts, characterized by cross-modal associations, and a model of their acquisition. An example is the concept *dog* with its visual, acoustic, acoustic as language, and tactile manifestations. Although this notion of concept may seem very limited, we claim that the purpose of building such model is indeed to better explore these limits.

In practice, we build experiments that involve two phases: an unsupervised learning phase where the system observes raw perception, and a behavioral evaluation where it solves a task. This separation opens a perspective on the relation between the properties of the perceptual signal available during learning and the nature of the learnt concepts that is specified by the evaluation task. In particular it raises the question of the drives and cues that enable the self-organization of multimodal perception. In the case of language learning, experiments on children performed by [Akhtar & Montague \(1999\)](#), and [Smith & Yu \(2008\)](#) demonstrate that *cross-situational learning*, which focuses on elements that are persistent in the environment across different uses of a word, might be used by children to learn the meaning of words. Most of this work relies on cross-situational learning to explain or model the acquisition of lexicons of concepts. Other mechanisms such as *the whole object assumption*, *mutual exclusivity* (see [Markman, 1990](#)), and *conceptual reasoning* ([Landau et al., 1998](#)) are also known to play a role in the process of associating linguistic labels to concepts, but are out of the scope of this article.

C. Structure and complexity

Natural perception generally consists in *complex* visual or acoustic scenes rather than in isolated pictures of objects or isolated sounds of words. In other words, an essential aspect of these scenes is their structure; and an essential capability for a learner is to leverage that structure to overcome the complexity of its perception. Such structure can however take several forms: for example a visual scene may contain several objects at various positions; a spoken sentence is composed of a sequence of words, the words themselves are made from basic phonemes and the sequence obey to a specific grammar.

A common intuition is that the complexity of the learnt concepts gradually increases along learning. In particular *cumulative learning* consists in gradually acquiring a lexicon of elements of increasing complexities, such that new elements can be obtained as the combination of simpler ones (see sec. 3–4 from [Cangelosi et al., 2010](#)). Actually the definition of what *complex* and *simple* mean may by itself be ambiguous;

therefore we focus in this article on a specific structure that is the combination of several basic elements in complex scenes such as objects in an image or words in a sentence.

Motor synergies and motion primitives are other examples of building blocks that can be combined to model complex perception and action; they were introduced by both motor control theorists and roboticists (see [Konczak, 2005](#)). For example [Mussa-Ivaldi & Bizzi \(2000\)](#) interpret a group of experiments on the control system of frogs and rats as giving strong evidence that the brain uses a set of primitive force fields that are combined linearly to generate a diversity of motor behaviors. [Tresch & Jarc \(2009\)](#) provide a more detailed review on that subject.

The structure of *complex* perceptual stimuli also adds more ambiguity to the problem of the emergence of concepts. We already mentioned examples of that ambiguity such as word segmentation (see [Brent, 1999](#)). One way to handle that ambiguity is to frame the learning on multimodal concepts, thus relating the structure in one modality with the structure in others. For example [Tuci et al. \(2011\)](#) provides a model of multimodal learning for symbolic language and real actions. Their experiment demonstrates that learning a compositional structure shared between action and language can allow robotic agents to achieve better generalization of the acquired motor knowledge. More precisely the linguistic input received by the system shapes a model of the structure of actions and makes the system capable of achieving behaviors that were not encountered in training.

Another interesting question arises from the notion of structured concepts and their acquisition and challenges the meanings of the words *simple* and *complex*. [Wrede et al. \(2012\)](#) contrast *compositional* understanding, that describes an agent that is aware of the local components and their combination into a global perception or action, and *teleological* understanding, that accounts for an agent that understands perception globally. More precisely the term teleological refers to a pragmatic emphasis on using the global knowledge toward a goal, even without refined understanding of its structure. According to [Wrede et al. \(2012\)](#) the developmental path of infants goes first through teleological understanding before reaching compositional understanding. This developmental path is to contrast to the one stating that compositional understanding occurs first before any usage of the knowledge (for example as suggested by [Cangelosi et al., 2010](#)).

D. Contribution

In this work we introduce a model named *Multimodal concept acquisition with non-negative matrix factorization* (MCA-NMF), of the learning of cross-modal concepts through the formation of structure in multimodal low-level signals (vision, speech sounds, gestural motions). We then present experiments combining the learning of dance gestures from human demonstrations, of words from full spoken sentences, and of visual objects from images.

We demonstrate the acquisition of grounded complex concepts from raw continuous signal only, without relying on symbols to train the artificial agent. Our learning experiment

exploits cross-situational association in the training perceptual data. Language learning, and in particular the learning of words, is in MCA-NMF treated as an instance of multimodal learning. This means that the linguistic data (here speech) is not handled in a specific way but rather in a symmetric manner with respect to other perceptual data. Finally we explore the question of the structure and (de)composition of concepts. In particular we show that it is possible for an artificial system to discover subcomponents of perception, such as words in spoken sentences, although the system is only exposed to a task that requires to associate whole sentences to corresponding scenes.

The next section presents the experimental framework as well as the datasets used in the experiments. section III details the algorithms that forms the MCA-NMF learner, including the low-level processing of sensor data. section IV provides the results and their analysis for the experiments on learning multimodal concepts as cross-modal associations. It also provides two additional experiments on the discovery of words inside sentences and the explicit representation of concepts by the learner. Finally we discuss our contribution and its articulation to previous work in section V.

II. MATERIALS

This section presents the experimental framework that we use to explore the question of learning multimodal concepts from perception. It first details the kind of concepts we consider, then explains the experimental setup, and finally introduces the datasets we use.

A. Target concepts

This article presents a system that learns to discover patterns that characterize associations in multiple modalities: that are patterns in some modality that are systematically associated with other patterns in other modalities. We perform several experiments that explore how the learner manages to represent semantic relations between the modalities. In practice we consider semantic relations that may correspond to either an essential relation, as the one relating the barking to the image of the dog, or a conventional relation as the one relating the name ‘*dog*’ to images of dogs.

The essential relation arises from the reality of an object that has manifestations in several modalities. There exists such a thing as a dog that has manifestations in the visual modality as images of the dog, in the touch modality as the touch of the dog’s fur or its claws, or in the acoustic modality as the sound of the dog barking. Although not all of these manifestations occur each time the dog is encountered, they are often perceived simultaneously since they corresponds to the actual presence of the dog. On the other side, the conventional relation characterises language: the word ‘*dog*’ is often pronounced when a dog is present and is the object of attention. It is extensively used by parents to teach new words to children.

Importantly, both relations manifest through the joint occurrence of frequent patterns in several modalities; therefore a mechanism leveraging such cross-situational information

would be able to learn both. In the following we denote by *semantic concept* the set of manifestations of such an object in perceptual modalities, either related essentially or by convention. Additionally a semantic concept may have several manifestations in a single modality. For instance a dog is associated to both the touch of its fur and claws, or to the sound of the dog barking and the word ‘dog’. The semantic relations we consider actually include Peirce’s icon, index, and symbol (see 3.1 in Steels, 2008). In the following, the only cue about the semantic relations in the stimuli is that related elements occur simultaneous in the various modalities; this corresponds to the cross-situational information.

B. Experimental framework

We consider the situation in which an intelligent system perceives a scene composed for example of objects or motions while hearing sentences that describe the scene. Such a setup is illustrated in fig. 1.

In the experiments we describe below, the modalities that the system observes can vary from one experiment to the other and it is not necessary that one modality is linguistic; however a semantic relation always exists between some elements of the different modalities. These elements might be of several natures: gestures in motions, object in visual scenes, or words in spoken utterances. We consider semantic relations as mappings between these elements: for example a word is related to a gesture, or a gesture to an object in a scene.

In the following, each experiment consists in two phase: one training phase during which the system is exposed to multimodal data and an evaluation phase in which we test the success of concept learning.

During its training the learning agent observes examples of scenes; each example is a set of one observations in each modality. In each scene only one relevant multimodal concept is present and is observed in several modalities. For example a sentence is heard containing the word ‘dog’ and a picture of a dog is seen. Although only one multimodal concept is present in each example, not all perceived elements are meaningful, that is to say related to elements in other modalities. For instance many words appear in spoken utterances that are not semantically related to anything in other modalities. For example in the sentence “Look at the circles I do”, only the word ‘circle’ is related to the observed gesture. Similarly several objects may appear in the visual scene while only one is related to the subject of the sentence.

During its evaluation the learner observes a test example in only one modality. It then has to chose between several other examples in another modality, the one that best matches the test example. For example the system hears a sentence talking about a dog and has to chose between several pictures the one containing a dog. We denote this task *cross-modal classification*; it forms an evaluation protocol that does not require one modality to be symbolic.

Interestingly this experiment is very similar to the one performed by developmental psychologists to study the role of various heuristics used by children for the acquisition of words, including cross-situational information, as in the




Fig. 2. Dance motions were performed by a human dancer and perceived through skeleton tracking based on a 3D kinect sensor. The figure illustrates some of the gestures demonstrated to the learner.

works from Markman (1990), Yurovsky et al. (2012). Unlike many computational approaches presented in previous works on multimodal learning, we present an evaluation of the performance of the learner that is not on a regular machine learning classification task. Instead the learner is evaluated on its ability to relate elements from distinct non-symbolic modalities, similarly to how one would evaluate a children. In additional experiments presented in section IV-C, we also evaluate the emergence of the recognition of single words inside sentences. More precisely the learner is asked to recognize parts of sentences instead of full utterances and this enables to localize the perceived semantics inside the speech sentences. An experiment also studies the mapping between the internal representation of the system and the semantics information to detect the emergence of explicit representations of the learnt concepts.

C. Datasets


The following experiments involve three raw modalities: motion, sound, and image.

1) *Motions*: In the experiments we use a dataset of choreography motions demonstrated by a human and recorded through a *kinect* device. More precisely the dataset contains a total of 1100 records, each presenting one of 10 different gestures that are spanned over one or two limbs. Figure 2 and table II illustrate the kind gestures that compose the datasets. Typical gestures are the *walk* gestures, involving the two legs, in which the human demonstrator mimics walking in place, or the *left hello* in which the demonstrator moves his left arm pretending to say ‘hello’. The data is publicly available and presented in more details at <http://flowers.inria.fr/choreo2> and is used by Mangin & Oudeyer (2012, 2013).

2) *Speech*: The acoustic records used in the following experiments are taken from the Caregiver dataset (Altosaar et al., 2008), provided by the ACORNS project (Boves et al., 2007). It is composed of 1000 utterances containing 13 keywords. Each utterance is spoken by 4 speakers in English adult directed speech. In the experiments we only use utterances from one speaker. An example of sentences used in the dataset is *Angus is lazy today*. where the keyword is *Angus*. Other examples of transcriptions from utterances from the dataset are given in table I.

3) *Images*: Pictures used in the experiments were acquired as frames from an interaction with an iCub robot, through an RGBD sensor (red, green, and blue camera coupled with a depth sensor). The acquisition was performed and is described in more details by Lyubova & Filliat (2012). Each image contains one visual object and possibly the hand of the

a) *Training*: The learner observes a set of examples of images each of which is paired with a spoken descriptions of the object in the image.



b) *Testing*: The learner hears a new spoken description and must choose its best match among objects in a small set of images.




Fig. 1. Illustration of the cross-modal classification task on which one learner presented in this article is tested. The learner perceives multimodal signal, vision and sound in the figure, and learns the associations between the appearance of objects and the sound of their names. It then has to prove its ability to choose the right object among a small set when hearing its name. In practice we use images of objects, motion captures, and recorded spoken utterances in the experiments.

<p>We take a bath To put it in the bath isn't funny either The shoe is a symbol Now mummy is losing her patience Daddy comes closer Angus takes off her shoe Daddy never calls She sits on a nappy Now everybody is in the car Where is the nappy</p>
--

TABLE I

TRANSCRIPTIONS FROM TEN RANDOM EXAMPLES FROM THE ACORNS CAREGIVER DATASET FROM [ALTOSAAR ET AL. \(2008\)](#). KEYWORDS ARE IDENTIFIED IN BOLD FONT.



Fig. 3. Example frames from the image dataset. This frames feature the following objects: *blue whale*, *yellow car*, *teddy bear* and *moose*. Color circles corresponds to local descriptors detected by the system. Interestingly, the hand of the operator appears in some pictures as in the last one. The area observed by the system is actually larger than the one represented in the figure.

operator. In the experiments we use images from a subset of 10 objects each appearing in more than a thousand frames. During the acquisition, the objects are moved and rotated. Hence they are presented from distinct points of view and they may be partially cluttered by the hand of the operator. Examples of frames from the dataset are presented in fig. 3.

III. METHODS

A. Open-source code to reproduce the experiments

The implementation of the MCA-NMF and the code used in the experiments presented in this paper are available publicly and openly (BSD license) on <http://github.com/omangin/multimodal>. It consists in a set of tools that implement the

MCA-NMF system, including the NMF algorithm and code to achieve multimodal learning with it, as well as the scripts that corresponds to the experiments presented here and produce their results. It includes the methods to process sound and motion data. Features extracted from the datasets are also publicly available. The code is meant for the reproduction of the experiment we presented as well as for the development of new experiments based on the same framework.

B. The MCA-NMF model

This section presents the algorithmic tools behind MCA-NMF. They are based on the nonnegative matrix factorization algorithm presented by [Paatero & Tapper \(1994\)](#), [Lee & Seung \(1999\)](#), that we use very similarly to [ten Bosch et al. \(2008\)](#), [Driesen et al. \(2009\)](#), [Mangin & Oudeyer \(2012\)](#).

The first part of this section presents the learning of a multimodal dictionary; it is then explained how the learned dictionary provides a representation of data that is not bound to any modality; in the following this representation is referred to as the learner's *internal representation* of data. Finally we explain how the learner can transform data from one or several modalities to an internal representation or to an expected representation in unobserved modalities. Figure 4 provides an overview of the transformations from raw data to higher level internal representation by the learning agent.

We consider a setting in which the learner observes samples in several modalities. For example, the system visually observes objects while hearing a spoken description of the scene. We represent the perception of the samples in each modality by a vector v_a , where a denotes the modality (for example the system observes the objects as v_{image} and the sound description as v_{sound}). Details about such representations for the modalities used in the experiments are given in section III-C.

1) *Learning a dictionary of multimodal components*: We call *components* primitive elements that are mixed together into observations, in the same way that phonemes combine together into a word or a sentence. Compared to the common

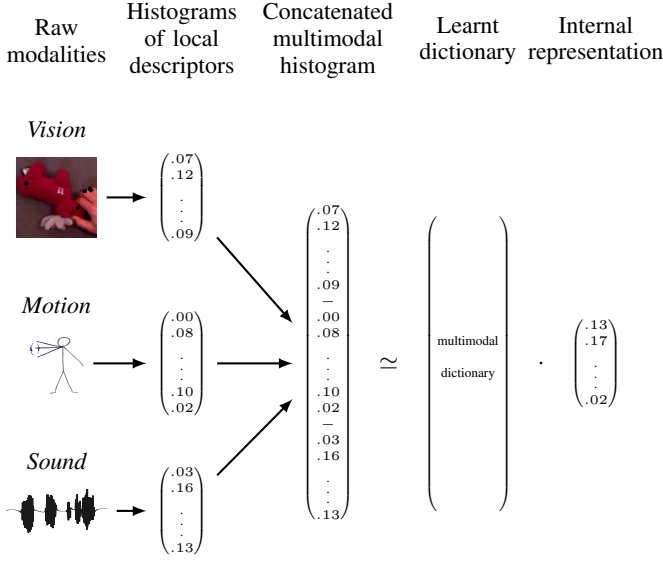


Fig. 4. The MCA-NMF system processing of data. First raw data from each modality is transformed into a histogram of local descriptors, that is a vector of nonnegative values. Then the histograms from each modality are concatenated into a histogram representation of the multimodal perception. During its training the MCA-NMF system uses a set of perceived example, each represented by such an histogram, to learn a multimodal dictionary that captures multimodal patterns. The system then uses the dictionary to transform perception into an internal representation. The internal representation can be obtained with all modalities observed (as on the figure) or with only a subset of modalities observed as explained further.

context of clustering, this notion of component is more general: observations are mixtures of several components at the same time, instead of being just a noisy observation of one centroid.

The learner implemented by MCA-NMF builds a dictionary of multimodal components according to the following model: it searches k components, each represented by a vector w_j (j from 1 to k), such that each observed example v^i verifies:

$$v^i \simeq \sum_{j=1}^k h_i^j w^j \quad (1)$$

where h_i^j are coefficients and \simeq denotes a notion of similarity between matrices that is defined below. This is equivalent to clustering when the w_j are the centroids and for each i only one h_i^j is nonzero and equals 1. We consider a more general case where w_j and h_i^j are only constrained to be nonnegative.

In the following, the set of n examples is represented by a matrix V of shape $d \times n$ (each example is a column of V), the set of components by a matrix W of shape $d \times k$, called *dictionary*, and the coefficients by a matrix H of shape $k \times n$. The previous equation, that models the objective of our learner, can thus be re-written for all observations as:

$$V \simeq W \cdot H \quad (2)$$

In order to fully define the reconstruction error between V and $W \cdot H$, we use a variant of the Kullback-Leibler divergence often called generalized Kullback-Leibler or I-divergence. The Kullback-Leibler divergence is originally an information theoretic measure of similarity between probability distributions.

The I-divergence is defined, for two matrices A and B of same shape, as $D_I(A||B)$ given by equation (3).

$$D_I(A||B) = \sum_{i=1}^d \sum_{j=1}^n \left(A_{i,j} \ln \left(\frac{A_{i,j}}{B_{i,j}} \right) - A_{i,j} + B_{i,j} \right) \quad (3)$$

In this paper in order to minimize $D_I(V||W \cdot H)$, we use the algorithm, based on multiplicative updates of W and H , that was originally presented in Lee and Seung's paper [Lee & Seung \(1999\)](#). This algorithm consists in alternating the two update steps from equation (4) where \otimes and $/$ denote Hadamard's (coefficient-wise) product and division on matrices.

$$H \leftarrow H \otimes \frac{W^T \cdot V}{W^T \cdot \mathbf{1}} \quad W \leftarrow W \otimes \frac{V}{\mathbf{1} \cdot H^T} \quad (4)$$

2) *NMF to learn mappings between modalities*: Previous section explains how, for a given set of observations from several modalities that is represented by a matrix V , the NMF algorithm can learn a dictionary W and a coefficient H matrices such that training examples are well approximated by the product $W \cdot H$.

We actually consider the case of data coming from several modalities (three in the example). More precisely we assume the data matrix V is composed of column vectors v such that:

$$v = \begin{pmatrix} v_{mod1} \\ v_{mod2} \\ v_{mod3} \end{pmatrix} \text{ and thus } V = \begin{pmatrix} V_{mod1} \\ V_{mod2} \\ V_{mod3} \end{pmatrix}.$$

The minimization of the I divergence induces a trade-off between error in one modality relatively to others. In order for the error in each modality to be treated on a fair level by the algorithm it is important that the average values in the representations are of similar magnitude. It can be easily obtained by normalizing data in each modality. In the following experiment data from each modality is normalized according to its average 1-norm.

Since the observations, that is to say the columns of V , are composed of several modalities, the dictionary W also split into several parts each corresponding to one modality. In other words each components can be seen as the concatenation of several parts: one for each modality. For example if the data is composed of three modalities: *mod1*, *mod2*, and *mod3*, there exist matrices W_{mod1} , W_{mod2} , and W_{mod3} such that:

$$W = \begin{pmatrix} W_{mod1} \\ W_{mod2} \\ W_{mod3} \end{pmatrix}.$$

In the following we interpret the columns of the matrix H , as an internal representation of the data by the learner. For example, an internal representation h is induced by an observation in modality one such that $v_{mod1} = W_{mod1} \cdot h$, or one in both modalities one and three by:

$$\begin{pmatrix} v_{mod1} \\ v_{mod3} \end{pmatrix} = \begin{pmatrix} W_{mod1} \\ W_{mod3} \end{pmatrix} \cdot h.$$

Interestingly, it is possible to use the learned dictionary to compute an internal representation of an example, even if the example is only observed in a subset of the modalities. Given an example observed only in one modality, v_{mod1} , one can search for an h such that v_{mod1} is well approximated as $W_{mod1} \cdot h$. More precisely this is equivalent to finding an h solution of:

$$\arg \min_h D_I(v_{mod1}, W_{mod1} \cdot h) \quad (5)$$




Fig. 5. Once the system has learnt the dictionary (W_{mod1} and W_{mod2}), given an observation v_{mod1} in one modality it can reconstruct the corresponding internal representation as well as the expected perception in another modality.

The NMF algorithm used in these experiments actually alternates steps minimizing $D_I(V||W \cdot H)$ with respect to W and H . Solving eq. (5) is equivalent to the NMF problem with respect to H only; therefore, it can be obtained with the same algorithm, but only using the steps that update H . In theory this approach scales to any number of modalities although the experiments presented here only test it on numbers from two to four.

Finally it is also possible to reconstruct a representation of the data that the system would expect in a modality, given observations in other modalities. For that, from an observation featuring a subset of the modalities, the system fits an internal representation h using the method described previously. Then it can reconstruct the expected representation in an unobserved modality (for example the third modality, $mod3$) by computing the product $W_{mod3} \cdot h$. This forms a framework, illustrated in fig. 5, that uses a learned multimodal dictionary to transform data from modalities to internal representations or expected data in other modalities. It enables a large set of experiments as illustrated in the following.

3) *Cross-modal classification without symbols*: The system is trained on various combinations of either two or three modalities. The modalities might be denoted as *Motion* or M , *Sound* or S , and *Image* or I . After being exposed to a set of training multimodal examples, the system is tested as follows: it observes a new example, called *test example* in a subset of its modalities and has to chose the best match among several examples observed in other modalities, which are denoted as *reference examples*. An illustration of that process is given by fig. 1. For example, the system is trained on sound and image and tested by hearing a sentence (the test example) and having to chose among a set of images (the reference

examples) the one that is best described by the heard sentence. Another possibility is to train the system on motions, sounds, and images, and test it on its ability to chose from several sentences the one that best describes a pair of a motion and an image that it observes. We denote such settings by the notation: $M1 \rightarrow M2$, where $M1$ represents the modality or modalities in which the test example is observed, called *test modalities*, and $M2$ the modality or modalities, denoted as *reference modalities*, in which a best matching example must be chosen among a set of reference examples. For example hearing a sentence and choosing the best matching object from images is denoted by Sound \rightarrow Image or $S \rightarrow I$. Viewing an object and a gesture and finding the best matching sentence amongst examples is denoted by $M, I \rightarrow S$. The testing process is illustrated in fig. 6.

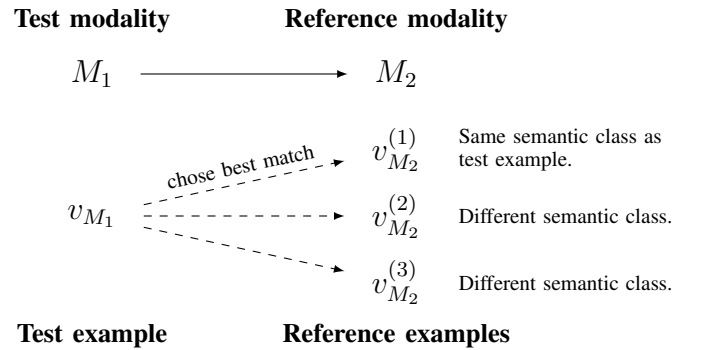



Fig. 6. The learner is tested on its ability to relate an observation of a test example in one modality to the right reference example in another modality.


Section III-B explains how to use NMF on multimodal data, to learn a dictionary and the associated *internal representation* and finally how to transform data either from one modality to another, or from a modality to the internal representation (see also fig. 5). We use that mechanism as a basis to implement a classification behavior for the learner. More precisely this refers to the ‘*chose best match*’ operation illustrated in fig. 6. In particular the issue is that the data to compare, that is representations of test and reference examples, are of different natures. Because it comes from distinct modalities, we use the ability of the learner to convert perception to its internal representation as well as to other modalities to achieve that behavior. To perform the comparison the MCA-NMF system can either:

- compute an internal representation of the test example, compute internal representations of the reference examples, and then compare these internal representations.




- compute an internal representation of the test example,

use it to generate an expected representation in the reference modality, and compare it to the reference examples.



- compute internal representations of reference examples, for each of them compute an expected representation in the test modality, and compare them to the test example.



The choice of one of these methods is referred as *the modality of comparison*. This is however not sufficient to fully define the system: in order to be able to chose a best matching reference example, the system needs a metric to perform the comparison. Several alternative metric could be chosen to perform the comparison. More importantly, the choice of the metric and its efficiency is highly dependent on the modality of comparison, as shown by the following results. We considered the following common metrics.

- **Euclidean distance**
- **Kullback-Leibler or I-divergence** The Kullback-Leibler and I-divergences are introduced by eq. (3). In the following we denote its usage as I-divergence (I-div.). By default the divergence from the test example to a reference example is computed; however since it is not symmetrical, we also experimented with the reversed divergence (that is to say the divergence from a reference example to the test example) and a symmetrized divergence obtained as: $D_{sym}(x||y) = \frac{1}{2}[D(x||y) + D(y||x)]$. None of the three approaches was systematically better in our experimentation.
- **Cosine similarity**¹ The cosine similarity is not a metric but can be used to compare vectors; it ranges between -1 and 1 and the biggest the value is, the most similar the vectors are. It is defined for two vectors x and $y \in \mathbb{R}^d$, denoting the scalar product, as:

$$\text{cosine_similarity}(x, y) = \frac{x \cdot y}{\|x\| \|y\|}$$

Other choices are possible. In our experiments, many modalities are represented by histograms, or concatenation of histograms, that are of high dimension. In these modalities the Euclidean norm is not necessary meaningful, this is why we

use other measures of similarity such as the I-divergence and the cosine similarity.

In the following, MCA-NMF is evaluated on its recognition success rate. It is defined as the proportion of correct recognition by the system: a recognition is correct when the system choses a reference example matching semantic concept from the test example.

C. Signal representation

Each of the three raw modalities used in the experiments is represented in a specific way. However all the representations are similar in their approach, all data is thus represented as nonnegative vectors, and share the important property of being additive that section III-C1 describes. The nonnegative and additive properties are required by the NMF algorithms and are thus limitations to the scope of its applications.

1) *Additive property*: Importantly, the following representations have a common additive property, that directly comes from the use of histograms of local events. For example in the acoustic modality, if two words which representation are w_1 and w_2 are concatenated into an utterance, which representation is denoted as s , then

$$s \simeq \lambda w_1 + (1 - \lambda) w_2$$

where $0 < \lambda < 1$. The approximation ignores the events coming from the border between the words. This important property transforms the sequencing operation into a convex combination. It therefore transforms a sentence into a mixture of its words, and similarly a word into a mixture of phonemes. Similarly the juxtaposition of several parts of an object is represented as the convex combination of the representations of each part. Several gestures combined in a motion are also approximately represented as the convex combination of the representations of the gestures.

2) *Motion*: In order to represent the recorded motion as vectors of nonnegative values in a way that makes it possible to use the algorithm presented in previous section we introduced the *histograms of motion velocity* representation in Mangin & Oudeyer (2012). In particular this representation consists in a simple histogram based representation of motion that can be seen as a rough approximation of the phase diagram of the dynamics of one body joint.

A kinect device captures the motions of a human demonstrator as trajectories in angle and angle velocity spaces of several articulations of the human body. Each trajectory on a specific body articulation (or degree of freedom) is considered separately and the entire sequence of angles and velocities is transformed into a histogram, represented by a fixed length non-negative vector. Vectors obtained for each degree of freedom are then concatenated into a larger vector as illustrates fig. 7. The device only captures angles and delayed velocities are computed to achieve better robustness to noise in the angle sequences. More precisely $\dot{x}_t = x_t - x_{t-d}$ is used to compute the velocities, instead of being restrained to the case where $d = 1$.

In Mangin & Oudeyer (2012), we explore and compare various alternative approaches to transform angles and velocities sequences into histograms. In the following experiments

¹http://en.wikipedia.org/wiki/Cosine_similarity

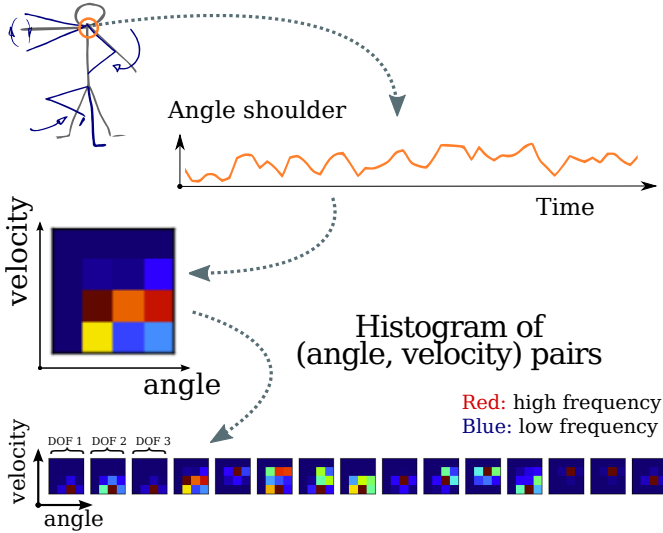


Fig. 7. Human motions are abstracted as concatenated histograms on joint position and velocities. In the final histograms, frequencies are represented through colors, x and y axis correspond respectively to values of angles and velocities. (Best seen in color)

we use angle-velocity histograms, that is to say histograms built on the joint space of angle and velocity. Furthermore we build histograms over a vector quantization that is an adaptive binning process using a k-means algorithm.

Representing motion data by separate histograms on each degree of freedom leads to two approximations: 1) for a given measurement in the trajectory, information about dependency between different degrees of freedom is dropped; 2) the sequential information between measures for a given degree of freedom is dropped. These simplifications are actually similar to successful ones from other fields, some of which we use for the sound and visual modalities as detailed further. Indeed [ten Bosch et al. \(2008\)](#) have demonstrated that, even if sequential information may appear necessary in language, and especially in speech utterances, very good word discovery can be achieved without considering this sequential information. Both in text classification and in computer vision *bag-of-words* techniques also achieve good performances by dropping positional information of extracted local features ([Joachims, 1997](#), [Sivic & Zisserman, 2008](#)). Finally using histograms built on joint angle positions and velocities is similar to representing transitions in angle space. By representing the sequence through its transition we approximate it by a Markovian process. Such an approximation is quite common in the gesture recognition and motion planning literature ([Calinon & Billard, 2009](#), [Kulic et al., 2009](#)). In the experiments, this process leads to a representation of motions that is of dimensionality 450.

3) *Sound*: In the following experiments we use the same representation of sound as in the works of [ten Bosch et al. \(2008\)](#), [Van Hamme \(2008\)](#), [Driesen et al. \(2012\)](#). Histograms of acoustic co-occurrences (HAC) were introduced as a representation of sound that is based on acoustic events. Similarly to the motion representation we have presented, it discards most of the sequential information of the acoustic events; it however consider co-occurrences of pairs of acoustic events

and uses a static approach to codebook construction.




Fig. 8. Sequence of transformations from raw (time sequence) acoustic signal to histograms of acoustic co-occurrence (HAC) representation.

The outline of the transformation from raw sound to HAC representation is given in fig. 8. The main steps are explained in more details below. The perception mechanism first segments the acoustic signal into a sequence of short time windows; it then computes Mel-frequency cepstral coefficients (MFCC) for each window. This transforms the original signal into a sequence of MFCC vectors. Additionally we consider dynamic information on top of the sequence of MFCC vectors of dimension 22; discrete derivatives are computed as defined by [Driesen \(2012\)](#) and denoted by the Δ operator to form the Δ MFCC and $\Delta\Delta$ MFCC. This transformation is analogous to the delayed velocities introduced previously to represent motions.

Then, the system computes three codebooks, for basic MFCC vectors and their Δ and $\Delta\Delta$ transformations, with a k-means algorithm, on a dataset of spoken language, as described in [Driesen \(2012\)](#). Following the implementation of [ten Bosch et al. \(2008\)](#), the sizes of the codebooks are $k = 150$ for MFCC vectors and Δ and $k = 100$ for the $\Delta\Delta$ vectors. The codebooks are used to convert the three sequences of MFCC vectors and their Δ and $\Delta\Delta$ transformations into a sequence of acoustic events: each cluster, that is to say each elements of a codebook defines an acoustic event; each time window is thus transformed into three discrete events corresponding to the clusters in which fall the three vectors associated to that time window.

The last step consists in removing most of the temporal information by building histograms of event occurrences and co-occurrences. This process happens on top of the three sequences of discrete acoustic events obtained from vector quantization. More precisely, co-occurrence histograms are simply histograms of the successive occurrences of pairs of events. What is denoted as HAC representation in the following is actually the concatenation of co-occurrence histograms for each one of the events categories, that is to say MFCC, Δ MFCC, and $\Delta\Delta$ MFCC events. The final dimensionality of the representation is more than 100,000; more detail on the HAC representation is given by [Driesen \(2012\)](#), [Van Hamme \(2008\)](#), [Mangin \(2014\)](#).

4) *Image and video*: The visual frame are transformed with the same tools as in [Lyubova & Filliat \(2012\)](#). More precisely, two types of *local features* are extracted from the pictures: SURF descriptors ([Bay et al., 2006](#)) and HSV (hue, saturation, value) of superpixels.

As described by ([Bay et al., 2006](#)), SURF features are

Keyword	Object	Motion	
		Limb(s)	Description
shoe	blue octopus	both legs	squat
nappy	teddy bear		walk
book	pink octopus	right leg	raise heel toward left knee
daddy	yellow car	both arms	clap
mummy	blue whale		mimic paddling left
Angus	blue-eyes-green-yellow orange fish	right arm	mimic punching
bath			starts horizontal and goes from side to front
bottle	squirrel	left arm	starts horizontal, forearm goes down to 90°
telephone car	mouse cube		waving arm hello sign

TABLE II

WE FORM ARBITRARY MULTIMODAL CONCEPTS BY ASSOCIATING AN OBJECT WITH A KEYWORD AND A GESTURE. THE TABLE PRESENTS A LIST OF SUCH ASSOCIATIONS. THE LIMBS ON WHICH THE MOTIONS OCCUR ARE ALSO MENTIONED. THE SYSTEM THEN OBSERVES THE KEYWORDS AS SPOKEN SENTENCES CONTAINING THE KEYWORD, OBJECTS AS IMAGES, AND GESTURES AS HUMAN DEMONSTRATIONS PERCEIVED BY A MOTION CAPTURE SYSTEM.

descriptors computed around points of interest in an image. They enable to transform an entire image into a set of small descriptors that are invariant to change in scale and rotation of the image. To obtain the HSV descriptors, areas with relatively homogeneous colors are detected into what is called *superpixels* (Micusik & Kořecka, 2009); then each area is represented as the triplet of its hue, saturation and value.

Both methods leads to the representation of an image as a set of local descriptors. Then the descriptors of both type are quantized by incrementally learning growing dictionaries of features. This algorithm used is a variant of k-means, very similar to the one presented in Filliat (2008) and for sound in Mangin et al. (2010). This process is the analogous to the grouping of MFCC vectors into acoustic events described in previous section.

Finally the occurrences of local event are counted and summarized in a histogram, that is a nonnegative vector of dimensionality around 50,000.

IV. EXPERIMENTS

We consider semantic associations between elements of the acoustic, visual, and motion modalities; more precisely we define a semantics as an artificial mapping between acoustic words, visual objects, and gestures. An example of such a mapping is given in table II. Each triplet of word, gesture, object forms a semantic concept. The data used to train the MCA-NMF model is composed of sentences, motions, and images; each sentence contains one of the keywords, each motion features one gesture, each image an object. Finally the gesture, the word, and the object from an example belong to the same semantic concept, which implements the cross-situational manifestation of the semantics.

In the testing phase of the following experiments the system is presented ten potential referent, each of a different class among the ten possible semantic classes. Thus a random

Modality			Success rate		
Test	Reference	Comparison	I-div.	Euclidean	Cosine
Sound	Motion	Internal	0.608	0.612	0.646
		Motion	0.552	0.379	0.444
		Sound	0.238	0.126	0.208
Motion	Sound	Internal	0.610	0.704	0.830
		Sound	0.106	0.090	0.186
		Motion	0.676	0.642	0.749

TABLE III

SUCCESS RATES OF RECOGNITION OF THE RIGHT REFERENCE EXAMPLE FROM A TEST EXAMPLE. THE VALUES ARE GIVEN FOR MANY CHOICES OF THE REFERENCE TEST AND COMPARISON MODALITIES AND VARIOUS MEASURES OF SIMILARITY. THE RESULTS ARE OBTAINED BY AVERAGING ON A TEN FOLD CROSS-VALIDATION, BASELINE RANDOM IS IN THAT CASE 0.11.

choice would lead to a rough 10% success rate. When not specified otherwise, we use a default value of $k = 50$ as the number of elements in the dictionary learnt by the NMF algorithm with 50 iterations, although a number of 10 is generally already close to convergence.

A. Motion and spoken utterances

The first experiment compares the various evaluation methods, that is the alternative modalities of comparison and metrics. In this experiment the learner is trained on two modalities: demonstrated motions and spoken utterances. It also validates the capabilities of the MCA-NMF learner to acquire multimodal concepts that appear in a cross-situational way in training data. Table III presents classification success rate for various combinations of comparison modalities and metrics. These results demonstrate that the system is capable of learning aspects of the semantic associations. If the system is trained on a dataset where no semantic association exists between the two modalities (such a dataset can be obtained by choosing a random motion and a random utterance for each demonstration), it generally scores around 0.11². Additionally the result show for example that the sound modality, mainly because of its very high dimension, is not a good choice for the comparison, specially when the comparison is performed with the Euclidean metric. Indeed the choice of the metric to use is highly dependent on the nature of the data in the modality; therefore, using the internal representation is a way to only use one metric. Therefore the following experiments focus on performing the comparison on the internal representations, using cosine similarity as a metrics. The main interest of proceeding that way is that the comparison is the same, regardless of what the test and reference modalities are.

Table IV presents analogous results in the case were symbolic labels are included in the training data. Indeed one could expect that adding symbols to the training data would increase its structure and thus ease the learning of the concepts in comparison to purely non-symbolic data. The labels are included similarly to the experiments from Mangin & Oudeyer (2012), Driesen et al. (2010): the label of each example is transformed to a vector of 10 binary values with zeros

²This is not 0.1 because the distribution of sound examples from the Caregiver dataset is not exactly uniform.

Modality			Success rate		
Test	Reference	Comparison	I-div.	Euclidean	Cosine
Sound	Motion	Internal	0.387	0.699	0.721
		Motion	0.543	0.261	0.424
		Sound	0.136	0.089	0.131
Motion	Sound	Internal	0.573	0.620	0.702
		Sound	0.114	0.090	0.122
		Motion	0.519	0.469	0.552

TABLE IV

THERE IS NO SIGNIFICANT IMPROVEMENT OF THE RECOGNITION RATE WHEN UNAMBIGUOUS SYMBOLS ARE ADDED TO THE TRAINING DATA. THE TABLE REPRESENTS THE SAME SUCCESS RATES AS PREVIOUSLY (SEE TABLE III) BUT WITH A LEARNER THAT OBSERVED SYMBOLIC LABELS REPRESENTING THE SEMANTIC CLASSES DURING TRAINING. THE RESULTS ARE OBTAINED BY AVERAGING ON A TEN FOLD CROSS-VALIDATION, BASELINE RANDOM IS IN THAT CASE 0.11.

everywhere except for a one at the index corresponding to the label. This binary vector is concatenated to the vector representing the example. Thus this setup is equivalent to adding a third modality, which contains unambiguous symbols, in order to improve the learning. The symbols are said to be unambiguous in comparison to utterances that contain several sounds where only some sequences of specific sounds form words, and generally only one word per sentence is relevant. The results from table IV illustrate the fact that the system does not clearly take advantage of this additional information. An interpretation of these results is that the system is already capable of dealing with the ambiguity and is not helped by such additional symbolic information. However the relevance of such comments is limited to the current algorithm and implementation of MCA-NMF.

B. Sound, motion, and image

The experiments in this section compare the performance of learners trained on various combinations of two or three modalities among sound, motion and images. In particular it explores the impact of adding modalities during training that do not occur during testing. In the following we only use the comparison on internal representation. The main interest of proceeding that way is that the comparison is the same, regardless of what the test and reference modalities are. Results are presented together with box plots corresponding to 20 repetitions of the experiment with random label associations, test set, train set, and reference examples.

More precisely several setup are presented, including learning from motion and sound, as well as from image and sound, as previously, but also learning from motion and image, and finally learning from the three modalities at the same time. For each of these choices of learning modalities, several setup are possible for the test phase, specially when the three modalities are present during training: these include testing on the recognition of one modality from another (for example Image \rightarrow Sound) but also from two modalities to another (for example Image, Motion \rightarrow Sound), or conversely one modality to two (as in Motion \rightarrow Sound, Image).

Figure 9 compares various setup of type $A \rightarrow B$, for both for the case where A and B are the only modalities

present in the training and the case where a third modality was also present during training. The results demonstrate that MCA-NMF is capable of learning the semantic concepts event when more than two modalities are present. There is no significant difference between the cases of two and three modalities: the system neither benefit noticeably from the third modality nor does it suffer from the increased dimensionality of the data. However, since the number k of elements in the NMF dictionary is fixed, the results could come from the fact that when the system is trained on three modalities, the dimension of the dictionary becomes insufficient to encode non-meaningful aspects of the three modalities. Therefore fig. 10 present the same experiment for various values of k in order to interpret more precisely the previous result. The comparison confirms the fact that the system mainly behaves similarly with two or three modalities.




Fig. 9. An additional modality during training does not significantly impact result. The box plot represent classification success rates for various experiments where two or three modalities are used for training. Each plot corresponds to the use of a subset of modalities during training: the first three plots use two modalities and the last one use three modalities. Each plot contains boxes representing the average success as well as quantiles and extreme values through cross-validation for various testing setups, using only two modalities. There are only two testing setups when only two modalities are used for training, and six when three modalities are used for training.

Figure 11 present the results on many possible test setups in the case where all modalities are present during training. The results demonstrate that the system is capable of using information contained in more than one modality in the test or reference example. Although the results are slightly better when using more modalities as input (as in Motion, Image \rightarrow Sound in comparison to Motion \rightarrow Sound or Image \rightarrow Sound), the improvement in performance is not really significant in the experiment.

A last experiment leaves the non-symbolic setup considered previously, in order to compare properties of MCA-NMF with results obtained in previous works, regular classification with the symbolic modality. For example Ngiam et al. (2011) present a learner that is trained on multimodal examples of phonemes, either perceived through their acoustic manifestation or through the motions of the lips that pronounce them. In




Fig. 10. With both two (full lines) and three (dashed) modalities during training, the classification success rates are similar and good for high enough value of the number k of elements in the NMF dictionary. The plots demonstrate that the success rate is quite stable above a minimum value of k .

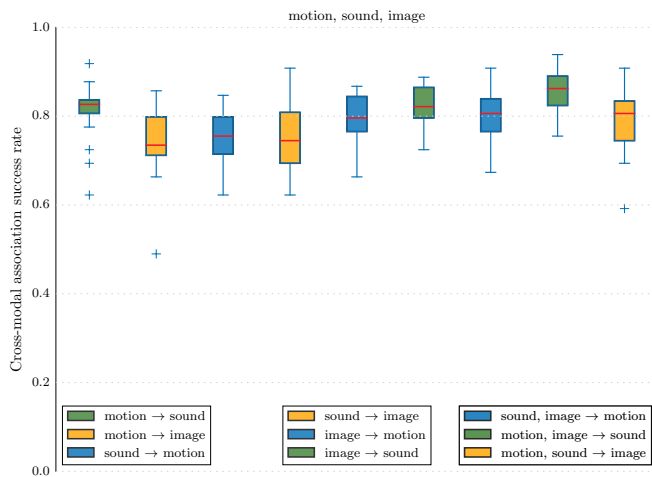


Fig. 11. MCA-NMF is capable of relating information from many modalities to one. There is however no substantial improvement in performance from the use of two modalities as input for the recognition. The figure presents box plots of classification success rates for various experiments where three modalities are used for training. There are boxes representing the average success as well as quantiles and extreme values through cross-validation for various testing setups.

their experiment they show that the learner can benefit from the observation of several modalities and improve its recognition success in comparison to the case where only one modality is observed.

We consider a regular classification setup, similar to the one presented by Mangin & Oudeyer (2012), Driesen et al. (2010). More precisely we introduce a symbolic modality represented by a binary vector as already explained in section IV-A. The system is trained by observing examples both in the symbolic modality and in one or several other modalities. Then results are compared between various testing setups to explore the

ability of the learner to improve its classification performance in the case where several modalities are observed. Such an experiment is actually a classification task with multimodal input unified through *sensor fusion*.

Table V present the results where the sound and motion modalities are combined to a symbolic modality, denoted as L. Interestingly training with the two modalities (sound and motion) does not significantly change the performance of the learner, when tested on sound, motion or both. In that case the benefit of having two non-symbolic modalities is not an increase in performance, but rather that the same learner can use either acoustic perception or motion perception to classify an example.

Training	Testing	Success rates
S + L	S → L	0.916 ± 0.034
M + L	M → L	0.906 ± 0.052
S + M + L	S → L	0.896 ± 0.043
S + M + L	M → L	0.910 ± 0.054
S + M + L	S + M → L	0.917 ± 0.055

TABLE V

SUCCESS RATE FOR THE LABEL RECOGNITION EXPERIMENT. IN THIS EXPERIMENT AN ADDITIONAL MODALITY CONTAINING LABELS, L , IS CONSIDERED. THE RESULTS ARE COMPUTED ON AVERAGE FOR A CROSS-VALIDATION OF THE TRAIN AND TEST SETS; STANDARD DEVIATIONS ARE ALSO GIVEN.

C. Word acquisition and recognition

The previous experiments demonstrate that the artificial learner we present is capable of learning the semantic connection between utterances and the objects or motions they describe. The meaning of the sentences is modelled in our experiment by the presence of a keyword; more precisely the association between sentences and images of objects or motions are based on the presence of keywords in the utterances. However the learner is not aware of the fact that all the meaning of the sentence is actually localized in one word; instead it only exploits cross-situational learning to discover relations between acoustic and, say, visual features. The task solved by the learner actually only involve holistic understanding and classification of the sentences. Therefore it is not completely clear what information the learner actually exploits in the sentence and whether the learner discovers word-like units from the acoustic stream. Indeed the previous experiments only demonstrate that the learner achieves teleological understanding of the sentences; however the question remains to know if it starts to understand compositionally the sentences. We further explore this question in the experiment presented in section IV-C1, that can be seen as extending the one presented by Stouten et al. (2008, 4.C) to the multimodal setup instead of only speech and labels. The system's behavior also display that the semantic classes have been encoded in its perception but not how it is encoded. Hence section IV-C2 takes a further look at this question.

1) *Words' location in sentences*: The experiment presented in this section explores whether the meaning of sentences is localized around the keyword that bears this meaning. In this section the system is trained as in the previous experiments but

we change the evaluation procedure. The system does not have to relate full utterances to visual objects any more. Instead we extract sliding windows of a given length from the utterances from the test set. The sliding windows extraction process actually takes two parameters: the width of the windows and the shift between two windows. In the following we use sliding windows of width 0.5s with a 0.05s shift. Once a sliding window has been extracted, it is represented using similar features than regular sound, that is that the subsequence of sound is extracted and converted to the HAC representation. Then each sound window is compared by the system to several examples in another modality, here images of objects. By evaluating the system on small parts of utterance we can visualize which part of sentences are associated with each object, that is with each semantic underlying concept.

In fig. 12, we present five typical behaviors of MCA-NMF on test sentences. The example have been chosen to illustrates these typical situations but do not correspond to what four random examples would look like. All these examples are taken from the same learner.

The first example (top left), corresponding to the sentence ‘*The book is always nice.*’, features a high similarity between the *book* image and sound windows that intersect the location of the word *book* in the sentence (between 0.8 and 1.0s).

In the second sentence (top right), the correct meaning is guessed by the system but it is not located specifically in the sentence. Because the dataset used for training is rather small in comparison to the variety of word combinations, many non-keywords appears only or most often together with the same keyword. In this case the word ‘*join*’ and the phrase ‘*join us*’ are always observed together with ‘*daddy*’, while the occurrences of ‘*does*’ are more spread among the various keywords. The grammar used to generate the utterances is described quickly by [Altosaar et al. \(2008\)](#) and in more details by [Driesen \(2012, chap. 2\)](#). This situation is similar to the case of natural sentences where contexts highly modify the distributions of words and may explain why such de-localization of the meaning of sentences. In particular this provides an example where the system achieves teleological understanding, which here means guessing the general meaning of a sentence, without compositional understanding.

The third sentence (bottom left) is very specific because it is one of the rare ambiguous sentences of the dataset. Indeed the sentence ‘*Angus takes off his shoe.*’ features both the keyword ‘*Angus*’ and the keyword ‘*shoe*’. Furthermore the learner is capable in this sentence to recognize both words and localize them roughly at the start and end of the utterance.

Finally, the fourth sentence illustrates the many cases where there are false positives for other meanings or (not in the example) no clear meaning is found by the learner in the sentence.

2) *Emergence of localized representations of concepts:* In previous experiments we evaluated the learner on concrete tasks that emphasizes its ability to relate information from one modality to another. While this demonstrates that the internal representation built by MCA-NMF encodes these concepts, it is however not trivial how it encodes information. In this section we explore the possibility that at least some

components of the dictionary matrix are more specialized into some of the semantic classes.

In order to investigate that aspect we quantified the mutual information between the semantic concepts and the coefficients of the internal representations of samples featuring the concepts. For each semantic concept l and sample i we consider the random variables X_l such that $X_l^i = 1$ if and only if the concept l appears in sample i . For each dimension j of the internal representation and each sample i we define the random variable $Y_j = h_j^i$. We then assume that $(X_l^i)_i$ are independent and identically distributed, as well as the $(Y_j^i)_i$. In the following we quantify the dependency between these two variables by looking at the mutual information between them. In information theory, the mutual information I is an information theoretic measure defined for two random variables X and Y as “the relative entropy [or Kullback-Leibler divergence] between the joint distribution $[p(x, y)]$ and the product distribution $p(x)p(y)$ ” by [Cover & Thomas \(1991\)](#).

$$I(X; Y) = D_{KL}(p(x, y) || p(x)p(y))$$

The X_l variables takes binary values but the Y_j are continuous. Therefore we use a quantization of each coefficients of h into 10 discrete values in order to estimate the probability distributions $p(X_l)$, $p(Y_j)$, and $p(X_l, Y_j)$ by using the samples for $1 \leq i \leq N$. Then we compute the mutual information between the discrete, approximated, probability distributions. From this process we obtain a value $I(X_l; Y_j)$ for each pair (l, j) that quantifies how much information the coefficient j captures from the concept l .




Fig. 13. Some components are more associated with some semantic labels. The figure represents the mutual information between (vertically) semantic classes (that are not observed by the learner) and (horizontally) each internal coefficient used by the learner to represent pairs of motion demonstration and acoustic descriptions from the training set. A value of $k = 15$ was used in this experiment.

Figure 13 represents, for each semantic class and each coefficient of the internal representation, the mutual information between the belonging of examples to that class and




Fig. 12. These examples illustrates various distributions of meaning in utterances, in the learner’s point of view. More precisely the horizontal axis represents time: the dots correspond to the mean time of sound windows, which width are materialized by an horizontal grey bar. Thus each plot represents the similarity between the small sound windows and various pictures which semantic class correspond to colors. The four sentences were chosen because they illustrate typical situations. Each record of utterance starts and ends with approximately half a second of silence. Importantly the similarities are taken between sound and images, that is, when a chunk of sound is similar to ‘mummy’, it actually means that the learner’s associates it strongly with an image of the object representing the *mummy* concept. The location of key words, in grey areas, have been indicated by manually detecting their boundaries in the utterance.

the value of a given coefficient of the internal representations of these examples. To emphasize the specialisation of some internal coefficients we re-ordered internal coefficients so that classes and coefficients that have high mutual information are on the diagonal. More precisely, the best alignment was computed by a Kuhn-Munkres algorithm (Munkres, 1957) and we plotted first the coefficients that are highly associated to one class and then the one that are less meaningful. The figure shows that some coefficients are highly specialized in one label. However, it does not display a perfect one to one relationship between labels and coefficients: the information about other labels is spread over several coefficients, and some information is also not clearly localized. For comparison with these results, random labels with uniform probability and a deterministic value of a coefficient knowing the label yields a mutual information of approximately 0.325.

V. DISCUSSION AND PERSPECTIVE

In this paper we present MCA-NMF, a framework for learning multimodal concepts from sensorimotor data, that is based on nonnegative matrix factorization. This algorithm builds a compressed representation of the sensor signals. This representation is used as an internal representation of the multimodal perception and allows to reconstruct any whole

or part of modalities given information on other parts or other modalities.

The experiments demonstrates that this procedure is capable of learning concepts from sub-symbolic input only. It can overcome several forms of ambiguity such as discriminating between relevant and non-relevant patterns to encode (for example relevant and non-relevant speech words). It also discusses the association between patterns from distinct modalities into multimodal concepts.

We then explain how the structure learnt from the observation of full, unsegmented, sentences can be used to localize *a posteriori* the semantics of the concepts into areas of the sentences such as words. Finally the last experiment explores how the system internally represent semantic classes. It shows that constraining the representation to be highly compressed leads to a relative specialisation of encoding units in semantic concepts.

A. Relation to previous work

Several previous works have addressed similar questions to the ones discussed here. As detailed below, the main novelty of our work consists in the facts that 1) we do not rely on symbols for the learning or evaluation phases, 2) the utterances are not segmented into words or phonemes using *a priori* or

expert knowledge, 3) more generally the acoustic modality that contains language is treated symmetrically to the others modalities, 4) the concepts co-organize across modalities, in particular words and their semantics shape each others.

In their seminal work, Roy (1999), Roy & Pentland (2002) introduce a learning architecture called *Cross-channel early lexical learning* (CELL), together with an example implementation, that demonstrates how the problems of learning *linguistic units, semantic categories*, and their relations (in the form of *lexical units*) can be achieved at the same time. CELL involves three stages. It first segments both linguistic information and contextual information, each of which may come from several sensory channels, according to saliency cues such as utterance boundaries or changes in motions. In a second stage, a model of short term memory filters pairs of recurrent co-occurring linguistic and contextual events. Finally a clustering algorithm builds models of linguistic units and semantic categories; they combine clustering of similar language stimuli as well as contextual stimuli and optimize the mutual information between language and context. The pairs of linguistic units and semantic categories with the highest mutual information are kept as lexical units, that correspond to concepts. Thus, in comparison to MCA-NMF, the CELL system first segments sentences in word-like units and thus words and semantics concepts are first formed independently and then matched, without modelling co-organization of words and their meanings.

Yu & Ballard (2004) have presented work addressing a similar problem but focusing more precisely on user-centric and multimodal information. They present a learning architecture that is capable of forming semantic models of both actions and observed objects by using unsupervised learning techniques. First, models of actions are formed by fitting a mixture of hidden Markov models on the observations and models of objects result from an agglomerative clustering algorithm. The models of objects and actions define concepts and together form the contextual information. Then, this contextual information is used to extract word-like units related to these concepts from phoneme transcriptions of the recorded utterances. More precisely longest common phonetic sequences are extracted from all utterances related to the same object or action. Then an alignment techniques, that comes from the field of automatic translation, is used to form the lexical units composed of words and concepts. This approach differs from MCA-NMF since concepts are formed beforehand and the model thus do not take into account the shaping of semantics by the linguistic modality. Furthermore their model relies on mechanisms specific for the linguistic modality.

In Iwahashi (2003) the studied language is related to an (object, action, position) semantics which appears to be closely related to the language grammar. More precisely a lexicon is built from data: the lexicon actually represents a mixture of word and meaning pairs, where meanings can either be objects or actions. Specific probability models are implemented to represent the acoustic modality as well as the modality of visual objects and the one of visual actions. The number of elements in the lexicon is automatically chosen in order to maximize the mutual information between the speech and

contextual modalities. Because the acoustic pattern and object recognitions abilities are separately acquired in a preliminary stage, this model does not take into account mutual shaping of linguistic and semantics patterns as MCA-NMF. It however features the acquisition of a grammar of the language, learnt by identifying in which order the linguistic elements corresponding to the eventual object, action, and landmark appear.

Sugita & Tani (2005) introduce a recurrent neural network architecture that learns to relate a basic language to corresponding behaviors of a robot. The system is capable of both understanding the words composing the language, that in their experiment are represented by symbols, and their composition, that is to say the syntactic structure of the language. As already mentioned in introduction, another aspect of learning action related to language is explored by Tuci et al. (2011) who demonstrates that learning a compositional structure shared between action and language can allow better generalization of the motor knowledge. Furthermore Massera et al. (2010) have demonstrated that providing linguistic instructions can facilitate the acquisition of a behavioral skill, in comparison to pure motor learning. Although these experiments are limited to symbolic language, they are good illustrations of the implication of learning multimodal actions and grammars.

In our previous work (Mangin & Oudeyer, 2012), we use the MCA-NMF framework in a multi-label classification experiment. More precisely, instead of observing several continuous modalities, the learner observes human demonstrations of motions and a symbolic linguistic modality. However each demonstration mixes several motions together, forming a complex choreography. Each motion demonstration is thus described in the linguistic modality by several symbols; in other word there is a joint grammar between the motion and the symbolic modalities. In that experiment, the MCA-NMF system learns the relation between the symbols and the gestures. This relation is initially ambiguous since the gestures are always observed mixed together as well as the symbols that describe them. The experiment then demonstrates that the learner is able to provide correct linguistic descriptions of combinations of gestures that were not observed during training. This result is thus an example of joint grammar acquisition in two modalities, similar to the ones from Sugita & Tani (2005), Tuci et al. (2011), Massera et al. (2010).

Ten Bosch et al. (2008), Driesen et al. (2009) have presented a similar experiment where the contextual modality is the one that is symbolic and the linguistic one is continuous. Similarly Lienhart et al. (2009), Akata et al. (2011), BenAbdallah et al. (2010), Srivastava & Salakhutdinov (2012) use the NMF, probabilistic latent semantic association (PLSA), or deep Boltzmann machine algorithms to learn from a continuous and a symbolic modality. Driesen et al. (2010) have also used the NMF algorithm to learn from two continuous modalities. Their model is very close to MCA-NMF and directly inspired it; however they only evaluate it on the reconstruction of a symbolic modality. An interesting aspect of these approaches is that they all use feature learning algorithms, that are some kind of unsupervised algorithms, instead of relying on explicit models of the lexical units and their relations to language and context.

Another example of the use of feature learning techniques is given by [Ngiam et al. \(2011\)](#) who also present an experiment based on a similar multimodal setup. They introduce an architecture based on sparse restricted Boltzmann machines that learns from two continuous modalities: one is acoustic and the other corresponds to the observation of the speaker's lips. They demonstrate how in certain conditions the algorithm reproduces the McGurk effect. Their algorithm actually learns a new representation of the input in an unsupervised setup and is then evaluated combined with a standard supervised classifier trained on top of this representation. Their work can also be described as a sensor-level multimodal fusion: several modalities are used to build a common representation that is later used to solve a classification problem. Actually multimodal fusion has already been used to improve supervised classification: [Potamianos et al. \(2003\)](#) discuss the use of both sensor-level fusion and decision-level fusion for speech recognition. [Saenko & Darrell \(2007\)](#) also implement decision-level fusion and demonstrate that it improves the recognition of objects. Although closely related to the one of [Driesen et al. \(2010\)](#) and similar to the one from [Ngiam et al. \(2011\)](#) the setting of the last experiments we presented differs. Indeed, we do not evaluate the learning through a standard classification task: instead of testing the reconstruction of symbolic labels, the system is tested on a behavior based classification task, as encountered by children. We show with MCA-NMF that fitting an explicit representation of a lexicon is not necessary to produce behaviors that are considered on children as evidence of the mastering of lexicons understanding. That aspect is an important novelty in comparison of the aforementioned previous work.

The interactions mechanism between the learning agent implemented by MCA-NMF and its caregiver, who provides the demonstration, actually shares many similarities with the one from the *Talking Heads* experiment as described by [Steels \(1999\)](#), [Steels & Kaplan \(2002\)](#). More precisely the agent we present in this chapter plays the role of the *hearer* from the talking heads, while the caregiver takes the role of the *speaker*. There are however important differences between our setup and the one from [Steels \(1999\)](#). First, there is no turn in the role taken in our experiment: the learner only plays the hearer and the caregiver only plays the speaker. Importantly this means that the language is taught to the learner by the caregiver, instead of evolving and emerging from their interaction. Also, in our experiments, the naming game, that consists for the hearer in guessing which object the speaker is talking about, is only played during the evaluation stage. During the training the learner passively observes the caregiver teaching and does not receive any other feedback. However, this makes MCA-NMF an alternate implementation of a naming game hearer, that does not rely on the success of the naming to update its knowledge, but rather uses cross-modal information to do so. This means here that the cross-modal heuristics may be an alternative to feedback on the game success.

B. Interpretations and limitations

In this paper we demonstrate, with MCA-NMF, the possibility of multimodal learning from a flow of non-symbolic sensory input by using variations of the non-negative matrix factorization. The idea we present are however not restricted to this specific algorithm and its implementation. For examples algorithms based on deep belief networks have been shown to perform well on problems similar to the one addressed here, though only on symbolic language, by [Srivastava & Salakhutdinov \(2012\)](#). It is thus an interesting subject for future work to investigate the behavior of these algorithms or other feature learning algorithms on setups similar to this one, as started by [Droniou et al. \(2014\)](#).

In particular we propose a focus shift in the notion of perceptual decomposition. Regarding the structure of speech, we do not consider a preliminary segmentation of utterances into words in order to learn their meaning. Similarly, in our previous work on learning composite choreographies ([Mangin & Oudeyer, 2012](#)), the MCA-NMF framework does not target the representation of individual gestures. Instead, it builds a representation of full body motions and simultaneously learns its correspondence to a linguistic modality. This approach achieves the recognition of composed gestures thanks to a key property: the internal representation and the mapping between the internal representation and the modalities are compatible with the target structure. Furthermore the strong structure in the symbolic modality that describes the motions is sufficient to shape the learnt representation of motions so that it encodes the same structure. Here we demonstrate similarly that the structure in a non-symbolic modality containing images of objects is sufficient to shape the perception of speech in order to recognizing individual words that correspond to the objects. Thus, in both the speech and motion cases the explicit representation of combination is not set as a feature of the algorithm. It is instead induced by the data.

This idea is also largely in line with the position that [Wrede et al. \(2012\)](#) expose: accuracy of decomposition is often a complicated objective to target at first. In contrast, simpler goals, such as understanding the broad meaning of a sentence, can first be achieved that later enable to improve on the decomposition capabilities. Results from section IV-C1 are actually in line with this perspective. This is also coherent with the fact that the apparent evidence of decomposition in perception is often illusory as revealed by many perceptual experiments (see for example [Brent, 1999](#), [Blake, 1989](#), [Leopold & Logothetis, 1999](#)).

In terms of learning the grammars of motion and language, a top down approach would consist in a strong a priori on the structure to be learned that the algorithms must encode. The question would be: 'What is the best algorithm to learn or encode the grammar from the data?' Instead we propose to explore another question: 'Which learning mechanism can lead to the emergence of which grammatical structure?' In MCA-NMF, the operations of word or image juxtaposition or motion combination are translated into nonnegative linear combinations of words or images representations. The results from [Mangin & Oudeyer \(2012\)](#) demonstrate that this property

enables MCA-NMF to learn the grammar of simultaneous motion. However, this form of grammar is obviously limited; thus future work includes the investigation of the ability of this mechanism to represent and let emerge other types of motion or word grammars. From this investigation follows the question of the data that can lead to the learning of new grammars. In particular which grammatical structure are common to several modalities such as speech, vision, or motion perception. Similar future work must also investigate the same question with other combination mechanisms and learning algorithms.

In this paper we show that MCA-NMF is able to acquire word knowledge from cross-situational information. The experiment we present are however limited to data that has been artificially forged to encode cross-situational information. It is thus an essential extension to experiment this learner on realistic data. This constitutes a way to assert the strength of cross-situational information in infant-like perception, and hence to investigate that hypothesis. Another important extension is to develop experiments on realistic data that use other similar computational models to better delimit what cross-situational information can and cannot explain in word learning. Indeed many classes of words and meanings may be impossible to learn solely from that information. Finally, the exploration of other word learning heuristics is also a major subject for future work. Indeed, other rules such as mutual exclusivity or whole object assumption have been identified to play an important role in the process of learning words [Markman \(1990\)](#). Their study raises new challenges such as building computational models of learners sensible to this information and experiment their capabilities.

The MCA-NMF model we present in this paper also suffers important limitations. First in its current state it does not consider the developmental dimensions of concept acquisition. For example it is based on a batch learning algorithm; although incremental versions can be implemented the present article does not cover this aspect. In particular, we illustrated various behaviors of the learners in the experiments with sliding windows and interpret them as teleological understanding potentially enabling compositional understanding. Further analysis of the occurrences of such behaviors may benefit from being performed in a dynamical setting, in order to focus on the evolution of the learner through its development. There no limitation in principle to implement the MCA-NMF learner incrementally, which would enable new experimental possibilities. As an example it would enable to study the evolution and refinement of a concept. In our approach to learning from speech, we do not address the question of attention. More precisely we implicitly consider two basic models of attention: full utterances and sliding windows. Such restriction on time scales may prevent to account for several aspects of human speech perception and further investigation in this direction also opens perspectives for future work. Finally the sensorimotor aspect of perception and other aspects of the interplay between action and perception are out of the scope of this article but of great interest ([Cangelosi et al., 2010](#)). In particular attention mechanisms play an important role in perception, at different time scales, to shape the

data perceived by the learner. While such aspects are not explored in this study the proposed MCA-NMF framework or its implementation may further be used as a building block to study these questions.

ACKNOWLEDGEMENT

We would like to thank members of the former ACORNS European project for providing access to the Caregiver databases. We would like to express our gratitude to Louis ten Bosch for the many inputs and insights he provided at the early stage of this project. This research was partially supported by the ERC Explorers 240007 funding.

Drawings from fig. 1 were adapted from flaticon.com authors icons8 and freepik.

REFERENCES

- Akata, Z., Thureau, C., & Bauckhage, C. (2011). Non-negative matrix factorization in multimodality data for segmentation and label prediction. In *Computer Vision Winter Workshop*, number 16 Mitterberg, Autriche.
- Akhtar, N. & Montague, L. (1999). Early lexical acquisition: The role of cross-situational learning. *First Language*, 19(57), 347–358.
- Altosaar, T., ten Bosch, L., Aimetti, G., Koniaris, C., Demuyneck, K., van den Heuvel, H., Proc, S., Box, P. O., & Tkk, F. (2008). A speech corpus for modeling language acquisition: Caregiver. In *Language Resources and Evaluation - LREC* (pp. 1062–1068).
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and brain sciences*, 22(04), 577–660.
- Bay, H., Tuytelaars, T., & Van Gool, L. (2006). Surf: Speeded up robust features. In *Computer Vision—ECCV 2006* (pp. 404–417). Springer.
- Belpaeme, T. & Morse, A. (2012). Word and category learning in a continuous semantic domain: Comparing cross-situational and interactive learning. *Advances in Complex Systems*, 15(03n04).
- BenAbdallah, J., Caicedo, J. C., Gonzalez, F. A., & Nasraoui, O. (2010). Multimodal image annotation using non-negative matrix factorization. In *International Conference on Web Intelligence and Intelligent Agent Technology* (pp. 128–135).: IEEE / WIC / ACM.
- Blake, R. (1989). A neural theory of binocular rivalry. *Psychological review*, 96(1), 145–167.
- Boves, L., ten Bosch, L., & Moore, R. (2007). Acorns—towards computational modeling of communication and recognition skills. In *Cognitive Informatics, 6th IEEE International Conference on* (pp. 349–356).: IEEE.
- Brent, M. R. (1999). Speech segmentation and word discovery: A computational perspective.
- Calinon, S. & Billard, A. G. (2009). Statistical learning by imitation of competing constraints in joint space and task space. *Advanced Robotics*, 23, 2059–2076.
- Cangelosi, A. (2010). Grounding language in action and perception: from cognitive agents to humanoid robots. *Physics of life reviews*, 7(2), 139–151.

- Cangelosi, A., Metta, G., Sagerer, G., Nolfi, S., Nehaniv, C., Fischer, K., Tani, J., Belpaeme, T., Sandini, G., Nori, F., Fadiga, L., Wrede, B., Rohlfing, K., Tuci, E., Dautenhahn, K., Saunders, J., & Zeschel, A. (2010). Integration of action and language knowledge: A roadmap for developmental robotics. *IEEE Transactions on Autonomous Mental Development*, 2(3), 167–195.
- Cederborg, T. & Oudeyer, P.-Y. (2013). From language to motor gavage: Unified imitation learning of multiple linguistic and non-linguistic sensorimotor skills. *Transactions on Autonomous Mental Development (TAMD)*, 5.
- Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the acoustical society of America*, 25, 975.
- Cover, T. M. & Thomas, J. A. (1991). *Elements of information theory*. John Wiley & Sons.
- Driesen, J. (2012). *Discovering words in speech using matrix factorization*. PhD thesis, KU Leuven, ESAT.
- Driesen, J., Gemmeke, J. F., & Van Hamme, H. (2012). Data-driven speech representations for nmf-based word learning. In *SAPA-SCALE Portland, USA*.
- Driesen, J., ten Bosch, L., & Van Hamme, H. (2009). Adaptive non-negative matrix factorization in a computational model of language acquisition. In *Interspeech* (pp. 1–4).
- Driesen, J., Van Hamme, H., & Kleijn, B. W. (2010). Learning from images and speech with non-negative matrix factorization enhanced by input space scaling. In *IEEE Spoken Language Technology Workshop (SLT)* (pp. 1–6). Berkeley, California, USA: IEEE.
- Droniou, A., Ivaldi, S., & Sigaud, O. (2014). Deep unsupervised network for multimodal perception, representation and classification. *Robotics and Autonomous Systems*.
- Filliat, D. (2008). Interactive learning of visual topological navigation. In *International Conference on Intelligent Robots and Systems (IROS 2008)* (pp. 248–254).: IEEE.
- Glenberg, A. M. & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic bulletin & review*, 9(3), 558–65.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1), 335–346.
- Iwahashi, N. (2003). Language acquisition through a human-robot interface by combining speech, visual, and behavioral information. *Information Sciences*, 156(1), 109–121.
- Joachims, T. (1997). *Text categorization with support vector machines: Learning with many relevant features*. Technical Report 23, Universitat Dortmund, LS VIII-Reportitat Do.
- Konczak, J. (2005). On the notion of motor primitives in humans and robots. In *International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 123 (pp. 47–53).: Lund University Cognitive Studies.
- Kulic, D., Imagawa, H., & Nakamura, Y. (2009). Online acquisition and visualization of motion primitives for humanoid robots. *Symposium on Robot and Human*, (pp. 1210–1215).
- Landau, B., Smith, L., & Jones, S. (1998). Object perception and object naming in early development. *Trends in cognitive sciences*, 2(1), 19–24.
- Lee, D. D. & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755), 788–91.
- Leopold, D. A. & Logothetis, N. K. (1999). Multistable phenomena: changing views in perception. *Trends in cognitive sciences*, 3(7), 254–264.
- Lienhart, R., Romberg, S., & Horster, E. (2009). Multilayer pls for multimodal image retrieval. In *International Conference on Image and Video Retrieval - CIVR*, number April (pp. 1). New York, New York, USA: ACM Press.
- Lupyan, G., Rakison, D. H., & McClelland, J. L. (2007). Language is not just for talking: redundant labels facilitate learning of novel categories. *Psychological Science*, 18(12), 1077–1083.
- Lyubova, N. & Filliat, D. (2012). Developmental approach for interactive object discovery. In *Proceedings of the 2012 IJCNN International Joint Conference on Neural Networks*.
- Mangin, O. (2014). *The Emergence of Multimodal Concepts: From Perceptual Motion Primitives to Grounded Acoustic Words*. PhD thesis, Universite de Bordeaux.
- Mangin, O., Filliat, D., & Oudeyer, P.-Y. (2010). A bag-of-features framework for incremental learning of speech invariants in unsegmented audio streams. In B. Johansson, E. Sahin, & C. Balkenius (Eds.), *Tenth International Conference on Epigenetic Robotics* (pp. 73–80). orenaas Slott, Sweden.
- Mangin, O. & Oudeyer, P.-Y. (2012). Learning to recognize parallel combinations of human motion primitives with linguistic descriptions using non-negative matrix factorization. In *International Conference on Intelligent Robots and Systems (IROS 2012)* Vilamoura, Algarve (Portugal): IEEE/RSJ.
- Mangin, O. & Oudeyer, P.-Y. (2013). Learning semantic components from subsymbolic multimodal perception. In *the Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)*, number 3.
- Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14, 57–77.
- Massera, G., Tuci, E., Ferrauto, T., & Nolfi, S. (2010). The facilitatory role of linguistic instructions on developing manipulation skills. *IEEE Computational Intelligence Magazine*, 5(3), 33–42.
- McGurk, H. & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264(5588), 746–748.
- Micusik, B. & Koecka, J. (2009). Semantic segmentation of street scenes by superpixel co-occurrence and 3D geometry. In *IEEE Workshop on Video-Oriented Object and Event Classification (VOEC), held jointly with International Conf. on Computer Vision (ICCV) Japan*: IEEE.
- Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial & Applied Mathematics*, 5(1), 32–38.
- Mussa-Ivaldi, F. A. & Bizzi, E. (2000). Motor learning through the combination of primitives. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 355(1404), 1755–69.
- Ngiam, J., Khosla, A., Kim, M., Nam, J., & Ng, A. Y. (2011).

- Multimodal deep learning. In *International Conference on Machine Learning*, number 28 Bellevue, Washington, USA.
- Paatero, P. & Tapper, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2), 111–126.
- Potamianos, G., Neti, C., Gravier, G., Garg, A., & Senior, A. W. (2003). Recent advances in the automatic recognition of audiovisual speech. *Proceedings of the IEEE*, 91(9), 1306–1326.
- Quine, W. V. O. (1960). *Word and object*. MIT press.
- Roy, D. (1999). *Learning from Sights and Sounds: A Computational Model*. PhD thesis, Massachusetts Institute of Technology.
- Roy, D. K. & Pentland, A. P. (2002). Learning words from sights and sounds: A computational model. *Cognitive science*, 26(1), 113–146.
- Saenko, K. & Darrell, T. (2007). Object category recognition using probabilistic fusion of speech and image classifiers. In *Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms - MLMI*, number 4 Brno, Czech Republic.
- Schwartz, J.-L. (2010). A reanalysis of mcgurk data suggests that audiovisual fusion in speech perception is subject-dependent. *The Journal of the Acoustical Society of America*, 127, 1584–1594.
- Schwartz, J.-L., Berthommier, F., & Savariaux, C. (2004). Seeing to hear better: evidence for early audio-visual interactions in speech identification. *Cognition*, 93(2), B69–B78.
- Schwartz, J.-L., Grimault, N., Hupé, J.-M., Moore, B. C., & Pressnitzer, D. (2012). Multistability in perception: binding sensory modalities, an overview. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1591), 896–905.
- Sivic, J. & Zisserman, A. (2008). Efficient visual search for objects in videos. *Proceedings of the IEEE*, 96(4), 548–566.
- Smith, L. & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106(3), 1558–1568.
- Sodoyer, D., Girin, L., Jutten, C., & Schwartz, J.-L. (2004). Developing an audio-visual speech source separation algorithm. *Speech Communication*, 44(1), 113–125.
- Srivastava, N. & Salakhutdinov, R. R. (2012). Multimodal learning with deep boltzmann machines. In F. Pereira, C. Burges, L. Bottou, & K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25* (pp. 2222–2230).: Curran Associates, Inc.
- Steels, L. (1999). *The Talking Heads Experiment*, volume 1. Laboratorium.
- Steels, L. (2008). The symbol grounding problem has been solved. so what's next? In M. de Vega, A. M. Glenberg, & A. C. Graesser (Eds.), *Symbols and Embodiment: Debates on Meaning and Cognition* chapter 12. Oxford: Oxford University Press.
- Steels, L. & Kaplan, F. (2002). Bootstrapping grounded word semantics. In T. Briscoe (Ed.), *Linguistic Evolution through Language Acquisition: Formal and Computational Models* chapter 3. Cambridge University Press.
- Stouten, V., Demuynck, K., & Van Hamme, H. (2008). Discovering phone patterns in spoken utterances by non-negative matrix factorization. *Signal Processing Letters, IEEE*, 15, 131–134.
- Sugita, Y. & Tani, J. (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior*, 13(1), 33.
- ten Bosch, L. F., Van Hamme, H., & Boves, L. W. (2008). Unsupervised detection of words questioning the relevance of segmentation. In *Speech Analysis and Processing for Knowledge Discovery*, ITRW ISCA: Bonn, Germany : ISCA.
- Tresch, M. C. & Jarc, A. (2009). The case for and against muscle synergies. *Current opinion in neurobiology*, 19(6), 601–7.
- Tuci, E., Ferrauto, T., Zeschel, A., Massera, G., & Nolfi, S. (2011). An experiment on behaviour generalisation and the emergence of linguistic compositionality in evolving robots. *IEEE Transactions on Autonomous Mental Development*, 3(2), 1–14.
- Van Hamme, H. (2008). Hac-models: a novel approach to continuous speech recognition. In *Interspeech ISCA* (pp. 2554–2557).
- Waxman Sandra, R. & Markow Dana, B. (1995). Words as invitations to form categories: Evidence from 12-to 13-month-old infants. *Cognitive psychology*, 29(3), 257–302.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic bulletin & review*, 9(4), 625–636.
- Wrede, B., Rohlfing, K., Steil, J., Wrede, S., Oudeyer, P.-Y., & Tani, J. (2012). Towards robots with teleological action and language understanding. In E. Ugur, Y. Nagai, E. Oztop, & M. Asada (Eds.), *Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?* Osaka, Japon.
- Yu, C. & Ballard, D. H. (2004). A multimodal learning interface for grounding spoken language in sensory perceptions. *Transactions on Applied Perception*, (1), 57–80.
- Yurovsky, D., Yu, C., & Smith, L. B. (2012). Statistical speech segmentation and word learning in parallel: scaffolding from child-directed speech. *Frontiers in psychology*, 3.