Perceptual Abstraction and Attention - D4.1

Deliverable D4.1 of the EU-funded Integrated Project
"IM-CLeVeR – Intrinsically Motivated Cumulative Learning Versatile Robots",
contract n. FP7-ICT-IP-231722

FIAS: Pramod Chandrashekhariah, Christos Dimitrakakis, Luca Lonini, Constantin Rothkopf, Gabriele Spina, Jochen Triesch
CNR-ISTC: Dimitri Ognibene, Gianluca Baldassarre
USFD: Alec Cope, Jon Chambers, Kevin Gurney
UU: Y. Gatsoulis, C.Burbridge, L.Riano, S. Ratnasingam, J. Condell, T. M. McGinnity
AU: Patricia Shaw, James Law, Mark Lee
IDSIA-SUPSI: Michael Frank, Alexander Förster, Juergen Schmidhuber, Varun Kompella, Matt Luciw, Dan Ciresan, Ueli Meier, Jonathan Masci, Davide Migliore, Leo Pape, Mark Ring, Faustino Gomez, Vincent Graziano

Lead Contractor for this Deliverable: FIAS
Other contributing Beneficiaries: CNR-ISTC, USFD, UU, AU, IDSIA-SUPSI
Deliverable due: 30 April 2011
Deliverable submission: 30 April 2011

Funding Institution: European Commission (European Union)
Project Officer: Cécile Huet
Funded under: Seventh Framework Programme
Work Programme Theme 3: ICT – Information and Communication Technology; Call identifier: FP7-ICT-2007-3; Challenge 2: Cognitive Systems, Interaction, and Robotics; Objective: ICT-2007.2.2 Cognitive System, Interaction, Robotics
Start: 01/01/2009 Termination: 30/04/2013

Cite this document as:

Chandrashekhariah, P., Dimitrakakis, C., Lonini, L., Rothkopf, C., Spina, G., Triesch, J., Ognibene, D., Baldassarre, G., Cope, A., Chambers, J., Gurney, K., Gatsoulis, Y., Burbridge, C., Riano, L., Ratnasingam, S., Condell, J., McGinnity, T.M., Shaw P., Law., J., Lee, M., Frank, M., Förster, A., Schmidhuber, J., Kompella, J., Luciw, M., Ciresan, D., Meier, U., Masci, J., Migliore, D., Pape, L., Ring, M., Gomez, F., Graziano V. (2011). Perceptual Abstraction and Attention, Deliverable 4.1 of the EU-funded Integrated Project "IM-CLeVeR – Intrinsically Motivated Cumulative Learning Versatile Robots". contract n. FP7-ICT-IP-231722.

INTRODUCTION

This is a report on the preliminary achievements of WP4 on abstraction for cumulative learning, in particular directed to: (1) producing algorithms to develop abstraction features under top-down action influence; (2) algorithms for supporting detection of change in motion pictures; (3) developing attention and vergence control on the basis of locally computed rewards; (4) searching abstract representations suitable for the LCAS framework; (5) developing predictors based on information theory to support novelty detection. The report is organized around these 5 tasks that are part of WP4. Below we provide a synthetic description of the work done for each task by the partners. A detailed description for each work is provided in the attached papers/documents. The relevant papers/documents are highlighted in bold within the references.

Task 4.1 (**FIAS**): Visual pre-processing based on feature abstraction

The work for this task is composed of 2 main contributions that are listed in the following:

Directed exploration and motivations

Link and relevance for the project of the presented work

When an agent acts within a partially observable environment, it is crucial that it can direct its attention to environmental features that will give the maximum information relevant to the task (Taylor and Rogers, 2002; Seth, 2000; Corchs and Deco, 2000; van de P. Laar et al., 1997). This is generally a hard decision theoretic problem (Ross et al., 2008; DeGroot, 1970). One special version of this problem is visual (Corchs and Deco, 2000) attention, including proper movement and vergence of the eyes. Related to this, is the question of whether it is possible to learn heuristic exploration strategies for visual attention tasks and whether such strategies can be generalized to more general exploration problems. The final open issue is how to most effectively take the interaction between tasks and attention into account.

Specific topic and problems tackled by the work

Our main goal is to develop simple formal models of attention as near-optimal exploration heuristics, with predictions that take into account the actions. Ideally this model should be able to scale up relatively easily to larger problems. Our second aim is to develop architectures for task-specific vergence control by integrating work done.

Methods and key results

Our aim is to develop formal models which have low computational and sample complexity. These models shall perform inference grounded in statistical decision theory and information theory. In order to achieve practical systems, approximate models of inference and decision-making will be investigated. We have used context models (Dimitrakakis, 2011, 2010c) for the predictive modeling of features. This is a hierarchical distributed approach, which also enables the implementation of distributed value functions for actions. For planning (rather than purely reactive behavior) we are using approximate decision-theoretic methods. In addition, we are addressing the question of whether formal settings exist under which curiosity heuristics arise naturally as near-optimal solutions (Rothkopf and Dimitrakakis, 2011).

We currently have results on hierarchical modelling (Dimitrakakis, 2011, 2010c,b), where it is shown that simple, closed-form, hierarchical distributed inference can be used for prediction and action in a number of standard partially observable problems.
In addition, we have results on approximate planning under uncertainty (Dimitrakakis and Lagoudakis, 2008; Dimitrakakis, 2010a), where we derive nearly optimal planning algorithms that use a relatively small amount of computation.
Finally, we have some experimental and theoretical results linking intrinsic motivations and attention to decision theory through an interesting class of abstract problems (Rothkopf and Dimitrakakis, 2011; Dimitrakakis et al., 2011).

Relevance of the results for the project

In future work we would like to scale up these models to visual tasks. The main challenge in doing so is the selection of an appropriate class of statistical image models. This is a necessary step in order to be able to even define exploration heuristics such as the information gain. However, our current practical experience in applied vision systems should greatly facilitate this step.

Selected bibliography:

Silvia Corchs and Gustavo Deco. A neurodynamical model for selective visual attention using oscillators. Neural Networks, 14(8):981–990, 2000.

Morris H. DeGroot. Optimal Statistical Decisions. John Wiley & Sons, 1970.

Christos Dimitrakakis. Complexity of stochastic branch and bound for belief tree search in Bayesian reinforcement learning. In 2nd international conference on agents and artificial intelligence (ICAART 2010), pages 259–264, Valencia, Spain, 2010a. ISNTICC, Springer.

Christos Dimitrakakis. Bayesian variable order Markov models. In Yee Whye Teh and Mike Titterington, editors, Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS), volume 9 of JMLR : W&CP, pages 161–168, Chia Laguna Resort, Sardinia, Italy, 2010b.

Christos Dimitrakakis. Context model inference for large or partially observable MDPs. In ICML workshop on reinforcement learning and search in very large spaces, 2010c.

Christos Dimitrakakis. Context models on sequences of covers. Technical Report arXiv:1005.2263v2, FIAS, 2011.

Christos Dimitrakakis and Michail G. Lagoudakis. Rollout sampling approx- imate policy iteration. Machine Learning, 72(3):157–171, September 2008. doi: 10.1007/s10994-008-5069-3. Presented at ECML'08.

Christos Dimitrakakis, Ronald Ortner, and Constantin A. Rothkopf. Sparse reward processes. Technical report, FIAS, 2011.

van de P. Laar, T. Heskes, and S. Gielen. Task-dependent learning of attention. Neural Networks, 10(6):981–992, 1997.

Stephane Ross, Joelle Pineau, Sebastien Paquet, and Brahim Chaib-draa. On-line planning algorithms for POMDPs. Journal of Artificial Intelligence Re- sesarch, 32:663–704, July 2008.

Seth, Anil K. Evolving action selection and selective attention without actions, attention or selection, 2000.
J.G. Taylor and M. Rogers. A control model of the movement of attention. Neural Networks, 15(3):309–326, 2002

Specificity and transfer in visual perceptual learning

Link and relevance for the project of the presented work

Learning of multiple tasks and generalization to novel contexts is a key feature of human learning and provide the basis for cumulative learning. Understanding the basis of generalization and how this is affected by the number or the order of tasks is a fundamental issue in IM-CLEVER. Results on this issue could provide inputs for modeling of the board experiment (WP6) as well as interpreting the results of the experiments on humans and monkeys (WP3).

Specific topic and problems tackled by the work

Improvement in discriminating orientations is one form of visual perceptual learning, a long-term enhancement of performances on a visual task as a result of training. Orientation discrimination is known to be specific for the trained orientation [1], such that the improvement will not transfer to a new orientation. Only recently, different experimental paradigms are challenging the existing theories [2,3]. Using a novel paradigm, we tested how transfer is affected by 1) the number and 2) the order of presentation of training stimuli. Understanding how learning generalizes across stimuli is of fundamental importance to understand how cumulative learning is accomplished in perceptual learning. Also, common mechanisms between motor and perceptual learning are being investigated. This will in turn produce an input for the development of bio-inspired algorithms to optimize generalization of learning given the stimuli.

Methods and key results

Participants are shown a first bar (reference) on a computer display, followed by a mask, and are asked to adjust the orientation of the second bar to match the orientation of the first (Fig. 4.1.1). Orientation of the second bar relative to the first is drawn from a uniform distribution ranging in (-25,-15) and (15,25) deg. The reference bar can appear at either 15,30,75 or 135 deg. Adjustments are performed by pressing the up/down arrow keys on a standard computer keyboard. Subjects are randomly assigned to 1 of 5 groups: 2 groups are trained on 3 orientations either in a fixed schedule or in a random schedule (297 trials in total); other 2 groups receive training on 1 orientation only, either for 99 or 297 trials. A control group does not receive any training session. During training, feedback on the correct orientation (red bar) is provided at the end of the adjustment. Performance is measured by precision (variable error or standard deviation) of subjects in the reproduction of the reference bar orientation.

Transfer to the untrained orientation occurred after participants were trained on 3 different orientations (15,30 and 75 deg) in the same day. Remarkably, participants trained only on 1 orientation (30 deg) also showed transfer to all the untrained orientations (Figure 5). Furthermore, we observed a further increase of subject precision across the 3 testing sessions, performed in separate days. Finally, we found an effect of training order on transfer for female subjects, such that only by training one orientation per session enhanced improvements in the subsequent testing sessions (Fig.6).

Relevance of the results for the project

These preliminary results suggest that specificity in perceptual learning can be strongly affected by the training paradigm being used as well as by the practice structure. We envisage the possibility of identifying some computational factors for generalization and interference in perceptual learning that could explain some mechanisms for cumulative learning. Specifically the work will provide inputs to Task 4.3 and WP6 (Hierarchical architectures) for the development of bio-inspired models for cumulative learning.

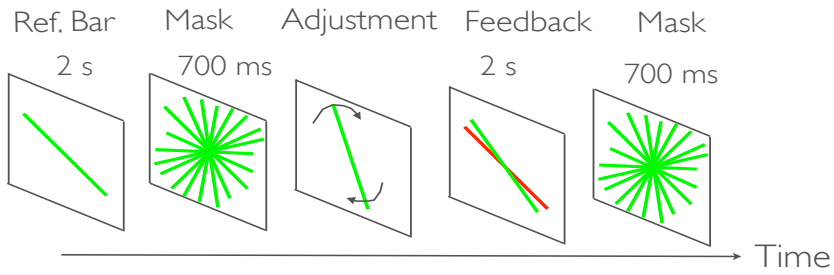

Figure 4.1.1: Experimental design – different phases of a trial

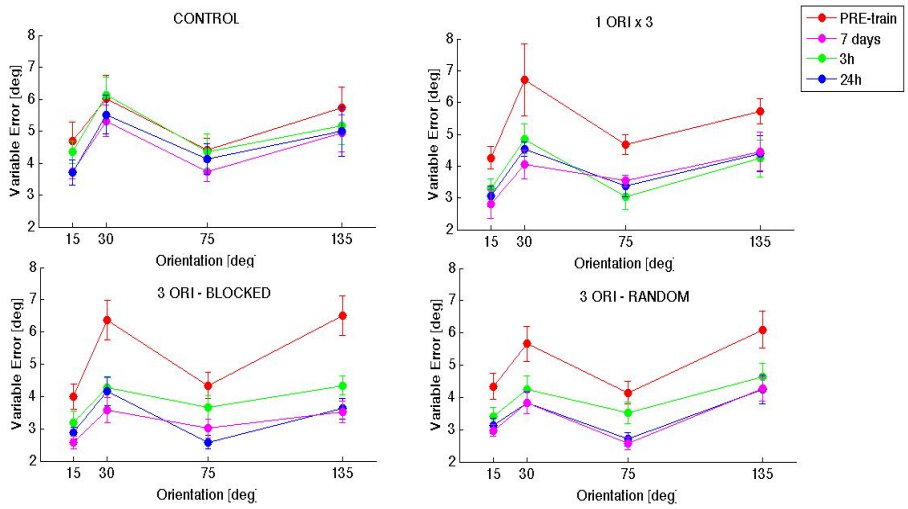

Figure 4.1.2: Precision (variable error) at each orientation across different sessions of testing – All the groups, except for the control group that did not receive any training, improved precision both for trained (15,30 and 75) and untrained orientations (135 deg) (ANOVA repeated measures, $p < .001$)

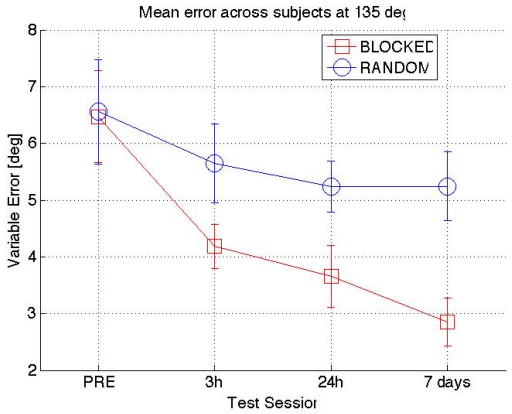

Figure 4.1.3: Training order affects transfer in female subjects across subsequent test phases (3h to 1 week). Blocked presentation of training stimuli favors transfer at 7 days (post hoc analysis, Duncan test, $p < .01$).

Selected bibliography:

[1] Ahissar and Hochstein (1997) Task difficulty and the specificity of perceptual learning. Nature, 387:401-406.

[2] Zhang et al., (2010) Rule-Based Learning Explains Visual Perceptual Learning and Its Specificity and Transfer, Journal of Neuroscience, 30(37):12323--12328.

[3] Green C.S. , Transfer in perceptual learning as extrapolation, Journal of Vision, 10(7): 1121, 2010.

[4] Tschopp-Junker C., Gentaz, E. and Viviani, P., Effect of selective and distributed training on visual identification of orientation. Exp Brain Res, 202:605-611, 2010.

[5] Shea, J.B. and Morgan, R.L. (1979). Contextual interference effects on the acquisition, retention, and transfer of a motor skill. Journal of Experimental Psychology: Human Learning and Memory, 5(2): 179-187.

Task 4.2 (**UU**): Visual pre-processing based on temporal abstraction

Link and relevance to project:
This work package involves strong collaborations between FIAS, UU and IDSIA. In particular we have made the software developed in this task available during visits to Aberystwyth and IDSIA as part of the integration of the CLeVeR-K2 demo.

Specific topic and problems tackled:
This task presents our work on the visual pre-processing methods used primarily in the novelty detection algorithms (Task 5.3) and also in the hierarchical structures (Task 6.5). In particular it discusses segmenting objects from the background, detecting features and recognising objects and finally learning the outcome of actions.

Methods and key results:
Figure 5.3.1 in Deliverable 5 shows a block diagram that describes the overall approach. In brief, a list of regions of interest are first extracted. The feature descriptors for each one of the ROI are passed to the novelty detector module, which decides whether the object is novel and which action should be performed. Corresponding data are obtained and the memory of the system is trained.
Segmenting objects from the background using 2D and 3D image data.
The aim here is to segment the image and provide regions of interest (ROI) corresponding to objects. This is passed to the novelty detector module to determine which objects will gain the robot's attention. All experiments use real robot image data.
- *2D image data:* In initial experiments, only a single camera was used to locate objects. As such, only 2D information about the scene was available and simple 2D segmentation techniques detected contours of objects. As these methods are not very reliable in natural scenes, we used an artificial uniform background to assist segmentation. Reliable detection and extraction was difficult to achieve. The process was time consuming and unsuitable for long unsupervised experiments over a number of days, thus we explored using 3D image data.
- *3D image data:* Initially we used an RGB-D camera (Kinnect) on a robotic system that consists of a 7 degree of freedom Schunk manipulator mounted on a Scitos G5 mobile base. The camera is capable of providing both a 2D image and the distance to every pixel. From this information, a 3D point cloud is constructed. The experimental configuration consisted of a table with objects on it (Figure 4.2.1). The 3D information aids the segmentation of objects from the surrounding scene. Although constrained (one object on a table at a time), this approach to object extraction has proved reliable enough to drive the unsupervised learning of objects in the experiments we conducted, where the robot was moving around the table with the object on top of it (**Burbridge et al., 2011**). Following these experiments, in a different configuration we used a Willow Garage PR2 robot. This exploited the PR2 on-board high resolution Prosilica camera to obtain a *5MP* 2D image and the on-board wide stereo cameras for obtaining a 3D point cloud of the viewable area. Segmentation of objects followed. The system is now capable of extracting ROIs of multiple objects simultaneously.
Detecting features and recognising objects. Following the detection of the ROI corresponding to an object, we detect descriptive image points on the object to allow learning and recognition. For the features describing the object we decided to use Speeded Up Robust Features (SURF) (Bay et al., 2008), a scale and rotation invariant interest point detector and descriptor. It approximates or even outperforms previously proposed schemes with respect to repeatability, distinctiveness, and robustness, yet can be computed and compared much faster (Bay et al., 2008). The SURF features have provided a reliable representation of the objects that has produced effective learning and recognition of objects (**Gatsoulis et al., 2011; Burbridge et al., 2011**). Our main use of the features is within the bag-of-words (BoW) model, which has been used in a number of studies (e.g. Leung

and Malik, 2001; Sivic and Zisserman, 2006). It is a compact and efficient method for feature learning and object classification. BoW here describes and classifies images, by using histograms of the frequencies of "visual words" from a dictionary that exist in the image. The steps involved are (see Figure 4.2.2):

1. Extract a set of feature descriptors, such as SIFT, SURF, etc., from perceived image.
2. Learn visual vocabulary; train unsupervised structure e.g. SOM on extracted features of perceived image.
3. For given vocabulary and image feature descriptors, compute histogram of frequency of visual words.
4. Train classifier (e.g. support vector machine, SOM, etc.) with produced histograms.

BoW has particular limitations (**Gatsoulis et al., 2011**), such as being non-expandable and limited to working offline. Accordingly we have extended the approach to address these issues (**Burbridge et al., 2011**). More details are presented in the report of Task 5.3 and in the paper.

Figure 4.2.1: Experimental setup

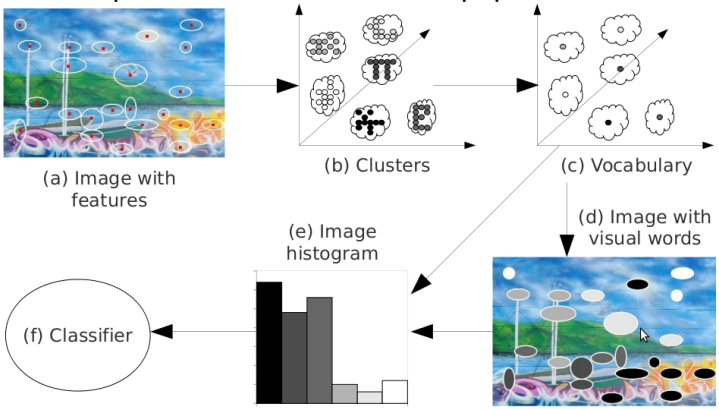

Figure 4.2.2: Bag-of-words model

Learning the outcome of actions

For effective skill learning and acquisition it is useful to know the link between actions and outcomes when performing actions. In its canonical form, it is observing the pose of the objects before and after the action. Actions exploited were push, pull, topple and rotate, as agreed at the Visual Perception and Novelty Detection Workshop at UU in November 2010. We also implemented "pick". We conducted a number of experiments to test the observation of the initial and final pose of objects after performed actions using a PR2 robot:

- The robot is in front of a table with objects on the table at reachable distances.
- The wide stereo cameras build a coarse 3D representation of the table and the objects. RANSAC (Fischler & Bolles, 1981) is used to segment a plane (the table) out of the stereo reconstructed point cloud. The points lying above the plane are clustered so that each cluster represents an object. The final object to manipulate is chosen by the novelty detection algorithm, but for this specific task the robot chooses a random action. Once an object has been chosen, the robot head focuses on it and the narrow stereo cameras create a new, denser, 3D reconstruction of the object (Hsiao et al., 2010). PCA is applied to this point-cloud to find the major and minor axis of the object (the object's z-axis is supposed to be normal to the table, Holz et al., 2011). The position and the main axis of the object are recorded at this stage.
- After an action is performed the robot's head is aimed at the end of the trajectory where the robot expects to find again the object. The segmentation and pose estimation described is repeated and, if the object has not been pushed off the table it is discovered again. The robot obtains an estimate of the displacement effect of the performed action by comparing the current position of the object with the previous one. Observing which axis is the longest allows determination of

whether the object has been toppled or not. If the longest axis has changed after the action, we can assume that the object has toppled. For example, if the object is a bottle, before toppling the longest axis is *z*. If after the action the robot finds that the longest axis is either *x* or *y*, the bottle has toppled over.

- To allow the robot to use this information to build knowledge of action-consequence links, the recognition of actions based on semantic graphs was investigated (Erdal et al., 2010). To test the applicability of the algorithm we used a video stream of a human arm manipulating tabletop objects. In the real experiment a robot arm is used to grasp and manipulate an object (**Ratnasingam & McGinnity, 2011a; 2011b**).
- The first step of the action recognition process is segmenting the video into small clips (each with one action, see Figure 4.2.3). Each clip was segmented into three segments in such a way that the first segment of the clip has frames that do not have the human hand present (frames of the scene before the action is performed) and the second segment has frames in which human hand is present (while the action is being performed) and the third segment of the clip has frames just after the human hand finished the action (hand is not present in the scene). Illuminant invariant skin colour was used as the cue. Once the clip has been segmented into three segments each of these segments were analysed to recognise the action. As the actual action occurs in the second segment, the object that is moved by the hand was identified using illuminant invariant colour features in the second segment. Particularly, the moving object was recognised by calculating the position of the objects between frames. The frames in the last segment were investigated for any movement of objects to identify any unstable states after the action has been performed. Thus recognition occurs of actions based on the object moved and consequential stability of the state.

Figure 4.2.3: Human hand placing the soap on top of a ball. (left) initial position of the objects, (middle) when hand is placing the soap on top of the ball and (right) Instability of the state caused the soap to fall and the ball has moved from its initial position.

Relevance of the results for the project:

Segmenting objects from the background: For the effective operation of the perceptual learning and novelty detector in Task 5.3, it is important to equip the robot with the ability to reliably locate objects and to reliably focus its attention on a single training object at a time (Holz et al 2011).

Detecting features and recognising objects: We experimented with a standard bag-of-words approach and integrated a visual perception learning novelty detector (**Gatsoulis et al., 2011**). This system still only operated offline and was unable to expand its vocabulary. These issues were successfully addressed in further work (**Burbridge et al., 2011**). This work links Tasks 4.2 and 5.3.

Learning the outcome of actions: Based on visual pre-processing implemented, and our novelty detection algorithms (Task 5.3) we enabled the robot to carry out actions on perceived objects and identify the outcome of these actions so that basic affordances of the objects can be associated with particular events.

Selected bibliography

Gatsoulis, Y., E. Kerr, J.V. Condell, N.H. Siddique, and T.M. McGinnity. "A visual novelty detection filter based on bag-of-words and biologically-inspired networks", In AISB'11 Convention. York, 2011.

Burbridge, C., Gatsoulis, Y., McGinnity, M. "Cumulative online unsupervised learning of objects by visual inspection", Submitted to IROS 2011.

Herbert Bay, Andreas Ess, Tinne Tuytelaars, Luc Van Gool, "SURF: Speeded Up Robust Features", Computer Vision and Image Understanding (CVIU), Vol. 110, No. 3, pp. 346--359, 2008.

T. Leung and J. Malik, "Representing and Recognizing the Visual Appearance of Materials using Three-dimensional Textons", International Journal of Computer Vision, vol. 43, pp. 29-44, 2001.

J. Sivic and A. Zisserman, "Video Google: Efficient visual search of videos", Toward Category Level Object Recognition, Springer Berlin, 2006.

Hsiao K, Chitta S, Ciocarlie M, Jones EG. Contact-Reactive Grasping of Objects with Partial Shape Information. In: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ:1228--1235.

Holz, Dirk., Holzer, Stefan., Rusu, Radu Bogdan., and Behnke, Sven. Real-Time Plane Segmentation using RGB-D Cameras. RoboCup Symposium, 2011, (2011).

A. E. Erdal, A. Alexey, W. Florentin, and D. Babette. "Categorizing Object-Action Relations from Semantic Scene Graphs". IEEE International Conference on Robotics and Automation (ICRA2010), Alaska, USA.

S. Ratnasingam, M. McGinnity. "Object Recognition Based on Tactile Form Perception", in IEEE Symposium on Computational Intelligence, Paris, France, April, (2011a).

S. Ratnasingam, M. McGinnity. "A Spiking Neural Network for Tactile form Based Object Recognition" Accepted in IEEE,IJCNN, San Jose, California, July 31 - August 5, (2011b).

Task 4.3 (**FIAS**): Learning of attention and vergence control
Bio-inspired robotic vision system for active visual perception based on intrinsic motivation

Link and relevance for the project of the presented work

We used iCub robot head as robotic platform to implement a general-purpose vision system. This will allow an easy integration with Clever-B (*Task 7.4*) and Clever-K (*Task 7.5*) demos.

Specific topic and problems tackled by the work

A central issue in robot vision is how to independently explore an unfamiliar environment, what to learn and how. To do this the robot should be able first to locate interesting objects in the scene and then tag and track them in order to build an internal representation of the surrounding environment and then perform complex tasks. The ultimate goal of this research is to create a biologically inspired autonomous vision system, capable of exploring the environment and determining the subset of relevant objects that will be subsequently learnt for future recognition.

Methods and key results

Observations and experimental data from the "Familiarity-to-novelty shift" (FTNS) task in infants have been used to build a conceptual and a computational model motivated by learning (Wang et al. 2011). By analyzing the data, we propose that infants tested on FTNS prefer the stimulus that provides them the highest learning rate. The learning rate, according to other works (Schmidhuber 2009), has been proposed to be the intrinsic motivation that drives the attention selection in infants. We use, hence, a prior condition (*intrinsic knowledge*) that we are looking for something that provides a higher learning rate in order to bias the saliency of a scene obtained by bottom-up mechanisms (Itti and Koch, 1998, Migliore et al. 2011). The robot then starts to move its head while simultaneously controlling its vergence in order to focus this salient point onto the plane of fixation, in a completely autonomous fashion. Democratic cue integration (Triesch and v.d. Malsburg, 2001) is used for real-time tracking of location of the target object. The object in the plane of fixation is then segmented from the rest of the scene for further learning by stereo matching using Gabor features. Gabor features are calculated at interest points (Harris corner points in our work) on both left and right camera images and corresponding matching point pairs are found out based on a similarity measure. We then compute the "complexity" of the object (that we choose proportional to the number of features found on the object) such that the fixation time will depend on the object complexity. The system is hence ready to save the features, that consist of 40-dimensional Gabor-jets that were found on the segmented object at all interest points, and starts to cumulatively build a model (learn) for the online recognition that follows. The system is able to autonomously select and learn the objects in the scene. By comparing the performances with a not intrinsically motivated exploration (without considering the aspect of searching for regions in the space that provide a high learning rate) our system is able to scan the critical points of the scene much faster.

Relevance of the results for the project

Following learning, the system will be able to identify objects, based on the learned models, independent of their pose and environmental conditions (different backgrounds, illumination, contrast, occlusion, etc). Future work will target on-line learning and provide top-down information to the current saliency map. Furthermore, we developed a separate control module that allows eye movements during single joint movements of the head, in order to maintain the object in the center of the field of view. This allows iCub to perform neck movements according to the object complexity in order to save features at multiple poses (orientations) of the object. This in turn will add a set of

action repertoire that can be exploited in WP7 (the robot is capable of autonomously developing an *action repertoire* on the basis of intrinsic motivations).

Selected bibliography

L. Itti and C. Koch. Computational modelling of visual attention. Nature reviews neuroscience, 2(3):194-203, 2001.

J. Triesch and C. v.d. Malsburg. Democratic integration: Self-organized integration of adaptive cues. Neural Computaiton , 13(9):2049-2074, 2001.

J. Schmidhuber. Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. In Anticipatory Behavior in Adaptive Learning Systems. From Psychological Theories to Artificial Cognitive Systems, volume 5499 of LNCS, pages 4876. Springer, 2009.

Pierre-Yves Oudeyer, Frédéric Kaplan, and Verena V. Hafner Intrinsic Motivation Systems for Autonomous Mental Development IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION, VOL. 11, NO. 2, APRIL 2007

Q. Wang, P. Chandrashekhariah and G. Spina Attention shift in infant habituation: a conceptual and computational model motivated by learning, submitted to ICDL 2011 (accepted)

Davide Migliore, Emmett Kerr, Gabriele Spina, Pramod Chandrashekhariah, Richard Veale, Yiannis Gatsoulis, Jochen Triesch, A NeuroRobotic Model of Infant Looking Behavior, First IM-CLeVer Spring School Project Report, 1-7 May, 2011, Capocaccia, Italy

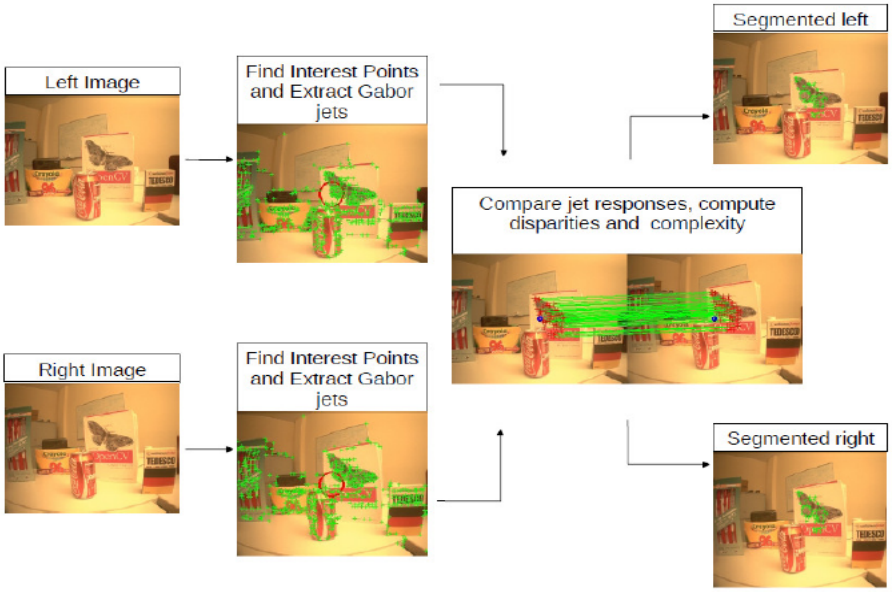

Figure 4.3.1: Block diagram of stereo computation used in object detection.

Task 4.3 (**CNR-ISTC**): Learning of attention and vergence control
Learning to control attention in an eye-arm robot:

Link, relevance, and input to the project of this work

CLEVER-B2 (Task 7.5) and various models built within the project (e.g., see models of Task 5.1), use attention to guide the action of an arm, for example for reaching suitable targets ("pragmatic actions"): in these models the eye control is, in most cases, hardwired. Other tasks, e.g. various tasks of this WP, show systems where attention control is learned to search pre-defined visual targets, to guide vergence, ecc. ("epistemic actions"), but it is not used to guide an actuator such as an arm or hand directed to manipulate and change the environment. When attention control is learned to guide a manipulation actuator, a number of interesting problems arise. The study of these problems is the target of this work. Such problems are studied on the basis of a bio-inspired model that allowed us to focus on computational problems and to discover principles that might be later used to model and understand biological systems. The output of this work furnishes insights relevant for CLEVER-B and other tasks where attention is used to guide the pragmatic action of the iCub robot. This work built on the research carried out by CNR-ISTC within another EU project named "MINDRaces": this is a relevant example of synergy between EU projects. The work produced three publications: Ognibene et al., 2009, 2010a, 2010b (the latter one is attached to this Deliverable).

The specific topic and problems tackled by the model

The active-vision/attention-for-action paradigm claims that only information relevant to accomplish the task at hand should be searched and processed. Overt visual attention, based on the control of a highly-sensitive but space-limited fovea, is a powerful mechanism that allows implementing this principle, but its use poses challenging learning problems. These involve the guidance of attention itself before suitable representations are formed, the introduction of a strong partial observability of the environment, the need of learning processes that are at the same time fast and capable of generalising the acquired capabilities to different contexts.

The model was tested with the following task. To test the model we used a custom eye-arm robot engaged in learning attentional skills in order to guide the performance of targeted reaching movements (see Figures below). The setup used to test the model is based on a kinematic simulation of the robotic setup illustrated in the figure below. This robotic setup is formed by a 4 degrees of freedom custom robotic arm and a fixed camera. In the tests the system sees an image formed by various coloured squares posed on the vertexes of a regular 6x8 grid. The system has to learn to find and touch, as quickly as possible, a red target in scenes which include also green elements (cues) forming a vertical line of variable size and position and blue scattered elements (distractors). The target is posed on the left side of the cues in any position along the line they form. The images acquired by the camera are used to simulate a moving camera by suitably cutting a sub-image, centred on the "eye gaze", from the whole camera image. Reaching actions result in an external reinforcement when they lead to reach the target, whereas they produce a small punishment (equivalent to energy consumption) when they lead the system to reach a cue or a distractor.

Methods and key results

The model has the following architecture and functioning. The architecture encompasses two neural controllers: (a) an attention controller, which integrates bottom-up and top-down information; (b) a controller for the 4 degrees of freedom arm. The attention controller selects the next point to gaze by integrating information produced by three components. The first component forms a bottom-up attentional sub-system whereas the second and third components form a top-down attention component.

The three components are as follow: (a) a pre-wired mechanism (bottom-up saliency map), which is task-independent and is fed by the peripheral input; (b) an adaptive mechanism, based on an actor-critic reinforcement learning architecture, which is fed by the fovea input and, as output (actor), "votes" for performing a saccade in one or more points of the environment (the output of the actor is a neural map that encodes the possible position in space to look at based on a population code); note how these votes are formulated on the basis of the currently foveated object: this allows the system to learn to look at different positions in space on the basis of such object; (c) a memory (Potential Action Memory - PAM), that stores potential saccade targets suggested by the actor in the latter cycles of the simulation. The PAM is formed by leaky neurons and is continuously updated on the basis of the "votes" of the actor. The PAM reference system is anchored to the body of the system, whereas the actor reference system is anchored to the gaze direction, so the actor output is suitably shifted before activating the PAM so to keep the PAM representation coherent in time.
The arm controller is trained before the experiments with attentional learning and is less relevant for the research (it was indeed developed in previous works). Within the arm controller, a neural competition between possible hand positions corresponding to recently gazed points takes place in order to decide the target for reaching. This competition is based on a mechanism for which the longer a target is foveated the highest the probability that it is selected as a target of reaching (this is based on a map of leaky neurons: when some neurons activate above a certain threshold, then a reaching action is triggered).
When a reaching action with a target corresponding to the currently foveated point in space is selected, the arm control model computes the visuo-motor transformations suitable for the arm, in particular it transforms the gaze direction (eye postures) to the desired hand posture encoded in terms of arm joint angles corresponding to the final position of reaching (these angles are then fed into the servos of the robotic arm for the execution of the reaching movement). The arm component was trained before the experiment where the system learns to perform the top-down attention control. The rewards and punishments are used by the reinforcement-learning top-down attention component to adapt the behaviour of attention so to explore the environment, find the target and trigger a reaching actions.

The results show that the system achieves a high performance very quickly, notwithstanding the partial observability that characterise the task and the limited available computational resources of the system. This results are made possible by: (a) the constraints posed on the architecture by the active-vision approach followed; (b) the successful interplay between the bottom-up and the top-down attentional components that lead to an interesting staged development.
The results show that the PAM plays a fundamental role for both exploration and learning processes. In particular, it allows achieving a representation of the spatial regularities of the objects in the environment, so permitting to obtain successful behavior with significantly less learning episodes compared to a system which uses only inhibition of return or no memory mechanisms. This is possible because the PAM produce action hypotheses for actions that the agent has never experienced before (e.g.: look at the left of any green cue you see), and this allows testing novel positions to look at. In this respect, the PAM gives rise to an active learning strategy based on the spatial relationships and regularities of the environment objects.
The developmental trajectory of the system exhibits these phases. The system first starts to learn to act with respect to the the most perceptually salient features of its environment (cues), in particular by generating an emergent inhibition of return, then it learns to exploit the cues to foveate the target (by integrating the top-down probability to explore the area at the left of the cues), and finally also learns to look up or down the distractors in the case it foveates one of them after the cues. By incrementally structuring knowledge in a task dependent way, these phases of learning allow the system to exploit bottom-up saliency to bootstrap and scaffold top-down learning, and then build a suitable top-down strategy on top of it.

Relevance of the results

Overall, the results show how learning attentional systems used to guide action do not need to learn whole scene statistics to pursue their goals, but they can rather develop, in a cumulative fashion and by "bootstrapping" on the basis of bottom-up saliency-based processes, task-specific exploratory strategies based only on the relevant regularities of the environment. The agent acquires easily this top down bias because of its structure: (a) the top-down fovea control aids learning local knowledge and at the same time eases generalization to other contexts; (b) the bottom up saliency map restrict the set of possible perceptions, selecting part of the environment which are similar and on which the acquired local knowledge can be applied.

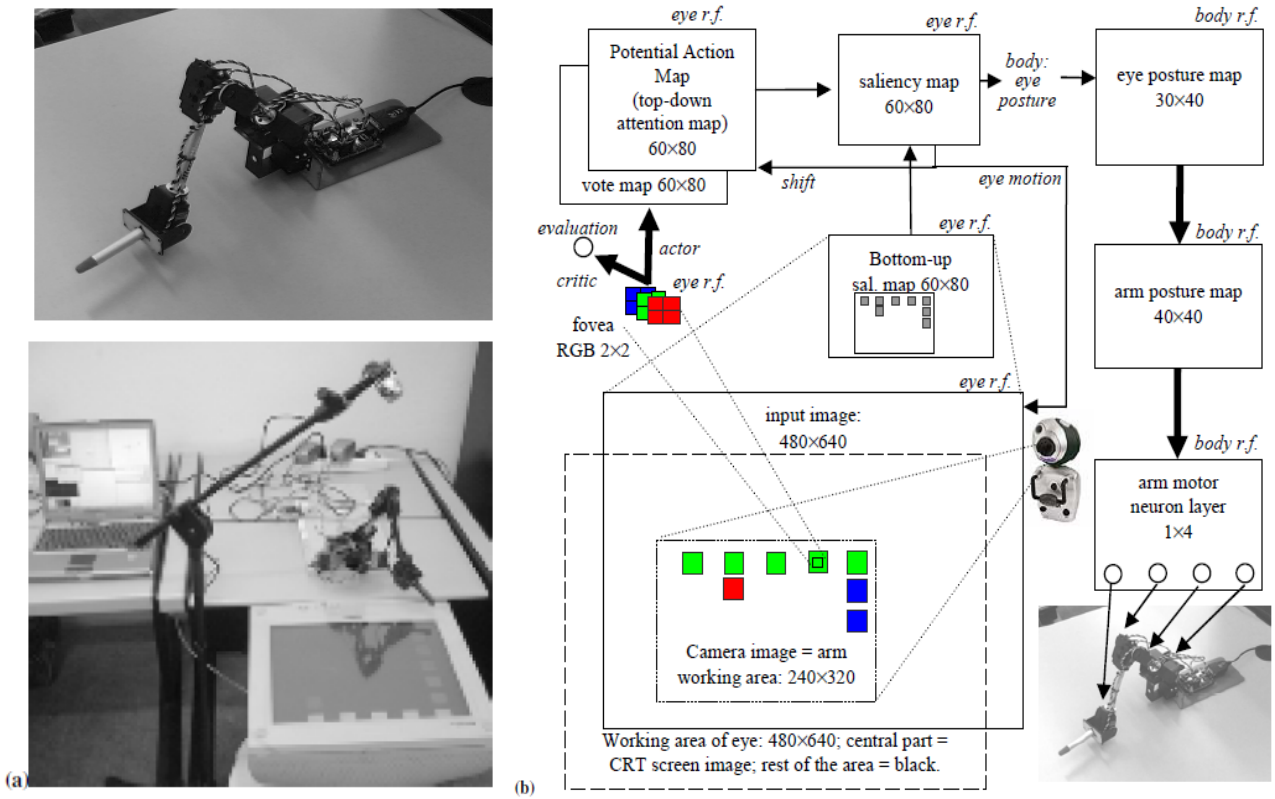

Figure 4.3.2 (a) The robotic arm built at CNR (top), and the experimental set-up used to test the system (bottom). The arm is a customised 3-link robotic arm, th at with its tip attempts to reach a coloured target projected onto a computer screen (bottom: the screen is the working plane of the robot arm). The screen also contains green cues (having a regular, but variable, spatial relation with the target) and blue distractors (having a random spatial relation to the target). The scene is observed from a camera mounted on top of the arm and looking down towards the screen (bottom). A sub-image is suitably "cut" from the whole camera image, with a position given by the attentional system (in this way we could simulate the image grabbed by a fast moving eye). (b) The architecture of the system. The camera image is used to build a larger image with external annulus assumed to be formed by black pixels. The image sent into the system is an image cut from the such larger image, and is formed by a large periphery and a very-small fovea. The pictures shows, at the top left, the bottom-up attention component (receiving the periphery input), the top down reinforcement learning component (receiving the fovea input), the PAM (activated by the reinforcement learning component), and a saliency map that sums the two and decides the saccade to perform. The right

part of the picture shows the arm controller, that maps a foveated target (the action of the arm is triggered when the area is foveated for long) to the corresponding arm posture.

Selected bibliography

Ognibene Dimitri, Pezzulo Giovanni, Baldassarre Gianluca (2009). How Are Representations Affected by Scene Statistics in an Adaptive Active Vision System? In Canamero Lola, Oudeyer Pierre-Yves, Balkenius Christian (eds.), Proceedings of the Ninth International Conference on Epigenetic Robotics (EpiRob'09). pp. 229-230. Lund University Cognitive Studies, 146. Lund: Lund University.

Ognibene Dimitri, Pezzulo Giovanni, Baldassare Gianluca (2010a). Learning to look in different enviroments: an active-vision model which learns and readapts visual routines. In Doncieux Stephane, Girard Benoit, Guillot Agnes, Hallam John, Meyer Jean-Arcady and Mouret Jean-Baptiste (Eds) From Animals to Animats 11 – Proceedings of the 11th International Conference on Simulation of Adaptive Behavior, SAB 2010. Paris – Clos Luce', France, August 2010

Ognibene Dimitri, Pezzulo Giovanni, Baldassarre Gianluca (2010b). How can bottom-up information shape learning of top-down attention-control skills? In Kuipers Benjamin, Shultz Thomas, Stoytchev Aalex, Yu Chec (eds.), IEEE International Conference on Development and Learning (ICDL2010). pp. 231-237. Piscataway, NJ: IEEE.University of Michigan – Ann Arbor, MA, August 18-21 2010

Task 4.3 (**USFD**): Learning of attention and vergence control
Using visual attention to obtain 'where' from 'what': a model of the visual system

Topic
Biologically plausible models of task directed gaze and attention. Integration of oculomotor control and visual attentiuon.

Specifc problems tackled
How can attentionally primed, spatially invariant object representations in higher order visual areas (like inferotemporal cortex) influence spatially topographic representations of gaze direction?
How can this competence be integrated with simpler visual gaze control directed by luminance and phasic signals?

Methods/Model
As a primate investigates the world around it, locating and directing gaze to objects in the visual scene performs a vital role. In the primate brain, however, the systems devote to recognising objects and to directing gaze diverge into separate cortical pathways (Ungeleider and Mishkin, 1982). The pathway for gaze redirection contains no processing to identify objects, and the pathway for object recognition identifies objects in the visual scene with little spatial information, thus avoiding a combinatorial explosion of positions and object identities. Previous work has shown that visual attention can provide a key role in recovering the spatial locations of objects that are relevant for the current task, and several models have demonstrated working versions of such architectures (Corbetta and Shulman, 2002;Fox et al., 2006; van de Velde and de Kamps, 2001; Deco and Rolls, 2004;Hamker, 2004).
We have extended this work in two main directions. Firstly we have provided constraints on the mechanisms of attention in the recognition pathway for biologically based models, by seeking the simplest reliable mechanism that allows attention to recover the visual location of objects. Secondly, we have embedded this model in a complete biologically based model of the oculomotor system, allowing the effects of this visual attention to direct gaze. This work has allowed us to both validate the model against a wide range of literature, from single cell recordings though to behavioural psychophysics data, and to make predictions on, and explanations of, the mechanisms and dynamical processes involved in object-based gaze redirection in primates. Further work has been completed in embodying this model in robotic hardware, with the aim of testing the performance of the model in a moreunconstrained, real-world, environment.

Key results
Attention requires the use of separate interacting channels for feedforward and feedback in order to perform best, and the use of a fundamentally multiplicative mechanism to combine the feedforward and feedback information.
The effects of attention in recovering spatial information when objects in the visual world are either perceptually similar or distinct, combined with the dynamics of selection in the oculomotor system, can reproduce and explain a range of the behavioural results in visual search.
By combining the phasic information in object onsets detected by the subcortex in the oculomotor system of the model, with the task information from the object recognition pathway, the effects of onset in visual search can be reproduced and explained.
This work is now being continued under funding from the IM-CleVer project, and will be written up with acknowledgement to the project.

Selected bibliography

Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. Nature reviews. Neuroscience, 3(3):201–15.

Deco, G. and Rolls, E. T. (2004).A neurodynamical cortical model of visual attention and invariant object recognition. Vision Research, 44(6):621–642.

Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., and Raichle, M. E. (2006). Spontaneous neuronal
activity distinguishes human dorsal and ventral attention systems. Proceedings of the National Academy of Sciences of the United States of America, 103(26):10046–51.

Hamker, F. (2004). A dynamic model of how feature cues guide spatial attention. Vision Research, 44(5):501.
Ungerleider, L. G. and Mishkin, M. (1982). Two cortical visual systems, pages 549–586. Cambridge, MA: MIT Press.

derVelde, F. and de Kamps, M. (2001). From knowing what to knowing where: Modeling object-based
attention with feedback disinhibition of activation. Journal of Cognitive Neuroscience, 13:479–491.

Task 4.4 (**AU**) Development of visual, tactile and motor abstract representations

Link and relevance for the project of the presented work

This work forms the foundation of the CLEVER-B architecture, defining the sensory motor control system for the iCub. Through this control, the decision making components will be able to interact with the robotic platform. This work also integrates with the work from WPs 5 and 6, using novelty to drive the development of learning and hierarchical architectures to build greater representations of the world from the sensory and motor abstraction.

Specific topic and problems tackled by the work

When interacting with the environment it is important to be able to build a model of the interactions in order to understand and reuse them. This model must reach from the sensory input information acquired, through their correlations with the motor spaces then to a higher level of representing objects in the world. By abstracting away from the details it is possible to define a model that can be reused at each of the different levels and is extendible to incorporate different sensory systems and understanding of the world around us (4).

Methods and key results

The model is built using mappings to provide a "content neutral" data structure that allows for rapid learning and adaptation. These mappings link sensory input systems with the motor control systems and make use of field structures to give abstractions in terms of resolution and provide scope for defining features between specific sets of points. The use of fields also allows for extrapolation where there are gaps in the mappings, allowing the mappings to rapidly build up over a small number of trials (1).

The results obtained illustrate the rapid sensorimotor development required for learning the basic motor control of the iCub robot. All the tests were performed on the real robot without requiring the use of long training sessions in a simulator.

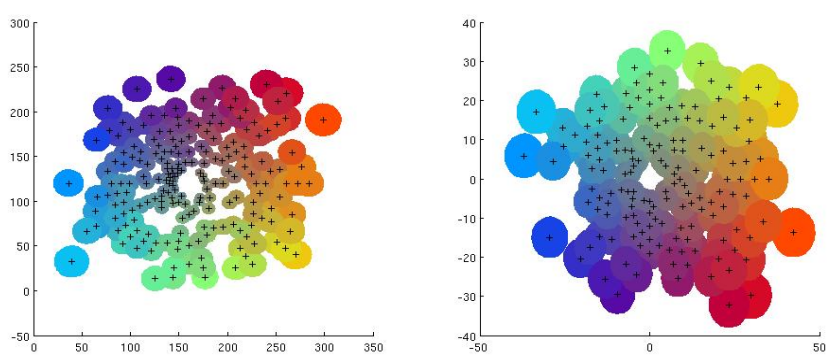

Figure 4.4.1: Sensorimotor mapping linking the retina (left) and eye pan/tilt motors (right)

The eye sensorimotor mapping in Figure 1 shows the fields that have been linked, shown here by colour, between the retina map (left) and the eye motor map (right) after a period of learning in which 50 saccades were completed. On average, after 20 saccades, the eye-motor map is sufficiently populated to reach most targets in 1 or 2 steps, after which the addition of new links slows, adding links only when gaps are encountered. The success of eye saccades is used as an indicator for the lifting or applying of a constraint to the next stage of development,in line with infant development (3)

introducing neck movements that give rise to a larger gaze space, as discussed in Task 6.1. Further experiments are looking into whether this staged development can be emergent (2).

Relevance of the results for the project

The abstract representation developed here allows mappings to be learnt quickly and updated dynamically for a wide range of sensorimotor mappings, or higher level maps. Due to their nature, the maps can be learnt on the real robot without prior training or calibration. This will help to speed up the developmental process and give the higher level decision making components developed in WP6 a solid foundation on which to operate.

Selected bibliography

1. Lee, M., Law, L., Huelse, M. A Developmental Framework for Versatile, Intrinsically-Motivated, Learning Robots. Book chapter, under review.
2. Huelse, M., Lee, M. Adaptation of coupled sensorimotor mappings: An investigation towards developmental learning of humanoids. In: Doncieux, S. et al. (Eds.) SAB 2010, Springer, LNAI 6226, pp. 468-477.
3. Law, J., Lee, M., Huelse, M.: Infant development sequences for shaping sensorimotor learning in humanoid robots. In: Johansson, B., Sahin, E., Balkenius, Ch. (Eds) Proc. 10th Int. Conf. on Epigenetic Robotics 2010 (EpiRob 2010), Lund University Cognitive Studies, 149, Sweden, 2010, 65-72.
4. **Shaw, P., Development of Visual and Motor Abstract Representations, internal document**

Task 4.5 (**IDSIA-SUPSI**): Learning abstract predictive world models.

Link and relevance for the project of the presented work

The goal of this task is to develop predictors (world models), based on powerful abstraction internal capabilities, capable of predicting the consequences of systems' actions. These predictors will be exploited in the project for these purposes: (1) develop key components for implementing the novelty detectors developed in WP5 (2) scale and extend recent methods that combine on-line reinforcement learning with the simultaneous acquisition of world models: this topic, investigated in this task, is expected to furnish algorithms and insights useful for the hierarchical models developed in task 6.4; (3) show how, by abstracting away irrelevant details, predictors can overcome problems generated by partially observable environments in efficient ways.

Specific topic and problems tackled by the work

The goals we worked toward are the development of fast and robust algorithms for abstract world prediction that can handle high-dimensional vision data, but still operate *in real time* for experiments with the iCub robot. We focused on two complementary approaches based on computer vision, and machine learning.

Methods and key results

Based on the K2 framework for learning through intrinsic motivations (see Pape et al., 2010), we divided the learning of abstract predictive world models into two components: (1) computer vision and machine-learning abstraction algorithms and (2) adaptive predictors. This division facilitates the use of simple, efficient learning of the predictor algorithms, because the irrelevant details are abstracted away. For the abstraction algorithms we used two complementary approaches; (1a) computer vision modules that detect the location and orientation of objects in the world, and places them in a fast abstract world simulator (see Frank et al., 2011) that can be used by various additional methods (e.g. evolutionary or RL algorithms) for motion planning, or learning to predict the outcome of the robot's actions. (1b) machine learning methods for spatial (compression) and temporal (prediction) abstraction.

In the computer vision approach, we combined the saliency map (Itti 1998) described in T4.1 with shape (Ciresan et al.) and texture classification (Corani et al. 2010), for robust objects identification. The positions of objects identified in this way can be compared with of the already viewed objects (T4.2), and for creating an abstract 3D representation (Frank et al., 2011) of the surrounding environment (for obstacle avoidance, grasping, etc.). Limitation of the robot related to low resolution of the cameras and the limited parallax are tackled with a novel solution called Nested Extended Kalman Filter (NEKF), an evolution of the classic approach, based on Extended Kalman Filter used in Simultaneous Localization and Mapping (SLAM) (Durrant-Whyte, 2006).

We used and developed several unsupervised machine learning methods for reducing the dimensionality of visual input data. Incremental slow features (a novel algorithm developed at IDSIA; Kompella et al., 2011), were combined with unsupervised autoencoders. The resulting method, called AutoIncSFA learns to represents abstract slow features for moving objects, such as the robots arms or objects, from vision input in an unsupervised manner. See figures 4.5.1 and 4.5.5.

Second, we developed a fully unsupervised algorithm based on deep belief networks (Hinton et al., 2006) that can overcome the catastrophic forgetting problem that is often encountered in connectionist architectures (Pape et al., 2011). We expect that this approach will be useful for visual preprocessing in learning systems where unsupervised feature abstraction performed on raw camera images is presented to subsequent reinforcement learning algorithms that control the behavior of the robot. Such a system could safely focus on learning novel features, without the

disadvantage of forgetting the abstract features learned before.
Third, we developed a reinforcement learning approach that is able to deal with rapidly changing intrinsic curiosity reward, using a planning strategy in an abstract model of the environment (Luciw et al., 2011). The approach uses vector quantization to build an internal state layer that represents the environment. Transition tables of this internal layer are then learned as the agent moves through its environment. Each internal state-action pair is associated to a pair of predictors, one with a long-term memory and the other short-term. These predictors keep a weighted average of the sensory reconstruction error and the state transition error. The reducible error — the difference in the two averages — in the reconstruction and state-transitions gives a valid measure of interestingness. That is, places where there is something to be *learned* will be considered interesting. The system was tested in high-dimensional, noisy, visual environment: an overhead view of a 2-dimensional maze, see Figure 4.5.2. The system was able to choose actions that improved both its perceptual and cognitive (including value prediction) systems. This enabled the agent to quickly learn (Figure 4.5.3) how to act so as to maximize external reward. Unlike other implementations (Oudeyer, 2007), our system explores the environment successfully without requiring any random, uninformed actions.

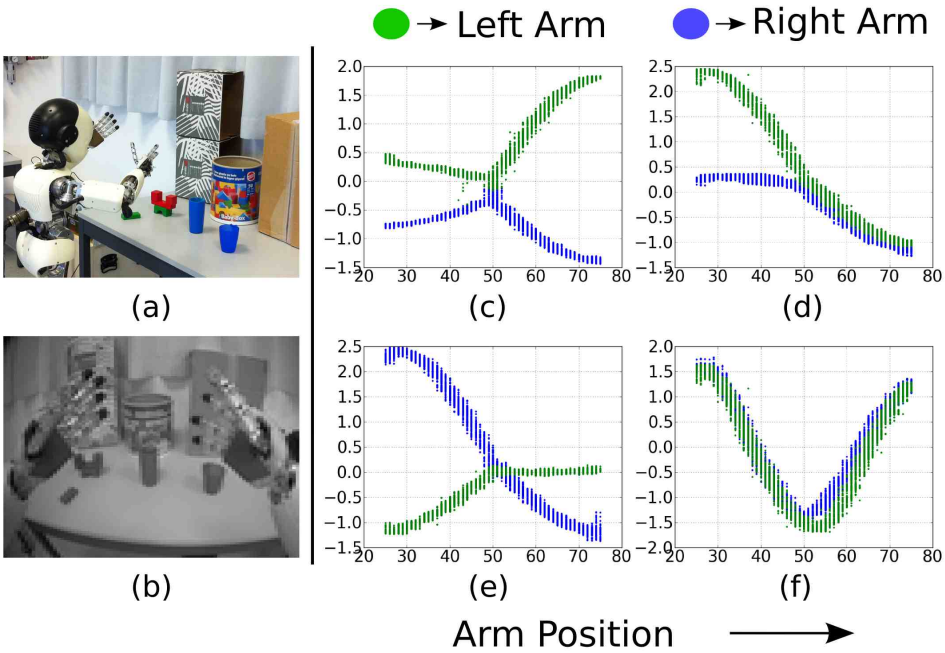

Figure 4.5.1: iCub-Arms Experiment: (a) Experimental setup: iCub observing its arms' movements. (b) Sample image from the dataset. (c) Unit 1 encodes the identity of each arm; (d)-(e) units 2 and 3 encode the left and the right arm's position respectively; (f) unit 4 encodes the arm position, and is invariant to identity.

Relevance of the results for the project

The results on standard benchmarks for the machine learning algorithms we developed are reported in several conference publications (Luciw et al., 2011, Kompella et al., 2011; Pape et al., 2011). We intend to further develop, integrate and test these methods on the iCub robot for learning a hierarchy of skills based on intrinsic motivation.
The computer vision and machine learning methods based on slow features were also tested on the iCub robot. Figure 4.5.4 shows the extracted location and orientation of a blue cup on the table with the computer vision module. This module outputs its findings over YARP to subsequent modules,

such as a neural network predictor, or the Virtual Skin abstract world model, which can be used for path planning and safety. Figure 4.5.5 shows the change in the slowest feature of the AutoIncSFA method, corresponding to a change in vertical orientation of the blue cup. In the context of WP7, the abstract features extracted by computer vision or AutoIncSFA are learned with linear and neural network predictor as a function of the robot's pose, its intended actions, and the current state of the objects in the world. The error and decrease in error made by this abstract predictive world model serves as the reward for a reinforcement learning algorithm action selector.

In the learning tasks described here, the abstract prediction methods consisted of multiple modules with state of the art computer vision and machine learning methods, which greatly reduced the dimensionality of the prediction task. Based on the excellent results achieved with these methods, we expect it will be possible to use the learned abstract features as input to recurrent neural networks, such as LSTM (Hochreiter and Schmidhuber, 1997), which can learn to overcome problems generated by partially observable environments in efficient ways.

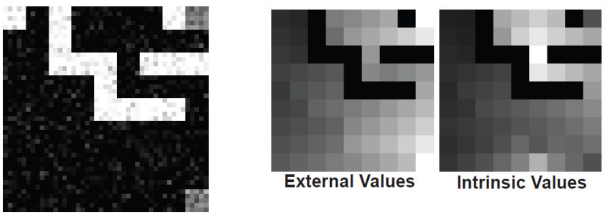

Figure 4.5.2 Left: Example high-dimensional observation in "duck-world". The agent sees its own position (white square in the upper left) and the entire environment, including goals (upper and lower right), walls (other white areas), and passable areas (black areas) as a noisy image taken from above. The agent can move up, down, left or right. Every time the agent touches a goal area, it teleports back to the upper left. The upper right goal is hard to reach using exploration based on randomness. **Right:** Examples of external and intrinsic state values for an agent that has reached both goals. The external values are stable, but the intrinsic ones are constantly changing. The agent handles the changing reward landscape using a planning technique (LSPI).

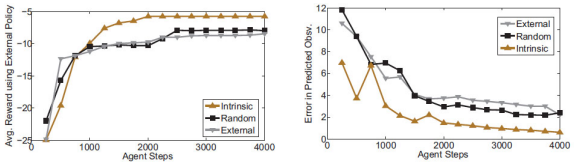

(a) Action selection methods (b) Average error prediction

Figure 4.5.3 Left: Average performance of the agent's computed external policy from all possible start positions over time, using three different action selection methods. The intrinsically motivated agent with planning achieves nearly optimal performance. **Right:** Average error in the predicted next observation for a sampling over all possible state-actions over time. Through intrinsic motivation, the agent has developed the capability to both effectively perceive and predict its environment.

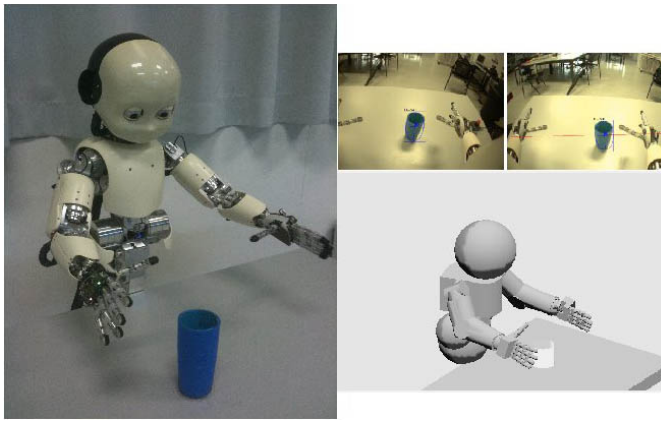

Figure 4.5.4: The computer vision module extracts the 3D location and orientation of the blue cup from visual input (top right), and places it in the abstract world simulator (bottom right). All modules communicate via YARP.

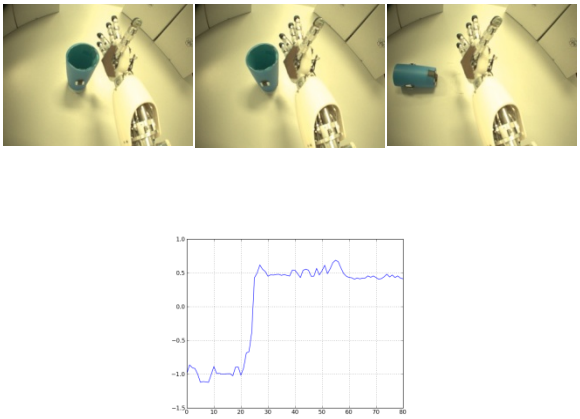

Figure 4.5.5: The slowest AutoIncSFA feature as a function of time (bottom) shows a step change when the blue cup is toppled (top).

Selected bibliography

M. Frank, A. Förster, J. Schmidhuber. Virtual Skin - Infrastructure for Curious Robots, 2011.

S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation, 9(8):1735-1780, 1997.*

V. R. Kompella, M. Luciw, J. Schmidhuber. Incremental slow feature analysis. IJCAI-2011, Barcelona, 2011.

L. Itti et al. A model of saliency-based visual attention for rapid scene analysis. PAMI 1998, 20(11):1254-1259.

D. Ciresan, U. Meier, J. Masci, L. M. Gambardella, J. Schmidhuber flexible high performance convolutional neural networks for image classification, IJCAI 2011.

G. Corani, A. Giusti, D. Migliore, J. Schmidhuber: "Robust Texture Recognition Using Credal Classifiers". Proc. of BMVC 2010.

G. E. Hinton, S. Osindero, and Y. W. Teh. A fast learning algorithm for deep belief nets. Neural Computation, 18:1527–1554, 2006.

M. Luciw, V. Graziano, M. Ring, and J. Schmidhuber, Artificial Curiosity with Planning for Autonomous Perceptual and Cognitive Development, ICDL2011.

P.-Y. Oudeyer, F. Kaplan, and V. Hafner. Intrinsic motivation systems for autonomous mental development. IEEE Transactions on Evolutionary Computation, 11(6), 2007.

L. Pape, M. Ring, F. Gomez, K. Frank, A. Förster, D. Migliore and J. Schmidhuber, (2010), IM-CLeVeR-K2 functional architecture, internal report of the EU-funded Integrated Project “IM-CLeVeR – Intrinsically Motivated Cumulative Learning Versatile Robots”.

L. Pape, F. Gomez, M. Ring and J. Schmidhuber, Modular deep belief networks that do not forget. IJCNN, San Jose, CA, 2011.

Durrant-Whyte, H.; Bailey, T. Simultaneous Localization and Mapping (SLAM): Part I The Essential Algorithms. Robotics and Automation Magazine 13: 99–110, 2006.