

Visually Guided Whole Body Interaction

Bram Bolder, Mark Dunn, Michael Gienger, Herbert Janssen, Hisashi Sugiura and Christian Goerick
Honda Research Institute Europe
Carl-Legien-Straße 30
D-63073 Offenbach am Main, Germany
{bram.bolder, mark.dunn, michael.gienger, herbert.janssen, hisashi.sugiura, christian.goerick}@honda-ri.de

Abstract—The design and implementation of a system is introduced that enables a humanoid robot to interact with its visual environment using its whole body. To allow this interaction, several new concepts were developed and implemented. Targets for the interaction are based on visually extracted proto-objects. An innovative control system allows the robot to increase its range of interaction, achieve multiple targets simultaneously, and avoid undesired postures. Two different selection mechanisms were used to switch between different kinds of behaviors at any time and posture.

I. INTRODUCTION

Research on humanoid robots is increasingly focusing on interaction in complex environments, including autonomous decision making and complex coordinated behavior.

Several interactive robot systems were already introduced. A complete architecture for a small humanoid (Sony QRIO) including perceptual abilities that uses a central action selection driven by so called *behavior values* provided by the individual behaviors is described in [1], [2]. This robot realizes some impressive abilities including emotional/motivational control and learning and humanoid multi-degree of freedom control.

Kismet [3] also realizes several impressive interaction abilities and contains both a powerful vision and attention system and behavior selection. It also integrates low level vision representations as feature maps and higher level representations for faces. The main focus of his system is child-like interaction and developmental learning.

Orabona et al. [4], [5] developed a system that uses proto-objects — in their case colored blobs — to let a robot learn the notion of an object consisting of possibly multiple proto-objects using statistical means.

Proto-objects are a concept originating from psychophysical modeling [6], [7], [8]. These can be thought of as coherent regions in the field of view that are trackable and can be pointed or referred to without identification. In this work, the term *proto-objects* is used in a similar manner. They are extended to three dimensions to cope for egomotion and for large changes in size due to depth changes.

The need for non-monolithic internal representation of the environment as well as non-monolithic control [9] inspired the distinction of proto-objects and object hypotheses. Object hypotheses interpret proto-objects as originating from physical objects according to some generic models to allow different behaviors to interact. Some behaviors don not even need object hypotheses, they can operate only using proto-objects.

In the field of strategy selection, Wolpert et al. [10] propose a learning architecture that consists of several inverse and forward models. According to the prediction error, the most adequate strategy is selected.

Regarding the challenge of realizing a fully autonomous interactive humanoid robot, an intermediate step is presented in this work: ASIMO is able to interact with the environment driven by visual perception using simple decision making and coordinated whole body motion.

Our approach is to build a system that — for now — uses a relatively simple definition of visual target objects, namely any elongated colored object. This system implements the fundamental elements of an architecture that is easily extendible to cope with more long term targets. Integrating a more elaborate computer vision system such as [11] is of course possible, but here the focus lies on closing the interaction loop.

Novel in the context of humanoid robots are the following key points:

- Relatively low level perceptual information (blobs) is represented as proto-objects in a short-term sensory memory. These proto-objects can be used both in raw form to e.g. visually track them in 3d, or to form stable object hypotheses as they are needed for reaching and grasping.
- Decision mechanisms that evaluate behavioral alternatives based on sensory information and internal prediction.
- A motion control system that is able to be driven by a wide range of possible target descriptions and that ensures smooth well coordinated whole body movements (using a set of cost functions as null space criteria).

The general design of the system is depicted in Fig. 1. The perception system uses color and stereo based 3d information to detect relevant visual stimuli. It keeps this information as proto-objects in a short-term sensory memory. This sensory memory is then used to derive targets for visual tracking and to form stable object hypotheses from which movement targets for reaching movements can be derived. A prediction based decision system selects the best movement strategy and executes it in real time. The internal prediction as well as the executed movements use an integrated control system that uses a flexible target description in task space in addition to cost-functions in null space to achieve well coordinated and smooth whole body movements. The system is implemented using the real time environment RTBOS [12].

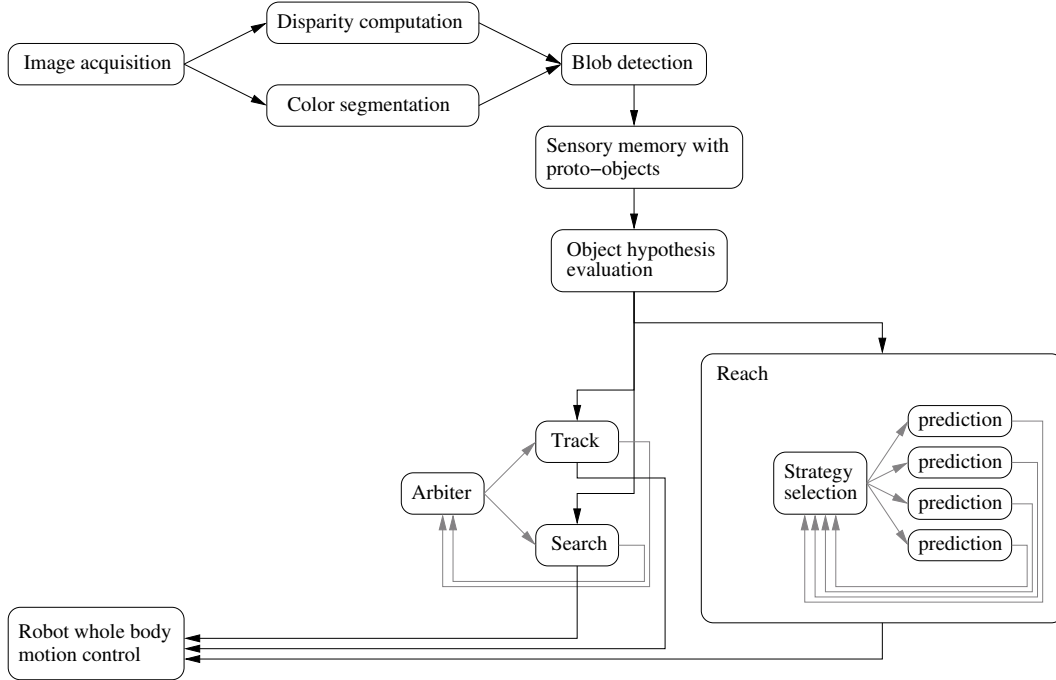


Fig. 1. Overview of the system design.

II. DESIGN OF THE SYSTEM

Several components of the system are discussed in more detail, since they contain the actual concepts needed to create the whole interactive system.

A. Image Preprocessing

To be able to generate proto-objects the image processing has to be able to find entities in the environment that are dynamically stable in position and extent. Thus the best candidates are segmentation algorithms such as color segmentation, texture segmentation, or feature extractors that find unique salient points. To obtain 3d information stereo disparity calculations or other stereo algorithms can be used.

Pairs of color images are used, labeled with the time of acquisition. These images are processed in two parallel paths:

- 1) **Stereo disparity computation:** Intensity image pairs are generated from the color image pairs. The images are rectified, i.e. transformed so that the result corresponds to images captured with two pin-hole cameras with collinear image rows. The horizontal disparities between corresponding features in both images are computed if the features are sufficiently prominent.
- 2) **Color segmentation:** One color image of the pair is rectified as above and converted to the HLS color space. All pixels are evaluated as to whether they lie in a certain volume in HLS space and the result is subjected to morphological operations that eliminate small regions of one class of pixel. The resulting pixels that lie in the HLS volume are grouped into regions that are contiguous in the image plane. The largest resulting

groups that exceed a minimum size are selected for further processing.

B. Blob Detection

Based on the 3d and image segmentation information from the image preprocessing a symbolic representation as 2d ellipsoids in 3d space (blobs) is generated.

For each of the groups from the color segmentation the center of area in the image plane x_p, y_p and the median of the disparities d of all its pixels are computed. Also it is detected whether the group region touches the image boundaries; if so the data is labeled as inaccurate since parts of the real world object corresponding to the region are probably outside the field of view. The orientation of the principal axis ω_p and the standard deviations $(\sigma_{p1}, \sigma_{p2})$ of the pixels in the image plane are computed for each group using a principal component analysis of the correlation matrix of the pixel positions. The extracted ellipse and its axes of a region of specified color can be seen in Fig. 2. Using the camera system geometry, vectors (x_p, y_p, d) and $(\sigma_{p1}, \sigma_{p2})$ are transformed to the metric coordinates (x_c, y_c, z_c) and the metric standard deviations σ_{c1} and σ_{c2} . Using the time label of the images, the robot posture and position at the time of image capture is derived. This is then used to transform the position to world coordinates $\vec{r}_b = A \cdot (x_c, y_c, z_c)$ and the principal axis orientation ω_p to an orientation vector ω_c in world coordinates. We define a blob as a set of data consisting of the time label, the position \vec{r}_b , the orientation ω_c , the standard deviations σ_{c1} and σ_{c2} , and the label whether the data is accurate or inaccurate as described above.

Fig. 2. Extracting a blob in form of an oriented ellipse.

C. Proto-Objects in Sensory Memory

To be able to form stable object hypotheses, the sensory information needs to be buffered and organized consistently. This is done in form of proto-objects in sensory memory. Thus the blobs are mapped one-to-one to proto-objects in sensory memory which themselves contain the most recent blob data.

The incoming blob data has to be mapped to the proto-objects before the data can be stored consistently.

If the memory is empty, a proto-object is generated from blob data by simply assigning a unique identifier to the new proto-object and inserting the incoming blob data into it.

If the sensory memory already contains one or more proto-objects, a prediction for each proto-object is generated as a blob data. This predicted blob data is based on all blob data that is contained in the proto-object and is generated for the current time. The current implementation assumes noisy input and slow moving objects and therefore the prediction uses a weighted low-pass filter for the positions, orientations, and standard deviations (σ_{c1} , σ_{c2}).

Each incoming blob is either inserted into an existing proto-object or a newly generated one based on a minimum distance between incoming and predicted blob so that all incoming blobs are assigned unique identifiers. The metric for the distance computation is based on both euclidean distance and rotation angle. The blob orientation description is ambiguous with respect to 180 degree flips since it is derived from the principal axis of a 2d distribution. To be able to track the blob orientation, the orientation of newly inserted blob data is modified so that the orientation distance of the new blob is always less than or equal to 90 degrees.

Every time new incoming blob data is generated by the processing, its time label is compared to the time labels of the blob data inside all proto-objects and all blob data that are older than a certain threshold are deleted. In order to keep the system in a consistent state, this deletion is done even if the image processing does not find any blobs in the image pairs. If a proto-object does not contain any blob data it is also deleted from sensory memory.

D. Object Hypothesis Generation

The interaction system needs to evaluate the proto-objects in sensory memory using different criteria. These evaluations are also based on the blob data predictions of all proto-objects. The label of this prediction is set to "memorized" if the latest blob data in the proto-object is older than the prediction time. Otherwise, it is set to the label of the latest blob data in the proto-object.

A minimum criterion that is already sufficient for the behavior of fixation and tracking is a blob labeled as inaccurate. Any blob labeled as accurate can be used for an approaching behavior. If more severe criteria such as stable values σ_{c1} and σ_{c2} and a maximum distance to avoid relying on insufficient vision data are considered, stable object hypotheses can be

extracted. To implement manipulation behaviors like "poke balloon", additional constraints can be put on the stable object hypotheses. These could be roughly spherical shape ($(\sigma_{c1} - \sigma_{c2})/\sigma_{c1} < \text{threshold}$) and easiest execution of the behavior (minimum distance to a behavior specific reference point, e.g. for poking in front of the body). A behavior like "power grasp object" could require a minimum elongation for grasp stability ($(\sigma_{c1} - \sigma_{c2})/\sigma_{c1} > \text{threshold}$) and a suitable diameter ($\text{threshold} < \sigma_{c2} < \text{threshold}$).

The output from the sensory memory is therefore evaluated with respect to the object criteria (in this case distance, size and elongation) and the raw output together with the evaluation is sent to the 3 behaviors (search, track, reach) in a compact form. Each behavior can then easily extract the relevant information, since e.g. inaccurate blobs that do not math the object criteria can be tracked, but only stable elongated objects will be reached for.

E. Behavior Selection: Tracking and Searching

The output of the sensory memory is used to drive two different head behaviors: 1) searching for objects and 2) gazing at or tracking objects or blobs. Separate from these behaviors is a decision instance or *arbiter* that decides which behavior should be active at any time. The decision of the arbiter is solely based on a scalar value that the behaviors provide, which we call a *fitness value*, but which other authors [1], [2] refer to as *behavior value*. This fitness value describes how well a behavior can be executed at any time. In this concrete case tracking needs at least an inaccurate blob position to point the gaze direction at, but of course can also use a full object hypothesis. Thus the tracking behavior will output a fitness of 1 if any blob or object is present and a 0 otherwise. The search behavior has no prerequisites at all and thus its fitness is fixed to 1. For the case of this very simple behavior setup, the arbitration is of course trivial. However we implemented it as a competitive dynamical system similar to the one described in [13] for extensibility. Thus the arbiter uses the vector of fitness values from all behaviors, as an input to a competition dynamics that calculates an activation value for each behavior. The competition dynamics uses a pre-specified inhibition matrix that can be used to encode directed inhibition — behavior A inhibits behavior B but not vice versa — to specify behavior prioritization and even behavior cycles. In this case tracking is prioritized to searching by such a directed inhibition.

The search behavior is realized by means of a very low resolution (5 by 7) inhibition of return map with a simple relaxation dynamics. If the search behavior is active and new vision data is available it will increase the value of the current gaze direction in the map and select the lowest value in the map as the new gaze target. Additionally the whole map is subject to a relaxation to 0 and a small additive noise. This generates a visual search pattern with a random sequence of fixations that takes into account all visual information immediately and results in an efficient and fast finding of relevant objects. The size of the inhibition of return map is

derived from the field of view of the cameras relative to the pan/tilt movement range. Higher resolutions will not change the searching significantly. The relaxation time constant is set in the second range so that motion of the robot, which will effectively invalidate the inhibition map, is not a problem.

The tracking behavior is realized as a multi-tracking of 3-dimensional points. The behavior takes all relevant proto-objects and object hypotheses into account and calculates the pan/tilt angles for centering them in the field of view. Then a cost function with a trapezoidal shape in pan/tilt coordinates is used to find the pan/tilt angle that will keep the maximum number of objects in the effective field of view of the cameras and this is sent as the pan/tilt command. Since the tracking behavior always uses the stabilized output of the sensory memory the robot will still look at a certain position even if a blob disappears for a short time. This significantly improves the performance of the overall system.

The two visual interaction behaviors together with the arbiter switching mechanism show very short reaction times and have proven efficient to quickly find and track objects.

F. Strategy selection: Reaching

Similarly to the search and track behaviors, the reaching behavior is driven by the sensory memory. As shown in Fig. 1, the proto-object information is evaluated, and the position and orientation of the target points is sent to the reaching behavior. This behavior is composed of a set of internal predictors and a strategy selection instance. Each predictor includes a whole body motion controller and a cost function evaluation.

The underlying whole body control model is depicted in Fig. 3. The geometry and kinematic topology matches the humanoid robot ASIMO [14]. The first link corresponds to the heel coordinate system comprising three degrees of freedom. It is centered between the feet and aligned with the heel edge. Its degrees of freedom are translations in forward and lateral direction as well as a rotation about the vertical axis. The consecutive links correspond to the body segments of the robot. The pelvis is undergoing three translations and rotations with respect to the heel frame. The head is connected to the upper body with pan and tilt joints. Further, the two arms comprise 5 dof each. An additional coordinate system with some offset to the hand origin defines a hand reference point. All together, the model comprises 21 dof.

The robot motion is generated with the "redundancy resolution" framework by Liégeois [15], [16], [17] for redundant systems. The basis for the following calculations is the task Jacobian \mathbf{J} , which relates task to state velocities: $\dot{\mathbf{x}} = \mathbf{J}\dot{\mathbf{q}}$. A detailed derivation of the equations for the presented kinematic description is given in [18], [19]. The state motion rates are

$$\dot{\mathbf{q}} = \mathbf{J}^\# \dot{\mathbf{x}} - \alpha \mathbf{N} \mathbf{W}^{-1} \left(\frac{\partial H}{\partial \mathbf{q}} \right)^T \quad (1)$$

with

$$\begin{aligned} \mathbf{J}^\# &= \mathbf{W}^{-1} \mathbf{J}^T (\mathbf{J} \mathbf{W}^{-1} \mathbf{J}^T)^{-1} \\ \mathbf{N} &= \mathbf{E} - \mathbf{J}^\# \mathbf{J} \end{aligned} \quad (2)$$




Fig. 3. Kinematic model for whole body motion.

where $\mathbf{J}^\#$ is a weighted generalized pseudo-inverse of \mathbf{J} with weighting matrix \mathbf{W} . Matrix \mathbf{E} denotes the identity matrix. Scalar H is an arbitrary optimization criterion, whose gradient is mapped into the null space with projection matrix \mathbf{N} and scalar α defining the step width. In this work, a joint limit avoidance criterion is employed.

The control system allows to give commands in a highly flexible way. For this, a vector \mathbf{l} with the number of task variables is defined. Its elements define if a respective task variable is active ("1") or inactive ("0"). In the current implementation, the task vector is

$$\mathbf{x}_{task} = (\mathbf{x}_{ht,l} \ \varphi_{ht,l} \ \mathbf{x}_{ht,r} \ \varphi_{ht,r} \ \varphi_{head})^T \quad (3)$$

where \mathbf{x}_{ht} denotes the hand tip reference position of the respective hand and φ the orientation of the hands and the head. Based on vector \mathbf{l} , the equation system is set up according to the following procedure: Starting with the task element 0, the respective task Jacobian is added to the overall task Jacobian if the corresponding element of \mathbf{l} equals "1", otherwise it is skipped. Vector \mathbf{l} is updated every sampling interval through a command interface, so that the task vector can be changed dynamically at run-time.

Internal prediction architecture

In the following, a step from a single whole body controller towards a parallel simulation architecture that consists of several controllers will be made. The idea is to evaluate many strategies that solve the task in different ways. In the remainder and without loss of generality, the task of reaching towards




Fig. 4. “Floating” heel frame: The stance position of the feet is the result of the null space motion. It is a local optimum regarding the given task.




Fig. 5. Parallel simulation architecture.

an object and aligning the robot’s palm with the objects longitudinal axis will be regarded.

In a first step, the visual target is split up into different motion commands, with which the task can be achieved. Four commands are chosen: Reaching towards the target with the left and right hand, both while standing and walking. In the following, a motion control simulation relating to one of those commands will be called a “strategy”.

In the strategies that reach from a fixed stance, the degrees of freedom describing the heel coordinate system are constrained. For the strategies that involve walking, the kinematic constraints on the heel degrees of freedom are released. This leads to a “floating” heel frame that will converge to a position and orientation that is a local optimum with respect to the null space criterion (see Fig. 4). This leads to a very interesting property of the control scheme: The control algorithm will automatically compute the optimal stance position and orientation with respect to a given target. On the robot, the floating frame is set as the target for a step pattern generator, which generates appropriate steps to reach the computed heel position and orientation.

Now each strategy computes the motion and an associated cost according to its specific command. The cost describes the suitability of the strategy in the current context. It is composed of a set of penalties that will be described later. The costs are evaluated by the strategy selection process, and the strategy with the lowest cost is identified. The corresponding command is redirected to the physical robot (See Fig. 5). The robot is controlled with the identical whole body motion controller that is employed for the internal simulations.

An interesting characteristic of the system is the temporal




Fig. 6. “Ideal” (wireframe) and “clipped” (solid) state vectors lead to a task error that is used as a measure of reachability.

decoupling of real robot control and simulation. The strategies are sped up by a factor of 10 with respect to the real-time control, so that each strategy has converged to the target while the physical robot still moves. Therefore, the strategies can be regarded as prediction instances, since they look some time ahead of the real robot. Nevertheless, the control algorithms running within the strategies and on the robot are identical.

Selection mechanism

The selection of the most appropriate strategy is based on the evaluation of a multi-criteria cost function. The cost function encodes the following heuristics: Standing will be preferred over walking, walking over doing nothing. For this, the following criteria are incorporated in the cost function:

Reachability of the target: In the control algorithm, two state vectors are computed. One regards the system as “ideal”, ignoring joint limits. The state vector that is commanded to the physical robot is clipped so that the joint limits are not violated. In normal operation, both vectors are identical. If however, one or more joint limits are violated, a more or less large error between commanded and actual task will emerge. This is illustrated in Fig. 6. This error is used as a measure of reachability.

To prevent the strategy selection mechanism to oscillate between two adjacent solutions, the target is surrounded by two interval regions ξ_1 and ξ_2 , where $\xi_1 < \xi_2$. Now, the two following cases are regarded:

- 1) The hand has not reached the target: The target is considered as reached, if the effector gets inside ξ_1
- 2) The hand has reached the target: The target is considered as not reached, if the effector gets outside ξ_2

Postural discomfort: The weighted least squares distance of the joints to their center positions defines a “discomfort”

penalty. Whenever the target is moved from the one side to the other, this penalty will make the system switch the reaching arm. The hysteresis ensures that there is no oscillation between left and right arm when both penalties have the same value.

“Laziness”: Both walking strategies receive a constant penalty, so that standing will be preferred over walking.

Time to target: This penalty is added to the walking strategies. It is a measure for the estimated “steps to target”. This penalty makes the robot select the strategy that brings it towards stance with the minimum number of steps.

If a walking strategy has been selected, the real robot will start making steps toward the computed heel coordinate system. In this case, the standing strategies will not be considered, until the step pattern generator of the real robot has reached the converged target heel position.

G. Collision Detection

In order to ensure safety of the robot during operation, a real time collision detection algorithm is used. The collision detection uses an internal hierarchical description of ASIMO’s body in terms of spheres and sphere-swept lines that is used together with the kinematic information to calculate the distances between the segments (limbs and body parts) of the robot. If any of these distances falls below a threshold the high-level motion control will be disabled, so that only the dynamic stabilization of the bipedal walking is active. The collision detection acts as a last safety measure and is not triggered during normal operation of the robot.

Additionally a simple collision avoidance limits the position of all movement targets so that e.g. wrist target positions inside or very close to the body are never generated.

III. IMPLEMENTATION

The vision and control processing is divided up into several smaller modules that interact in a data driven way in the real time environment RTBOS [12]. This component based subdivision of processing is similar to that of other RTBOS based systems, for instance the Brain-like Active Sensing System BASS described in [11], [20] that combines active vision with object recognition. Since several components are shared between both systems, an additional object recognition can be easily added to the system described here.

For the implementation the system and work load was distributed over four different computers interconnected using fast ethernet connections. Onboard of the robot, the vision host is used to acquire images (half-frame NTSC) and send them to the offboard vision PC. The control host receives whole body motion targets, generates postures taking the collision detection into account, and sends the current posture every 5 ms to the offboard control PC.

All vision processing was done on images of 384 by 240 pixels, i.e. a horizontally scaled version of the NTSC images. This allowed faster stereo calculations for which the Small Vision System SVS software of SRI was used.

Two offboard PCs were used to split the computationally expensive vision calculations (from the raw images to the proto-objects in sensory memory) from the control part running




Fig. 7. Progression of fitness values over time.

many parallel components with small cycle times (processing from the proto-object evaluation to the motion targets).

IV. RESULTS

The system as described above was tested many times with different people interacting with ASIMO with a variety of target objects. The scenario was always to have a human interaction partner who had an elongated object that was shown or hidden in various ways to ASIMO.

The system shows a very good interaction performance. Objects are quickly found and reliably tracked even when moved quickly. The reaching behavior will reach for any elongated object of appropriate size that is presented within a certain distance — from about 20cm to about 3m. ASIMO will also track the object orientation correctly, e.g. if a bottle is presented first with the bottleneck up, ASIMO will reach with the proper hand orientation following the object orientation in real time even if the bottleneck points downwards. ASIMO will even continue to track the orientation correctly if the object is rotated to a horizontal or inverted position.

Fig. 7 shows the cost functions of the four moving strategies over time. The cost function with the lowest value is considered the best. The periods within which the respective strategy is active, is shown under the time axis. When moving the object from the left to the right, the robot will first (period 1) try to reach it with the right hand and then dynamically switches to the left hand. When the object gets out of reach, it will start walking and follow the object (period 3). In period 4, the robot walks and tries, at the same time, to reach the object with the right hand. If all costs are above the gray region, no command is given to the robot.

ASIMO switches between reaching with the right and left hand according to the relative object position with some hysteresis. Also walking is used only when necessary and the robot can be driven to show rather extreme postures when an object is presented accordingly (see image sequence). The whole body motion control generates smooth motions and is able to handle even the extreme postures which gives a very natural and human-like impression even to the casual observer.

The visual system is capable of running at 12.5 fps (80 ms loop time), its main limitations are the stereo disparity calculations (50 ms) and the transmission time of the color images via the fast ethernet connection (maximum 15 fps). The data connection between the onboard and offboard control hosts exchanged postures each 5 ms. The total reaction is mainly given by three parts: the latency between frame acquisition and the arrival of data on the offboard vision PC, the latency between its arrival and the start of the vision loop and the computation side of the vision loop. The latency of the control loop and the sending to the onboard host can be neglected. The total latency between time of acquisition and arrival of vision data on the offboard control PC was measured to be between 150 and 210 ms, i.e. roughly 2 frames.

Due to this latency and the fact that the field of view is not so large, motion speeds of the interaction object have to be limited, especially in the range where the robot is able to reach the object. Furthermore, the interaction object is restricted to have a narrow color distribution and no possible distractors should be too close and have similar colors.

Some limitations of the current system are that the reaching is done with a vertical offset of about 10 cm to avoid occlusions of the object by ASIMO's hand and the limited field of view of ASIMO's current cameras and the limited pan/tilt range which in this scenario limit the interaction range especially in vertical direction.

Obviously, since the object segmentation used is based on color, only certain objects can be used for interaction. Green objects were used mostly, but also tested with different colors like yellow and red - adjusting the color segmentation accordingly. These limitations can be reduced when incorporating a system like BASS [11], [20].

We had mixed results when using distractingly colored backgrounds, as most of the time the background was ignored since it was outside the interaction range, but in some cases the disparity calculation failed and produced spurious targets which confused the robot.

Overall we are satisfied with the performance and hope to build on this system in our future research.

V. OUTLOOK

The separation of the behavior selections is of a preliminary status, a unification of the selections will be made in future.

The collision detection scheme will be extended by a collision avoidance system that will be integrated into the whole body interaction mechanism.

Vision capabilities will be substantially researched and extended, building on the current concept of proto-objects. Furthermore we are interested in including a detection mechanism for support surfaces based on vision in order to enrich the interaction possibilities of the system. The presented selection mechanism is based on a pre-defined heuristics. Future work will focus on learning such selection criteria, and on making a step towards learning behaviors from demonstration.

ACKNOWLEDGMENT

The authors would like to thank the members of Honda's ASIMO and robot research teams in Japan for their support.

REFERENCES

- [1] Fujita, Takagi, and Hasegawa, "Ethological modeling and architecture for an entertainment robot," in *ICRA*, 2001.
- [2] —, "An ethological and emotional basis for human-robot interaction," *Robotics and Autonomous Systems*, no. 3-4, 2003.
- [3] C. Breazeal and B. Scassellati, "How to build robots that make friends and influence people," in *IROS*, 1999.
- [4] F. Orabona, G. Metta, and G. Sandini, "Object-based visual attention: a model for a behaving robot," in *CVPR*, 2005.
- [5] L. Natale, F. Orabona, F. Berton, G. Metta, and G. Sandini, "From sensorimotor development to object perception," in *Humanoids*, 2005.
- [6] R. A. Rensink, "Seeing, sensing, and scrutinizing," *Vision Research*, vol. 40, pp. 1469–1487, 2000.
- [7] A. Clark, "Feature-placing and proto-objects," *Philosophical Psychology*, no. 4, pp. 443–469, December 2004.
- [8] Z. W. Pylyshyn, "Visual indexes, preconceptual objects, and situated vision," *Cognition*, no. 1, pp. 127–158, June 2001.
- [9] R. A. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. M. Williamson, "The cog project: Building a humanoid robot," *Lecture Notes in Computer Science*, vol. 1562, pp. 52–87, 1999.
- [10] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Networks*, vol. 11, no. 7-8, pp. 1317–1329, 1998. [Online]. Available: citeseer.ist.psu.edu/wolpert98multiple.html
- [11] C. Goerick, H. Wersing, I. Mikhailova, and M. Dunn, "Peripersonal space and object recognition for humanoids," in *Proceedings of the IEEE/RSJ International Conference on Humanoid Robots (Humanoids 2005)*, Tsukuba, Japan, 2005.
- [12] A. Ceravola, F. Joubin, M. Dunn, J. Eggert, and C. Goerick, "Integrated research and development environment for real-time distributed embodied intelligent systems," in *International Conference on Intelligent Robots and Systems IROS*, 2006, submitted.
- [13] T. Bergener, C. Bruckhoff, P. Dahm, H. Janssen, F. Joubin, R. Menzner, A. Steinhage, and W. von Seelen, "Complex behavior by means of dynamical systems for an anthropomorphic robot," *Neural Networks*, 1999.
- [14] M. Hirose, Y. Haikawa, T. Takenaka, and K. Hirai, "Development of humanoid robot asimo," in *IEEE/RSJ International Conference on Intelligent Robots and Systems – Workshop 2, Hawaii, USA*, 2001.
- [15] A. Liégeois, "Automatic supervisory control of the configuration and behavior of multibody mechanisms," *IEEE Transactions on Systems, Man, and Cybernetics*, no. 12, December 1977.
- [16] J. D. English and A. A. Maciejewski, "On the implementation of velocity control for kinematically redundant manipulators," *IEEE Transactions on Systems, Man, and Cybernetics*, pp. 233–237, May 2000.
- [17] Y. Nakamura, *Advanced Robotics: Redundancy and Optimization*. Addison-Wesley Publishing, 1991.
- [18] M. Gienger, H. Janssen, and C. Goerick, "Task oriented whole body motion for humanoid robots," in *Proceedings of the IEEE-RAS/RSJ International Conference on Humanoid Robots*, 2005.
- [19] —, "Exploiting task intervals for whole body robot control," in *IEEE/RSJ*, 2006, submitted.
- [20] H. Wersing, S. Kirstein, M. Götting, H. Brandl, M. Dunn, I. Mikhailova, C. Goerick, J. Steil, H. Ritter, and E. Körner, "A biologically motivated system for unconstrained online learning of visual objects," in *Proc. Int. Conf. Art. Neur. Netw. ICANN*, 2006, submitted.