

Learning Coupled Dynamical Systems from Human Demonstration for Robotic Eye-Arm-Hand Coordination

Luka Lukic, José Santos-Victor and Aude Billard

Abstract—Efficient, adaptive and reliable visuomotor control system is crucial to enable robots to display flexibility in the face of changes in the environment. This paper takes inspiration in human eye-arm-hand coordination pattern to develop an equivalently robust robot controller. We recorded gaze, arm, hand, and trunk data from human subjects in reaching and grasping scenarios with/without obstacle in the workspace. An eye-arm-hand controller is developed, based on our extension of Coupled Dynamical Systems (CDS). We exploit the time-invariant properties of the CDS to allow fast adaptation to spatial and temporal perturbations during task completion. CDS global stability guarantees that the eye, the arm and the hand will reach the target in retinal, operational and grasp space respectively. When facing perturbations, the system can re-plan its actions almost instantly, without the need for an additional planning module. Coupling profiles for eye-arm and arm-hand systems can be modulated allowing to adjust the behavior of each slave system with respect to control signals flowing from the corresponding master system. We show how the CDS eye-arm-hand control framework can be used to handle the presence of obstacles in the workspace. The eye-arm-hand controller is validated in a series of experiments conducted with the iCub robot.

I. INTRODUCTION

Humans can adapt their visuomotor actions to various perturbations within milliseconds and also reuse and combine existing skills to be able to solve new and more challenging problems. A vast number of physiological studies reports on strong evidence that visual and motor systems mutually modulate actions, even when doing simple and well known tasks [1]. For instance, the gaze precedes the hand and locks on the contact points before the hand and the fingers close on the target [2], estimates motor goals for preparation of limb movements [3] and for movement initialization [4], and provides visual feedback for motion correction [5], etc.

Solutions to robotic visual-based reaching follow either of two well-established approaches: model-based control (open loop) [6] and visual servoing techniques (closed loop) [7], [8]. Model-based techniques are very dependent on the accuracy of the model. Besides, the eye-arm mapping is “rigid” (there is no way to modulate the profile of coordination) and the arm movement is not initiated before the gaze fixates the target. On the other hand, solutions that use visual servoing

are limited by the requirement for the simultaneous view of the end-effector and the target object.

In this paper, we extend the Coupled Dynamical Systems (CDS) framework for arm-hand coordination [9] to learn robotic eye-arm-hand coordination from human demonstration. Our approach contributes to the state of the art in visual-based reaching and grasping, embracing the best of the model-based and visual servoing techniques. It uses human demonstrations, which provide a model to guide dynamics of motion as in model-based control. Similarly to visual servoing, it performs close-loop control and hence guarantees to reach the target under perturbation. Hand control in robotic reaching and grasping usually operates separately from eye-arm control. Here we build a single eye-arm-hand controller, which drives synchronously and in coordination all three systems. Furthermore, we show how the model can be extended to handle the presence of obstacles. We estimate the consequences of future visuomotor actions by integrating the eye-arm CDS. The objects which are tagged as obstacles after propagating the forward model are treated as intermediary targets for the visuomotor system. The proposed approach is evaluated in robotic visuomotor manipulation experiments using the humanoid robot iCub [10].

The rest of the paper is structured as follows. Section II reviews related work on dynamical systems, eye-arm coordination and path planning. Section III describes the experimental setup, the procedure of data recording with human subjects, and provides an analysis of the human data. Section IV explains our approach to eye-arm-hand coordination, planning and obstacle handling. Section V reports validations of the approach in experiments on the iCub robot. In Section VI we conclude and propose future extensions of this work.

II. RELATED WORK

Methods for visually-guided reaching that build estimates of eye-arm transformation are known as model-based. One approach learns the model of transformation from the eye state (e.g. eye-neck angles) to the arm state (e.g. arm joints) [6]. This model is used in a straightforward manner, such that the arm state is computed by using the obtained mapping after the target is visually fixated. An alternative approach is to build a model that transforms the arm state to the eye state, and find the goal arm state as a solution to an optimization problem [8]. In general, model-based methods

Luka Lukic and Aude Billard are with Learning Algorithms and Systems Laboratory (LASA), EPFL, Switzerland: {luka.lukic, aude.billard}@epfl.

José Santos-Victor is with Computer Vision Laboratory (VISLAB), IST, Portugal: jasv@isr.ist.utl.pt.

are very appealing because of their simplicity and practical applications. However, these methods suffer from several drawbacks: the accuracy of the reaching movement is limited by the accuracy of the eye-hand mapping estimate. Moreover, the learned mapping profile cannot be modulated. Finally, the reaching path is generated by relying on interpolation between the starting arm state and the computed goal arm state. By using model-based approaches the arm movement is initiated after the gaze fixates the target. This differs from the natural profile of visuomotor coordination since reaching movements in humans are initialized before the gaze locks at the target [1].

Visual servoing approaches control the speed of the arm based on measurement of the visual error between the hand and the target [7], [8]. This approach ensures zero-error reaching, but it requires to have the target object and the hand simultaneously in the field of view. Visual servoing does not allow to produce a family of human-like motion profiles in reaching tasks.

Recent efforts combined model-based reaching and visual servoing [8]: the arm is driven by a model-based open loop controller until it reaches close to the target, and then visual servoing takes over the control of the arm motion, driving the arm towards the target. This method has shown good results in terms of the reaching accuracy. However, it relies on an external, adhoc mechanism to switch between the two modes.

Our recent work on CDS [9] has demonstrated benefits of learning explicitly the coupling between two dynamical systems over using a naive approach to model motions of both physical systems with a single dynamical system. The major strength of the CDS framework is due to the fact that it relies on encoding motion with time-invariant Dynamical Systems (DS) [11]. This approach ensures robustness against spatial and temporal perturbations, while producing trajectories similar to those demonstrated by humans. In this paper we extend the original CDS framework such that the dynamical system governing the gaze motion acts as the master controller to the arm dynamical system, and the arm dynamical system is the master to the dynamical system governing the hand. Next, we review briefly the most relevant approaches to obstacle avoidance in the context of visuomotor coordination.

Robots operating in cluttered environments have to be able to plan their motion avoiding collisions with objects in the workspace. Sampling-based algorithms [12] are currently one of the most reliable techniques for path planning. The time needed to plan motion with these techniques is tolerable for robotic applications where fast adaptation to perturbations is not critical for successful accomplishment of a task. Sampling-based algorithms cannot meet demands for rapid motion planning which humans perform almost effortlessly within a timescale of only a few milliseconds. In contrast, our approach uses the ability of Dynamical Systems to instantly re-plan motion with the presence of perturbations. Motion generation with time-invariant DS is an alternative

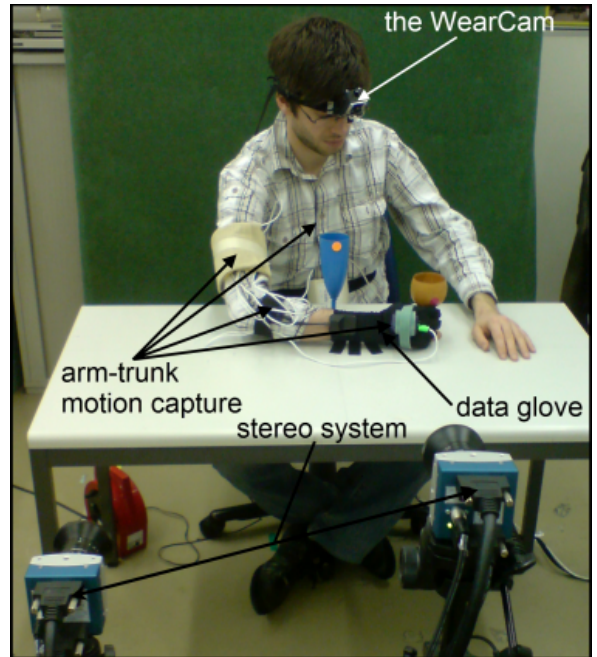


Figure 1. Experimental setup to record eye-arm-hand coordination from human demonstrations in grasping tasks with/without the presence of obstacles. In this trial, the human subject is grasping the target object (wine glass) avoiding the obstacle (champagne glass).

to classical path planning algorithms.

Approaches taken inspiration in neuroscience and on the role of cerebellum in planning for the timing of motion [13] state that forward models play a very significant role in motor control by predicting consequences of future actions. Our work exploits this biologically inspired notion of forward models in motor control, and uses a model of the dynamics of the reaching motion to predict collisions with objects in the workspace when reaching for the target object.

III. EXPERIMENTAL PROCEDURE AND DATA ANALYSIS

We first describe the experimental procedure to record data from human subjects when performing reach and grasp motions. We provide an analysis of visuomotor coordination in human subjects, which supports the development of similar coordinated patterns of motion in our robot control, which will be described in Section 4.

A. Methods and materials

Three right-handed subjects participated in this study. The experimental setup is illustrated in Fig. 1. The initial positions of the hands, the target object and the obstacle object were predetermined, as indicated by markers on the table. The wine glass is the object to be grasped (target) and the champagne glass is the object to be avoided (obstacle). Grasping during all trials was conducted with the right hand. The left hand is rested on the table, contributing to reduction of movements of the trunk. Initially, prior to the onset of motion, the subject's body is quasi stationary and the




Figure 2. Typical profiles of eye-arm coordination in human demonstrations. Distance from gaze to the target and from the arm to the target (calculated from the hand center) in: obstacle (a) and no-obstacle (b) scenarios. Distance from gaze to the obstacle and from the arm to the obstacle (c).

subject is instructed to look at his/her hand. A sound signal indicates the start of a trial. In the absence of the obstacle, the subject moves the hand towards the target object and grasps it. A trial is finished when the target is grasped. If the obstacle is present, the procedure is generally the same, but with obstacle avoidance during manipulation. Each subject performed the tasks eight times for both obstacle and no-obstacle conditions in randomized ordering.

A head-mounted eye-tracker designed in our lab, the WearCam system [14], was used for gaze tracking and for recording the scene in the subject’s field of view. Video and gaze position from the WearCam were recorded at 25 Hz. XSensTM IMU motion capture system was used for recording of trunk motion and the arm motion. It provided information about three joints of trunk motion (roll, pitch and jaw), three joints that modeled shoulder (flexion-extension, abduction-adduction and circumduction), two joints in the elbow (flexion-extension and pronation-supination) and two wrist angles (abduction-adduction and flexion-extension). 5DTTM data glove was used for recording of finger joint angles. Data from XSensTM IMU motion capture sensors and 5DTTM data glove were recorded at 25 Hz. Stereo system was used for tracking of a 3D position of the hand and the objects on the scene. The speed of data recording from the stereo camera was 100 Hz.

All signals were filtered with a preprogrammed peak-removal technique, and with a moving average filter. Finally, piecewise spline fitting was done, which did additional smoothing as well. All data were synchronized and re-sampled at 25 Hz.

B. Analysis of recordings from human trials

We started from the hypothesis that the eyes precede the arm’s motion, so as to guide planning of arm transport. There is ample evidence of such saccadic eye motion toward the target during reaching; see e.g. [1], [15]. To determine whether the same mechanism was at play in our experiment, we plotted the relative displacement of the eyes compared to the arm approaching the target in both the obstacle and without obstacle condition, see Fig. 2(a) and (b).

Since we did not control for the timing of motion during our trials, we normalized the duration of the motion in each trial. We see that the eyes precede motion of the arm in both the obstacle and no-obstacle scenarios. When the obstacle is present the gaze reaches the target object at $44.45\% \pm 9.55\%$

of the total time needed for the arm to reach the target object. When the obstacle is not present the gaze reaches the target object at $46.69\% \pm 14.73\%$ of the total duration of the motion. There is a systematic phase lag between the eyes and the arm, which suggest that eye and arm motion towards the target object are synchronized even when the obstacle is present. We observed large variation in the time to complete motion (std. of 12.35% of the mean time to complete motion in the obstacle case and 13.62% in the no-obstacle case). This supports the hypothesis that this eye-arm coordination is due to a coupling between the two systems that is independent of the time to perform the motion. Two-way ANOVA¹ (factors: subjects, obstacle presence) yields a significant effect for factor of subjects ($F(2, 47) = 6.99$, $p = 0.0024$), but there is no statistically significant effect of the presence of the obstacle ($F(1, 47) = 0.58$).

To see whether the same mechanism was also at play when the arm passes the obstacle, we plotted the relative path of the eyes and arm motion with respect to the obstacle, see Fig. 2(c). We see that the gaze reaches the obstacle at $48.33\% \pm 14.51\%$ of the total time needed for the arm end-effector to reach the closest distance to the obstacle object. The lag between gaze and the arm with respect to the obstacle is almost the same as the lag between gaze and the arm with respect to the target. This observation suggests that the eyes and the arm are driven to both the obstacle and the target by the same mechanism of coordination, and that the obstacle acts as an intermediary target for both the eyes and the arm when moving towards the final target.

In summary, the mechanism of the eyes leading the arm is observed in both obstacle and no-obstacle trials. The phase lag between the eyes and the arm with respect to the target object is kept in both obstacle and no-obstacle scenarios, indicating a high degree of synchronization of the eye-arm manipulation. It can be seen that the lag between gaze and the arm with respect to the obstacle is very close to the value of the lag between gaze and the arm with respect to the target. Such evidence of the same coordinated pattern of eye-arm motion when reaching the obstacle and the target, supports the development of a biologically inspired controller

¹ANOVA (analysis of variance) is a statistical method which compares the variances around two or more means, to determine whether significant differences exist between distinct conditions of an experiment. See [16] for more.

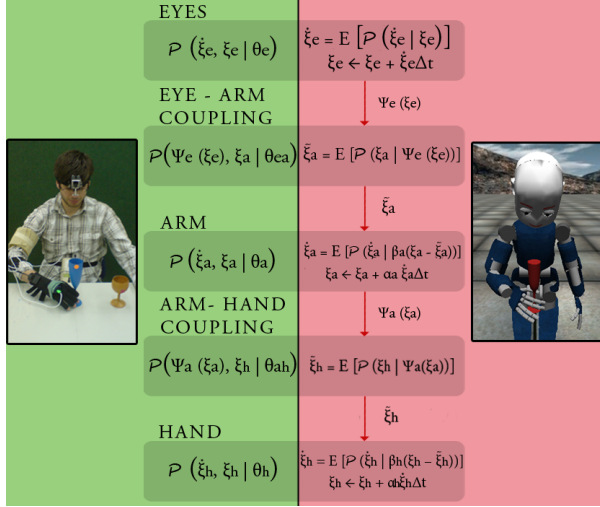


Figure 3. CDS-based robotic eye-arm-hand coordination. Left (green) part of the figure shows how the CDS model is learned. Reproduction of motion on the robot is shown on the right side of the figure (red part). CDS consists of five building “blocks”: three dynamical systems (the eyes, the arm and the hand) and two coupling models: eye-arm coupling and arm-hand coupling.

for coordinating eye-arm-hand in robots, which we describe next.

IV. CDS EYE-ARM-HAND CONTROL AND OBSTACLE HANDLING

In the first part of this section, we introduce the principle of robot control using time-invariant DS and the probabilistic approach taken to estimate the parameters of the system. Extension of this formulation to model and control coupled dynamics is presented. Furthermore, we show how the basic model of eye-arm-hand coordination in the obstacle-free grasping can be extended to handle the obstacle in the workspace.

A. Single Dynamical System and GMM/GMR

The motion of our system is represented through the state variable $\xi \in \mathbb{R}^d$. N recorded demonstrations of the task yield the data set $\{\xi_t^n, \dot{\xi}_t^n\}$, $\forall t \in [0, T_n]$; $n \in [1, N]$, of robot’s states and state derivatives at particular time steps t , where T_n is the number of samples in the n -th demonstration. We posit that recorded data samples are instances of motion that is governed by a first-order autonomous differential equation:

$$\dot{\xi} = f(\xi) + \epsilon \quad (1)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a continuous and continuously differentiable function, with a single equilibrium point $\xi^* = f(\xi^*) = 0$. ϵ is a zero-mean Gaussian noise. The noise term encapsulates both sensor inaccuracies and errors inherited from human demonstrations. Time-invariance provides inherent robustness to temporal perturbations. In order to achieve robustness to spatial perturbations, the robot’s state variable ξ is represented in the target’s reference frame.

We use Gaussian Mixture Model (GMM) to encode dynamics in a probabilistic framework. GMM define a joint

probability distribution function $\mathcal{P}(\xi_t^n, \dot{\xi}_t^n)$ over the set of data from demonstrated trajectories, as a mixture of K Gaussian distributions (with π^k , μ^k and Σ^k being the prior probability, the mean value and the covariance matrix of the k -th Gaussian, respectively):

$$\mathcal{P}(\xi_t^n, \dot{\xi}_t^n) = \sum_{k=1}^K \pi^k \mathcal{N}(\xi_t^n, \dot{\xi}_t^n; \mu^k, \Sigma^k), \quad (2)$$

where each Gaussian probability distribution is defined as:

$$\mathcal{N}(\xi_t^n, \dot{\xi}_t^n; \mu^k, \Sigma^k) = \frac{1}{\sqrt{(2\pi)^{2d} |\Sigma^k|}} e^{-\frac{1}{2}([\xi_t^n, \dot{\xi}_t^n] - \mu^k)^T (\Sigma^k)^{-1} ([\xi_t^n, \dot{\xi}_t^n] - \mu^k)}, \quad (3)$$

where the mean and the covariance matrix are defined as:

$$\mu^k = \begin{pmatrix} \mu_{\xi}^k \\ \mu_{\dot{\xi}}^k \end{pmatrix} \text{ and } \Sigma^k = \begin{pmatrix} \Sigma_{\xi\xi}^k & \Sigma_{\xi\dot{\xi}}^k \\ \Sigma_{\dot{\xi}\xi}^k & \Sigma_{\dot{\xi}\dot{\xi}}^k \end{pmatrix}. \quad (4)$$

Stable Estimator of Dynamical Systems (SEDS) [11] is used to compute the GMM parameters. SEDS ensures global stability of the noise-free estimate of the underlying dynamics, which will be denoted as \hat{f} .

Taking the posterior mean estimate of $\mathcal{P}(\xi_t^n | \xi_t^n)$ yields an estimate of $\hat{\xi} = \hat{f}(\xi)$, function that approximates the model dynamics:

$$\hat{\xi} = \sum_{k=1}^K \frac{\pi^k \mathcal{N}(\xi; \mu^k, \Sigma^k)}{\sum_{i=1}^K \pi^i \mathcal{N}(\xi; \mu^i, \Sigma^i)} \left(\mu_{\xi}^k + \Sigma_{\xi\xi}^k (\Sigma_{\xi\xi}^k)^{-1} (\xi - \mu_{\xi}^k) \right). \quad (5)$$

B. Coupled Dynamical Systems

Our recent work [9] has shown the benefits of learning explicitly a coupling between the arm DS and the finger DS over modeling motions of the physical systems with a single, extended DS. The problem associated with learning one high-dimensional dynamical model that guides motion of two physical systems is the fact that explicit following of correlations shown in demonstrations between the two coupled dynamics is not guaranteed. This could be a problem if the robot is perturbed far from the region of demonstrated motion, since the behavior of the dynamical systems may not be correctly synchronized. An approach adopted in [9], is to learn separately two dynamics and then learn a coupling between them. This approach assures that CDS follows the same synchronized pattern of motion as observed in human demonstrations.

1) *Extended CDS architecture and learning*: We extend the original CDS architecture, having in total five building blocks: three dynamical systems and two coupling blocks between them. They are organized in the following order: eyes \rightarrow eye-arm coupling \rightarrow arm \rightarrow arm-hand coupling \rightarrow hand. The gaze DS is the master to the arm DS, and the arm DS is the master to the hand DS. There is a

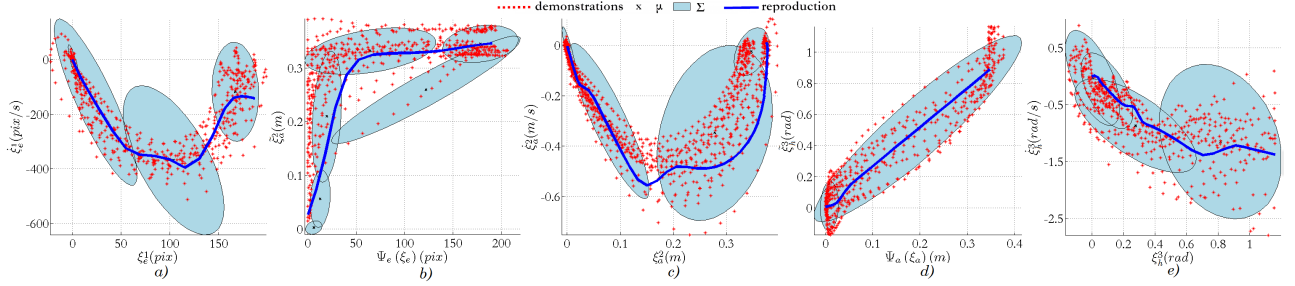


Figure 4. Learned CDS eye-arm-hand coordination model: a) eye dynamics, b) eye-arm coupling, c) arm dynamics, d) arm-hand coupling and e) hand dynamics. For simplicity of graphical representation, we plotted CDS model for one gaze position, one arm position and one hand position. The eye state is presented with horizontal gaze coordinate, denoted as ξ_e^1 . The arm state is presented with Cartesian coordinate that corresponds to the direction of the major hand displacement in the task, denoted as ξ_a^2 . The hand state is represented with thumb proximal joint, denoted as ξ_h^3 . Superposed to the datapoints, we see the regression signal (plain line) and the different Gaussian distributions (elliptic envelopes).

coupling block between each master and its slave. The major assumption is that the modulation signals between them flow only in the direction from the master to the corresponding slave, i.e. dynamics of the slave is modulated with control signals coming from its master, not vice versa. The master system evolves independently of its slave. Fig. 3 illustrates the architecture of CDS, and the principles of learning and reproduction of coordinated motion.

The state of the eyes is denoted with $\xi_e \in \mathbb{R}^2$, the state of the arm is $\xi_a \in \mathbb{R}^3$, and the state of the hand is $\xi_h \in \mathbb{R}^9$. The eye state ξ_e is represented as the distance between the position of the gaze and the position of a visual target in retinal coordinates (i.e. retinal error). The arm state ξ_a is represented as the distance in Cartesian coordinates between the arm end-effector (the palm center) and the target object. The hand state ξ_h is expressed as the difference between the current hand configuration and the goal hand configuration, i.e. grasp configuration. In other words, the attractors are placed at the target projection in the retinal plane, its Cartesian position in the workspace and at the corresponding hand configuration when the target is reached, which is formally expressed as: $\xi_e^* = 0$, $\xi_a^* = 0$ and $\xi_h^* = 0$. The arm state vector can be extended to encompass both position and orientation variables of the end-effector, which could be beneficial for more complex coordinated tasks.

CDS model of eye-arm-hand coordination is built in the following manner: joint probability distributions that encode the eye dynamics $\mathcal{P}(\dot{\xi}_e, \xi_e | \theta_e)$, the arm dynamics $\mathcal{P}(\dot{\xi}_a, \xi_a | \theta_a)$ and the hand dynamics $\mathcal{P}(\dot{\xi}_h, \xi_h | \theta_h)$ are first learned separately. Then we learn the joint distribution for eye-arm coupling $\mathcal{P}(\Psi_e(\xi_e), \xi_a | \theta_{ea})$ and arm-hand coupling $\mathcal{P}(\Psi_a(\xi_a), \xi_h | \theta_{ah})$, where $\theta_e, \theta_a, \theta_h, \theta_{ea}$ and θ_{ah} denote the GMM parameters, and $\Psi_e(\xi_e)$ and $\Psi_h(\xi_h)$ denote the coupling functions. GMMs that encode dynamics of the eyes, the arm dynamics and the hand dynamics are learned using the SEDS algorithm, for more details [11]. GMMs that model coupling are learned with Expectation-Maximization (EM) algorithm [17].

Two open parameters, α and β allow additional fine-tuning of the characteristics of the slave response (a and h

subscripts denote whether they modulate arm or hand motion, respectively). The speed is modulated with the scalar α , and the amplitude of motion is tuned by changing the value of the scalar β . Fig. 4 illustrates CDS model learned from demonstrations.

Algorithm 1 CDS eye-arm-hand coordination

do

 General:

- query frames from cameras
- read the current hand position from forward kinematics
- read the hand joints from encoders
- recognize and segment the target object
- estimate the position of the target in both retinal and Cartesian coordinates
- compute ξ_e , ξ_a and ξ_h

 Gaze:

 if gaze is not at target then

$$\dot{\xi}_e \leftarrow \mathbb{E} \left[\mathcal{P} \left(\dot{\xi}_e | \xi_e \right) \right]$$

$$\xi_e \leftarrow \xi_e + \dot{\xi}_e \Delta t$$

- look at new gaze point

 end if

 Eye – arm coupling:

$$\tilde{\xi}_a \leftarrow \mathbb{E} \left[\mathcal{P} \left(\xi_a | \Psi_e \left(\xi_e \right) \right) \right]$$

 Arm:

 if the arm is not at target then

$$\Delta \xi_a \leftarrow \xi_a - \tilde{\xi}_a$$

$$\dot{\xi}_a \leftarrow \mathbb{E} \left[\mathcal{P} \left(\dot{\xi}_a | \beta_a \Delta \xi_a \right) \right]$$

$$\xi_a \leftarrow \xi_a + \alpha_a \dot{\xi}_a \Delta t$$

- solve inverse kinematics

- move the arm and the torso to new joint conf.

 end if

 Arm – hand coupling:

$$\tilde{\xi}_h \leftarrow \mathbb{E} \left[\mathcal{P} \left(\xi_h | \Psi_a \left(\xi_a \right) \right) \right]$$

 Hand:

 if the hand is not at target then

$$\Delta \xi_h \leftarrow \xi_h - \tilde{\xi}_h$$

$$\dot{\xi}_h \leftarrow \mathbb{E} \left[\mathcal{P} \left(\dot{\xi}_h | \beta_h \Delta \xi_h \right) \right]$$

$$\xi_h \leftarrow \xi_h + \alpha_h \dot{\xi}_h \Delta t$$

- move the hand to new joint conf.

 end if

until object grasped

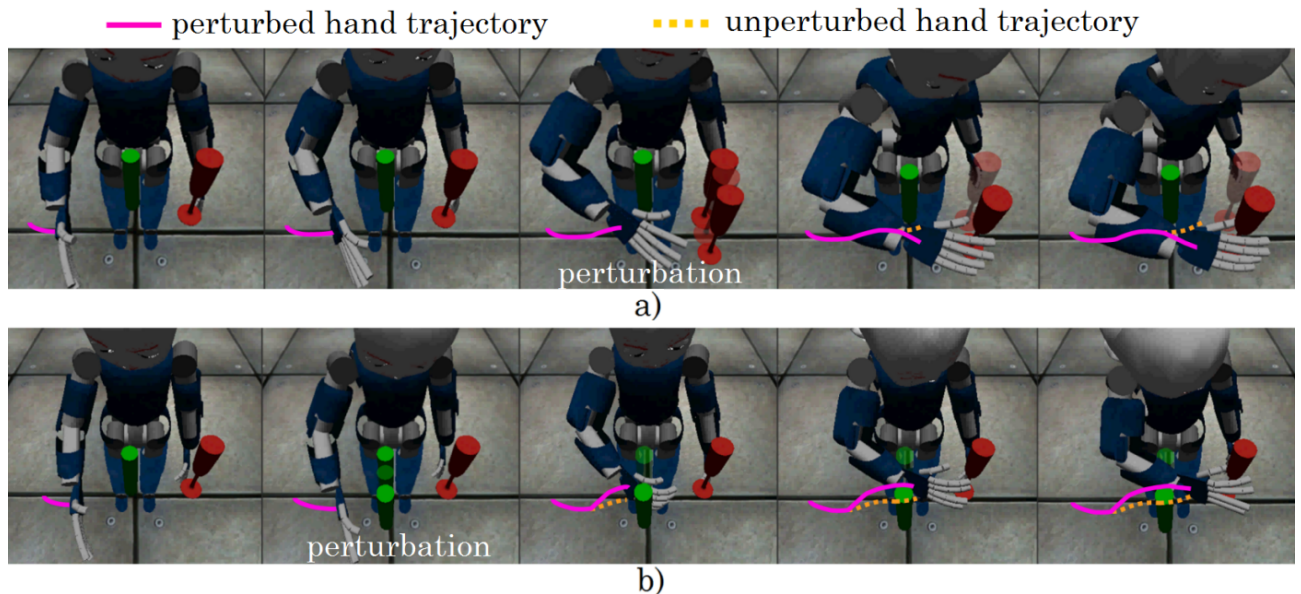


Figure 5. Experiments of visually-guided reaching and grasping in the iCub’s simulator, with the presence of the obstacle and with perturbations. The obstacle is an intermediary target for visuomotor system, hence obstacle avoidance is divided in two sub-tasks: from the start position to the obstacle (via-point), and from the obstacle to the grasping object. Figures show execution of eye-arm-hand coordination from the start of the task (left) until the successful grasp completion (right). Figures in the upper row (a) present scenario when the target object (red champagne glass) is perturbed during motion (perturbation occurs in the third frame from left). Visuomotor coordination when the obstacle is perturbed during manipulation is shown in the bottom row (perturbation in the second frame). The orange line shows the trajectory of the hand if perturbation did not happen, and the purple line is the actual trajectory of the hand from the start of unperturbed motion, including the path of the hand after perturbation until successful grasping. In both scenarios (target perturbed and obstacle perturbed), visuomotor system instantly adapts to the perturbation, and drives motion of the eyes, the arm and the hand to new position of the object.

2) *CDS reproduction*: The DS that drives the eyes evolves independently in time, and leads the whole system. The eye state velocity $\dot{\xi}_e$ is generated by conditioning the learned distribution of eye dynamics on the current eye state, as shown in Eq. 5. The eye state variable is incremented by adding the velocity multiplied by the time step Δt to its current value ξ_e . The eye system modulates the behavior of the arm system via the coupling function Ψ_e . The “desired” arm state value $\tilde{\xi}_a$ is inferred from the eye-arm coupling model by conditioning on $\Psi_e(\xi_e)$. The arm velocity $\dot{\xi}_a$ is generated by conditioning the distribution of arm dynamics on the difference between the current arm state and the “desired” arm state.

The desired hand state value $\tilde{\xi}_h$ is obtained by conditioning arm-hand coupling model on $\Psi_a(\xi_a)$. Finally, the hand velocity $\dot{\xi}_h$ is inferred by conditioning the hand model on the difference between ξ_h and $\tilde{\xi}_h$. Algorithm 1 shows how the robotic eye-arm-hand coordination is performed with CDS.

3) *Eye-arm-hand coordination for obstacle avoidance*: The use of a CDS-based visuomotor controller for obstacle avoidance is grounded on our hypothesis that the obstacle acts as an intermediary target for visuomotor system in reaching and grasping tasks, see Section III.

In order to define which objects in the workspace are obstacles for the realization of the intended reach-and-grasp tasks, we use a planning scheme to estimate consequences of future actions. More specifically, the motion of the arm towards the target is estimated by integrating the dynamics

of the extended CDS until each DS reaches its attractor. We integrated only the eye-arm part of the whole CDS, ignoring the hand’s DS, since our collision checking scheme is relatively simple. The arm end-effector is modeled as a point which moves along the estimated trajectory, and each object in the workspace is modeled as a cylinder. An object is tagged as an obstacle when the trajectory of the end-effector intersects with a cylinder modeling the object (certain collision), or when the cylinder lies within the area where it is very likely that it may collide with the forearm (very likely collision). For motions we consider here and observing the iCub’s body, this area is defined as the slice of the workspace enclosed by the estimated trajectory of the end-effector and the coronal plane of the body.

As suggested earlier on, we consider the eye-arm-hand coordination as a composition of two segments: motion from the starting position towards the obstacle and from the obstacle towards the target object. Individual segments of coordinated motion (from the starting point to the obstacle, and from the obstacle to the target) are performed in a manner presented in Algorithm 1. In the first part of the task, the arm DS moves under the influence of the attractor placed at the via-point. The hand DS is driven by the attractor placed at the hand configuration when the palm reaches the closest point (along the trajectory computed ahead of time) to the obstacle. The position of the via-point is determined with respect to the obstacle, such that its displacement vector from the obstacle position is oriented in either anterior or

ventral direction, for the length that corresponds to some safety distance d_{safety} between the palm and the obstacle. The direction of a displacement of the via-point (anterior or ventral) is chosen to correspond to a side of the obstacle where a collision is estimated to occur. In the second part of the task, after the obstacle is passed, CDS is driven towards the object to be grasped. Hand adaptation with respect to the obstacle provides a configuration of the hand such that collisions between the fingers and the obstacle are eluded during obstacle avoidance manipulation, whereas hand adaptation with respect to the target assures coordinated and stable grasping of the target as the arm reaches it.

The arm end-effector passing through the via-point at d_{safety} from the obstacle and hand adaptation with respect to the obstacle ensures that the hand will not collide with the obstacle. However, the end-effector obstacle avoidance has no means to modulate the arm (and torso) posture such that collisions with the remaining arm links are avoided. We benefit from controlling the arm in Cartesian coordinates, and having an efficient inverse kinematics (IK) solver [18] which is able to handle two tasks: to find suitable joint configuration (primary task) while keeping solutions as close as possible to a desired arm rest posture (secondary task). By having IK that can solve for reaching Cartesian position trying to keep joints close to a given rest posture, one can modulate the robot’s motion in operational space by providing joint rest postures suitable for obstacle avoidance. Our approach to the problem of finding suitable joint postures is to learn these joint postures from human demonstrations, since human demonstrations in obstacle avoidance tasks encode inherently favorable joint configurations.

Here we learn correlations between those joints, which provide major contributions in obstacle avoidance manipulation, and arm position in the operational space. The joints, for which we observed statistically significant differences in trunk and shoulder joint effort² with respect to the factor of obstacle presence, define the rest position, namely: torso pitch and yaw, and shoulder joints corresponding to adduction-abduction and flexion-extension. Hence, we proceed with learning the joint probability distribution $\mathcal{P}(q, x)$, where $q \in \mathbb{R}^4$ denotes the joint rest posture and $x \in \mathbb{R}^3$ denotes the Cartesian position of the arm.

Adaptation of the arm posture for obstacle avoidance is done in the following manner. When performing reaching for a visuomotor target (the obstacle object or the grasping object) the CDS system infers the state velocities, as explained earlier. By integrating the arm velocity, a new arm state is obtained. By taking the posterior mean estimate of $\mathcal{P}(q | x)$, we infer the favorable rest posture. Finally, the IK solver optimizes for joint angles that correspond to the desired Cartesian position, while trying to keep the four joints as close as possible to the suggested values from the model.

²Joint effort is measured as the amplitude of recorded joint motion.




Figure 6. Visuomotor adaptation to perturbation during the task, generated by a sudden displacement of the target object. The upper part of the graph shows how the eye state variable, represented with ξ_e^1 , adapts to perturbation. The middle graph part of the graph shows the arm state variable denoted with ξ_a^2 , and the lower part shows the hand state variable ξ_h^3 . Gaze DS adapts independently to spatio-temporal perturbations, whereas DS guiding arm motion is modulated via the coupling function $\Psi_e(\xi_e)$, and arm motion modulates hand DS via $\Psi_a(\xi_a)$. The figure shows that all three systems successfully reach the target when perturbed.

V. VALIDATION OF THE MODEL FOR ROBOT CONTROL

The CDS model is learned using the data gathered during the human trials, described in Section III. The parameters of SEDS algorithm (i.e. maximum number of iterations, optimization criterion, etc) and a number of Gaussians in mixtures are determined by using a grid-search with 10-fold crossvalidation on RMSE between the recorded motion and retrieved trajectories from the model. In all our experiments, we used $\Psi_e(\xi_e) = \|\cdot\|$, $\Psi_a(\xi_a) = \|\cdot\|$, and the values of parameters α_a , α_h , β_a and β_h were set to 1.

We validate the ability of the CDS controller on the iCub robot to reproduce the same task of visually guided obstacle-free reaching and grasping to the one that humans performed in our trials, together with the advocated robustness of the model to perturbations and the ability to handle the obstacles in the workspace. In each run, the object to be grasped is placed at a randomly computed position within a 15 cm-cube in the workspace. Fig. 5 shows performed coordinated manipulation in obstacle scenario with sudden perturbations of the target object and the obstacle, respectively. The robot’s end-effector avoids the obstacle during reaching for grasping in two task segments: 1. start position \rightarrow via-point at d_{safety} from obstacle, and 2. via-point at d_{safety} from obstacle \rightarrow grasping object. This safety distance in our human trials is $d_{safety} = 0.142 \pm 0.01$ m. We rescaled the safety distance from human trials by 2, since the dimensions of the iCub are similar to those of a 3.5 year old child, hence it has smaller workspace than our adult subjects. Once the obstacle has been reached, the target for the visuomotor system is

changed, and the eye-arm-hand motion is directed to the object to be grasped. The IK solver adapts the arm rest posture to be as close as possible to the output inferred from the model learned from human demonstrations.

Since the eye state is the distance between the position of gaze and the position of a visual target in retinal coordinates, and the arm state is represented with respect to the position of the object in Cartesian space, both variables are instantly updated when the perturbation occurs, see Fig. 6. DS of the eyes adapts independently to the perturbation. Behavior of DS of the arm is modulated via eye-arm coupling function, and hand DS is modulated via arm-hand coupling. Such modulation ensures that the learned profile of eye-arm-hand coordination will be preserved, and that the hand will re-open as the object is perturbed away from it. Besides the anthropomorphic profile of visuomotor coordination, the gaze-arm lag allows enough time to foveate at the object, re-estimate object's pose and compute suitable grasp configuration for the hand before it approaches too close to the object. Results presented here, with several additional experiments in simulation and experiments with the real robot are recorded in the video submitted with this paper.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we presented a CDS approach for learning robotic eye-arm-hand coordination from human demonstration. This approach harvests the major benefits of encoding motion with time-invariant DS: robustness to spatio-temporal perturbations and instant re-planning of motion when perturbations occur. We show that the CDS framework makes use of learned motion constraints as model-based visuomotor coordination does and assures parametrized reaching of the target with the zero-error like visual-servoing techniques.

We used forward planning by integrating the CDS model in order to estimate which objects collide with the arm when reaching for the target. The extension of the eye-arm-hand coordination for obstacle avoidance is based on our hypothesis that the obstacle acts as an intermediary target in reach-and-grasp motions. It is important to mention that our obstacle avoidance scheme does not have the full strengths of methods like RRTs for solving reaching in very complex workspaces, but it endows the visuomotor system with instant reactions to perturbations, providing means for rapid handling of a relatively simple obstacle in the workspace.

Our approach relies on using an efficient IK solver, calibrated stereo cameras and reliable encoder readings, but these days the demands for the aforementioned functionalities are not very difficult to achieve, especially when using robotic platforms like the iCub.

The flow of modulating signals in our controller is mono-directional, from the eyes to the arm, and from the arm to the hand. There is, however, strong evidence of shared control between the different modalities in humans, meaning that control signals also flow from the hand to the eyes, not only in the opposite direction [19]. We will investigate potential

benefits of such bidirectional visuomotor control schemes for robots in our future work.

VII. ACKNOWLEDGMENTS

This work was supported in part by EU Project First-MM (FP7/2007-2013) under grant agreement number 248258 and FCT doctoral grant (SFRH/BD/51072/2010).

REFERENCES

- [1] M. Hayhoe, A. Shrivastava, R. Mruczek, and J. Pelz, "Visual memory and motor planning in a natural task," *Journal of Vision*, vol. 3, no. 1, 2003.
- [2] A. Brouwer, V. Franz, and K. Gegenfurtner, "Differences in fixations between grasping and viewing objects," *Journal of Vision*, vol. 9, no. 1, 2009.
- [3] C. Prablanc, J. E. Echallier, M. Jeannerod, and E. Komilis, "Optimal response of eye and hand motor systems in pointing at a visual target. ii. static and dynamic visual cues in the control of hand movement." *Biological Cybernetics*, vol. 35, no. 3, pp. 183–187, 1979.
- [4] C. Prablanc, D. Pelisson, and M. Goodale, "Visual control of reaching movements without vision of the limb," *Experimental Brain Research*, vol. 62, no. 2, pp. 293–302, 1986.
- [5] H. Zelaznik, B. Hawkins, and L. Kisselburgh, "Rapid visual feedback processing in single-aiming movements." *Journal of Motor Behavior*, 1983.
- [6] L. Natale, G. Metta, and G. Sandini, "A developmental approach to grasping," in *Developmental Robotics AAAI Spring Symposium*, vol. 44, 2005.
- [7] F. Chaumette and S. Hutchinson, "Visual servoing and visual tracking," *Handbook of Robotics*, pp. 563–583, 2008.
- [8] L. Natale, F. Nori, G. Sandini, and G. Metta, "Learning precise 3d reaching in a humanoid robot," in *Development and Learning, 2007. ICDL 2007. IEEE 6th International Conference on*. IEEE, 2007, pp. 324–329.
- [9] A. Shukla and A. Billard, "Coupled dynamical system based arm-hand grasping model for learning fast adaptation strategies," *Robotics and Autonomous Systems*, 2011.
- [10] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. Von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor *et al.*, "The icub humanoid robot: An open-systems platform for research in cognitive development," *Neural Networks*, vol. 23, no. 8-9, pp. 1125–1134, 2010.
- [11] S. Khansari-Zadeh and A. Billard, "Learning stable nonlinear dynamical systems with gaussian mixture models," *IEEE Transactions on Robotics*, vol. 27, no. 5, pp. 943–957, 2011.
- [12] J. Kuffner Jr and S. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 2. IEEE, 2000, pp. 995–1001.
- [13] D. Wolpert, R. Miall, and M. Kawato, "Internal models in the cerebellum," *Trends in cognitive sciences*, vol. 2, no. 9, pp. 338–347, 1998.
- [14] B. Noris, J. Keller, and A. Billard, "A wearable gaze tracking system for children in unconstrained environments," *Computer Vision and Image Understanding*, 2010.
- [15] R. Johansson, G. Westling, A. Bäckström, and J. Flanagan, "Eye-hand coordination in object manipulation," *the Journal of Neuroscience*, vol. 21, no. 17, pp. 6917–6932, 2001.
- [16] D. Montgomery and G. Runger, *Applied statistics and probability for engineers*. Wiley, 2010.
- [17] C. Bishop, "Pattern recognition and machine learning (information science and statistics)," *Pattern Recognition*, vol. 4, no. 2, 2007.
- [18] U. Pattacini, F. Nori, L. Natale, G. Metta, and G. Sandini, "An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 2010, pp. 1668–1674.
- [19] S. Neggers and H. Bekkering, "Ocular gaze is anchored to the target of an ongoing pointing movement," *Journal of Neurophysiology*, vol. 83, no. 2, pp. 639–651, 2000.