

Learning postures through sensorimotor training: a human simulation case study

Belonia Gabalda Lionel Rigoux Olivier Sigaud

Institut des Systèmes Intelligents et de Robotique

Université Pierre et Marie Curie - Paris 6, CNRS FRE2507

4 Place Jussieu, F-75252 Paris Cedex 05, France

Belonia.Gabalda@laposte.net, Lionel.Rigoux@laposte.net, Olivier.Sigaud@lip6.fr

Abstract

In this paper, we present a case study showing why the modeling effort of computational neuro-sciences must call upon interactions with an environment when developmental phenomena are investigated. More specifically, in the context of a model of the spatial organization of a repertoire of postures, we show that calling upon a realistic simulation of the human kinematics (i) results in raising the question of the nature of information encoded in this repertoire ; and (ii) reveals that some previous assumptions about the functional organization of this repertoire were not necessary. Finally, we point out that conversely, robotics can benefit from such studies by using the suggested principles to design more adaptive control architectures.

1. Introduction

Given the complexity of tasks that humans and animals can address in contrast to the limited cognitive and motor capabilities of robots, an emerging line of research consists in building developmental and neuro-computational models of biological systems, with the hope to understand some underlying principles that may then be re-used to design robot control architectures. In particular, a way of understanding the principles of the generation of movement in the corresponding specific neural structures (some parts of the motor cortex, the cerebellum and basal ganglia) consists in building computational models implementing these principles and investigating their functional properties. But, in order to evaluate these models, we will exemplify in this paper the importance of embedding the neuro-computational model into a physical body or at least a realistic simulation, particularly when the model calls upon developmental processes. Indeed, we will show through a case study how taking into account the kinematics of the physical system can both reveal important questions ignored by models that did not simulate this kinemat-

ics and point out unnecessary assumptions made by these previous models.

More specifically, our starting point is Aflalo and Graziano's model (Aflalo and Graziano, 2006b), which addresses the emergence of a repertoire of motor primitives representing final arm postures corresponding to experimental observations. Indeed, Graziano and his team have observed through intracranial stimulations of the Macaque monkey that the hand reaches particular spatial configurations independently of its initial configuration when particular groups of neurons are stimulated (Graziano et al., 2002a, Graziano et al., 2002b, Graziano et al., 2003, Graziano et al., 2004). The stimulated area lies within the precentral gyrus. It contains parts of the primary motor cortex (M1), the caudal part of the dorsal premotor cortex (PMdc) and the ventral premotor cortex (areas F4 and F5 in the monkey), as shown in Figure 1.




Figure 1: The areas involved in the monkey's brain: primary motor area and subparts of the premotor area. The stimulated areas are in grey [from Graziano et al., 2005]

Later on, through recording the neural activity in the monkey during free movement of superior limbs, (Aflalo and Graziano, 2006a) have shown that the relevant neurons were mainly, though non exclu-

sively, tuned for the final posture in terms of articular angles. Other parameters previously considered as encoded in the premotor cortex, such as movement direction (Georgopoulos et al., 1986), speed, torques, hand position, seem less correlated to the variability observed in the neural firing rate.

In (Graziano et al., 2004), the authors claim that groups of neurons involved in arm movement form a map of final hand positions, where the tuning of neurons reflects the behavioral repertoire of the animal. Finally, it is suggested in (Graziano et al., 2005) that the organization of this map is not only based on the spatial position of the hand, but also on the somatotopy and the “ethological category” of the posture (i.e. the monkey “reaches”, “eats”, “climbs” etc.). Those three parameters would compete and the competition would result in a complex organization where all three representations overlap.

Aflalo and Graziano’s model is based on a Kohonen map (Kohonen, 1984, Kohonen, 2001). Very shortly, Kohonen maps are a family of neural networks with self-organizational properties, where the activated cell is the one that matches some input the most closely. Training a map with a set of inputs composed of a vector of values results in the property that neighbor cells according to some low-dimensional topology get activated for similar data in the higher dimensional space of inputs.

In the case of Aflalo and Graziano’s model, each input of the Kohonen map is a vector in three parts representing the limbs involved, the ethological category of the corresponding movement and the position of the hand. The first part is a binary vector of length 10 (one bit for each limb) where a 1 indicates that the corresponding limb is involved. The second part is a binary vector of length 5 that codes for the 5 possible categories (one bit for each ethological category). The third part is a real number vector of length 3 coding for the spatial coordinates of the hand in 3 dimensions.

Aflalo and Graziano initialize their Kohonen map with a simple somatotopic representation of the primary motor cortex drawn from (Woolsey et al., 1952). They show how training the map with a large set of handcrafted inputs results in a spatial organization similar to the one they observed on the surface of the cortex of the monkey by stimulating it. Nevertheless, the *ad hoc* encoding of information they used is questionable. In particular, one may wonder how the encoding of the ethological category of the movement could be so explicit in the neural system of the monkey.

Based on this work, (Ognibene et al., 2006a) have proposed another model of the constitution of a repertoire of postures. In (Ognibene et al., 2006b), they used a Kohonen map as well, but their model is trained in the context of a system learning to interact

with its environment.

Our work shares some similarities with Ognibene et al.’s but, where their system only consists of a simple arm model with two degrees of freedom (dofs) in the context of a simple reaching task, we use a more complicated 24 dofs human-like system realizing a more complete repertoire of movements. Addressing this more complex system raises the issue of the most efficient encoding of the postural information within the Kohonen map. In particular, we compare an encoding based on articular parameters with an encoding based on spatial positions of the extremities of relevant body parts. More importantly, we show that, just by rewarding some postures in some contexts and combining a reinforcement learning process with a self-organizing process, we can obtain a topological organization of the map reflecting the ethological categories of movement that Aflalo and Graziano were giving explicitly to their system.

The paper is organized as follows. First, we present the global architecture and principles of our model. We describe how these principles are implemented and our experimental procedure in section 3. Results comparing the coding with angles to the coding with positions are described in section 4. We discuss these results and more general features of our model in section 5. Finally, we conclude and present some directions for future research.

2. Our computational model

Our model consists of three elements (see Figure 2):

- an environment specifying a sequence of contexts where our system must fulfil successive constraints on postures in order to get rewarded (see table 1) ;
- a human-like manikin simulation providing a realistic kinematics where one can check at any moment whether the current posture corresponds to a rewarded posture in a given context or not ;
- a control architecture based on a Kohonen map, which must be trained to code for the relevant rewarded postures.

2.1 Environment

Our model of the environment reflects very abstractly the situation of a monkey motivated by hunger in the presence of a fruit. The monkey must successively catch the fruit, manipulate it, bring it to its mouth and defend itself against an opponent (Figure 3 right). This simplified environment is directly inspired by (Graziano et al., 2005) (see Figure 3 left). The constraints that a posture must fulfil to be rewarded are given in table 1.

Context	Goal of action	Constraints on final posture
(C1) fruit in range	catch fruit	$d_{hb} > 64cm$
(C2) fruit in hand	manipulate fruit	$d_{hh} < 16cm, d_{hb} < 40cm, h$ in front of b
(C3) fruit peeled	eat fruit	$d_{hm} < 14cm$
(C4) aggression	defense	$d_{hh} < 15cm, d_{aa} < 22cm, d_{af} < 19cm$

Table 1: Specification of rewarded postures corresponding to the contexts. In the last column, d = distance, h = hand, m = mouth, b = body, a = arm, f = face, thus d_{af} means distance from arm to face, for instance

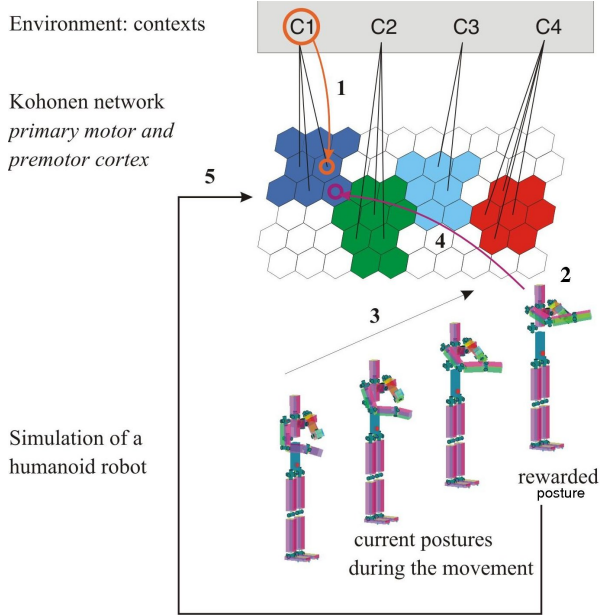


Figure 2: Model architecture: environment, Kohonen map and simulated manikin. Context C1 (see table 1) activates the marked neuron (1) which codes for the represented goal posture (2). Movement from the initial posture to the goal posture is shown at the bottom from left to right (3). The forth posture satisfies the context rewarding constraints, thus the link between C1 and the most active cell is reinforced (4). Finally, the Kohonen map is trained with this posture and a subsequent one towards the goal posture corresponding to C2 (5)

2.2 Simulated body: Arboris

We use a human manikin simulation named “Arboris” (Collette, 2007). The poly-articulated model of the system contains 5 branches: the trunk, both legs and both arms (see Figure 4 left). The trunk is split into two bodies, arms and legs are split into three bodies. There are 36 dofs, expressed in articular parameters corresponding to angles. All bodies and articular limitations are defined according to anthropometric data adapted from the HuMAnS manikin documentation (Wieber et al., 2006). We restrict the manikin to 24 dofs on the upper part of the model to avoid dealing with equilibrium constraints. The movement from one posture to another




Figure 3: Left: categories of movement targets [from Aflalo and Graziano 2006]. Light blue: hand-to-mouth; dark blue:reach; red:protect oneself; green:manipulate; pink:climb. Right: equivalent in our simulation. Dark blue:catch; green:manipulate; light blue:eat; red:defense




Figure 4: Virtual manikin used in our simulation. left: Numbers indicate the numbers of degrees of freedom. right: Orientation of the reference frame defining positions of body part extremities

is obtained by a mere spline interpolation, imposing a bell-shaped velocity profile as observed in natural motion. The number N of intermediate postures between the initial and final null speed postures is a parameter of our system (here we take $N = 5$).

2.3 Posture representation

Like (Aflalo and Graziano, 2006b) and (Ognibene et al., 2006b), our encoding of a repertoire of postures is based on Kohonen maps (Kohonen, 1984, Kohonen, 2001). But our use of a complex virtual manikin highlights a problem that neither of these authors addressed. Indeed, building

our model implies finding a good match between three languages: the specification of rewards associated to each context, the encoding of goal postures in the repertoire, and the encoding of the current posture in the manikin simulation.

The goal of the repertoire learning process is to ensure that Kohonen cells representing goal postures code as precisely as possible for rewarded postures of the manikin. Aflalo and Graziano did not face this question because their model does not call upon a simulation, and Ognibene’s two-dimensional model being too simple, the transformation from angles to positions and vice-versa is trivial.

In our case, the movement of the manikin is defined by the articular parameters of the initial and final postures, thus coding postures with angles, as we do in a first model, seems the most appropriate. But representing rewarded postures is simpler with positions of body parts or cartesian distances between these parts than with angles. For instance, when the manikin must eat, the relevant feature is the relative position of one hand with respect to the mouth and specifying that with a chain of articular angles is complex. The underlying question is whether the brain codes for postures in a language adapted for controlling the body or adapted for specifying the goals, or even something else.

As a consequence, we designed a second model where the topology of the Kohonen map is based on positions rather than on angles. We extracted 24 coordinates corresponding to the spatial position of eight body parts: both shoulders, elbows, wrists and hand tips. These positions are expressed in a reference frame based on the eyes, as shown in Figure 4 right, so as to be able in future work to integrate visuomotor transformations in the context of a wider model calling upon a visual loop and a motor loop as suggested by (Nakahara et al., 2001).

In order to implement this second model, we must transform a goal position defined with spatial coordinates into a goal posture defined with angles, which is necessary to control the manikin. In our context, we cannot directly use inverse kinematics solution schemes, since we specify the position of only a subset of body parts: there exists infinitely many corresponding angular configurations. (Ognibene et al., 2006b) solved this problem with a “direct inverse modeling” method by learning simultaneously two associations through a large set of 2D postures: first, the association between articular angles (outputs) and the activation of the Kohonen map and second, the association between a perceived position and the activation of the Kohonen map. The articular angles can then be retrieved by a mere activation of the map. However, what they did on a simple two-dimensional problem cannot be straightforwardly extended to our 24 dofs context.




Figure 5: Second model architecture: one Kohonen map coding for spatial positions of body parts extremities is added with respect to Figure 2 and the correspondence between angles and positions is ensured (see text)

Thus we rather used a second Kohonen map, as shown in Figure 5. The second map codes for positions, and a cell in this map can be addressed from a cell of the map coding for angles through a direct geometric transformation computing the position from the angles. The neighborhood relationship is established in the map coding for positions. Then, during training, this relationship is “mapped” onto the map coding for angles through a 25^{th} variable coding for distances between postures expressed as a set of positions of body parts. Then, the map coding for angles is trained and the result is transferred again to the map coding in positions, so as to preserve the one-to-one correspondence. Through this process, we can both use our interpolation mechanism and constrain the modification of angles so as to ensure that articular limitations are not violated.

3. Model implementation and experimental design

We used an implementation of Kohonen maps available on internet¹. The cells of the Kohonen map in our model code for goal postures (noted P_g) of the manikin, which means that the vector of weights associated to each cell contains information about a posture (either with angles or with positions). These maps of $16 \times 24 = 384$ cells (noted c) are initialized by training them with angles with 2 million random postures, taken in different order, which results in different maps.

After initialization of the Kohonen map and selec-

¹at <http://www.cis.hut.fi/projects/somtoolbox>

Parameters: $\{C_1, \dots, C_n\}, N$ (here $N = 5$)

1. $P_c(0) \leftarrow$ random posture
 2. for $k = 1$ to n //index of current context
 3. for $i = 1$ to N //index of current posture
 - (a) $P_g \leftarrow \max_c (e^{\frac{Q(c, C_k)}{T}}) / (\sum_j e^{\frac{Q(c_j, C_k)}{T}})$
 - (b) $P_c(i+1) \leftarrow$ interpolate($P_c(i), P_g$)
 - (c) if ($P_c(i+1)$ rewarded)
 - $Q(P_c(i+1), C_k) = (1 - \alpha)Q(P_c(i+1), C_k) + \alpha R$ with $R = 1$
 - train Kohonen map,
 - (d) if ($(P_c(i+1)$ not rewarded) and ($P_c(i+1) = P_g$))
 $Q(P_c(i+1), C_k) = (1 - \alpha)Q(P_c(i+1), C_k) + \alpha R$
with $R = -0.3$
 4. end for
 5. end for
-

Figure 6: The main algorithm (implemented in Matlab)

tion of an initial random posture (Figure 6, line 1), we present to the system the sequence of contexts (noted C) that appears in table 1, starting on the top line. The system must learn to reach the corresponding rewarded postures. Once the first rewarded posture is reached, the second context is presented and so on up to the fourth. If the system fails, a new iteration starts: a new random posture is selected and the system is reinitialized to the first context. We consider that the goal posture corresponding to a context is learned as soon as the system reaches it during 10 successive iterations. We perform $N_i = 1000$ iterations of the algorithm in Figure 6 for each test.

We do not model context perception: a “context cell” is simply set to “on” in the corresponding context. In order to learn the correspondence between contexts and goal postures, we define a set of links (noted $Q(c, C)$) between each context cell and each cell in the Kohonen map. The active goal posture P_g will tend to be the one whose link to the current context cell is the strongest, provided that we use some exploration. Here, our exploration mechanism is a classical Boltzmann law (Figure 6, line 3a) that explores all the more than the constant temperature T is higher and we take $T = 0.03$. All links are initialized to 0 at the beginning of a test.

Given an initial posture and a goal posture, the manikin control system will generate a sequence of current postures P_c through a simple interpolation mechanism. If a current posture satisfies the constraints corresponding to the current context, the link $Q(P_c, C)$ is reinforced (Figure 6, line 3c). If the current posture is the goal posture and the posture is not rewarded, the link $Q(P_c, C)$ is weakened (Figure 6, line 3d).

Furthermore, a second learning process consists in training the Kohonen map so as to represent opti-

mally the set of postures that are rewarded in a given context. The training data are the postures corresponding to the first rewarded posture in a given context and the first posture after the next change of context if this posture would have been rewarded in the previous context. This latter posture is added so as to drive the current goal posture towards the next goal posture when it is possible. For instance, when one manipulates before eating, it may be more efficient to manipulate with the hands as close to the mouth as possible. The system performs 15 standard training iterations for each of these postures.

4. Results

We studied the influence of the type of coding, either angles or positions, on the learning speed and resulting organization.

Figures 7 and 8 were obtained from the same initial map, but in the former we coded postures with angles whereas in the latter they were coded with positions. From these figures, one can see that there are already in the initial map a few cells (between 1 and 8 cells with a slightly different topology each time) that code for a rewarded posture in each context. And we noted that the convergence time is highly dependent upon this number of initial “correct” cells.

Our learning process results from the combination of two tightly coupled learning and self-organization processes. First, the system learns to bind to each context the cell representing an adequate goal posture in this context so as to drive the manikin towards this posture. Second, the Kohonen map is trained so that each goal cell codes as precisely as possible for the domain in which the corresponding postures are rewarded.

If we restricted the model to the first learning process, it would consist in finding through a random trial-and-error process that activating those initially adequate goal cells in adequate contexts results in receiving some reward. Then the link between the corresponding cell and context would be reinforced and this cell would get a higher probability of getting activated when the same context is presented again. Thus if there was no such cells coding for rewarded postures in the initial map, this first learning process would fail. This explains why the convergence time depends so much on the initial conditions.

The second learning process comes into play because, each time the current posture is rewarded, the Kohonen map is trained so that the active cell codes more centrally for this rewarded posture in the future (Figure 6, line 3c). In Figures 7 and 8, one can observe that this self-organization process results in a global reduction of distances coded by neighbor cells and in the formation of clusters of goal cells that code adequately for rewarded postures. This is a key el-

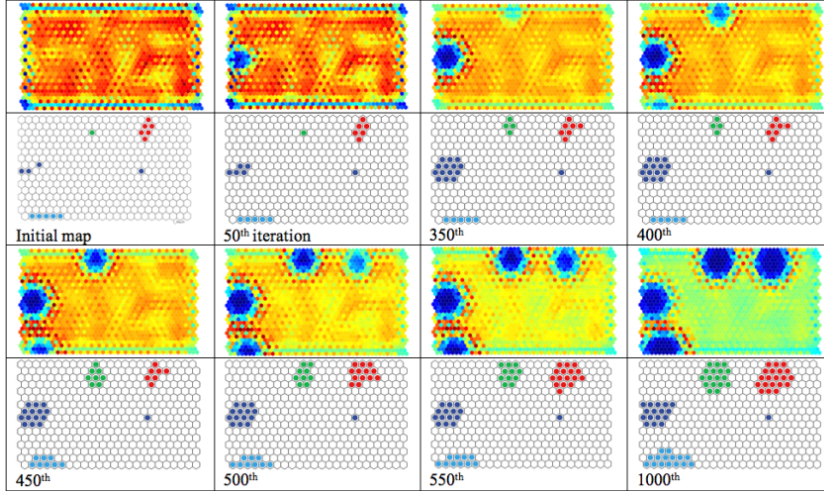


Figure 7: Evolution of the repertoire of postures along iterations when postures are coded with articular parameters. First and third lines: average distances between the postures represented by neighbor neurons (dark blue = very close, red = very far). Second and fourth lines: cells coding for rewarded postures in different contexts are shown in color: dark blue: catch ; green: manipulate ; light blue: eat; red: protect oneself .

ement in our work: thanks to this second process, clusters corresponding to the ethological categories of Aflalo and Graziano’s model are formed without giving explicitly these categories as inputs.

Globally, whatever the encoding used, our algorithm learns steadily to reach a sequence of rewarded postures that are not given explicitly to the system, but discovered through a trial-and-error process. On average over 40 runs, it takes 100 iterations to learn to reach steadily the next rewarded posture in the sequence once the previous one is already learned.

The location of the cluster of neurons corresponding to the ”catching” context is different between Figures 7 and 8. These figures were chosen to illustrate the fact that a cluster tend to be formed around the neurons coding initially for a rewarded posture. In the case of ”catching”, we can see that in the initial map two regions are relevant for this context, but with our learning algorithm only one region is trained per run. So in two different runs, this cluster can be at two different positions; but in the same run, even if two different occurrences of the same ethological posture are presented in the sequence, only one cluster is formed. By comparing these figures, one can also observe that this cluster formation process is slightly faster and clearer when postures are represented with articular parameters, though the difference in speed is not significant.

5. Discussion

While (Aflalo and Graziano, 2006b) were trying to reproduce the complex spatial organization of neurons coding for a repertoire of postures in the pre-central gyrus of the monkey that they had observed

through stimulation and recording, our standpoint was rather to look for general organizational principles that may explain the constitution of this repertoire through learning in interaction with an environment.

As a consequence, our work led us to question some of the methodological assumptions underlying Aflalo and Graziano’s model. In particular, the organization they get in their Kohonen map directly results from an *ad hoc* encoding of the inputs of this map. In that respect, we have shown how one can obtain a spatial organization into clusters coding for relevant ethological categories without explicitly representing these categories in the inputs of the repertoire.

Furthermore, whereas in Aflalo and Graziano’s model the cells of the Kohonen map only keep track of the position of one hand, the fact that we considered a complex representation of postures with 24 dofs led us to face the problem, biologically relevant, of the nature of information used to code for these postures.

Indeed, we observed that coding with articular parameters results in a clearer organization of clusters than with body parts positions. The distances between encoded goals outside of clusters is more uniform and this distance decreases faster, resulting in a faster cluster formation process in the former case. As a result, coding with angles seems more advantageous, since it is also computationally less expensive: it does not call upon the costly transformation from articular angles to positions and back.

These arguments in favor of coding with angles obtained in a simulation study must be reconsidered in the light of the neuro-physiological debate. Indeed, some experimental studies

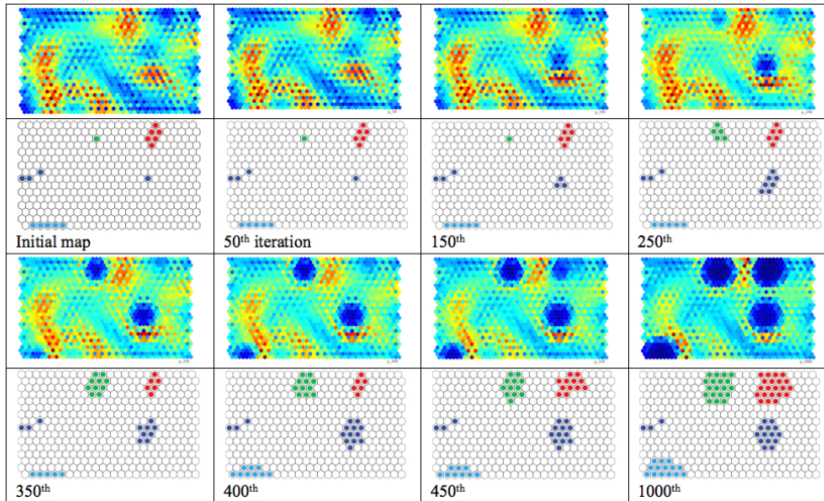


Figure 8: Same information as in Figure 7, but cells code for body parts positions instead of articular parameters.

inspired by the “target position hypothesis” (McNeilage, 1970, Russel, 1976) clearly speak in favor of a “global” reference system expressed with positions (Georgopoulos et al., 1981, Baud-Bovy and Viviani, 1998). But other studies like (Scott and Kalaska, 1997) or (Rosenbaum et al., 1999) have shown that motor cortical neurons fire in function of initial and final angular postures of a movement. Thus it seems that both kinds of information are used in the brain. The kind of representation used may depend on the task, or there might be some redundant encoding in the motor system, as (Haggard et al., 1995) and (Kawato, 1996) suggest. A control strategy based on this redundancy still awaits to be studied.

Another possibility would be that the postures are coded in a visual reference frame. There is an important body of literature about the different reference frames and transformations between them that are used in visuomotor control (Brotchie et al., 1995, Pouget et al., 2002). Taking this dimension into account would also open the possibility to distinguish two learning mechanisms, one in the visual system and one in the motor system (Nakahara et al., 2001).

6. Conclusion and future work

In this paper, we have shown that using a simulation of a sophisticated human-like manikin may be necessary to design more accurate models of the neuro-physiological processes that take place in the control of complex movements in humans and animals. First, such a choice revealed the importance of the nature of information encoded in the repertoire of postures that was not addressed by simpler models lacking this simulation dimension. Second, it was made possible to show that the assumption of an explicit encoding of the ethological category was not

necessary to obtain the expected organization. Indeed, since the repertoire of postures obtained was resulting from the combination of a self-organization process and a reinforcement learning process involving interactions with an environment, a simulation of this interaction was necessary to observe the emergence of these ethological categories.

Thus we have shown how human-like simulations can contribute to modeling activities in life sciences. Conversely, our study illustrates how drawing inspiration from neuro-physiological studies to model developmental processes that take place in motor control can result in the design of efficient control architecture for virtual manikin and humanoid robots. Indeed, from the repertoire of postures that we obtained, one can design or learn complex sequences of simple actions that may result in the achievement of diverse robotics tasks.

Like Affalo and Graziano’s, our model was focused on learning a repertoire of postures. Even if our study is concerned with sensorimotor learning, here we did not address numerous aspects of the motor learning system such as the constitution of sequences of actions or motor adaptation.

In order to address biological relevance further, in the immediate future we have to replace the simple interpolation mechanism used in our model by a more accurate low level control law taking into account the dynamics of the system and some equilibrium constraints that must be satisfied to prevent the manikin from falling. In parallel, we are investigating the automatic learning of a dynamical model that may take place in the cerebellum.

Furthermore, the structure of the sequence of goals learned here is extremely simple, it is a mere succession of associations between a set of predefined contexts and constraints on postures. In order to

address more general robotics tasks, we will have to focus in the future on a more accurate model of sequence learning and goal processing architecture, drawing inspiration from computational models of basal ganglia that try to elucidate the mechanisms of reinforcement learning and sequence formation. A true robotic implementation will then be necessary to qualify the robustness of the global architecture to sensors and actuators noise.

Acknowledgments

The authors thank Sébastien Barthelemy, Cyrille Collette and Alain Micaelli for providing Arboris.

References

- Aflalo, T. N. and Graziano, M. S. A. (2006a). Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. *PNAS*, 103:2909–2914.
- Aflalo, T. N. and Graziano, M. S. A. (2006b). Possible origins of the complex topographic organization of motor cortex: Reduction of a multidimensional space onto a two-dimensional array. *Journal of Neuroscience*, 26(23):6288–6292.
- Baud-Bovy, G. and Viviani, P. (1998). Pointing to kinesthetic targets in space. *Journal of Neuroscience*, 18(4):1528–1545.
- Brotchie, P. R., Andersen, R. A., Snyder, L. H., and Goodman, S. J. (1995). Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature*, 375:232–235.
- Collette, C. (2007). Etude et réalisation d’algorithmes de contrôle d’humains virtuels dynamiques. Technical report, Rapport interne CEA.
- Georgopoulos, A. P., Kalaska, J. F., and Massey, J. T. (1981). Spatial trajectories and reaction times of aimed movements: Effects of practice, uncertainty, and change in target location. *Journal of Neurophysiology*, 46:725–743.
- Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419.
- Graziano, M. S. A., Cooke, D. F., Taylor, C. S. R., and Moore, T. (2003). Distribution of hand location in monkeys during spontaneous behavior. *Experimental Brain Research*, 155:30–36.
- Graziano, M. S. A., Patel, K. T., and Taylor, C. S. R. (2004). Mapping from motor cortex to biceps and triceps altered by elbow angle. *Journal of Neurophysiology*, 94:4209–4223.
- Graziano, M. S. A., Taylor, C. S. R., and Moore, T. (2002a). Complex movements evoked by microstimulation of precentral cortex. *Neuron*, 34:841–851.
- Graziano, M. S. A., Taylor, C. S. R., Moore, T., and Cooke, D. F. (2002b). The cortical control of movement revisited. *Neuron*, 36:1–20.
- Graziano, M. S. A., Tyson, N. S., and Cooke, D. F. (2005). Arm movements evoked by electrical stimulation in the motor cortex of monkeys. *Journal of Neurophysiology*, 94:4209–4223.
- Haggard, P., Hutchinson, K., and Stein, J. (1995). Patterns of coordinated multi-joint movement. *Experimental Brain Research*, 107:254–266.
- Kawato, M. (1996). Bidirectional theory approach to integration. In an McClelland J. L., I. T., (Ed.), *Attention and performance XVI: information integration*, pages 335–367. MIT Press, Cambridge, MA.
- Kohonen, T. (1984). *Self-Organizing Maps and Associative Memory*. Springer, Berlin.
- Kohonen, T. (2001). *Self-Organizing Maps*. Springer, Berlin.
- McNeilage, P. F. (1970). Motor control of serial ordering of speech. *Psychological Review*, 77:182–196.
- Nakahara, H., Doya, K., and Hikosaka, O. (2001). Parallel cortico-basal ganglia mechanisms for acquisition and execution of visuomotor sequences. *Journal of Cognitive Neuroscience*, 13:626–647.
- Ognibene, D., Mannella, F., Pezzulo, G., and Baldassare, G. (2006a). Integrating reinforcement learning, accumulator models and motor primitives to study action selection and reaching in monkeys. In *ICCM 06*, pages 214–219.
- Ognibene, D., Rega, A., and Baldassare, G. (2006b). A model of reaching that integrates reinforcement learning and population encoding of postures. In *SAB 06*, pages 381–393.
- Pouget, A., Ducom, J. C., Torri, J., and Bavelier, D. (2002). Multisensory spatial representations in eye-centered coordinates for reaching. *Cognition*, 83:1–11.
- Rosenbaum, D. A., Meulenbroek, R. G. J., and Vaughan, J. (1999). Remembered positions : Stored locations or stored postures? *Experimental Brain Research*, 124:503–512.
- Russel, D. G. (1976). Spatial location cues and movement production. In E., S. G., (Ed.), *Motor control: issues and trends*, pages 67–83. Academic, New York.
- Scott, S. H. and Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientations: I. activity of individual cells in motor cortex. *Journal of Neurophysiology*, 77:826–852.
- Wieber, P.-B., Billet, F., Boissieux, L., and Pissard-Gibollet, R. (2006). The humans toolbox, a homogeneous framework for motion capture, analysis and simulation. In *Proceedings of the ninth ISB Symposium on 3D analysis of human movement*, Valenciennes, France. Academic.
- Woolsey, C. N., Settlage, P. H., Meyer, D. R., Sencer, W., Hamuy, T. P., and Travis, A. (1952). Pattern of localization in precentral and supplementary motor areas and their relation to the concept of a premotor area. *Association for research in nervous and mental disease*, 30:238–264.