Mastering Growth while Bootstrapping Sensorimotor Coordination

Matthias Rolf, Jochen Steil, Michael Gienger

2010

Preprint:

This is an accepted article published in Int. Conf. on Epigenetic Robotics (EpiRob). The final authenticated version is available online at: https://doi.org/[DOI not available]

Mastering Growth while Bootstrapping Sensorimotor Coordination

Matthias Rolf* **Jochen J. Steil*** **Michael Gienger****
*CoR-Lab, Bielefeld University
{mrolf,jsteil}@CoR-Lab.Uni-Bielefeld.DE
**Honda Research Institute Europe GmbH
michael.gienger@Honda-RI.de

Abstract

The bodily change in infancy due to growth is a fundamental challenge for the bootstrapping of sensorimotor coordination. We argue that learning by doing, and thus a babbling of goals instead of motor commands provides an appealing explanation for the success of infants in that bootstrapping. We show that Goal Babbling allows to bootstrap reaching skills during different growth patterns on a robot arm with five degrees of freedom and on the infant like humanoid iCub.

1. Introduction

While infants bootstrap their repertoire of sensorimotor skills, their bodies undergo massive changes in overall size and weight, segment lengths as well as mass distribution. Until infants master reaching around the age of one year, they have grown (on average) by 50% of their body length at birth. The change of morphology is non-linear in many ways. The growth proceeds very rapidly in the first few months before it slows down (Kuczmarski et al., 2002). Also, different body segments grow with different speeds, changing the body proportions. For instance, the upper arm typically grows faster than the forearm (Wells et al., 2002).

How can infants develop stable sensorimotor skills during such ongoing and non-linear change in morphology? The control of tasks like reaching can be well understood with the notion of internal models (Wolpert et al., 1998). Internal models describe relations between motor commands and their consequences. Once internal models are established for a certain task, a forward model predicts the consequence of a motor command, while an inverse model suggests a motor command necessary to achieve a desired outcome. Before internal models can be applied for goal-directed control, experience must be gained by exploration. The question how, and when to explore is particularly delicate in the case of growth. Any experience that is gathered by exploration is, strictly speaking, void as soon as the body changes – and so are internal models.

Piaget suggested that human development progresses in several stages (Piaget, 1953). At first infants react purely reflexive. Meltzoff and Moore (Meltzoff and Moore, 1997) suggested the concept of "body babbling" as an initial stage in which experience is gathered. Infants can then use this knowledge to attempt goal-directed actions and finetune their skills on the fly. This idea of an explicit, early exploration strategy has been frequently picked up in computational models for the development of reaching skills. A random exploration of all possible motor commands is performed under the notion of "motor babbling" (Demiris and Dearden, 2005). Several models also include a later fine-tuning which is done on the fly (Jordan and Rumelhart, 1992, D'Souza et al., 2001), while performing goal-directed movements.

However, models that separate exploration from goal-directed movement can hardly explain how infants master growth during sensorimotor development. It is rather implausible that infants spend their entire first months babbling randomly, while growth outdates the gathered experience. In particular, they cannot explain findings of early goal-directed movements in newborns: Statistics revealed that already days after birth, infants attempt goal-directed action by means of arm and finger movements (von Hofsten, 1982). This finding contradicts the Piagetian view of purely reflexive movements and also purely random movements. Von Hofsten has repeatedly highlighted the role of goal-directed action for infant motor development. "Before infants master reaching, they spend hours and hours trying to get the hand to an object in spite of the fact that they will fail, at least to begin with." (von Hofsten, 2004)

Early goal-directed movements in infants suggest an important role of "learning by doing": infants learn to reach by trying to reach. In the case of voluntary control we previously denoted this kind of learning as *Goal Babbling* (Rolf et al., 2010). Different goals are tried to achieve and thereby sensorimotor coordination is bootstrapped. Such goal-directed bootstrapping has been previously possible with Feedback-Error Learning (Kawato, 1990), but only with prior knowledge that is unlikely to be avail-

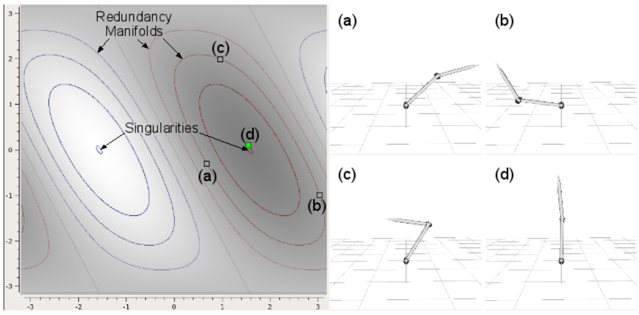

Figure 1: Robot arm (length $1m$) with two joints. The left side shows the joint space. Multiple configurations (see postures a-c) can be used to apply the same height of the end effector, but can not be averaged without leaving the desired height (see posture d). The sets of joint angles that apply the same height are marked by colored contours in the joint space.

able for complex motor skills. Goal-directed bootstrapping appears to be a "chicken or egg" dilemma: Goal-directed movements require an inverse model, but inverse models require previous exploration in order to work. However, we have previously shown that even an entirely untrained inverse model is enough for starting a successful bootstrapping.

This paper extends previous lines of research on Goal Babbling. We show that Goal Babbling can bootstrap reaching skills while the morphology changes, i.e. while the body is growing. Our algorithm can do so without explicit knowledge of when a change occurs or what that change is, which shows the conceptual strength of Goal Babbling, or learning by doing in general. In the following, we introduce the learning problem in Sec. 2. and shortly describe our algorithm for Goal Babbling. We present experimental results with different growing speeds and unproportional body changes on redundant morphologies in Sec. 3. and conclude in Sec. 4.

2. Goal Babbling

In the present work, we investigate the kinematic control of redundant systems. Formally, we consider the relation between joint angles $q \in \mathbf{Q} \subset \mathbb{R}^m$ and effector poses $x \in \mathbf{X} \subset \mathbb{R}^n$ (e.g. the position of the hand). Thereby m is the number of degrees of freedom (DOF) and n is the dimension of the target variable (e.g. n =3 for the spatial position of a hand). The forward kinematics function $f(q) = x$ uniquely describes the causal relation between both sizes. If the hand needs to be positioned at some desired coordinate x^*, an inverse model $g(x^*) = q$ is needed to find appropriate joint angles q ($f(g(x^*)) = x^*$). Such a function is not uniquely defined if the number of joint angles m exceeds the number of controlled dimensions n. An example is shown in Fig. 1: a robot arm with two joints ($m = 2$) and a total length of 1m. Since we want to consider a redundant structure, the goal here is to control only the height of the effector ($n = 1$). The redundancy appears in form of manifolds through the joint-space, on which all joint-angles apply the same effector height. An inverse kinematics function in this example must suggest joint angles $q \in \mathbb{R}^2$ for each desired effector height $x^* \in \mathbb{R}^1$.

2.1 Related work

Existing approaches to the exploration and learning of inverse kinematics split into two groups: error-based and example-based methods. *Error-based* methods follow the "learning by doing" approach. An estimate $g(x^*)$ of the inverse kinematics is used for trying to reach for a target position. Using the joint angles $q = g(x^*)$ suggested by the inverse estimate, the resulting position of the effector is evaluated with the forward kinematics function $x = f(q)$. One group of mechanisms is based on the "motor error", which is a correction Δq of the joint angles in order to improve the performance. In *Feedback-error learning* (Kawato, 1990) it is simply assumed that a mechanism to compute that motor error is already available. In *Learning with distal teacher* (Jordan and Rumelhart, 1992) a forward model $\hat{f}(q)$ must be learned beforehand which requires a not goal-directed exploration. A motor error can be derived analytically by differentiating the forward model. Both methods can in principle deal with redundant systems. The critical problem is that the motor error is not directly observable, and on its own subject to redundancy. Thus current motor error schemes can not explain infants ability to fully bootstrap sensorimotor coordination.

Example-based methods use example configurations $(f(q), q)$ for the learning of an inverse estimate $g(x)$. The existing approaches differ in the way how such examples are generated. Motor babbling (Demiris and Dearden, 2005) is a pure random form of exploration. It has been proposed as an implementation of the "body babbling" introduced by Meltzoff and Moore, but was used also before body babbling was introduced (Bullock et al., 1993). Joint angles are randomly chosen from the set of all possible configurations $q_i \in \mathbf{Q}$. This approach can find solutions for all possible targets, if enough examples are generated. However, it is subject to the non-convexity problem: Non-convex solution sets (see Fig. 1) prohibit learning from multiple solutions. Also goal-directed exploration approaches have been investigated (Oyama and Tachi, 2000, Sanger, 2004), which we discuss and extend in the next section. The approach – as previously discussed in literature – does not find appropriate inverse estimates in a reliable fashion. Example-based learning of inverse kinematics has only been shown to be successful if training data without inconsistent solutions is already available (Rolf et al., 2009).

2.2 Goal Babbling Algorithm

With "Goal Babbling" we generally refer to the successful bootstrapping of some motor skill by the (*i*) repeated process of (*ii*) trying to accomplish (*iii*) multiple goals related to that skill. Goal babbling means learning by doing from scratch. We use this terminology in order to highlight the similarities but also differences to previous concepts. The exploration process focuses on the goals of action instead of the means (motor-commands). The emphasis in this approach is on "trying to accomplish", which means to generate paths towards the given goal with the currently learned system and to evaluate samples along this path. In the present work we use target positions as a rather simple notion of goals and select randomly which goal is tried to achieve next. In an integrated, physical system this might be done by presentation of external stimuli which the agent is motivated to reach for.

Goal-directed exploration In a goal-directed exploration (Oyama and Tachi, 2000, Sanger, 2004), examples $(f(q), q)$ are generated with an untrained or inaccurate inverse estimate $g(x, \theta)$, where θ are the parameters adaptable by learning. If the parameters are not necessary for the discussion, we will write $g(x)$ for short. Initially, a target motion is chosen and represented as a temporal sequence of target positions: $x^*_t \in \mathbf{X}^* \subseteq \mathbf{X}, t = 1...T$. The inverse estimate is then used for trying to reach for those targets: $q_t = g(x^*_t)$, $x_t = f(q_t)$. After the adaption of the parameters, the process is repeated.

Inconsistencies In this approach inconsistent examples can be detected, and excluded by considering paths on the n dimensional manifold inside the joint space on which examples are generated. This manifold is defined by the inverse estimate and spanned by the set of target positions: $\mathbf{Q}^{expl} = g(\mathbf{X}^*)$.

Two samples (x_1, q_1) and (x_2, q_2) are inconsistent, if they represent the same effector pose $x_1 = x_2$ but different joint angles $q_1 \neq q_2$. We have previously shown (see (Rolf et al., 2010) for the details) that such examples can only be generated by goal-directed exploration, if there are either unintended changes of movement direction or inefficient movements. Thus, if we generate examples with goal-directed exploration and exclude both unintended changes of movement direction and inefficient movements, the remaining examples must not contain inconsistencies.

In order to realize this exclusion, we assign weights

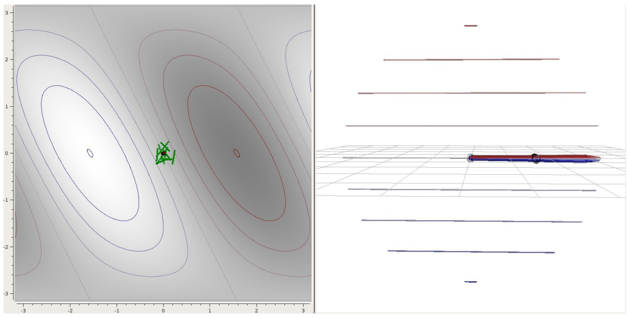

(a) The inverse estimate is initialized around the home posture.

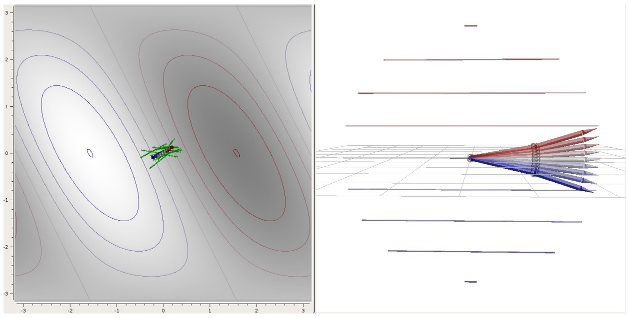

(b) The inverse estimate has aligned with the optimal movement direction and starts to expand.

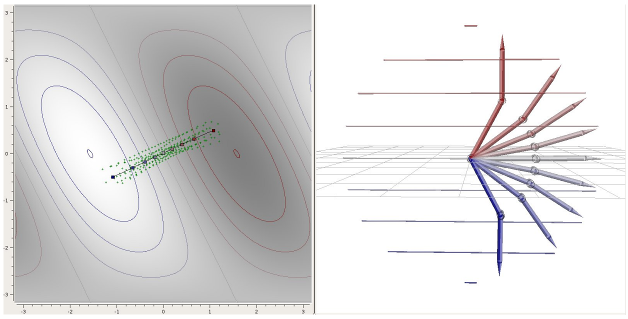

(c) The performance increases rapidly. Until the ridge of the forward function is hit.

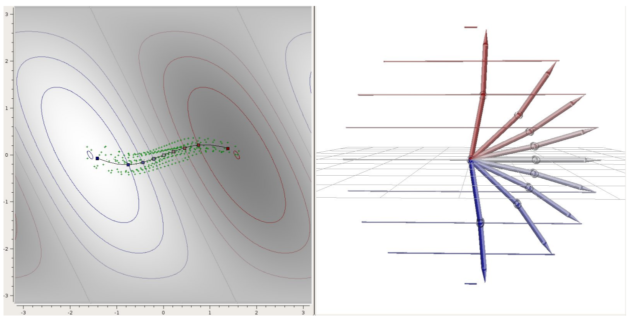

(d) The inverse estimate finds the necessary non-linearities to reach for extreme positions.

Figure 2: Inverse kinematics learning with Goal Babbling. The images show successive stages of the learning process. The inverse estimate is initialized around the home posture. It spreads successively and ends up with an accurate solution.

$w_t \in \mathbb{R}$ for each example (x_t, q_t). Unintended changes of movement direction can be tackled with the following scheme:

$$w_t^{dir} = \frac{1}{2}\left(1 + cos\sphericalangle(x_t^* - x_{t-1}^*,\ x_t - x_{t-1})\right). \quad (1)$$

Thereby $\sphericalangle(x_t^* - x_{t-1}^*, x_t - x_{t-1})$ is the angle between the intended and actual movement direction of the effector. If both are identical the angle is 0.0° and the weight becomes $w_t^{dir} = 1.0$. If the observed movement has the exact opposite direction, the angle is 180.0° and the weight becomes $w_t^{dir} = 0.0$. Inconsistencies evoked by unintended movement directions can therefore be broken.

Inefficient movements can be excluded by weighting with the ratio of effector motion and joint motion, which is 0.0 if the joints move without effector motion:

$$w_t^{eff} = \frac{||x_t - x_{t-1}||}{||q_t - q_{t-1}||} \quad . \quad (2)$$

Since both weights are necessary for inconsistency resolution, they are combined by multiplication, such that an example is ignored if any of the two criteria assigns a weight zero:

$$w_t = w_t^{dir} \cdot w_t^{eff} \quad . \quad (3)$$

The weighting scheme relies on time, since the actual and the last sample is taken into account. In particular, it relies on goals: unintended changes of movement direction can only be detected if there is an intended direction.

Structured Variation for Efficient Exploration So far, only those examples are explored that are exactly on the manifold of the inverse estimate. Such behavior is highly unrealistic for human motor development. If a motor command is sent twice, neural and muscular noise as well as external perturbations can cause slightly different outcomes. In our simulation we need to introduce such perturbations artificially. Therefore we add a small disturbance term $E^v(x)$ to the inverse estimate:

$$g^v(x) = g(x) + E^v(x). \quad (4)$$

Examples are then generated with this variation instead of the actual inverse estimate: $q_t^v = g^v(x_t^*)$, $x_t^v = f(q_t^v)$. For a set of examples, generated with a variation $g^v(x)$, the weighting scheme can be applied as proposed above. The index v is added to identify weights for examples of a specific variation: $w_t^v = w_t^{v^{dir}} \cdot w_t^{v^{eff}}$.

Although exploration is fundamental in infancy, infants do not try to reach for an object forever. At a time, they stop exploration, relax their muscles and rest. Learning is possible from such a "neutral" motor command, since there is still a resulting effector pose. At the level of kinematics, we denote a home posture q^{home} as neutral motor command. The result $f(q^{home})$ can be observed and be used for learning as any other example. We add the example $q_0^v = q^{home}, x_0^v = f(q^{home}) = x^{home}$ to each set generated with goal-directed exploration:

$$D^v \longleftarrow \{(f(q^{home}), q^{home})\} \cup D^v \quad (5)$$

The "home" example receives the full weight $w_0^v = 1.0$.

A home posture is a stable point in exploration, and thus in learning. The inverse estimate will generally tend to reproduce the connection between q^{home} and x^{home} if it is used for learning: $g(x^{home}) \approx q^{home}$. This stable point largely prevents the inverse estimate to drift away. Learning can start around the home posture and proceed to other targets.

Learning Example data (and corresponding weights) from multiple different variations $g^v(x^*), v = 1...V$ is combined for learning, where $V \in \mathbb{N}$ is the number of different variations. In the learning step, the parameters θ of the inverse estimate $g(x, \theta)$ are updated using the generated examples $(x_t^v, q_t^v), t = 0...T$ (including the home posture) and weights w_t^v in a reward weighted regression manner. Thereby the weighted command error

$$E_w^Q(\theta) = \sum_v \sum_t w_t^v \cdot (g(x_t^v, \theta) - q_t^v)^2 \quad (6)$$

is minimized. Any regression algorithm can be used for this step (e.g. linear regression schemes).

The overall procedure works in epochs. The inverse estimate is initialized with some parameters θ. We generally use a random initialization, but such that the inverse estimate generates joint configurations closely around the home posture for all goal positions. There is no a priori knowledge about the structure of the kinematics. Within one epoch, examples are generated from multiple variations, weights are assigned and the learning is done with the examples. The next epoch repeats the procedure with the updated inverse estimate.

An example of inverse kinematics learning with Goal Babbling on the 2 DOF arm (see Fig. 1) is shown in Fig. 2. The inverse estimate is visualized by a one dimensional manifold through the joint space. For several target heights x^*, the joint angle estimates are shown by colored markers on the manifold (the joint angles are furthermore visualized by corresponding postures in the 3D simulation). An accurate inverse estimate places all colored markers on the redundancy manifold with the same color. Small green markers show the examples used for learning. The inverse estimate is initialized in a small region around the home posture, which we set to $q^{home} = (0.0, 0.0)$. The next images show the progress of the method after several epochs. The aim is to control the effector's height within the full range from -1.0m to 1.0m. Initially, only heights around $f(q^{home}) = 0$m are reachable.

3. Experiments

In this section we present results of Goal Babbling for various growth patterns on a simulated 5 DOF robot arm and a simulated growth on the iCub humanoids robot. In all experiments we use polynomial regression (Poggio and Girosi, 1990) to represent the inverse estimate $g(x^*, \theta)$. The input vector $x \in \mathbb{R}^n$ is expanded by a feature mapping $\Phi^P(x) \in \mathbb{R}^p$ which calculates all polynomial terms of the entries of x. Thereby we use a polynomial degree of $P = 3$. A standard linear regression with parameters $\theta = \mathbf{M}$ operates on these features:

$$g(x^*, \mathbf{M}) = \mathbf{M} \cdot \Phi^P(x^*), \quad \mathbf{M} \in \mathbb{R}^{p \times m}. \qquad (7)$$

The entries of the regression matrix $\mathbf{M}$ are adapted during learning with a gradient descent of the weighted command error as defined in equation 6. Before exploration and learning proceed, we first set $\mathbf{M}$ to zero and make some random adaptions such that $g(x^*, \mathbf{M})$ produces joint angles in a range of 0.1 radian around the home posture. For the exploration we use linear disturbance terms:

$$E^v(x) = \mathbf{A} \cdot x + b, \quad \mathbf{A} \in \mathbb{R}^{m \times n}, b \in \mathbb{R}^m \qquad (8)$$

The values of $\mathbf{A}$ and b are chosen randomly, such that the disturbance never exceeds a range R within the bounded set of target positions $\mathbf{X}^*$:

$$E^v(x) = (e_1, ..., e_m)^T, \ |e_i| <= R \ \forall \ i = 1 \ldots m, x \in \mathbf{X}^*$$

3.1 2D proportional Growth

We start with a simulated arm with five joints ($m = 5$). The aim is to reach for a subset of all reachable positions ($n = 2$) as shown in Fig. 3a. The home posture has a slightly curved shape. A new sequence of targets x_t^* is generated in each epoch. $K = 15$ positions $x_{k \cdot L}^* \in \mathbf{X}^*, k = 0 \ldots K-1$ are randomly selected from the target grid shown in Fig. 3b. One after the other is connected by a linear target motion with $L = 7$ intermediate target positions ($l = 0 \ldots L-1$):

$$x_{k \cdot L + l}^* = \frac{L - l}{L} \cdot x_{k \cdot L}^* + \frac{l}{L} \cdot x_{(k+1) \cdot L}^* \qquad (9)$$

We used an exploration range of $R = 0.2$ and $V = 20$ variations. In the first epoch the arm is $0.5m$ long. During the first T_G epochs, the arm's length increases to $1.0m$ linearly and remains constant afterwards. All segments of the arm grow proportionally. For the duration of 100.000 epochs we evaluated the absolute performance error, indicating the average positioning error in meters:

$$E_{abs}^X(\theta) = \frac{1}{N} \sum_t (x_t - x_t^*)^2 \qquad (10)$$

Results for different growth periods $T_G \in \{10, 100, 1000, 10000\}$ are shown in Fig. 4. As reference, performance trajectories are also shown for Goal Babbling without growth: one with a constant short length ($0.5m$) and one with a constant length of $1.0m$. Both with and without growth the Goal Babbling algorithm finds an accurate solution with an accuracy of approx. $2cm$. A significant temporary increase of the performance error can be seen for $T_G \in \{10, 100, 1000\}$ during the growth period. This effect does, however, not display a degeneration of the inverse model, but is caused by more distant targets. During growth, we scale the set of target positions together with the overall length of the arm (compare Figures 3a and 3b). A linear scaling of the targets therefore automatically increases the absolute performance error. In order to account for this effect, we also evaluated the relative performance error, i.e. the position error relative to the body size:

$$E_{rel}^X(\theta) = \frac{1}{N} \sum_t \frac{(x_t - x_t^*)^2}{S} \qquad (11)$$

where we write $S = 1.0$ for the full length of the arm and $S = 0.5$ for half of the length correspondingly. Results are shown in Fig. 4b. The graph shows a very smooth blending between the performance errors of the short reference curve and long reference curve for very rapid growth ($T_G = 10$) as well as slow growth ($T_G = 10000$). There is no degeneration and the error decreases continuously, resulting in an accurate solution. An example of the performance after proportional growth is shown in Fig. 3b.

3.2 2D un-proportional Growth

The results for proportional growth show that our Goal Babbling algorithm can deal with growth of different speeds on the fly. Proportional growth is rather simple in the sense, that such growth could be equally compensated by an interal re-scaling of the effector positions. If all segments are scaled by a factor α, also the effector positions scale with α. Of course, this knowledge is not build into the algorithm and the inverse estimate has to change constantly while the body grows. However, we can complement these results with un-proportional growth. In this case, we let only the first, third and fifth segment grow (see Fig. 3c). We repeated the same evaluation for un-proportional growth as for the previous, proportional pattern. The results are shown in Fig. 5. The learning curves have the exact same characteristics as in the previous experiment and the bootstrapping is successful in all cases.

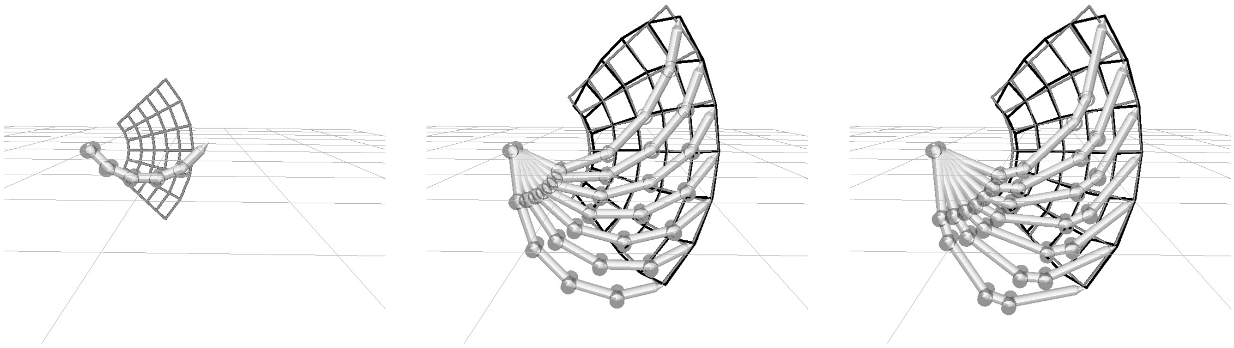

(a) Initial conditions. (b) Result after proportional growth. (c) Result after un-proportional growth.

Figure 3: Goal Babbling on a growing 5 DOF arm. Grey grids show the set of target positions. Black grids show the finally reached positions. Several postures produced by the inverse estimate show how the redundancy is resolved.

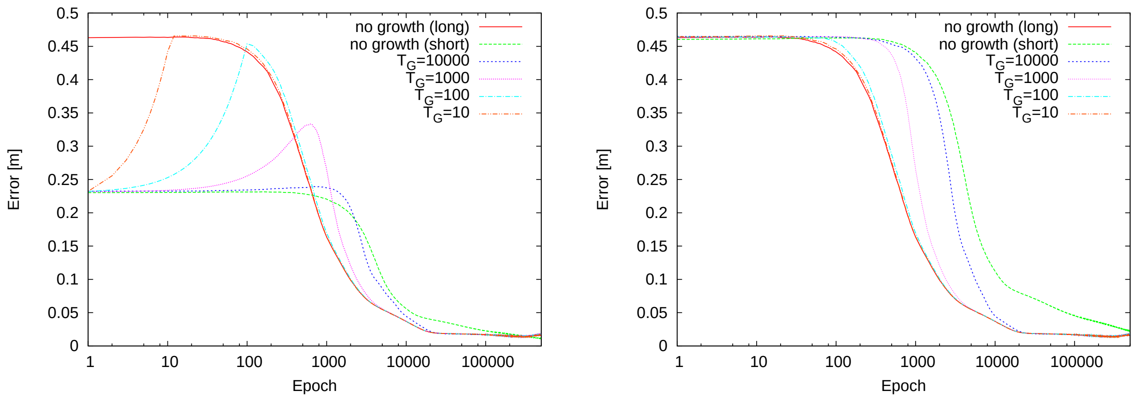

(a) Absolute performance error E^X_{abs} (b) Relative performance error E^X_{rel}

Figure 4: Goal Babbling performance on a planar arm $(n = 2, m = 5)$ for *proportional* segment growth from a total length of 0.5m to 1.0m. Results are averaged over 20 independent trials.

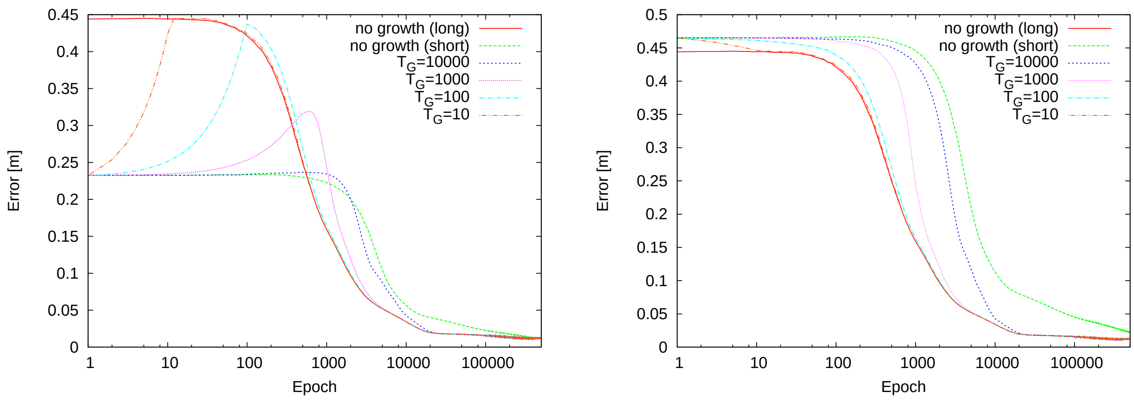

(a) Absolute performance error E^X_{abs} (b) Relative performance error E^X_{rel}

Figure 5: Goal Babbling performance on a planar arm $(n=2, m=5)$ for *un-proportional* segment growth from a total length of 0.5m to 1.0m. Results are averaged over 20 independent trials.

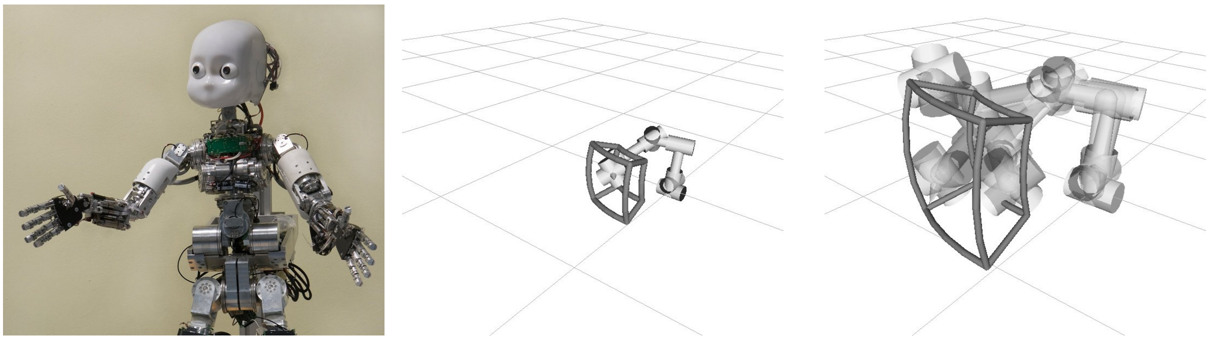

(a) iCub (b) Initial conditions. (c) Result after proportional growth.

Figure 6: Goal Babbling on a growing iCub

3.3 iCub Growth

We complete our experimental evaluation with a kinematic simulation of the humanoid robot iCub (Tsagarakis et al., 2007) (see Fig. 6a). The goal is the 3D position control of the right hand ($n = 3$). Ten degrees of freedom ($m = 10$) are controlled: three degrees of freedom in the hip, three in the shoulder, one in the elbow and three in the forearm. The model therefore resembles iCub's kinematic chain reaching from the hip to the right hand[1]. Since the ranges of the possible angles differ significantly between different joints, we normalize the range to $q_i \in [-1.0;\ 1.0]\ \forall i = 1...10$. Initially, the kinematic model has only half of iCub's original size (see Fig. 6b). During the growth period, all joints grow proportionally to the full size.

For iCub we used $V = 20$ variations with range $R = 0.1$. Target sequences are generated between $K = 30$ different targets with $L = 10$ intermediate steps each. Fig. 7 shows results over time for the same growth periods as used in the planar arm experiments. Again, the change induced by growth is clearly visible by an increase of the absolute performance error. However, the bootstrapping is successful as the error reaches a stable, low level. The relative performance errors the same smooth blending between the two reference curves that has been observed on the planar arm. The final performance errors are in the range of 6mm. Fig. 6c shows several postures produced by an inverse estimate in order to show how the redundancy is resolved. The results show that Goal Babbling also allows to track growth on a complex, three dimensional morphology.

4. Discussion

Our Goal Babbling algorithm (Rolf et al., 2010) allows to bootstrap reaching skills in an entirely goal-directed manner. We have shown that different growth rates and unproportional growth patterns during the bootstrapping can be handled and presented results for the infant-like humanoid iCub.

Bodily change obviously requires to keep on learning once a sensorimotor skill is established. The hierarchical, lifelong acquisition of new skills "on top" of each other (Demiris and Dearden, 2005, Prince et al., 2005) is only possible, if the performance of older skills is stable. Given an already developed skill, models that include a later on fine-tuning can explain the adaptation to growth. But what if a skill is not yet fully developed and the body undergoes continuous change? The batch character of staged models lacks the flexibility to deal with such changes. In contrast, goal-directed exploration is entirely incremental which permits adaptation to changes at any time – without explicit knowledge of when changes occur and what those changes are.

While models with separate exploration might be formulated more incrementally, the major explanatory strength of Goal Babbling is its simplicity and efficiency. The model can explain how infants learn to reach with a single mechanism: the only necessity is to keep on trying to reach. The approach is thereby not only compatible with the evidence of early goal-directed action in infants. It explains the function and relevance of these reaching attempts because this strategy leads to the successful acquisition of skills. Trying to reach is a very efficient strategy as it allows to focus on behaviorally relevant motor commands. It can explain how infants can explore their very high-dimensional motor system, while unstructured approaches like motor babbling are unapplicable in high-dimensional domains. This efficiency contributes to the mastery of growth because only relevant motor commands need to be re-explored once the body has changed.

Acknowledgements

Matthias Rolf gratefully acknowledges the financial support from Honda Research Institute Europe for

[1]For the kinematic parameters see: `http://eris.liralab.it/wiki/ICubFowardKinematics_right`

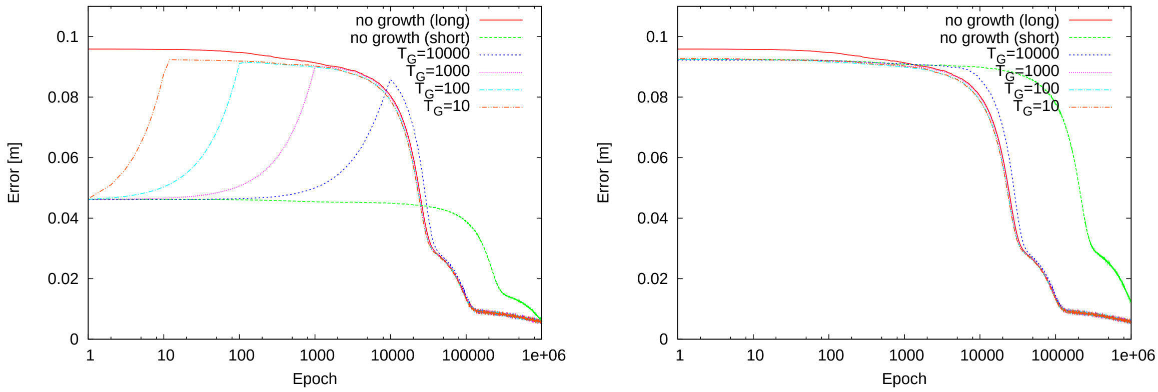

(a) Absolute performance error E^X_{abs} (b) Relative performance error E^X_{rel}

Figure 7: Goal Babbling performance on growing iCub simulation ($n = 3, m = 10$). Results are averaged over 5 independent trials. The inverse estimate uses hip and right arm joints in order to position the right hand.

the Project "Neural Learning of Flexible Full Body Motion".

References

Bullock, D., Grossberg, S., and Guenther, F. H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Cognitive Neuroscience*, 5(4).

Demiris, Y. and Dearden, A. (2005). From motor babbling to hierarchical learning by imitation: A robot developmental pathway. In *EpiRob*.

D'Souza, A., Vijayakumar, S., and Schaal, S. (2001). Learning inverse kinematics. In *IROS*.

Jordan, M. and Rumelhart, D. (1992). Forward models: supervised learning with distal teacher. *Cognitive Science*, 16:307–354.

Kawato, M. (1990). Feedback-error-learning neural network for supervised motor learning. In *Advanced Neural Computers*. Elsevier.

Kuczmarski, R., Ogden, C., and Guo, S. (2002). 2000 cdc growth charts for the united states. *Vital Health Stat*, 11(246).

Meltzoff, A. and Moore, M. (1997). Explaining facial imitation. *Early Develop. and Parenting*.

Oyama, E. and Tachi, T. (2000). Goal-directed property of online direct inverse modeling. In *IJCNN*.

Piaget, J. (1953). *The Origin of Intelligence in the Child*. Routledge and Kegan Paul.

Poggio, T. and Girosi, F. (1990). Networks for approximation and learning. *Proc. IEEE*, 78(9).

Prince, C., Helder, N., and Hollich, G. (2005). Ongoing emergence: A core concept in epigenetic robotics. In *EpiRob*.

Rolf, M., Steil, J. J., and Gienger, M. (2009). Efficient exploration and learning of whole body kinematics. In *ICDL*.

Rolf, M., Steil, J. J., and Gienger, M. (2010). Goal babbling permits direct learning of inverse kinematics. *IEEE Trans. Auto. Mental Development*, 2(3).

Sanger, T. D. (2004). Failure of motor learning for large initial errors. *Neural Computation*, 16(9).

Tsagarakis, N., Metta, G., and Sandini (2007). icub: the design and realization of an open humanoid platform. *Advanced Robotics*.

von Hofsten, C. (1982). Eye-hand coordination in the newborn. *Developmental Psychology*, 18(3).

von Hofsten, C. (2004). An action perspective on motor development. *Trends in Cog. Sci.*, 8(6).

Wells, J., Hyler-Both, D., and Danley, T. (2002). Biomechanics of growth and development in the healthy human infant. *The Journal of the American Osteopathic Association*, 102(6).

Wolpert, D., Miall, R. C., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cog. Sci.*, 2(9).