



HAL
open science

Sensorimotor learning in a Bayesian computational model of speech communication

Marie-Lou Barnaud, Jean-Luc Schwartz, Julien Diard, Pierre Bessière

► **To cite this version:**

Marie-Lou Barnaud, Jean-Luc Schwartz, Julien Diard, Pierre Bessière. Sensorimotor learning in a Bayesian computational model of speech communication. ICDL-EpiRob 2016 - 6th Joint IEEE International Conference Developmental Learning and Epigenetic Robotics, Sep 2016, Cergy-Pontoise, France. hal-01371719

HAL Id: hal-01371719

<https://hal.science/hal-01371719v1>

Submitted on 27 Sep 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sensorimotor learning in a Bayesian computational model of speech communication

Marie-Lou Barnaud
and Jean-Luc Schwartz
Univ. Grenoble Alpes, Gipsa-lab,
CNRS, Gipsa-lab,
F-38000 Grenoble, France
marie-lou.barnaud@gipsa-lab.grenoble-inp.fr

Julien Diard
Univ. Grenoble Alpes, LPNC,
CNRS, LPNC,
F-38000 Grenoble, France

Pierre Bessière
CNRS - SORBONNE Universités
UPMC - ISIR, Paris, France

Abstract—Although sensorimotor exploration is a basic process within child development, clear views on the underlying computational processes remain challenging. We propose to compare eight algorithms for sensorimotor exploration, based on three components: “*accommodation*” performing a compromise between goal babbling and social guidance by a master, “*local extrapolation*” simulating local exploration of the sensorimotor space to achieve motor generalizations and “*idiosyncratic babbling*” which favors already explored motor commands when they are efficient. We will show that a mix of these three components offers a good compromise enabling efficient learning while reducing exploration as much as possible.

I. INTRODUCTION

Acting efficiently in the environment requires a precise knowledge of the relationship between motor commands and sensory outputs. Cognitive theories of motor control as well as computational approaches in robotics generally consider that agents have at their disposal internal models linking motor and sensory information [1]. There are two kinds of internal models: forward and inverse models. To implement a forward model, cognitive or computational agents have to learn to predict the sensory output s corresponding to a given motor command m . Conversely, within an inverse model, agents must learn to select motor commands m corresponding to a given sensory output s .

Sensorimotor learning, for either forward or inverse models, typically involves sensorimotor exploration supposed to occur at an initial stage in child development, and accordingly called “babbling”. Exploration enables the agent to visit parts of the sensorimotor environment, so as to be able to learn the association between sensory and motor variables inside the visited regions, and to hopefully be able to generalize the learned patterns to other regions of the sensory or motor space. The key question in this process is to select adequate regions to explore, so as to make the learning process as rapid and efficient as possible.

There are two main families in sensorimotor exploration algorithms: motor and goal babbling. In motor babbling, the agent selects a motor command m and observes its sensory effect s . In goal babbling, the agent selects an effect s and tries to find a motor command m able to generate this effect. This

research can be performed by some kind of estimation process, like inference [2] or optimization [3]–[5], among others.

Both approaches can be decomposed in two steps: element selection and learning. In the first step, the agent selects an element to learn: a motor command m in motor babbling or a sensory percept s in goal babbling. The choice of adequate elements to select in order to increase learning speed and accuracy is crucial. Indeed, selection can be random but a number of studies show that a thoughtful choice can considerably improve learning [2], [6], [7].

This choice can be made internally according to criteria decided by the agent, or externally tuned by a master or by social guidance. Internal criteria should enable the agent to develop adequate sensorimotor exploration according to considerations about the interest of some regions of the space, e.g. for improving the performance of the internal models or to increase the knowledge of the possible relationships between motor and sensory variables [6], [8].

Social guidance uses the experience of other agents to guide learning. For instance, selection can be supervised by an imitation process during which a learning agent tries to reproduce elements provided by a teacher agent.

Another dimension characterizing the selection step concerns whether or not it is influenced by previously selected items. When it is the case, it is called active learning [8], [9].

After the selection step, an (m, s) pair is available. It is used by the agent to update parameters of its internal models; this is the learning step proper.

Despite its apparent simplicity, sensorimotor exploration is a challenging problem, particularly if the sensorimotor space to explore is dimensionally large. In this paper, in the context of sensorimotor exploration in speech communication, we present a new algorithm called *accommodation goal babbling* (AGB) under social guidance by a master. We also introduce two additional mechanisms likely to improve exploration and learning: *local extrapolation* which generalizes sensorimotor observations to a neighborhood of each selected motor command, and *idiosyncratic babbling* which aims at favoring motor commands that are already known to efficiently approximate the selected sensory goal.

We evaluate the efficiency of the three proposed components, in comparison with classical sensorimotor exploration algorithms, around three questions:

- Does the AGB algorithm efficiently learn the motor-to-sensory relationship in the portions of the sensory space provided by the master?
- Does local extrapolation improve learning?
- Does idiosyncratic babbling improve learning?

The proposed algorithms are implemented in the framework of a Bayesian model of speech communication, named COSMO, in which we study the learning of the relationship between articulatory motor gestures and their corresponding auditory outputs (typically formants of an auditory signal).

In Section II, we detail the different algorithms. In Section III, we present the Bayesian model COSMO and provide implementation details. In Section IV, we describe experimental results of the comparison of the algorithms we study. Finally, we discuss these results and give some limitations and possible improvements in Section V.

II. SENSORIMOTOR LEARNING

Within a Bayesian framework, we consider a motor space M composed of motor commands m and a sensory space S composed of sensory percepts s ; both spaces are described as probabilistic variables. Learning agents are defined by the joint probability distribution $P(S | M)$, decomposed as a product of a forward internal model $P(S | M)$ and a prior distribution on motor gestures $P(M)$. We note Prod the algorithm that performs the transformation from a motor gesture m to a sensory percept s :

$$s = \text{Prod}(m) . \quad (1)$$

Prod is the “ground-truth” transformation that learning agents aim at identifying.

In this paper, we implement altogether eight algorithms, all based on the two steps described before: selecting an element to learn and determining a method to learn it.

A. Random motor babbling and random goal babbling

We first propose two exploration-by-babbling algorithms providing a baseline for further comparisons. In the *random motor babbling* (RMB) algorithm, the agent selects a uniformly random motor command m . In the *random goal babbling* (RGB) algorithm, the agent selects a uniformly random sensory percept s . Then, it infers a motor command m thanks to $P(M | [S = s])$ where:

$$P(M | [S = s]) \propto P(M)P([S = s] | M) . \quad (2)$$

In both cases, the agent then “performs” m , which is transformed by Prod (Eq. 1), resulting in a sensory output s' . This provides a couple $\langle m, s' \rangle$, which enables the learning agent to update its forward model $P(S | M)$.

Notice that this implementation of RGB follows a recent definition of goal babbling [2], that does not use a distance measure to compare s and s' , contrary to previous versions [3], [4]. Indeed, the agent just modifies its forward model based

on the observed $\langle m, s' \rangle$ couple whatever the success of the inference process (described in Eq. (2)). This opportunistic behavior will also be a crucial component of the accommodation algorithm introduced in the next section.

B. Accommodation and idiosyncratic goal babbling

The proposed algorithm, called *accommodation goal babbling* (AGB), can be seen as goal babbling with social guidance. The agent receives a sensory stimulus s provided by the master. Then, the inference and learning processes are identical to the ones in the random goal babbling algorithm. We expect that the agent should first draw motor commands uniformly randomly at the beginning of the learning process. Then, it should progressively focus on sensory stimuli provided by the master and only improve its forward model around the corresponding sensory regions.

On this basis, we add an “idiosyncratic babbling” component. The model, called *idiosyncratic goal babbling* (IGB), is exactly the same as the previous one (AGB) except that once the $\langle m, s' \rangle$ couple is obtained, the agent updates not only its forward model $P(S | M)$ thanks to the couple $\langle m, s' \rangle$, but also its prior distribution of motor gestures $P(M)$ thanks to m . We expect that motor exploration should focus around adequate motor commands already selected, and therefore convergence could be quicker, by keeping an already found good motor solution (possibly idiosyncratic, that is, different from one agent to another), rather than searching for all possible solutions ¹.

C. Local extrapolation

Finally, we explore a variant of each of these four algorithms incorporating a “local extrapolation” component. This consists in updating $P(S | M)$ with couples $\langle m_i, s' \rangle$ at each learning step, where m_i corresponds to several motor commands in the neighborhood of m , and where s' is the sensory output computed with Prod for m . We expect that local extrapolation should increase the speed and efficiency of exploration around the target regions provided by the master by providing motor generalization around the explored motor regions, at the cost of potential inaccurate local learning. Indeed, associating m_i with s' might be incorrect with respect to Prod; the larger the generalization neighborhood volume or the less smooth Prod around m , the greater this inaccuracy. We hence obtain four new algorithms with this variant: RMB-LE (for RMB with Local Extrapolation), RGB-LE, AGB-LE, and IGB-LE.

III. APPLICATION TO A BAYESIAN MODEL OF COMMUNICATION

In this paper, we focus on sensorimotor exploration occurring during speech development. More specifically, the algorithms we compare are incorporated into a Bayesian model of communication named COSMO (for “Communicating Objects using Sensory-Motor Operations”) that we have been developing in the last years. This model has already been applied to

¹Assuming here, as is commonly the case, a many-to-one motor-to-sensory transformation Prod.

the emergence of sound systems in human languages [10], the study of speech perception in adverse conditions [11], [12], and it is currently expanded towards modeling speech production [13]. Here, we only implement the portion of the COSMO model involving sensorimotor exploration and we test our algorithms in the context of learning vowel production.

A. Model definition

The motor space M corresponds to articulatory configurations of the vocal tract. Configurations and commands are defined from a realistic articulatory model of the vocal tract, called VLAM (for “Variable Linear Articulatory Model”) [14]–[16], which can transform configurations of articulatory positions into sounds. Each configuration is described by seven parameters, but only three are essential to model vowels: lip height (L_H) controlling the distance between the lips, tongue body (T_B) corresponding to the advancement/retraction of the tongue, and tongue dorsum (T_D) controlling the tongue height.

Sensory percepts S are described by components of the acoustic signal, called formants, corresponding to frequency peaks of the signal spectrum. We keep only the two first formants, $F1$ and $F2$, which are essentially sufficient to characterize vowels.

The motor and sensory spaces are discretized, so that sensorimotor exploration occurs inside a motor space M containing 15,625 articulatory positions ($25 \times 25 \times 25$) and a sensory space S containing 4,703 (59×73) couples of formants $\langle F1, F2 \rangle$. In a realistic fashion, there are more motor than sensory configurations and different motor configurations may result in the same sensory effect (this is referred to as the “many-to-one problem”).

As described in Section II, learning agents have a forward model $P(S | M)$ and a prior about motor commands $P(M)$. The conditional probability distribution $P(S | M)$ is defined by a set of Gaussian distributions²: one for each motor gestures, that is, 15,625 Gaussian distributions, with a mean μ and a covariance matrix Σ for each distribution. μ is initialized in the middle of the sensory space and with a large variance, so as to approximate a uniform distribution. In the RMB, RGB, AGB and IGB algorithms, $P(S | M)$ is updated with the couple $\langle m, s' \rangle$ at each learning step. This update is weighted by a constant value equal to 1. In algorithms with local extrapolation (RMB-LE, RGB-LE, AGB-LE and IGB-LE), $P(S | M)$ is updated at each learning step with a set of couples $\langle m_i, s' \rangle$, which are also weighted: weights follow a Gaussian distribution around mean m with a small variance. This allows fair comparison between variants with and without local extrapolation: both receive experimental observations weighted, overall, by 1, and this unitary weight is concentrated in a single m value without local extrapolation, and spread over the neighborhood of m with local extrapolation.

In IGB and IGB-LE, the prior $P(M)$ is represented by a histogram over the 15,625 motor values, with one bin for each

motor value. This histogram is initialized as uniform with an initial frequency f_i and updated with a weight 1 for each incoming observation, at each learning step. We tested f_i for three values : 1, 0.1 and $1/15625$. In all other algorithms, $P(M)$ does not change from a uniform distribution.

B. Selection and production phases

Motor commands and sensory percepts are uniformly randomly selected respectively in the motor babbling RMB/RMB-LE and the goal babbling RGB/RGB-LE algorithms. On the contrary, AGB/AGB-LE and IGB/IGB-LE algorithms use the sensory goals s provided by a master. More precisely, the master follows a probability distribution of sensory percepts corresponding to the vowels /a, i, u/ (e.g. in American English, /a/ as in “far”, /i/ as in “feed” and /u/ as in “food”). We choose these vowels as they correspond to vowels present in nearly all languages [17].

C. Experimental evaluation

We simulate sensorimotor exploration for 100,000 learning steps. At the end of the learning process, we evaluate our algorithms in two manners: quality of learning and amount of exploration.

Quality of learning is assessed by the capacity of the learning agent to reproduce sounds provided by its master. This aims at checking that the agent has learned the forward model $P(S | M)$ accurately in the useful sensory regions, that is, where vowels lie. Therefore, a confusion matrix $P(S_{prod} | S_{given})$ is computed, in which S_{given} is a given sound to reproduce and S_{prod} is the sound achieved by the agent at the end of a given sensorimotor learning algorithm, when it tries to reproduce S_{given} . This matrix can be computed with Bayesian inference:

$$P(S_{prod} | S_{given}) = \sum_M P(S_{prod} | M)P(M | S_{given}), \quad (3)$$

where $P(S_{prod} | M)$ is the real production system in the environment, computed thanks to the VLAM model, and $P(M | S_{given})$ is the inversion of the forward model of the learning agent. Thereafter, we only conserve the diagonal of the confusion matrix, noted $\text{Diag}_{\text{ConfMat}}$, giving for each sensory value the probability that it is correctly reproduced. From $\text{Diag}_{\text{ConfMat}}$ evaluated at each learning step, we look at sensory values drawn according to the master distribution and compute their weighted average providing an accuracy measure. The contrary probability gives us an error measure.

For quantifying exploration, we count the number of different selected motor commands m to assess the relative volume of the motor space explored during learning. For algorithms with local extrapolation (RMB-LE, RGB-LE, AGB-LE and IGB-LE), we consider that neighborhood exploration does not increase this count since the total volume of local explored configurations was controlled to count for 1 (see Section III-A).

²More precisely, as S is a discrete space, $P(S | M)$ corresponds to truncated, bell-shaped distributions that approximate Gaussian distributions.




Fig. 1: Error measure (Top) and amount of exploration (Bottom) as a function of time, during learning, for RMB/RMB-LE, RGB/RGB-LE and AGB/AGB-LE algorithms. See text for the definitions of the error and exploration measures.

IV. RESULTS

Let us describe how accommodation, idiosyncratic babbling and local extrapolation affect both quality of learning and the amount of exploration during sensorimotor learning.

A. Accommodation: comparing AGB with RMB and RGB

Let us first compare the AGB algorithm to RMB and RGB. We display the evolution of the error measure during learning for each algorithm (colored dotted lines in the top of Figure 1). We observe that, although error measures all appear to converge towards zero, convergence is quicker for RGB (around 45,000 learning steps) than for AGB (around 80,000 learning steps) and for RMB (approaching convergence at the end the 100,000 simulated learning steps).

In terms of exploration, we compare the number of selected motor commands m for the three algorithms during learning (colored dotted lines in the bottom of Figure 1). Overall, the three algorithms explore all the motor space M at the end of the learning phase, though a little faster, that is with more exploration, at each step for RGB.

Let us now focus on the portions of the motor space explored during learning. As a reference, consider the distributions of the sensory stimuli provided by the master (top left of Figure 2); there are three regions of the sensory space with high probabilities, that correspond to the three vowels /a, i, u/ (/i/ on the top left, /a/ on the bottom and /u/ on the top right).

Then, we display the values of $\text{Diag}_{\text{ConfMat}}$ for all sensory values at the end of learning simulations (top right and bottom




Fig. 2: **Top Left:** Distribution of the sensory stimuli provided by the master. Values of $\text{Diag}_{\text{ConfMat}}$ represented in the sensory space for **(Top right)** the AGB algorithm, **(Bottom left)** the RGB algorithm and **(Bottom right)** the RMB algorithm. The x -axis is reverse $F2$ and the y -axis is reverse $F1$. Values spread from dark blue (low probability) to red (high probability)

plots of Figure 2). Non-zero values form a rounded triangle in the sensory space. This corresponds to the classical vocalic space, that is, the portion of the sensory space which can be produced by motor gestures. In this area, the AGB algorithm is especially efficient on regions corresponding to the master's stimuli, with high probability in the areas corresponding to /a/ and /i/, and to a lesser extent for /u/. The RMB and RGB algorithms are naturally more diffuse since they do not rely on the master's guidance. Interestingly, they focus more on the left part of the vocalic space, especially RMB. This suggests that right portions of the sensory space are harder to learn: there are more motor gestures corresponding to the left portions of the vocalic space (probably due to redundancy in the motor space). Therefore, vowels /a/ and /i/ are easier to find than /u/, which probably explains why the AGB algorithm displays higher probability on the two vowels /a/ and /i/.

The proposed interpretation is that, at the onset of learning, all algorithms first explore more or less randomly the motor space in the same way until around 20,000 learnings steps. As more motor gestures correspond to the left portion of the vocalic space, they have higher probability in this area. Then, RGB continues to explore randomly and constantly improves its performance, resulting in Figure 1 in a linear decrease of error towards convergence. At the end of learning, it is efficient on the whole vocalic space but keeps higher performance on the left area, that was more explored in the initial learning phase (Figure 2). Meanwhile, the AGB algorithm succeeds to focus on the master's portions of the space, which reduces its exploration but somewhat slows down its error decrease

(Figure 1). Finally, since the RMB algorithm randomly selects motor gestures, it has to deal with the redundancy of the motor space, hence it becomes efficient on highly redundant portions of the space but decreases slowly its error on the remaining space (Figures 1 and 2).

In summary, the AGB algorithm succeeds to focus on the master’s sensory stimuli. However, it obtains similar performance as the RGB and RMB algorithms at the end of learning, and it actually converges more slowly than the RGB algorithm.

B. Local extrapolation: comparing RMB, RGB, AGB with RMB-LE, RGB-LE, AGB-LE

The top plot of Figure 1 also shows error during learning for the three variants RMB-LE, RGB-LE, AGB-LE, i.e., that include local extrapolation. We observe that error decrease for these algorithms appear exponential rather than linear, and that they converge faster than without local extrapolation (RMB, RGB, AGB), with convergence around 20,000 learning steps. However, this is at the expense of a residual error which remains around 20% higher than without local exploration. At the bottom of Figure 1, we observe that AGB-LE explores marginally less than the other algorithms.

Altogether, this suggests that adding a local extrapolation component somewhat accelerates learning, though of course at the cost of a weak error at the end of the learning phase. This sustained error is due to the update of the forward model with $\langle m_i, s' \rangle$ pairs where s' actually corresponds to the production of m and not to the production of the actual m_i .

C. Idiosyncratic behavior: comparing IGB/IGB-LE with AGB/AGB-LE

In the same manner as in Figure 1, we compare in Figure 3 AGB/AGB-LE with IGB/IGB-LE algorithms. The curves for the AGB/AGB-LE algorithms are the same as in Figure 1 and are used here as a reference. We display three versions for the IGB/IGB-LE algorithms differing in their initial frequency f_i .

We observe that for a high initial frequency f_i (increasing the initial trend for uniform M selection) performance and exploration are close to those with the AGB/AGB-LE algorithms, particularly for IGB-LE. With smaller initial frequencies (increasing the role of idiosyncrasies), the IGB/IGB-LE algorithms explore less but are also less efficient.

A more detailed comparison shows that the ratio between performance and exploration is different between the AGB-LE and IGB-LE algorithms. Indeed, the IGB-LE algorithm achieves performance close to AGB-LE by exploring much less. For instance, with f_i equal to 0.1, IGB-LE is 10% less efficient but performs 30% less exploration of the motor space.

This suggests that an adequately tuned idiosyncratic babbling, coupled with local extrapolation, could offer a good compromise between performance and motor space exploration for sensorimotor learning.

V. CONCLUSION

In this paper, we implemented and compared eight algorithms for sensorimotor learning within the Bayesian model




Fig. 3: Error measure (**Top**) and amount of exploration (**Bottom**) as a function of time, during learning, AGB/AGB-LE algorithms and IGB/IGB-LE algorithms with different initial frequency f_i .

COSMO. We introduced and tested three components for sensorimotor learning: “accommodation babbling”, “local extrapolation” and “idiosyncratic babbling”.

The first result is that the agent efficiently learns the sensory stimuli provided by its master with the AGB algorithm. Importantly, the agent mainly focuses on regions relevant to the master, after however acquiring some knowledge of the other sensory regions thanks to the initial, uninformed stage of the learning process. The second result is that the addition of a “local extrapolation” process (RMB-LE, RGB-LE, AGB-LE) accelerates learning: algorithms converge more quickly, at the expense of a small residual error at the end of learning. The third and last result suggests that “idiosyncratic babbling” coupled with “local extrapolation” (IGB-LE) offers leverage towards a good compromise. IGB-LE provides an efficiency similar to simple “local extrapolation” algorithms, but with less exploration of the motor space.

Of course, these results depend on theoretical and implementation choices which can be discussed. Let us notice four points.

First, we have chosen to learn a forward model $P(S | M)$. Alternatives would be to learn an inverse model or a global internal model of sensorimotor representations [2]. Formally, and without simplifying assumptions, these three representations are equivalent:

$$P(S | M) = P(S)P(M | S) = P(M)P(S | M) .$$

We use a forward model for implementation reasons. Indeed, because of the many-to-one problem, the distributions

$P(M | S)$ involved in an inverse model are likely to be complex and presenting several modes, requiring sophisticated parametric forms, difficult to specify on theoretical grounds. For the same reasons, the number of parameters of $P(S | M)$ can be difficult to evaluate a priori, since a given motor gesture results in a single sensory output. The Gaussian parametric form of $P(S | M)$ appears simpler and appropriate.

Second, “accommodation babbling” relies on social guidance by a master, enabling the learning agent to focus on relevant sensory regions. This choice is reasonable in the framework of speech development. Indeed, children focus their learning on the phoneme repertoire of their native language (as displayed by perceptual narrowing, [18]). Of course, internal motivation could be used as a complement of social guidance, in a hybrid algorithm; how social guidance and self-motivated explorations combine in speech development in children is an open question.

Third, “local extrapolation” is built on the hypothesis that one can associate motor gestures m_i , close to a selected motor gesture m , with the sensory output s' produced with the selected m . We have observed that this motor smoothing process can accelerate learning, although it induces a residual error at the end of the learning process. This could be improved by conceiving a dynamic process in which “local extrapolation” would only appear at the beginning of the learning phase to accelerate exploration of the sensorimotor space and then decrease progressively along learning, in a manner similar to classical simulated annealing [19].

Finally, “idiosyncratic goal babbling” appears as a promising component for sensorimotor learning. However, we focused here on a specific implementation of how the prior about motor gestures $P(M)$ is learned. For instance, we considered that this distribution was uniform at the onset of the learning process, which is of course questionable. We have seen that vowel /u/ seemed more difficult to learn than vowels /a/ and /i/, since less motor gestures are related to this vowel. However, /u/, as /a/ and /i/, is a quasi universal. It could be assumed that motor gestures related to /u/ have a higher prior, thus facilitating their learning, e.g. because they are related to sucking behaviors in the first months of age [20].

Whatever the limitations of the studied algorithms, we believe that their comparison within our Bayesian framework is promising, and likely to shed interesting light on the computational processes at hand in sensorimotor exploration in speech development.

ACKNOWLEDGMENT

Research supported by a grant from the European Research Council (FP7/2007-2013 Grant Agreement no. 339152, “Speech Unit(e)s”).

REFERENCES

[1] M. Kawato, “Internal models for motor control and trajectory planning,” *Current opinion in neurobiology*, vol. 9, no. 6, pp. 718–727, 1999.
 [2] C. Moulin-Frier and P.-Y. Oudeyer, “Exploration strategies in developmental robotics: a unified probabilistic framework,” in *Development and Learning and Epigenetic Robotics (ICDL)*, 2013 *IEEE Third Joint International Conference on*. IEEE, 2013, pp. 1–6.

[3] M. Rolf, J. J. Steil, and M. Gienger, “Goal babbling permits direct learning of inverse kinematics,” *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 3, pp. 216–229, 2010.
 [4] —, “Online goal babbling for rapid bootstrapping of inverse models in high dimensions,” in *Development and Learning (ICDL)*, 2011 *IEEE International Conference on*, vol. 2. IEEE, 2011, pp. 1–8.
 [5] M. Murakami, B. Kroger, P. Birkholz, and J. Triesch, “Seeing [u] aids vocal learning: babbling and imitation of vowels using a 3d vocal tract model, reinforcement learning, and reservoir computing,” in *Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2015 *Joint IEEE International Conference on*. IEEE, 2015, pp. 208–213.
 [6] A. Baranes and P.-Y. Oudeyer, “Intrinsically motivated goal exploration for active motor learning in robots: A case study,” in *Intelligent Robots and Systems (IROS)*, 2010 *IEEE/RSJ International Conference on*. IEEE, 2010, pp. 1766–1773.
 [7] M. Rolf and M. Asada, “Autonomous development of goals: From generic rewards to goal and self detection,” in *Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2014 *Joint IEEE International Conferences on*. IEEE, 2014, pp. 187–194.
 [8] A. Baranes and P.-Y. Oudeyer, “Active learning of inverse models with intrinsically motivated goal exploration in robots,” *Robotics and Autonomous Systems*, vol. 61, no. 1, pp. 49–73, 2013.
 [9] R. Saegusa, G. Metta, and G. Sandini, “Active learning for multiple sensorimotor coordination based on state confidence,” in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*. IEEE, 2009, pp. 2598–2603.
 [10] C. Moulin-Frier, J. Diard, J.-L. Schwartz, and P. Bessière, “COSMO (“Communicating about Objects using Sensory-Motor Operations”): A Bayesian modeling framework for studying speech communication and the emergence of phonological systems,” *Journal of Phonetics*, vol. 53, pp. 5–41, 2015.
 [11] C. Moulin-Frier, R. Laurent, P. Bessière, J.-L. Schwartz, and J. Diard, “Adverse conditions improve distinguishability of auditory, motor, and perceptuo-motor theories of speech perception: An exploratory Bayesian modelling study,” *Language and Cognitive Processes*, vol. 27, no. 7-8, pp. 1240–1263, Sep 2012. [Online]. Available: <http://dx.doi.org/10.1080/01690965.2011.645313>
 [12] R. Laurent, J.-L. Schwartz, P. Bessière, and J. Diard, “A computational model of perceptuo-motor processing in speech perception: learning to imitate and categorize synthetic CV syllables,” in *Proceedings of Interspeech 2013*, F. Bimbot, Ed. Lyon, France: International Speech Communication Association (ISCA), Aug 2013, pp. 2796–2800. [Online]. Available: <http://hal.archives-ouvertes.fr/hal-00827885>
 [13] J.-F. Patri, J. Diard, and P. Perrier, “Optimal speech motor control and token-to-token variability: a bayesian modeling approach,” *Biological cybernetics*, vol. 109, no. 6, pp. 611–626, 2015.
 [14] S. Maeda, “Compensatory articulation during speech: Evidence from the analysis and synthesis of vocal-tract shapes using an articulatory model,” in *Speech production and speech modelling*. Springer, 1990, pp. 131–149.
 [15] L.-J. Boë and S. Maeda, “Modélisation de la croissance du conduit vocal,” in *Journées d’Études Linguistiques, La voyelle dans tous ses états*, 1998, pp. 98–105.
 [16] L. Ménard, J.-L. Schwartz, L.-J. Boë, S. Kandel, and N. Vallée, “Auditory normalization of French vowels synthesized by an articulatory model simulating growth from birth to adulthood,” *The Journal of the Acoustical Society of America*, vol. 111, no. 4, p. 1892, 2002. [Online]. Available: <http://dx.doi.org/10.1121/1.1459467>
 [17] J.-L. Schwartz, L.-J. Boë, N. Vallée, and C. Abry, “The dispersion-focalization theory of vowel systems,” *Journal of phonetics*, vol. 25, no. 3, pp. 255–286, 1997.
 [18] J. F. Werker and R. C. Tees, “Influences on infant speech processing: Toward a new synthesis,” *Annual review of psychology*, vol. 50, no. 1, pp. 509–535, 1999. [Online]. Available: <http://dx.doi.org/10.1146/annurev.psych.50.1.509>
 [19] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed. Pearson Education, 2003.
 [20] L. Boë, J. Schwartz, and R. Laboissière, “Integrating articulatory constraints in the prediction of sound structures,” in *4th Speech Production Seminar, 1st ETRW on Speech Production Modeling: from control strategies to acoustics*, 1996.