

The role of intrinsic motivations in learning sensorimotor vocal mappings: a developmental robotics study

Clément Moulin-Frier, Pierre-Yves Oudeyer

► To cite this version:

Clément Moulin-Frier, Pierre-Yves Oudeyer. The role of intrinsic motivations in learning sensorimotor vocal mappings: a developmental robotics study. Interspeech, ISCA, Aug 2013, Lyon, France. hal-00860655

HAL Id: hal-00860655

<https://hal.inria.fr/hal-00860655>

Submitted on 11 Sep 2013

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.

The role of intrinsic motivations in learning sensorimotor vocal mappings: a developmental robotics study

Clément Moulin-Frier¹, Pierre-Yves Oudeyer¹

¹INRIA, FLOWERS team, Bordeaux, France

{clement.moulin-frier, pierre-yves.oudeyer}@inria.fr

Abstract

Learning complex mappings between various modalities (typically articulatory, somatosensory and auditory) is a central issue in computationally modeling speech acquisition. These mappings are generally nonlinear and redundant, involving high dimensional sensorimotor spaces. Classical approaches consider two separate phases: a relatively pre-determined exploration phase analogous to infant babbling followed by an exploitation phase involving higher level communicative motivations. In this paper, we consider the problem as a developmental robotics one, in which an agent actively learns sensorimotor mappings of an articulatory vocal model. More specifically, we show how intrinsic motivations can allow the emergence of efficient exploration strategies, driving the way a learning agent will interact with its environment to collect an adequate learning set.

Index Terms: speech acquisition, sensorimotor learning, forward and inverse modeling, developmental robotics, intrinsic motivations, active learning, Gaussian mixture model.

1. Introduction

In the early stage of speech acquisition, the infant has to learn complex bidirectional sensorimotor mappings between various modalities: articulatory, somato-sensory and auditory spaces. The learning of such sensorimotor mappings is a central issue in most of the existing computational models of speech acquisition. For example, the DIVA model [1, 2] proposes an architecture partly inspired by neurolinguistics. The sensorimotor mapping learning phase is analogous to infant babbling and corresponds to semi-random articulator movements producing auditory and somato-sensory feedbacks. This is used to tune a neural network between representation maps. The Elija model [3] also deals with sensorimotor mapping learning, which is driven by a reward function (including sound salience and diversity, as well as articulatory effort). Such computational models of speech acquisition pre-determine the way the agent will explore its sensorimotor space to learn the corresponding mappings (e.g. DIVA involves a semi-random articulatory exploration and Elija a hand-coded reward function). Alternatively, we show in this paper how an active exploration of the sensorimotor space can significantly improve the learning process. This is to rely to the concept of *intrinsic motivations* in developmental psychology [4, 5], which refer to the ability to do something because it is inherently interesting or enjoyable.

Let us formalize the problem of learning complex sensorimotor mappings, including those involved in speech acquisition models. A learning agent interacts with a surrounding environment through motor commands M and sensory perceptions S . We call $f : M \rightarrow S$ the unknown function defining the physical properties of the environment, such that when the

agent produces a motor command $m \in M$, it then perceives $s \in S$, although (m, s) is generally only partially observable. Classical robotic problems are e.g. the prediction of the sensory effect of an intended motor command through a forward model $\tilde{f} : M \rightarrow S$, or the control of the motor system to reach sensory goals through an inverse model $f^{-1} : S \rightarrow M$. The agent has to learn such models by collecting (m, s) pairs through its interaction with the environment. These learning processes are often difficult for several reasons:

- the agent has to deal with uncertainties both in the environment and in its own sensorimotor loop;
- M and S can be highly dimensional, such that random sampling in M to collect (m, s) pairs would lead to too sparse data for an efficient learning;
- f can be strongly non-linear, such that the learning of \tilde{f} from experience is not trivial;
- f can be redundant (many M to one S), such that the learning of f^{-1} is an ill-posed problem.

When a learning process faces these issues, random motor exploration (or motor babbling) in M is not a realistic exploration strategy to collect (m, s) pairs. Due to high dimensionality, data are precious whereas, due to non-linearity and/or redundancy, data are not equally useful to learn an adequate forward or inverse model.

As we have seen, computational models of speech acquisition such as DIVA or Elija generally pre-determine the way the agent will collect (m, s) pairs. Rather than considering a pre-determined exploration, we are interested in the internal mechanisms which can drive adaptive exploration and learning at the early stage of spontaneous vocal exploration.

In the next section we describe several exploration strategies proposed in the literature to efficiently learn complex sensorimotor mappings. Section 3 is our first contribution, where we integrate these strategies into a unified probabilistic framework. In Section 4 we implement this general formal framework using Gaussian mixture models. Section 5 describes our developmental robotics experiment using an articulatory synthesizer able to compute auditory features from articulatory commands, and validates the approach by quantitatively comparing various exploration strategies to show that intrinsically motivated learning based on an empirical measure of the competence progress is an efficient way to rapidly learn to achieve a control task.

2. Exploration strategies

Two main principles can allow a learning agent to explore its sensorimotor space in order to efficiently collect (m, s) pairs. The first one concerns the space in which the agent chooses

points to explore, what we will call the *choice space*. Previous models [6, 7] have shown that learning redundant inverse models could be achieved more efficiently if exploration was driven by goal babbling in S , triggering reaching, rather than direct motor babbling in M . The reason is that goal babbling drives the agent to find the efficient motor commands allowing to reach goals uniformly distributed in a choice space corresponding to S , rather than exploring all the possible motor commands uniformly in a choice space corresponding to M . The second principle comes from the field of active learning, where exploration strategies are conceived as an optimization process, e.g. by minimizing the expected variance of a predictor in the case of forward model learning [8].

Combining both principles, recent works grounded in developmental psychology has been interested by defining *empiric measures* of interest in choice spaces. Computational studies have shown the importance of developmental mechanisms guiding exploration and learning in high-dimensional M and S spaces and with highly redundant and non-linear f [9, 7]. Among these guiding mechanisms, intrinsic motivations, generating spontaneous exploration in humans [10, 4], have been transposed in curiosity-driven learning machines [11, 12, 13] and robots [9, 7] and shown to yield highly efficient learning of inverse models in high-dimensional redundant sensorimotor spaces [7]. Efficient versions of such mechanisms are based on the active choice of learning experiments that maximize learning *progress*, for e.g. improvement of predictions or of competences to reach goals [11, 9]. This automatically drives the system to explore and learn first easy skills, and then explore skills of progressively increasing complexity. Such intrinsically motivated exploration was also shown to generate automatically behavioural and cognitive developmental structures sharing interesting similarities with infant development [9, 14, 15, 16]. This approach is grounded in psychological theories of intrinsic motivations [10, 17], explores several fundamental questions about curiosity-driven open-ended learning in robots [9], and allows to generate some novel hypotheses for the explanation of infant development, regarding behavioural [15], cognitive [14] and brain circuitry [18] issues.

Systematic comparisons of various exploration strategies have been performed [7], which differ on the way the agent iteratively collects (m, s) pairs to learn forward and/or inverse models (comparing random vs. competence progress based exploration, in either motor M or sensory S choice spaces). These strategies are summarized below (the original name of the corresponding algorithm appears in parenthesis).

- **Random motor exploration (ACTUATOR-RANDOM):** at each time step, the agent randomly chooses an articulatory command $m \in M$, produces it, observes $s = f(m)$ and updates its sensorimotor model according to this new experience (m, s) .
- **Random goal exploration (SAGG-RANDOM):** at each time step, the agent randomly chooses a goal $s_g \in S$ and tries to reach it by producing $m \in M$ using an inverse model f^{-1} learned from previous experience. It observes the corresponding sensory consequence $s = f(m)$ and updates its sensorimotor model according to this new experience (m, s) .
- **Active motor exploration (ACTUATOR-RIAC):** at each time step, the agent chooses a motor command m by maximizing an interest value in M based on an empirical measure of the competence progress in prediction in its recent experience. The agent uses a forward model

\tilde{f} learned from its past experience to make a prediction $s_p \in S$ for the motor command m . It produces m and observe $s = f(m)$. The agent updates both its sensorimotor models (forward and inverse) according to the new experience (m, s) . A measure of competence is computed from the distance between s_p and s , which is used to update the interest model in the neighborhood of m .

- **Active goal exploration (SAGG-RIAC):** at each time step, the agent chooses a goal s_g by maximizing an interest value in S based on an empirical measure of the competence progress to reach goals in its recent experience. It tries to reach s_g by producing $m \in M$ using a learned inverse model f^{-1} . It observes the corresponding sensory consequence $s \in S$ and updates its model according to this new experience (m, s) . A measure of competence is computed from the distance between s_g and s , which is used to update the interest model in the neighborhood of s_g .

In the two active strategies, the measure of interest was obtained by recursively splitting the choice space (M in ACTUATOR-RIAC, S in SAGG-RIAC) into sub-regions during the agent life. Each region maintains its own empirical measure of competence progress from its competence history in a relative time window. The competence is defined as the opposite of the distance between s_p and s in the active motor strategy, and s_g and s in the active goal one.

In a seek of unification, we can extract the following general principles from these strategies.

- Whatever the strategy used, the agent has to sample points in a given space. This space is M for the first and the third strategy, S for the second and the fourth. We call it the *choice space* X .
- In all but the first strategy, the agent has to make an inference from the choice space X to its “complement” in $M \times S$ (which is S if $X = M$ and M if $X = S$). We call this latter the *inference space* Y .
- In the active exploration strategies, the agent has to maintain an empirical measure of interest in the choice space X . In the other strategies, the agent makes a random sampling in X .

We thus suggest to classify these exploration strategies along two dimensions: the choice space X (either M or S) and the measure of interest in X (either uniform or competence progress based). The next section describes a generic exploration algorithm encompassing each particular strategy.

3. Probabilistic modeling

We use a probabilistic framework where the notations and principles are inspired by [19, 20]. Upper case A denotes a probabilistic variable, defined by its continuous, possibly multidimensional and bounded domain $\mathcal{D}(A)$. The conjunction of two variables $A \wedge B$ can be defined as a new variable C with domain $\mathcal{D}(A) \times \mathcal{D}(B)$. Lower case a will denote a particular value of the domain $\mathcal{D}(A)$. $p(A | \omega)$ is the probability distribution over A knowing some preliminary knowledge ω (e.g. the parametric form of the distribution, a learning set ...). Practically, ω will serve as a model identifier, allowing to define different distributions of the same variable, and we will often omit it in the text although it is useful in the equations. $p(A B | \omega)$ is the probability distribution over $A \wedge B$. $p(A | [B = b] \omega)$ is the conditional distribution over A knowing a particular value b of

another variable B (also noted $p(A | b \omega)$ when there is no ambiguity on the variable B).

Considering that we know the joint probability distribution over the whole sensorimotor space, $p(M S | \omega_{SM})$, Bayesian inference provides the way to compute every conditional distribution over $M \wedge S$. In particular, we can compute the conditional distribution over Y knowing a particular value x of X , as long as X and Y correspond to two complementary subdomains of $M \wedge S$ (i.e. they are disjoint and $X \wedge Y = M \wedge S$). Thus, the prediction of $s_p \in S$ from $m \in M$ in the active motor exploration strategy, or the control of $m \in M$ to reach $s_g \in S$ in the active or random goal exploration strategies, correspond to the probability distributions $p(S | M \omega_{SM})$ and $P(M | S \omega_{SM})$, respectively. More generally, whatever the choice and inference spaces X and Y , as long as they are disjoint and that $X \wedge Y = M \wedge S$, Bayesian inference allows to compute $p(Y | X \omega_{SM})$.

Such a probabilistic modeling is also able to express the interest model, that we will call ω_I , such that the agent draws points in the choice space X according to the distribution $p(X | \omega_I)$. In the random motor and goal exploration strategies, this distribution is uniform, whereas it is a monotonically increasing function of the empirical interest measure in the case of the active exploration strategies. We will provide more details about the way to iteratively compute $p(M S | \omega_{SM})$ and $p(X | \omega_I)$ from the experience of the agent in the next section.

Given these probabilistic framework, Algorithm 1 describes our generic exploration algorithm.

```

1: while true do
2:    $x \sim p(X | \omega_I)$ 
3:    $y \sim p(Y | x \omega_{SM})$ 
4:    $m = Proj_M((x, y))$ 
5:    $s = exec(m)$ 
6:    $update(\omega_{SM}, (m, s))$ 
7:    $update(\omega_I, (x, y, m, s))$ 
8: end while

```

Algorithm 1: Generic exploration algorithm

Line 2, the agent draws a point x in the choice space X according to the current state of its interest model ω_I , through the probability distribution $p(X | \omega_I)$ encoding the current interest over X . Line 3, the agent draws a point y in the inference space Y according to the distribution $p(Y | x \omega_{SM})$, using Bayesian inference on the joint distribution $p(M S | \omega_{SM})$. Line 4, the agent extracts the motor part m of (x, y) , i.e. x if $X = M$, y if $X = S$. Line 5, the agent produces m and observe $s = exec(m)$, i.e. $s = f(m)$ with possible sensorimotor constraints and noises. Line 6 the agent updates its sensorimotor model according to its new experience (m, s) . Line 7 the agent update its interest model according to the choice and inference (x, y) it made and its new experience (m, s) .

In this framework, we are able to more formally express each algorithm presented in Section 2. The random motor strategy (ACTUATOR-RANDOM) is the simpler case where the choice space is $X = M$ and the interest model of line 2 is set to a uniform distribution over X . Inference in line 3 is here useless because motor extraction (line 4) will return the actual choice x and that there is no need to update the interest model in line 7. The active motor strategy (ACTUATOR-RIAC) differs from the previous one by the interest model of line 2 which favors regions of $X (= M)$ maximizing the competence progress in prediction. This latter is computed at the update step of line 7

using the history of previous competences defined as the opposite difference between the prediction $y \in Y$ and the actual realization $s \in S$ (with $Y = S$). The random goal strategy (SAGG-RANDOM) is the case where the interest model is uniform and the choice space is S , implying that the inference corresponds to a control task to reach $x \in X$ by producing $y \in Y$ (with $X = S$ and therefore $Y = M$). Finally, the active goal strategy (SAGG-RIAC) differs from the previous one by the interest model which favors regions of $X (= S)$ maximizing the competence progress in control. This latter is computed in the same way that for ACTUATOR-RANDOM, except that the opposite difference is here between the chosen goal $x \in X$ and the actual realization $s \in S$ (with $X = S$).

4. Implementation with Gaussian mixture models

In the present paper, we only provide the principles of our implementation of the sensorimotor $p(M S | \omega_{SM})$ and interest $p(X | \omega_I)$ distributions, and leave its detailed description to a further paper. Both the sensorimotor and the interest distributions involves learning of Gaussian mixture models (GMM) using the Expectation-Maximization (EM) algorithm [21]. The values of the parameters we will use in the experiment of the next section appear in parenthesis.

$p(M S | \omega_{SM})$ involves K_{SM} ($=28$) components (i.e. it corresponds to a weighted sum of K_{SM} Gaussian distributions). It is learnt using an online version of EM proposed by [22] where incoming data are considered in lots in an incremental manner. Each update corresponds to line 6 but is executed once each sm_step ($=400$) iterations of Algorithm 1. The ω_{SM} model is thus refined incrementally during the agent life, updating it each time sm_step new (m, s) pairs are collected. Moreover, we adapted this online version of EM to introduce a *learning rate* parameter α (from 0.1 to 0.01 in a logarithmic decreasing manner), allowing to set a relative weight of the new learning data with respect to the old ones.

$p(X | \omega_I)$ is a uniform distribution in the random strategies, whereas it has to reflect an interest measure in the active strategies, which is here related to an empirical measure of the *competence progress*. For this aim, we compute a measure of competence for each 4-tuple (x, y, m, s) collected during the execution of Algorithm 1. We define the competence c of each iteration as $c = e^{-\|(x,y)-(m,s)\|_2}$, i.e. the exponential of the opposite of the distance between the concatenation of the choice and infer points (x, y) and the actual realization (m, s) . Thus, each episode is associated with a tuple (t, x, c) , where t is the (normalized) time index of the iteration. We then consider the competence progress as a correlation between time and competence (the higher the correlation, the higher the competence progress). We learn the joint distribution of this data $p(T X C | \omega_{TXC})$ (where T and C are the mono-dimensional variables defining the time and competence domains, with values in \mathbb{R}^+) using a classical version of EM on a GMM of K_I ($=12$) components using the last $sm_step * im_step$ tuples (t, x, c) , on the time window corresponding to the last im_step ($=12$) updates of the sensorimotor model. After convergence of the EM algorithm, we bias the result by setting the a priori distributions of the model ω_{TXC} (i.e. the weight of each Gaussian) to the resulting value of the covariance between t and c (normalized to sum up to 1 and considering only the positive correlations). Finally, the interest model $p(X | \omega_I)$ corresponds to the Bayesian inference $p(X | [T = t^+] \omega_{TXC})$, where t^+ is

the time index of the future update of the sensorimotor model (e.g. if the t values of the learning set are $\{1, \dots, n\}$, then $t^+ = n + 1$). This allows line 2 of Algorithm 1 to sample values in regions of X which maximize the expected competence progress at the next update of the sensorimotor model.

5. Results

We ran computer experiments using the articulatory synthesizer of the DIVA model described in [23], which thus corresponds to the f unknown function defined in the introduction. This synthesizer is based on the Maeda’s model [24], using 13 articulatory parameters: 10 from a principal component analysis (PCA) performed on sagittal contours of images of the vocal tract of a human speaker, plus glottal pressure, voicing and pitch. It is then able to compute the formants of the signal (among other auditory and somato-sensory features). In the present study, we only use the 7 first parameters of the PCA and the two first formants, approximately normalized in the range $[-1, 1]$. We ran the implementation of the algorithm described in the previous sections with different choice spaces and interest distributions corresponding to the four strategies ACTUATOR-RANDOM, ACTUATOR-RIAC, SAGG-RANDOM and SAGG-RIAC. We evaluate the efficiency of the obtained sensorimotor models to achieve a control task, i.e. to reach a test set of goals uniformly distributed in the reachable auditory space.

Figure 1A shows the performance results of the four exploration strategies on a control task during the life time of learning agents. We observe that the strategies with S as the choice space (random and active goals) are significantly more efficient than those with M (random and active motor), i.e. both convergence speed (say around 100 updates) and generalization at the end of the simulation (500 updates) are better. Moreover, generalization is relatively more efficient for the active than for the random strategies. These results are similar (though less significant) to those obtained in previous experiments [7] in other sensorimotor spaces (e.g. a arm reaching points on a plan), and we refer to the corresponding paper for a thorough analysis of these results.

Figure 1B shows the state of the sensorimotor model ω_{SM} projected in the sensory space S at the end of one simulation for each of the four exploration strategies. We observe that the position of the Gaussians are relatively disorganized when M is the choice space, whereas some structure appears when it is S . Self-organization seems to spontaneously appear in the choice space where points are sampled from the interest model (either uniformly or actively). When this latter is S , this provides a sensorimotor model allowing to better control the vocal tract to reach auditory goals. Another observation is that the auditory space is covered more uniformly in the active than in the random goal exploration strategy. The reason is that the random strategies more often choose goals outside the reachable space, thus favoring reaching at the borders of the sensory space S . To summarize, using S as the choice space is more efficient than using M because self-organization in S is adequate to achieve a control task, and the active goal strategy is more efficient than the random goal one because it allows to focus on the reachable part of S (and perhaps to set goals of increasing difficulties).

6. Conclusions

We have shown in this paper that two principles of developmental robotics (exploration in the sensory space and active learning based on an empirical measure of the competence progress) can




Figure 1: *A: Performance comparison of the four exploration strategies. X axis: number of update of the sensorimotor model. Y axis: Mean error distance on a control task where an agent has to reach 30 test points uniformly distributed in the reachable area of S . For each evaluation point $s_g \in S$, the agent infers 10 motor commands in M from the distribution $p(M | s_g \omega_{SM})$, where ω_{SM} is the state of the sensorimotor model at the corresponding time step (number of update on the X axis). The error of an agent at a time step is the mean distance between the sensory points actually reached by the 10 motor commands and the evaluation point s_g . Each curve plots the mean and standard deviation of the error for 10 independent simulations with different random seeds, for each of the four exploration strategies described in the previous sections. B: State of the sensorimotor model at the end of the simulations for the four exploration strategies. Auditory parameters are the two first formants computed by the articulatory synthesizer. Yellow (or grey) area is the auditory area reached by the agent at the end of the simulation. Ellipses represents the Gaussians of the sensorimotor GMM $p(M | \omega_{SM})$ projected on S .*

improve the learning of the complex sensorimotor mappings involved in speech acquisition. We proposed an integrated probabilistic framework able to express these principles in a compact and unified manner, as well as an original implementation of the underlying algorithm using GMMs and where the competence progress is measured as a statistical correlation between time and competence. Finally, we showed that this modeling is able to match performance comparison results obtained in previous works, and have interesting properties in terms of self-organization.

Further works should focus on an online adaptation of the choice space, e.g. by considering only the dimensions of the sensorimotor space in which the competence progress is maximal, on the integration of maturational constraints and social guidance, as well as on the analysis of emerging developmental sequences confronted with infant experimental data (see [16] for preliminary results on this latter perspective using the same principles of developmental robotics).

7. Acknowledgements

This work was partially financed by ERC Starting Grant EXPLORERS 240 007.

8. References

- [1] F. H. Guenther, M. Hampson, and D. Johnson, "A theoretical investigation of reference frames for the planning of speech movements," *Psychological Review*, vol. 105, no. 4, pp. 611–633, Oct. 1998, PMID: 9830375.
- [2] F. H. Guenther, "Cortical interactions underlying the production of speech sounds," *Journal of Communication Disorders*, vol. 39, no. 5, pp. 350–365, Sep. 2006.
- [3] I. Howard and P. Messum, "Modeling the development of pronunciation in infant speech acquisition," *Motor Control*, vol. 15(1), pp. 85–117, 2011.
- [4] E. Deci and R. M. Ryan, *Intrinsic Motivation and self-determination in human behavior*. New York: Plenum Press, 1985.
- [5] R. M. Ryan and E. L. Deci, "Intrinsic and extrinsic motivations: Classic definitions and new directions," *Contemporary educational psychology*, vol. 25, no. 1, pp. 54–67, 2000.
- [6] M. Rolf, J. Steil, and M. Gienger, "Goal babbling permits direct learning of inverse kinematics," *IEEE Trans. Autonomous Mental Development*, vol. 2, no. 3, pp. 216–229, 09/2010 2010.
- [7] A. Baranes and P.-Y. Oudeyer, "Active learning of inverse models with intrinsically motivated goal exploration in robots," *Robotics and Autonomous Systems*, 2012.
- [8] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active learning with statistical models," *arXiv preprint cs/9603104*, 1996.
- [9] P.-Y. Oudeyer, F. Kaplan, and V. Hafner, "Intrinsic motivation systems for autonomous mental development," *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 2, pp. 265–286, 2007.
- [10] D. E. Berlyne, "A theory of human curiosity," *British Journal of Psychology*, vol. 45, pp. 180–191, 1954.
- [11] J. Schmidhuber, "A possibility for implementing curiosity and boredom in model-building neural controllers," in *Proc. SAB'91*, J. A. Meyer and S. W. Wilson, Eds., 1991, pp. 222–227.
- [12] A. Barto, S. Singh, and N. Chentaz, "Intrinsically motivated learning of hierarchical collections of skills," in *Proc. 3rd Int. Conf. Dvp. Learn.*, San Diego, CA, 2004, pp. 112–119.
- [13] J. Schmidhuber, "Formal theory of creativity, fun, and intrinsic motivation (1990-2010)," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 230–247, 2010.
- [14] F. Kaplan and P.-Y. Oudeyer, "The progress-drive hypothesis: an interpretation of early imitation," in *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions*, K. Dautenhahn and C. Nehaniv, Eds. Cambridge University Press, 2005.
- [15] P.-Y. Oudeyer and F. Kaplan, "Discovering communication," *Connection Science*, vol. 18, no. 2, pp. 189–206, 06 2006.
- [16] C. Moulin-Frier and P.-Y. Oudeyer, "Curiosity-driven phonetic learning," in *International Conference on Development and Learning, Epirob, San Diego, USA*, 2012.
- [17] M. Csikszentmihalyi, *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins, 1997. [Online]. Available: <http://books.google.fr/books?id=aci.Ea4c6woC>
- [18] F. Kaplan and P.-Y. Oudeyer, "In search of the neural circuits of intrinsic motivation," *Frontiers in neuroscience*, vol. 1, no. 1, p. 225, 2007.
- [19] E. T. Jaynes, *Probability Theory: The Logic of Science*, G. L. Bretthorst, Ed. Cambridge University Press, June 2003.
- [20] O. Lebeltel, P. Bessiere, J. Diard, and E. Mazer, "Bayesian robot programming," *Autonomous Robots*, vol. 16, p. 4979, 2004, see <http://www.springerlink.com/content/lh76585657n21244/> See <http://emotion.inrialpes.fr/bibemotion/2004/LBDM04/>.
- [21] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, Jan. 1977, ArticleType: research-article / Full publication date: 1977 / Copyright © 1977 Royal Statistical Society.
- [22] S. Calinon, *Robot Programming by Demonstration*. CRC, 2009.
- [23] F. H. Guenther, S. S. Ghosh, and J. A. Tourville, "Neural modeling and imaging of the cortical interactions underlying syllable production," *Brain and language*, vol. 96, no. 3, pp. 280–301, 2006.
- [24] S. Maeda, "Compensatory articulation during speech: Evidence from the analysis and synthesis of vocal tract shapes using an articulatory model," *Speech production and speech modelling*, pp. 131–149, 1989.

T