

Efficient Bootstrapping of Vocalization Skills Using Active Goal Babbling

Anja Kristina Philippsen¹, René Felix Reinhart², Britta Wrede¹

¹Cognitive Interaction Technology Center (CITEC), Bielefeld University, Germany

²Research Institute for Cognition and Robotics (CoR-Lab), Bielefeld University, Germany

{anja.philippsen, freinhart}@uni-bielefeld.de, bwrede@techfak.uni-bielefeld.de

Abstract

We use goal babbling, a recent approach to bootstrapping inverse models, for vowel acquisition with an articulatory speech synthesizer. In contrast to motor babbling, goal babbling organizes exploration in a low-dimensional goal space. While such a goal space is naturally given in many motor learning tasks, the difficulty in modeling speech production lies within the complexity of acoustic features. Formants can serve as low-dimensional features, but richer acoustic features are too high-dimensional to allow for an efficient goal-directed exploration. We propose to generate a low-dimensional goal space from high-dimensional features by applying dimension reduction. In this way the goal space adapts to a set of speech sounds, which models the influence from ambient speech on the speech acquisition process. Instead of pre-defining targets in this goal space, we estimate a target distribution with a Gaussian Mixture Model. We demonstrate that goal babbling can be successfully applied in this goal space in order to learn a parametric model of vowel production. By augmenting goal-directed exploration with an active selection of targets, we achieve a significant speed-up in learning.

Index Terms: speech motor learning, goal-directed exploration, acoustic-to-articulatory inversion, dimension reduction, active learning

1. Introduction

In order to learn to speak, infants have to explore the capabilities of their vocal tract by executing articulatory configurations and observing the auditory outcome. This babbling produces articulatory-acoustic examples which can be used to gradually build up an inverse model of speech production containing information on which articulatory commands are necessary in order to achieve a specific auditory goal.

While a random exploration of motor configurations is not feasible in high-dimensional motor spaces, reinforcement learning methods can be applied to guide the exploration process by rewarding “good” examples. This has been implemented in the speech domain for the purpose of modeling spontaneous vocalization [1, 2, 3, 4] or imitation learning [5, 6, 7, 8, 9, 10].

These approaches, however, do not directly yield an inverse model that maps from auditory targets to motor commands. In the context of learning sensorimotor coordination, goal babbling was introduced for efficiently bootstrapping an inverse model [11, 12, 13]. The key feature of goal babbling is that exploration is not organized in motor space, but in the space of desired outcomes (goals). This has several advantages: it is more efficient, as the task space (in contrast to the high-dimensional motor control space) is usually low-dimensional, and goal babbling is capable of directly bootstrapping a parametric model. This accounts for the fact that sounds develop not separately, but

in conjunction with each other. Finally, it is developmentally plausible, as infants perform goal-directed movements even at very early stages of their development [14, 15].

Recently, Moulin-Frier et al. used the concept of goal babbling for modeling vocal development in the context of curiosity-driven learning [16, 17]. They achieved promising results by effectively simplifying the problem: [16] limited learning to the acquisition of vowels, [17] developed an intrinsically motivated robot learner that gradually improves from unarticulated to articulated speech sounds. Although it successfully models the emergence of syllables by defining two sub-goals, it does not learn to produce a set of distinguishable utterances due to the limited acoustic feature representation. [18] uses goal babbling to bootstrap a model to control F0 contour in speech sounds.

Up to date, bootstrapping a set of complex distinguishable syllables is still an unsolved problem. There are two important reasons for that. One reason is the difficulty to find an appropriate goal space. While for goal babbling inverse kinematics a low-dimensional continuous goal space is naturally given by the space of 3D coordinates [11], speech can be represented in various features spaces, most of which have no advantage over the motor space because they are similarly high-dimensional. The space of the first and second formant is low-dimensional and effectively captures the differences between vowel sounds. It is, however, not adaptable to different inputs and captures consonant characteristics only to a limited degree. A second reason is that for speech production, time is an important additional dimension as it distinguishes e.g. via voice-onset time between voiced and unvoiced consonants. Including this dimension into the goal space makes the problem even more high-dimensional and goal-directed exploration less efficient. For these reasons we argue that there is the need for a low-dimensional space that can be used for goal-directed exploration in the context of speech production.

In this paper we propose to first learn such a goal space in an unsupervised way based on ambient speech sounds. Inspired by a recent approach to organizing motor skills along meaningful dimensions, the Parameterized Skill Memory [19], we embed speech sounds into a low-dimensional space by applying dimension reduction techniques. In this way, we reduce the dimensionality of the goal space drastically such that an entire speech sequence is mapped onto a single point, e.g. in 2D. This solves the above mentioned problems by providing a low-dimensional goal space which captures the variance in the ambient speech sounds.

In this goal space, we apply goal-directed exploration along linear paths as implemented by Rolf [12] to bootstrap a set of vowel sounds. We use a learned model of ambient target distribution instead of pre-defined targets. To accelerate the bootstrapping process, we replace the random target selection in [12]




Figure 1: *Initialization phase: the goal space is generated from ambient language sounds. Production and perception loop: after training, the inverse model $g(x)$ estimates an articulatory parameter configuration q for a selected target x in the goal space such that the forward model $f(q)$ embeds the produced acoustics close to the desired target in goal space.*

with a competence-based selection inspired by [17].

In this first study we demonstrate this concept for bootstrapping vowel production skills. The influence of ambient language on vocal development is modeled similarly to [4], but as their system uses random motor babbling, they vary only 2 or 6 vocal tract parameters. Using goal babbling in this study, such a reduction is not necessary.

2. Embedding speech sounds

Research on speech acquisition in children typically suggests that infants are not influenced by ambient speech sounds during the first 10 months [20], but in fact they are exposed to the ambient language even before birth [21]. In the same way that an infant’s early movements are goal-directed from the beginning [14, 15], their acoustic targets could also arise from the sounds they perceive from the environment. Evidence from developmental research supports the view that speech perception influences speech learning in young infants, as deaf children fail to produce well-formed syllables within the first 10 months [22].

Based on this idea, we assume that a set of speech examples of the ambient language is available to the system before it starts to explore. This data set contains only acoustic examples, as no knowledge about articulation can be assumed for ambient language. For speech production, we used the Maeda speech synthesizer [23] as implemented in the DIVA model [8]. The set of ambient speech sounds was created with the DIVA model by requesting articulatory trajectories with specific formant frequencies. In this way we obtained articulatory configurations for [a], [e], [i], and (using the default vocal tract posture) the neutral “schwa” [ə]. Articulatory postures are represented by 10 parameters (parameter values $\in [-1, 1]$) describing the vocal tract configuration (the 3 source parameters were omitted and fixed to values such that phonation occurs) and extended in time such that the generated speech signals are 600 ms long. We generated the acoustic consequences of 100 variations of each of the four vowel sounds by applying normally distributed noise (variance 0.05) to the articulatory parameters.

As acoustic features we use cochleograms as calculated in the Auditory Toolbox by Lyon’s Passive Ear Model, a biologically inspired model which models the hair cell response in the




Figure 2: *Frequency responses of Lyon’s Passive Ear Model for an audio sample rate of 11025 Hz.*

human inner ear (cochlea) [24, 25]. These features change continuously over time which might be beneficial for dimension reduction. However, in principle any acoustic feature representation can be used. Default parameters from [25] were used for the calculation. The filters are automatically generated; for an audio sample rate of 11025 Hz the number of filters is 74. Figure 2 shows the frequency responses of each fifth filter.

From this set of ambient speech sounds our system learns a goal space in an unsupervised way via dimension reduction (see initialization phase in Figure 1). The acoustic sequences (downsampled to 120 time steps \times 74 feature dimensions) were first transformed into 8880-dimensional vectors, then a simple dimension reduction technique, namely Principal Component Analysis (PCA), was applied to randomly selected 90% of the ambient speech samples. The resulting 2D representation (see goal space in Figure 1) captures approx. 73% of the variance in the ambient speech data.

The mapping from the high-dimensional acoustic features to the goal space, which PCA provides, characterizes the goal space, as it can be used to map other acoustic perceptions into this space as well. Points in the goal space correspond to targets that the system should learn to achieve. To obtain a representation of the distribution of these targets, a Gaussian Mixture Model (GMM) [26] was trained on the embedded ambient speech data. Mean and covariance of the four mixture components are displayed as ellipses in Figure 1.

3. Goal-directed exploration

After the goal space and the target distribution are generated, learning to speak can be defined as learning the inverse mapping from the goal space to the articulatory parameters. The aim is to close the production and perception loop (see Figure 1). The forward model f includes sound production via the DIVA model, acoustic feature extraction, and the mapping into the goal space. This stays fixed during exploration. In contrast to that, the inverse model g is adapted after each exploration step. After training it should be capable of imitating acoustic sounds, represented as a position in goal space, by estimating an articulatory posture that leads to an acoustically similar result. In other words, the loop is closed if for a desired target position \mathbf{x}^* the estimated articulatory posture $\mathbf{q} = g(\mathbf{x}^*)$ produces an outcome in goal space $\mathbf{x} = f(\mathbf{q})$ that is close to the desired target position. If all target positions in goal space can be successfully reached, the system has learned how to speak with respect to the ambient language.

Goal babbling implements a way of bootstrapping the inverse model $g(\mathbf{x}, \theta)$ by continuously trying to reach targets and updating the inverse model parameters θ . Section 3.1 explains our implementation of goal babbling, which is a slightly modified version of [11, 12]. Section 3.2 explains how we integrated intrinsic motivation for an active selection of targets.

3.1. Goal babbling

We adopted the goal babbling method from Rolf that explores along linear paths towards targets [11, 12]. The inverse model is implemented as a Radial Basis Function (RBF) network [27] with an underlying clustering algorithm that can be updated with weighted samples in an online fashion similar to [11, 12]. While in [11, 12] targets were set manually, we obtained a statistic representation of targets in the form of a GMM from the ambient speech sounds set as described in Section 2. We assume the utterances in the acoustic space to be Gaussian distributed within the generated goal space, thus, the target distribution is defined as

$$P(\mathbf{x}^*) = \sum_{n=1}^N \pi_n \mathcal{N}(\mathbf{x}^* | \mu_n, \Sigma_n), \quad (1)$$

where π_n are the prior probabilities for the $N = 4$ target clusters and μ_n and Σ_n are the parameters of the Gaussian distribution obtained from GMM training.

The overall bootstrapping process can be divided into two major steps: exploration in goal space and adaptation of the inverse model. In the beginning, the inverse model g is initialized with $(\mathbf{x}^{home}, \mathbf{q}^{home})$, where $\mathbf{q}^{home} \in \mathbb{R}^{10}$ is the default vocal tract posture and $\mathbf{x}^{home} \in \mathbb{R}^2$ is the corresponding position in goal space determined by the forward model as $\mathbf{x}^{home} = f(\mathbf{q}^{home})$.

3.1.1. Exploration in goal space

Targets are drawn from the target distribution $\mathbf{x}_k^* \sim P(\mathbf{x}^*)$ in each iteration k . To gradually teach the learner to achieve a target \mathbf{x}_k^* , sub-targets $\mathbf{x}_{k,l}^*$, $l = [0 \dots L]$ are defined by dividing the path between $\mathbf{x}_{k,0}^*$ and $\mathbf{x}_{k,L}^*$ into equally spaced exploration steps, such that $\mathbf{x}_{k,L}^* = \mathbf{x}_k^*$. In original goal babbling [11, 12], the linear movement towards a target \mathbf{x}_k^* starts from the target of the previous iteration, i.e. $\mathbf{x}_{k,0}^* = \mathbf{x}_{k-1,L}^*$. In this study, the movement towards a new target starts from the position that the learner actually managed to reach, i.e. $\mathbf{x}_{k,0}^* = f(g(\mathbf{x}_{k-1,L}^*))$.

In this way we make sure that new exploratory movements always start at a point that the learner is already able to reach. This also alleviates the problem that articulatory postures drift away from the home posture and renders it unnecessary to perform homeward movements as described in [11, 12].

After target selection, the inverse model estimate is consulted and exploratory noise is added to the estimated vocal tract posture in order to obtain an articulatory estimation:

$$\mathbf{q}_{k,l} = g(\mathbf{x}_{k,l}^*, \theta_{k,l-1}) + E(\mathbf{x}_{k,l}^*), \quad (2)$$

where $E(\mathbf{x}_{k,l}^*)$ is a structured continuous variation (see [12] for details). By applying the forward model, the actually reached outcome is identified: $\mathbf{x}_{k,l} = f(\mathbf{q}_{k,l})$.

3.1.2. Adaptation of inverse model

After each exploration step, the inverse model parameters are updated with the new training pair $(\mathbf{x}_{k,l}, \mathbf{q}_{k,l})$, weighted according to $w_{k,l} = w_{k,l}^{dir} \cdot w_{k,l}^{eff} \cdot w_{k,l}^{tar}$.

$w_{k,l}^{dir}$ measures the direction and is defined as in [11, 12].

$w_{k,l}^{eff}$ measures the effectiveness of the movement and is slightly adjusted here to be 0 if a small change in posture leads to a large change in position (due to the complexity introduced by the learned mapping from acoustic features to goals):

$$w_{k,l}^{eff} = \begin{cases} \min\left(\frac{\|\mathbf{x}_{k,l} - \mathbf{x}_{k,l-1}\|}{\|\mathbf{q}_{k,l} - \mathbf{q}_{k,l-1}\|}, 1\right) & \text{if } \frac{\|\mathbf{x}_{k,l} - \mathbf{x}_{k,l-1}\|}{\|\mathbf{q}_{k,l} - \mathbf{q}_{k,l-1}\|} \leq 2 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Additionally, we introduce $w_{k,l}^{tar}$ in this study, which expresses how well the target was approximated:

$$w_{k,l}^{tar} = \begin{cases} \exp(-2\|\mathbf{x}_{k,l}^* - \mathbf{x}_{k,l}\|) & \text{if } \|\mathbf{x}_{k,l}^* - \mathbf{x}_{k,l}\| \leq 0.5 \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

By collecting a new training pair in each exploration step the inverse model gradually learns to estimate articulatory configurations in order to reach targets in the goal space with a low reproduction error. We define the reproduction error of the updated inverse model with parameters $\theta_{k,l}$ for a desired target \mathbf{x}^* as:

$$e_{k,l}(\mathbf{x}^*) = \|\mathbf{x}^* - f(g(\mathbf{x}^*, \theta_{k,l}))\|. \quad (5)$$

3.2. Active target selection

With the above described version of goal babbling, the system selects the next target \mathbf{x}_k^* randomly according to the distribution $P(\mathbf{x}^*)$ (see Eq. (1)). In fact it could happen that some of the targets can already be effectively reached. The learner then would lose valuable time by further exploring these targets. To accelerate learning, we implement a simple variant of intrinsic motivation: the next target is selected actively by integrating information about the current learning progress. Such active goal-directed exploration was found to be superior to random exploration schemes [16, 17, 28, 29, 30].

We measure the current learning progress at the end of a movement towards target \mathbf{x}_k^* by calculating the reproduction errors of the N GMM cluster centers μ_n according to Eq. (5) as $e_{k,L}(\mu_n)$. Before selecting a new target \mathbf{x}_{k+1}^* , the priors π_n of the GMM in Eq. (1) are adjusted according to:




Figure 3: Average reproduction error in goal space plotted for 3000 speaking attempts with random (light gray) or active (dark gray) selection of targets. Average and standard deviation over 10 independent trials are displayed.

$$\pi_n = \frac{e_{k,L}(\mu_n)}{\sum_j e_{k,L}(\mu_j)}. \quad (6)$$

Prior probabilities π_n in Eq. (6) take higher values for target clusters that are poorly approximated. Updating $P(\mathbf{x}^*)$ with these new priors, the learner concentrates mainly on targets that it cannot produce yet, and only occasionally repeats already mastered sounds.

4. Bootstrapping a set of speech sounds

We applied goal babbling with random or active selection of targets in the goal space generated in Section 2. $K = 120$ movements with $L = 25$ steps each, resulting in 3000 exploratory speaking attempts in total, were performed in each run. The level of exploratory noise was set to $\sigma^2 = 0.1$, $\sigma_\Delta^2 = 0.1$ (cf. [12]). The learning rate for the adaptation of the inverse model was 0.9.

Figure 3 shows the average reproduction errors after each exploration step which are assessed by averaging the reproduction errors of the GMM cluster centers (cf. (5)):

$$e_{k,l} = \frac{1}{N} \sum_{n=1}^N e_{k,l}(\mu_n) \quad (7)$$

Mean and standard deviation were computed over 10 runs of the experiments. It can be observed that the error decreases faster and reaches a lower level after 3000 exploration steps if targets are selected based on the competence-based measure. Additionally, the lower standard deviation in the case of active target selection suggests more stable results.

To assess the performance and generalization capability of the trained inverse model, we evaluated the production and perception loop with one randomly selected inverse model trained via active goal babbling for 2000 exploratory steps. The inverse model estimated articulatory parameters for 41×41 equally spaced target positions $\mathbf{x}^* = [x_1^*, x_2^*]$ in goal space, where $x_1^*, x_2^* \in [-1, 1]$. These articulatory configurations were then executed and mapped back into the goal space by the forward

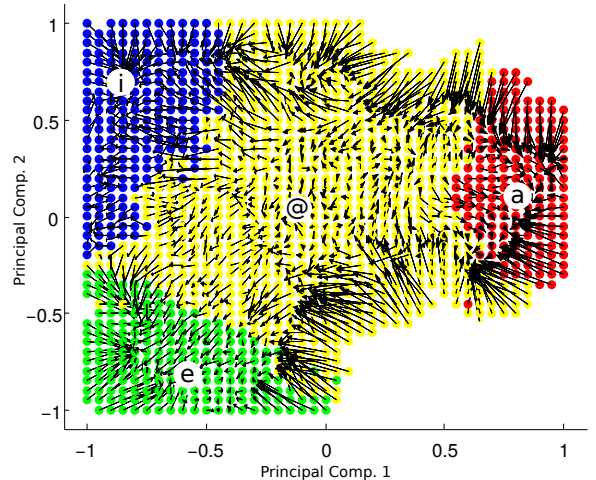


Figure 4: Reproduction error of the inverse model after training. Arrows point from desired targets to reached targets.

model (cf. Figure 1). In Figure 4 the deviations of the reproduction from the original target are depicted in goal space with arrows pointing from the desired target positions to the actually reached positions. The colors of the points at the position of the requested targets indicate how the reproduction is perceived in goal space, i.e. to which of the four target clusters the newly embedded point is assigned. Small reproduction errors occur for targets near the cluster centers (cf. goal space in Figure 1). The further away from the ambient speech distribution a target is requested, the higher is the deviation in goal space. For the purpose of clarity, targets \mathbf{x}^* that are reproduced with an error $\|\mathbf{x}^* - f(g(\mathbf{x}^*))\| > 0.3$ are omitted in this figure.

5. Conclusion & outlook

In this study, we demonstrated that it is possible to apply goal babbling for the learning of vowel sounds in a goal space that was generated from high-dimensional acoustic features. An active selection of targets based on competences accelerates learning such that the inverse model can be learned in less than 2000 speaking attempts. More advanced measures of competence that are selective towards specific regions of the goal space could facilitate even quicker bootstrapping.

A major advantage of the proposed method is that, in contrast to previous studies, it does not require low-dimensional acoustic features, where often only formants are an option, but can easily be used with a variety of different speech features. Furthermore, the goal space adapts to the ambient speech, which could help to investigate the influence from the ambient language on speech acquisition in future studies.

As next steps, we want to test the method with other acoustic features, embedding methods or vocal tract models. We also plan to extend it towards bootstrapping of syllables by representing articulatory trajectories.

6. Acknowledgements

This research was supported by the Cluster of Excellence Cognitive Interaction Technology 'CITEC' (EXC 277) at Bielefeld University, which is funded by the German Science Foundation (DFG), and has been conducted in the framework of the European Project CODEFROR (FP7-PIRSES-2013-612555).

7. References

- [1] A. S. Warlaumont, "A spiking neural network model of canonical babbling development," in *IEEE Second Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 2012, pp. 1–6.
- [2] —, "Salience-based reinforcement of a spiking neural network leads to increased syllable production," in *IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 2013, pp. 1–7.
- [3] A. S. Warlaumont, G. Westermann, E. H. Buder, and D. K. Oller, "Prespeech motor learning in a neural network using reinforcement," *Neural Networks*, vol. 38, pp. 64–75, 2013.
- [4] G. Westermann and E. R. Miranda, "A new model of sensorimotor coupling in the development of speech," *Brain and language*, vol. 89, no. 2, pp. 393–400, 2004.
- [5] I. S. Howard and P. Messum, "Modeling the development of pronunciation in infant speech acquisition," *Motor Control*, vol. 15, no. 1, pp. 85–117, 2011.
- [6] F. H. Guenther, "Speech sound acquisition, coarticulation, and rate effects in a neural network model of speech production," *Psychological review*, vol. 102, no. 3, p. 594, 1995.
- [7] —, "Cortical interactions underlying the production of speech sounds," *Journal of communication disorders*, vol. 39, no. 5, pp. 350–365, 2006.
- [8] J. A. Tourville and F. H. Guenther, "The DIVA model: A neural theory of speech acquisition and production," *Language and Cognitive Processes*, vol. 26, no. 7, pp. 952–981, 2011, source code available at: <http://www.bu.edu/speechlab/software/diva-source-code/>.
- [9] B. J. Kröger, J. Kannampuzha, and C. Neuschaefer-Rube, "Towards a neurocomputational model of speech production and perception," *Speech Communication*, vol. 51, no. 9, pp. 793–809, 2009.
- [10] M. Murakami, B. Kröger, 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 *IEEE Fifth Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 2015.
- [11] M. Rolf, J. J. Steil, and M. Gienger, "Goal babbling permits direct learning of inverse kinematics," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 216–229, 2010.
- [12] —, "Online goal babbling for rapid bootstrapping of inverse models in high dimensions," in *IEEE International Conference on Development and Learning (ICDL)*. IEEE, 2011.
- [13] M. Rolf, "Goal babbling with unknown ranges: A direction-sampling approach," in *IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 2013.
- [14] A. Van der Meer, F. Van der Weel, D. N. Lee *et al.*, "The functional significance of arm movements in neonates," *SCIENCE-NEW YORK THEN WASHINGTON-*, pp. 693–693, 1995.
- [15] C. Von Hofsten, "An action perspective on motor development," *Trends in cognitive sciences*, vol. 8, no. 6, pp. 266–272, 2004.
- [16] C. Moulin-Frier and P.-Y. Oudeyer, "Curiosity-driven phonetic learning," in *IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 2012, pp. 1–8.
- [17] C. Moulin-Frier, S. M. Nguyen, and P.-Y. Oudeyer, "Self-organization of early vocal development in infants and machines: the role of intrinsic motivation," *Frontiers in psychology*, vol. 4, 2013.
- [18] H. Liu and Y. Xu, "Learning model-based F0 production through goal-directed babbling," in *9th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2014, pp. 284–288.
- [19] R. F. Reinhart and J. J. Steil, "Efficient policy search with a parameterized skill memory," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*. IEEE, 2014, pp. 1400–1407.
- [20] P. K. Kuhl, "Early language acquisition: cracking the speech code," *Nature reviews neuroscience*, vol. 5, no. 11, pp. 831–843, 2004.
- [21] A. J. DeCasper and M. J. Spence, "Prenatal maternal speech influences newborns' perception of speech sounds," *Infant behavior and Development*, vol. 9, no. 2, pp. 133–150, 1986.
- [22] D. K. Oller and R. E. Eilers, "The role of audition in infant babbling," *Child development*, pp. 441–449, 1988.
- [23] 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.
- [24] R. Lyon, "A computational model of filtering, detection, and compression in the cochlea," in *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP'82.*, vol. 7. IEEE, 1982, pp. 1282–1285.
- [25] M. Slaney, "Auditory toolbox," *Interval Research Corporation, Tech. Rep.*, vol. 10, p. 1998, 1998, source code available at: <https://engineering.purdue.edu/%7Emalcolm/interval/1998-010/>.
- [26] S. Calinon, F. Guenther, and A. Billard, "On learning, representing, and generalizing a task in a humanoid robot," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 37, no. 2, pp. 286–298, 2007, source code available at: <http://www.calinon.ch/sourcecodes.php> (GMM-GMR).
- [27] J. A. Freeman and D. Saad, "Online learning in radial basis function networks," *Neural Computation*, vol. 9, no. 7, pp. 1601–1622, 1997.
- [28] 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.
- [29] S. M. Nguyen, "A curious robot learner for interactive goal-babbling: Strategically choosing what, how, when and from whom to learn." Ph.D. dissertation, Université Sciences et Technologies – Bordeaux I, 2013.
- [30] S. M. Nguyen and P.-Y. Oudeyer, "Socially guided intrinsic motivation for robot learning of motor skills," *Autonomous Robots*, vol. 36, no. 3, pp. 273–294, 2014.