

Self-supervised Body Image Acquisition Using a Deep Neural Network for Sensorimotor Prediction

Alban Laflaquière
AI Lab
SoftBank Robotics Europe
Paris, France
alafaquiere@softbankrobotics.com

Verena V. Hafner
Adaptive Systems Group
Humboldt-Universität zu Berlin
Berlin, Germany
hafner@informatik.hu-berlin.de

Abstract—This work investigates how a naive agent can acquire its own body image in a self-supervised way, based on the predictability of its sensorimotor experience. Our working hypothesis is that, due to its temporal stability, an agent’s body produces more consistent sensory experiences than the environment, which exhibits a greater variability. Given its motor experience, an agent can thus reliably predict what appearance its body should have. This intrinsic predictability can be used to automatically isolate the body image from the rest of the environment. We propose a two-branches deconvolutional neural network to predict the visual sensory state associated with an input motor state, as well as the prediction error associated with this input. We train the network on a dataset of first-person images collected with a simulated Pepper robot, and show how the network outputs can be used to automatically isolate its visible arm from the rest of the environment. Finally, the quality of the body image produced by the network is evaluated.

I. INTRODUCTION

Going through different developmental phases, human children continuously acquire control over their own bodies [1]. Eventually, they can distinguish between self and others and predict the consequences of their own actions. This pre-reflective identification is often called the minimal self, and constitutes of two major components: a sense of agency, and a sense of body ownership [2]. These notions have recently become relevant also in developmental robotics. Firstly, computational models of the self can provide interesting insights into the processes of self-development in humans, and secondly an adaptive self-model can be crucial for intuitive and adaptive interaction in robotics.

Different models have been proposed in the past few years to address one or more facets of this problem. In a recent paper, Lang, Schillaci and Hafner [3] presented a study on the sense of agency and object permanence in which a robot predicted its own arm movements using a convolutional neural network. Hoffman et al. [4] investigated the formation of a body representation in a humanoid robot equipped with artificial skin during experiments on touch. Another study investigates the effects of self-touch in human infants related to their own motor actions, and suggests robotic models [5]. An information-theoretic approach to the formation of body maps in robots has been proposed in [6], in which the structure of the body map resulted from information distances between sensory data during a particular behaviour of the robot. An

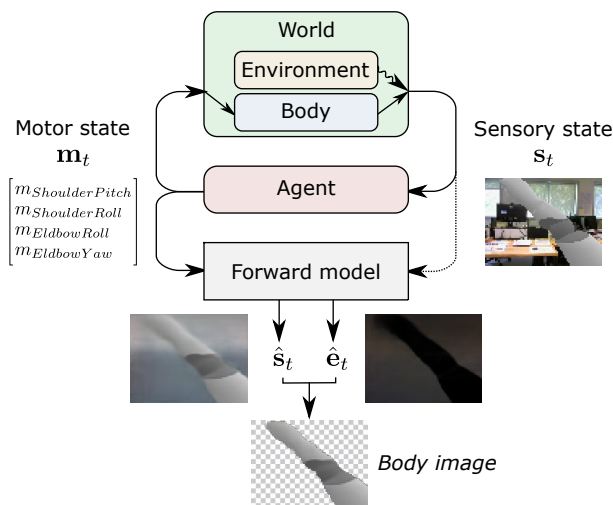


Fig. 1. Interaction between the agent and its environment, and forward model learned by the agent. Based on the predictability of the sensory inputs, the agent can isolate its body image from the rest of the environment.

approach for body and non-body discrimination in robotics has been proposed by Yoshikawa et al. [7], where they propose a method to approximate posture sensations by Gaussian distributions. A predictive coding approach to generate visuo-proprioceptive patterns has been implemented by Hwang et al. [8] on a simulated iCub robot. In this study, the robot was trained to imitate gestures of another robot or its mirror image displayed on a screen. Hinz et al. [9] present a study investigating prediction errors in tactile-visual data in both humans and robots. They use a rubber-hand illusion setup typically used to study properties of body ownership.

Gallagher [10] argues that body image and body schema are separate concepts. He suggests that the body image is a conscious representation of the body, whereas the body schema is a prepersonal structural mapping. In artificial systems, a clear distinction between these two concepts is more difficult to make, and their respective definition often varies. In this paper, we use the term “body image” in a very literal sense, as the “appearance of the body in the visual flow”, and we argue that body image acquisition depends on the predictive capabilities of the agent. Extending the studies of Lang et al.




Fig. 2. The neural network architecture used to predict an output image \hat{s}_t and a prediction error \hat{e}_t , given an input motor state \mathbf{m}_t . Two (de)convolutional branches are dedicated to produce these two respective outputs. In the diagram, FC stands for a fully-connected layer, R stands for a reshaping operation, D stands for a deconvolutional layer, and C stands for a convolutional layer.

[3], we suggest a mechanism for an artificial agent to acquire its body image from scratch through predictive processes in self-supervised way.

The paper is structured as follows: section II presents the methodology and neural network used in the approach. Section III presents the experimental setup. Section IV presents the qualitative and quantitative results. Finally, we discuss the results and conclusions in section V.

II. METHODOLOGY

A. Approach

The objective of this work is to study how an agent could learn its own body image (appearance) in a self-supervised way. Our working hypothesis is that when predicting sensori-motor experiences, the body appears as a primary source of predictability, as it is always present and does not change, or only at a very slow pace. As a consequence, when reaching motor states, the “appearance” of the body in the sensory flow is consistent throughout the exploration of the environment. Compared to the rest of the sensory flow induced by the environment, which the agent does not directly control, this *body image* should thus be significantly more predictable.

We can formalize this intuition by considering the mapping between the agent’s motor state and sensory state. We denote $\mathbf{m}_t = [m_{1,t}, \dots, m_{N_m,t}]$ the proprioceptive motor state, or posture, in which the agent is at time t , where each $m_{n,t}$ corresponds to the static configuration of a joint. Similarly, we denote $\mathbf{s}_t = [s_{1,t}, \dots, s_{N_s,t}]$ the sensory state that the agent receives from its exteroceptive sensors at time t , where each $s_{n,t}$ is an individual sensory component (see Fig. 1).

We assume in this work that the sensors are such that they provide an instantaneous reading of the state of the world, and that, for any motor state \mathbf{m}_t , the sensory state \mathbf{s}_t can be divided into two subsets: $\mathcal{B}_{\mathbf{m}_t}$ and $\overline{\mathcal{B}_{\mathbf{m}_t}}$. The subset $\mathcal{B}_{\mathbf{m}_t}$ gathers all components $s_{i,t}$ associated with the body, while its complement gathers the ones associated with the

environment. Said otherwise, we assume that the elementary sensory excitations $s_{i,t}$ are not due to a mixed contribution of the body and the environment. This is typically the case with a camera as, for a given body posture, a subset of pixels (and their respective channels) corresponds to body parts in the field of view, while other pixels correspond to the environment.

Without a priori knowledge about the state of the environment, the mapping $\mathbf{m}_t \rightarrow \mathbf{s}_t$ is not deterministic. Putting aside potential sensorimotor noise, the uncertainty about the state of the environment limits the ability to predict \mathbf{s}_t . However, if the body appearance stays temporally consistent, we hypothesize that the subset $\mathcal{B}_{\mathbf{m}_t}$ should exhibit a significantly lower variability in time than $\overline{\mathcal{B}_{\mathbf{m}_t}}$, i.e.:

$$\begin{aligned} s_{i,t} \in \mathcal{B}_{\mathbf{m}_t} &\Rightarrow \text{Var}(s_{i,t}|\mathbf{m}_t) \rightarrow 0, \\ s_{i,t} \in \overline{\mathcal{B}_{\mathbf{m}_t}} &\Rightarrow \text{Var}(s_{i,t}|\mathbf{m}_t) \rightarrow \epsilon \gg 0, \end{aligned} \quad (1)$$

where $s_{i,t}$ is treated as a random variable over time. In an environment which exhibits a sufficient amount of variability, this intrinsic difference can be used, in a data-driven way, to distinguish the sensory components which belong to the body image and to the environment (symmetrizing the implication arrows of Eq.(1)). Note that the typical RGB excitations of a pixel are here considered as separate components $s_{i,t}$.

More formally, the agent can learn a forward model mapping from \mathbf{m}_t to a prediction $\hat{\mathbf{s}}_t$ of \mathbf{s}_t , as well as to a prediction \hat{e}_t of the error $|\hat{\mathbf{s}}_t - \mathbf{s}_t|$ (see Fig. 1). According to Eq. 1, the elementary prediction errors $\hat{e}_{i,t}$ should be significantly lower if $s_{i,t} \in \mathcal{B}_t$ than if $s_{i,t} \in \overline{\mathcal{B}_t}$, allowing to distinguish the two subsets based on the accuracy of the forward model.

B. Neural network architecture

We use a deep neural network to learn both the forward mapping $\mathbf{m}_t \rightarrow \hat{\mathbf{s}}_t$ and the error prediction mapping $\mathbf{m}_t \rightarrow \hat{e}_t$. As displayed in Fig. 2, the network first passes the input motor state \mathbf{m}_t through two fully-connected layers. The last of these layers is then reshaped to form a low resolution




Fig. 3. The qibullet simulator is used to generate first person view images of the Pepper robot’s right arm moving in front of a green background. This background is then replaced by office-type images collected by a real Pepper robot. The final training dataset consist in these composite images and their corresponding motor states (posture).

image-like representation that can be spatially processed for deconvolution. This intermediary representation is then fed to two distinct branches which respectively predict the output image \hat{s}_t and the prediction error \hat{e}_t . Each branch is composed of three successive deconvolutional layers and three convolutional layers. The three deconvolution layers upscale the low resolution representation to the final size of the image. Following the good practice proposed in [11], the deconvolution operation consists in an upscaling operation, followed by a convolution. The next three convolutional layers perform typical convolution operations [12], and progressively reduce the depth of the input from 32 channels to 3, corresponding to the RGB channels of the pixels. In each branch, all layers use the SeLu activation function [13] —which performs an internal normalization of the neurons’ activation—, except the last convolutional layer which uses the ReLu activation function [14], in order to guarantee the positivity of each $\hat{s}_{i,t}$ and $\hat{e}_{i,t}$. Throughout the network, all convolutions are done using kernels of size 3×3 , with a stride of 1.

A different loss is associated with each branch of the network. The first branch (top one in Fig. 2) outputs a predicted image \hat{s}_t , and its associated reconstruction loss is:

$$L_{rec} = \frac{1}{NN_s} \sum_{k=1}^N \sum_{i=1}^{N_s} |\hat{s}_{i,k} - s_{i,k}|, \quad (2)$$

where $s_{i,k}$ is a sensory component for input k , and $|\cdot|$ is the absolute operator. It corresponds to the mean value of the L^1 norm between \hat{s}_t and s_t , normalized by the number of sensory components (three times the number of pixels).

The second branch (bottom one in Fig. 2) outputs the predicted absolute prediction error \hat{e}_t between the predicted image \hat{s}_t and the ground-truth image s_t , and its associated loss is:

$$L_{err} = \frac{1}{NN_s} \sum_{k=1}^N \sum_{i=1}^{N_s} |\hat{e}_{i,k} - |\hat{s}_{i,k} - s_{i,k}||. \quad (3)$$

It corresponds to the mean value of the L^1 norm between \hat{e}_t and the actual prediction error $|\hat{s}_t - s_t|$, normalized by the number of sensory components. Finally, the total loss is a weighted composition of the these two losses:

$$L = L_{rec} + \alpha \cdot L_{err}, \quad (4)$$

where α denotes a scalar relative weight.

This total loss is minimized using the ADAM optimizer [15], for 5×10^4 iterations, and with a learning rate linearly decreasing from 10^{-3} to 10^{-5} during training. At each iteration a random mini-batch of $N = 100$ samples is fed to the network to compute the (stochastic) gradient. Finally, α is set to increase linearly from 0 to 1 during the first 2.5×10^4 iterations. It helps stabilizing the convergence by first focusing the optimization on the image prediction branch, so that its output can be used as a reliable target for the second branch.

III. EXPERIMENTS

A. Sensorimotor data generation

In order to test our approach, we created a dataset of images in which a simulated Pepper robot [16] observes its own right arm in different environments. In order to quickly generate large datasets without directly using the physical robot and to easily assess a ground-truth body image mask, we created a synthetic experimental dataset by composing images.

First, the realistic Pepper simulator qibullet [17], was used to quickly generate 8000 right arm configurations \mathbf{m}_t . They were randomly generated by uniformly drawing the 4 following joint orientations $m_{ShouldPitch} \in [-1, 1]$, $m_{ShouldRoll} \in [-1, 0]$, $m_{ElbRoll} \in [0, 1]$, $m_{ElbYaw} \in [-1, 1]$ (radians). The motor exploration thus resembles a typical motor babbling [18]. During this exploration, the simulated Pepper is placed in front of a green wall, with the head oriented towards the right arm (downward pitch of 0.253, rightward yaw of 0.7 (radians)). For each \mathbf{m}_t , the (240×320) pixels image captured by the robot’s camera is recorded. As can be seen in Fig. 3, these images correspond to a first person view of the arm in front of a uniform green surface. This clear bimodal structure allows use to easily replace the green surface with any desired background. We filled it with random images from a previous dataset collected while a real Pepper robot moved in an office-like environment. The robot’s body is not visible in this previous dataset, which means that the arm images generated using qibullet can be embedded in them without ambiguity. The whole creation of the dataset is illustrated in Fig. 3, where one can see the arm configurations and the background images used to compose the experimental dataset.




Fig. 4. Evaluation of the network outputs on a random subsampling of the training dataset. First row: ground-truth composite image from the dataset. Second row: image \hat{s}_t predicted by the network, given the corresponding motor state \mathbf{m}_t . Third row: prediction error \hat{e}_t predicted by the network, given the motor state. Fourth row: body mask automatically derived from \hat{e}_t . Fifth row: predicted body image after applying the mask to \hat{s}_t . Note that each plot is displayed as a RGB image, including the mask, as the 3 channels are considered independently in the data processing.

Note that in this compositional approach, the ambient lighting of the scene (background) does not affect the appearance of the robot’s arm. Due to its probabilistic nature, we however expect our approach to be robust to small changes in lighting conditions, such as the ones observed in the office-type background dataset. This however needs to be confirmed by future experiments.

Before being fed to the network, both the motor states and sensory states (images) are pre-processed. The motor states are normalized such that each component spans the $[-1, 1]$ interval, and the sensory states are normalized such that all components lie in $[0, 1]$ (instead of $[0, 255]$ originally). Finally, the input images are downsampled to a (60×80) resolution, in order to limit the network size and computation time (although no theoretical limitation prevents the approach from scaling to greater resolutions).

B. Body image extraction

Our working hypothesis is that for each motor state \mathbf{m}_t it should exist a subset of sensory components for which predictability is significantly higher than for other ones. Those components should correspond to the consistent part of the robot’s visual experience: its body image. We thus propose to automatically isolate this body image by looking at the predicted prediction error $\hat{e}_t | \mathbf{m}_t$. If its distribution is indeed bimodal, we propose to set a threshold \mathcal{T} between these two modes and to consider any component with a predicted error $\hat{e}_{i,t} | \mathbf{m}_t$ inferior to \mathcal{T} as part of the body. This way we can create a binary mask to isolate the body image from the rest of the input image.

The code used to generate the data, and train and test the network is available at: <https://github.com/alafquiere/learn-masked-body-image>.

IV. RESULTS

After training, we qualitatively and quantitatively evaluate the mappings learned by the neural network. Figure 4 shows

the predicted image \hat{s}_t and the predicted prediction error \hat{e}_t for 11 random samples from the training set.

Firstly, we can see that the predicted images contain a meaningful approximation of the appearance of the arm. Note that the network also predicts the absence of the arm when the motor state moves it out of the field of view (see the last column of Fig. 4). The arm appears in front of a background made of two relatively uniform horizontal stripes. The brighter upper stripe seems to statistically correspond to walls and windows which tend to be white in the background dataset. The darker lower stripe seems to statistically correspond to the floor which tends to be darker in the dataset. Apart from this statistical distinction, the background of the predicted images does not contain any specific pattern from the original ground-truth images. In order to minimize its prediction error, the network thus learned to output the expectation of each unpredictable background sensory component.

Secondly, the predicted error displays a similar structure. The area of the image corresponding to the arm is associated with a very low predicted prediction error (darker), while the background is associated with greater errors (lighter). It thus seems that our working hypothesis was correct: based on the predictability of the visual input sub-components, it is possible to differentiate the body image from the rest of the environment in a self-supervised way.

Figure 5 shows a normalized histogram of all the predicted prediction errors $\hat{e}_{i,k}$ produced by the network for 100 random motor states from the training dataset. As expected, the distribution appears to be bimodal. We fit it with a 2-component Gaussian Mixture Model and define a threshold at the intersection of the two components, i.e. $\mathcal{T} = 0.056$. This threshold is used to automatically distinguish the sensory components $\hat{s}_{i,t}$ belonging to the body image \mathcal{B}_t ($\hat{e}_{i,t} \leq \mathcal{T}$) from the ones belonging to the environment $\bar{\mathcal{B}}_t$ ($\hat{e}_{i,t} > \mathcal{T}$). It allows us to compute a mask on the sensory components to isolate the body image, as displayed in Fig. 4. Note that in this process, the R, G, and B channels of each pixel are




Fig. 5. Normalized histogram of all the predicted prediction errors $\hat{e}_{t,k}$ over 100 random inputs from the dataset. The distribution appears to be bimodal, with a lower mode corresponding to highly predictable sensory components, and a higher mode corresponding to unpredictable ones. The threshold $\mathcal{T} = 0.056$ is set at the intersection of the two modes.

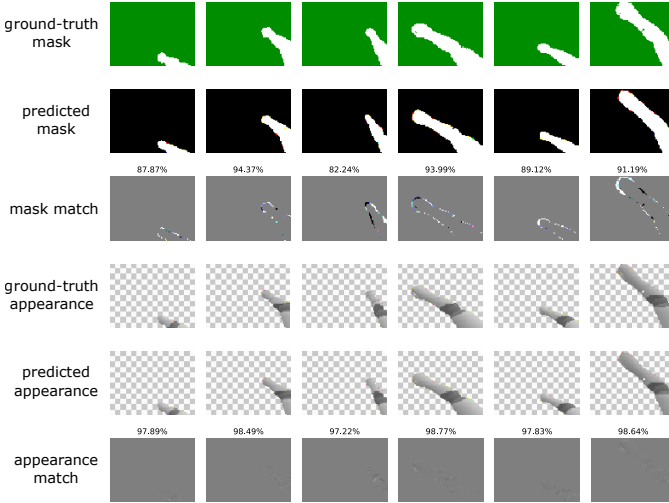


Fig. 6. Visualization of the match between the ground-truth body mask and the estimated body mask (first three rows), and of the match between the ground-truth appearance of the arm and the predicted appearance of the arm after application of the mask (last three rows).

treated independently, assuming minimal a priori knowledge about the structure of the sensory state. This potentially allows a mismatch between the different color components of the body image. Finally, this mask of depth 3 is applied to the predicted image \hat{s}_t by making transparent the components not in the mask. The resulting body image, cleaned from the poorly predictable background, is displayed in Fig. 4 as well. We can see that the appearance of the arm in the masked predicted images is very close to the one in the ground-truth images. The biggest disparities are located at the arm tip (hand), which appears to be the most difficult part to reconstruct. This can be explained by the fact that the appearance and localization of the hand in the image changes the most rapidly as a function of the joint configuration. On the contrary, the upper arm, closest to the shoulder and to the camera, is the most consistently reconstructed part, as its appearance changes the least as a function of the joint configuration.

We introduce two measures to quantify the quality of the

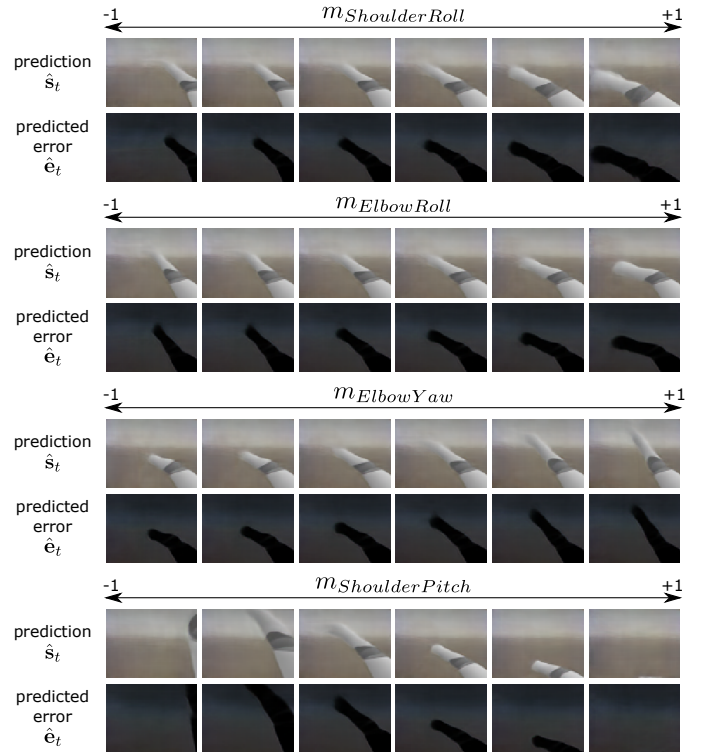


Fig. 7. Sampling of the motor space. Starting from the reference motor state $[0, 0, 0, 0]$, each motor dimension is individually crossed from -1 to 1, and the corresponding network outputs are displayed.

learned body image. First, the *mask match* corresponds to the Intersection over Union (IoU) of the mask derived from the network’s output and the ground-truth mask (non-green pixels in the image before composition).

$$\text{mask match} = \#(\hat{\mathcal{B}}_{\mathbf{m}_t} \cap \mathcal{B}_{\mathbf{m}_t}) / \#(\hat{\mathcal{B}}_{\mathbf{m}_t} \cup \mathcal{B}_{\mathbf{m}_t}). \quad (5)$$

Second, the *appearance match* is defined as:

$$\text{appearance match} = 1 - \frac{1}{N_{\mathcal{B}}} \sum_{i=1}^{N_{\mathcal{B}}} \mathbf{1}_{\mathcal{B}} |\hat{s}_{i,t} - s_{i,t}|, \quad (6)$$

where $\mathbf{1}_{\mathcal{B}}$ is equal to 1 if $s_{i,t} \in \hat{\mathcal{B}}_{\mathbf{m}_t} \cap \mathcal{B}_{\mathbf{m}_t}$ and 0 otherwise, and $N_{\mathcal{B}}$ is the number of components in this mask intersection. Note that in each measure, the three RGB channels of each pixel are considered independently.

Figure 6 displays 6 visualizations of mask and appearance matches for random inputs of a testing dataset generated the same way as the training dataset. We can see that most errors happen at the edge of the arm, where pixels values are the most likely to quickly change as a function of the motor joints, due to the environment being unpredictable. We can also see that the errors in the arm appearance are insignificant, and barely visible as a consequence (see last row). Over the whole testing dataset of 2000 samples, the mask match is equal to $88.8 \pm 13.3\%$, and the appearance match is equal to $98.3 \pm 0.7\%$. The match between the ground-truth body image and the one produced by the network is thus very good.

Finally, the network can act as a (deterministic) conditional generative model. We can sample the input motor space and let the network generate an output image and output prediction error. Figure 7 shows the result of such a sampling when varying each motor component independently, starting from the reference arm configuration $[0, 0, 0, 0]$. We can see that the arm appearance and the associated low prediction error mask stay consistent throughout the motor space. The neural network thus seems to generalize and to accurately predict the arm appearance over the whole motor space.

V. CONCLUSION

We presented an algorithm and experiments for autonomous body image acquisition. This work extends the studies of Lang et al. [3] by a mechanism to distinguish sensory components that belong to the body image from those that belong to the environment, from scratch, in a self-supervised way. It relies on the hypothesized intrinsic difference in variability of those two kinds of components, that a robot can capture when trying to predict its sensory state given an input motor state.

Like in the previous one, only movements of a single arm have been considered in this study. However this work could potentially be extended to more complex movements, including other limbs or the head itself. As shown by Schmerling et al. [19], the resulting change in the visual image caused by the head motion could indeed be included in the predictive learning process. The only theoretical limitation of such an extension is the number of samples required to correctly estimate the conditional statistics of the sensory experience as the dimension of the motor space increases.

We have shown that the prediction of the visual information associated with a motor state is crucial for the formation of the body image. From a predictive coding perspective, the body image would in this way correspond to the component of the sensory experience which is reliably and quickly predictable, on a developmental scale, given the motor experience that the agent generates itself. The importance of motor information in this formative process suggests a strong connection between the sense of agency, body ownership, body schema, and body image. Another possible extension of this work would be to emphasize this motor aspect even more by introducing a dynamical system formulation of the problem, like has been proposed in [3]. A network could for instance be optimized to predict the future sensory state \mathbf{s}_{t+1} , given the current sensory state \mathbf{s}_t and a motor command \mathbf{m}_t , instead of a static posture like in the current formulation. It would also be interesting to investigate if it is possible for an extended version of the network to simultaneously learn to infer its current body posture \mathbf{m}_t , using an auto-encoder paradigm, and to isolate its body image, based on the latent code learned by this auto-encoder.

Finally, it is important to note again that we used in this work the term “body image” in a literal sense to refer to the *appearance* that the body has in a visual flow. The same term is often used to refer to more complex notions covering different psychological concepts, and whose definition can vary. Many

open questions remain regarding the complete modeling of the notion of body image, and its coupling with other notions like body schema, body ownership, or the sense of agency.

REFERENCES

- [1] P. Rochat, “What is it like to be a newborn?” in *The Oxford Handbook of the Self*, S. Gallagher, Ed. Oxford University Press, 2011, pp. 57–79.
- [2] S. Gallagher, “Philosophical conceptions of the self: implications for cognitive science,” *Trends in cognitive sciences*, vol. 4, no. 1, pp. 14–21, 2000.
- [3] C. Lang, G. Schillaci, and V. V. Hafner, “A deep convolutional neural network model for sense of agency and object permanence in robots,” in *8th Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*. IEEE, 2018, pp. 260–265.
- [4] M. Hoffmann, Z. Straka, I. Farkaš, M. Vavrečka, and G. Metta, “Robotic homunculus: Learning of artificial skin representation in a humanoid robot motivated by primary somatosensory cortex,” *IEEE Trans. on Cognitive and Developmental Systems*, vol. 10, no. 2, pp. 163–176, 2018.
- [5] M. Hoffmann, L. K. Chinn, E. Somogyi, T. Heed, J. Fagard, J. J. Lockman, and J. K. O’Regan, “Development of reaching to the body in early infancy: From experiments to robotic models,” in *IEEE Int. Conf. on Development and Learning and Epigenetic Robotics*. IEEE, 2017.
- [6] F. Kaplan and V. V. Hafner, “Information-theoretic framework for unsupervised activity classification,” *Advanced Robotics*, vol. 20, no. 10, pp. 1087–1103, 2006. [Online]. Available: <https://doi.org/10.1163/156855306778522514>
- [7] Y. Yoshikawa, Y. Tsuji, K. Hosoda, and M. Asada, “Is it my body? body extraction from uninterpreted sensory data based on the invariance of multiple sensory attributes,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 3, 2004, pp. 2325–2330.
- [8] J. Hwang, J. Kim, A. Ahmadi, M. Choi, and J. Tani, “Predictive coding-based deep dynamic neural network for visuomotor learning,” in *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics, ICDL-EpiRob 2017, Lisbon, Portugal, September 18-21, 2017*, 2017, pp. 132–139. [Online]. Available: <https://doi.org/10.1109/DEVLRN.2017.8329798>
- [9] N.-A. Hinz, P. Lanillos, H. Mueller, and G. Cheng, “Drifting perceptual patterns suggest prediction errors fusion rather than hypothesis selection: replicating the rubber-hand illusion on a robot,” in *IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)*. IEEE, 2018. [Online]. Available: <https://arxiv.org/pdf/1806.06809.pdf>
- [10] S. Gallagher, “Body image and body schema: A conceptual clarification,” *Journal of Mind and Behavior*, vol. 7, pp. 541–554, 11 1985.
- [11] A. Odema, V. Dumoulin, and C. Olah, “Deconvolution and checkerboard artifacts,” *Distill*, 2016. [Online]. Available: <http://distill.pub/2016/deconv-checkerboard>
- [12] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel, “Handwritten digit recognition with a back-propagation network,” in *Advances in neural information processing systems*, 1990, pp. 396–404.
- [13] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, “Self-normalizing neural networks,” in *Advances in neural information processing systems*, 2017, pp. 971–980.
- [14] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 807–814.
- [15] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [16] SoftBank Robotics. (2014) Pepper robot. [Online]. Available: [https://en.wikipedia.org/wiki/Pepper_\(robot\)](https://en.wikipedia.org/wiki/Pepper_(robot))
- [17] B. Maxime and C. Maxime. (2019) qibullet: A bullet-based python simulation for softbank robotics’ robots. [Online]. Available: <https://github.com/ProtolabSBRE/qibullet>
- [18] A. N. Meltzoff and M. K. Moore, “Explaining facial imitation: A theoretical model,” *Infant and child development*, vol. 6, no. 3-4, pp. 179–192, 1997.
- [19] M. Schmerling, G. Schillaci, and V. V. Hafner, “Goal-directed learning of hand-eye coordination in a humanoid robot,” in *2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, Aug 2015, pp. 168–175.