

# How to Sense the World: Leveraging Hierarchy in Multimodal Perception for Robust Reinforcement Learning Agents

Miguel Vasco  
INESC-ID & Instituto Superior Técnico,  
Universidade de Lisboa  
Lisbon, Portugal

Francisco S. Melo  
INESC-ID & Instituto Superior Técnico,  
Universidade de Lisboa  
Lisbon, Portugal

Hang Yin  
KTH Royal Institute of Technology  
Stockholm, Sweden

Ana Paiva  
INESC-ID & Instituto Superior Técnico,  
Universidade de Lisboa  
Lisbon, Portugal

## ABSTRACT

This work addresses the problem of *sensing the world*: how to learn a multimodal representation of a reinforcement learning agent’s environment that allows the execution of tasks under incomplete perceptual conditions. To address such problem, we argue for *hierarchy* in the design of representation models and contribute with a novel multimodal representation model, MUSE. The proposed model learns a hierarchy of representations: low-level modality-specific representations, encoded from raw observation data, and a high-level multimodal representation, encoding joint-modality information to allow robust state estimation. We employ MUSE as the perceptual model of deep reinforcement learning agents provided with multimodal observations in Atari games. We perform a comparative study over different designs of reinforcement learning agents, showing that MUSE allows agents to perform tasks under incomplete perceptual experience with minimal performance loss. Finally, we also evaluate the generative performance of MUSE in literature-standard multimodal scenarios with higher number and more complex modalities, showing that it outperforms state-of-the-art multimodal variational autoencoders in single and cross-modality generation.

## KEYWORDS

Multimodal Representation Learning; Reinforcement Learning; Un-supervised Learning

### ACM Reference Format:

Miguel Vasco, Hang Yin, Francisco S. Melo, and Ana Paiva. 2022. *How to Sense the World: Leveraging Hierarchy in Multimodal Perception for Robust Reinforcement Learning Agents*. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Perceiving the world is a core skill for any autonomous agent. To act efficiently, the agent must be able to collect, process and interpret the perceptual information provided by its environment.

*Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.*




**Figure 1: Acting under incomplete perceptual conditions.**  
Adapted from George Morland’s *Blind Man’s Buff* (1788).

This information can be of considerable complexity, such as the high-dimensional pixel information provided by the agent’s camera or the sounds collected by the agent’s microphones. To reason efficiently over such information, agents are often endowed with mechanisms to encode low-dimensional internal representations of their complex perceptions. Such representations enable more efficient learning of how to act [6, 33] and facilitate the adaptation to distinct (albeit similar) domains [8, 11].

In this work we pivot on the design of representations for sensory information. In particular, we address the problem of *sensing the world*: how can we process an agent’s sensory information in such a way that it is robust to changes in the perceptual conditions of the environment (such as the removal of input modalities), mimicking animal behavior [22] and the human experience (Fig. 1).

Recently, deep reinforcement learning (RL) agents have shown remarkable performance in complex control tasks when provided with high-dimensional observations, such as in Atari Games [1, 18, 20]. Early approaches considered training a controller, instantiated as a neural network, directly over sensory information (e.g. pixel data), thus learning an *implicit* representation of the observations of the agent within the structure of the controller. However, in a multimodal setting, such agents struggle to act when provided with incomplete observations: the missing input information flowing through the network degrades the quality of the agent’s output policy. To provide some robustness to changing sensory information, other agents learn *explicit* representation models of their sensory information, often employing variational autoencoder (VAE) models, before learning how to perform a task in their environment. These explicit representations allow agents to perceive their environment and act in similar environments, such as in the DARLA



**Figure 2: Explicit representation models for RL agents in multimodal scenarios: (a) VAE model with a flat latent representation  $z_\pi$  and fused observations  $x_{[1:M]}$  [8, 9, 11]; (b) Multimodal VAE (MVAE) with a flat latent representation  $z_\pi$  [25]; (c) the proposed MUSE model, with hierarchical levels: modality-specific ( $z_{1:M}$ ) and a multimodal ( $z_\pi$ ) representations. We highlight the encoders (orange, dashed) and decoders (blue, full) networks of the models.**

framework [11], or to act from observations generated by the representation itself, such as in the World Models framework [8].

In this work, we focus on the design of explicit representation models that are robust to changes in the perceptual conditions of the environment (i.e. with unavailable modalities at execution time) for RL agents in multimodal scenarios. In such conditions, traditional approaches employing VAE models with a fixed *fusion* of the agent’s sensory information struggle to encode a suitable representation: the missing information flowing within the network degrades the quality of the encoded state representation and, subsequently, of the agent’s policy [4]. To address such problem, multimodal variational autoencoders (MVAE) have been recently employed to learn representations in a multimodal setting [25]. By considering the *independent* processing of each modality, these models are able to overcome the problems of fusion solutions, as no missing information is propagated through the network [25]. However, by design, the models encode information from all modalities into a single *flat* representation space of finite capacity, often leading the agent to neglect information from lower-dimensional modalities [24], hindering the estimation of the state of the environment when high-dimensional observations are unavailable.

To provide robust state estimation to the agent, regardless of the nature of the provided sensory observations, we argue for designing perceptual models with *hierarchical* representation levels: raw observation data is processed and encoded in low-level, modality-specific representations, which are subsequently merged in a high-level, multimodal representation. By accounting for the intrinsic complexity of each input modality at a low-level of the model, we are able to encode a multimodal state representation that better accounts for all modalities and is robust to missing sensory information. We instantiate such design in a novel framework, contributing the *Multimodal Unsupervised Sensing* (MUSE) model.

We evaluate MUSE as an explicit representation model for RL agents in a recently-proposed deep reinforcement learning scenario, where an agent is provided with multimodal observations of Atari games [25]. We perform a comparative study against other architectures of reinforcement learning agents and show that only RL agents employing MUSE are able to perform tasks with missing observations at test time, incurring in a minimal performance loss. Finally, we assess the performance of MUSE in scenarios with increasing complexity and number of modalities and show, quantitatively and qualitatively, that it outperforms state-of-the-art MVAEs.

In summary, the main contributions of this work are threefold: (i) MUSE, a novel representation model for multimodal sensory information of RL agents that considers hierarchical representation levels; (ii) a comparative study on different architectures of RL agents showing that, with MUSE, RL agents are able to perform tasks with missing sensory information, with minimal performance loss; (iii) an evaluation of MUSE in literature-standard multimodal scenarios of increasing complexity and number of modalities showing that it outperforms state-of-the-art multimodal variational autoencoders.

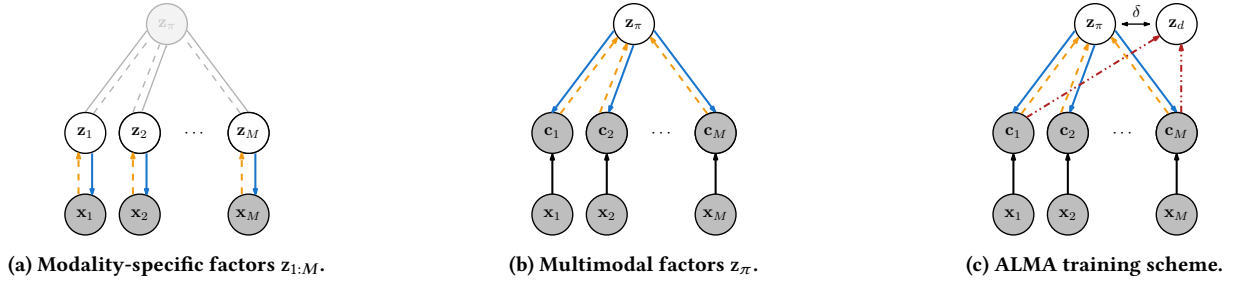
## 2 BACKGROUND

Reinforcement learning (RL) is a computational framework for decision-making through trial-and-error interaction between an agent and its environment [26]. RL problems can be formalized using Markov decision processes (MDP) that describe sequential decision problems under uncertainty. An MDP can be instantiated as a tuple  $\mathcal{M} = (\mathcal{X}, \mathcal{A}, P, r, \gamma)$ , where  $\mathcal{X}$  and  $\mathcal{A}$  are the (known) state and action spaces, respectively. When the agent takes an action  $a \in \mathcal{A}$  while in state  $x \in \mathcal{X}$ , the world transitions to state  $y \in \mathcal{X}$  with probability  $P(y | x, a)$  and the agent receives an immediate reward  $r(x, a)$ . Finally, the discount factor  $\gamma \in [0, 1)$  defines the relative importance of present and future rewards for the agent.

The reward function  $r$ , typically unknown to the agent, encodes the goal of the agent in its environment, whose dynamics are described by  $P$ , often unknown to the agent as well. Through a trial-and-error approach, the agent aims at learning a policy  $\pi : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ , where  $\pi(a | x)$  is the probability of performing action  $a \in \mathcal{A}$  in state  $x \in \mathcal{X}$ , that maximizes the expected reward collected by the agent. Such a policy can be found from the optimal  $Q$ -function, defined for every state-action pair  $(x, a) \in \mathcal{X} \times \mathcal{A}$  as

$$Q^*(x, a) = r(x, a) + \gamma \sum_{y \in \mathcal{X}} P(y | x, a) \max_{a' \in \mathcal{A}} Q^*(y, a'). \quad (1)$$

To compute the optimal  $Q$ -function, multiple methods can be employed [26], such as  $Q$ -learning [30]. Recently, RL has been applied to scenarios where the agent collects high-dimensional observations from its environment, leading to new methods and extensions of classical reinforcement learning methods [1]. For discrete action spaces the Deep  $Q$  Network (DQN) was proposed, a variant of the  $Q$ -learning algorithm, that employs a deep neural network to approximate the optimal  $Q$ -function [20]. For continuous action spaces the Deep Deterministic Policy Gradient (DDPG) algorithm



**Figure 3: Learning hierarchical representations in the MUSE framework: (a) the modality-specific factors  $z_{1:M}$ , employing the loss in Eq. (2); (b) the multimodal factors  $z_\pi$ , encoded from the low-dimensional representations  $c_{1:M}$ , employing the loss in Eq. (3) (c) the Average Latent Multimodal Approximation (ALMA) training scheme, where we approximate the complete multimodal latent variable  $z_\pi$  with partial multimodal latent variables  $z_d$ , encoded from observations with missing modalities.**

allows agents to perform complex control tasks [18], for example in robotic manipulation [29]. However, in a multimodal setting, both algorithms learn an implicit representation of the fused sensory information provided to the agent, which may hinder their performance in scenarios with missing modality information.

### 3 MUSE

In this work, we consider that a reinforcement learning agent is provided with information regarding the state of environment through  $M$  different channels,  $\{\mathbf{x}_1, \dots, \mathbf{x}_M\}$ , where  $\mathbf{x}_m$  is the information provided by an input “modality”. Each modality may correspond to a different type of information (e.g., image, sound).

To design RL agents for such scenario, we can employ several distinct architectures: one can employ a neural-network controller that learns a policy  $\pi$  directly from the fixed fusion of all observations  $\mathbf{x}_{[1:M]}$ , thus learning an implicit representation of sensory information, as in DQN or DDPG [18, 20]. One can also employ the same fusion mechanism to learn an explicit representation of the agent’s observations using a VAE model (Fig. 2a) and, subsequently, learn a policy over the flat latent representation  $z_\pi$  [8, 9, 11]. Another architecture uses a multimodal VAE (Fig. 2b), able to consider the independent processing of each modality to encode  $z_\pi$  [24, 27, 31, 32]. However, two issues arise from such design choices:

- The fusion solution employed by both implicit and explicit representation models is not robust to missing observations: the missing input propagates through the network reducing its performance, thus leading to an inaccurate state representation and to an incorrect policy;
- The flat representation solution with independent processing encodes information from all modalities into a finite-capacity representation, often disregarding the information provided by low-dimensional observations, as shown in Shi et al. [24]. Thus, this option struggles to encode a robust state representation solely from low-dimensional modalities.

To provide robust state representation, regardless of the number and nature of the available modalities, we argue for considering a *hierarchical* architecture that accommodates:

- Low-level *modality-specific* representations, encoding information specific to each modality in (low-level) latent variables whose capacity (dimensionality) is individually defined considering the intrinsic complexity of each modality;

- A high-level *multimodal* representation, merging information from the available low-level representations. By using low-dimensional representations as input, the multimodal representation can better balance the information provided by the distinct modalities and, thus, encode a robust representation regardless of the nature of the available observations.

#### 3.1 Learning Hierarchical Representations

We instantiate the previous hierarchical design in the *Multimodal Unsupervised Sensing* (MUSE) model, depicted in Fig.2c. To train the modality-specific representations  $z_{1:M}$ , we assume that each modality  $\mathbf{x}_m$  is generated by a corresponding latent variable  $z_m$  (Fig. 3a). We can employ a loss function  $\ell_b(\mathbf{x}_{1:M})$ , similar to the single-modality VAE loss [13], to learn a set of independent, single-modality, generative models  $p_\theta(\mathbf{x}_{1:M}) = \prod_{m=1}^M p_\theta(\mathbf{x}_m)$ , i.e.,

$$\ell_b(\mathbf{x}_{1:M}) = \sum_{m=1}^M \left( \mathbb{E}_{q_\phi(z_m|\mathbf{x}_m)} [\lambda_m \log p_\theta(\mathbf{x}_m|z_m)] - \alpha_m \mathbb{KL}(q_\phi(z_m|\mathbf{x}_m) \parallel p(z_m)) \right), \quad (2)$$


where  $\lambda_m$  affects the reconstruction of modality-specific data and  $\alpha_m$  controls the regularization of the corresponding latent variable.

At a high-level, we encode modality-specific information to learn a multimodal representation,  $z_\pi$ . As shown in Fig. 3b, we employ the modality-specific encoder networks to sample an efficient low-level representation of modality data  $\mathbf{c}_{1:M} \sim [q_\phi(z_{1:M}|\mathbf{x}_{1:M})]^1$ . We assume that the representations  $\mathbf{c}_{1:M}$  are generated by a multimodal latent variable  $z_\pi$  and employ a loss function  $\ell_t(\mathbf{x}_{1:M})$ , similar to the standard MVAE loss [31], to learn the multimodal representation,

$$\ell_t(\mathbf{x}_{1:M}) = \mathbb{E}_{q_\phi(\mathbf{c}_{1:M}|\mathbf{x}_{1:M})} \left( -\beta \mathbb{KL}(q_\phi(z_\pi | \mathbf{c}_{1:M}) \parallel p(z_\pi)) + \sum_{m=1}^M \mathbb{E}_{q_\phi(z_\pi|\mathbf{c}_{1:M})} [\gamma_m \log p_\theta(\mathbf{c}_m | z_\pi)] \right), \quad (3)$$

where  $\gamma_m$  affects the reconstruction of the modality-specific representations and  $\beta$  controls the regularization of the multimodal latent variable.

<sup>1</sup>At test time, we compute deterministically the low-level representation  $\mathbf{c}_{1:M} = \mathbb{E}[q_\phi(z_{1:M}|\mathbf{x}_{1:M})]$



**Figure 4: Designing reinforcement learning agents provided with multimodal observations  $\mathbf{x}_{1:M}$ : (a) the Multimodal Controller solution; (b) the VAE + Controller solution; (c) the MVAE + Controller solution; (d) our proposed MUSE + Controller solution. We highlight the representation learning (orange, dashed) and the policy learning (purple, full) components of the agents.**

### 3.2 Multimodal Training Scheme

To encode a multimodal representation that is both agnostic to the nature of the modalities and scalable to large number of modalities, we approximate the joint-modality posterior distribution  $p(\mathbf{z}_\pi | \mathbf{c}_{1:M})$  using a product-of-experts (PoE), where,

$$p(\mathbf{z}_\pi | \mathbf{c}_{1:M}) = p(\mathbf{z}_\pi) \prod_{m=1}^M q(\mathbf{z}_\pi | \mathbf{c}_m). \quad (4)$$

This solution is able to scale to a large number of modalities [31]: assuming that both the prior and posteriors are Gaussian distributions, the product-of-experts distribution is itself a Gaussian distribution with mean  $\mu = (\sum_m \mu_m T_m) (\sum_m T_m)^{-1}$  and covariance  $\Sigma = (\sum_m T_m)^{-1}$ , where  $T_m = \Sigma_m^{-1}$  and  $\Sigma_m$  is the covariance of  $q(\mathbf{z}_\pi | \mathbf{c}_m)$ . However, the original PoE solution is prone to learning overconfident experts, neglecting information from low-dimensional modalities, as shown in [24].

To address this issue, we introduce the *Average Latent Multimodal Approximation* (ALMA) training scheme for PoE. With ALMA, we explicitly enforce the similarity between the multimodal distribution encoded from all modalities with the “partial” distributions encoded from observations with missing modalities, as depicted in Fig. 3c. In particular, during training we encode a latent distribution  $q(\mathbf{z}_\pi | \mathbf{c}_{1:M})$  that considers all modalities and, in addition, we also encode distributions with missing modality information,  $q(\mathbf{z}_d | \mathbf{x}_d)$ , one for every possible combination of modalities. For example, in scenarios with two input modalities ( $\mathbf{x}_1, \mathbf{x}_2$ ) we encode  $D = 2$  partial distributions, corresponding to  $\mathbf{z}_d^1 \sim q(\cdot | \mathbf{c}_1)$  and  $\mathbf{z}_d^2 \sim q(\cdot | \mathbf{c}_2)$ . To encode a multimodal representation robust to missing modalities, we force all the multimodal distributions to be similar by including  $|D|$  additional loss terms, yielding a final loss function

$$\begin{aligned} \ell(\mathbf{x}_{1:M}) &= \ell_b(\mathbf{x}_{1:M}) + \ell_t(\mathbf{x}_{1:M}) \\ &+ \frac{\delta}{|D|} \sum_{d \in D} \mathbb{KL}^*(q_\phi(\mathbf{z}_\pi | \mathbf{c}_{1:M}) \| q_\phi(\mathbf{z}_d | \mathbf{c}_d)), \end{aligned} \quad (5)$$

where the parameter  $\delta$  governs the impact of the approximation loss term and  $\mathbb{KL}^*$  is the symmetrical KL-divergence, following [32]. Employing the loss function of (5), we train both representation

levels concurrently in the same data pass through the model.<sup>2</sup> In Appendix we show how the ALMA term plays a fundamental role in addressing the overconfident expert phenomena of PoE solutions.

### 3.3 Learning To Act

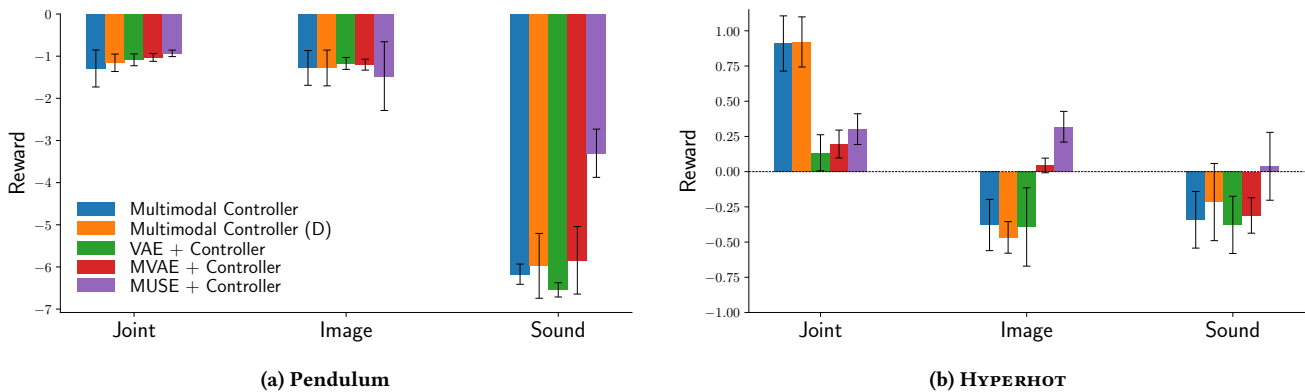
To employ MUSE as a sensory representation model for RL agents, we follow the three-step approach of [25]: initially, we train MUSE on a previously-collected dataset of joint-modality observations  $\mathcal{D}(\mathbf{x}_{1:M}) = \{\mathbf{x}_{1:M}^0, \dots, \mathbf{x}_{1:M}^N\}$ , using the loss function of Eq. (5). After training MUSE, we encode the perceptual observations of the RL agent in the multimodal latent state  $\mathbf{z}_\pi$  to learn a policy  $\pi : \mathcal{Z} \rightarrow \mathcal{A}$  that maps the latent states  $\mathbf{z}_\pi \in \mathcal{Z}$  to actions of the agent  $a \in \mathcal{A}$ . To learn such a policy  $\pi$  over the latent state, one can employ any continuous-state space reinforcement learning algorithm, such as DQN [20] or DDPG [18].

## 4 EVALUATION

We evaluate MUSE addressing the following two questions: **(i)** What is the performance of an RL agent that employs MUSE as a sensory representation model when provided with observations with missing modality information? **(ii)** What is the standalone performance of MUSE as a generative model in scenarios with larger number and more complex modalities?

To address **(i)** we perform on Section 4.1 a comparative study of different architectures of RL agents against a novel agent that employs MUSE as a sensory representation model. We evaluate the agents in the recently proposed multimodal Atari game scenario, where an agent is provided with multimodal observations of the game [25]. The results show that—with MUSE—the RL agents are able to perform tasks with missing modality information at test time, with minimal performance loss. To gain further insight into these results, we also address **(ii)**, evaluating the standalone performance of MUSE against other state-of-the-art representation models across different literature-standard multimodal scenarios. We provide quantitative and qualitative results on Section 4.2 that show that MUSE outperforms other models in single and cross-modality generation. Code available at <https://github.com/miguelsvasco/muse>.

<sup>2</sup>We stop the gradients of the top loss (3) from propagating to the bottom-level computation graph by cloning and detaching the codes  $\mathbf{c}_{1:M}$  sampled from the modality-specific distributions.



**Figure 5: Performance after zero-shot policy transfer in (a) the Pendulum scenario; and (b) the HYPERHOT scenario. At test time, the agent is provided with either image ( $x_I$ ), sound ( $x_S$ ), or joint-modality ( $x_I, x_S$ ) observations (Higher is better). All results averaged over 10 randomly seeded runs. Complete numerical results are presented in Appendix.**

#### 4.1 Acting Under Incomplete Perceptions

We evaluate the performance of an RL agent that employs MUSE as a sensory representation model in a comparative study of different literature-standard agent architectures. In particular, we evaluate how the nature of the state representation (implicit vs explicit), the method of processing sensory information (fusion vs independent) and the type of representation model employed (flat vs hierarchical) allows agents to act under incomplete perceptions. As shown in Fig. 4, we consider different design choices for RL agents:

- *Multimodal Controller* (Fig 4a), where we train a controller directly over the multimodal observations, in an end-to-end fashion (implicit, fusion, flat). This is the case of models such as DQN [20] and DDPG [18];
- *VAE + Controller* (Fig 4b), where we initially train a representation model of the (fused) multimodal observations of the environment using a VAE and afterwards train the controller over the state representation  $z_\pi$  (explicit, fusion, flat). This is the case of frameworks such as World Models [8], DARLA [11] and Dreamer [9];
- *MVAE + Controller* (Fig 4c), where we initially train a representation model of sensory information using the MVAE model [31] and afterwards train the controller over the state representation  $z_\pi$  (explicit, independent, flat). This is the case of the models in Silva et al. [25];
- Our proposed **MUSE + Controller** (Fig 4d), where we employ MUSE to learn a representation model of the environment and subsequently train a controller over the state representation  $z_\pi$  (explicit, independent, hierarchical);

To evaluate the role of the training scheme in the agent’s performance, we also include a variation of the implicit agent *Multimodal Controller (D)* that employs a training method similar to MUSE, in which modalities are randomly dropped while learning the policy.

We evaluate the agents on two recently-proposed multimodal scenarios for deep reinforcement learning<sup>3</sup>: the multimodal Pendulum [2] and the HYPERHOT scenarios, described in Appendix.


<sup>3</sup>The multimodal Atari games are taken from <https://github.com/miguelsvasco/multimodal-atari-games>

We adopt the same RL algorithms, training hyper-parameters and network architectures of [25], due to the similarity in evaluation scenarios. We compare the performance of the RL agents when directly using the policy learned from joint-modality observations in scenarios with possible missing modalities, without any additional training. Fig 5a and 5b summarize the total reward collected per episode, for the Pendulum and HYPERHOT scenarios, respectively. Results are averaged over 100 episodes and 10 randomly seeded runs. Numerical results shown in Appendix.

The results show that the MUSE + Controller agent is the only agent able to act robustly in incomplete perceptual conditions, regardless of the modalities available. In the Pendulum scenario, the performance of our agent is similar when provided with image or sound observations, as both modalities hold information to fully describe the state of the environment. In the HYPERHOT scenario, the sound modality is unable to describe the complete state of the environment and, as such, the MUSE + Controller agent provided with image observations outperforms the one provided with sound observations. However, even in the latter case our agent is able to successfully complete the majority of episodes, shown by the positive average reward.

**4.1.1 Implicit vs. Explicit.** The results in Fig. 5 regarding the performance of the agents when provided with joint-modality observations highlight the benefits (and challenges) of learning an explicit state representation. In the HYPERHOT scenario, the implicit representation agent Multimodal Controller, which learns a policy directly from data, outperforms all other agents when provided with joint-modality observations. This result reveals the challenge of learning an accurate state representation in complex scenarios, suitable for RL tasks. However, in the simpler Pendulum scenario, the explicit representation agents are able to learn a robust low-dimensional state representation, performing on par with the Multimodal Controller agent when provided with joint-modality information.

**4.1.2 Fusion vs. Independent.** The Multimodal Controller and the VAE + Controller agents struggle to act when provided with only



**Figure 6: Cross-modality image generation from sound information in the multimodal Atari Game scenarios. MUSE is the only explicit representation model able to generate coherent image information from sounds. Best viewed with zoom.**

sound observations in both scenarios. Such behaviour confirms the robustness issue raised in Section 3: solutions that consider the fixed fusion of perceptions, both with implicit and explicit representations, are unable to robustly estimate the state of the environment when the agent is provided with incomplete perceptions. Moreover, the results of the variations with dropout show that the improved performance of the independent models is not due to their training scheme: the improvement by considering dropout with fused observations is limited for the Multimodal Controller (D) agent.

**4.1.3 Flat vs. Hierarchical.** The MVAE + Controller agent also struggles to act when provided only with sound, despite an improvement in performance in comparison with the fusion solutions. This highlights the issue raised in Section 3 regarding multimodal representation models with a flat latent variable: by employing a single representation space (of finite capacity) to encode information from all modalities, the model learns overconfident experts, neglecting information from lower-dimensional modalities [24]. On the other hand, the hierarchical design of MUSE allows to balance the complexity of each modality by encoding a joint-representation from low-dimensional, modality-specific, representations.

**4.1.4 Qualitative Evaluation.** We can evaluate qualitatively the performance of explicit representation models under incomplete perceptual experience by observing the image state generated from sound information, as shown in Fig. 6. The samples show that the MUSE + Controller solution is the only able to encode a robust estimation of state from low-dimensional information. In the Pendulum scenario, the MUSE agent is able to perfectly reconstruct the image observation of the game state, when only provided with sound information. This is coherent with the results in Fig. 5 regarding the similar performance when the agent is provided observations from either modality. In the HYPERHOT scenario, sound information is unable to perfectly account for the image state of the environment: the images generated from sound information are unable to precisely describe the state of the environment, e.g. the number and position of enemies. This is also coherent with the results of Fig. 5 regarding the lower performance by the agent when provided with sound observations in comparison to the agent provided with

image observations. However, the image reconstructions show that, even when encoded from sound, the multimodal representation contains fundamental information for the actuation of the agent, such as the position of projectiles near the agent. On the other hand, in both scenarios, the VAE + Controller and the MVAE + Controller agents are unable to reconstruct the (image) game state from sound information, evidence of their lack of robustness to missing modalities.

## 4.2 Generative Performance of MUSE

We now evaluate the potential of MUSE in learning a multimodal representation in scenarios with more complex and numerous modalities available to the agent. We do so by evaluating the generative performance of our model against two state-of-the-art multimodal variational autoencoders,<sup>4</sup> the MVAE [31] and MMVAE [24], in literature-standard multimodal scenarios: the MNIST dataset [15], the CelebA [19] and the MNIST-SVHN scenario [21].

We employ the authors’ publicly available code<sup>5</sup> and training loss functions, without importance-weighted sampling, as well as the suggested hyper-parameters (when available). In each scenario, we employ the same representation capacity (total dimensionality of representation spaces) in all models. Code, model architectures and training hyper-parameters are presented in Appendix.

We compute standard likelihood-based metrics (marginal and joint), using both the joint variational posterior and the single variational posterior, averaged over 5 independently-seeded runs. Our results are summarized in Table 1 and we present image samples generated from label information in Fig. 7 (more in Appendix).

**4.2.1 Single-Modality Performance.** In the two-modality scenarios, MUSE outperforms all baselines in terms of the image marginal likelihood,  $\log p(\mathbf{x}_1)$ . Regarding the attribute marginal likelihood  $\log p(\mathbf{x}_2)$ , MUSE performs on par with MMVAE in the MNIST dataset, despite employing a modality-specific latent representation 16 times smaller. The results in the three-modality scenario again

<sup>4</sup>We selected these models as both are agnostic to the nature of the modalities and scalable to large number of modalities, similarly to MUSE (as shown in Section 5)

<sup>5</sup>The MVAE model is taken from <https://github.com/mhw32/multimodal-vae-public> and the MMVAE model is taken from <https://github.com/iffid/mmvae>.

**Table 1: Standard metrics for generative performance in the different datasets (best results in bold). We used 5000 importance samples for the MNIST dataset and 1000 importance samples for the CelebA dataset and the MNIST-SVHN scenario. All results averaged over 5 independent runs. Higher is better.**

(a) MNIST ( $x_1$ - Image; $x_2$ - Label)					
Model	$\log p(x_1)$	$\log p(x_2)$	$\log p(x_1, x_2)$	$\log p(x_1   x_2)$	$\log p(x_2   x_1)$
<b>MUSE (Ours)</b>	<b>-24.06 ± 0.03</b>	<b>-2.41 ± 0.02</b>	-34.72 ± 1.47	-33.97 ± 1.16	- 4.73 ± 0.15
MVAE	-24.45 ± 0.14	<b>-2.38 ± 0.01</b>	<b>-25.46 ± 0.28</b>	<b>-29.23 ± 1.00</b>	- 9.60 ± 0.38
MMVAE	-25.96 ± 0.11	-2.82 ± 0.07	-	-39.94 ± 0.48	-10.01 ± 0.23
(b) CelebA ( $x_1$ - Image; $x_2$ - Attributes)					
Model	$\log p(x_1)$	$\log p(x_2)$	$\log p(x_1, x_2)$	$\log p(x_1   x_2)$	$\log p(x_2   x_1)$
<b>MUSE (Ours)</b>	<b>-163.63 ± 0.13</b>	-19.69 ± 0.09	-288.25 ± 3.70	<b>-394.65 ± 2.22</b>	- 47.48 ± 0.46
MVAE	-165.88 ± 0.40	<b>-15.21 ± 0.04</b>	<b>-181.56 ± 0.44</b>	-471.50 ± 1.67	- 72.36 ± 0.53
MMVAE	-548.11 ± 10.46	-28.43 ± 0.25	-	-787.24 ± 4.07	-108.50 ± 0.76
(c) MNIST-SVHN ( $x_1$ - MNIST image; $x_2$ - SVHN image; $x_3$ - Label)					
Model	$\log p(x_1)$	$\log p(x_2)$	$\log p(x_1   x_2)$	$\log p(x_1   x_2, x_3)$	$\log p(x_2   x_1)$
<b>MUSE (Ours)</b>	<b>-24.23 ± 0.02</b>	<b>-36.03 ± 0.04</b>	<b>-39.93 ± 1.45</b>	<b>-37.66 ± 1.42</b>	- 64.99 ± 3.35
MVAE	-24.34 ± 0.07	<b>-36.17 ± 0.22</b>	-43.44 ± 0.19	-40.96 ± 0.86	- 55.51 ± 0.18
MMVAE	-28.21 ± 0.06	-41.33 ± 0.23	-54.84 ± 0.47	-	-166.14 ± 0.76

show that MUSE outperforms the baselines in single-modality generation. The hierarchical design of MUSE allows the adjustment of the capacity of each modality-specific latent space to the inherent complexity of the corresponding modality.

**4.2.2 Joint-Modality Performance.** Regarding the joint-modality performance,  $\log p(x_1, x_2)$ , in MUSE,  $z_\pi$  encodes low-dimensional codes abstracted from modality data. This leads to a multimodal reconstruction process that loses some information for higher-dimensional modalities (e.g., image), yet accentuates distinctive features that define the observed phenomena (e.g., digit class), leading to lower generative performance. This phenomena can also be seen in reconstruction from higher latent variables in single-modality hierarchical generative models [10]. However, as seen in Section 4.1, such abstraction leads to improved state representation when the agent is provided with complex observations.

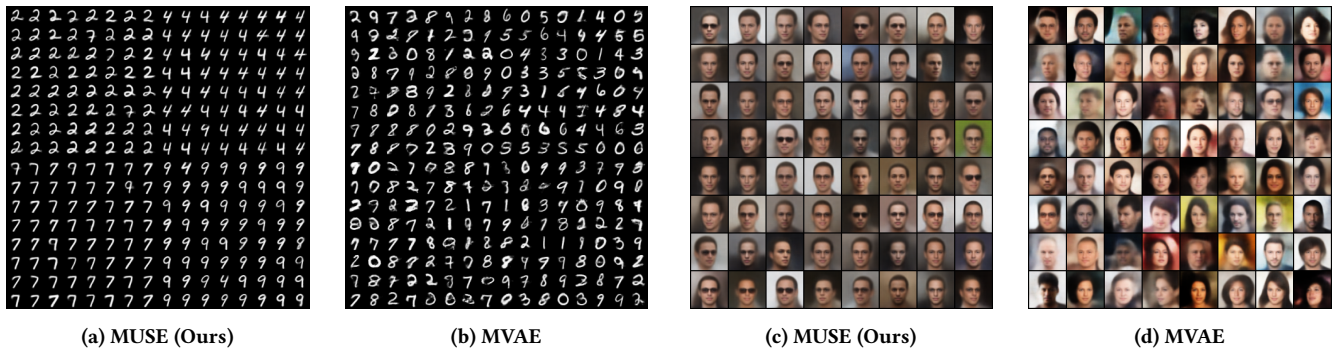
**4.2.3 Cross-Modality Performance.** In the three-modality MNIST-SVHN scenario, MUSE, in contrast with MMVAE, is able to encode information provided by two modalities ( $x_2$  and  $x_3$ ) for cross-modality generation, shown by the increase in performance from  $\log p(x_1 | x_2)$  to  $\log p(x_1 | x_2, x_3)$ . MUSE also outperforms all other baselines in the CelebA dataset in regards to the cross-modality metrics,  $\log p(x_1 | x_2)$  and  $\log p(x_2 | x_1)$ . The quantitative results are aligned with the qualitative evaluation of the generated images in Fig. 7: MUSE is the only model able to generate high-quality and diverse image samples, semantically coherent with attribute information. However, in the MNIST dataset, the quantitative results for the label-to-image conditional generation,  $\log p(x_1 | x_2)$ , are at odds with the qualitative assessment of the image samples (Fig. 7): MVAE seems to outperform MUSE in terms of  $\log p(x_1 | x_2)$ , yet

fails to generate coherent image samples from label information, as shown in Fig. 7. This apparent contradiction motivates the need for the development of more suitable metrics to evaluate the performance of generative models when provided with missing modality information.

## 5 RELATED WORK

Reinforcement Learning (RL) agents often employ representation models to learn to act in their equally often complex environments. Such representations can encompass the structure of the environment (state representations) [16] or of the actions of the agent (action representations) [3]. Low-dimensional representations of the environment allow agents to learn efficiently how to perform tasks. Initial works considered state aggregation algorithms employing bisimulation metrics to provide sample-efficient learning [5, 7, 17]. However, such methods are not easily scalable to scenarios with large-scale state spaces or complex dynamics. Variational autoencoders were first employed by Higgins *et al.* to learn a low-dimensional representation of observations provided by the environment that allows for zero-shot policy transfer to a target domain [11]. Similarly, VAEs are employed as a representation model in the *World Model* [8] and *Dreamer* [9] frameworks, learning to encode a low-dimensional code for image observations of the agent, at each time-step. However, all the previously discussed works assume that the agent is provided only, and always, with image observations.

Recently, the problem of transferring policies across different perceptual modalities has been proposed and addressed for Atari Games with two modalities [25]. However, such work assumes that the agent only has access to a single modality during training and



**Figure 7: Cross-modality image generation from label information: in the MNIST dataset considering  $x_2 = \{2, 4, 7, 9\}$ , in the CelebA dataset considering  $x_2 = \{\text{Male, Eyeglasses, Black Hair, Receding Hairline, Goatee}\}$ . The MUSE model is the only able to generate high-quality, varied and coherent image samples from label information. Best viewed with zoom.**

task execution. In this work, we uplift such restriction and show that agents using MUSE as a sensory representation model are able to learn policies considering joint-modality observations and directly reuse such policies, without further training, in scenarios of missing modality information.

In order to learn representations of multimodal data, multimodal VAE (MVAE) models are widely employed, extending the original model [13]. Early approaches, such as the AVAE and JMVAE models, considered forcing the individual representations encoded from each modality to match [27, 32]. However, these models require individual neural-networks for each modality and each combination of modalities, limiting their applicability in scenarios with more than two modalities [14]. Other approaches proposed the factorization of the multimodal representation into separate, independent, representations, such as the MFM model [12, 28]. However, such models require explicit semantic (label) information to encode the multimodal representation, thus unable to perform cross-modal generation considering all possible sets of available modalities.

Recently, two approaches were proposed to learn a scalable multimodal representation, differing on the function responsible for merging multimodal information: the Multimodal VAE (MVAE) [31], employing a Product-of-Experts (PoE) solution to merge modality-specific information, and the Mixture-of-Experts MVAE (MMVAE), which employs a Mixture-of-Experts (MoE) solution to encode the multimodal representation [24]. However, the MVAE model struggles to perform cross-modal generation robustly across all modalities due to the overconfident experts issue, neglecting information from lower-dimensional modalities, as shown in [24]. On the other hand, the MMVAE is computationally less efficient, relying on importance sampling training schemes and  $M^2$  passes through the decoder networks to train the MoE encoder, hindering its application in scenarios with arbitrary large number of modalities. Moreover, the model assumes that the modalities are of comparable complexity and, due to the MoE sampling method, is unable to merge information from multiple modalities for downstream tasks.

We summarize the main differences that distinguish our work. MUSE is simultaneously: (i) able to scale to a large number of input modalities without requiring combinatorial number of networks, unlike latent-representation approximation methods [14, 27, 32]; (ii)

agnostic to the nature of the provided modalities, unlike factorized approaches [12, 28] and MMVAE [24]; (iii) able to robustly perform cross-modality inference regardless of the complexity of the target modality, unlike MVAE [31]; (iv) is computationally efficient to train, unlike MMVAE [24].

## 6 CONCLUSION

We addressed the question of *sensing the world*: how to learn a multimodal representation of a RL agent’s environment that allows the execution of tasks under missing modality information. We contributed MUSE, a novel multimodal sensory representation model that considers hierarchical representation spaces. Our results show that, with MUSE, RL agents are able to perform tasks under such incomplete perceptions, outperforming other literature-standard designs. Moreover, we showed that the performance of MUSE scales to more complex scenarios, outperforming other state-of-the-art models. We can envision scenarios where MUSE would allow agents to act with unexpected damaged sensors (e.g. autonomous cars) or respecting privacy concerns (e.g. virtual assistants).

The introduction of hierarchy in multimodal generative models broadens the design possibilities of such models. In future work, we will exploit the modularity arising from having modality-specific representations, exploring the use of pretrained representation models. Additionally, we will further explore evaluation schemes to assess the quality of multimodal representations beyond standard likelihood metrics, such as following Poklukar et al. [23]. Finally, we will also explore how to provide agents with robustness to noisy observations and develop mechanisms to reason about their confidence on the information provided by different sensors.

## 7 ACKNOWLEDGEMENTS

This work was partially supported by national funds through the Portuguese Fundação para a Ciência e a Tecnologia under project UIDB/50021/2020 (INESC-ID multi annual funding) and project PTDC/CCI-COM/5060/2021. In addition, this research was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No. 952215. This work was also supported by funds from Europe Research Council under project BIRD 884887. The first author acknowledges the Fundação para a Ciência e a Tecnologia PhD grant SFRH/BD/139362/2018.



## REFERENCES

- [1] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*, 2017.
- [2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [3] Yash Chandak, Georgios Theodorou, James Kostas, Scott Jordan, and Philip Thomas. Learning action representations for reinforcement learning. In *International Conference on Machine Learning*, pages 941–950. PMLR, 2019.
- [4] Mark Collier, Alfredo Nazabal, and Christopher KI Williams. Vaes in the presence of missing data. *arXiv preprint arXiv:2006.05301*, 2020.
- [5] Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In *UAI*, volume 4, pages 162–169, 2004.
- [6] Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G Bellemare. Deepmdp: Learning continuous latent space models for representation learning. In *International Conference on Machine Learning*, pages 2170–2179, 2019.
- [7] Robert Givan, Thomas Dean, and Matthew Greig. Equivalence notions and model minimization in markov decision processes. *Artificial Intelligence*, 147(1-2): 163–223, 2003.
- [8] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- [9] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- [10] Jakob D Havtorn, Jes Frelsen, Søren Hauberg, and Lars Maaløe. Hierarchical vaes know what they don't know. *arXiv preprint arXiv:2102.08248*, 2021.
- [11] Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. Darla: Improving zero-shot transfer in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1480–1490. JMLR. org, 2017.
- [12] Wei-Ning Hsu and James Glass. Disentangling by partitioning: A representation learning framework for multimodal sensory data. *arXiv preprint arXiv:1805.11264*, 2018.
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [14] Timo Korthals, Daniel Rudolph, Jürgen Leitner, Marc Hesse, and Ulrich Rückert. Multi-modal generative models for learning epistemic active sensing. In *2019 IEEE International Conference on Robotics and Automation*, 2019.
- [15] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [16] Timothée Lesort, Natalia Díaz-Rodríguez, Jean-François Goudou, and David Filliat. State representation learning for control: An overview. *Neural Networks*, 108: 379–392, 2018.
- [17] Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a unified theory of state abstraction for mdps. *ISAIM*, 4:5, 2006.
- [18] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [19] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [20] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [21] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [22] Sarah R Partan. Multimodal shifts in noise: switching channels to communicate through rapid environmental change. *Animal Behaviour*, 124:325–337, 2017.
- [23] Petra Poklukar, Anastasia Varava, and Danica Kragic. Geomca: Geometric evaluation of data representations. *arXiv preprint arXiv:2105.12486*, 2021.
- [24] Yuge Shi, N Siddharth, Brooks Paige, and Philip Torr. Variational mixture-of-experts autoencoders for multi-modal deep generative models. In *Advances in Neural Information Processing Systems*, pages 15692–15703, 2019.
- [25] Rui Silva, Miguel Vasco, Francisco S Melo, Ana Paiva, and Manuela Veloso. Playing games in the dark: An approach for cross-modality transfer in reinforcement learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1260–1268, 2020.
- [26] Richard Sutton and Andrew Barto. *Reinforcement Learning: An Introduction*. MIT press Cambridge, 1998.
- [27] Masahiro Suzuki, Kotaro Nakayama, and Yutaka Matsuo. Joint multimodal learning with deep generative models. *arXiv preprint arXiv:1611.01891*, 2016.
- [28] Yao-Hung Hubert Tsai, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency, and Ruslan Salakhutdinov. Learning factorized multimodal representations. In *International Conference on Representation Learning*, 2019.
- [29] Mel Vecerik, Todd Hester, Jonathan Scholz, Fumin Wang, Olivier Pietquin, Bilal Piot, Nicolas Heess, Thomas Rothörl, Thomas Lampe, and Martin Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*, 2017.
- [30] Christopher Watkins. *Learning from delayed rewards*. PhD thesis, Cambridge University, 1989.
- [31] Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5575–5585, 2018.
- [32] Hang Yin, Francisco S Melo, Aude Billard, and Ana Paiva. Associate latent encodings in learning from demonstrations. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [33] Marvin Zhang, Sharad Vikram, Laura Smith, Pieter Abbeel, Matthew Johnson, and Sergey Levine. Solar: Deep structured representations for model-based reinforcement learning. In *International Conference on Machine Learning*, pages 7444–7453. PMLR, 2019.