

Crossmodal Attentive Skill Learner

Shayegan Omidshafiei

Massachusetts Institute of Technology
Cambridge, Massachusetts
shayegan@mit.edu

Jason Pazis*

Amazon Alexa
Cambridge, Massachusetts
pazisj@amazon.com

Dong-Ki Kim

Massachusetts Institute of Technology
Cambridge, Massachusetts
dkkim93@mit.edu

Jonathan P. How

Massachusetts Institute of Technology
Cambridge, Massachusetts
jhow@mit.edu

ABSTRACT

This paper introduces the Crossmodal Attentive Skill Learner (CASL), integrated with the recently-introduced Asynchronous Advantage Option-Critic (A2OC) architecture [15] to enable hierarchical reinforcement learning across multiple sensory inputs. We provide concrete examples where the approach not only improves performance in a single task, but accelerates transfer to new tasks. We demonstrate the attention mechanism anticipates and identifies useful latent features, while filtering irrelevant sensor modalities during execution. We modify the Arcade Learning Environment [7] to support audio queries, and conduct evaluations of crossmodal learning in the Atari 2600 games H.E.R.O. and Amidar. Finally, building on the recent work of Babaeizadeh et al. [4], we open-source a fast hybrid CPU-GPU implementation of CASL.¹

KEYWORDS

hierarchical learning; options; reinforcement learning; attention

ACM Reference Format:

Shayegan Omidshafiei, Dong-Ki Kim, Jason Pazis, and Jonathan P. How. 2018. Crossmodal Attentive Skill Learner. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 8 pages.

1 INTRODUCTION

Intelligent agents should be capable of disambiguating local sensory streams to realize long-term goals. In recent years, the combined progress of computational capabilities and algorithmic innovations has afforded reinforcement learning (RL) [39] approaches the ability to achieve this desiderata in impressive domains, exceeding expert-level human performance in tasks such as Atari and Go [30, 36]. Nonetheless, many of these algorithms thrive primarily in well-defined mission scenarios learned in isolation from one another; such monolithic approaches are not sufficiently scalable missions where goals may be less clearly defined, and sensory inputs found salient in one domain may be less relevant in another.

How should agents learn effectively in domains of high dimensionality, where tasks are durative, agents receive sparse feedback,

and sensors compete for limited computational resources? One promising avenue is hierarchical reinforcement learning (HRL), focusing on problem decomposition for learning transferable skills. Temporal abstraction enables exploitation of domain regularities to provide the agent hierarchical guidance in the form of options or sub-goals [21, 40]. Options help agents improve learning by mitigating scalability issues in long-duration missions, by reducing the effective number of decision epochs. In the parallel field of supervised learning, temporal dependencies have been captured proficiently using attention mechanisms applied to encoder-decoder based sequence-to-sequence models [5, 24]. *Attention* empowers the learner to focus on the most pertinent stimuli and capture longer-term correlations in its encoded state, for instance to conduct neural machine translation or video captioning [43, 44]. Recent works also show benefits of spatio-temporal attention in RL [29, 37].

One can interpret the above approaches as conducting dimensionality reduction, where the target dimension is *time*. In view of this insight, this paper proposes an RL paradigm exploiting hierarchies in the dimensions of *time* and *sensor modalities*. Our aim is to learn rich skills that attend to and exploit pertinent crossmodal (multi-sensor) signals at the appropriate moments. The introduced crossmodal skill learning approach largely benefits an agent learning in a high-dimensional domain (e.g., a robot equipped with many sensors). Instead of the expensive operation of processing and/or storing data from all sensors, we demonstrate that our approach enables such an agent to focus on important sensors; this, in turn, leads to more efficient use of the agent’s limited computational and storage resources (e.g., its finite-sized memory).

In this paper, we focus on combining two sensor modalities: audio and video. While these modalities have been previously used for supervised learning [31], to our knowledge they have yet to be exploited for crossmodal skill learning. We provide concrete examples where the proposed HRL approach not only improves performance in a single task, but accelerates transfer to new tasks. We demonstrate the attention mechanism anticipates and identifies useful latent features, while filtering irrelevant sensor modalities during execution. We also show preliminary results in the Arcade Learning Environment [7], which we modified to support audio queries. In addition, we provide insight into how our model functions internally by analyzing the interactions of attention and memory. Building on the recent work of Babaeizadeh et al. [4], we open-source a fast hybrid CPU-GPU implementation of our framework. Finally, note that despite this paper’s focus on audio-video sensors, the

*Work done prior to Amazon involvement of the author, and does not reflect views of the Amazon company.

¹Code available at <https://github.com/shayegano/CASL>

framework presented is general and readily applicable to additional sensory inputs.

2 RELATED WORK

Our work is most related to crossmodal learning approaches that take advantage of multiple input sensor modalities. Fusion of multiple modalities or sources of information is an active area of research. Works in the diverse domains of sensor fusion in robotics [12, 25, 33], audio-visual fusion [6, 8, 31, 38], and image-point cloud fusion [1, 9] have shown that models utilizing multiple modalities tend to outperform those learned from unimodal inputs. In general, approaches for multimodal fusion can be broadly classified depending on the means of integration of the various information sources. Filtering-based frameworks (e.g., the Extended Kalman Filter) are widely used to combine multi-sensor readings in the robotics community [12, 25, 33]. In machine learning, approaches based on graphical models [6, 8] and conditional random fields [22] have been used to integrate multimodal features [1, 9]. More recently, deep learning-based approaches that learn a representations of features across multiple modalities have been introduced [14, 19, 31, 38].

A variety of attentive mechanisms have been considered in recent works, primarily in application to supervised learning. Temporal attention mechanisms have been successfully applied to the field of neural machine translation, for instance in [5, 24], where encoder-decoder networks work together to transform one sequence to another. Works also exist in multimodal attention for machine translation, using video-text inputs [10]. Spatial attention models over image inputs have also been combined with Deep Recurrent Q-Networks [16] for RL [37]. Works have also investigated spatially-attentive agents trained via RL to conduct image classification [3, 29]. As we later demonstrate, the crossmodal attention-based approach used in this paper enables filtering of irrelevant sensor modalities, leading to improved learning and more effective use of the agent’s memory.

There exists a large body of HRL literature, targeting both fully and partially-observable domains. Our work leverages the options framework [40], specifically the recent Asynchronous Advantage Option-Critic (A2OC) [15] algorithm, to learn durative skills. HRL is an increasingly-active field, with a number of recent works focusing on learning human-understandable and/or intuitive skills. In [2], annotated descriptors of policies are used to enable multitask RL using options. Multitask learning via parameterized skills is also considered in [13], where a classifier and regressor are used to, respectively, identify the appropriate policy to execute, then map to the appropriate parameters of said policy. Construction of *skill chains* is introduced in [20] for learning in continuous domains. In [26], option discovery is conducted through eigendecomposition of MDP transition matrices, leading to transferable options that are agnostic of the task reward function. FeUdal Networks [41] introduce a two-level hierarchy, where a manager defines sub-goals, and a worker executes primitive actions to achieve them. A related HRL approach is also introduced in [21], where sub-goals are hand-crafted by a domain expert. Overall, the track record of hierarchical approaches for multitask and transfer learning leads us to use them

as a basis for the proposed framework, as our goal is to learn scalable policies over high-volume input streams.

While the majority of works utilizing multi-sensory and attentive mechanisms focus on supervised learning, our approach targets RL. Specifically, we introduce an HRL framework that combines cross-modal learning, attentive mechanisms, and temporal abstraction to learn durative skills.

3 BACKGROUND

This section summarizes Partially Observable Markov Decision Processes (POMDPs) and options, which serve as foundational frameworks for our approach.

3.1 POMDPs

This work considers an agent operating in a partially-observable stochastic environment, modeled as a POMDP $(\mathbb{S}, \mathbb{A}, \mathbb{O}, \mathcal{T}, \mathcal{O}, \mathcal{R}, \gamma)$ [18]. \mathbb{S} , \mathbb{A} , and \mathbb{O} are, respectively, the state, action, and observation spaces. At timestep t , the agent executes action $a \in \mathbb{A}$ in state $s \in \mathbb{S}$, transitions to state $s' \sim \mathcal{T}(s, a, s')$, receives observation $o \sim \mathcal{O}(o, s', a)$, and reward $r_t = \mathcal{R}(s, a) \in \mathbb{R}$. The value of state s under policy $\pi : \text{Dist}(\mathbb{S}) \rightarrow \mathbb{A}$ is the expected return $V_\pi(s) = \mathbb{E}[\sum_{k=0}^H \gamma^k r_{t+k+1} | s_t = s]$, given timestep t , horizon H , and discount factor $\gamma \in [0, 1)$. The objective is to learn an optimal policy π^* , which maximizes the value.

As POMDP agents only receive noisy observations of the latent state, policy π typically maps from the agent’s belief (distribution over states) to the next action. Recent work has introduced Deep Recurrent Q-Networks (DRQNs) [16] for RL in POMDPs, leveraging recurrent Neural Networks (RNNs) that inherently maintain an internal state h^t to compress input history until timestep t .

3.2 Options

The framework of *options* provides an RL agent the ability to plan using temporally-extended actions [40]. Option $\omega \in \Omega$ is defined by initiation set $\mathbb{I} \subseteq \mathbb{S}$, intra-option policy $\pi_\omega : \mathbb{S} \rightarrow \text{Dist}(\mathbb{A})$, and termination condition $\beta_\omega : \mathbb{S} \rightarrow [0, 1]$. A policy over options π_Ω chooses an option among those that satisfy the initiation set. The selected option executes its intra-option policy until termination, upon which a new option is chosen. This process iterates until the goal state is reached. Recently, the Asynchronous Advantage Actor-Critic framework (A3C) [28] has been applied to POMDP learning in a computationally-efficient manner by combining parallel actor-learners and Long Short-Term Memory (LSTM) cells [17]. Asynchronous Advantage Option-Critic (A2OC) extends A3C and enables learning option-value functions, intra-option policies, and termination conditions in an end-to-end fashion [15]. The option-value function models the value of state $s \in \mathbb{S}$ in option $\omega \in \Omega$,

$$Q_\Omega(s, \omega) = \sum_a \pi_\omega(a|s) \left(r(s, a) + \gamma \sum_{s'} \mathcal{T}(s'|s, a) U(s', \omega) \right), \quad (1)$$

where $a \in \mathbb{A}$ is a primitive action and $U(s', \omega)$ represents the option utility function,

$$U(s', \omega) = (1 - \beta_\omega(s')) Q_\Omega(\omega, s') + \beta_\omega(s') (V_\Omega(s') - c). \quad (2)$$

A2OC introduces deliberation cost, c , in the utility function to address the issue of options terminating too frequently. Intuitively,

the role of c is to impose an added penalty when options terminate, enabling regularization of termination frequency. The value function over options, V_Ω , is defined,

$$V_\Omega(s') = \sum_{\omega} \pi_\Omega(\omega|s') Q_\Omega(\omega, s'), \quad (3)$$

where π_Ω is the policy over options (e.g., an epsilon-greedy policy over Q_Ω). Assuming use of a differentiable representation, option parameters can be learned using gradient descent. Readers are referred to [15] for more details.

4 APPROACH

Our goal is to design a mechanism that enables the learner to modulate high-dimensional sensory inputs, focusing on pertinent stimuli that may lead to more efficient skill learning. This section presents motivations behind attentive skill learning, then introduces the proposed framework.

4.1 Attentive Mechanisms

Before presenting the proposed architecture, let us first motivate our interests towards attentive skill learning. One might argue that the combination of deep learning and RL already affords agents the representation learning capabilities necessary for proficient decision-making in high-dimensional domains; i.e., why the need for crossmodal attention?

Our ideas are motivated by the studies in behavioral neuroscience that suggest the interplay of attention and choice bias humans' value of information during learning, playing a key factor in solving tasks with high-dimensional information streams [23]. Works studying learning in the brain also suggest a natural pairing of attention and hierarchical learning, where domain regularities are embedded as priors into skills and combined with attention to alleviate the curse of dimensionality [32]. Works also suggest attention plays a role in the intrinsic curiosity of agents during learning, through direction of focus to regions predicted to have high reward [27], high uncertainty [34], or both [35].

In view of these studies, we conjecture that crossmodal attention, in combination with HRL, improves representations of relevant environmental features that lead to superior learning and decision-making. Specifically, using crossmodal attention, agents combine internal beliefs with external stimuli to more effectively exploit multiple modes of input features for learning. As we later demonstrate, our approach captures temporal crossmodal dependencies, and enables faster and more proficient learning of skills in the domains examined.

4.2 Crossmodal Attentive Skill Learner

We propose Crossmodal Attentive Skill Learner (CASL), a novel framework for HRL. One may consider many blueprints for integration of multi-sensory attention into the options framework. Our proposed architecture is primarily motivated by the literature that taxonomizes attention into two classes: *exogeneous* and *endogeneous*. The former is an involuntary mechanism triggered automatically by the inherent saliency of the sensory inputs, whereas the latter is driven by the intrinsic and possibly long-term goals, intents, and beliefs of the agent [11]. Previous attention-based neural architectures take advantage of both classes, for instance, to solve natural

language processing problems [42]; our approach follows a similar schema.

The CASL network architecture is visualized in Figure 1. Let $M \in \mathbb{N}$ be the number of sensor modalities (e.g., vision, audio, etc.) and x_m denote extracted features from the m -th sensor, where $m \in \{1, \dots, M\}$. For instance, x_m may correspond to feature outputs of a convolutional neural network given an image input. Given extracted features for all M sensors at timestep t , as well as hidden state h^{t-1} , the proposed crossmodal attention layer learns the relative importance of each modality $\alpha^t \in \Delta^{M-1}$, where Δ^{M-1} is the $(M-1)$ -simplex:

$$z^t = \tanh \left(\underbrace{\sum_{m=1}^M (W_m^T x_m^t + b_m)}_{\text{Exogeneous attention}} + \underbrace{W_h^T h^{t-1} + b_h}_{\text{Endogeneous attention}} \right) \quad (4)$$

$$\alpha^t = \text{softmax} \left(W_z^T z^t + b_z \right) \quad (5)$$

$$x_{\oplus} = \begin{cases} \sum_{m=1}^M \alpha_m^t x_m^t & (\text{Summed attention}) \\ \left[(\alpha_1^t x_1^t)^T, \dots, (\alpha_M^t x_M^t)^T \right]^T & (\text{Concatenated attention}) \end{cases} \quad (6)$$

Weight matrices W_m , W_h , W_z and bias vectors b_m , b_h , b_z are trainable parameters and nonlinearities are applied element-wise.

Both exogeneous attention over sensory features x_m^t and endogeneous attention over LSTM hidden state h^{t-1} are captured in Equation (4). The sensory feature extractor used in experiments consists of 3 convolutional layers, each with 32 filters of size 3×3 , stride 2, and ReLU activations. Attended features $\alpha_m^t x_m^t$ may be combined via summation or concatenation (Equation (6)), then fed to an LSTM cell. The LSTM output captures temporal dependencies used to estimate option values, intra-option policies, and termination conditions (Q_Ω , π_ω , β_ω in Figure 1, respectively),

$$Q_\Omega(s, \omega) = W_{Q, \omega} h^t + b_{Q, \omega}, \quad (7)$$

$$\pi_\omega(a|s) = \text{softmax}(W_{\pi, \omega} h^t + b_{\pi, \omega}), \quad (8)$$

$$\beta_\omega(s) = \sigma(W_{\beta, \omega} h^t + b_{\beta, \omega}), \quad (9)$$

where weight matrices $W_{Q, \omega}$, $W_{\pi, \omega}$, $W_{\beta, \omega}$ and bias vectors $b_{Q, \omega}$, $b_{\pi, \omega}$, $b_{\beta, \omega}$ are trainable parameters for the current option ω , and $\sigma(\cdot)$ is the sigmoid function. Network parameters are updated using gradient descent. Entropy regularization of attention outputs α^t was found to encourage exploration of crossmodal attention behaviors during training.

5 EVALUATION

The proposed framework is evaluated on a variety of learning tasks with inherent reward sparsity and transition noise. We evaluate our approach in three domains: a door puzzle domain, a 2D-Minecraft like domain, and the Arcade Learning Environment [7]. These environments include challenging combinations of reward sparsity and/or complex audio-video sensory input modalities that may not always be useful to the agent. The first objective of our experiments is to analyze performance of CASL in terms of learning rate and transfer learning. The second objective is to understand relationships between attention and memory mechanisms (as captured in the LSTM cell state). Finally, we modify the Arcade Learning

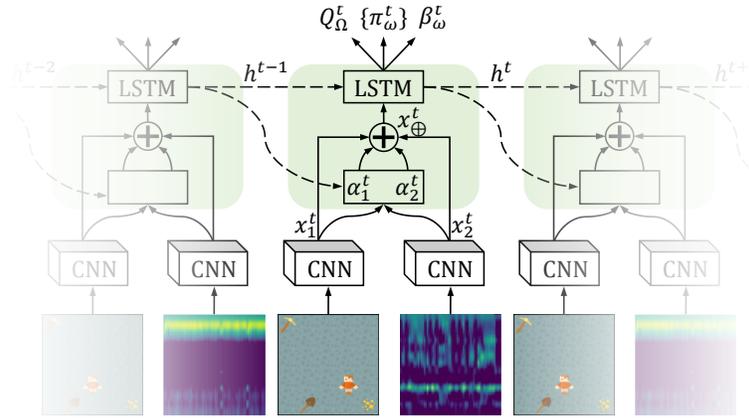


Figure 1: CASL network architecture enables attention-based learning over multi-sensory inputs. Green highlighted region indicates crossmodal attention LSTM cell, trained via backpropagation through time.

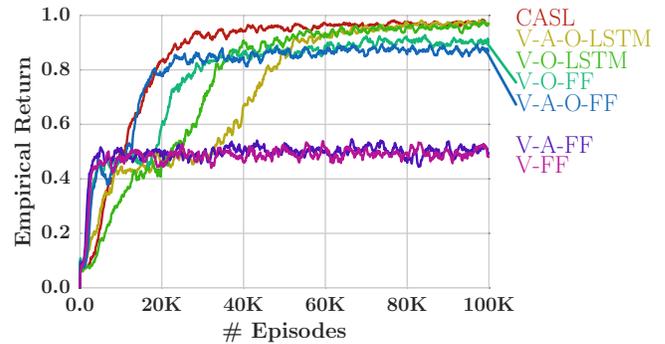
Environment to support audio queries, and evaluate crossmodal learning in the Atari 2600 games H.E.R.O. and Amidar.

5.1 Crossmodal Learning and Transfer

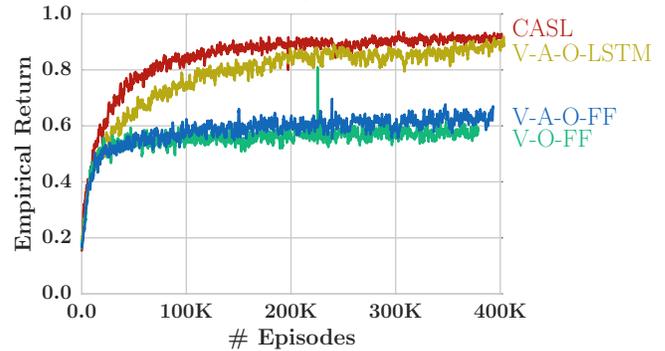
We first evaluate crossmodal attention in a sequential door puzzle game, where the agent spawns in a 2D world with two locked doors and a key at fixed positions. The key type is randomly generated, and its observable color indicates the associated door. The agent hears a fixed sound (but receives no reward) when adjacent to the key, and hears noise otherwise. The agent must find and pick up the key (which then disappears), then find and open the correct door to receive +1 reward ($\gamma = 0.99$). The game terminates upon opening of either door. The agent’s sensory inputs x_m^t are vision and audio spectrogram. This task was designed in such a way that audio is not necessary to achieve the task – the agent can certainly focus on learning a policy mapping from visual features to open the correct door. However, audio provides potentially useful signals that may accelerate learning, making this a domain of interest for analyzing the interplay of attention and sensor modalities.

Attention Improves Learning Rate. Figure 2a shows ablative training results for several network architectures. The three LSTM-based skill learners (including CASL) converge to the optimal value. Interestingly, the network that ignores audio inputs (V-O-LSTM) converges faster than its audio-enabled counterpart (V-A-O-LSTM), indicating the latter is overwhelmed by the extra sensory modality. Introduction of crossmodal attention enables CASL to converge faster than all other networks, using roughly half the training data of the others. The feedforward networks all fail to attain optimal value, with the non-option cases (V-A-FF and V-FF) repeatedly opening one door due to lack of memory of key color. Notably, the option-based feedforward nets exploit the option index to implicitly remember the key color, leading to higher value. Interplay between explicit memory mechanisms and use of options as pseudo-memory may be an interesting line of future work.

Attention Accelerates Transfer. We also evaluate crossmodal attention for transfer learning (Figure 2b), using the more promising option-based networks. The door puzzle domain is modified to



(a) Door puzzle domain.



(b) Door puzzle transfer learning.

Figure 2: CASL improves learning rate compared to other networks. Abbreviations: V: video, A: audio, O: options, FF feedforward net, LSTM: Long Short-Term Memory net.

randomize the key position, with pre-trained options from the fixed-position variant used for initialization. All networks benefit from an empirical return jumpstart of nearly 0.2 at the beginning of training, due to skill transfer. Once again, CASL converges fastest, indicating more effective use of the available audio-video data. While the asymptotic performance of CASL is only slightly higher

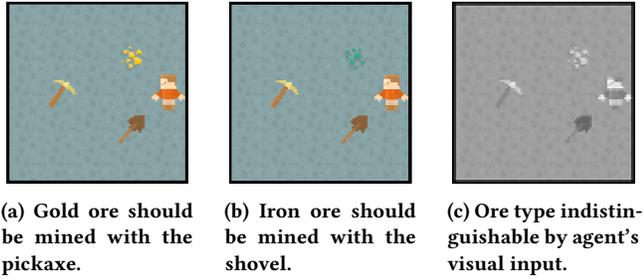


Figure 3: Mining domain. Ore type is indistinguishable by the grayscale visual observed by the agent.

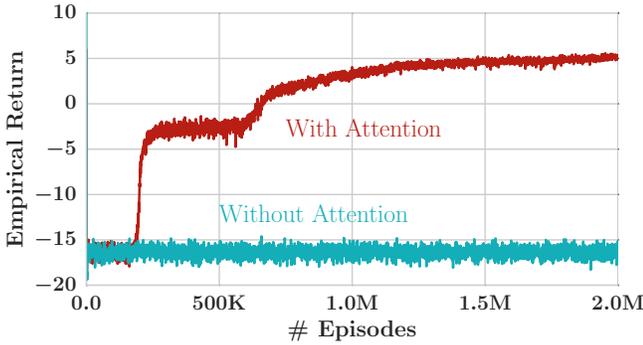


Figure 4: In the mining domain, the non-attentive network fails to learn, whereas the attentive network succeeds.

than the V-A-O-LSTM network, the reduction in number of samples needed to achieve a high score (e.g., after 100K episodes) makes it advantageous for domains with high sampling cost.

Attention Necessary to Learn in Some Domains. Temporal behaviors of the attention mechanism are also evaluated in a 2D Minecraft-like domain, where the agent must pick an appropriate tool (pickaxe or shovel) to mine either gold or iron ore (Figures 3a to 3c). Critically, the agent observes identical images for both ore types, but unique audio features when near the ore, making long-term audio storage necessary for selection of the correct tool. The agent receives +10 reward for correct tool selection, -10 for incorrect selection, and -1 step cost. Compared to the door puzzle game, the mining domain is posed in such a way that the interplay between audio-video features is emphasized. Specifically, an optimal policy for this task must utilize both audio and video features: visual inputs enable detection of locations of the ore, agent, tools, whereas audio is used to identify the ore type.

Visual occlusion of the ore type, interplay of audio-video features, and sparse positive rewards cause the non-attentive network to fail to learn in the mining domain, as opposed to the attentive case (Figure 4). Figure 5a plots a sequence of frames where the agent anticipates salient audio features as it nears the ore at $t = 6$, gradually increasing audio attention, then sharply reducing it to 0 after hearing the signal.

5.2 Interactions of Attention and Memory

While the anticipatory nature of crossmodal attention in the mining domain is interesting, it also points to additional lines of investigation regarding interactions of attention and updates of the agent's internal belief (as encoded in its LSTM cell state). Specifically, one might wonder whether it is necessary for the agent to place any attention on the non-useful audio signals prior to timestep $t = 6$ in Figure 5a, and also whether this behavior implies inefficient usage of its finite-size memory state.

Motivated by the above concerns, we conduct more detailed analysis of the interplay between the agent's attention and memory mechanisms as used in the CASL architecture (Figure 1). We first provide a brief overview of LSTM networks to enable more rigorous discussion of these attention-memory interactions. At timestep t , LSTM cell state C^t encodes the agent's memory given its previous stream of inputs. The cell state is updated as follows,

$$f^t = \sigma(W_f[x_\oplus, h^{t-1}] + b_f), \quad (10)$$

$$i^t = \sigma(W_i[x_\oplus, h^{t-1}] + b_i), \quad (11)$$

$$C^t = f^t \odot C^{t-1} + i^t \odot \tanh(W_C[x_\oplus, h^{t-1}] + b_C), \quad (12)$$

where f^t is the forget gate activation vector, i^t is the input gate activation vector, h^{t-1} is the previous hidden state vector, x_\oplus is attended feature vector, and \odot denotes the Hadamard product. Weights W_f , W_i , W_C and biases b_f , b_i , b_C are trainable parameters. The cell state update in Equation (12) first forgets certain elements (f^t term), and then adds contributions from new inputs (i^t term). Note that a forget gate activation of 0 corresponds to complete forgetting of the previous cell state element, and that an input gate activation of 1 corresponds to complete throughput of the corresponding input element.

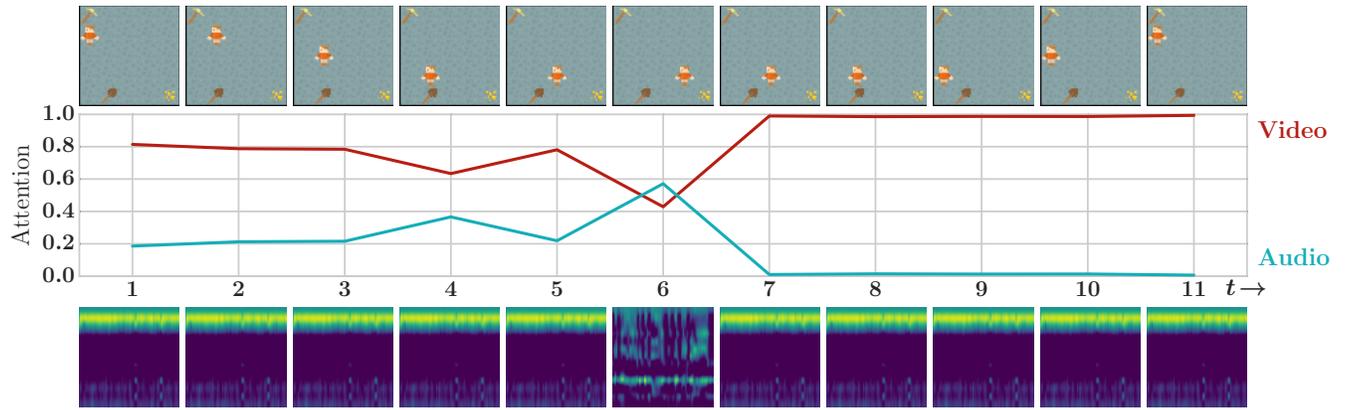
Our goal is to not only analyze the overall forget/input activations throughout the gameplay episode, but also to quantify the relative impact of each contributing variable (audio input, video input, hidden state, and bias term) to the overall activations. Many methods may be used for analysis of the contribution of explanatory variables in nonlinear models (i.e., Equations (10) to (12)). We introduce a means of quantifying the correlation of each variable with respect to the corresponding activation function. In the following, we focus on the forget gate activation, but the same analysis applies to the input gate. First, expanding the definition of forget gate activation (Equation (10)) assuming use of concatenated attention (Equation (6)) yields,

$$f^t = \sigma([W_{f_a}, W_{f_v}, W_{f_h}, b_f][\alpha_a x_a, \alpha_v x_v, h^{t-1}, I]), \quad (13)$$

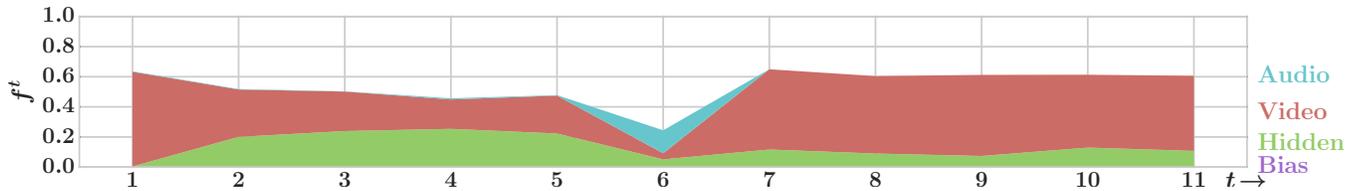
where x_a and x_v are, respectively, the audio and video input features, and I is the identity matrix. Define \hat{f}_m^t as the forget gate activation if the m -th contributing variable were removed. For example, if audio input x_a were to be removed, then,

$$\hat{f}_a^t = \sigma([W_{f_v}, W_{f_h}, b_f][\alpha_v x_v, h^{t-1}, I]). \quad (14)$$

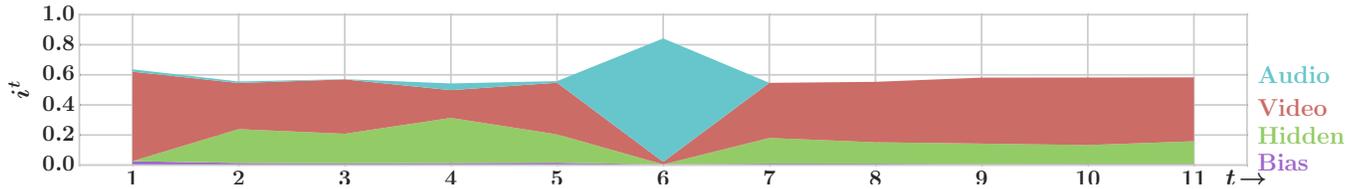
Define the forget gate activation residual as $\tilde{f}_m^t = |f^t - \hat{f}_m^t|$ (i.e., the difference in output resulting from removal of the m -th contributing variable). Then, let us define a 'pseudo correlation' of the m -th



(a) Agent anticipates salient audio features as it nears the ore, increasing audio attention until $t = 6$. Audio attention goes to 0 upon storage of ore indicator audio in the LSTM memory. Top and bottom rows show images and audio spectrogram sequences, respectively. Attention weights α^t plotted in center.



(b) Average forget gate activation throughout episode. Recall $f^t = 0$ corresponds to complete forgetting of the previous cell state element.



(c) Average input gate activation throughout episode. Recall $i^t = 1$ corresponds to complete throughput of the corresponding input element.

Figure 5: Interactions of crossmodal attention and LSTM memory. At $t = 6$, the attended audio input causes forget gate activation to drop, and the input gate activation to increase, indicating major overwriting of memory states. Relative contribution of audio to the LSTM forget and input activations drops to zero after the agent hears the necessary audio signal.

contributing variable with respect to the true activation,

$$\rho(\tilde{f}_m^t) = \frac{\tilde{f}_m^t}{\sum_m \tilde{f}_m^t}. \quad (15)$$

This provides an approximate quantification of the relative contribution of the m -th variable (audio input, video input, hidden unit, or bias) to the overall activation of the forget and input gates.

Armed with this toolset, we now analyze the interplay between attention and LSTM memory. First, given the sequence of audio-video inputs in Figure 5a, we plot overall activations of the forget and input LSTM gates (averaged across all cell state elements), in Figure 5b and Figure 5c, respectively. Critically, these plots also indicate the relative influence of each gate’s contributing variables to the overall activation, as measured by Equation (15).

Interestingly, prior to timestep $t = 6$, the contribution of audio to the forget gate and input gates is essentially zero, despite the positive attention on audio (in Figure 5a). At $t = 6$, the forget gate

activation drops suddenly, while the input gate experiences a sudden increase, indicating major overwriting of previous memory states with new information. The plots indicate that the attended audio input is the key contributing factor of both behaviors. In Figure 5a, after the agent hears the necessary audio signal, it moves attention entirely to video; the contribution of audio to the forget and input activations also drops to zero. Overall, this analysis indicates that the agent attends to audio in anticipation of an upcoming pertinent signal, but *chooses* not to embed it into memory until the appropriate moment. Attention filters irrelevant sensor modalities, given the contextual clues provided by exogeneous and endogeneous input features; it, therefore, enables the LSTM gates to focus on learning when and how to update the agent’s internal state.

5.3 The Arcade Learning Environment

Preliminary evaluation of crossmodal attention was also conducted in the Arcade Learning Environment (ALE) [7]. We modified ALE to support audio queries, as it previously did not have this feature.

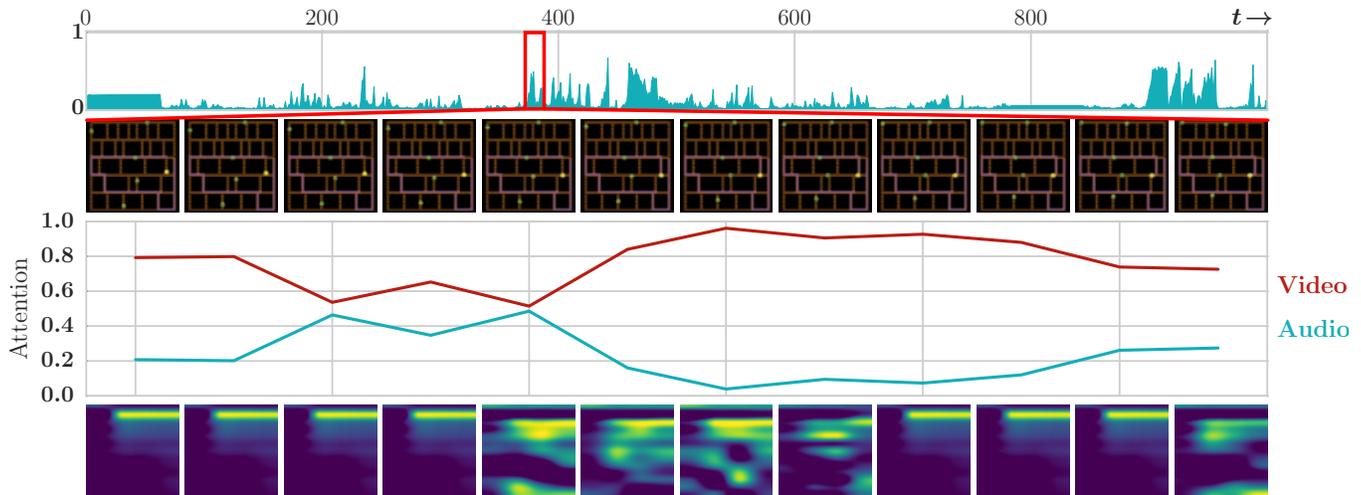


Figure 6: In Amidar, pathway vertices critical for avoiding enemies make an audible sound if not previously crossed. The agent anticipates and increases audio attention when near these vertices. Top row shows audio attention over 1000 frames, with audio/video/attention frames highlighted for zoomed-in region of interest.

Table 1: Preliminary results for learning in Atari 2600 games. The crossmodal attention learner, even *without* options, achieves high score for non-hierarchical methods. We emphasize these are not direct comparisons due to our method leveraging additional sensory inputs, but are meant to highlight the performance benefits of crossmodal learning.

Algorithm	Hierarchical?	Sensory Inputs	Amidar Score	H.E.R.O. Score
DQN [30]	✗	Video	740	19950
A3C [28]	✗	Video	284	28766
GA3C [4]	✗	Video	218	-
Ours (no options)	✗	Audio & Video	900	32985
A2OC [15]	✓	Video	880	20100
FeUdal [41]	✓	Video	>2500	~36000

Experiments were conducted in the Atari 2600 games H.E.R.O. and Amidar (Table 1). This line of investigation considers impacts of crossmodal attention on Atari agent behavior, even without use of multiple (hierarchical) options; these results use CASL with a single option, hence tagged “non-hierarchical” in the table. Amidar was one of the games in which Deep Q-Networks failed to exceed human-level performance [30]. The objective in Amidar is to collect rewards in a rectilinear maze while avoiding patrolling enemies. The agent is rewarded for painting segments of the maze, killing enemies at opportune moments, or collecting bonuses. Background audio plays throughout the game, and specific audio signals play when the agent crosses previously-unseen segment vertices. Figure 6 reveals that the agent anticipates and increases audio attention when near these critical vertices, which are especially difficult to observe when the agent sprite is overlapping them (e.g., zoom into video sequences of Figure 6).

Our crossmodal attentive agent achieves a mean score of 900 in Amidar, over 30 test runs, outperforming the other non-hierarchical methods. A similar result is achieved for the game H.E.R.O., where our agent beats other non-hierarchical agents. Note agent also beats the score of the hierarchical approach of [15]. We emphasize these are not direct comparisons due to our method leveraging additional sensory inputs, but are meant to highlight the performance benefits of crossmodal learning. We also note that the state-of-the-art hierarchical approach FeUdal [41] beats our agent’s score, and future investigation of the combination of audio-video attention with their approach may be of interest.

6 CONTRIBUTION

This work introduced the Crossmodal Attentive Skill Learner (CASL), integrated with the recently-introduced Asynchronous Advantage Option-Critic (A2OC) architecture [15] to enable hierarchical reinforcement learning across *multiple* sensory inputs. We provided concrete examples where CASL not only improves performance in a single task, but accelerates transfer to new tasks. We demonstrated the learned attention mechanism anticipates and identifies useful sensory features, while filtering irrelevant sensor modalities during execution. We modified the Arcade Learning Environment [7] to support audio queries, and evaluations of crossmodal learning were conducted in the Atari 2600 games H.E.R.O. and Amidar. Finally, building on the recent work of Babaeizadeh et al. [4], we open-source a fast hybrid CPU-GPU implementation of CASL. This investigation indicates crossmodal skill learning as a promising avenue for future works in HRL that target domains with high-dimensional, multimodal inputs.

ACKNOWLEDGEMENTS

This work was supported by Boeing Research & Technology, ONR MURI Grant N000141110688, and BRC Grant N000141712072.

REFERENCES

- [1] C. De Alvis, L. Ott, and F. Ramos. 2017. Online learning for scene segmentation with laser-constrained CRFs. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 4639–4643. <https://doi.org/10.1109/ICRA.2017.7989539>
- [2] Jacob Andreas, Dan Klein, and Sergey Levine. 2016. Modular multitask reinforcement learning with policy sketches. *arXiv preprint arXiv:1611.01796* (2016).
- [3] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. 2014. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755* (2014).
- [4] Mohammad Babaeizadeh, Iuri Frosio, Stephen Tyree, Jason Clemmons, and Jan Kautz. 2017. Reinforcement Learning through Asynchronous Advantage Actor-Critic on a GPU. In *ICLR*.
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [6] Matthew J. Beal, Hagai Attias, and Nebojsa Jovic. 2002. *Audio-Video Sensor Fusion with Probabilistic Graphical Models*. Springer Berlin Heidelberg, Berlin, Heidelberg, 736–750. https://doi.org/10.1007/3-540-47969-4_49
- [7] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research* 47 (06 2013), 253–279.
- [8] Samy Bengio. 2002. An Asynchronous Hidden Markov Model for Audio-visual Speech Recognition. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*. MIT Press, Cambridge, MA, USA, 1237–1244. <http://dl.acm.org/citation.cfm?id=2968618.2968772>
- [9] C. Cadena and J. Košecká. 2014. Semantic segmentation with heterogeneous sensor coverages. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 2639–2645. <https://doi.org/10.1109/ICRA.2014.6907237>
- [10] Ozan Caglayan, Loïc Barrault, and Fethi Bougares. 2016. Multimodal Attention for Neural Machine Translation. *arXiv preprint arXiv:1609.03976* (2016).
- [11] Marisa Carrasco. 2011. Visual attention: The past 25 years. *Vision research* 51, 13 (2011), 1484–1525.
- [12] Andrew D. Chambers, Sebastian Scherer, Luke Yoder, Sezal Jain, Stephen T. Nuske, and Sanjiv Singh. 2014. Robust Multi-Sensor Fusion for Micro Aerial Vehicle Navigation in GPS-Degraded/Denied Environments. In *In Proceedings of American Control Conference, Portland, OR*. Pittsburgh, PA.
- [13] Bruno Da Silva, George Konidaris, and Andrew Barto. 2012. Learning parameterized skills. *arXiv preprint arXiv:1206.6398* (2012).
- [14] Andreas Eitel, Jost Tobias Springenberg, Luciano Spinello, Martin Riedmiller, and Wolfram Burgard. 2015. Multimodal Deep Learning for Robust RGB-D Object Recognition. In *Proc. of the IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*. Hamburg, Germany. <http://ais.informatik.uni-freiburg.de/publications/papers/eitel15iros.pdf>
- [15] Jean Harb, Pierre-Luc Bacon, Martin Klissarov, and Doina Precup. 2017. When Waiting is not an Option: Learning Options with a Deliberation Cost. *arXiv preprint arXiv:1709.04571* (2017).
- [16] Matthew Hausknecht and Peter Stone. 2015. Deep recurrent Q-learning for partially observable MDPs. *arXiv preprint arXiv:1507.06527* (2015).
- [17] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural comp.* 9, 8 (1997), 1735–1780.
- [18] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1 (1998), 99–134.
- [19] Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. 2014. Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models. *CoRR abs/1411.2539* (2014). [arXiv:1411.2539](http://arxiv.org/abs/1411.2539) <http://arxiv.org/abs/1411.2539>
- [20] George Konidaris and Andrew G. Barto. 2009. Skill discovery in continuous reinforcement learning domains using skill chaining. In *Advances in neural information processing systems*. 1015–1023.
- [21] Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in Neural Information Processing Systems*. 3675–3683.
- [22] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 282–289. <http://dl.acm.org/citation.cfm?id=645530.655813>
- [23] Yuan Chang Leong, Angela Radulescu, Reka Daniel, Vivian DeWoskin, and Yael Niv. 2017. Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron* 93, 2 (2017), 451–463.
- [24] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025* (2015).
- [25] S. Lynen, M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart. 2013. A robust and modular multi-sensor fusion approach applied to MAV navigation. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 3923–3929. <https://doi.org/10.1109/IROS.2013.6696917>
- [26] Marlos C Machado, Marc G Bellemare, and Michael Bowling. 2017. A Laplacian Framework for Option Discovery in Reinforcement Learning. *arXiv preprint arXiv:1703.00956* (2017).
- [27] Nicholas J Mackintosh. 1975. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological review* 82, 4 (1975), 276.
- [28] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*. 1928–1937.
- [29] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. 2014. Recurrent models of visual attention. In *Advances in neural information processing systems*. 2204–2212.
- [30] 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. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [31] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. 689–696.
- [32] Yael Niv, Reka Daniel, Andra Geana, Samuel J Gershman, Yuan Chang Leong, Angela Radulescu, and Robert C Wilson. 2015. Reinforcement learning in multidimensional environments relies on attention mechanisms. *Journal of Neuroscience* 35, 21 (2015), 8145–8157.
- [33] S. Nobili, M. Camurri, V. Barasuol, M. Focchi, D. G. Caldwell, C. Semini, and M. Fallon. 2017. Heterogeneous Sensor Fusion for Accurate State Estimation of Dynamic Legged Robots. In *Proceedings of Robotics: Science and Systems*. Boston, USA. to appear.
- [34] John M Pearce and Geoffrey Hall. 1980. A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological review* 87, 6 (1980), 532.
- [35] John M Pearce and Nicholas J Mackintosh. 2010. Two theories of attention: A review and a possible integration. *Attention and associative learning: From brain to behaviour* (2010), 11–39.
- [36] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of Go without human knowledge. *Nature* 550, 7676 (2017), 354–359.
- [37] Ivan Sorokin, Alexey Seleznev, Mikhail Pavlov, Aleksandr Fedorov, and Anastasiia Ignateva. 2015. Deep attention recurrent Q-network. *arXiv preprint arXiv:1512.01693* (2015).
- [38] Nitish Srivastava and Ruslan Salakhutdinov. 2014. Multimodal Learning with Deep Boltzmann Machines. *Journal of Machine Learning Research* 15 (2014), 2949–2980. <http://jmlr.org/papers/v15/srivastava14b.html>
- [39] Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge.
- [40] Richard S Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112, 1-2 (1999), 181–211.
- [41] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. 2017. Feudal networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1703.01161* (2017).
- [42] Oriol Vinyals, Lukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. Grammar as a foreign language. In *Advances in Neural Information Processing Systems*. 2773–2781.
- [43] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J Smola, and Edward H Hovy. 2016. Hierarchical Attention Networks for Document Classification. In *HLT-NAACL*. 1480–1489.
- [44] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. 2015. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision* (2015), 1–15.