Developing hierarchical anticipations via neural network-based event segmentation

Christian Gumbsch¹² Maurits Adam³ Birgit Elsner³ Georg Martius² Martin V. Butz¹

Abstract

Humans can make predictions on various time scales and hierarchical levels. Thereby, the learning of event encodings seems to play a crucial role. In this work we model the development of hierarchical predictions via autonomously learned latent event codes. We present a hierarchical recurrent neural network architecture, whose inductive learning biases foster the development of sparsely changing latent state that compress sensorimotor sequences. A higher level network learns to predict the situations in which the latent states tend to change. Using a simulated robotic manipulator, we demonstrate that the system (i) learns latent states that accurately reflect the event structure of the data, (ii) develops meaningful temporal abstract predictions on the higher level, and (iii) generates goal-anticipatory behavior similar to gaze behavior found in eye-tracking studies with infants. The architecture offers a step towards the autonomous learning of compressed hierarchical encodings of gathered experiences and the exploitation of these encodings to generate adaptive behavior.

1. Introduction

Humans are able to generate hierarchical predictions on various time scales. How does this ability develop? Developmental psychology has shown that goal-predictive eye gaze develops during the first year of life: Infants start to look at the goal of an action before it is concluded. The generation of this goal-anticipatory gaze depends on various factors such as agent familiarity, action familiarity, the com-



Figure 1. Hypothesis for modeling infants' goal anticipations. (a): Infants first learn event representations (top). When they see familiar movements, these representations get re-activated (middle), enabling the anticipation of action consequences (bottom). (b): We implement this process in artificial systems, which learn event codes in a self-supervised manner. Figure is adapted from (Gumbsch et al., 2021a) with screenshots from (Adam & Elsner, 2020).

petence to execute the action themselves, agency cues, and action effects (Adam et al., 2016; Adam & Elsner, 2018; 2020; Adam et al., 2021; Elsner & Adam, 2021; Gredebäck & Melinder, 2010; Kanakogi & Itakura, 2011).

Theories on event-predictive cognition (Butz, 2016; Butz et al., 2021; Zacks et al., 2007) have led to the development of a theoretical framework, which attempts to explain these findings (Elsner & Adam, 2021) and which is illustrated in Fig. 1. The main assumption is that infants learn to compress their action experiences into action-event schemata. These schemata encode perceivable features of the entities, their movements and the goal. Once learned, the observation of a known event, e.g., a hand moving towards an object, results in the activation of the matching event schema. This activated schema leads to the generation of goal-predictive gaze shifts, striving to minimize uncertainty about the predicted end of the event, e.g., the hand touching the object.

Elsewhere (Gumbsch et al., 2021a), the development of

¹Neuro-Cognitive Modeling Group, University of Tübingen, Tübingen, Germany ²Autonomous Learning Group, Max-Planck Institute for Intelligent Systems, Tübingen, Germany ³Developmental Psychology, University of Potsdam, Potsdam, Germany. Correspondence to: Christian Gumbsch <christian.gumbsch@uni-tuebingen.de>.

Accepted at 2022 IEEE International Conference on Development and Learning (ICDL 2022)

such goal-predictive gaze shifts was modeled via performing active inference, i.e. predicted uncertainty minimizing behavior, on the level of events. First, event schemata are learned. Then, aiming at minimizing anticipated uncertainty within an event and between events, goal-predictive eye gaze develops mimicking infant behavior. However, the existing implementation has various restrictions: It learns event models based on supervised segmentations and requires fully-observable, linear simulations. This prevents the application of the model to more realistic scenarios, where motions are boundedly complex but non-linear, information can be hidden, and events are unknown.

Here we introduce a hierarchical recurrent neural network architecture, which can be trained end-to-end. The architecture learns to compress sensorimotor sequences into sparsely changing latent states. Based on these compressions, it learns temporal abstract predictions to anticipate the observations encountered upon latent state changes. Striving to minimize uncertainty, the system develops goal-predictive behavior similar to the way it develops in infants.

2. Event cognition

We humans are exposed to a continuous stream of perceptual information. However, a large body of psychological, neurological, and linguistic evidence suggests that we perceive, memorize, and predict information in terms of discrete events (Zacks et al., 2007; Zacks & Tversky, 2001; Radvansky & Zacks, 2014; Butz et al., 2021). According to Event Segmentation Theory (EST) (Zacks et al., 2007), and in line with various theories on predictive processing (Friston, 2010; Hohwy, 2013; Clark, 2015), humans continuously attempt to predict future perceptual input. EST argues that event encodings develop to better predict the next immediate perceptions as well as to anticipate event successions (Butz, 2016; Zacks et al., 2007; Radvansky & Zacks, 2014). In fact, the resulting hierarchical structures may constitute the foundations for abstract thought, counterfactual reasoning, and versatile planning (Butz et al., 2021; Eppe et al., 2022; Botvinick et al., 2009). Meanwhile, the Theory of Event Coding (TEC) (Hommel et al., 2001) argues that humans encode both actions and perceptions in event codes, emphasizing the close connection between action and perception.

Using these insights to develop event representations within an artificial system is not straight forward. Recurrent neural networks (RNNs) are commonly used in deep learning to process sensorimotor time series. However, they encode experience in a fundamentally different way: Namely, they update their latent representation of the ongoing activity in every time step. More abstract latent codes can emerge in hierarchical RNNs with levels that operate on different time scales (Tani, 2016). Previously, this was implemented by



Figure 2. GateLORD (Gumbsch et al., 2021b): (a) Internal components, (b) its update gate activation Λ and binarization via the step function Θ .

continuously updating higher level layers at a slower, predefined update rate, e.g. (Yamashita & Tani, 2008; Koutnik et al., 2014). In contrast to that, theories on event cognition (Zacks et al., 2007; Zacks & Tversky, 2001; Radvansky & Zacks, 2014; Butz et al., 2021) suggest that event encodings are mostly updated upon encountering critical situations, i.e. *event boundaries*, not based on fixed time scales. Recently, Gumbsch et al. (Gumbsch et al., 2021b), proposed Gated L0-Regularized Dynamics (GateL0RD), an RNN that only sparsely updates its latent state when the task demands it. In this work, we demonstrate that embedding GateL0RD in a suitable sensorimotor learning architecture fosters the self-supervised development of latent event codes.

3. Hierarchical event-inference architecture

Our architecture learns to process time series with sparsely changing latent states, which modulate predictive forwardinverse models. It furthermore learns to predict situations when the latent states tend to change. These predictive abilities enable goal-anticipatory behavior. We now detail our architecture and show how its inductive biases enable it to develop (i) event-compressing latent states, (ii) the prediction of situations for which the latent codes are expected to change, and (iii) goal-anticipatory behavior.

3.1. Sparsely changing memory

We describe the environment by a partially observable Markov decision process (POMDP) with state, action, and observation spaces S, A, O, and hidden transitions $S \times A \rightarrow$ S. Our system encodes the unobservable information at time t by means of a latent state $\mathbf{h}_t \in \mathbb{R}^{H,1}$ To learn the latent states \mathbf{h}_t we use GateLORD (Gumbsch et al., 2021b). Similar to other RNNs, GateLORD implements a general mapping $(\hat{\mathbf{y}}_t, \mathbf{h}_t) = f_{\theta}(\mathbf{x}_t, \mathbf{h}_{t-1})$, with inputs $\mathbf{x} \in \mathbb{R}^N$, outputs $\mathbf{y} \in \mathbb{R}^M$, and learnable parameters θ . Additionally, though, GateLORD's modularization and a regularizing loss

¹Notation: We denote vectors with bold lowercase letters (e.g. \boldsymbol{x}) and matrices as bold uppercase letters (e.g. \boldsymbol{W}). Vector dimensions are denoted by superscript and additional information by subscript (e.g. $\boldsymbol{x}_t = [\boldsymbol{x}_t^1, ..., \boldsymbol{x}_t^n]$).

term yield the inductive bias to develop piecewise constant latent states h_t .

GateLORD's internal structure is sketched-out in Fig. 2a. It uses three functions, or subnetworks, a *recommendation* function r, a gating function g, and an output function o, which together route sensorimotor prediction dynamics and maintain an inner latent state h_t , as follows:

$$\hat{\boldsymbol{h}}_t = r(\boldsymbol{x}_t, \boldsymbol{h}_{t-1}) \tag{1}$$

$$\mathbf{\Lambda}_t = \max(\mathbf{0}, \tanh(g(\boldsymbol{x}_t, \boldsymbol{h}_t))) \tag{2}$$

$$\boldsymbol{h}_t = \boldsymbol{\Lambda}_t \odot \hat{\boldsymbol{h}}_t + (1 - \boldsymbol{\Lambda}_t) \odot \boldsymbol{h}_{t-1}, \quad (3)$$

$$\hat{\boldsymbol{y}}_t = o(\boldsymbol{x}_t, \boldsymbol{h}_t) \tag{4}$$

The recommendation function r proposes a new latent state \hat{h}_t (Eq. 1). The gating function g computes an *update gate* $\Lambda_t \in [0, 1]$ (Eq. 2) with a rectified tanh activation function (Λ shown in Fig. 2b). If $\Lambda_t^i > 0$ the latent state is updated at dimension i (Eq. 3). The network output is then determined using the output function o (Eq. 4). To enforce sparse latent state updates, GateL0RD is trained using the loss:

$$\mathcal{L}_{\boldsymbol{\theta}} = \mathbb{E}\Big[\sum_{t} \mathcal{L}_{\text{task}}(\hat{\boldsymbol{y}}_{t}, \boldsymbol{y}_{t}) + \lambda \Theta(\boldsymbol{\Lambda}_{t})\Big], \quad (5)$$

where \mathcal{L}_{task} is the task-based loss, e.g. MSE and Θ is the Heaviside step function (shown in Fig. 2b). $\Theta(\Lambda_t)$ punishes latent state changes and thus fosters the tendency to develop stable, piecewise constant latent states h_t . The hyperparameter λ modifies the influence strength of the regularization term.

3.2. Forward-inverse model

We use GateLORD as a memory module, which is embedded into a predictive forward-inverse model structure as shown in Fig. 3. In every time step the model receives an observation o_t and an action a_t as its input. At time t the sensorimotor inputs (a_t, o_t) are fed into GateLORD and a network output \hat{y}_t and a new latent state h_t are computed as $(\hat{y}_t, h_t) = f_{\theta}(o_t, a_t, h_{t-1})$. The network outputs \hat{y}_t are processed by a forward model f_{FM} , which predicts the next sensory observation \hat{o}_{t+1} :

$$\hat{\boldsymbol{o}}_{t+1} = f_{\mathrm{FM}}(\hat{\boldsymbol{y}}_t). \tag{6}$$

Based on the updated latent state h_{t+1} and the next observation o_{t+1} , an inverse model f_{IM} predicts the next action \hat{a}_{t+1} :

$$\hat{\boldsymbol{a}}_{t+1} = f_{\mathrm{IM}}(\boldsymbol{o}_{t+1}, \boldsymbol{h}_{t+1}). \tag{7}$$

GateL0RD's initial latent state h_0 is determined by an initialization model f_{init} based on the first sensorimotor inputs:

$$\boldsymbol{h}_0 = f_{\text{init}}(\boldsymbol{a}_1, \boldsymbol{o}_1). \tag{8}$$

Our system implements a fully predictive sensorimotor model, that attempts to constantly predict its next sensory input as well as the next agentive action. However, in many situations the prediction accuracy varies drastically across states of the environment or the agent. For example, depending on the agent's gaze, different parts of the observation may be noisy or focused. Similarly, while in some state the model can be very certain about the agent's next actions, in other states a variety of potential actions could follow. Thus, we want the sensorimotor predictions of our network to be probabilistic, reflecting the heteroscedastic uncertainty in the world. Instead of directly predicting the next observation \hat{o}_{t+1} , the forward model f_{FM} predicts a probability distribution $p_{\theta}(\hat{o}_{t+1}|o_t, a_t, h_t)$, from which an observation prediction \hat{o}_{t+1} is sampled. We model $p_{\theta}(\hat{o}_{t+1}|o_t, a_t, h_t)$ as a normal distribution with

$$p_{\boldsymbol{\theta}}(\hat{\boldsymbol{o}}_{t+1}|\boldsymbol{o}_t, \boldsymbol{a}_t, \boldsymbol{h}_t) = \mathcal{N}(\boldsymbol{o}_t + \boldsymbol{\mu}_{t,\Delta o}, \boldsymbol{\Sigma}_{t,o}).$$
(9)

The mean $\mu_{t,\Delta o}$ and a diagonal covariance matrix $\Sigma_{t,o}$ are predicted by the forward model f_{FM} . Equivalently, for the distribution over actions $p_{\theta}(\hat{a}_{t+1}|o_{t+1}, h_t)$ the inverse model f_{IM} outputs a mean $\mu_{t,a}$ and diagonal covariance matrix $\Sigma_{t,a}$. We model the probability distribution as

$$p_{\boldsymbol{\theta}}(\hat{\boldsymbol{a}}_{t+1}|\boldsymbol{o}_{t+1},\boldsymbol{h}_t) = \mathcal{N}(\boldsymbol{\mu}_{t,a},\boldsymbol{\Sigma}_{t,a}). \tag{10}$$

We can train the overall architecture to minimize the negative log likelihood (NLL) loss. The overall loss of our forward-inverse model is defined as

$$\mathcal{L}_{\boldsymbol{\theta}} = \mathbb{E} \Big[\sum_{t} -\log \left(p_{\boldsymbol{\theta}}(\boldsymbol{o}_{t+1} | \boldsymbol{o}_{t}, \boldsymbol{a}_{t}, \boldsymbol{h}_{t}) \right) \\ \underbrace{-\log \left(p_{\boldsymbol{\theta}}(\boldsymbol{a}_{t+1} | \boldsymbol{o}_{t+1}, \boldsymbol{h}_{t}) \right)}_{\text{prediction loss}} + \underbrace{\lambda \Theta(\Lambda_{t})}_{\text{regularization}} \Big].$$
(11)

This corresponds to the original loss from GateLORD (Eq. 5) with the task-based loss $\mathcal{L}_{\text{task}}$ being the summed NLL of actions and observations (prediction loss). In practice, we use the recently proposed β -NLL loss² (Seitzer et al., 2022) ($\beta = 0.5$), to avoid instabilities during training and achieve high-quality predictions. We model all model components apart from GateLORD, i.e. f_{init} , f_{FM} , and f_{IM} , as feed-forward neural networks.

Can we expect that the resulting latent states h_t will resemble event codes? EST proposes that event models are only updated upon transient prediction errors, or surprise (Zacks et al., 2007). Equation 11 essentially phrases this idea in the form of a loss function: the system is trained to minimize

²The β -NLL loss scales the NLL gradients by a β -exponentiated per-sample variance for a hyperparameter $\beta \in [0, 1]$. $\beta = 0$ corresponds to the normal NLL loss. For $\beta = 1$ the mean μ is learned as when using MSE loss.



Figure 3. The overall system: The forward-inverse model f_{θ} and its latent states h_t are unrolled over time (top). The skip network f_{ϑ} (bottom) predicts input observation o_{τ} for which the latent states change. The attention focus input α_t is only used for modeling infants' goal anticipation (see Sec. 3.4).

prediction errors of future sensorimotor inputs. Thereby, the regularization loss punishes changes in its latent state. Thus, the system learns to change the latent state only when the transient prediction errors outweigh the regularization loss. In accordance to TEC, the latent states h_t will encode both sensory- and motor-predictive aspects, because they are used to predict both the next sensory states (Eq. 6) and the next actions (Eq. 7).

3.3. Event boundary predictions

A crucial feature of event cognition is the ability to generate temporally abstract predictions. We thus add a neural network f_{ϑ} with learnable parameters ϑ , which we call the *skip network*. The skip network is designed to predict the observation that will occur at the next event boundary, that is, at the end of the current event and at start of a next event. For example the sensation of ones hand touching an object could characterize the ending of a reach-event and the start of a grasp-event.

Our system should encode an ongoing event with the same latent state h_t . At an event boundary, the latent state changes, which is caused by the opening of an update gate $\Lambda_t^i > 0$ (Eq. 3). We can use this segmentation of the time series to generate the training data for the skip network in a self-supervised fashion. The process is illustrated in Fig. 3.

To generate the training data we feed a sensorimotor sequence $((o_1, a_1), (a_2, o_2), ..., (o_T, a_T))$ through our forward-inverse model f_{θ} to receive a sequence of latent states and update gate openings $((h_1, \Lambda_1), (h_2, \Lambda_2), ..., (h_T, \Lambda_T))$. We can then determine all event boundaries as

$$\mathcal{T} = \{ t' | \exists_i \Lambda_i^t > 0, 0 < t' \le T \}.$$
(12)

Thus, \mathcal{T} is the set of time steps for which at least one gate opens. For every time step t, we can then determine the next event boundary as

$$\mathcal{T}(t) = \min\left(\{t' \in \mathcal{T} | t' > t\}\right). \tag{13}$$

We train the skip network to predict $o_{\mathcal{T}(t)}$ from the observation o_t and latent state h_t . In other words, the skip network is trained to predict the final observation of an event from an arbitrary starting point within this event. We treat the last time step T of a sequence also as an event boundary.

To enable the skip network to predict uncertainties we model its output as a distribution of observations $p_{\vartheta}(\hat{o}_{\mathcal{T}(t)}|o_t, h_t)$ and use a multivariate Gaussian distribution parametrization:

$$p_{\boldsymbol{\vartheta}}(\hat{\boldsymbol{o}}_{\mathcal{T}(t)}|\boldsymbol{o}_t, \boldsymbol{h}_t) = \mathcal{N}(\hat{\boldsymbol{o}}_{\mathcal{T}(t)}|\boldsymbol{\mu}_{t, o_{\mathcal{T}(t)}}, \boldsymbol{\Sigma}_{t, o_{\mathcal{T}(t)}}) \quad (14)$$

with mean $\mu_{t,o_{\mathcal{T}(t)}}$ and diagonal covariance matrix $\Sigma_{t,o_{\mathcal{T}(t)}}$ predicted by the skip network. The skip network is trained to minimize the loss \mathcal{L}_{ϑ} :

$$\mathcal{L}_{\boldsymbol{\vartheta}} = \mathbb{E}\Big[-\log\big(p_{\boldsymbol{\vartheta}}(\boldsymbol{o}_{\mathcal{T}(t)}|\boldsymbol{o}_t,\boldsymbol{h}_t)\big)\Big], \qquad (15)$$

for its learnable parameters ϑ . As before, we use the β -NLL loss (Seitzer et al., 2022) ($\beta = 0.5$). We model the skip network as an MLP.

3.4. Modeling infants' goal anticipation

As previously done in (Gumbsch et al., 2021a), we use the structure of events and event boundaries to model experimental findings on infants' goal-predictive gaze. For modeling infants' gaze (experiments in Sec. 5.4) we modify the inputs of the two networks. Here, the forward-inverse model f_{θ} and the skip network f_{ϑ} receive an attention focus α_t as an additional input. We implement the focus as a noise mask on the input observations. When the system attends to an entity e, e.g., the agent's hand, all dimensions of the input observation o_t^e concerning this entity receive no sensory noise, e.g. the position of the hand. All other dimensions of o_t , for example the position of an object on a table, are masked by sensory noise. In this way we model simplified gaze behavior: The system can focus on entities, akin to looking at them, to gain clear sensory observations, whereas unattended entities receive sensory noise.

Gumbsch et al. (Gumbsch et al., 2021a) proposed that infants' goal-predictive gaze emerges from selecting gaze behavior that minimizes uncertainty about the future on multiple time scales. According to their model, infants direct their gaze to minimize uncertainty within the ongoing event to better predict the next sensory observations (intraevent). Additionally, infants attempt to minimize uncertainty about future event boundaries to better predict when, where, and which future events will follow (inter-event). In our system we can express this idea as

$$\boldsymbol{\alpha}_{t} = \underset{\boldsymbol{\alpha}}{\operatorname{arg\,min}} \left[\underbrace{\underbrace{U\left(p_{\boldsymbol{\theta}}(\hat{\boldsymbol{o}}_{t+1} | \boldsymbol{o}_{t}, \boldsymbol{a}_{t}, \boldsymbol{h}_{t}, \boldsymbol{\alpha})\right)}_{\text{intra-event uncertainty}} + \underbrace{U\left(p_{\boldsymbol{\vartheta}}(\hat{\boldsymbol{o}}_{\mathcal{T}(t)} | \boldsymbol{o}_{t}, \boldsymbol{h}_{t}, \boldsymbol{\alpha})\right)}_{\text{inter-event uncertainty}} \right],$$
(16)

where U is an uncertainty measurement and p_{θ} and p_{ϑ} are the Gaussian distributions over observations generated by the forward-inverse-model f_{θ} and skip network f_{ϑ} , respectively. We use the following simple uncertainty estimation U

$$U(p) = \sum_{i \in I} \Sigma_{t,p}^{i,i},\tag{17}$$

which is the sum of variances³, with $\Sigma_{t,p}$ the predicted covariance matrix of distribution p, i.e. either p_{θ} or p_{ϑ} , and I a set of relevant dimensions. I allows us to specify for which parts of the observation the system should aim to minimize uncertainty, e.g., focusing on future hand positions only.

4. Scenario and data

We apply our system in the **Fetch Pick & Place** environment (OpenAI Gym v1), a benchmark reinforcement learning problem in which a robotic manipulator should move a randomly placed box on a table to a random goal position. The "hand" of the robot is a gripper with two fingers. The action $a_t \in \mathbb{R}^4$ position-controls the movement of the hand and the opening of the gripper. In our experiments the observation $o_t \in \mathbb{R}^{11}$ is composed of the three-dimensional position of three entities, i.e., the hand, the object, and the goal, and the fingers' opening. We modify the simulation such that the height of the table can vary across simulations.

The system is trained and evaluated on two datasets \mathcal{D} of 9.2k sensorimotor sequences $((o_1, a_1), ...(o_T, a_T)) \in \mathcal{D}$ of length T = 25. The dataset $\mathcal{D}_{\text{script}}$ contains three types of scripted movements. In **reach-grasp-transport** sequences the hand moves to the object until it is located between the fingers. The hand then closes its fingers and, once the object is fully grasped, lifts the object to the goal position where it is held until the sequence is over. In **point-ing sequences** the robot "points" to the goal by moving its hand to the goal position. In stretching sequences the



Figure 4. Results on \mathcal{D}_{script} : Test prediction error of observations (a) and actions (b) and mean number of gate openings (c), i.e., latent updates. Shaded areas show \pm one standard deviation.

robot stretches its arm by repeatedly performing the same randomly-generated motor command. In all sequences goaldirected motions follow a fixed direction and decrease their velocity when approaching the target to avoid overshooting. All actions were perturbed by normal distributed motor noise ($\sigma = 0.05$).

The \mathcal{D}_{APEX} dataset contains sequences that were generated by the policy-guided model-predictive control method APEX (Pinneri et al., 2021) trained to move the object to the goal position. APEX finds various ways to interact with the object, including lifting, pushing, or flicking it. In \mathcal{D}_{APEX} table height is fixed.

When modeling infants' gaze behavior (experiment Sec. 5.4), we provide an additional attention focus α_t . This corresponds to attending to one of the three entities, i.e., hand, object, or goal. When attending to one entity, sensory information about the other entities' positions are masked by normally distributed noise ($\sigma = 0.05$). During training, the attentional focus is randomly shifted 5 times per sequence.

The forward-inverse model and the skip network were trained independently using Adam ($lr_{\theta} = .0005$, $lr_{\vartheta} = .0001$). Each experiment was run with 10 different random seeds.

5. Results

We now first evaluate the semantics of the learned event encodings, then the quality of the skip predictions, the robustness of the learning procedures on more challenging data, and finally the emergence of goal-anticipatory behavior.

5.1. Learned event segmentation

To analyze prediction accuracy and the learned latent state of the forward-inverse model, we first trained our model on the dataset \mathcal{D}_{script} of scripted movements. We compare our system (with GateLORD) with different regularization strengths λ to an ablated version without sparsity regular-

³In (Gumbsch et al., 2021a) entropy was used for the uncertainty U. Computing the entropy of Gaussian distributions involves multiplying variances. When some variances are very close to zero, we found the sum of variances to be more stable.

Developing hierarchical anticipations via neural network-based event segmentation



Figure 5. Exemplary sequences of $\mathcal{D}_{\text{script}}$ for reach, grasp, and transport of the object (a), pointing to the goal (b), and stretching the arm (c). The topmost row shows hand and object positions over time. The middle row shows actions $(a_t^{\{1,2,3\}} \text{ control hand movements}, a_t^4 \text{ controls gripper closing})$. The bottom row shows the latent states relative to initialization, i.e. $h_t - h_0$.

ization that uses a GRU (Chung et al., 2014) as its internal RNN cell.

Fig. 4a and Fig. 4b show the prediction errors for the next observations and actions, respectively. The networks learn to improve their predictions over training. For a regularization of $\lambda \leq 1$ the regularized versions reach the same prediction accuracy as the GRU ablation. Thus, the sparsity regularization via GateLORD, does not necessarily degrade performance. Figure 4c shows the mean number of gate openings, i.e., latent state changes. With a higher λ , fewer latent states are updated. Our system with $\lambda = 1$ learns to only adapt its latent state on average a handful of times during a sequence⁴, with the same prediction abilities as less regularized variants. Thus, for further analysis we focus on our system with $\lambda = 1$.

Figure 5 shows exemplary sequences for the three types of scripted movements. The bottom row illustrates the latent states for one exemplary model. For the reach-grasp-transport sequence (Fig. 5a), the model changes its latent state when the hand has reached the object, when the object is fully grasped, and when hand and object arrive at the goal. Thus, the model seems to encode *reaching*, *grasping*, *transporting*, and *holding* the object using four different latent states. For the pointing sequence (Fig. 5b), the model changes its latent state once the hand has reached the goal. Interestingly, the same dimension is updated when reaching the object (Fig. 5a) and the goal (Fig. 5b), indicating the similarity between the two events. For the stretching sequence (Fig. 5c), no clear latent update is visible. This implies that stretching is encoded as one event.

5.2. Skip predictions

To systematically analyze the skip predictions, we fed the scripted sequences into the forward-inverse model up to a



Figure 6. Skip predictions for $\mathcal{D}_{\text{script}}$. Euclidean distance between skip-predicted hand position in $\hat{o}_{\mathcal{T}(2)}$ to the position of all three entities in o_2 for a skip at time t = 2. Shown for reach-grasp-transport (a), pointing (b), and stretching (c) sequences. Shaded areas denote standard deviation.

fixed point in time at t = 2. We then used (o_2, h_2) as inputs to the skip network to produce a skip prediction $\hat{o}_{\mathcal{T}(2)}$. Thus, we essentially queried the skip network at time t = 2 which observation it predicts at the next event boundary.

We compare the predicted position of the hand from $\hat{\sigma}_{\mathcal{T}(2)}$ to the current position of the entities in σ_2 . Figure 6 shows the Euclidean distances between the hand position of the skip prediction and the current positions of all entities. For reaching sequences (Fig. 6a) the distance between predicted hand position and current object position decreases, while the other distances stay the same or increase. Thus, the skip network learns that the hand tends to move to the object, where the next gate opening, or event boundary, will occur. Similarly, for pointing sequences (Fig. 6b) the system learns that the hand will move to the goal. For stretching sequences (Fig. 6c), on the other entities show high variances without any regularities.

Figure 7 shows the hand position predictions of all fullytrained skip networks for three exemplary sequences. One blue circle marks the predicted position of the hand for one trained skip network. We show the predictions of all trained networks (10 random seeds) overlayed. For reach-grasp-

 $^{{}^{4}\}mathbb{E}_{t}[\Theta(\Lambda_{t})] = 10^{-2}$ corresponds to changing one dimension of h_{t} on average 4 times during a sequence (for T = 25 and 16-dimensional h_{t}).



Figure 7. Exemplary skip predictions of the hand. Skips at t = 2 shown for reaching (a), pointing (b), and stretching (c) sequences of $\mathcal{D}_{\text{script}}$. One blue circle shows one skip-predicted hand position. Predictions of all 10 random seeds are shown together. The goal position is depicted as a red sphere.

transport sequences (Fig. 7a), the skip network predicts that the hand will be close to the object at the next event boundary. For pointing sequences (Fig. 7b), the skip network predicts that the hand will be at the goal location (red sphere). For both sequences the predicted positions strongly overlap, implying a consistent segmentation and hierarchical prediction across different random initializations. For stretching sequences (Fig. 7c), the skip network predicts that the hand will be close to positions when the arm is fully stretched-out, or somewhere in-between the current position and the final position.

5.3. Training on more diverse sequences

To investigate learning robustness, we trained our system on sequences of the \mathcal{D}_{APEX} dataset, which contains various object interaction sequences. Again, we compare our system (with GateL0RD, $\lambda \in \{1, 5, 10\}$) to a GRU ablation.

Fig. 8a and Fig. 8b show the prediction errors for the observations and actions, respectively. While our system outperforms the unregularized GRU ablation for predicting observations, the ablation learns to better predict actions. For further evaluations we focus on our system with $\lambda = 5$.

Figure 8c shows an exemplary reach-grasp-and-transport sequence from \mathcal{D}_{APEX} . The latent states, depicted in the bottom row, changes strongly at two points in the sequence. First, when the object is grasped, and later, when the object was successfully transported to the goal. Thus, the latent states seem to encode a similar event structure as for the scripted data (cf. Fig. 5a), even for the much noisier actions of \mathcal{D}_{APEX} .

To evaluate the system's ability to make temporal abstract predictions, we input (o_2, h_2) to the skip network and analyze the resulting predictions $\hat{o}_{\mathcal{T}(2)}$. Figure 8d shows the Euclidean distance between the skip-predicted hand position in $\hat{o}_{\mathcal{T}(2)}$ and the current positions of all entities at the time t = 2. Over the course of training, the distance between predicted hand position and current object position decreases. Thus with increasing experience the skip network learns that the hand will move to the object. Additionally, the increasing distance between current and predicted hand position suggests that with more training experience the skip predictions reach farther into the future. Figure 8e shows the skip-predicted hand positions for one exemplary sequence. Across all random seeds, the system reliably learns to predict that the hand will be close to the object position at the next event boundary.

5.4. Modeling infants' goal anticipations

Lastly, we model experimental findings of goal anticipations in infants. We train our system on \mathcal{D}_{script} with an additional attentional focus α_t . We test the system similarly to eyetracking studies on infant goal anticipation, e.g. (Adam & Elsner, 2020). We show the system observation sequences⁵ and let it choose it's focus α_t to minimize uncertainty about future positions of the hand (specified as indices *I* in Eq. 17). Similar to how the developmental studies track the first gaze to a goal area, we track the time step t_e , when the system first attended to entity *e* (hand, object, or goal)⁶. Akin to infants' goal-predictive gaze shifts, we want to see if our system develops a *goal-predictive attention shift*: Does it attend to the goal of an event, e.g., the to-be-grasped object, before the goal is reached?

We compare three different versions of attention selection for reaching events: Minimizing uncertainty within the current event (intra-event uncertainty in Eq. 16), minimizing uncertainty about the next event boundary (inter-event uncertainty in Eq. 16), or minimizing both uncertainties (full Eq. 16). When minimizing intra-event uncertainty for reaching (Fig. 9a), the system first attended to the hand (blue) long before hand-object-contact (dashed black line) with little variance across simulations. Naturally, attending to the hand helps to predict immediate hand movements during reaching. The system on average attended to other entities only after hand-object-contact with higher variance across simulations. When minimizing inter-event uncertainty (Fig. 9b), the system on average first attended to the object (green) before it was reached by the hand (dashed black line). Apparently, the system has learned that attending to the object helps to minimize uncertainty about the end of a reaching event.

When putting both uncertainties together, the system exhibits goal-predictive attention shifts: For reach-grasptransport sequences (Fig. 9c), the system on average first

⁵Action inputs are generated by the inverse model f_{IM} . Only the first action a_1 is provided to initialize the latent state h_0 and start the process.

⁶Attention is selected as in Eq. 16. For each entity e, we determine the first time step t_e for which $\alpha_{t_e}^e = 1$, with α_t the one-hot encoded attentional focus. No focus on an entity is treated as $t_e = 25$, i.e., maximum length T.

Developing hierarchical anticipations via neural network-based event segmentation



Figure 8. Results on \mathcal{D}_{APEX} : Test prediction error of observations (a) and actions (b). Exemplary reach-grasp-transport sequences (c) showing hand and object position (top), actions (middle) and latent states relative to initialization, i.e. $h_t - h_0$ (bottom). Euclidean distance between skip-predicted hand position to the current position of all three entities for a skip at time t = 2 (d). Exemplary skip predictions (e) at t = 2 with circles depicting predicted hand positions overlayed for 10 random seeds.

attended to the hand (blue), followed by attending to the object (green). The focus on the object happened on average before hand-object-contact (dashed black line). Similarly, for pointing sequences (Fig. 9d) the system on average first attended to the hand (blue) and then to the goal (red) before the hand arrived there (dashed line). Thus, in both cases, the system shifted its attention to the end of the event before the event was over. Crucially, this behavior developed with training experience.

This behavior mimics gaze-behavior found in infants: Figure 9e shows the mean gaze arrival time for infants watching movies of a hand grasping and lifting a toy (Adam & Elsner, 2020) depending on age. Mean gaze arrival times were calculated by subtracting the time when the hand entered the target object area of interest (AOI), from the time of the first fixation to that AOI. Negative gaze arrival times thus indicate goal-predictive gaze. Only trials in which the hand was first fixated (for a minimum of 200 ms) were considered. As shown in Fig. 9e, 6-month-olds tend to follow the hand with their gaze. Older infants (7 and 11 months) shift their gaze from the hand to the object before it is reached. These results are qualitatively matched by our systems' behavior (cf. Fig. 9c, green line).

6. Discussion

We propose a hierarchical deep learning system that learns to encode its sensorimotor experience in event-compressing, latent states, which are used for probabilistic forward- and inverse predictions. We show that the learned codes can uncover suitable sensorimotor sequence segmentations. Additionally, the system can learn meaningful temporal abstract predictions of latent state changes, that is, event boundaries.

Previously, (Gumbsch et al., 2021a) used an explicit structure of events and their boundaries to model the development of goal-predictive gaze in infants. However, in their setup events unfolded linearly and the segmentation was supervised. Our system learns event segmentation in a completely self-supervised fashion. This allowed us to apply uncertainty minimizing gaze inference strategies without providing any event information explicitly.

We tested our model similarly to how goal anticipation was studied in infants (Kanakogi & Itakura, 2011; Adam & El-



Figure 9. Goal anticipation of our system and infants. The y-axis shows the time t_e of first attending to entity e relative to the time t_{EB} of the first event boundary of a sequence, i.e. $t_e - t_{EB}$. Thus, negative values for an entity indicate that the system attended to this entity before the first event was over. The first event boundaries are marked as dashed black lines (reaching: hand at object; pointing: hand at goal). For reaching events we evaluate attention selection for minimizing uncertainty within the event (a), about the next event boundary (b), or both (c). For pointing sequences we only consider minimizing both uncertainties (d). In (e) the mean gaze arrival time for infants watching grasping sequences are shown (from (Adam & Elsner, 2020)). Gaze arrival times below the dashed line (hand object contact) mark anticipatory gaze. Shaded areas and error bars show standard error.

sner, 2020). Our system automatically learned to shift its focus from the moving entity (hand) to the next goal as it became more and more mature. In this way it replicates the development of goal-predictive gaze shifts in infants. Hypotheses that event familiarity and motor experience enable goal-predictive gaze shifts (Kanakogi & Itakura, 2011; Elsner & Adam, 2021) align with our model.

Besides general neural network related hyperparameters, our model only has one hyperparameter that needs to be set: the regularization strength λ , controlling the trade-off between prediction accuracy and sparsity in latent state changes. The parameter is important for developing robust event segmentation and thus the emergence of meaningful goal-predictive attention shifts. While GateLORD works rather robustly under a range of values for λ , in the future it may be interesting to investigate a battery of gates with their own λ values to foster the development of hierarchical event segmentation (Gumbsch et al., 2017).

Seeing that our model learns event segmentation best when there is clear event-separating sensorimotor information available, our model predicts that salient segmentation cues, e.g., a sound at the event boundary, should ease event segmentation and, thus, boost goal anticipations. Future infant studies could test whether goal anticipation for newly learned event sequences occurs at an earlier age if during learning of sequences salient segmentation cues are provided.

We believe our system is not limited to modeling goalpredictive gaze. Event-based anticipations are ubiquitous (Butz et al., 2021). Moreover, we expect to show that our system is able to plan hierarchically, to exhibit complex one-shot learning, and to solve challenging reasoning tasks (Sutton et al., 1999; Eppe et al., 2022).

Acknowledgment

We acknowledge the support from the Machine Learning Cluster of Excellence, EXC number 2064/1–project number 390727645, the Tübingen AI Center (FKZ: 01IS18039B), the Volkswagen Stiftung (No. 98 571), the DFG Priority-Program "The Active Self" SPP 2134 (project BU 1335/11-1, EL 253/8-1) and the IMPRS-IS.

References

- Adam, M. and Elsner, B. Action effects foster 11-montholds' prediction of action goals for a non-human agent. *Infant Behavior and Development*, 53:49–55, 2018.
- Adam, M. and Elsner, B. The impact of salient action effects on 6-, 7-, and 11-month-olds' goal-predictive gaze shifts for a human grasping action. *Plos one*, 15(10):e0240165, 2020.

- Adam, M., Reitenbach, I., Papenmeier, F., Gredebäck, G., Elsner, C., and Elsner, B. Goal saliency boosts infants' action prediction for human manual actions, but not for mechanical claws. *Infant Behavior and Development*, 44: 29–37, 2016. ISSN 0163-6383. doi: https://doi.org/10. 1016/j.infbeh.2016.05.001.
- Adam, M., Gumbsch, C., Butz, M. V., and Elsner, B. The impact of action effects on infants' predictive gaze shifts for a non-human grasping action at 7, 11, and 18 months. *Frontiers in Psychology*, pp. 3373, 2021.
- Botvinick, M. M., Niv, Y., and Barto, A. G. Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3): 262–280, 2009.
- Butz, M. V. Towards a unified sub-symbolic computational theory of cognition. *Frontiers in Psychology*, 7(925), 2016. ISSN 1664-1078. doi: 10.3389/fpsyg.2016.00925.
- Butz, M. V., Achimova, A., Bilkey, D., and Knott, A. Event-predictive cognition: A root for conceptual human thought. *Topics in Cognitive Science*, 13(1):10–24, 2021. doi: https://doi.org/10.1111/tops.12522.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555, 2014.
- Clark, A. Surfing uncertainty: Prediction, action, and the embodied mind. Oxford University Press, 2015.
- Elsner, B. and Adam, M. Infants' goal prediction for simple action events: The role of experience and agency cues. *Topics in Cognitive Science*, 13(1):45–62, 2021.
- Eppe, M., Gumbsch, C., Kerzel, M., Nguyen, P. D., Butz, M. V., and Wermter, S. Intelligent problem-solving as integrated hierarchical reinforcement learning. *Nature Machine Intelligence*, pp. 1–10, 2022.
- Friston, K. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010. doi: 10.1038/nrn2787.
- Gredebäck, G. and Melinder, A. Infants' understanding of everyday social interactions: A dual process account. *Cognition*, 114(2):197–206, 2010.
- Gumbsch, C., Otte, S., and Butz, M. V. A computational model for the dynamical learning of event taxonomies. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, pp. 452–457, 2017.
- Gumbsch, C., Adam, M., Elsner, B., and Butz, M. V. Emergent goal-anticipatory gaze in infants via event-predictive learning and inference. *Cognitive Science*, 45(8):e13016, 2021a.

- Gumbsch, C., Butz, M. V., and Martius, G. Sparsely changing latent states for prediction and planning in partially observable domains. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Hohwy, J. *The predictive mind*. Oxford University Press, 2013.
- Hommel, B., Müsseler, J., Aschersleben, G., and Prinz,W. The theory of event coding (tec): A framework for perception and action planning. *Behavioral and brain sciences*, 24(5):849–878, 2001.
- Kanakogi, Y. and Itakura, S. Developmental correspondence between action prediction and motor ability in early infancy. *Nature communications*, 2(1):1–6, 2011.
- Koutnik, J., Greff, K., Gomez, F., and Schmidhuber, J. A clockwork RNN. In *International Conference on Machine Learning (ICML)*, 2014.
- Pinneri, C., Sawant, S., Blaes, S., and Martius, G. Extracting strong policies for robotics tasks from zero-order trajectory optimizers. In *International Conference on Learning Representations*, 2021.
- Radvansky, G. A. and Zacks, J. M. Event Cognition. Oxford University Press, 2014. doi: 10.1093/acprof: oso/9780199898138.001.0001.
- Seitzer, M., Tavakoli, A., Antic, D., and Martius, G. On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. In *International Conference on Learning Representations (ICLR)*, 2022.
- Sutton, R. S., Precup, D., and Singh, S. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1): 181–211, 1999. ISSN 0004-3702. doi: https://doi.org/10. 1016/S0004-3702(99)00052-1.
- Tani, J. Exploring robotic minds: actions, symbols, and consciousness as self-organizing dynamic phenomena. Oxford University Press, 2016.
- Yamashita, Y. and Tani, J. Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS computational biology*, 4(11):e1000220, 2008.
- Zacks, J. M. and Tversky, B. Event structure in perception and conception. *Psychological Bulletin*, 127(1):3–21, 2001. doi: 10.1037/0033-2909.127.1.3.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., and Reynolds, J. R. Event perception: a mind-brain perspective. *Psychological Bulletin*, 133(2):273–293, 2007. doi: 10.1037/0033-2909.133.2.273.

A. Implementation details

We use GateLORD (Gumbsch et al., 2021b) with a 16dimensional latent state h_t . The subnetworks g and r are implemented by multilayer perceptrons (MLPs) with three layers (64 \rightarrow 32 \rightarrow 16 feature neurons, tanh hidden activation), whereas o is implemented as two single-layered neural networks (16 feature neurons per layer), whose outputs are multiplied element-wise. When using the GRU we found that a 16-dimensional latent state h_t performs much worse on \mathcal{D}_{APEX} , which is why we used the GRU with 32 latent dimensions. The other components of the forward-inverse model, i.e. f_{init} , f_{FM} , and f_{IM} , are also implemented as feed-forward neural networks. For f_{init} and $f_{\rm FM}$ we use MLPs with three layers (64 \rightarrow 32 \rightarrow 16 feature neurons, \tanh hidden activation). The inverse model $f_{\rm IM}$ has an additional mutliplicative layer, akin to the GateL0RD output function o (16 features per layer), followed by a threelayered MLP (64 \rightarrow 32 \rightarrow 16 feature neurons , tanh hidden activation). We add this extra layer, such that $f_{\rm IM}$ has the same computational power to compute its outputs based on o_t and h_t as $f_{\rm FM}$. We model the skip network as a deep MLP with five layers $(512 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 32$ feature neurons, tanh hidden activation). Networks predicting Gaussian distributions (i.e. $f_{\rm FM}$, $f_{\rm IM}$ and f_{ϑ}) use separate read-out layers for predicting the mean and predicting the covariance. The read-out layers for mean predictions have a linear activation function, whereas the read-out layers for the covariance matrix predictions use the ELU activation function shifted by 1 to avoid covariances < 0.

We trained and tested on datasets of 9.2k sequences using a batch size of 192. The forward-inverse model and the skip network were trained independently using Adam ($lr_{\theta} = 0.0005$, $lr_{\vartheta} = .0001$, $\epsilon = 10^{-4}$) and gradient norm clipping (max= 0.1). The code to run our experiments can be found at https://github.com/ CognitiveModeling/HierarchicalGateLORD.

B. Skip predictions of goal and object positions

In Sec. 5.2 we analyzed the skip predictions and showed that our system learns to make meaningful temporal abstract predictions of future hand positions. How are the skip predictions for the positions of the other two entities?

We investigated the skip-predictions for object and goal positions as before by feeding the scripted sequence of $\mathcal{D}_{\text{script}}$ into the forward-inverse model to generate inputs for the skip network. Figure 10 (a)-(c) shows the skip predictions for three exemplary sequences based on the inputs (o_2, h_2) . For all sequences the skip network predicts that the object position (green) and the goal position (red) will not change (cf. Fig. 7 for hand predictions and the ground truth goal positions). This is reasonable for the given events, since the goal position never changes and for reaching, pointing,



Figure 10. Exemplary skip predictions for object and goal positions. Skips at t = 2 shown for reaching (a), pointing (b), and stretching (c) sequences. Skips at t = 12 shown for a transport event (d). One circle shows one skip-predicted position of the object (green) or the goal (red). Predictions of all 10 random seeds are shown overlayed.

and stretching events the positions of the object does not change. However, for reach-grasp-transport sequences, the object position changes once the object is grasped (typically at $t \leq 12$). As shown exemplary in Fig. 10 (d), if we used (o_{12}, h_{12}) as inputs to the skip network for a reach-grasp-transport sequence, the skip network indeed predicts that the object will be close to the goal at the next event boundary. Thus, the skip network can also make temporal abstract predictions of the object and the goal position.