

Neural Reasoning about Agents’ Goals, Preferences, and Actions

Matteo Bortoletto, Lei Shi, Andreas Bulling

University of Stuttgart, Germany
 {matteo.bortoletto, lei.shi, andreas.bulling}@vis.uni-stuttgart.de

Abstract

We propose the *Intuitive Reasoning Network* (IRENE) – a novel neural model for intuitive psychological reasoning about agents’ goals, preferences, and actions that can generalise previous experiences to new situations. IRENE combines a graph neural network for learning agent and world state representations with a transformer to encode the task context. When evaluated on the challenging Baby Intuitions Benchmark, IRENE achieves new state-of-the-art performance on three out of its five tasks – with up to 48.9% improvement. In contrast to existing methods, IRENE is able to bind preferences to specific agents, to better distinguish between rational and irrational agents, and to better understand the role of blocking obstacles. We also investigate, for the first time, the influence of the training tasks on test performance. Our analyses demonstrate the effectiveness of IRENE in combining prior knowledge gained during training for unseen evaluation tasks.

Introduction

Common-sense reasoning refers to a broad class of abilities that enable agents to understand basic facts about events, objects, beliefs, or desires (McCarthy 1989). Studies in developmental cognitive science have demonstrated that even young infants have these abilities (Needham and Baillargeon 1993; Gergely et al. 1995; Aguiar and Baillargeon 1999) while their absence can be linked to developmental disorders, such as autism (Baron-Cohen 1997). Given its fundamental importance for human social cognition and behaviour, it is imperative that artificial intelligent (AI) agents possess similar capabilities to effectively be understood by humans and in turn, understand them. This is crucial for a wide variety of applications, such as robotics (Mota and Sridharan 2019) or human-machine collaboration (Conti, Varde, and Wang 2022). Previous research in machine common-sense reasoning has mainly focused on evaluating language processing (Bhagavatula et al. 2020; Huang et al. 2019; Zellers et al. 2019; Bisk et al. 2020; Sap et al. 2019; Sakaguchi et al. 2021) or visual scene understanding (Yi et al. 2020; Smith et al. 2019; Ates et al. 2022) in a task-specific manner. More recently, several new benchmarks have been introduced that allow to assess the

more general ability of AI systems to reason about unexpected events or situations (Riochet et al. 2020; Gandhi et al. 2021; Dasgupta et al. 2021; Shu et al. 2021; Piloto et al. 2022; Weihs et al. 2022). Most of them focus on *intuitive physics* (Yi et al. 2020; Smith et al. 2019; Ates et al. 2022; Riochet et al. 2020; Dasgupta et al. 2021; Piloto et al. 2022; Weihs et al. 2022; Piloto et al. 2018) in which computational models have to reason about properties and interactions of physical macroscopic objects. In contrast, comparably little advances have been achieved in methods for *intuitive psychology*, i.e. common-sense reasoning about other agents’ *mental states* from their observed behaviour (Gandhi et al. 2021; Shu et al. 2021).

Similar to studies in developmental cognitive science, evaluation of common-sense reasoning abilities in these benchmarks uses a violation of expectation paradigm (Baillargeon 1987, VoE). In a set of familiarisation trials, an observer model has to form an expectation about a particular agent behaviour. In a subsequent test trial, given the context extracted from the familiarisation trials, the observer has to judge how expected the behaviour of the agent is. Expectedness is defined as the observer’s prediction error with respect to a ground truth (e.g. agent’s next action). The core idea behind the VoE paradigm is to pair test trials that may not differ much in traditional error metrics but differ in terms of human reasoning. As such, VoE allows probing AI capabilities by comparing scenarios that humans can differentiate based on those capabilities.

Inspired by behavioural experiments with infants, Gandhi et al. have recently introduced the Baby Intuitions Benchmark (Gandhi et al. 2021, BIB) – a set of tasks that require an observer model to reason about agents’ goals, preferences, and actions by observing their behaviour in a grid-world environment. The BIB poses two key challenges. First, despite being performed in the same grid-world environment, training and evaluation tasks differ in their number and setting. Moreover, test trials in the training set present only expected outcomes, requiring observer models to generalise to unseen situations by combining different pieces of knowledge gained during training. Second, the benchmark not only challenges models’ ability to predict future actions but also to use them to quantify the expectedness of novel situations. As such, the BIB can be seen as posing a meta-learning problem (Rabinowitz et al. 2018) in which

an observer model “learns to learn” about other agents’ behaviour. While initial models are based on the machine theory of mind network (Rabinowitz et al. 2018), the most recent model (Hein and Diepold 2022, VT) is based on a video transformer (Neimark et al. 2021). VT is rather successful in modelling agents’ preferences, but it fails in binding them to specific agents. Moreover, VT struggles with understanding the role of blocking barriers and that rational agents, in contrast to irrational ones, move efficiently towards their goal.

To address these limitations we introduce the *Intuitive Reasoning Network* (IRENE) – a novel neural network for core intuitive psychology. IRENE uses a graph neural network (GNN) to obtain rich state embeddings by processing graphs extracted from video frames as well as a transformer to encode the familiarisation trials in a context vector. We show that IRENE achieves new state-of-the-art performance on three out of five reasoning tasks defined on the BIB. Our model performs particularly well on tasks that existing models (Gandhi et al. 2021; Hein and Diepold 2022) struggle with: binding preferences to specific agents, differentiating between rational and irrational agents, and understanding how to deal with obstacles that block the agents’ goal object. We also show that for selected tasks, its predictions are in line with infants’ responses collected on a subset of the BIB (Stojnić et al. 2023). Finally, for the first time, we investigate the choice of training tasks on generalisation performance. Our analyses demonstrate our model’s ability to combine knowledge gained during training to solve unseen evaluation tasks. Moreover, we show that training on one type of task does not necessarily improve performance on similar evaluation tasks but that training on unrelated tasks can lead to improvements.

In summary, our contributions are two-fold:

- We propose IRENE – a novel model for intuitive psychology that combines a GNN and a transformer to learn rich state and context representations. IRENE achieves new state-of-the-art performance on three out of five BIB reasoning tasks. In particular, it is capable of binding preferences to specific agents and of modelling blocking obstacles and irrational agents better than existing models.
- We are first to provide a detailed analysis of the influence of the chosen training tasks on performance. IRENE can achieve new state-of-the-art reasoning performance only when trained on all training tasks, showing its ability to combine knowledge gained during training to solve unseen evaluation tasks.

Related Work

Common-Sense Reasoning Benchmarks

Benchmarks for common-sense reasoning are becoming increasingly popular as they enable the first steps towards intelligent, collaborative agents that reason like humans. Most research has focused on evaluating language processing (Bhagavatula et al. 2020; Huang et al. 2019; Zellers et al. 2019; Bisk et al. 2020; Sap et al. 2019; Sakaguchi et al. 2021) or visual scene understanding (Yi et al. 2020; Smith et al. 2019; Ates et al. 2022) through task-specific

assessments, e.g. predictive accuracy. However, a number of benchmarks that aim to evaluate the general ability of AI systems to reason about unexpected events or situations have recently emerged (Riochet et al. 2020; Shu et al. 2021; Gandhi et al. 2021; Dasgupta et al. 2021; Piloto et al. 2022; Weihs et al. 2022). Most of these benchmarks focus on intuitive physics, targeting physical concepts such as continuity, solidity, object persistence and gravity (Smith et al. 2019; Riochet et al. 2020; Dasgupta et al. 2021; Piloto et al. 2022; Weihs et al. 2022). Conversely, benchmarks that test models’ ability to reason about other agents have received less attention. A notable exception is the recent BIB (Gandhi et al. 2021) that is based on the findings that infants expect other agents to have goals, preferences and engage in instrumental actions. The benchmark includes a set of tasks to evaluate whether an AI observer can exhibit the same capabilities. Like the BIB, AGENT tests whether models can predict that agents have object-based goals and act efficiently (Shu et al. 2021). In contrast to the BIB, however, AGENT does not evaluate whether models can reason about multiple agents, inaccessible goals, instrumental actions, or distinguish between rational and irrational agents. Moreover, whilst AGENT’s training and evaluation sets present only minor differences and training is done using different leave-out splits, the BIB provides a single canonical split to maximise the evaluation of models’ generalisability. We evaluate our model on the BIB given its accessibility and relevance as a benchmark for assessing intuitive psychological reasoning.

Models for Intuitive Psychology

Models for reasoning about agents’ behaviour and mental states can be grouped into models based on Bayesian theory (Baker, Saxe, and Tenenbaum 2011; Baker et al. 2017) or deep learning (Rabinowitz et al. 2018). Shu et al. have introduced Bayesian Inverse Planning and Core Knowledge (BIPaCK) (Shu et al. 2021) that combines Bayesian inverse planning (Baker et al. 2017) and physics simulation (Battaglia, Hamrick, and Tenenbaum 2013), evaluating it on AGENT. More recently, Zhi et al. have evaluated a Bayesian Theory of Mind model with hierarchical priors over agents’ preference and efficiency (HBToM) on the BIB (Zhi-Xuan et al. 2022). However, this method makes strong assumptions and uses a tailored definition of expectedness. Gandhi et al. define expectedness in terms of the mean square error between the model prediction and the ground-truth (Gandhi et al. 2021). Instead, HBToM computes a *plausibility* score by training a set of logistic regression classifiers on a synthetic dataset similar to the BIB evaluation set. Therefore, in this work we do not compare our results with those of HBToM.

On the BIB, Gandhi et al. (Gandhi et al. 2021) have proposed a model based on the Theory of Mind neural network (ToMnet) introduced by (Rabinowitz et al. 2018). More recently, Hein et al. have proposed a method (VT) based on a video transformer (Neimark et al. 2021) that encodes frames using a CNN and performs cross- and self-attention over frames (Hein and Diepold 2022). In this work, we introduce a novel method that uses a GNN (Gori, Monfardini,

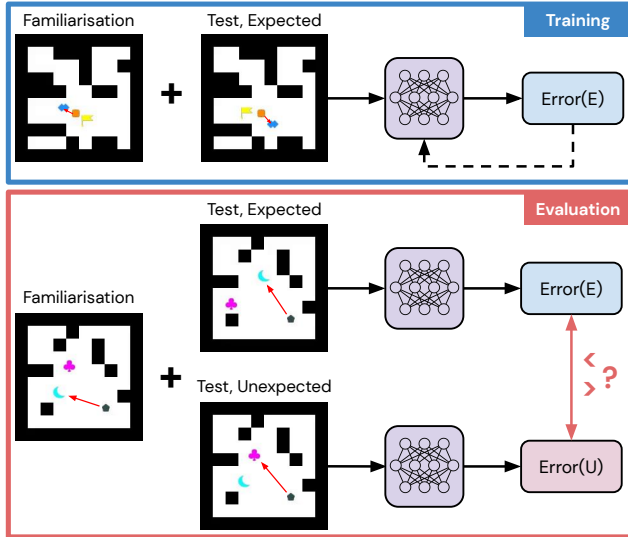


Figure 1: Training and evaluation on the Baby Intuitions Benchmark (Gandhi et al. 2021). During training, a model conducts eight familiarisation trials and a test trial, which is always expected. The model’s weights are updated using backpropagation on the error, computed with respect to the expected test trial ground truth. Evaluation employs a violation of expectation paradigm: based on the familiarisation trials, the model makes predictions on expected and unexpected test trials and both errors are compared. The model is successful if the error on the expected trials is smaller than that on the unexpected trials.

and Scarselli 2005) to encode graphs built from frames and a transformer (Vaswani et al. 2017) to generate context embeddings. On AGENT, Shu et al. have also used a GNN to encode states, where graphs connect the agent node to all the other nodes (Shu et al. 2021). Differently, our GNN performs message passing on heterogeneous graphs with edges representing different spatial relations.

The Baby Intuitions Benchmark

The BIB consists of 2D videos of an agent moving and interacting with different objects in a grid-world environment (see Figure 1). Both the agent and objects are represented as geometric shapes of different colours. The benchmark proposes five common-sense reasoning tasks (*Preference*, *Multi-Agent*, *Inaccessible Goal*, *Efficient Action*, and *Instrumental Action*) derived from research on infant intuitive psychology, that require an observer model to reason about agents’ goals, preferences, and actions by observing their behaviour in a grid-world environment. The *Efficient Action* and *Instrumental Action* present three sub-tasks each: *Efficiency Path Control*, *Efficiency Time Control*, *Efficiency Irrational Agent* and *Instrumental No Barrier*, *Instrumental Blocking Barrier*, *Instrumental Inconsequential Barrier*, respectively. Task descriptions are provided in the Appendix.

An episode in the BIB includes nine trials with trajectories $\{\tau_i\}_{i=1,\dots,9}$, where each trajectory consists of a series of

state-action pairs $\tau_i = \{(s_{ij}, a_{ij})\}_{j=1,\dots,T}$, with s_{ij} being video frames and T the trial length. $\{\tau_i\}_{i=1,\dots,8}$ are familiarisation trials and τ_9 is the test trial. Familiarisation trials serve to give a context to the model while test trials are used to make predictions. A test trial can be consistent with the familiarisation examples (expected outcome) or inconsistent (unexpected outcome). According to the VoE paradigm, if the observer is more surprised by the unexpected outcome, this means that what they believed or predicted would happen is not in line with what actually occurred. Expectedness is defined as the observer model’s prediction error: a model is successful if the prediction error on the unexpected outcome is higher than the error on the expected outcome. In practice, the prediction error is quantified by the mean squared error (MSE) with respect to a ground truth (e.g. next frame or action).

The authors provide a single canonical split for training and evaluation, as they differ in terms of tasks and sample distributions. The evaluation set presents the tasks mentioned above. Familiarisation and test trials follow different distributions: expected and unexpected trials are perceptually and conceptually different from familiarisation trials, respectively. Consider, for example, the *Preference* task (see Figure 1). In expected trials, the preferred object is located differently than in familiarisation trials but the agent still moves towards it. In the unexpected trials, the agent moves towards the non-preferred object whose location is the same as familiarisation.

The training set presents four tasks: *Single-Object*, *No-Navigation Preference*, *Single-Object Multiple-Agent* and *Agent-Blocked Instrumental Action*. The training tasks are more trivial and less informative than the evaluation tasks, compelling models to combine and generalise the knowledge acquired from the different training tasks to solve the evaluation tasks. For example, for the *Instrumental Blocking Barrier* task, a model has to put together knowledge of navigation (*Single-Object*) and instrumental actions (*Agent-Blocked Instrumental Action*). However, at the same time, the model has to generalise its knowledge: for example, in the training trials (*Agent-Blocked Instrumental Action*), the agent is confined in the barrier, whereas in the evaluation trials (*Instrumental Blocking Barrier*) it is the object which is confined. For a detailed description of the dataset we refer the interested reader to the original paper (Gandhi et al. 2021). We provide a brief summary in the Appendix.

Method

Graph Generation. In the BIB, each video is paired with a json file containing information about the grid-world layout, which we use to build graphs. Videos are sampled at 3 FPS and a graph $G_{ij} = (V_{ij}, E_{ij})$ is built from each sampled frame $j = 1, \dots, T$ in each trial $i = 1, \dots, 9$, where V_{ij} is a set of nodes and E_{ij} a set of edges. Each entity in a frame is assigned to a graph node $v_{ij} \in V_{ij}$, which has four features: *type* (e.g. “agent”), *position* (x and y coordinates), *colour* (GBR channels) and *shape* (e.g. “pentagon”). We represent categorical variables (*type* and *shape*) as one-hot vectors while *position* and *colour* are normalised between $[-1, 1]$ and $[0, 1]$, respectively.

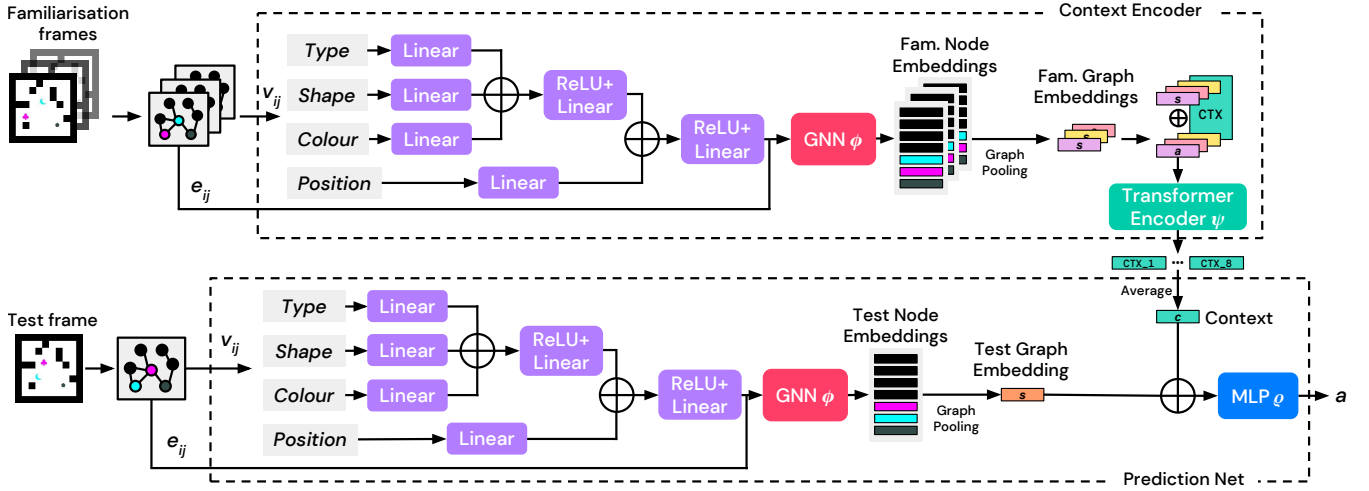


Figure 2: Architecture of IRENE. Inputs are graphs representing entities in a video frame. In the *context encoder*, the state encoder includes a feature fusion module that combines the four node features into a single one, which is input to a GNN. The encoded states are concatenated to the corresponding actions and to a special learnable CTX token. A transformer context encoder produces a context embedding as the mean of the CXT embedding vectors of each familiarisation trial. In the *prediction net*, the encoded test state is concatenated to the context and input into a MLP policy that outputs a prediction for the agent’s next action.

Edges $e_{ij} \in E_{ij}$ represent spatial relationships and are defined following (Jiang et al. 2021). Specifically, we consider *local directional relations*, which identify the relative position of two adjacent entities, *remote directional relations*, which do not require adjacency, and two *non-directional relations*, adjacent and aligned (see the Appendix for formal definitions).

Model Architecture. The architecture of our *Intuitive Reasoning Network* (IRENE) is shown in Figure 2. A *context encoder* parses the agent’s past trajectories (i.e. state-action pairs) into a context vector, and a *prediction net* predicts the future behaviour of the agent based on the context and the current state. In the state encoder, a feature fusion module combines the node features (type, position, colour, shape). For each node $v_{ij} \in V_{ij}$, each feature v_{ij}^k is first embedded using a linear layer f_k . Then, type, colour, and shape are concatenated and passed through a first fusion layer consisting of a ReLU activation (Fukushima 1975; Nair and Hinton 2010) and a linear layer f_{F_1} . The output is concatenated with the position embedding and passed through the second fusion layer, analogous to the first:

$$v_{ij}^{t,s,c} = f_{F_1}(\text{ReLU}(\bigoplus_{k=\text{type}, \text{shape}, \text{colour}} f_k(v_{ij}^k))), \quad (1)$$

$$v'_{ij} = f_{F_2}(\text{ReLU}(f_{\text{position}}(v_{ij}^{\text{position}}) \oplus v_{ij}^{t,s,c})), \quad (2)$$

where \oplus indicates concatenation. The resulting vector constitutes the input feature for the GNN that performs message passing to update the node embeddings and produces a state embedding h_{ij} by applying a final graph average pooling operation:

$$h_{ij} = \text{AvgPooling}(\phi(v'_{ij}, e_{ij})). \quad (3)$$

As different edges represent different relations, ϕ is a Relational GNN (Schlichtkrull et al. 2018) which uses different weights for each edge type. In particular, we use GraphSAGE layers with LSTM aggregation (Hamilton, Ying, and Leskovec 2017).

In the context encoder, the state encoder outputs an embedding vector h_{ij} for each frame graph G_{ij} , obtained by applying average pooling to the nodes. Then, the encoded states $\{h_{ij}\}_{j=1,\dots,T}$ are concatenated to the corresponding actions $\{a_{ij}\}_{j=1,\dots,T}$ and projected to the transformer input dimension by a linear layer f_{proj} . A learnable CTX token is concatenated to each embedding vector. Thereafter, positional embedding is added, followed by layer normalisation (Ba, Kiros, and Hinton 2016). The result is input into a transformer encoder ψ :

$$\text{CTX}'_i = \psi(f_{\text{proj}}(\{h_{ij}\}_{j=1}^T \oplus \{a_{ij}\}_{j=1}^T) \oplus \text{CTX}_i) \quad (4)$$

and the output CTX'_i are taken as trial representations. The context encoder outputs a single context embedding obtained by computing the mean of the eight familiarisation trial representations,

$$c = \frac{1}{8} \sum_{i=1}^8 \text{CTX}'_i. \quad (5)$$

In the prediction net, a test frame graph $G_{9,j} = (V_{9,j}, E_{9,j})$ is encoded by the same state encoder (Eq. 1, 2, 3). The resulting state embedding $h_{9,j}$ is concatenated to the context embedding c and input into an MLP policy ρ that outputs the next action prediction:

$$a_{\text{pred}} = \rho(c \oplus h_{9,j}), \quad (6)$$

which in our case is represented by the agent’s next position in the grid-world.

BIB Task	BC-MLP	BC-RNN	Video-RNN	VT	IRENE
Preference	26.3	<i>48.3</i>	47.6	80.8	48.5
Multi-Agent	48.7	48.2	<u>50.3</u>	<i>49.2</i>	74.9
Inaccessible Goal	76.9	<i>81.6</i>	74.0	<u>85.5</u>	85.8
Eff. Path Control	94.0	92.8	99.2	<i>97.5</i>	<u>98.1</u>
Eff. Time Control	99.1	99.1	<u>99.9</u>	<i>99.7</i>	100.0
Eff. Irrational Agent	<u>73.8</u>	<i>56.5</i>	50.1	34.1	85.7
Eff. Action Average	<u>88.8</u>	82.5	<i>83.1</i>	77.1	94.7
Inst. No Barrier	98.8	98.8	99.7	<i>97.9</i>	78.4
Inst. Incons. Barrier	55.2	<u>78.2</u>	<i>77.0</i>	91.9	52.4
Inst. Blocking Barrier	47.1	56.8	<i>62.9</i>	64.2	83.5
Inst. Action Average	67.0	<i>77.9</i>	<u>79.9</u>	84.7	71.5

Table 1: VoE accuracy of existing models and IRENE on the BIB evaluation set. Best score is in bold, second best score is underlined, third best score is italic.

Experiments

Technical Details

IRENE’s feature fusion module encodes the node features using linear layers of hidden dimension 96. The state encoder consists of two GraphSAGE layers for each relation, with hidden dimension 96 and ELU activation (Clevert, Unterthiner, and Hochreiter 2015). The transformer encoder consists of a stack of six layers with four attention heads, feedforward dimension 512 and GELU activations (Hendrycks and Gimpel 2016). The prediction net uses the same GNN and feature fusion module used in the context encoder. The MLP policy has hidden dimensions 256, 128 and 256 and output dimension two, corresponding to the (x, y) coordinates of the agent in the next frame. Additional training details are reported in the Appendix.

Results

We compare IRENE with the three models originally proposed by Gandhi et al. (Gandhi et al. 2021) – BC-MLP, BC-RNN, and Video-RNN – as well as the more recent VT model by Hein et al. (Hein and Diepold 2022). The VoE accuracy scores on all evaluation tasks, averaged over three different runs, are shown in Table 1. Given that results did not vary much across different runs, we only report the errors for all evaluations in the Appendix. In line with previous work (Gandhi et al. 2021; Hein and Diepold 2022), we calculated expectedness as the maximum prediction error. We additionally evaluated our model using the mean prediction error and report these results in the Appendix. As can be seen from Table 1, our model achieves state-of-the-art results on three out of five tasks (*Multi-Agent*, *Inaccessible Goal*, *Efficient Action*). Moreover, when also considering the *Efficient Action* and *Instrumental Action* sub-tasks, IRENE achieves state-of-the-art results on five out of nine tasks. We conducted t-tests to compare IRENE’s performance with the baselines. All results were significant ($\alpha = 0.05$, $p < 0.01$) with only two exceptions: *Preference* between IRENE and BC-RNN; *Time Control* between IRENE and Video-RNN.

BIB Task	LSTM	GCN	Local	Remote	IRENE
Preference	48.2	49.7	49.8	50.7	48.5
Multi-Agent	49.7	50.3	98.2	50.0	79.4
Inaccessible Goal	84.8	58.1	41.1	80.6	85.8
Eff. Path Control	97.3	94.7	31.7	98.2	98.1
Eff. Time Control	99.9	98.5	37.6	99.8	100.0
Eff. Irrational Agent	52.4	89.3	99.7	83.6	85.7
Eff. Action Average	83.2	94.2	56.3	93.9	94.7
Inst. No Barrier	78.5	64.6	51.6	78.7	78.4
Inst. Incons. Barrier	53.3	52.1	52.4	52.7	52.4
Inst. Blocking Barrier	83.2	48.0	48.9	83.8	83.5
Inst. Action Average	71.7	54.8	51.0	71.7	71.5

Table 2: VoE accuracy for ablated versions of IRENE. “LSTM” makes use of an LSTM context encoder instead of the transformer; “GCN” substitutes GraphSAGE with GCN layers; “Local” takes as input relational graphs with only local directional relations; and “Remote” takes as input relational graphs with only remote directional relations.

The larger improvements are in the *Multi-Agent*, *Instrumental Blocking Barrier* and *Efficiency Irrational Agent* tasks. In the *Multi-Agent* task, IRENE dramatically outperforms the other models, improving over the previous best score (Video-RNN) by 48.9%. In the *Instrumental Blocking Barrier* sub-task, our model improves by 30% on the previous best score (VT). In the *Efficiency Irrational Agent* task, IRENE outperforms BC-MLP by 16%. This results in an improvement in the *Efficient Action* task of 6.6% with respect to the previous best model (BC-MLP). Similar to the baseline methods, our model performs well on the *Path Control* and *Time Control* sub-tasks. Remarkably, the score on the *Time Control* sub-task is perfect. Scores on the *No Barrier* and *Inconsequential Barrier* tasks are lower than those of the other methods. Overall, our model struggles the most in the *Preference* (48.5) and *Instrumental Action* tasks (average 71.5), especially in the *Inconsequential Barrier* (52.4).

Ablation Studies

To investigate how different components of our method contribute to these performance improvements, we performed a series of ablation studies summarised in Table 2.

Graph Relations. We trained IRENE on graphs whose edges represent only local or remote directional relations. Using only local directional relations, the performance on the *Multi-Agent* and *Efficiency Irrational Agent* tasks improved to an almost perfect score. However, performance on other tasks became worse, especially in the *Time* and *Path Control* sub-tasks that other models solved almost perfectly. This decrease in performance was to be expected as using local relations alone leaves many nodes isolated, including the agent’s node. As a consequence, the message passing is lacking important information, such as the presence of obstacles that are not adjacent to the agent or to the objects. Using only remote relations, performance is comparable to the original model except for *Multi-Agent*, which is at chance level. In

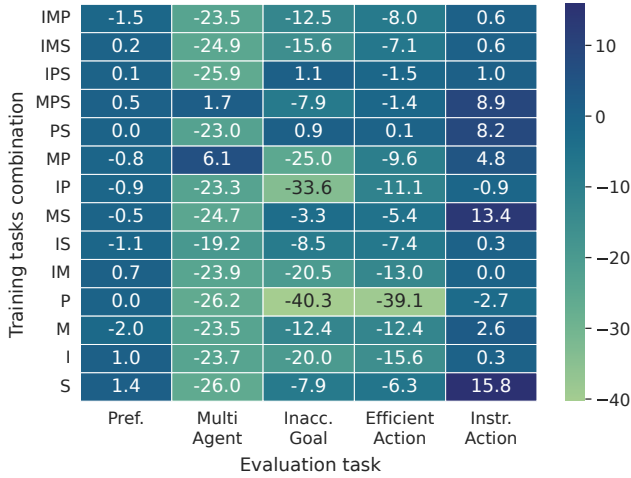


Figure 3: Relative difference of VoE scores obtained on the evaluation set by training on all possible combinations of training tasks with respect to training on all tasks.

combination, this suggests that remote relations contribute more to the final scores than local ones.

State Encoder. We replaced the GraphSAGE layers with GCN layers in the state encoder. The resulting model achieved considerably worse scores in the *Multi-Agent*, *Inaccessible Goal*, *Instrumental No Barrier* and *Blocking Barrier* (sub-)tasks. This suggests that GraphSAGE is more effective for modelling obstacles and, as such, is key for our model to better understand the role of blocking barriers and to bind preferences to specific agents.

Context Encoder. We also replaced the transformer encoder ψ with an LSTM. Performance remained mostly unchanged except for the *Multi-Agent* and *Efficiency Irrational Agent* tasks where performance dropped to chance level. This is in line with (Gandhi et al. 2021) who showed that models with an LSTM failed to adapt their predictions according to whether an agent was rational or irrational during familiarisation. In combination with the results obtained when substituting GraphSAGE with GCN, this shows that good performance on the *Multi-Agent* task can be obtained only by including both GraphSAGE and a transformer in the model architecture.

Analysis of the Training Tasks

Prior work on the BIB has focused on improving performance (Gandhi et al. 2021; Hein and Diepold 2022) while the question of if and how the choice of training task(s) impacts performance remains under-explored. This is surprising given that, similar to human learning (Clarke and Roche 2010), it is reasonable to assume that computational models also benefit more from some training tasks than others. To fill this gap, we performed the first investigations into the influence of training tasks on evaluation performance. To this end, we trained IRENE on the four individual tasks as well as on all of their possible combinations. To indi-

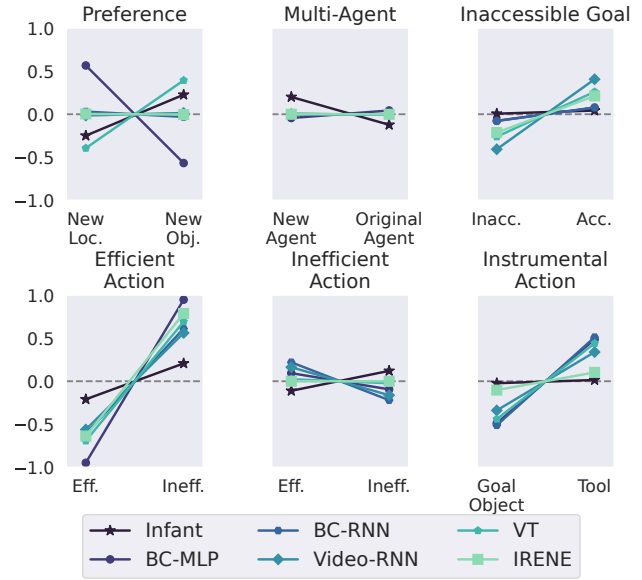


Figure 4: Z-scored means of infants’ looking times and models’ scores for expected and unexpected outcomes in the BIB evaluation episodes. Positive values indicate expectedness, negative values indicate unexpectedness.

cate a combination of training tasks, we used the following notation: P = *No-Navigation Preference*, M = *Single-Object Multi-Agent*, I = *Agent-Blocked Instrumental Action*, S = *Single-Object*. For example, IMP indicates joint training on *Agent-Blocked Instrumental Action*, *Single-Object Multi-Agent*, and *No-Navigation Preference* tasks.

Relative differences in performance scores compared to training on all tasks are shown in Figure 3 (absolute scores are reported in the Appendix). Training only on a subset of tasks generally leads to a decrease in performance – in some cases drastic. This demonstrates the importance but also the effectiveness of IRENE in extracting and combining knowledge gained from different training tasks. One notable exception is when training on MPS: in this case the total average performance is comparable to full training (75.1 vs. 75.4). In particular, the score on the *Inaccessible Goal* task is worse (77.7) while the one on the *Instrumental Action* task is better (80.4), despite the model not having been trained on I. Such improvement is due to better scores on sub-tasks in which blocking barriers are absent or irrelevant: *No Barrier* (84.8) and *Inconsequential Barrier* (92.7).

Training on S improved scores both on the *No Barrier* and *Inconsequential Barrier* sub-tasks, resulting in a higher score in the *Instrumental Action* task. When training on M, performance on the *Multi-Agent* task is considerably worse and training on P does not lead to improvements in the *Preference* task. This is in line with the general observation that training on one type of task does not always improve performance for similar types of tasks in the evaluation set.

When trained on a subset of the tasks, performance scores show a similar pattern. When IRENE is trained on P and M but not on I (i.e. MPS, MP), performance on *Multi-Agent* is

good. Considering each evaluation task, while performance on the *Preference* task remains unchanged, it degrades notably for *Multi-Agent*, *Inaccessible Goal*, and *Efficient Action* in most cases. The score on the *Instrumental Action* task increased for selected combinations of training tasks (MPS, PS, MS, S), especially in MS and S.

Comparison to Infants’ Intuitions

Figure 4 shows a comparison between the z-scored means of infants’ looking times as collected by Stojnić et al. (2023) and the corresponding models’ scores for expected and unexpected evaluation trials. As can be seen from the figure, our model’s expectations generally align with those of the infants, specifically for the *Inaccessible Goal*, *Efficient Action*, *Inefficient Action*, and *Instrumental Action* tasks. IRENE performs better in the *Inefficient Action* task in which the other models are more surprised at the expected outcome than the unexpected one. One notable outlier is *Multi-Agent* for which infant behaviour differs from all models. As reported by (Stojnić et al. 2023), this can be explained by infants not always reacting as expected.

Discussion

Our experiments demonstrate the effectiveness of our model in addressing intuitive psychology tasks. IRENE outperforms the state of the art for three out of five reasoning tasks on the challenging Baby Intuitions Benchmark – in some cases with significant improvements (see Table 1). More specifically, IRENE better learns to bind preferences to specific agents and to model the role of blocking barriers. This suggests that, in contrast to existing models (Gandhi et al. 2021; Hein and Diepold 2022), IRENE relies less on heuristics, such as directly moving towards the goal object. However, this also leads to lower scores on sub-tasks for which such simple heuristics are sufficient and which, consequently, allow models that employ them to perform better. One example of this is the *Instrumental Action* tasks, where models that did not learn the role of barriers during training apply a simple heuristic – i.e. directly moving towards the goal object – that works on the *Instrumental No Barrier* and *Inconsequential Barrier* sub-tasks but not on the more challenging *Blocking Barrier* sub-task. In contrast, IRENE has proven to be effective in handling complex tasks that cannot be tackled by heuristics alone. This highlights the importance of training models that can truly reason, rather than just applying shortcuts or heuristics, and resonates with the well known problem of reward hacking in deep reinforcement learning (Amodei et al. 2016). IRENE represents an important advancement on this challenging endeavour. Our model also obtains near-perfect or perfect scores on the *Path Control* and *Time Control* sub-tasks, demonstrating that it can effectively find the shortest path to an object goal. The improved score on the *Efficient Action* task suggests that IRENE can also better model rational agents’ behaviour. In particular, in the *Irrational Agent* sub-task, our model can distinguish between rational agents that move efficiently towards their goal and irrational agents that are not as efficient. The *Irrational Agent* is also the task in which the IRENE z-scored mean is the closest to the infants’ (see Fig. 4).

Our ablation studies (see Table 2) demonstrate the effectiveness of the proposed combination of GraphSAGE for agent and world state encoding with a transformer for task context encoding. We speculate GraphSAGE performs well thanks to inductivity: learning an aggregator allows the model to effectively generate embeddings for nodes which the model only sees during the evaluation trials. In addition, the non-sequential nature of the transformer together with its self-attention mechanism allow it to overcome LSTM limitations. We also performed an extensive analysis to better understand the importance of the chosen training tasks on evaluation performance. Our results confirm the requirements set by Gandhi et al. (2021) who have argued that models have to combine knowledge from different training tasks: IRENE performs best when trained on all training tasks while training only on a subset of tasks generally leads to a decrease in performance (see Fig. 3). Exceptions are mainly due to a lack of knowledge of blocking barriers, which results in the effective heuristic of ignoring them.

Limitations and Future Work. By working on IRENE we identified several points that we believe are crucial for future work. First, IRENE does not perform better on all tasks – just like the previous models we compared it to. This suggests that further advances are still needed. To fully solve the BIB tasks, it may also be necessary to learn complementary basic concepts, such as from intuitive physics. Second, besides notable exceptions such as the BIB, advances in intuitive psychology are currently slowed down by the lack of accessible and well-maintained benchmarks. In particular, we would have liked to evaluate our model on AGENT (Shu et al. 2021) – that covers an interesting environment and complementary set of challenging reasoning tasks – but AGENT does not provide any benchmark or model code. Despite our significant efforts, re-implementing the method and reproducing the results based solely on the information provided in the paper turned out to be an extremely challenging task. Likely because of this also no other papers have evaluated on AGENT to date. There is an urgent need for the community to design and create new benchmarks to foster the development of “general neural common-sense reasoners”. Third, despite the general trend of designing benchmarks to be used only for evaluation (Weihs et al. 2022), future work should also explore how introducing new training tasks allows models to learn more effectively.

Conclusion

In this work we have proposed IRENE – a novel neural network for reasoning about agents’ goals, preferences, and actions. IRENE sets new state-of-the-art on three out of five tasks on the Baby Intuitions Benchmark, with improved modelling of agents’ preferences, obstacle handling, and distinguishing between rational and irrational agents. We also demonstrated the effectiveness of IRENE in combining knowledge gained during training for unseen evaluation tasks. These results are not only promising for advancing human-like reasoning in AI systems but also shed new light on the importance of the choice of training tasks for good generalisation performance.

Acknowledgements

M. Bortoletto and A. Bulling were funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 801708. L. Shi was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2075 – 390740016. The authors would like to especially thank Pavel Denisov, Hsiu-Yu Yang, Ekta Sood, and Manuel Mager for numerous insightful discussions, and thank Ann-Sophia Müller and Constantin Ruhdorfer for their technical support.

References

- Aguiar, A.; and Baillargeon, R. 1999. 2.5-month-old infants’ reasoning about when objects should and should not be occluded. *Cognitive Psychology*, 39(2): 116–157.
- Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete Problems in AI Safety. *arXiv preprint arXiv:1606.06565*.
- Ates, T.; Atesoglu, M. S.; Yigit, C.; Kesen, I.; Kobas, M.; Erdem, E.; Erdem, A.; Goksun, T.; and Yuret, D. 2022. CRAFT: A Benchmark for Causal Reasoning About Forces and Interactions. *Findings of the Association for Computational Linguistics*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450*.
- Baillargeon, R. 1987. Object Permanence in 3 1/2- and 4 1/2-Month-Old Infants. *Developmental Psychology*, 23(5): 655–664.
- Baker, C.; Saxe, R.; and Tenenbaum, J. 2011. Bayesian Theory of Mind: Modeling Joint Belief-Desire Attribution. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33.
- Baker, C. L.; Jara-Ettinger, J.; Saxe, R.; and Tenenbaum, J. B. 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4): 1–10.
- Baron-Cohen, S. 1997. *Mindblindness: An Essay on Autism and Theory of Mind*. MIT Press.
- Battaglia, P. W.; Hamrick, J. B.; and Tenenbaum, J. B. 2013. Simulation as an engine of physical scene understanding. In *Proceedings of the National Academy of Sciences*, volume 110, 18327–18332. National Acad Sciences.
- Bhagavatula, C.; Bras, R. L.; Malaviya, C.; Sakaguchi, K.; Holtzman, A.; Rashkin, H.; Downey, D.; Yih, W.; and Choi, Y. 2020. Abductive Commonsense Reasoning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Bisk, Y.; Zellers, R.; Bras, R. L.; Gao, J.; and Choi, Y. 2020. PIQA: Reasoning about Physical Commonsense in Natural Language. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 7432–7439. AAAI Press.
- Clarke, D.; and Roche, A. 2010. Teachers’ Extent of the Use of Particular Task Types in Mathematics and Choices behind That Use. *Mathematics Education Research Group of Australasia*.
- Clevert, D.-A.; Unterthiner, T.; and Hochreiter, S. 2015. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). *International Conference on Learning Representations*.
- Conti, C. J.; Varde, A. S.; and Wang, W. 2022. Human-Robot Collaboration With Commonsense Reasoning in Smart Manufacturing Contexts. *IEEE Transactions on Automation Science and Engineering*.
- Dasgupta, A.; Duan, J.; Ang Jr, M. H.; Lin, Y.; Wang, S.-h.; Baillargeon, R.; and Tan, C. 2021. A Benchmark for Modeling Violation-of-Expectation in Physical Reasoning Across Event Categories. *arXiv preprint arXiv:2111.08826*.
- Fukushima, K. 1975. Cognitron: A self-organizing multi-layered neural network. *Biological Cybernetics*, 20(3): 121–136.
- Gandhi, K.; Stojnic, G.; Lake, B. M.; and Dillon, M. R. 2021. Baby Intuitions Benchmark (BIB): Discerning the goals, preferences, and actions of others. *Advances in Neural Information Processing Systems*, 34: 9963–9976.
- Gergely, G.; Nádasdy, Z.; Csibra, G.; and Bíró, S. 1995. Taking the intentional stance at 12 months of age. *Cognition*, 56(2): 165–193.
- Gori, M.; Monfardini, G.; and Scarselli, F. 2005. A new model for learning in graph domains. In *IEEE International Joint Conference on Neural Networks*, volume 2, 729–734.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. *Advances in Neural Information Processing Systems*, 30.
- Hein, A.; and Diepold, K. 2022. Comparing Intuitions about Agents’ Goals, Preferences and Actions in Human Infants and Video Transformers. In *Shared Visual Representations in Human and Machine Intelligence Workshop at Advances in Neural Information Processing Systems*.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415*.
- Huang, L.; Le Bras, R.; Bhagavatula, C.; and Choi, Y. 2019. Cosmos QA: Machine Reading Comprehension with Contextual Commonsense Reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2391–2401.
- Jiang, Z.; Minervini, P.; Jiang, M.; and Rocktäschel, T. 2021. Grid-to-Graph: Flexible Spatial Relational Inductive Biases for Reinforcement Learning. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi Agent Systems*, 674–682.
- McCarthy, J. 1989. Artificial Intelligence, Logic and Formalizing Common Sense. In *Philosophical Logic and Artificial Intelligence*, 161–190. Springer.
- Mota, T.; and Sridharan, M. 2019. Commonsense Reasoning and Knowledge Acquisition to Guide Deep Learning on Robots. In *Robotics: science and systems*.

- Nair, V.; and Hinton, G. E. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *International Conference on Machine Learning*.
- Needham, A.; and Baillargeon, R. 1993. Intuitions about support in 4.5-month-old infants. *Cognition*, 47(2): 121–148.
- Neimark, D.; Bar, O.; Zohar, M.; and Asselmann, D. 2021. Video Transformer Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3163–3172.
- Piloto, L.; Weinstein, A.; TB, D.; Ahuja, A.; Mirza, M.; Wayne, G.; Amos, D.; Hung, C.-c.; and Botvinick, M. 2018. Probing Physics Knowledge Using Tools from Developmental Psychology. *arXiv preprint arXiv:1804.01128*.
- Piloto, L. S.; Weinstein, A.; Battaglia, P.; and Botvinick, M. 2022. Intuitive physics learning in a deep-learning model inspired by developmental psychology. *Nature Human Behaviour*, 1–11.
- Rabinowitz, N.; Perbet, F.; Song, F.; Zhang, C.; Eslami, S. A.; and Botvinick, M. 2018. Machine Theory of Mind. In *International Conference on Machine Learning*, 4218–4227. PMLR.
- Riochet, R.; Castro, M. Y.; Bernard, M.; Lerer, A.; Fergus, R.; Izard, V.; and Dupoux, E. 2020. IntPhys: A Framework and Benchmark for Visual Intuitive Physics Reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2021. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. *Communications of the ACM*, 64(9): 99–106.
- Sap, M.; Rashkin, H.; Chen, D.; Le Bras, R.; and Choi, Y. 2019. Social IQa: Commonsense Reasoning about Social Interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4463–4473.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Berg, R. v. d.; Titov, I.; and Welling, M. 2018. Modeling Relational Data with Graph Convolutional Networks. In *European Semantic Web Conference*, 593–607. Springer.
- Shu, T.; Bhandwaldar, A.; Gan, C.; Smith, K.; Liu, S.; Gutfreund, D.; Spelke, E.; Tenenbaum, J.; and Ullman, T. 2021. AGENT: A Benchmark for Core Psychological Reasoning. In *International Conference on Machine Learning*, 9614–9625. PMLR.
- Smith, K.; Mei, L.; Yao, S.; Wu, J.; Spelke, E.; Tenenbaum, J.; and Ullman, T. 2019. Modeling Expectation Violation in Intuitive Physics with Coarse Probabilistic Object Representations. *Advances in Neural Information Processing Systems*, 32.
- Stojnić, G.; Gandhi, K.; Yasuda, S.; Lake, B. M.; and Dillon, M. R. 2023. Commonsense psychology in human infants and machines. *Cognition*, 235: 105406.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention Is All You Need. *Advances in Neural Information Processing Systems*, 30.
- Weihs, L.; Yuile, A. R.; Baillargeon, R.; Fisher, C.; Marcus, G.; Mottaghi, R.; and Kembhavi, A. 2022. Benchmarking Progress to Infant-Level Physical Reasoning in AI. *Transactions on Machine Learning Research*.
- Yi, K.; Gan, C.; Li, Y.; Kohli, P.; Wu, J.; Torralba, A.; and Tenenbaum, J. B. 2020. CLEVRER: Collision Events for Video Representation and Reasoning. *International Conference on Learning Representations*.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4791–4800.
- Zhi-Xuan, T.; Gothoskar, N.; Pollok, F.; Gutfreund, D.; Tenenbaum, J. B.; and Mansinghka, V. K. 2022. Solving the Baby Intuitions Benchmark with a Hierarchically Bayesian Theory of Mind. *Robotics: Science and Systems Workshop on Social Intelligence in Humans and Robots*.