

Cognitive Inductive Prejudice for Corporal Edifice in Hominids and Contraption

Chandra Bhim Bhan Singh
Department of Computer Engineering,
National Institute of Technology, Kurukshetra
Kurukshetra, Haryana - 136119
bhimbhan@gmail.com

ABSTRACT

A strong and insightful interpretation of scientific knowledge and practice must take into consideration how human cognitive skills and constraints enable as well restrict the scientific enterprise's activities and products. While existing deep learning systems are outstanding in functions such as object classification, language processing, and gameplay but few can create or transform a complex system like a Frame Pyramid. Assume that what these systems lack is a "Cognitive Inductive Prejudice": an ability to justify inter-object relationships and make decisions about an organized description of the incident. In order to assess this premise, this paper concentrated on a work involving stapling together stacks of frames to balance a castle and quantify how well hominids are doing. Then for analyzing contraption capability, our work introduce the Significant Stimulus Learning Tool that utilizes object-and interaction-centered scene and policy representations, these apply to the task. Our results shows that these structural portrayals enable the tool to perform both hominids and contraption for more naive methods, indicating that cognitive inductive effect is a significant element in solving structured reasoning issues and building more intelligent also flexible for machines.

Keywords: Cognitive science, Inductive reasoning, Graph network, Stimulus Learning Tool, Cognitive Inductive Prejudice.

Date of Submission: Feb 24, 2020

Date of Acceptance: March 07, 2020

I. INTRODUCTION

Human physical reasoning and more widely cognition is rooted in a wealthy structure of object and relationship understanding that can be comprised to promote strong forms of combinatorial generalization. One of the major types of reasoning in cognitive process is inductive reasoning. This type of reasoning includes making projections based on current information about novel scenarios [1]. These projections are probabilistic necessarily. For instance, if you are told that grizzly bears have some kind of enzyme, you may be mildly optimistic, but not sure that this property will generalize to other bears.

Induction is much of the reasoning individuals make in their daily lives. Predicting whether tomorrow's rain is probable, how your partner will respond to the chocolate box you bought as a donation, or whether inventory prices will increase in the next six months all require some sort of induction. More broadly, induction involves a variety of cognitive operations such as categorization, assessment of probability, analogical reasoning, science inference, and decision-making [2]. One of the reasons why it has become a significant field of research for cognitive scientists is the pervasive nature of induction. Another reason is that inductive reasoning seems to tackle one of the key cognitive science issues, namely how knowledge from known to unknown instances is generalized.

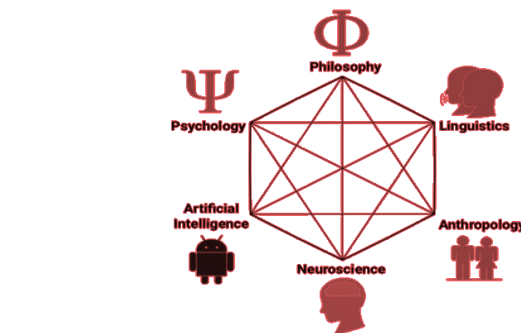


Figure 1: Included Techniques in Cognitive Science

The profound reinforcement learning (RL) community has shown that a broad variety of difficult assignments can be mastered by well-tuned deep RL algorithms. On board games like go and Chess video games like Atari and custom 3D navigation assignments, human level performance has been approached or exceeded [3]. These findings testify to the general approach's generality. However, sometimes the excitement for the constant stream of new domains being mastered by suitable RL algorithms may have overshadowed the dependence on these agents' inductive biases and the amount of tuning that is often required for them to perform effectively in new domains [4]. The Alpha Zero algorithm is a straightforward illustration of the advantages of generality. By removing all dependencies on go-specific inductive biases and human information, Alpha Zero varies from the previous Alpha Go algorithm. Alpha Zero not only

achieved greater efficiency in the Go game after removing these biases, but was also able to know how to play Chess and Shogi efficiently [5]. In particular, when we inject inductive biases into our algorithms, there is a tradeoff between generality and efficiency. There are many types of inductive biases, including domain knowledge and pretuned parameters of teaching. Such biases can result in quicker and better learning if implemented closely. Crucially, most inductive biases are not free: significant effort may be needed to acquire appropriate domain knowledge or pretune parameters, for example. For example, this cost is often hidden, one could use established hyper parameters as good in prior work on the same domain, without knowing how much data or time was spent optimizing these, or how specific the settings are for the given domain [6]. Systematic studies on the effect of various inductive biases are uncommon, and the generality of these various biases is often uncertain.

Analogous to the description of language productivity by von Humboldt as making "unlimited use of finite means," objects and relationships are the construction blocks that assist clarify how our knowledge of the everyday scene can function over infinite situations. Likewise, by leveraging these same depictions, individuals interact with everyday scenes [7]. Among the most impressive human behavior is our omnipresent drive to construct stuff, an ability to compose items and components under relational limitations, from pyramids to space stations, resulting in our most notable accomplishments.

One of Artificial Intelligence (AI)'s basic goals is to be able to communicate as robustly and flexibly with the globe as individuals do [8]. Assume that this flexibility is partly provided by called as relational inductive bias. More usually, an inductive bias is the set of conditions of a learning algorithm that contributes to selecting one hypothesis over another independent of the information observed [9]. Such assumptions can be encoded in a Bayesian model or installed in a neural network through architectural assumptions. For instance, a convolutionary neural network's weight-sharing architecture induces an inductive prejudice of translational invariance might call a "spatial inductive bias" because it builds on the world's spatial structure in particular assumptions [10]. Similarly, in particular assumptions about the world's relational structure, a relational inductive bias develops.

Hence, practical and mathematical models obviously contain inductive relational biases, resulting from propositional or causal descriptions. Thus, existing methods often struggle with structured and combinatorial issues, so there is a great need to introduce a novel solution. The rest of paper is organized as follows. In Section 2, survey of previous literature is addressed. In section 3, our proposed frame work is presented and discussed. Section 4, gives the result and evaluate the performance of our proposed method under various scenarios. Finally, we conclude the paper in section 5.

II. LITERATURE SURVEY

Chang et al. [11] propose a factorization of a physical scene into composable object-based depictions and architecture of a neural network whose compositional design factorizes object dynamics into parity relationships. Like a symbolic physics engine, the NPE is equipped with generic concepts of objects and their interactions; realized as a neural network, it can be taught to adapt to particular object characteristics and dynamics of different worlds by means of stochastic gradient descent.

Li et al. [12] presented article, they contrast a more traditional strategy of taking a model-based path through an end-to-end strategy with explicit 3D representations and physical simulation that directly predicts stability from appearance. They are asking the question whether and to what extent and quality such a skill can be obtained directly in a data-driven manner bypassing the need for explicit simulation at runtime. They present a learning-based strategy oriented on simulated information that predicts the stability of towers consisting of wooden blocks under various circumstances and quantities associated with the towers ' prospective collapse.

Denil et al. [13] introduce a basic set of tasks requiring agents to estimate properties such as mass and object cohesion in an interactive simulated environment in which objects can be manipulated and the consequences observed. We discovered that techniques of learning deeply reinforced can learn to conduct the experiments needed to uncover such hidden characteristics.

Yildirim et al. [14] present an assignment that examines how individuals communicate with block towers and make decisions. In Experiment 1, a series of issues were solved by lab respondents in which they had to reconfigure three blocks from an initial to a final setup. Whether they used one or two hands to do so, we registered. In Experiment 2, they asked internet respondents to judge whether they believe that one or two hands were used by the individual in the laboratory.

wang et al. [15] propose Nerve Net to explicitly model an agent's structure that takes the form of a graph naturally. In particular, as the policy network of the agent, Nerve Net first propagates information about the agent's structure and then predicts actions for different parts of the agent. In the tests, we first demonstrate that our Nerve Net on normal MuJoCo settings is similar to state-of - the-art techniques. Fragkiadaki et al. [16] explore how an agent can be fitted with an internal model of external world dynamics and how this model can be used to schedule new activities by conducting various inner simulations (' visual imagination'). It processes raw visual input directly and uses a novel object-centered forecast formulation based on object-centered visual glimpses (fixations) to implement translational invariance of learned physical laws. The agent trains himself through random interaction with a set of distinct settings, and then the resulting model can be

used to schedule goal-oriented activities in new settings that the agent had never met before.

Lian et al. [17] Study indicates that with overlapping neural mechanisms susceptible to numerical inductive reasoning and calculation procedures, frontal-parietal network activity separates the cognitive differential processes between them. According to ACT-R cognitive architecture, the exchange of data between intermediate depictions and long-term declarative knowledge during the rule recognition procedures peculiar to numerical inductive reasoning is attributable to more activity in this network particular to numerical inductive reasoning.

In [11] provides neural networks for relational reasoning, these papers not consider the chart system for quantify an approach. [12] Provide for physical reasoning do not explicitly the notion of objects or relation [13] describes about learning physical simulation not for graphical chart simulant. [14] Does not include inductive bias relational learning. [15] Represents graphical network not for interaction. [16] Describes visual imagination, it does not include relational prejudice. [17] Indicate approach has numerical inductive reasoning, but there is a lack of learning algorithm. From the above discussion it is clear that there is a need to analyze and quantifies the performance of hominids as well as contraption in stapling pyramid.

III. COGNITIVE RELATIONAL INDUCTIVE PREJUDICE

Inductive inference is defined from the view of machine learning as a mechanism to generalize over observed regularities in instances. Each learning algorithm must create some a priori assumptions to allow generalization. This so-called inductive bias provides the rational basis for the transfer to fresh circumstances of the learned hypothesis. The restriction or language bias characterizes the language of representation of induced hypotheses. The bias of preference or search characterizes the hypothesis selection process. Every machine learning algorithm makes such (maybe implicit) assumptions, be it statistical, symbolic or neuronal. Certainly this fundamental proposal of machine learning also applies to human induction. That is, human learning is very strong, but learning is nevertheless limited by the manner in which knowledge can be represented and by some mechanism to prefer some generalization to others.

This paper objective is to claim advanced performance on the stapling challenge. Somewhat, the objective is to characterize the sort of inductive bias required to solve such physical building duties in particular. Our research builds on both the wider cognitive literature on relational reasoning using charts as well as classic methods such as stimulus learning tool and represents a step forward by demonstrating how relational knowledge can be disconnected from physical information through relational strategies approximated by profound neural networks.

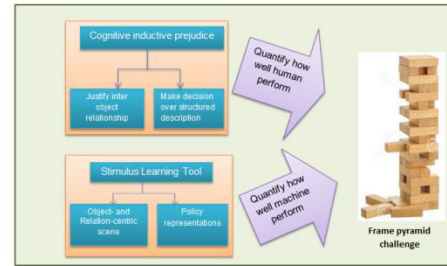


Figure 2: Hominid and machine stapling pyramid quantification

In processes such as object classification, language processing, and gameplay, existing deep learning systems are outstanding, few can create or transform a complex system such as a Frame Pyramid. Assume that the absence of these structures is a "Cognitive Inductive Prejudice": an ability to justify inter-object interactions and make choices on a structured description of the event. To evaluate this assumption, this article focused on a job involving fastening stacks of frames together to balance a castle and quantify how well people do. Then present the Significant Stimulus Learning Tool that uses scene and policy representations centered on objects and interactions, these apply to the assignment. Figure 2 illustrates the structure of how cognitive inductive prejudice and stimulus learning tool quantify the performance of hominids and contraption. Thus detailed explanation of our work is follows.

A. Cognitive Inductive Prejudice:

We enlisted 35 Deep-Mind participants. Each participant gets treated in accordance with UCL Research Ethics Committee protocols and finished 158 tests over an hour session. Within the allocated moment, two respondents failed to finish the assignment and were excluded from the assessment, leaving a total of 33 respondents. In this paper initially establish the stabling challenge, an interactive physical building issue that involves choices about object relationships; then to measure human efficiency in the stabling assignment; after to create a profound RL tool with an object- and relationship-centered representation and action strategy ; and to show the significance of relational inductive bias by comparing the performance of our method with several alternatives, as well as humans, on both the gluing task and several control tasks that isolate different aspects of the full problem.

For each pair of objects, they tried to attach; they lost one point and earned one point for each frame that remained unmoved after the application of gravitational force. As a side benefit, if participants used the minimum amount of glue required to keep the tower stable, 10 extra points are obtained. In reality and experimental sessions, the highest possible results are 191 points and 2077 points, respectively.




Figure 3: Stapling frame challenge

Each test usually consists of two stages: the phase of stapling and the phase of gravitational force. The trial started in the stapling phase, during which for an indefinite amount of time a static tower was displayed on the screen. Participants could pick it by dragging on one item (either a block or ground) and then another object by "stapling" the two together. Only when the two objects were in contact was Glue applied. When glue between the two objects had already been applied, then the glue is eliminated. Both of these actions providing the glue to semi-adjacent objects and ungluing a connection that has already been glued still cost one point. After completed this assignment over 35 participants, in this paper consider the chart system with neural network model for solving optimization issues.

B. Chart Systems With Stimulus Learning Tool:

A main characteristic of our profound RL agent is that it expresses its strategy of decision-making as a function over an object-and relationship-centered state representation, reflecting a powerful inductive relationship bias. In particular, a chart system (CS) is inside the agent, a neural network model that can be trained to approximate graph features. A CS is a generalization of latest neural network approaches to learning physics engines. CSs have been shown to be efficient in solving classic combinatorial optimization issues that inspire our architecture of agents to perform physical building functions.




Figure 4: Different graph representation

Here, a chart is defined as a collection of M nodes, F edges, and a worldwide H function. Nodes correspond to

blocks in the "tower chart" of the gluing task; edges correspond to block pairs; and worldwide characteristics could correspond to any worldwide piece of data, such as the tower's general stability. A CS requires a tower graph as an input and provides the same size and shape of a chart. Conceptual details encode the representation of nodes, edges, and globals: the representation of nodes corresponds to the position (p) and orientation (o), and the edges correspond to the presence of glue (e).

Our system architectures encode the block characteristics in an encoder, into a distributed node representation π_i .

$$m_i = \text{encode}_m(p_i; o_i; \theta_{\text{encodem}}) \quad (1)$$

For an edge ed_{ij} , the edge properties are similarly encoded in a distributed representation using another encoder,

$$ed_{ij} = \text{encode}_{ed}(e_{ij}; \theta_{\text{encodeed}}) \quad (2)$$

The worldwide characteristics are initially empty and set to zero, i.e. $g = 0$. The normal CS calculates features over pairs of nodes (e.g. to determine whether those nodes are in contact), edges (e.g. to determine the force acting on a block) and globals (e.g. to calculate general stability) with these node, edge, and worldwide representations. Specifically, the edge model is represented as ed'_{ij} , the node model is m'_i , and the global model is g' . The above models are derived by,

$$ed'_{ij} = f_{ed}(m_i, m_j, ed_{ij}, g; \theta_{fed}) \quad (3)$$

$$m'_i = f_m(m_i, \epsilon_j ed'_{ij}, g; \theta_{fm}) \quad (4)$$

$$g' = f_g(g; \epsilon_i m'_i, \epsilon_{ij} ed'_{ij}; \theta_{fg}) \quad (5)$$

The CS can be applied multiple times, recurrently, where ed'_{ij} , m'_i and g' are fed in as the new ed_{ij} , m_i and g on the next step.




Figure 5: Chart System

We can decode them into edge-specific projections, such as Q-values or un-normalized log probabilities, considering the modified edge, node, and global representations (Figure 2).

For the supervised setting, edges are glued with probability,

$$p_{ij} \propto \text{decode}_{ed}(ed'_{ij}; \theta_{\text{decodeed}}) \quad (6)$$

For the sequential decision making setting, we decode one action for each edge in the graph

$$\Theta_{ij} = \text{decode}_{ed}(ed'_{ij}; \theta_{\text{decodeed}}) \quad (7)$$

In addition a "stop" action to end the gluing phase is,

$$\Theta_{\sigma} = \text{decode}_g(g'; \theta_{\text{decodeg}}) \quad (8)$$

Applying the CS to calculate conditions of communication and update the nodes over and over again can be described as passing messages, which propagate data across the chart. Such acquired propagation of data in the stapling

assignment can parallel the propagation of forces and other structural limitations. Consider the tower in Figure 3 for observation. The CS should be able to determine which block pairs are locally volatile, such as the top-most block in the figure, after one implementation of the edge model, and therefore require glue. Though, it does not have sufficient data to decide that the bottommost block in Figure 3 often requires to be attached because it completely supports the above block. Frequent message-passing enables for the propagation of data about other blocks to the bottom-most, enabling for reasoning about non-local relationships. CS performs the solution for optimal problem, and then this paper analyzed stability of learning followed.

1) Stimulus learning:

Until considering the complete stapling assignment, we first analyzed how important sub-tasks could be performed in a controlled environment by parts of the graph network agent, such as anticipating stability or inferring which edges to glue. We used towers with variable amount of blocks to evaluate the stability projections of the CS, where the input edges were marked to show whether or not there was glue (1 for glue, 0 for no glue). For each scene, the glue was sampled randomly, and stability is described as not dropping blocks. Two settings were tested: completely linked graphs (where all block-to-block edges were included in the chart), and sparse graphs (where edges only existed between contact blocks). In both instances, CSs have learned to predict the stability of partly glued towers correctly, but sparse graph inputs have resulted in more effective learning. Results for 5 blocks are shown, but these results are consistent across 6-9 blocks towers as well. We also evaluated whether CSs could learn that a contact should be glued between two blocks. Some glue places, as mentioned earlier, involve reasoning as to how forces propagate throughout the structure. We hypothesized therefore that numerous message passing steps would be essential to propagate this data, and indeed we discovered that one recurrence was sufficient to dramatically enhance the precision of glue forecast. In this paper clearly shows how deep stimulus learning tool can be enhanced by embracing relational inductive biases such as those in human cognition, and human behaviors like physical scene building and interaction.

IV. RESULT AND DISCUSSION:

For the human participants, the stapling challenge was complicated, but they still performed far beyond chance. We also discovered so many developments in the behavior of people, such as working from top to bottom and spending more time before applying the first glue than before applying the following glue. The outcomes here represent provisional experimentation of people's behavior in construction tasks. Participants achieved an average score of 1025 points, with the lowest score being 534 points and the highest score being 1578 points (out of 2004). There was a low (although not quite important)

impact of learning, with a Pearson correlation of $r = 0.15$; 96% CI [0:01; 0:30] between the trial number and the average scaled reward (confidence intervals were calculated around the median using 12000 substitute bootstrap samples; 'scaled rewards' are calculated by standardizing rewards such that 0 was the reward received if no action was found.

Table 1: Accuracy Prediction For Our Work

Full (without contact information)	Sparse (with contact information)	Accuracy
1025	1350	0.75
1200	1400	0.79
1900	2100	0.8
2200	2900	0.85
2500	3000	0.89
2800	3800	0.9
3200	4000	0.91
4800	5000	0.92
5600	6000	0.93
6700	7000	0.95
7000	8000	0.96

Table 1 describes the accuracy prediction for different number of recurrent steps. Here eleventh type of performance attains the 96% of accuracy and without contact information as 7000 as well with contact information as 8000.

The reaction times of the participants disclosed that clicking on the first block in a couple was considerably slower than the second block, with a distinction of $t = 3.58s$; 96 percent CI [4:34s;4:62s]. This indicates that before pressing the first block, they had decided on which pair to glue. We estimated that individuals were considerably slower in choosing the first gluing action ($t = 5.49s$; 96% CI [4:30s; 4:56s]; averages calculated using the log RTs mean) than any later gluing action ($t = 2:07s$; 96% CI [2:00s; 2:15s]; $F(1; 12878) = 149.14, p < 0.001$). We also discovered an impact of block amount on response time ($F(1; 12878) = 401.58, p < 0.001$) as well as an interaction between block amount and whether or not the action was the first glue ($F(1; 12878) = 13.78, p < 0.001$), with the first action requiring more time per block than subsequent actions.

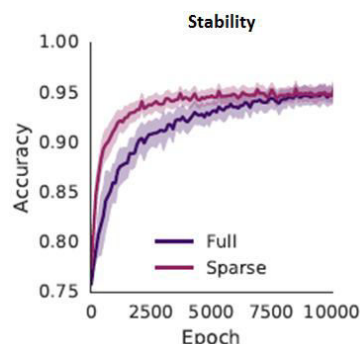


Figure 6: Stabilization for response charts with (sparse) or (full) contact data

Before researching the complete stapling assignment, in this paper initially analyzed, how important sub-tasks could be performed in a controlled environment by parts of the graph network agent, such as anticipating stability or inferring which edges to glue.

We used towers with a variable amount of blocks to evaluate the stability projections of the CS, where the input edges were marked to show whether or not there was glue (1 for glue, 0 for no glue). For each part, the glue was sampled randomly, and stability was described as not dropping blocks. Figure illustrates the completely linked graphs (where all block-to-block edges were included in the chart), and figure describes the sparse graphs (where edges only existed between contact blocks).

CSs have learned to estimate the stability of partly glued towers correctly in both instances, but the sparse graph inputs have resulted in more effective learning. It is denoted in figure 6. Results for 5 blocks are shown, but these findings are consistent across 6-9 blocks towers as well. We also assessed whether CSs could learn that a contact should be placed between two blocks. Some glue places, as mentioned earlier, involve reasoning as to how forces propagate throughout the structure.

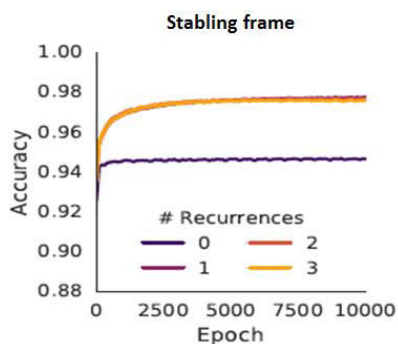


Figure 7: Estimate of glue for systems with distinct numbers of recurring steps.

In this paper, observed therefore that numerous message-passing steps would be essential to propagate this data, and indeed we discovered that one recurrence is sufficient to dramatically enhance the precision of glue forecast. Figure 7 describes optimal glue prediction for models with different numbers of recurrent steps.




Figure 8: Stapling level curve depends on frequency and participants

In figure 8 illustrates the stapling level of pyramid, in which depends on the frequency of movement blocks. These results indicate that individuals can either decide where to place glue until acting or at least engage in a costly encoding operation of the useful simulative representation. 10 out of 35 respondents reported making glue choices top-to-bottom on an open-ended strategy issue in the post-experiment analysis, and another 3 reported sometimes working top-to-bottom and sometimes bottom-to-top. We experimentally corroborated this by adding a line between the action amount and the height of the glue place for each trial and finding that their slopes were usually negative ($b = 0.06$; 96% CI [either 0.07; or 0.05]).




Figure 9: Time interval between the constructions of towers

Figure 9 shows that the time elapsed between tower constructions did not exceed 100 seconds for the first 28 towers. The elapsed time, however, differed significantly after that stage, with the maximum being 783 seconds. The average has been 28.86 seconds for the first 28 towers and the average is 20.5 seconds. The average is 120.45 seconds for the last eight towers and the average is 58.5 seconds.

We compared people's choice of glue set up to ideal glue settings and discovered that individuals were much more likely to apply glue when it was not essential (64% of mistakes) than when it was essential to not apply glue ($N = 4201$, $p < 0.001$). In addition, respondents are very good at preventing invalid activities: although they had the choice to collect pairs of blocks that were not in touch, they did so only 0.7 percent of the moment (out of $N = 7896$). Similarly, respondents did not commonly use the option to un-glue blocks (0.14% out of $N = 7896$), probably because a penalty was incurred.

A. Comparison Analysis:

We considered three agents: a multilayer perceptron (or MLP) agent, a fully-connected graph network (or GN-FC) agent, a chart system (or CS) agent, and a simulation agent. Since most deep RL agents are either implemented as MLPs or CNNs with no relational structure, our first agent chose actions based on a Q-function approximated by an MLP; since MLPs have a fixed input and output size, we chose actions based on a Q-function approximated by an MLP; The agents of CS and GN-FC (who had related expertise but no specific physical

expertise) also selected activities based on a Q-function and used 3 recurrent measures.

Table 2: Comparison Of Overall Reward For Humans And Agents

Rewards	H	MLP	GN-FC	GN	proposed
1	0.9	1	1	1	1
2	0.7	0.95	0.9	0.98	1
3	0.6	0.85	0.6	0.95	1
4	0.5	0.75	0.85	0.92	1
5	0.4	0.65	0.65	0.9	0.99
6	0.38	0.55	0.5	0.88	0.98
7	0.35	0.45	0.4	0.86	0.95
8	0.33	0.35	0.38	0.85	0.88
9	0.2	0.1	0.3	0.83	0.85
10	0.15	0.05	0.2	0.7	1

In table 2 describes the comparison of different agents with our proposed agent. H: human; MLP: MLP agent; GN-FC: GN agent operating over a fully-connected graph; GN: GN agent operating over a sparse graph; Sim: simulation agent.

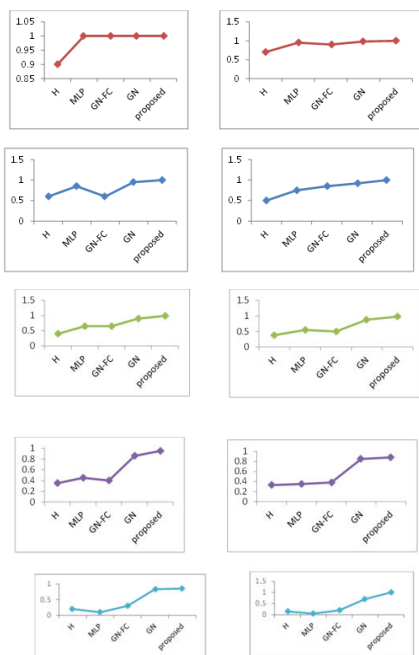


Figure 10: Comparison graph between humans and contraption from reward 1 to 10

The CS agent has been using a sparse graph framework with edges corresponding to the block contact points, while the GN-FC used a completely linked graph framework and had to know which edges corresponded to valid actions. Finally, using simulation, the simulation officer (who had both relational and physical understanding) chose activities. Specifically, the agent conducted a simulation for each unglued contact point to calculate how many blocks would drop if that point were glued, then selected the point that caused the fewest blocks to drop. Until no blocks fell, this operation was repeated. Note that the simulation agent is not optimal because it greedily selects glue points.




Figure 11: overall comparison graph for total reward between human and contraption

While our goal isn't to build a model of human cognition on the stapling challenge, in our work compared the behavior of people to the behavior of the CS agent in order to clarify any obvious differences. The average reward of the participants dropped between the agents of MLP and GN-FC. It is illustrated by figure 11.

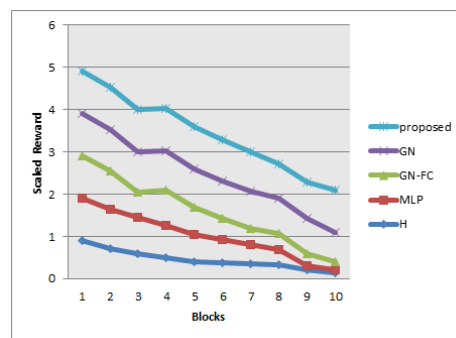


Figure 12: Comparison of scaled reward

As in Figure 12 Scaled reward comparison across towers of various dimensions, it is increasingly difficult for both agents and humans to solve the task as a function of tower size, although this is expected: as the number of blocks in the tower increases, the number of possible glue combinations increases exponentially. Rewards are scaled so that 0 corresponds to the reward when no action is taken and 1 corresponds to the ideal reward.

This paper realized that within scenes the CS agent made distinct patterns of mistakes than humans. For instance, while our work observed that individuals were more likely to create false positives (applying glue when none was required), this paper did not consider this true for the CS agent (38% of mistakes, $N=142$, $p<0.05$). This distinction could lead from human perceptual uncertainty, which leads to a tendency to overestimate tower instability.

V. CONCLUSION

Research on induction is becoming deeper and broader over the past decade. Through establishing more extensive and advanced computational models, a deeper understanding of induction has been accomplished. It has revealed unique dimensions of induction that are missing from previous theories, such as the role of information analysis. In this paper a novel solution is presented for building the "stapling challenge", it involves stapling the pairs of blocks to stabilize a block tower. Our experiment outcomes reveals that an agent with an object- and relationship-centered strategy is able to address the assignment even better than humans, while an agent does not have such relational inductive prejudice is doing even worse.

VI. REFERENCES

- [1] P.W. Battaglia, J.B. Hamrick, & J.B. Tenenbaum, Simulation as an engine of physical scene understanding, *Pnas*, 110(45), 2013, 18327–18332.
- [2] A.M. Collins, & E.F. Loftus, A spreading-activation theory of semantic processing, *Psychological Review*, 82(6), 1975.
- [3] H. Dai, E.B. Khalil, Y. Zhang, B. Dilkina, & L. Song, Learning combinatorial optimization algorithms over graphs, *Nips*, 30, 2017.
- [4] M. Denil, P. Agrawal, T.D. Kulkarni, T. Erez, P.W. Battaglia, & N. de Freitas, Learning to Perform Physics Experiments via Deep Reinforcement Learning, *ICLR*, 2017.
- [5] S. Dzeroski, L. De Raedt, & K. Driessens, Relational reinforcement learning, *Machine Learning*, 43(1-2), 2001, 7–52.
- [6] D. Gentner, Structure-mapping: A theoretical framework for analogy, *Cognitive Science*, 7(2), 1983, 155–170.
- [7] T. Gerstenberg, N.D. Goodman, D.A. Lagnado, & J.B. Tenenbaum, Noisy Newtons: Unifying process and dependency accounts of causal attribution, *In Cogsci*, 2012.
- [8] J.B. Hamrick, A.J. Ballard, R. Pascanu, O. Vinyals, N. Heess, & P.W. Battaglia, Metacontrol for adaptive imagination based optimization, *ICLR*, 2017.
- [9] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, D. Hassabis, Human-level control through deep reinforcement learning, *Nature*, 518(7540), 2015.
- [10] F. Scarselli, M. Gori, A. Tsoi, M. Hagenbuchner, & G. Monfardini, The graph neural network model, *IEEE Transactions On Neural Networks*, 20, 2009, 61–80.
- [11] M.B. Chang, T. Ullman, A. Torralba, and J.B. Tenenbaum, A compositional object-based approach to learning physical dynamics, *arxiv preprint arxiv:1612.00341*, 2016.
- [12] W. Li, A. Leonardis, and M. Fritz, Visual stability prediction for robotic manipulation, *In 2017 IEEE International Conference On Robotics And Automation (ICRA)*, 2017, 2606-2613.
- [13] M. Denil, P. Agrawal, T.D. Kulkarni, T. Erez, P. Battaglia, and N. de Freitas, Learning to perform physics experiments via deep reinforcement learning, *Arxiv Preprint Arxiv:1611.01843*, 2016.
- [14] I. Yildirim, T. Gerstenberg, B. Saeed, M. Toussaint, and J. Tenenbaum, Physical problem solving: Joint planning with symbolic, geometric, and dynamic constraints, *Arxiv Preprint Arxiv:1707.08212*, 2017.
- [15] T. Wang, R. Liao, J. Ba, and S. Fidler, Nerve net: Learning structured policy with graph neural networks, 2018.
- [16] K. Fragkiadaki, P. Agrawal, S. Levine, and J. Malik, Learning visual predictive models of physics for playing billiards, *Arxiv Preprint Arxiv:1511.07404*, 2015.
- [17] P. Liang, X. Jia, N.A. Taatgen, J.P. Borst, and K. Li. Activity in the fronto-parietal network indicates numerical inductive reasoning beyond calculation: An fMRI study combined with a cognitive model, *Scientific Reports*, 6, 2016, 25976.