Human causal reasoning guided autonomous object assembly under uncertainty

Semanti Basu[1], Semir Tatlidil[2], Moon Hwan Kim [1], Steven Sloman[2], and R. Iris Bahar[1,3]
[1]Brown University, Dept. of Computer Science, Providence, RI 02912
[2]Brown University, Dept. of Cognitive, Linguistic, and Psychological Sciences, Providence, RI 02912
[3]Colorado School of Mines, Dept. of Computer Science, Golden, CO 80401

ABSTRACT

In this paper we explore robot planning under partial observability with the guidance of human-generated causal reasoning. We hypothesize that human causal models, even when imperfect, can provide valuable cues that can improve a robot's decision making under uncertainty. To that end, we develop (1) an interface to collect causal models from human participants, and (2) a method to embed collected models as imperfect priors in a Partially Observable Markov Decision Process (POMDP). We further theorize that causal models can be a generalizable prior that can transfer to related tasks. We evaluate our methods on the problem of object assembly (framed as a POMDP) of different kinds of objects. We note over 2X improvement in reward on average when using user human generated causal models as priors. When transferring causal models to other objects, we note an average improvement of 1.25X in rewards.

CCS CONCEPTS

• **Computing methodologies → Cognitive robotics**.

KEYWORDS

Causal reasoning, planning under uncertainty, object assembly

1 INTRODUCTION AND RELATED WORK

Studies on human thinking and reasoning have demonstrated that people construct their mental models using causal information. For instance, causal knowledge affects how people categorize objects and update their beliefs [10, 12]. Experts tend to categorize systems based on causal information rather than surface-level features such as domain similarity [9]. People also rely on prior causal knowledge to make predictions in novel environments [5].

Human decision-making under uncertainty is often guided by causal reasoning. When complete knowledge about a task is not available, causal reasoning is often used by people to decide on potential actions that could lead to a successful outcome. When an action is taken, the person will observe its effect on the environment, which validates or refutes their original reasoning. This observation is then used to re-orient the person's thoughts before taking the next step.

In robotics, problems that require planning under uncertainty are often modeled as Partially Observable Markov Decision Processes (POMDPs). The uncertainty is represented as a probability distribution over system states, known as the belief. Solving a POMDP exactly is intractable, so Monte Carlo sampling based online planning methods are a popular approach to solving them. Using these methods, the agent takes an action in the environment, receives an observation, and uses it to update its belief.

We hypothesize that making generalized human causal knowledge accessible to robots can significantly improve specific planning outcomes in partially observable domains (represented as POMDPs). In this paper, we propose a causally enhanced approach to object assembly under partial observability. We describe our process of obtaining causal models from human participants and embedding them in an autonomous agent. We further explore generalization of causal models to assist assembly of new objects. We evaluate our approach on four different objects of varying complexity in simulation.

Human robot collaboration in object assembly has been studied in the context of Human intention tracking [1] [3], Task allocation [8] [4] and Learning from demonstrations [7]. In these approaches the robot and the human share a workspace to assemble an object together. Our work critically differs from these approaches by decoupling the human and robot contribution. The human contributes a mental model of the object by providing a causal graph, but does not physically participate in assembly. The robot is solely responsible for assembling an object but gets to leverage human causal reasoning for smarter decision making.

2 METHODS

2.1 Obtaining Causal Models from Human Participants

We showed participants a diagram of a light producing object, such as a desk lamp, alongside a table listing the parts and their functions, and told participants they would be tested on their understanding of how the object worked. Next, the participants were asked a series of questions in the form of: *"Remove [Part X]. Would the [Part Y] still perform [function of Y]?"*, for all pairs of object parts. The participants gave a binary answer to these interventional questions. If the answer is "No" that indicates a causal link from Part X to Part Y. The survey had 100 participants each of who entered causal information on 4 different objects.

2.2 Modeling furniture assembly as a POMDP

The goal of our POMDP planner is to discover which parts of an object should be connected together, given no prior information about geometry or part-compatibility. A piece of furniture with n parts can have a maximum of $\binom{n}{2} = \frac{n!}{2!(n-2)!}$ connections.

The agent cannot observe the shape or size of the parts or how many connections there are in the particular assembly. However, at each time step it is allowed to try to connect any two parts together. If the parts fit, it receives a positive observation, otherwise it receives a negative observation. It is illegal to try to connect the same combination of parts more than once. It receives a positive reward once all the connections have been discovered and a negative reward for every time step. The goal is to discover all part connections for a particular object in as few steps as possible.

We model our POMDP using the tuple $\langle S, A, \Omega, \text{BlackBoxModel}, d\rangle$, where the state space (S) describes the state of assembly for an object with n parts. There are $\binom{n}{2} = \frac{n!}{2!(n-2)!}$ unique part combinations if there are n parts (see Fig. 1). The state space describes the status of each of these part combinations. It has a binary indicator (α) for each combination to indicate whether the robot has tried to connect those two parts together yet or not. There is another binary indicator (β) per combination that describes whether or not that is a valid combination that should be part of the final assembly. The robot cannot observe the value of β unless it is trying that particular part combination. There is another variable per state (γ), which keeps track of the number of correct connections left to be found by the robot. The robot cannot directly observe the value of γ but can observe whether it is 0 or not. It terminates when $\gamma = 0$. The action space (A) has a size of $\binom{n}{2} = \frac{n!}{2!(n-2)!}$ for an n part object. It is essentially the set of unique combinations of two object parts that the agent can try as it assembles an object. The observation space (Ω) is binary — it is obtained upon trying a particular combination and observing whether it is a correct connection or not (reading the value of β). The BlackBoxModel is a way to define the successor state, observation, and reward based on the current state and action, without explicitly defining the transition, observation and reward model for the POMDP. In our case, the agent visits the connection defined by the action and the α value of that connection changes to 1 to mark it as a tried combination. An observation is then read, which is defined by the value of β for that connection. If the observation is positive, then the number of connections left to be discovered (γ) is updated. A negative reward is awarded for every time step. However, if the terminal state has been reached, then the robot receives a positive reward of $\binom{n}{2} = \frac{n!}{2!(n-2)!}$. The reward quantifies the number of steps saved by the agent. Hence, if the agent had to try every possible combination to find the correct ones, then it will receive a reward of 0.

$$
\begin{array}{c}
\begin{array}{cc}
 & \begin{array}{cccc} \quad BackPlate & Lampbody & Lightbulb & Cord \end{array} \\
\begin{array}{c} BackPlate \\ Lampbody \\ Lightbulb \\ Cord \end{array} &
\begin{bmatrix}
\cdots & \alpha_1, \beta_1 & \alpha_2, \beta_2 & \alpha_3, \beta_3 \\
\cdots & \cdots & \alpha_4, \beta_4 & \alpha_5, \beta_5 \\
\cdots & \cdots & \cdots & \alpha_6, \beta_6 \\
\cdots & \cdots & \cdots & \cdots
\end{bmatrix}
\end{array} \\
\gamma = k, 0 < k <= 6
\end{array}
$$

Figure 1: Sample nonterminal state for an object with 4 parts

2.3 POMDP Belief state modeling with causal information

The belief state in a POMDP is a probability distribution of system states at a particular time. In our case, it is represented using an unweighted particle filter. We use a set of Z particles, where each particle represents a sample state as defined in Fig. 1 to approximate the belief. Initially, none of the connections have been tried, so $\alpha = 0$ for all possible connections. The value of β is unknown. In the absence of any prior information about the connections, β is randomly set to 0 or 1. However, we have user generated causal graphs which we use as priors to inform the initial belief state of the POMDP, which essentially gives the agent a hint on the β values in the real assembly (see Algorithm 1). Object parts that are connected in the causal model are more likely to be physically connected in assembly, however they are not guaranteed to be. The initial belief is designed so that the majority of the states in the initial belief considers a causal connection to be a true physical connection in assembly. The β variable associated with a certain connection is modeled as a Bernoulli distribution that takes the value of 1 with probability of $conf$ if the connection if causal, otherwise takes the value of 1 with probability $1 - conf$. The value of $conf$ can be anything greater than 0.5 and can be set by the user — a higher value indicates greater trust in the probability that a causal relationship implies a physical connection.

Algorithm 1 Causal_Initial_Belief_State()

```
Z ← Number of particles used to approximate belief
B ← Initial belief
A ← Locations of possible unique connections
C ← Locations of causal connections
for i<=Z do
    P_i ← empty n x n matrix
    for a in A do
        if a ∈ C then
            P_i[a].α = 0
            P_i[a].β = Bernoulli(conf)
        else
            P_i[a].α = 0
            P_i[a].β = Bernoulli(1 − conf)
        end if
        if P_i[a].β==1 then
            P_i[a].γ = P_i[a].γ + 1
        end if
    end for
    B ← B ∪ P_i
end for
```

2.4 Generalizing causal models for transfer

We investigate if our method of causal model inference generalizes to new objects. We seek to answer the following question: If we have obtained causal models from human users for certain objects, can we leverage that data to automatically generate a causal model for a new object without involving a person?

To that end we reframe our interventional approach to causal model extraction as a text classification problem. A dataset is prepared from the human causal data obtained from the survey. Each interventional question asked in the survey coupled with the context (object and part descriptions) forms a single query in our dataset with two soft labels Yes and No. The value of the labels is the percentage of people who answered Yes and No to that question, respectively.

We fine-tune a custom pre-trained BERT Model [2] on the downstream task of probabilistic classification. Our custom model consists of a BERT pretrained model (bert-base-uncased), followed by a dropout layer and a linear layer to map it to our outputs. Cross-entropy loss was used for training. We used guidelines for stable training from [6] to set parameters for training and early stopping to prevent overtraining. Since we collected data for 4 objects, we used queries corresponding to 3 objects for training and tested our model on the 4th object. We performed cross validation to understand transfer across different types of objects. For the unseen test objects all pertinent interventional questions were asked of the trained model and the output probabilities were thresholded at 0.5. If the output probability of the label "No" exceeded 0.5, it was considered a causal link otherwise not. After asking all questions as described in Section 2.1 we obtained the transferred causal model for the test object.

3 EXPERIMENTAL RESULTS

3.1 Evaluating the quality of causal models

We generated our own causal models for each object, which we refer to as the expert model. We compared the causal models obtained from the participants to the expert causal models for each object. We used F1 score as our metric to assess the model, in order to combine precision and recall into a single metric. The F1 score is calculated as:

$$F1\_score = \frac{2 \cdot TruePositives}{2 \cdot TruePositives + FalsePositives + FalseNegatives}. \quad (1)$$

While F1 score is mostly used for predictive performance of classification models, it is also applicable for evaluating user causal models. Each causal link is treated as a prediction from the user — if the user "detected" a link present in the expert model, then it is a true positive. If the user detected a link to be causal when it is actually not, the link is treated as a false positive. If the user failed to detect a causal link that is present in the expert causal model, then it is considered a false negative.

Similarly for generalization, the causal models obtained from our fine-tuned BERT Model were compared to the expert models to evaluate their quality. We also asked our survey questions to Chat-GPT and evaluated the models obtained. In Fig. 3, we show the quality of causal models obtained through different processes across all 4 objects. We compare causal models obtained from 3 sources — user study, trained BERT Model and ChatGPT. The average F1-score for user models are plotted as well as the F1 score of the causal model obtained by aggregating all user answers for that object. It can be seen that the human-generated user models are of better quality that those obtained from ChatGPT or through transfer using BERT.

3.2 Evaluating the efficacy of user Causal models as priors for planning under partial observability

We next evaluate the effect of embedding causal models as priors in the object assembly POMDP. We solve our POMDP using Partially Observable Monte Carlo Planning(POMCP)[11]. We run the search with 16384 Monte Carlo simulations, and each data point is the average of 100 runs of the algorithm. We plot the reward obtained (expressed as a percentage) against causal model $F1\_score$ ranges (embedded as priors to obtain the said reward) for 4 different objects

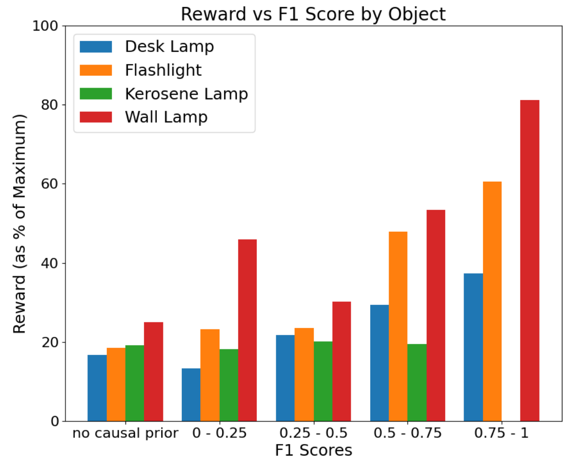

Figure 2: $F1\_score$ of causal models vs. reward .

(Fig. 2). It can be seen that the reward obtained steadily increases as the causal model quality increases, for all objects. Even for imperfect causal models, the reward obtained is better than starting without a prior. The causal prior results in a 1.8X improvement for desk-lamp, 2.7X improvement for flashlight, 1.02X for kerosene lamp and 2.73X for wall lamp.

3.3 Evaluating generalization of causal models for object assembly

We explored generalization by leveraging our user-data from the survey to fine-tune BERT. We train on survey data from three objects and test the ability of the model to generalize to the fourth object. The causal model for the unseen object obtained from the trained model is referred to as the "Transferred Model". In Fig. 4 we compare the effect of the transferred model in object assembly compared to user models, expert models and ChatGPT models. We can see that the transferred model is beneficial as a prior and improves reward over no causal prior for flashlight and wall lamp by 2.55X and 1.41X respectively. However, for desk-lamp and kerosene lamp it is misleading as a prior and worsens performance. In Fig. 3, it can be seen that the quality of the transferred models falls short of user models by a large margin, and is also worse than ChatGPT. These suggest that our method of generalization is not robust, and ChatGPT without any fine-tuning outperforms a fine-tuned BERT model on our causal-reasoning dataset. One of the reasons could be that our dataset is extremely small consisting of only 93 samples overall. Of these, all queries pertinent to the test object were not used in training. However, it also raises the question of whether or not language models can handle causal reasoning. It is interesting to note that ChatGPT had worse quality causal models than user models but performed similarly to them on the downstream task of assembly. This suggests that ChatGPT-generated models reflected physical causality, which helped with assembly, but weren't the goal of our method of causal model extraction. It is likely that models generated with ChatGPT were a product of its rigorous training (which likely had a lot of data on a common task like assembly) rather than reasoning.

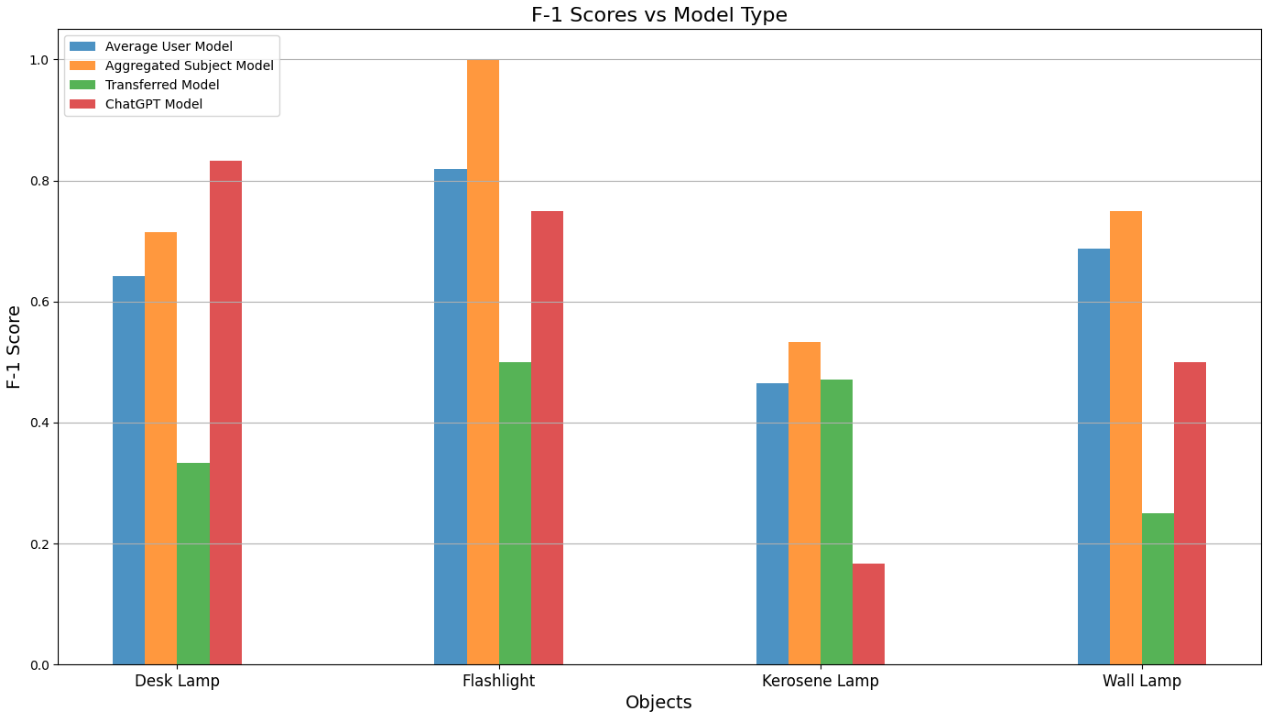

Figure 3: $F1\_score$ **of causal models for different objects.**

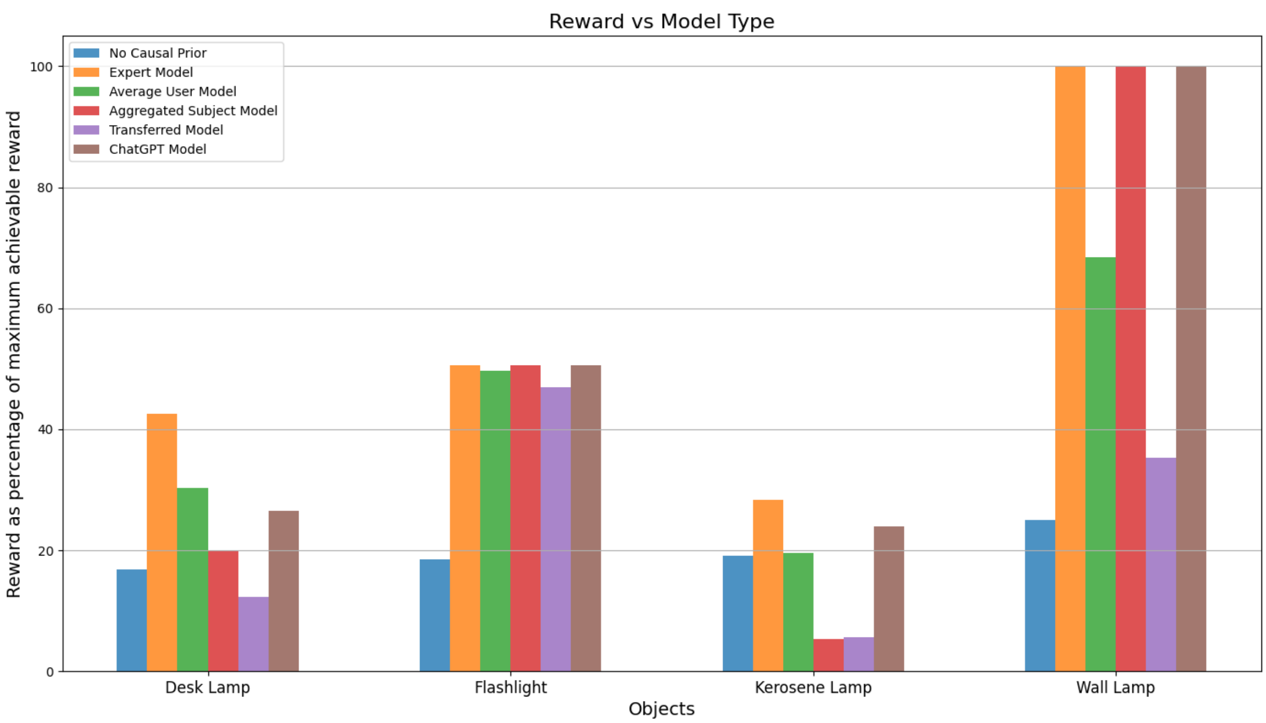

Figure 4: Comparison of Transfer models

REFERENCES

[1] Martijn Cramer, Karel Kellens, and Eric Demeester. 2021. Probabilistic Decision Model for Adaptive Task Planning in Human-Robot Collaborative Assembly Based on Designer and Operator Intents. *IEEE Robotics and Automation Letters* 6, 4 (2021), 7325–7332. https://doi.org/10.1109/LRA.2021.3095513

[2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[3] Zhe Huang, Ye-Ji Mun, Xiang Li, Yiqing Xie, Ninghan Zhong, Weihang Liang, Junyi Geng, Tan Chen, and Katherine Driggs-Campbell. 2023. Hierarchical Intention Tracking for Robust Human-Robot Collaboration in Industrial Assembly Tasks. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 9821–9828. https://doi.org/10.1109/ICRA48891.2023.10160515

[4] Lars Johannsmeier and Sami Haddadin. 2017. A Hierarchical Human-Robot Interaction-Planning Framework for Task Allocation in Collaborative Industrial Assembly Processes. *IEEE Robotics and Automation Letters* 2, 1 (2017), 41–48. https://doi.org/10.1109/LRA.2016.2535907

[5] Yunnwen Lien and Patricia W Cheng. 2000. Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology* 40, 2 (2000), 87–137. https://doi.org/10.1006/cogp.1999.0724

[6] Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2020. On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. *arXiv preprint arXiv:2006.04884* (2020).

[7] Heramb Nemlekar, Runyu Guan, Guanyang Luo, Satyandra K. Gupta, and Stefanos Nikolaidis. 2022. Towards Transferring Human Preferences from Canonical to Actual Assembly Tasks. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1161–1167. https://doi.org/10.1109/RO-MAN53752.2022.9900872

[8] Alessandro Roncone, Olivier Mangin, and Brian Scassellati. 2017. Transparent role assignment and task allocation in human robot collaboration. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 1014–1021. https://doi.org/10.1109/ICRA.2017.7989122

[9] Benjamin M. Rottman, Dedre Gentner, and Micah B. Goldwater. 2012. Causal Systems Categories: Differences in Novice and Expert Categorization of Causal Phenomena. *Cognitive Science* 36, 5 (2012), 919–932. https://doi.org/10.1111/j.1551-6709.2012.01253.x

[10] Benjamin Margolin Rottman and Reid Hastie. 2014. Reasoning about causal relationships: Inferences on causal networks. *Psychological bulletin* 140, 1 (2014), 109.

[11] David Silver and Joel Veness. 2010. Monte-Carlo Planning in Large POMDPs. In *Advances in Neural Information Processing Systems*, J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta (Eds.), Vol. 23. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2010/file/edfbe1afcf9246bb0d40eb4d8027d90f-Paper.pdf

[12] Steven A Sloman and David Lagnado. 2015. Causality in thought. *Annual review of psychology* 66 (2015), 223–247.

Received 9 February 2024