

# Empirical Estimates on Hand Manipulation are Recoverable: A Step Towards Individualized and Explainable Robotic Support in Everyday Activities

Alexander Wich  
Institute for Artificial Intelligence,  
University of Bremen  
Bremen, Germany  
awich@uni-bremen.de

Holger Schultheis  
Institute for Artificial Intelligence,  
University of Bremen  
Bremen, Germany  
schulth@uni-bremen.de

Michael Beetz  
Institute for Artificial Intelligence,  
University of Bremen  
Bremen, Germany  
beetz@cs.uni-bremen.de

## ABSTRACT

A key challenge for robotic systems is to figure out the behavior of another agent. The capability to draw correct inferences is crucial to derive human behavior from examples.

Processing correct inferences is especially challenging when (confounding) factors are not controlled experimentally (observational evidence). For this reason, robots that rely on inferences that are correlational risk a biased interpretation of the evidence.

We propose equipping robots with the necessary tools to conduct observational studies on people. Specifically, we propose and explore the feasibility of structural causal models with non-parametric estimators to derive empirical estimates on hand behavior in the context of object manipulation in a virtual kitchen scenario. In particular, we focus on inferences under (the weaker) conditions of partial confounding (the model covering only some factors) and confront estimators with hundreds of samples instead of the typical order of thousands. Studying these conditions explores the boundaries of the approach and its viability.

Despite the challenging conditions, the estimates inferred from the validation data are correct. Moreover, these estimates are stable against three refutation strategies where four estimators are in agreement. Furthermore, the causal quantity for two individuals reveals the sensibility of the approach to detect positive and negative effects.

The validity, stability, and explainability of the approach are encouraging and serve as the foundation for further research.

## KEYWORDS

causal inference; treatment effect estimation; robotics; human behavior

### ACM Reference Format:

Alexander Wich, Holger Schultheis, and Michael Beetz. 2022. Empirical Estimates on Hand Manipulation are Recoverable: A Step Towards Individualized and Explainable Robotic Support in Everyday Activities. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 9 pages.

*Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.



**Figure 1: A robot observes a person set the table in a virtual kitchen scenario: Why does this person sometimes choose the left hand over the right? Does the proximity of the grasp play a role in hand-selection? And if so, why?**

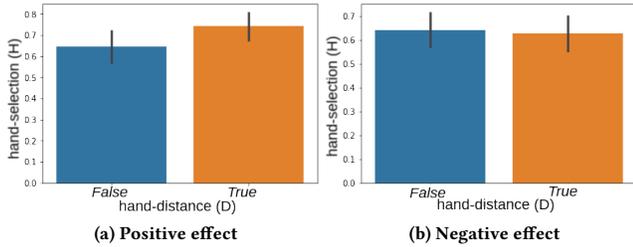
## 1 INTRODUCTION

The potential benefits for robots observing human behaviors are twofold. Having robots adopt people’s know-how has the potential to incorporate key knowledge for robust control systems in workspaces that are made by and for human beings. In turn, robots modeling human behavior are in a better position to anticipate a person’s intention, thereby improving assistance in joint tasks. In the same manner, robots adopting human behaviors act predictably for their human counterpart.

At the age of two toddlers already display capabilities of reasoning over cause and effect [15]. Inferring the cause of effects is a common capability in children, which enables them to understand their surroundings by simple observation, including the actions of other social agents [33]. Causal reasoning enables children to observe the actions of other social agents and learn causal relationships from probabilistic events [35].

We envision robots able to acquire knowledge simply by observing their surrounding as illustrated in Figure 1. To enable observational capabilities in robots that resemble those of children, requires mechanizing inferences that are robust against bias.

Ignoring confounding bias (a.k.a. reversals) can invalidate inferences in any domain [32]. For instance, reversals affect a random uniformly distributed  $2 \times 2 \times 2$  table with a probability of  $1/60$  [25]. The probability that reversals would occur at random in path models involving two predictors and one criterion variable is approximately ten percent [20].



**Figure 2: Given identical data, Figure 2a and 2b show opposing conclusions when segregating and aggregating trends. This confirms the existence of reversal in our data and motivates the effort to account for bias. This data is further described in Table 1.**

The data considered in this work is no exception to this. Figure 2 shows two opposing conclusions that are drawn on identical data. The disagreement of segregated and aggregated trends confirms the existence of reversal in the data considered here. Confounding can be confronted with a model as suggested in [27] as approached here. Section 2.2 further describes the data.

The main proposition of this paper is for the robots to, literally, conduct observational studies. We propose causal inference to equip robotic agents with the necessary observational skills. The central question raised in this article is whether robots equipped with capabilities to perform observational studies extract reliably and transparently information from complex behavioral examples.

To this end, we take on the task to recover a non-trivial relationship on a few examples of hand-manipulation originating from uncontrolled sources of evidence. This is essentially the challenge children overcome when looking over to role models.

The steps required to equip robots with the necessary observational skills are introduced in Section 2.1 to 2.3. Subsequently, the validity of inferences are verified (Section 3) and the stability of effects tested (Section 4). Then, inferences target evidence from two individuals whom are not influenced by any experimental design (Section 5). Last, individual insights on hand behavior are then presented (Section 6).

## 2 DRAWING CAUSAL CONCLUSIONS ON OBSERVATIONAL EVIDENCE

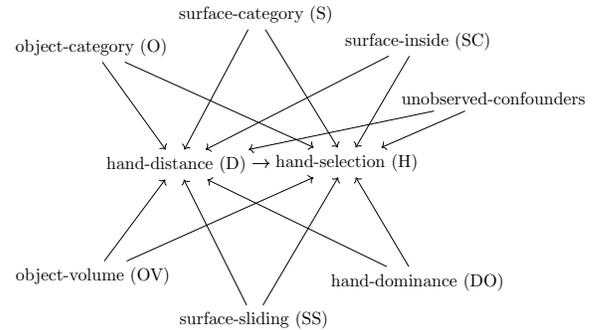
One of the necessary capabilities to derive behavior from evidence is for robots to draw causal conclusions. Because causal inferences mitigate bias, conclusions are closer to the truth than the correlational counterpart.

Drawing causal conclusions on uncontrolled evidence involves three steps. Step 1 defines a model for hand manipulation (described in Section 2.1), step 2 introduces the evidence on which inferences are drawn (Section 2.2), and step 3 formalizes causal reasoning (Section 2.3).

### 2.1 Step 1: Defining the Model

A graph encodes the model of hand manipulation as shown in Figure 3. In this diagram the nodes represent potential factors that

affect hand manipulation, while arrows indicate the direction these factors affect each other. For example, the arrow pointing from hand-distance to hand-selection states that the former affects the latter. Here, the choice of hand-selection refers to the hand, either left or right, a person grasps an object with. Similarly, the choice of hand-distance, either close or far from their body, refers to the straight-line distance between the head and hand at the time the person triggered a grasp.



**Figure 3: A causal graph modeling reach-to-grasp actions including two hand decisions along other relevant factors. The validation hypothesis of this work corresponds to the arrow directed from hand-distance to hand-selection.**

The arrows included in the causal graph are described as follows:

- (1) **Hand-distance on hand-selection:** Hand selection occurs according to object proximity. For example, when targeting an object for which the left hand is closer than the right, the left hand is chosen. This behavior is supported by the kinesthetic hypothesis [11, 12].
- (2) **Object on hand-distance:** The category of an object influences the distance a person would stretch their hand relative to their body when grasping objects. For example, dangerous objects (eg. hot soup or a knife) are manipulated farther away from the body than fragile objects.
- (3) **Object on hand-selection:** The category of an object can lead to the selection of a particular hand. For example, tools like knives or spoons are typically held with the dominant hand. [13, 22] show that the dominant hand crosses the midline more often when a task involves the manipulation of a tool.
- (4) **Surface on hand-distance:** Surfaces can affect the distance at which a person manipulates objects. For example, handling a pot on a hot stove lead to safety distance to avoid burns.
- (5) **Surface on hand-selection:** Spatial layouts can lead to hand selection. For example, a refrigerator’s door opens comfortably for the right hand but less for the left.
- (6) **Object-Volume (OV) on hand-selection:** The volume of objects can affect hand selection. Large objects, but still manageable for a single hand, are likely handled with the more dexterous hand. [30, 37] report that the non-dominant hand is for smaller objects while the dominant hand for larger.

- (7) **Object-Volume (OV) on hand-distance:** The volume of objects could affect the extent to which a person stretches the arm. For instance, the handle of a frying pan could lead to a near-body grasp.
- (8) **Surface-Sliding (SS) on hand-distance:** Sliding surfaces can affect the distance at which a person takes an object. For example, when reaching into a sliding drawer, pulling the drawer open could affect hand distance.
- (9) **Surface-Sliding (SS) on hand-selection:** Surfaces that slide can potentially drive hand selection. For example, preferences opening containers with a particular hand and reach into the container with another or the same hand.
- (10) **Surface-Inside (SC) on hand-distance:** Target objects located inside of containers are harder to reach than those on unconstrained surfaces. For example, reaching for a milk carton stored deep inside the refrigerator demands stretching the arm wide while on a table not.
- (11) **Surface-Inside (SC) on hand-selection:** Targeting objects placed inside of containers can lead hand-selection. For example, when grabbing an object inside of a drawer, one hand opens the drawer while other reaches into.
- (12) **Hand-dominance (DO) on hand-selection:** The preference of one hand over the other directly influences hand selection. [5, 23] show that the dominant hand is preferred for close to mid-line targets.
- (13) **Hand-dominance (DO) on hand-distance:** The preference on a particular hand could lead a person to extend the dominant arm further away than the non-dominant thereby affecting the proximity at which an object is grasped.

The mechanization of causal inference [26] allows agents to target any of the hypotheses (arrows) embedded in the graph in Figure 3. This article covers one of them to validate the approach.

The question addressed in this work is whether hand-distance drives hand-selection or not. This hypothesis involves two decisions people make (unconsciously) when grasping objects with their hands. In other words, the distance at which a person decides to grasp an object (i.e. the arm extension) potentially affects the choice of hand (i.e. left or right).

This question is challenging given that other factors could affect (even simultaneously) the choice of hand. First, hand-selection could be driven by hand-dominance as well as any of the other factors modeled in the graph. Second, inferences have to deal with class imbalances in raw data. Third, because not all confounding factors are captured by our model, bias-free guarantees are compromised. This is indicated explicitly in the graph by the node of unobserved confounding in Figure 3. For instance, we have not modeled the mental state of a person, which could play a relevant role in grasping actions and therefore be driving change.

Note that arrows in causal graphs only state that a potential influence exists, these do not assert certainty. Moreover, because causal graphs do not impose any form of functional assumption on the data-generating-process (e.g. linearity), models are expected to transfer across scenarios. For example, the graph in Figure 3 is expected to hold true for other actions than the modelled here, e.g. placing.

In summary, to account for confounding in unconstrained examples of hand manipulation, we propose the model of Figure 3. Among the many hypotheses this graph provides, the causal quantity of hand-distance on hand-selection validates this work. Before plugging this model into the machinery of causal inference, the collection of evidence are introduced next.

## 2.2 Step 2: Collecting the Evidence

A headset and two hand controllers (VR-devices) track user motions in three-dimensional space. A game engine (the unreal engine) is responsible for logging streams of data and rendering a near photo-realistic virtual kitchen scenario as shown in Figure 4. The gear users employ to interact with the virtual environment are the HTC Vive Pro and the original handheld controllers.

A headset and two hand controllers (VR devices) track user motions in three-dimensional space. A game engine is responsible for logging streams of data and rendering a near photo-realistic virtual kitchen scenario (Figure 4). The gear users employ to interact with the virtual environment are the HTC Vive Pro and the original handheld controllers. With the tool described in [4], raw streams of data are annotated with semantics such as collisions, changes in object states and so on. Robots can then access the semantically annotated data with [19]. The focus of this contribution is on the processing of such data.



**Figure 4: A table setting as experienced by the person in the virtual kitchen. The last image in the sequence shows the table set.**

A table setting involves a person moving objects from one place to another. This includes searching, transporting, and placing objects in the kitchen - opening and closing containers, and so forth.

The three collections of table settings considered in this work employ the same recording equipment. Likewise, the virtual kitchen layout, the initial position of objects, and the instruction to set the table for breakfast are the same. However, each collection of evidence differs considerably.

Data collection [28] was recorded with the purpose to study how people arrange objects on the table (denoted here as Ds-1). The dataset [17] studies the time spent setting the table under various modalities (in short Ds-2). The collection of data Ds-v validates this work. The key distinction is that Ds-1 and -2 originate from uncontrolled evidence (described in Section 5), whereas Ds-v stems from a controlled experiment (discussed in Section 3). Note that the unit of analysis are not participants (only 3) but rather grasping actions (order of hundreds).

Hand selection in each set of data is different in spatio-temporal aspects such as the preference of objects selected for the breakfast, the surfaces on which objects are placed, etc. The dashes in the

summary Table 1 are evidence of such asymmetries across data collections.

Features are either binary, categorical, or numeric as shown in Table 1. The distance of a grasp (denoted by  $D$ ) is measured when the user triggers a grasp with the motion controllers. Hand-distance is the only feature which requires a transformation because not all selected estimators (introduced in the next section) support continuous treatment values. Discretizing with a median split overcomes this issue.

A median split over the straight-line distance between the tracked headset and handheld VR-devices determines the distance. The resulting splits for the validation set are (0.56, 0.77) meters for a close grasp and (0.77, 0.85) meters for a distant grab. Likewise, for  $Ds-1$  the bins corresponding to close/far are (0.59, 0.73) and (0.73, 0.92). Similarly, the splits for  $Ds-2$  are (0.47, 0.63) and (0.63, 0.77), correspondingly. The frequencies for hand-distance are based on these splits as shown in Table 1.

### 2.3 Step 3: Drawing Conclusions

To mechanize inferences a Structured Causal Model (SCM) [26] is specified. The causal graph defined earlier in step 1 (Figure 3) is plugged into the SCM.

Based on the selected hypothesis and on the graph topology, this formalism derives the following (bias-free) expression also known as the estimand [26]:

$$\frac{\partial}{\partial D}(Expectation(H|O, OV, S, SS, SC)) \quad (1)$$

The symbols of expression (1) correspond to the features introduced in step 2 (Table 1). The treatment and outcome in this expression are determined by the hypothesis, i.e. hand-selection ( $H$ ) and hand-distance ( $D$ ). The backdoor criterion [26] identifies the remaining symbols as potential factors affecting the hypothesis.

The next element to define is the estimate reported in subsequent sections. The causal estimate is defined by the mean conditional average treatment effect (a.k.a. CATE mean):

$$E \left[ Y_{i,t=1}^e - Y_{i,t=0}^e | X_i \right] \quad (2)$$

where the outcome  $Y$  (hand choice) is computed by estimator  $e$  (described below) for data point  $i$  (instant of time the person triggered a grasp) conditioned on the treatment  $t$  (proximity distance) provided the set of observed features  $X$  (factors in Figure 3).

Because the data generating process in human hand-behavior is unknown, four estimators operating under differing assumptions cover the full spectrum:

- **linear treatment:** The outcome  $Y$  is a linear function of the treatment  $X$ . Estimator [7] Forest Double Machine Learning Estimator (FDML) imposes this assumption.
- **linear heterogeneity:** The outcome  $Y$  is linear on observable characteristics ( $Z$ ). Estimator [3] Linear Doubly Robust Learner (LDRL) embeds this assumption.
- **linear treatment and heterogeneity:** Estimator [6] LinearDML (LDML) embeds both of previous assumptions.
- **no assumption:** Estimator [10] Forest Doubly Robust Learner (FDRL) imposes no operational assumption.

Note that these four are non-parametric estimators suitable for high-dimensional data. Each of them is a composite two-stage ML-algorithm. The final-stage grants desirable properties and is not subject to hyper-parameter optimization unlike the first-stage models, which comprise cross-validated grid-search<sup>1</sup>.

The selected estimators operate under the unconfoundedness assumption and construct valid confidence intervals. Due to small sample size across data sets, statistical significance is set to 0.9. The computation of confidence intervals follows OLS<sup>2</sup> Inference for (LDML and LDRL) and Sub-sampled Honest Forest for (FDML and FDRL). Further information on these methods are available in [29].

### 3 VALIDATING THE APPROACH

A ground truth dataset is recorded to validate the approach. In particular, to verify that the setting derives correct conclusions, a controlled experiment is performed.

The validation set holds examples of a person setting the table for breakfast. The choice of hand in grasping actions is the only aspect that has been controlled. The experimenter instructs a person to use either the left, right, or both hands before engaging in the activity. The sequence of instructions is chosen randomly by the experimenter. The case of both hands is a distracting element to avoid suspicion by the participant on the matter being investigated. The data collected with the experiment is summarized under column  $Ds-v$  in Table 1 alongside the other collections of data.

Because hand-selection is randomized in the experiment, only weak (near to null) supporting evidence can be expected for hand-selection on the validation set. Therefore, estimators that find support for the hypothesis in the validation data would indicate a red flag on the validity of the setting. On the contrary, ideal results consist of estimators not being able to find evidence in such data.

After computing the effects on the validation set, four estimators recovered no evidence in support of the hypothesis. This is indicated in Table 2 by the confidence intervals including the zero value for all estimators. The result is clear - the setting recovers the expected effect.

This result is encouraging. The setting correctly detected the expected null effect considering the freedom with which the activity had been recorded (free choice in the selection of objects, order, placement, etc., see Table 1). Moreover, as not all confounding factors are captured by our model, bias cannot be guaranteed to be fully removed. Despite breaking unconfoundedness, the approach manages to recover the correct effect. Furthermore, the estimators have been mostly applied to operate at the level of thousands, but in this setting only hundreds of samples are considered. In spite of small quantities of data, the setting is sensible enough to detect the correct effect.

Regarding the methods, Table 2 shows that estimators operating under linear treatment and heterogeneity (top three rows) present narrower confidence intervals than the last variant (bottom row). Smaller confidence intervals for the top three rows makes sense

<sup>1</sup>With parameter grid: Random Forest Regressor (max\_depth:[3, None], min\_samples\_leaf: [10, 50]), Gradient Boosting Regressor (n\_estimators: [50, 100], max\_depth:[3], min\_samples\_leaf: [10, 30]), Random Forest Classifier (max\_depth: [3, 5], min\_samples\_leaf: [10, 50]), Gradient Boosting Classifier (n\_estimators: [50, 100], max\_depth: [3], min\_samples\_leaf: [10, 30]) .

<sup>2</sup>Ordinary Least Squares (OLS).

**Table 1: Summary of three datasets and frequencies of occurrences for different levels of features**

User interaction		Ds-1	Ds-2	Ds-v	
<b>Object (O):</b>	<b>OV (<math>m^3</math>)</b>	<b>384</b>	<b>174</b>	<b>137</b>	
Silverware	$8.2 \times 10^{-5}$	83	70	29	
Glass	$1.22 \times 10^{-3}$	83	-	23	
Milk	$1.40 \times 10^{-3}$	18	8	24	
Juice	$1.93 \times 10^{-3}$	88	-	21	
Bowl	$2.40 \times 10^{-3}$	19	36	16	
Cereal	$6.30 \times 10^{-3}$	93	28	24	
Tray	$9.95 \times 10^{-3}$	-	32	-	
<b>Surface (S):</b>	<b>SC</b>	<b>SS</b>	<b>384</b>	<b>174</b>	<b>137</b>
DiningTable	F	F	12	18	-
FrdgArea	T	F	-	1	1
FrdgDrBtmShlf	T	F	5	2	22
FrdgGlassShlf	T	T	41	-	21
IslndArea	F	F	-	8	21
IslndDrwBtmLft	T	T	-	16	-
LabFloor	F	F	5	-	2
OvenArea	F	F	-	2	-
OvenDrwRight	T	T	41	16	22
SinkArea	F	F	11	24	1
SnkDrwLftBtm	T	T	-	-	1
SnkDrwLftMid	T	T	41	-	20
SnkDrwLftTop	T	T	41	42	25
Tray	F	F	184	42	-
<b>Hand-selection (H):</b>	<b>DO</b>	<b>384</b>	<b>174</b>	<b>137</b>	
Left ( $H_0$ )	F	121	96	64	
Right ( $H_1$ )	T	263	78	73	
<b>Hand-distance (D):</b>		<b>384</b>	<b>174</b>	<b>137</b>	
Close ( $HD_0$ )	F	192	87	69	
Far ( $HD_1$ )	T	192	86	68	

Abbreviation *OV* refers to the volume of an object. *SC* corresponds to surfaces within containers such as the refrigerator’s shelf. *SS* refers to surfaces that slide such as drawers. *DO* refers to hand dominance. *SnkDrwLftMid* refers to the middle drawer out of three drawer stacked one over another spatially located left of the sink. *FrdgDrBtmShlf* refers to the bottom shelf on the fridge door. Last, *T* and *F* refer to true and false.

**Table 2: Effects on the validation set Ds-v**

Samples	Estimator	Effect	Conf. Intv.
137	LDML	-0.04	[-0.21, 0.13]
	LDRL	-0.08	[-0.25, 0.09]
	FDML	0.01	[-0.26, 0.27]
	FDRL	-0.54	[-1.57, 0.50]

Four non-parametric estimators report the effect on the validation set comprising 137 grasping actions of a person performing table settings. The confidence intervals for all estimators include the zero value.

because these estimators impose some structure and thereby deviations are (to some extent) partially restricted. Another reason the least restricted estimator (FDRL) presents the widest intervals is that it requires more data non-forest counterparts. However, beyond distinctions, all of them recovered the correct quantity.

Taken together, the approach recovers the expected null effect on the validation set. Moreover, the effect is recoverable with four estimators operating under different assumptions.

This concludes the validation of the setting on empirical data. Next the stability of the model is verified with the validation data, before studying individuals who freely choose their hands (Ds-2 and Ds-3).

**Table 3: Refutation strategies using validation data**

Estimator	Effect	Check-1(  $\Delta$  )	Check-2	Check-3
LDML	-0.04	-0.07 (0.03)	-0.03 (0.01)	-0.05 (0.01)
LDRL	-0.08	-0.11 (0.03)	-0.04 (0.04)	-3.43 (3.35)
FDML	0.01	0.05 (0.04)	0.01 (0.0)	-0.04 (0.05)
FDRL	-0.54	-0.22 (0.32)	-0.02 (0.52)	-4.82 (4.28)

The columns check-1 to -3 correspond to the re-computed effects after changing the setting (i.e. model and/or data). The numbers enclosed by parenthesis ( $|\Delta|$ ) are the differences between the original and re-computed effects. Differences close to zero values exhibit robustness against refutations.

#### 4 STABILITY OF THE APPROACH

Models are expected to be applicable beyond scenarios for which they were initially designed. Therefore, it is crucial to verify the stability of the setting under different conditions than the initially planned one.

Verifying the stability first involves computing baseline values using the original setting. Then, the original setting is modified and the effects are re-computed (e.g. introducing noise to the data, and/or even changing the structure of the model). Ideally, the re-computed effects differ only slightly from the original values. In general, the smaller the difference between re-computed and original estimates, the more stable the model is.

Table 3 summarizes the results for three refutation strategies. The column effect corresponds exactly to the original values presented earlier in Table 2, while the remaining columns are the re-computed estimates after inflicting changes with a refutation strategy.

Refutation strategy check-1 randomizes the treatment and outcome data. Check-2 changes the structure of the model by introducing unobserved random factors that affect the treatment and outcome<sup>3</sup>. Strategy check-3 computes the estimates on randomly drawn subsets of the original data<sup>4</sup>.

The re-computed effects after randomizing the treatment and outcome data (check-1) should not differ from the original values. Indeed this is the case as Table 3 reports near zero differentials except for FDRL. Likewise, when introducing unobserved random confounding factors (check-2) to the causal graph, close to zero  $\Delta$  values are reported for the first three estimators except for the last (FDRL). At last, removing random subset of data exhibits near zero differentials for LDML and FDML but not for LDRL and FDRL.

Overall, Table 3 reports that DML estimators outperform DRL across refutation checks as indicated by the near-to-zero differential values ( $|\Delta|$ ). DRL, on the other hand, recovers less stable effects.

DRL typically has a higher variance than DML variants as reported in [29], and our results confirm this. Due to imbalances in our datasets, treatment has a small probability of being assigned in the regions of the control space ( $Z$ ) which is critical for DRL estimators, also known as small-overlap. Under such conditions, the DML variants extrapolate better than DRL, and this is the case shown in Table 3.

<sup>3</sup>Configured with linear effect strength set to 0.02 for both treatment and outcome.

<sup>4</sup>Configured with fraction of data to 0.9.

**Table 4: Effects on observational data Ds-1 and -2**

Data set	Smpl.	Estimator	Effect	Conf. Intv.
1	384	LDML	0.09	[0.02, 0.17]
		LDRL	0.09	[0.01, 0.16]
		FDML	0.16	[0.05, 0.26]
		FDRL	0.08	[-0.01, 0.17]
2	174	LDML	-0.13	[-0.26, 0.01]
		LDRL	-0.18	[-0.42, 0.05]
		FDML	-0.13	[-0.29, 0.02]
		FDRL	-0.08	[-0.31, 0.15]

The effect for hand-distance on hand-selection for two individuals freely selecting their hands when setting the table. For one individual the effect is positive, whereas for the other negative. In all cases estimators agree except for FDRL on Ds-1, which is a weak positive effect as the confidence interval includes the null value.

#### 5 OBSERVING DIFFERENT INDIVIDUALS

This section analyses the effect of hand-distance on hand-selection for two individuals. Unlike the effects reported in Section 3, in the following subjects exercise the free choice in hand selection - all else remains equal (model, estimators, parameters).

Provided with examples from dataset Ds-1, the estimators recovered positive effects. This is indicated by the confidence intervals not including the zero value for three estimators except for one as shown in Table 4. A positive effect means that when targeting objects farther away from the body, the preferred choice is the right hand.

In dataset Ds-2 the four estimators derive a weak negative effect. This is indicated by the confidence intervals marginally containing the zero. Moreover, these intervals are less centered than those reported earlier on the validation set (Table 2) thus indicating a weaker null effect. A negative effect means that the individual prefers near-to-body grasps with the right hand.

To summarize, different than on the validation set, the resulting effects are not null on evidence which is uncontrolled (Ds-1 and Ds-2). Both individuals differ in the way hand-distance affects hand-selection. One individual engages with the right hand targets that are farther away, whereas the other one prefers nearby.

#### 6 EXPLAINING HAND MANIPULATION

So far we have described the effect (direction and strength) that hand-distance has on hand-selection for three individuals. But insights explaining the sensibility of these effects w.r.t factors are missing. Interpretability is especially important on inputs for which no ground truth data exists (Ds-1 and Ds-2).

To this end, the sensibility of the effects (reported earlier in Table 2 and Table 4) are interpreted using trees as shown in Figure 5. Every tree interprets the sensibility of a different person. A tree encodes the direction of sensibilities with colors - negative (red), neutral (white) or positive (green) - while the intensity of colors indicates the level of sensibility (either strong or weak). The splitting policy of the tree is in accordance with factors that influence the sensibility of effects. For instance, Figure 5a interprets the effect in terms of object volume (i.e. the bounding box of an object).

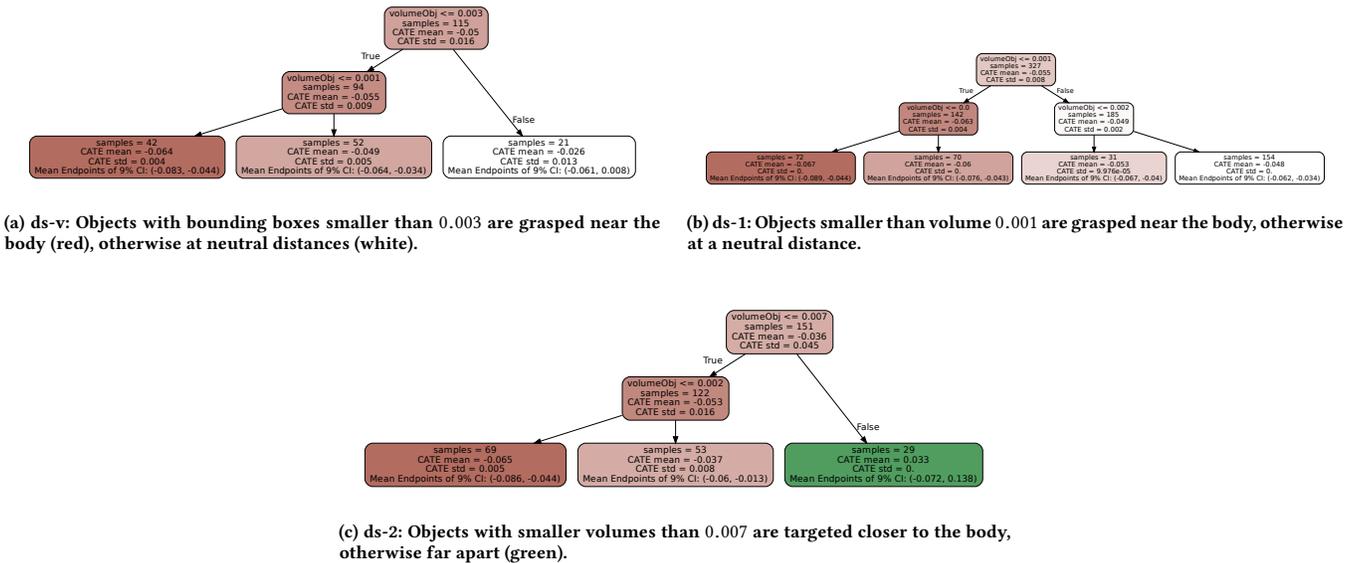


Figure 5: Interpreting the effect of hand-distance on hand-selection in terms of object volume on each collection of data.

Figure 5 shows that the smaller an object is, the closer to the person’s body it is grabbed. This is indicated by the colors ranging from dark red to lighter colors as the volume of objects increases. This behavior is consistent across the three individuals. However, the quantitative value of when an object is considered to be big/small differs for each individual as indicated by the root node of the trees.

To summarize, the effects presented in earlier sections (Table 2 and Table 4) are explainable in terms of sensibility to influential factors. The sensibility of the effect that hand-proximity exerts on hand-selection are explained in terms of object-volume. Different individuals present consistent trends but differ in terms of quantitative preferences of object volumes. Individualized insights like these have the potential to enable robots to provide personalized assistance.

## 7 RELATED WORK

Autonomous robots could one day spare the cost of programming them. One way to reduce the amount of programming is by providing robots with examples of the task space [2]. Another strategy is learning the control space showing exemplary movements to the robots [21]. Providing robots with examples instead of code shifts the burden from programming to teaching them. Teaching is demanding as high quality, successful, and sufficient examples are required.

One strategy to cut the cost of teaching robots, is to learn by doing in simulation environments as done in [14, 31]. Virtual settings offer nearly endless synthetic examples for an initial upfront cost. Computational resources leverage exhaustive explorations in simulation space. Besides sample inefficiency, solutions from simulated environments must overcome transfer issues to real-world scenarios [18].

Unlike the previous strategies, children approach things differently. They can learn quickly from a few ‘uncurated’ good-and-bad examples [16]. Moreover, children can explain their answers [36], an ability often overlooked in model-free data-driven approaches. More crucially, children transfer knowledge from one domain to another [24].

Developmental AI and Robotics are still in their infancy [8]. A key achievement in recent years has been the algorithmization of causal inference [26]. Such advancements offer opportunities to explore alternative reasoning capabilities for robots.

This contribution explores the feasibility to equip robot with child-like observational abilities. To prove the concept, we setup a framework and confront the observational capabilities to a few uncurated empirical examples of hand manipulation.

The authors in [9] build a causal and-or graph for objects and actions from videos. Unlike our contribution, this work focus on action detection and automating the discovery of causal models at the expense of transferability and interpretability.

[1] employs SCM to recover trajectories for different kinesthetic teaching styles. However, this work differs considerably. First, the end-effector of a robot is modeled whereas our contribution involves the human hand. Second, the authors address robot poses while our contribution targets human motor decisions. Third, their source of evidence is structured and curated when compared to the unconstrained activities considered here.

[34] employed causal inference to reduce the space of action pairs recorded in a virtual environment thereby boosting performance by six fold. This work puts the focus on causal dependencies of action sequences rather than on decision-making and context as presented here.

## 8 CONCLUSION

To unlock individualized robotic support in human-robot interactions, this contribution explores the validity of causal quantities on empirical evidence in the context of everyday hand manipulations.

The conditions under which the approach is examined are challenging. First, this contribution does not include all relevant factors in the model thus breaking the unconfoundedness assumption. Therefore, we can expect bias-reduced estimates but not bias-free. This setting is valuable to investigate because children cope with states of incompleteness. Second, the non-parametric estimators of this contribution typically target samples in the order of thousands but only hundreds are considered here.

The approach recovered the null effect on the validation data. Four estimators operating under differing assumptions agreed on this result. Moreover, the model is stable against three refutation strategies on two collections of data.

Unlike the validation set, targeting observational data does not derive a null effect. Instead, estimates result in positive and weak negative estimates for two individuals.

Taken both results, the approach is capable to detect null and non-null effects. The possibility of the framework being either too sensitive (to derive null effects) or insensitive (to detect non-null effects) is thus excluded. Moreover, the opposing estimates derived for two individuals manifest the feasibility to recover effects in different directions.

Lastly, we interpret the sensibility of effects w.r.t. contextual factors (i.e. object-volume) potentially affecting hand behavior (i.e. the relationship between hand-distance and hand-selection). The trees highlighted commonalities as well as individualities across three subjects. Such as the preference of the right hand for distant targets, and the personal threshold of objects considered small or large given by the root of the tree.

Our results are encouraging and foundational for further research. In the future, we plan to study other actions such as placing, opening, and closing. We look forward to extending our model with factors such as shape, handles, grasping points, and full-body tracking.

## ACKNOWLEDGMENTS

This work is partially funded by the Scholarship Programme Don Carlos Antonio López (BECAL) and DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 EASE - Everyday Activity Science and Engineering, University of Bremen (<http://www.ease-crc.org/>).

## REFERENCES

- [1] Daniel Angelov, Yordan Hristov, and Subramanian Ramamoorthy. 2019. Using Causal Analysis to Learn Specifications from Task Demonstrations. *CoRR* abs/1903.01267 (2019). <http://arxiv.org/abs/1903.01267>
- [2] Brenna D. Argall, Sonia Chernova, Manuela M. Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (2009), 469–483.
- [3] Heejung Bang and James M. Robins. 2005. Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics* 61, 4 (June 2005), 962–973. <https://doi.org/10.1111/j.1541-0420.2005.00377.x>
- [4] Michael Beetz, Daniel Bessler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoglu, and Georg Bartels. 2018. KnowRob 2.0 – A 2nd Generation Knowledge Processing Framework for Cognition-enabled Robotic Agents. In *International Conference on Robotics and Automation (ICRA)*. Brisbane, Australia.
- [5] D. V. M. Bishop, V. A. Ross, M. S. Daniels, and P. Bright. 1996. The measurement of hand preference: A validation study comparing three groups of right-handers. *British Journal of Psychology* 87, 2 (#may# 1996), 269–285. <https://doi.org/10.1111/j.2044-8295.1996.tb02590.x>
- [6] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, and Whitney Newey. 2017. Double/Debiased/Neyman Machine Learning of Treatment Effects. *American Economic Review* 107, 5 (May 2017), 261–65. <https://doi.org/10.1257/aer.p20171038>
- [7] Victor Chernozhukov, Denis Nekipelov, Vira Semenova, and Vasilis Syrgkanis. 2018. *Plug-in Regularized Estimation of High-Dimensional Parameters in Nonlinear Semiparametric Models*. Papers 1806.04823. arXiv.org. <https://ideas.repec.org/p/arx/papers/1806.04823.html>
- [8] Manfred Eppe, Stefan Wermter, Verena V. Hafner, and Yukie Nagai. 2021. Developmental Robotics and its Role Towards Artificial General Intelligence. *KI - Künstliche Intelligenz* 35, 1 (Feb. 2021), 5–7. <https://doi.org/10.1007/s13218-021-00706-w>
- [9] Amy Sue Fire and Song-Chun Zhu. 2016. Learning Perceptual Causality from Video. *ACM TIST* 7, 2 (2016), 23:1–23:22.
- [10] Dylan J. Foster and Vasilis Syrgkanis. 2019. Orthogonal Statistical Learning. *CoRR* abs/1901.09036 (2019). <http://arxiv.org/abs/1901.09036>
- [11] Carl Gabbard and Casi Rabb Helbig. 2004. What drives children’s limb selection for reaching in hemispace? *Experimental Brain Research* 156, 3 (June 2004), 325–332. <https://doi.org/10.1007/s00221-003-1792-y>
- [12] Carl Gabbard and Casi Rabb. 2000. What Determines Choice of Limb for Unimanual Reaching Movements? *The Journal of General Psychology* 127, 2 (#apr# 2000), 178–184. <https://doi.org/10.1080/00221300009598577>
- [13] Claudia L. R. Gonzalez, Jason W. Flindall, and Kayla D. Stone. 2015. Hand preference across the lifespan: effects of end-goal, task nature, and object location. *Frontiers in Psychology* 5 (#jan# 2015). <https://doi.org/10.3389/fpsyg.2014.01579>
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afcc3-Paper.pdf>
- [15] Alison Gopnik. 2010. How Babies Think. *Scientific American* 303, 1 (July 2010), 76–81. <https://doi.org/10.1038/scientificamerican0710-76>
- [16] Frank Guerin. 2011. Learning like a baby: a survey of artificial intelligence approaches. *Knowledge Eng. Review* 26, 2 (2011), 209–236.
- [17] Andrei Haidu and Michael Beetz. 2019. Automated Models of Human Everyday Activity based on Game and Virtual Reality Technology. In *International Conference on Robotics and Automation (ICRA)*. Montreal, Canada.
- [18] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Florian Golemo, Melissa Mozifian, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, C. Karen Liu, Jan Peters, Shuran Song, Peter Welinder, and Martha White. 2020. Perspectives on Sim2Real Transfer for Robotics: A Summary of the R-SS 2020 Workshop. arXiv:2012.03806 [cs.RO]
- [19] Gayane Kazhoyan, Alina Hawkin, Sebastian Koralewski, Andrei Haidu, and Michael Beetz. 2020. Learning Motion Parameterizations of Mobile Pick and Place Actions from Observing Humans in Virtual Environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. <https://doi.org/10.1109/IROS45743.2020.9341458>
- [20] Ned Kock. 2015. How Likely is Simpson’s Paradox in Path Models? *IjeC* 11, 1 (2015), 1–7.
- [21] Dana Kulic, Christian Ott, Dongheui Lee, Junichi Ishikawa, and Yoshihiko Nakamura. 2012. Incremental learning of full body motion primitives and their sequencing through human motion observation. *I. J. Robotics Res* 31, 3 (2012), 330–345.
- [22] Pascale Leconte and Jacqueline Fagard. 2006. Which factors affect hand selection in children’s grasping in hemispace? Combined effects of task demand and motor dominance. *Brain and Cognition* 60, 1 (#feb# 2006), 88–93. <https://doi.org/10.1016/j.bandc.2005.09.009>
- [23] Carla M. Mamolo, Eric A. Roy, Pamela J. Bryden, and Linda E. Rohr. 2004. The effects of skill demands and object position on the distribution of preferred hand reaches. *Brain and Cognition* 55, 2 (#jul# 2004), 349–351. <https://doi.org/10.1016/j.bandc.2004.02.041>
- [24] Katalin Oláh, Fruzsina Elekes, Réka Pető, Krisztina Peres, and Ildikó Király. 2016. 3-Year-Old Children Selectively Generalize Object Functions Following a Demonstration from a Linguistic In-group Member: Evidence from the Phenomenon of Scale Error. *Frontiers in Psychology* 7 (June 2016). <https://doi.org/10.3389/fpsyg.2016.00963>
- [25] Marios G. Pavlides and Michael D. Perlman. 2009. How Likely Is Simpson’s Paradox? *The American Statistician* 63, 3 (#aug# 2009), 226–233. <https://doi.org/10.1198/tast.2009.09007>
- [26] Judea Pearl. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, England.
- [27] Judea Pearl. 2014. Comment: Understanding Simpson’s Paradox. *The American Statistician* 68, 1 (Jan. 2014), 8–13. <https://doi.org/10.1080/00031305.2014.876829>

- [28] Lisset Salinas Pinacho, Alexander Wich, Fereshta Yazdani, and Michael Beetz. 2018. Acquiring Knowledge of Object Arrangements from Human Examples for Household Robots. In *KI 2018: Advances in Artificial Intelligence: 41st German Conference on AI, Berlin, Germany, September 24–28, 2018, Proceedings*, Vol. 11117. Springer, 131. [https://link.springer.com/content/pdf/10.1007%2F978-3-030-00111-7\\_12.pdf](https://link.springer.com/content/pdf/10.1007%2F978-3-030-00111-7_12.pdf)
- [29] Microsoft Research. 2019. EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/microsoft/EconML>. Version 0.x.
- [30] Christian Seegelke and Peter Wühr. 2018. Compatibility between object size and response side in grasping: the left hand prefers smaller objects, the right hand prefers larger objects. *PeerJ* 6 (Dec. 2018), e6026. <https://doi.org/10.7717/peerj.6026>
- [31] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nature* 550, 7676 (Oct. 2017), 354–359. <https://doi.org/10.1038/nature24270>
- [32] E. H. Simpson. 1951. The interpretation of interaction in contingency tables. *J. of the Royal Statistical Society, Series B* 13, 2 (1951), 238–241.
- [33] David M. Sobel, Joshua B. Tenenbaum, and Alison Gopnik. 2004. Children’s causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science* 28, 3 (2004), 303–333.
- [34] Constantin Uhde, Nicolas Berberich, Karinne Ramirez-Amaro, and Gordon Cheng. 2020. The Robot as Scientist: Using Mental Simulation to Test Causal Hypotheses Extracted from Human Activities in Virtual Reality. In *IROS 2020*. Las Vegas, USA.
- [35] Anna Waismeyer, Andrew N. Meltzoff, and Alison Gopnik. 2014. Causal learning from probabilistic events in 24-month-olds: an action measure. *Developmental Science* 18, 1 (July 2014), 175–182. <https://doi.org/10.1111/desc.12208>
- [36] Caren M. Walker, Tania Lombrozo, Joseph J. Williams, Anna N. Rafferty, and Alison Gopnik. 2016. Explaining Constrains Causal Learning in Childhood. *Child Development* 88, 1 (July 2016), 229–246. <https://doi.org/10.1111/cdev.12590>
- [37] Peter Wühr and Christian Seegelke. 2018. Compatibility between Physical Stimulus Size and Left-right Responses: Small is Left and Large is Right. *Journal of Cognition* 1, 1 (Feb. 2018), 17. <https://doi.org/10.5334/joc.19>