
Counterfactual Bayesian inference in the development of pre-reaching behaviour



Radboud Universiteit Nijmegen

By
Wouter EIJLANDER
s4243242

Supervised by
Johan KWISTHOUT

January 2019

Contents

1	Introduction	3
1.1	Developmental Robotics Models	5
1.2	Bayesian models of motor action	8
1.3	Goals and Hypotheses	12
2	Methods	13
2.1	Movement	13
2.1.1	Arm Model	13
2.1.2	Kinematics	14
2.1.3	Muscle Coactivation	18
2.2	Cognitive Model	19
2.2.1	Causal Model	19
2.2.2	Imaging over Limb Mechanics	22
2.2.3	Replacing X,Y and Z	23
2.2.4	Processing Timeline	25
2.3	Simulations	25
2.3.1	Development of Muscle Coactivation	26
2.3.2	Motor Babbling and the Learning Phase	26
2.3.3	Developmental Phase	26
2.3.4	Testing Phase	27
2.4	Data Acquisition	27
3	Results	29
4	Discussion of Results	31
5	Conclusion and future work	35
6	References	40

Abstract

The development of reaching behaviour in infants starts just days after birth, and generally follows a certain pattern of pre-reaching frequency which contains a notable drop in frequency around 7 weeks of age. The developmental mechanisms underlying this pattern vary, but a prominent theory is that the onset of muscle co-activation influences developmental pre-reaching behaviour. Developmental robotics studies are valuable in examining developmental theories' merits by reverse-engineering behavioural models from leading theories and observing resulting behaviour. The current study aims to use developmental robotics to provide additional insight into the muscle coactivation theory. The cognitive model used to represent processing of motor actions was realized as a counterfactual Bayesian agent, implementing a previously untested method, counterfactual imaging, as its inference strategy. The viability of the imaging procedure in cognitive modeling of motor control was evaluated based on task performance. Due to complexity issues inherent in the imaging procedure, its viability as a model for cognition was found to be limited. Even after inference simplification, performance in motor tasks was low. As such, its efficacy was found to be insufficient for it to be a viable method in modeling cognitive tasks. Influence of the onset of muscle coactivation was evaluated by comparing it to empirical results from behavioural studies in developmental psychology. The behavioural effects observed in human infants were not replicated by the simulated infants implementing an onset of muscle coactivation. Since imaging was found to be an insufficient model of motor control, the observed results currently only provide a preliminary indication that muscle coactivation may not be the cause of observed behavioural patterns. Further investigation using different models of motor cognition may be necessary.

1 Introduction

From birth, one of the first things all human infants do is exploring and learning their motor space. All of our limbs have their own functionalities, and their functioning must be learned before we can use them dependably and inattentively. Among these functionalities is the capability of reaching that our arms grant us, and infants start exploring this capability almost from birth. Before infants reach the age of reaching onset, they exhibit goal-oriented forward extensions of the arms. These movements are called prereaching movements or prereaching behaviour, and their development usually follows certain consistent patterns. In a longitudinal study, von Hofsten [1] found that prereaching behaviour is exhibited even in the first weeks of an infant's life in frequencies that generally follow the following pattern: in the first several weeks of life, the amount of prereaching behaviour is relatively stable, and happens somewhat frequently. Around 7 weeks of age, the number of prereaches overall drops, while visual fixation on objects of interest increases. As infants age, the frequency of prereaching behaviour starts increasing around 10 weeks of age onward, and keeps increasing. This temporary decrease in prereaching movements seems to indicate a change in a goal-directed action, but the mechanism behind this change still remains unknown. Additionally, von Hofsten found a pattern in hand posture during prereaching behaviour: prereaches performed were mostly performed with a closed hand in the first two months of life, and is especially predominant during the drop-off in overall prereaching behaviour around 7 weeks of age. After this dip occurs, the number of closed-hand prereaches decreases again, and prereaches where the hand is opened before or during the forward extension overtake them in frequency.

Von Hofsten proposed two possible explanations for the decrease in infant's prereaching movements at 7 week of age. The first is based on the theory of approach and withdrawal [2]. This explanation states that, as infants are subjected to stimuli, they may withdraw from a stimulus as its intensity increases. This could be visual withdrawal (i.e. turning eyes away, or even the entire head), but also proximal withdrawal (i.e. not wanting to get physically near the stimulus). This explanation, however, would suggest that alongside a decrease in forward extensions, infants' fixations on the target would also decrease, which was not the case in von Hofsten's study. The second explanation has von Hofsten's favour, and focuses on co-activation of agonist and antagonist muscles which commences between 1 and 2 months of age [3]. This hypothesis states that, between the age of 1 and 2 months, infants transition from using an agonistic muscle group to using both an agonistic and an antagonistic muscle group in goal-directed movement. Agonistic muscle groups are muscles that pull parts of a limb in the direction of a goal position, whereas antagonistic muscle groups do the opposite.

Contrary to this phenomenon is the restriction imposed by Descartes' law of reciprocal innervation [4]. According to this theory, whenever a movement involving contrary muscle groups is executed, the agonistic muscle group should contract to perform that movement, while the antagonistic muscle group should relax with the same magnitude. It is due to the inhibitory nature of this effect that it is often called *reciprocal inhibition*. The muscle coactivation as seen by Gatev [3]

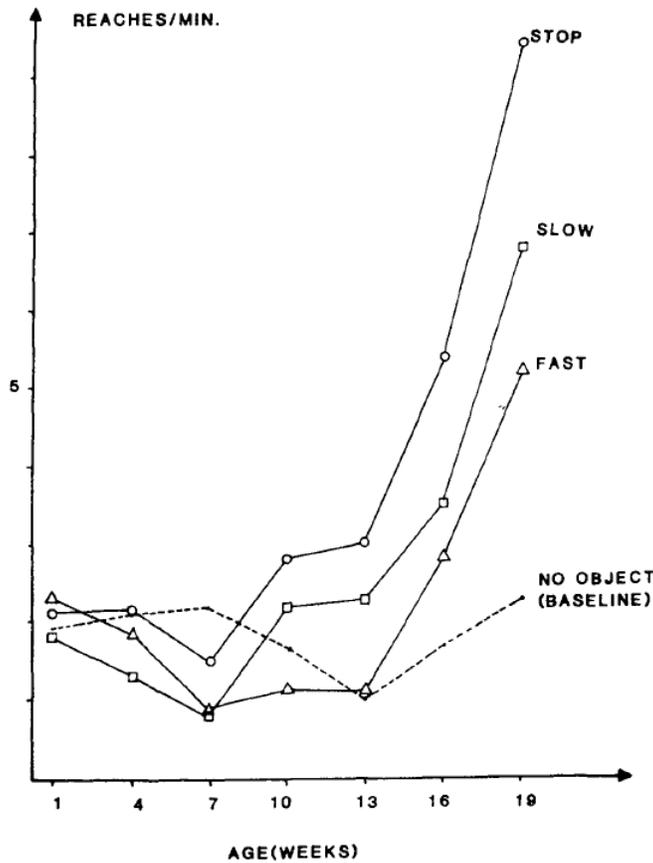


Figure 1: Frequencies of extended reaches across different conditions as found by von Hofsten. One can clearly see trends that are consistent across movement conditions (the target object movement speed: *stop*, *slow*, *fast*): Stark increasing slopes starting at 13 weeks of age, and small dips in extended reaching frequency at 7 weeks of age. This consistent behaviour does not hold under the indicated control condition, where no object was present. From [1]

runs contrary to this phenomenon, making it a notable effect in the developmental cycle of infants' motor control. The law of reciprocal innervation is also reported for the control of eye muscles, which is more commonly referred to as Sherrington's law of reciprocal innervation. Gatev's reports on muscle coactivation of the arms do not extend into the domain of oculomotor control, and as such no assumptions can be made about its presence or absence in this field. If muscle coactivation in the oculomotor system does not follow the same developmental pattern as it does in control of limbs, this may cause the observed difference between the two: The increased gaze percentages at the 7-week age category, in contrast to the apparent reaching withdrawal at the same age.

Frey-Law et al. [5] present research into muscular activation profiles in both the knee and elbow joints. They showed a definite, but small, amount of muscle coactivation to be present in healthy

adults during precise, goal-oriented movements. This may point towards coactivation onset in infants not being an ‘external’ effect, but rather a strategy involved in more precise, less reflexive goal-directed reaching. Frey-Law et al. note the source of this antagonist coactivation seems to be the premotor areas, separating it from reciprocal inhibition, which has “distinct pathways”. These distinct pathways are often stated to reside mostly in the spinal cord. However, these pathways seem to be linked not only to inhibitory, but also excitatory connections [6]. More concretely, Crone et al. show that the spinal source for reciprocal inhibition can be traced back to the inhibitory reciprocal Ia pathway [7]. This pathway stems from the muscle spindles in the agonist muscles, and follows the inhibitory interneurons to the antagonistic motor neurons. Additionally, Hoshiyama et al. show preliminary evidence in favour of a cortical influence on reciprocal muscle control [8]. Using Transcranial Magnetic Stimulation (TMS) over smaller intervals just before voluntary movements, they showed a gradual increase in agonist facilitation, which co-occurred with antagonist inhibition from 60ms before the voluntary movement. These results suggest a cortico-spinal influence on the reciprocal effects of muscle innervation.

Gribble et al. [9] measured the amount of muscle coactivation in goal-oriented reaching tasks in healthy adults in order to determine its involvement in movement accuracy. They measured electromyographic signals in the arm muscles during and after movements aimed at targets of varying sizes. The arm’s muscle coactivation was observed to be inversely related to the target’s size; for smaller targets, participants co-contracted their opposing muscles more. This suggests that, in healthy adults, muscle coactivation is a modulating force that facilitates accurate movement of the arms.

It stands to reason that the neurological sources of coactivation and reciprocal inhibition are the same in both infants and adults. As such, one can conclude that, as in adults, muscle coactivation in infants may be strongly involved in fine-grained motor control. Indeed, where the involvement of small measures of muscle coactivation in fine-grained control is to be expected, the amount of coactivation in infants is much increased compared to that in adults. The reason for this phenomenon is, as per this research, unknown. We hypothesize, however, that it is used by infants as a means to quickly explore and learn fine-grained control of the arm necessary for reaching, which is thus preceded by the coactivation phase.

1.1 Developmental Robotics Models

Empirical developmental studies show many behavioural trends and neural mechanisms that seem to be involved in the development of infants. Confirming or falsifying theories regarding causalities between behaviour and cognition can prove difficult due to a major restriction inherent in developmental psychology: Participants are (very) young. Selecting participants for specific abnormalities becomes near impossible considering such studies often require infants only a few weeks or months of age— after all, it is not often that a 2-week-old infant is diagnosed with specific neurological disorders. It is here that developmental robotics can provide insight into the relations between

physical and cognitive mechanics on one hand, and developmental trends on the other. Modeling a ‘growing’ agent that adheres (or runs counter) to the tested theory allows researchers to test the merit of said theory. In this specific case, modeling an artificial agent following the natural constraints set by von Hofsten’s coactivation hypothesis allows us to test the hypothesis by comparing simulation results to the empirical ones. Such insight into the basis underlying developmental trends can prove vital in the understanding of human cognition.

Before diving into our own model of development, some knowledge on the current state-of-the-art in developmental robotics must be established. Several models for the development of motor skills have been presented, some of which in direct response to von Hofsten’s findings. The following are several robotic models of motor-space exploration in developing infants. Savastano & Nolfi [10] present a biologically inspired approach to learning models of inverse kinematics. It learns in increments, and makes use of the reflexive nature of infants’ movements, as well as the maturation process involved in them. Their results closely resemble those found in empirical studies done with young infants, and show the importance of variation in factors that influence the complexity of tasks during motor-skill acquisition. This incremental learning process provides opportunities for infants’ motor skills to progress more effectively as they learn new, more complex variations as they improve. The results found in this study are noted to be analogous to the motor skill characteristics seen in 2 to 18-month old infants. This time interval does not include the prereaching stage, of which Savastano and Nolfi take note. They do, however, model the prereaching stage, characterizing it by its low visual acuity, its reflexive nature, and the reduced involvement of cortical areas. They implement these characteristics using only sensory-motor connections, and hand-coded weights and fixed learning rates.

Berthier [11] proposed a mathematical model for motor-skill acquisition based on ‘movement units’ described by von Hofsten in later research [12]. Using Q-learning, he trained a model that, using these smaller sub-movements, learned to reach for targets in a 2-dimensional movement space. This model was made on the assumption that infants use these submovements due to lack of control of the arms, and as such was made to exhibit them during low performance time windows. The resulting movement characteristics (i.e. use of submovements, end effector position, movement velocity) were compared to those collected from six human infants. The simulated infants indeed had less need for submovements as they ‘aged’, and generally finished goal-directed movements in the first move after having been trained. This follows the trends seen in humans, where fully in-control adults will tend to make arm movements without the need for submovements. Additionally, the underlying assumption that infants’ use of submovements is directly linked to lack of control was evaluated using measurements of variability and error in infants’ arm movements. These measures were compared to those found in the simulation, and were shown to be statistically similar.

Shaw et al. [13] provide robotic simulations of the exact experiment performed by von Hofsten as part of the babybot challenge organized for IEEE ICDL-Epirob in 2015. They derive some of von Hofsten’s major findings, and attempt to model three of four findings they find most

major: the 7-week dip in prereaching, the peak in closed-handed prereaches at 7 weeks, and the tendency in older infants to prefer stationary targets over moving targets. They do so by modelling infants' fixation, excitation, and reaching probability, while separating systems to learn fixation and reaching. The model is based on learning sensory-motor mappings of target locations and motor commands, which they used previously to model fixation learning [14]. They base the form of (pre)reaching at specific ages on empirical results regarding the use of several agonist-antagonist pairs. These are held accountable for certain reaching phenomena such as the locked-elbow reaches, and circumvent modelling of muscle coactivation, noting that “[muscle coactivation] only affects the type of reaching, and not the amount”. However, it should be noted that von Hofsten does not necessarily measure the total number of attempted prereaches, but those that exceed a certain length, which may still be impacted by restrictions made by certain types of reaching. The results found by Shaw et al. do seem to replicate the statistical characteristics of the results found by von Hofsten, but the timescale is shifted, showing the approximate expected results some 3 weeks late.

Zibner, Tekülve & Schöner [15] present a model of visual fixation, motor reaching, hand opening and closing, and returning to resting position, influenced by a model of the muscular system. They make use of three separate models to simulate the three stages of prereaching development identified by von Hofsten. Their results, although rudimentary, do resemble what would be expected, emulating the empirical results found by von Hofsten. This model does not perform any autonomous machine learning, and are thus separated into three 'snapshots' emulating the three developmental windows. In a later version, the model does perform machine learning to learn the most successful temporal combination of sub-actions (i.e. open hand, reach, close hand) in goal-directed reaching [16]. Indeed, their model still accounts for all stages of prereaching development, supporting the idea that these are caused precisely by the development of sequential organization highlighted by von Hofsten's- and their own results.

Despite the impressive body of psychological and behavioural research regarding the source of muscle coactivation and reciprocal inhibition in healthy adults, there is little such empirical work regarding infants. Aforementioned research from robotics models serve to provide insight into possible workings of motor learning. Due to the extensive observations presented by von Hofsten, current robotics research lifts the veil of what mechanics may be involved in the development of motor skill acquisition. Eijlander et al. [17] used a highly simplified model of the arm controlled by a small neural network to show the effects of muscle coactivation on the development of pre-reaching behaviour. Their results show similarities in simulated pre-reaching frequency and von Hofsten's results. This supports the muscle coactivation account on a low-dimensional, simplified scale. However, interpretation of these results must take into account the level of abstraction used in obtaining them. In particular, the decreased complexity of restricting simulations from a 3D to 2D space may already heavily impact behavioural results. The current study not only aims to expand the current body of knowledge, but also to extend the work from [17].

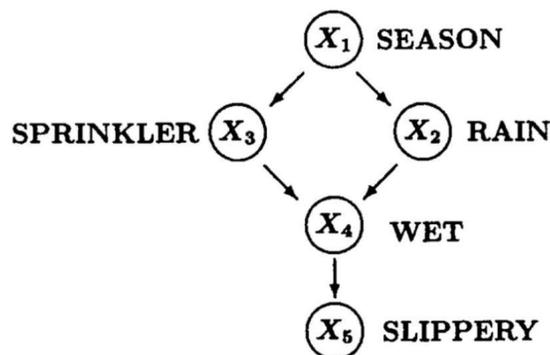


Figure 2: An example Bayesian Network encoding probabilistic relationships between variables. Variables are represented as nodes (X_1 through X_5), and probabilistic or causal relationships are the arrow connections between nodes. The network shows relationships between the season of the year (X_1), whether or not it is raining (X_2), whether or not a garden sprinkler is on (X_3), if a surface is wet (X_4), and if that surface is slippery (X_5). The season directly impacts the odds of rain (i.e. it rains more in autumn than in summer), and the odds of garden sprinklers being on (people water their gardens more in summer), but it does not *directly* impact the odds of a surface being wet. The other relationships also encode such direct relations- indirect relations are encoded as such, through other variables. From: Pearl, 2009 [18, p.15]

1.2 Bayesian models of motor action

There are many different theories dictating models of cognition. One such theory is rooted in Bayesian theory, modeling functions of the brain as processing information using Bayesian statistics. The basic form deals with structural (and sometimes causal) graphical models of the world known as (*Causal*) *Bayesian Networks*, or BNs. These networks keep track of probabilistic relationships between the variables that they represent in a *Directed Acyclic Graph* (DAG) structure. This means that the relationships between variables show a direction, and do not form directed cycles (where in any directed cycle, A affects B , which affects C , which feeds back into A). Such variables would be represented as *nodes*, with each pair of nodes in a network sharing a connection of they directly influence each other. A graphical example of a BN can be found in figure 2. Nodes can encode discrete or continuous values, encoding the probabilistic relations between variables as discrete probability distributions or continuous probability densities. Discrete variables can encode binary values representing a "yes or no" answer, or non-binary values, such as "no rain, slight drizzle, hard rain". As variables encode more possible values (with continuous nodes technically encoding infinitely many values within a certain range), the relationships between nodes in the causal network grow in complexity. From these models, *prior probabilities*, *joint probabilities* and *conditional probabilities* can be derived, and predictions over probabilities can be made based on the internal beliefs and possible observations.

Predictive Processing (PP, also called Predictive Coding) methods are often used to model cognition, providing a framework that utilizes top-down influences on sensory input to determine

the state of the outside world, and learn its regularities [19, 20]. Conceptually, this is done by capturing causalities of the outside world in a (hierarchical) statistical model [20], which can be concretely realised as a BN [21]. Whenever an estimation of a variable in that world is needed, the model makes a top-down prediction from the variables in the model, comparing that prediction to the model’s bottom-up input. If there is a difference between the prediction and the input, that is called the *Prediction Error* (PE). PEs are subsequently used to update the internal model to more accurately reflect the observed causalities, in an effort to minimize future PE. Top-down predictions in situations with high certainty can be used to cancel out bottom-up processing to efficiently make sense of the surrounding world. This makes the task of processing complex statistical surroundings less laborious. In situations where the model has not yet learned the relevant causal relations, PE will be high, and influences from bottom-up inputs will be large, steering the model towards a state that allows for accurate predictions. Imagine the following situation:

You are walking down a flight of stairs blindfolded. At first, you know little of these stairs, and will not be able to accurately predict what is in front of you. Within the first few steps, you will have learned there are downward steps with regular sizes, forming an *internal model* of what you expect next. Your *predictions* are accurate, and you can walk down the staircase with relative ease without having to feel how big the next step will be, no longer relying so heavily on the *bottom-up* detection of the next step. When reaching the bottom of the staircase, you stumble — you have encountered the floor on this next step much earlier than predicted; there is a *prediction error*. You will likely need to check another step, to ensure the floor is flat from here onward, *updating your beliefs* on what is around you, to once again be able to walk upright with relative ease.

Naturally, different applications of Bayesian architectures make use of various other components and algorithms than described above. Some applications use the beliefs inherent in the predictive model to infer the causes or effects of certain observations. Others make use of the internal beliefs to make what-if predictions about the effects of hypothetical scenarios within the model. When modeling the cognitive processes underlying motor control, a Bayesian model must act upon its surroundings based on its internal beliefs. The probabilistic mechanisms of acting upon one’s surroundings should be clear to the agent. It is here that a distinction needs to be made between *actions* and *acts*, where we will follow the definitions provided by Pearl [18, p.108]. These definitions describe an *act* as “[...] a consequence of an agent’s internal beliefs, disposition, and environmental inputs [...]”. An *action* is then described as a deliberative internal decision-making process, generally one that regards the consequences of an act that could be performed. The vital difference between these two is who can observe them, and how this is done: An act can be observed externally from the agent, while an action cannot- its observable consequence is an act, but the *action itself* is a process internal to the agent.

PP models are often used for modeling perception, or elements thereof, such as receptive fields (RFs) in the visual system. Rao & Ballard [22] make use of a tree-like structure to model

receptive fields in the visual cortex, and attempt to show additional *extra-classical effects* resulting from processing in cortical RFs. These extra-classical RF effects pertain to the modulation of RF responses by extension of sensory input to surrounding RFs, and they can be seen across the visual cortex [23, 24]. They implemented their model with a neuromorphic hierarchical structure in which the feedback connections propagate predictions from higher order-units to lower-order units, and the feedforward connections propagate the prediction error that arises from these predictions. When the model was exposed to natural images, they found subsets of their model that showed extra-classical RF effects, implying that the modulatory effects of extra-classical receptive fields are not merely a feedforward mechanism, but are influenced by feedback phenomena allowing the visual cortex to encode common image features.

In the effort of modeling motor control, PP theories have been proposed considering that movement actions stem from proprioceptive predictions made in order to minimize surprise, or free energy [25]. This idea of actively changing sensory inputs to minimize free energy differs from the way perception systems minimize prediction error, which is to change predictions. This stems from the idea that perception cannot change outside influences on an agents, but action can; Friston calls this *active inference*. This can be likened to *interventions*: taking action to force observations into a desired direction. Adams, Shipp & Friston expand this theory further, explaining the counter-intuitive mirroring of connections in models of the motor cortex compared to those of the visual cortex [26].

Modeling the cognitive processes underlying motor control involves dealing with interactions between variables (i.e. muscles that potentially act on the same joint) influencing the same spatial variables at once. As such, the Bayesian architecture in the present research must be able to predict the effects of such interactions when contemplating an action. Lewis introduced a theory for Counterfactual inference (CI) [27]. With decades of support for this theory, CI has grown into a family of inference methods used to perform actions, allowing acts of *intervention* upon the observed world to be premeditated based on belief, rather than observation. For any cognitive agent to be able to use modulated antagonistic muscles acting on several *Degrees of Freedom* (DOF) at once, it should be able to ‘imagine’ the effects of these interactions during the planning stage. CI provides methodologies fit for reasoning about these interactions within an action. The general workings of CI are as follows:

An agent A reasoning over a set of variables V with, among others, variables $X \in V$ and $Y \in V$. A has observed neither X nor Y , but (for whatever reason) would like to infer what would happen to the probability distribution over the values of Y if A were to set the value of X to x . The new distribution of Y cannot be accurately found by conditioning on $X = x$, since this is not an observation, but an interventional action by A. Thus, A cannot assume the new probability distribution over Y to be $P(Y|X = x)$. A must instead determine the effects of acting upon X ; *doing* $X = x$, or: $P(Y|do(X = x))$.

The distinction must be made between this so-called *do*-operator [28], and conditioning on its parameters. Support for the necessity of a *do*-operator can be exemplified from the problems that would arise when counterfactualizing by way of conditioning:

”[...] *workers should never hurry to work, to reduce the probability of having overslept; students should not prepare for exams, lest this would prove them behind in their studies.*” [18, p.108-109]

The practical application of CI in Bayesian inference for modeling such complex cognitive tasks as motor control has, to our knowledge, not been explored in practice. As such, it is reasonable to test not only the hypothesis that coactivation influences motor-skill acquisition, but also the viability of CI in motor control applications. Exact implementations and of the *do*-operator have been the subject of some debate within the Bayesian community. Pearl [29] proposes a mathematical account of the probabilistic mechanisms underlying CI as action to provide a concrete function as the *do*-operator. The proposed method, dubbed *Imaging* or *Bayesian Imaging*, provides an exact definition of the shift in probability mass caused by counterfactualizing a value for any variable in a model. Moreover, Pearl claims it allows for *disjunctive* CI: counterfactualizing on various disjoint values for any given variable (i.e. $X = x_1$ or $X = x_2$). Though this claim is followed by supporting arguments, Pearl also provides due warning for those considering disjunctive CI as a plausible theory of mind, stating that the underlying assumptions may be too grand to justify. Regardless of whether we have need for disjunctive CI, we can still benefit from the clarity with which Pearl describes the steps involved in imaging. The way he describes Imaging is as follows: If we have a causal model with a set of variables in which we want to speculate over the effects of an action, we divide the model’s variables into $\langle x, \mathbf{y}, \mathbf{z} \rangle$. Here, \mathbf{x} is the subject of the current action, where it is ‘set’ to $x = x^*$. \mathbf{z} is the set of variables that are in the past— they are known in the current moment. \mathbf{y} is the set of variables that are in the future, whose probability distributions will change given the action $do(x = x^*)$. Each possible set $\langle x, \mathbf{y}, \mathbf{z} \rangle$ is a so-called *world*; an instance of the causal network. Together, all worlds form the set \mathbf{W} . Before performing an action, each world w in \mathbf{W} is assigned a *probability mass* m which is equal to the world’s likelihood in the causal model:

$$m = P(x, y, z) \tag{1}$$

When the action $do(x = x^*)$ is performed¹, all worlds in \mathbf{W} with $x \neq x^*$ vanish, shifting their probability mass to the surviving worlds \mathbf{W}' with $x = x^*$. Each vanishing world w selects its *most similar* set of worlds in \mathbf{W}' : $\mathbf{S}_x(w)$. The most similar worlds are selected using a similarity measure that guarantees that all worlds w' that share a history \mathbf{z} are equally similar to w . Each world w' in $\mathbf{S}_x(w)$ receives mass from w proportionally to the prior probability of w' . This guarantees that even among equally similar worlds, the mass released by w is not shared equally, but proportionally. Each surviving world’s probability mass is the *post-interventional* probability over \mathbf{y} for that world: $P(\mathbf{y}|do(x = x^*))$.

It is important to note is that imaging does not dictate a set similarity measure for worlds, nor

¹Note that this action is part of an intervention, thus we cannot observe yet the new distributions over \mathbf{y} , and thus it is not an act. This is the distinction between actions and acts in imaging over multiple variables.

does it dictate an action selection measure; it describes the probabilistic effects of a counterfactual action. As such, CI with imaging preserves the possibility for *exploration* and *exploitation* [30], depending on the action selection criteria an agent upholds. Furthermore, the definition of worlds included in the use of imaging is best suited for causal models with discrete variables. Dividing worlds that include continuous variables would cause an implicit discretization. Indeed, the ‘borders’ between worlds serve the same probabilistic purpose as encoding discrete variables with a granularity (i.e., the number of values a variable can take) equal to the distance between a world’s borders.

1.3 Goals and Hypotheses

Taking into consideration current research and state-of-the-art, we propose the following research questions:

- *Can we employ imaging as counterfactual inference to model complex cognitive tasks such as motor-skill?*
- *Can the onset of muscle coactivation cause an impetus motor-skill-acquisition?*
- *Can the onset of muscle coactivation in 3-dimensional reaching behaviour cause observed developmental patterns?*

The indicated strengths of imaging seem to qualify it as a good candidate algorithm for modeling motor control. However, as often seen with models of reasoning and cognition, it may be too complex to qualify as a serious cognitive theory, or it may simply generate behaviour that is not observed in the real world. We surmise the effects of muscle coactivation may at least include improvements in movement accuracy, as it is shown to be involved in precise movements in healthy adults. In earlier research into this very topic, 2-dimensional simulations of a similar system did show statistically similar developmental trends of prereaching behaviour as a result of coactivation onset [17]. However, extending movement into the third dimension adds redundancy. This redundancy brings increased complexity, which may cause the onset of muscle coactivation alone to be insufficient to generate observed developmental patterns. The idea is that learning the intricacies between interacting muscles in a redundant system poses a sufficient challenge to simple models of cognition, disallowing the onset of muscle coactivation to really make a difference.

2 Methods

2.1 Movement

The first step in testing our hypotheses is to build a system that performs these movements while incorporating said constraints. Thus, we first set out to code a movement simulation framework that could simulate movements of limbs in a 3-dimensional space. Simulation of a 2-link, 4-DOF arm was coded in *Python 3.7*. It operates in a 3D space that encodes objects positions as a set of x,y, and z coordinates. The following sections will describe the characteristics of the arm model, explain the mechanics behind movement in this model, and cover the mechanical constraints implemented by emulating the muscular system.

2.1.1 Arm Model

It is simple to imagine a simulation of a human arm: It consists of the upper arm and the forearm, which are connected by a the elbow. The upper arm is connected to the shoulder, which we can assume to take a fixed position and orientation in the simulated space, based on von Hofsten’s experimental set-up. Those with an ambitious imagination may also include a hand, which connects to the forearm by the wrist joint. The upper arm rotates along its shoulder joint, and the forearm rotates along the elbow joint. Due to the limited contribution that movement of the hand and its connected wrist joint make towards reaching distance and direction, we opted against including them, instead lengthening the forearm by 20%.

Building a simulation to match this and allow for realistic motion forces us to take into account some additional constraints:

In order to move based on muscle (co-)contraction, each DOF must be articulated by a set of two muscle groups.

The joints in the human arm have their limits: One cannot bend their elbow backwards, nor does the shoulder joint allow us to scratch our own backs. Additionally, the forearm cannot intersect the upper arm, blocking the elbow joint.

Any rotation of the upper arm moves the forearm’s position and orientation, affecting the way it moves in space by rotating along the elbow joint.

The first issue is solved rather matter-of-factly: each DOF in the arm is actuated by two muscles, which are mediated by a coactivation coefficient (CC). The exact mathematics behind this process will be covered at the end of subsection 2.1.3. The net force on the joint in question is transformed into a joint angle between pre-defined bounds for each DOF. These bounds form the solution to the second issue — allowing each DOF to function within its own physical limits. The bounds of motion were based on results found by Rosen et al. [31], and take into account both types of limits (i.e. limits are either a hard rotation limit, or are blocked by limbs). The third problem outlined above describes the term *kinematics*: If we rotate each joint j in the set of joints

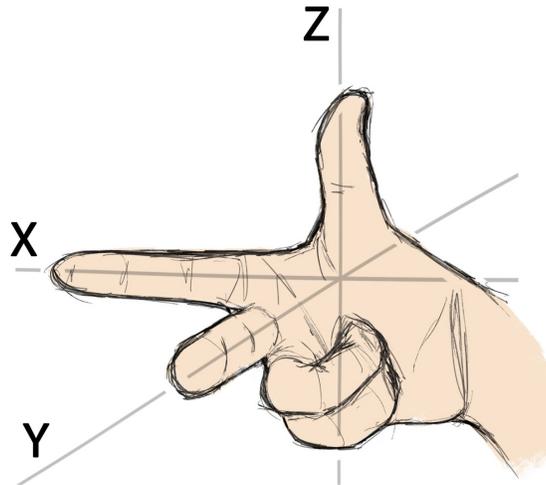


Figure 3: The right-hand rule commonly used in robotics and engineering approaches. It is easily remembered by holding one’s right hand square in front, and pointing the index finger, middle finger and thumb in orthogonal directions, indicating the x,y, and z-axis respectively. Curling the index finger in the direction of the middle finger represents motion along the y-axis. Some accounts of the right-hand rule interchange the y- and z-axis. We uphold the version as shown here throughout this research.

J connected by links (or vectors), around axes $[x, y, z]$ by angles θ_{xj} , θ_{yj} and θ_{zj} , then what are the new states of each j in J ? More simply put: If the joints in limb A do X , then what will be the new state of A ? This subject will be covered in subsection 2.1.2.

The simulated arm consists of two segments, which we will call *links* L1 and L2 for the upper arm and forearm respectively. L1 is jointed at (i.e. can rotate around and translate from) the origin of the representation space for simplicity. L2 is jointed at the end-point of L1 — it’s *frame of reference* is L1. The vectors representing L1 and L2 indicate positions in a 3D space based on their reference frames, with axes ordered $[x, y, z]$ following the right-hand rule used in many standard robotics applications such as ROS². Figure 3 shows an easy method to visualize this coordinate system, explained by the right-hand rule.

2.1.2 Kinematics

The subject of kinematics is a common problem in many engineering and robotics applications. Indeed, it is useful to be able to calculate the effect force inputs have on the movement and position of any moving system. More importantly, it is useful in these applications to be able to determine the right force inputs to arrive at a desired outcome state or extract a certain change in states. This process is aptly named *inverse kinematics* [32]. In contrast, our aim is to provide a kinematic system, and let our model infant ‘figure out’ the dependencies by itself. Therefore, we require the simulated system to be able to calculate the transformations of the simulated arm

²ROS: Robotic Operating System. <http://www.ros.org/>

given joint inputs. Additionally, our aim is not to calculate and learn the entire temporal profile (i.e. translational and rotational speed) of movements; we are merely interested in final posture of the arm, and the position of the hand; the *end-effector*.

Processing motions of the links in the *multi-link system* describing the arm model follows a set of matrix transformations imposed on the vectors representing the links. This allows for rotation along all three axes, and for translation within the space described by these axes. First, we will explain translational transformations, as they are the easiest to describe and understand. Then, we will describe rotational transformations along each axis, and subsequently describe how these are combined to allow for a full (*homogeneous*) transformation of a link in 3D space. Lastly, we will describe how these transformations can be propagated through a hierarchy of links in a multi-link system to explain the mechanics of the arm model. The following methodology is based on material from [33], and verified with reports of comparable methods in [34].

Translational motion of a *positional vector* $[x, y, z]$ in space can be defined as an offset of the very same vector. For convenience, we will not treat this as matrix additions, but matrix multiplication. The matrix encoding the translation will be referred to as the *translation matrix*. Calculating translation using the translation matrix will go as follows:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} x + d_x \\ y + d_y \\ z + d_z \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & d_x \\ 0 & 1 & 0 & d_y \\ 0 & 0 & 1 & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (2)$$

Equation (2), and the contents of its matrices is based on the following reasoning: The translation matrix must follow certain rules to be useful when combining multiple transforms later on, which also affect how we treat our positional vector $[x, y, z]$. The two main criteria that shape the matrices in equation (2) are as follows:

- Translation of a vector v by 0 unit distance (i.e. the vector stays the same) requires multiplication with translation vector t to return v . Thus we need to construct t such that no translation (t_0) yields $vt_0 = v$. t_0 is easily found — the identity matrix I of v does precisely as described.
- When performing any non-zero translation, the translation vector t must retain the shape of I , but contain translation distances inserted such that the translation distances are applied only to their respective axes. However, if we were to keep $v = [x, y, z]$, then I would be a 3×3 identity matrix. Subsequently, inserting the translational distance for the z -axis d_z into I would contradict the previous condition; $t_0 \neq I$, since d_z must be placed on the diagonal of I . To satisfy both conditions, we must extend both v and t to contain an extra dimension. Appending the value 1 to v to make $v = [x, y, z, 1]$ would extend I (and thus, t) to a 4×4 matrix, with space for inserting d_x , d_y , and d_z , without interfering with the first criterion.

The added dimension to both the positional vector and the translation vector do not affect the outcome values for any transformation we apply. Additionally, this appended value holds no

meaning for the position or orientation of a given link. As such, we can ignore the added 1 in the positional vector.

Rotation of a vector along any axis in a 3D space is akin to rotating a vector in a 2D plane that excludes said axis. Rotation of vectors on a 2D plane by θ° is described by:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (3)$$

Equation (3) follows the right-hand rule in that rotation using positive angles using this formulation is performed counterclockwise on the given plane. Note that, when used to rotate any vector by 0° , the *rotation matrix* is a 2×2 identity matrix (since $\sin(0) = 0$ and $\cos(0) = 1$). Indeed, the first criterion for translation matrices also applies to the rotation matrix. This property is what shapes the rotation matrices in whichever dimensions we require. However, the matrix as used in equation (3) cannot be used in our method for rotation in 3D, since they do not incorporate the z-axis. Additionally, when composing a rotation and translation later on, we require the rotation and translation matrices to share a 4-length dimension. When composing multiple rotations and a translation, we require all matrices to be 4×4 in keeping with the translation matrix. As such, we will shape the rotation matrix such that any non-rotation (i.e. 0°) forms a 4×4 identity matrix. For rotation along each axis, we define a separate rotation matrix that rotates a 2D plane along that axis. For each axis, we insert the *sin* and *cos* of the rotation angle into the rotation matrix at the dimensions that indicate the axes of the rotated 2D plane. The resulting rotation equations take the following forms for rotating along axis x, y and z respectively:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (4)$$

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (5)$$

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (6)$$

Much like equation (3) describes a rotational transformation of a vector in 2D, we can combine the rotation matrices in (4), (5), and (6) into a form that describes a rotational transformation of a vector in 3D space. This combination of all three dimensions can be written as follows:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (7)$$

Equation (7) assumes that we apply rotation transformations in the order x, y, z, but these rotation matrices can be swapped around to alter this order. For purposes of continuity and clarity, we perform these operations in the order our axes are mentioned, which follows equation (7). Combining equations (2) and (7) yields the following equation, showing us how to simultaneously rotate and translate a vector. We call this the *homogeneous transformation*:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & d_x \\ 0 & 1 & 0 & d_y \\ 0 & 0 & 1 & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (8a)$$

Where we can define a matrix H for the homogeneous transformation by:

$$H = D(d_x, d_y, d_z)R_x(\theta_x)R_y(\theta_y)R_z(\theta_z) \quad (8b)$$

which in (8a) is defined as:

$$H = \begin{bmatrix} 1 & 0 & 0 & d_x \\ 0 & 1 & 0 & d_y \\ 0 & 0 & 1 & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (8c)$$

Equation (8) shows us how to perform the homogeneous transformation on a single link, and provides us with the mathematical definition of the homogeneous transformation matrix H . Since the arm is a multi-link system, we also require an approach to propagate changes in global position and orientation through the hierarchy of links. While our upper-arm link $L1$ expresses its endpoint relative to the world (its $[x, y, z]$ values are the same as in our global coordinate system, or $L1 = L1^G$), our forearm link $L2$ expresses its endpoint relative to $L2$ ($L2 = L2^{L1}$). Thus, the end-effector position resulting from transformations performed on the system $L1, L2$ cannot be expressed directly from $L2$. Instead, we must define its position in $L2$ relative to the global positioning through $L1$. We can derive the global expression of $L2$ (denoted $L2^G$) as follows:

$$L2^G = L1 + H_{L1}H_{L2}L2^{L1} \quad (9)$$

Using the concepts of kinematics described above, we can define movements of the arm model concretely within our 3D coordinate system.

Rotational transformations in both links of the arm model are calculated from the arm model's muscle activation. The degrees of motion are calculated within each DOF's bounds proportional

to the net muscle activation within its bounds. That means that, after applying the effects of muscle coactivation (as will be described in subsection 2.1.3), the net force applied by each agonist-antagonist pair defines the angle θ of one rotational DOF in the connected link. The upper arm L1 can be rotated along all three axes, whereas the forearm link L2 is only rotated at the elbow along the x-axis, as the elbow only has one DOF. Transformation matrices H_{L1} and H_{L2} are calculated from the input angles, and the end-effector position is calculated in a global reference frame following equation (9). This provides us with a hand position in space as a function of muscle activations and a CC.

2.1.3 Muscle Coactivation

Implementing a model of the muscular system allows us to view muscle coactivation and reciprocal innervation (or reciprocal inhibition) are two sides of the same coin: Whenever a muscle group tenses up, what does its antagonistic muscle group do? Full muscle coactivation would imply that antagonist muscle groups contract at the same time without inhibition. Equal amounts of agonist and antagonist muscle activation is rare in healthy adults, where partial muscle coactivation is often used to facilitate high-precision movements. On the other hand, full reciprocal innervation implies that, when tensing a muscle group causes an intended movement, its antagonistic muscle group relaxes fully to facilitate the movement. By these definitions, we define both terms in a combined CC, scaling from 0 to 1, where 0 denotes full reciprocal innervation, and 1 denotes the *fraction of* full muscle coactivation. This distinction is crucial: the coactivation scale does not dictate antagonist activation based on agonist activation. Otherwise, an agent would relegate control over half of its muscles to a single control-related intent. Instead, the CC encodes a measure of reciprocity between the agonist and antagonist muscles. For example:

If an agonist-antagonist pair A, B receive muscle inputs 0.8, 0.3 respectively, this means that A is tensed up 80% of its maximum capacity, and B is tensed 30% of its maximum capacity. If paired with a CC of 0, B would be tensed by 0% of its maximum capacity instead, since a CC of 0 indicates full reciprocity. However, with a CC of 1, B would be tensed by the full 30% indicated. A CC of 0.5 would lead to B tensing up by 15% of its maximum capacity, and so forth.

This system allows our agent to activate all muscles individually, while retaining control over conflicting activity. Such conflict control allows use of coactivation for precise movements in each DOF separately, while only making use of one CC value per movement. We can define the net force for each DOF given agonist input A , antagonist input B , and coactivation CC as follows:

$$F(A, B, CC) = A - CC \times B \quad (10)$$

However, this requires us to pre-emptively define inputs A and B for the agonist and antagonist respectively. In order to make this decision implicit, we assume that the most strongly activated muscle group is the agonist. This is simply a matter of the definition of antagonist muscles; they

are the muscles that pull a DOF in the opposite direction of the intended movement. If, for any DOF, the antagonist were to contract more strongly than the agonist, the DOF would move in the other direction. This results in a movement for which our antagonist is technically the agonist — it is the muscle that pulls the DOF in the direction of the movement. Therefore, we can implicitly determine the agonist and antagonist by the activation, and arrive to a net joint force of:

$$F(A, B, CC) = \max(A, B) - CC \times \min(A, B) \quad (11)$$

This net muscle contraction receives some normally distributed noise ($\mu = 0, \sigma = 0.05$). This noise serves to introduce some randomness to the system, promoting richness in the causal relationships represented by the cognitive model. For each DOF in the arm, this joint force can be calculated and transformed into a joint angle by scaling the joint force proportionally to the DOF’s rotation limits. When a joint angle for each DOF has been determined, the kinematic system calculates the pose and end-effector position resulting from these rotations.

2.2 Cognitive Model

In order to simulate infants that utilize the described movement system to reach for objects, we require a model of cognition to determine the arm’s desired input signals to reach any given end-effector position. We construct a cognitive agent that uses the arm model, and envelops a probabilistic causal model encoding its knowledge about the arm model’s functions. It utilizes this causal model to process spatial information of a desired action to generate inputs for the arm model.

The following subsection describes the Bayesian network used to model the causalities between action and observation. Subsequently, the way imaging is performed with the Bayesian network is explained in detail, accompanied by additional considerations regarding its use. Lastly, the entire processing timeline is highlighted, showing how inputs for the causal model are processed to produce motor output.

2.2.1 Causal Model

To generate predictions for the ‘best’, ‘most likely’, or ‘least uncertain’ muscle activations to perform a desired reach, we constructed a causal Bayesian network. This network encodes the causal and probabilistic relationships between muscle (co)activations, arm postures and spatial locations as a DAG. To allow the simulated agent to perform inference over all these variables, its structure must contain the following:

- Pairs of nodes for representing the agonist-antagonist muscles for each DOF in the arm: ag_x and ant_x for every DOF x . Considering the arm model’s specifications, the causal network must contain 4 of these pairs, for 8 muscle nodes in total.
- A coactivation coefficient node CC to encode how heavily the agonists’ activations are reciprocally inhibited by the antagonists’ activations.

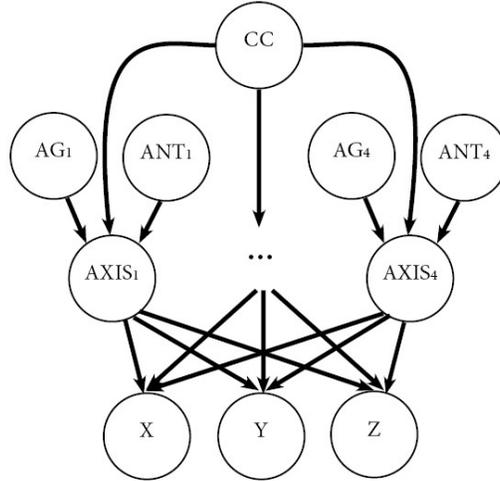


Figure 4: The causal model that codes for muscle (co)activation, DOF movement, and spatial position.

- Nodes that encode positions in each DOF, integrating activity from the corresponding agonist-antagonist pair, as well as the coactivation coefficient. Four of these nodes are required, one for each DOF.
- Nodes that combine DOF movement into spatial positions for the end-effector. Three of these nodes are required, one for each dimension on the physical space.

The resulting causal Bayesian network can be found in figure 4. Due to the discrete nature of imaging, each node is a discrete variable. However, representing these variables as binary values only allows them to be on/off switches. Encoding these variables as such does not necessarily allow for a natural model of motor control. Indeed, if muscles and CC values are only encoded as 0 or 1, the DOF nodes would only be able to encode $2^3 = 8$ different causal combinations. Such sparsity would ‘trickle down’ to the lower layers. Of course, encoding all variables as binary values to avoid probabilistic sparsity is no solution: reducing motor control to such lengths would constrain the resulting data to a point where comparing it to any empirical data is an exercise in futility. When observations over the BN are made, the internal beliefs (likelihoods over variables) contained therein must be updated to more accurately reflect past experience. Updating beliefs or hypotheses H given some evidence E is usually done using Bayes Theorem:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)} \quad (12)$$

An alternative, however, is to iteratively construct *hyperparameters* of each of these likelihoods, which can be collapsed into probabilities given any query in the BN. A hyperparameter is a parameter describing a probability distribution, such as the number of occurrences of a certain value in that distribution. One can think of a hyperparameter as a parameter used to describe prior

knowledge about an underlying distribution. This hyperparameter can be varied or updated to represent different prior knowledge about the mechanisms underlying observations. Furthermore, iterative approaches of updating a hyperparameter allows for construction of a distribution over the hyperparameter— a *hyperprior*. The practical applications of hyperparameters vary: They allow us to represent an arbitrary function in a relatively easy way. Moreover, hyperpriors constructed from hyperparameters reflect uncertainty about the correctness of a model’s distributions.

The hyperparameter’s ease of implementation in both construction and updating allows us to easily keep track of observations in high-granularity systems. Hyperparameters that keep track of observations of specific states can be collapsed into probabilities for any queried state. A belief update over hyperparameters then means to simply add an observation to whatever state was observed. Collapsing a hyperparameter to find the likelihood of a certain query simply consists of calculating the proportion of total value of the query over all observations:

$$P(q) = \frac{\sum_{q \in Q} q}{\sum H} \quad (13)$$

Where Q is a queried collection of states q , and H is the hyperparameter.

A note on complexity

The consideration of node granularity brings about a different problem pertaining to imaging: the number of worlds that must be considered when imaging over the network, and the number of actions that must be executed when imaging actions over multiple nodes in a network. Concretely, the space complexity, as well as the time complexity of imaging, quickly become problematic if the network it is leveraged on contains many nodes, or high-granularity nodes.

Assuming a BN with n nodes with consistent node granularity g , the number of worlds that must be assigned a probability mass at the start of the operation, M is g^n : The space complexity of imaging over an entire granularity-consistent network is exponential.

For time complexity, we must consider the number of actions that need to be made to image over a size- n granularity-consistent network. With each i -th action (where i is a counter of the nodes for which an action has been selected) that selects a value for a node, M shrinks to g^{n-i} . However, every excluded world must consider every surviving world to determine how to spread its probability mass. For each action, $g^{n-1}(g-1)$ worlds are excluded, all of which must consider the remaining g^{n-1} worlds. Multiplying the number of excluded worlds by the number of worlds each must consider, the complexity of performing an action in a model with n nodes with granularity g is $g^{2n-2}(g-1)$. More generally, inferring an action as imaging within a network is also an operation that increases in run-time exponentially in the input size: it is an intractable operation that is NP-hard. Intractability is a problem for theories of mind due to its inherent assumption of computational power or available time. For any form of computer, whether it be a laptop or a brain, to finish an intractable algorithm requires a computational speed or available time that grows exponentially in the input size. For any real-world scenario to be processes, the input size of a

problem can reasonably be assumed to be sufficiently large to pose a computational problem for any type of computer. Since processes of cognition are, in most cases, split-second processes, models of cognition cannot assume the unreasonably long computation times necessary to solve intractable problems. By extension, the speed at which computations of intractable algorithms must be made in the brain to restrict the algorithm’s run-time also grows to unreasonable proportions. The apparent complexity of imaging highlights a crucial problem in applying it as a theory of cognition. There are measures that can be taken to reduce the complexity of imaging over a network relative to the input size. These measures will be described in the following section, 2.2.2.

2.2.2 Imaging over Limb Mechanics

In order to be able to perform imaging over the variables involved in the mechanics of moving the arm, steps need to be taken to guarantee viable space- and time-requirements. The most obvious answer is to only apply imaging to parts of the causal model, separating nodes from one another, allowing for an action to include fewer imaging steps. The question then becomes which parts of the BN to split apart from one another, and how to do so. Naturally, this splitting of the BN into smaller subnetworks should respect the causal influences between parts of the network. The most straightforward way of causally separating subnetworks is called *d-separation* [35]. The d-separation criterion identifies whether two collections of nodes X and Y in a causal model are conditionally independent (d-separated) or conditionally dependent (d-connected), given some evidence E . This means that, for any evidenced nodes E , making an observation over any node in X can affect the probability distributions in X , but not in Y , and vice versa. Basic d-separation criteria can be found in figure 5, showing how evidence in causal structures d-separates and d-connects variables. This is the basis of conditional dependency— d-separated sets of nodes are *conditionally independent* of one another.

D-separating the causal model into smaller subnetworks cannot be done indiscriminately. In fact, although d-separation is a prerequisite for subnetworks to be imaged over independently, not all d-separated subnetworks are functional choices for imaging. Firstly, considering d-separated single nodes as subnetworks, and imaging over them as such does not decrease the algorithm’s complexity. Secondly, imaging in a top-down or bottom-up manner influences which d-separation criteria are fit qualifiers for subnetworks. Although an action is not probabilistically equivalent to conditioning, it does d-connect or d-separate subnetworks. Thus, imaging in a top-down fashion allows for d-separation of subnetworks only after the procedure has already started, and a set number of worlds is already considered. One could argue that, if the first imaging action d-separates the necessary subnetworks, one could separate them beforehand, and perform the same first action in all networks. Considering this idea in terms of our causal BN, one can conclude that top-down imaging can be used with separate subnetworks if it first performs an action on CC, which will d-separate all four axis nodes. However, top-down imaging cannot infer the best arm input parameters for any target position. Indeed, if our cognitive model aims to provide input parameters for the arm model given an X,Y,Z position, it must perform bottom-up inference over

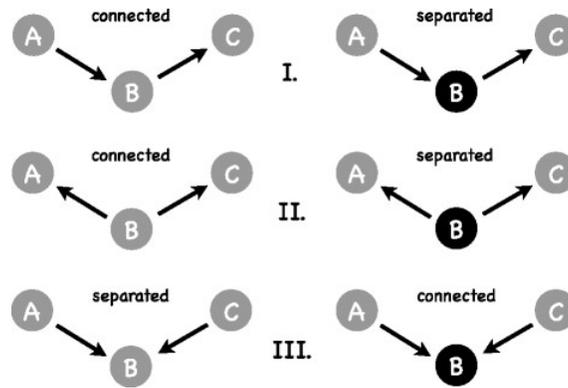


Figure 5: D-separation criteria for variables A and C in three different basic network structures: **I**): When B is not observed, A and C are d-connected through it. When B is observed, A can no longer influence C through B, and they are d-separated. **II**): When B is a common parent to A and C, they are d-connected since observing one provides information about B, influencing knowledge of the other. However, observing B d-separates A and C. **III**): When A and C have a common child in B, they are d-separated when no observations are made. However, observing B d-connects A and C. From Shriprakash, 2016 [36]

its nodes to find the input parameters that are most likely to match the positional values.

As with top-down imaging, bottom-up imaging in our causal model would require conditioning (or first acting) on CC to d-separate the axis nodes, and then imaging bottom-up. However, the first imaging action d-connects the axis nodes by performing an action on a common child node. This does not aid in solving the complexity of imaging over our causal model; the architecture as seen in figure 4 would not allow for d-separation of any subnetworks in an informative manner. As such, we made a structural change to the network: we removed the X,Y and Z nodes. The removal of these nodes is described in 2.2.3. It provides the additional benefit that spatial information no longer requires discretization to allow for imaging. Indeed, separating spatial information and cognitive function allows removes the cognitive model’s restrictions on spatial variables. The resulting BN can be found in figure 6.

2.2.3 Replacing X,Y and Z

With the positional nodes gone, the axis nodes (and their muscles) are d-separated from one another, and imaging bottom-up through them can be done separately without d-connecting them. Removing the positional nodes from the causal model can be likened to representing movement cognition in a body-oriented frame of reference. Movements, and by extension, DOF configurations, are coded relative to the body. The translation from target positions in space should also be coded in a body-oriented frame, and input to the DOF configurations through that. To represent this body-oriented translation from spatial positions to body positions, the architecture needs a way to determine a ‘goal posture’ that coincides with spatial positions.

Calculating a limb posture coinciding with a given end-effector position is the inverse of what

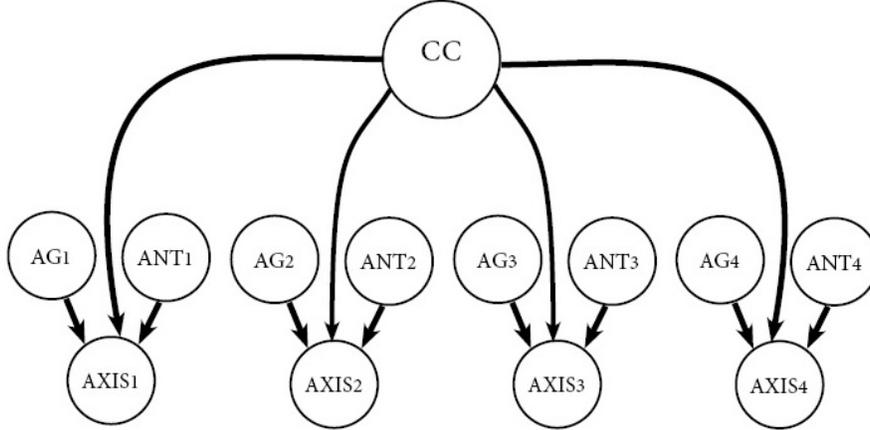


Figure 6: The causal model that codes for muscle (co)activation and DOF movement. By removing the common child relation between the axis nodes, this BN allows for easy separation of subnetworks for a bottom-up imaging procedure.

our arm model does: The arm model calculates the end-effector position given a set of joint positions, which we call kinematics. Finding the analytic solution to inverse kinematics in a redundant system in 3D is a non-trivial problem. However, we can approximate a posture solution (i.e. goal states for the axis nodes) given an end-effector position using several methods. Indeed, many methods for finding or approximating an inverse kinematic solution have been proposed in the past. However, the current model requires a relatively simple solution. Buss provides several simple iterative approaches that provide good approximations of an analytic solution [37]. The basis of these methods is the *Jacobian* J , a transformation matrix that contains all relevant rotational and translational information pertaining to the end-effector [38]. J can be determined as a propagation of homogeneous transformation matrices determined in equation (8), and is used in forward kinematics. For inverse kinematics, one should find the inverse Jacobian— finding this analytically is, again, non-trivial, but we can substitute it for something else. For example, Buss shows that using a weighted *transposed* Jacobian to iteratively approach an inverse solution can provide a good approximation given the right weighting. A more complex, but more flexible solution that Buss highlights, is the Damped Least Squares (DLS), which was first used for inverse kinematics approximation by Nakamura [39] and Wampler [40]. DLS is more stable when approximating singular inverse solutions— when the target end-effector position is at or beyond the possible reach length, where there is only one joint posture that provides that end-effector position. Due to the out-of-reach position of targets in von Hofsten’s experiment [1], this method was selected as the approximator due to this stability.

2.2.4 Processing Timeline

The timeline of processing a target-oriented reach is a composite of the various methods as described above. These methods are used in the same order for every reach, but with different targets. When a target is processed into desired node states, the following happens:

A target's position is sampled from a normal distribution with a confidence interval equal to the target's width. This sampling is a rudimentary form of visual noise, and serves to introduce variability into the system's inputs. The sampled target X,Y,Z position is processed by the DLS approximator until an approximation is determined within an error bound of 1 unit distance. For each iteration of the DLS approximation, only the arm model's four DOFs will be retained. After a sufficient approximation has been found, the desired joint rotations are transformed into their respective desired axis node activations.

Once the desired axis node activations are known, the current CC value is determined based on the agent's 'age'. The basis for this value will be highlighted further in section 2.3.3. The agent will then split its causal model into four subnetworks: one of each axis node with its corresponding muscles, and CC as its parents. For each subnetwork, the imaging procedure starts: Possible worlds are initiated, and receive their probability mass. For all subnetworks, the first imaging action that is performed is setting CC to the predetermined value, which is the same for all four subnetworks. This first action is dictated by the inference ordering; it d-separates the subnetworks, and only by imaging over CC first is the separation of subnetworks a valid operation. Probability mass from excluded worlds is spread to surviving worlds. Subsequently, each subnetwork's desired axis value is set using an imaging action, and probability mass is again spread. Once the CC and axis values have been set, the surviving worlds' total probability mass dictate probability distributions over the muscle nodes within the current action. The action selection criterion is not dictated by the imaging procedure. Based on the distribution over node values dictated by remaining worlds, values are sampled with their probability. In contrast to selecting the most likely action, this allows for some exploration of the action space when there is large uncertainty (high entropy) in the probabilistic effects of an action. A value is first selected over the agonist node, and probability mass is spread for the last time. Lastly, a value is selected over the antagonist node, after which a single world remains.

The selected world for each subnetwork dictates values for that subnetwork's muscle nodes. The selected values and CC are input into the arm model, which calculates a reach including its random muscle noise. Note that the executed reach may differ from the desired reach by virtue of prediction error and muscle noise.

2.3 Simulations

Much like the procedure for calculating a reach, the simulation protocol of infants always is always subject to the same development and testing. This section describes the various phases of a simulation.

2.3.1 Development of Muscle Coactivation

Over the course of development, an infant’s tendency to coactivate antagonistic muscles during a reach follows a set trend: In weeks 1 and 4, there is only agonist activation ($CC = 0.0$). Starting at week 7 there is strong inhibition from the antagonist muscle ($CC = 0.8$), which decreases back (through $CC = 0.6$ and $CC = 0.3$ at weeks 10 and 13 respectively) to an established ‘adult’ level ($CC = 0.2$) from week 16 onward. In every motor activity, be it learning phases or test-phase reaching, the coactivation coefficient that corresponds to the infant’s ‘age’ is selected. These figures, although based on reports from Gatev [3] and Spencer & Thelen [41], are still arbitrary to a degree. Indeed, reports of ‘proportional activity’ conform to these numbers, but are reported separately. In favour of keeping the architecture consistent, we opted to keep this abstraction in place. Further considerations and discussion on this topic are presented in section 5.

2.3.2 Motor Babbling and the Learning Phase

When an infant is generated, the beliefs implicit in the BN are not yet developed. In order to develop probabilistic and update relationships in the BN, we employ the following learning phase procedure: The infant’s CC is set to a value that corresponds to its age. For each muscle in the arm’s system, a random decimal value between 0 and 1 is sampled from a uniform distribution. The corresponding reach is calculated, and the causal model’s hyperparameters are updated to reflect the new observation. In order to develop a new infant’s beliefs, this random learning procedure is executed with 5000 trials. Such high numbers of random trials are often used in developmental robotics to jump-start motor learning in simulated or robotic agents. This is generally dubbed *motor babbling* as a blanket name for such random learning initiation, based on the equally named phenomenon seen in infants [42].

After the motor babbling phase has concluded, each 3-week age group elicits a developmental phase to learn to deal with its available system, and a testing phase to test its new capabilities as von Hofsten did. These phases will be briefly explained in 2.3.3 and 2.3.4.

2.3.3 Developmental Phase

For each 3-week age group, 2000 development phase trials are executed much like in the motor babbling phase: the coactivation coefficients are determined on age, which can now be values other than 0.0 (note that in the motor babbling phase, $CC = 0.0$ always holds), leading to new experiences for the agent to learn from. The number of learning trials is kept consistent for each age group to prevent contamination of results. After all, if certain age groups perform more learning trials than others, the difference (or lack thereof) in behaviour may be caused by this learning gap, rather than by mechanistic changes.

The developmental phase learns from such random activity rather than repeatedly calculating informed reaches for several reasons. First and foremost, learning by way of testing the outcomes of random activations builds probabilistic relationships that represent observed causalities. After all,

performing random reaches still adheres to the constraints set by the arm model. These constraints are what the agent must learn to deal with. Secondly, calculating a reach with CI is an operation that requires too much run-time to execute hundreds, let alone thousands of times in a row, even despite measures taken to reduce overall complexity. The extent of this run-time issue will be made concrete in 2.4.

2.3.4 Testing Phase

For each age group, the learning phase is followed by a testing phase, in which von Hofsten’s set-up is mimicked to gather comparable data. The simulated infant is assumed to be ‘strapped in’, locking its position in space. A set of 70 random targets are generated within a certain out-of-reach range: z-axis positions were normally distributed ($\mu = 0.0, \sigma = 0.5$). Values on the x-axis were uniformly distributed between -30 and 30 unit distances. Values on the y-axis were normally distributed at a distance from the reachable space threshold ($\mu = 5, \sigma = 2$).

During the testing phase, the 70 targets are presented to the infant. For each target, the infant infers a desired state to reach for that target, and predicts the best motor signals to reach the desired state. That is, given a desired state, a prediction is made using imaging for each muscle node to determine which value for that muscle is most likely to lead to the desired end-state. The most likely values are chosen to be executed by the motor system. When a reach is executed, data is gathered regarding the infant’s performance, and the number of extended reaches.

2.4 Data Acquisition

The testing phase gathers data regarding the following metrics:

- The total number of extended reaches.
- The sum over prediction errors in the causal model’s axis nodes per trial.
- The reach error between the end-effector and the reaching target per trial.

For determining the number of extended reaches, a reaching threshold is determined based on the unit length of the arm, to match von Hofsten’s metric. Despite the fact that von Hofsten does not mention the extent of his chose threshold, one can assume that it was chosen to only be exceeded when at least the elbow is flexed or extended such that the hand would point toward the target. As such, we can assume it must be at least the length of the forearm. So as not to exclude too many reaches, we set the threshold slightly over the length of the forearm, to a unit distance of 12 along the y-axis (forward in the movement space). This means that extending the arm downward (where no target will be present) would not be counted as an extended reach. Whenever the y-axis value of the end-effector position for a performed reach crosses the extension threshold, it is added to the number of extended reaches for that testing phase. Each infant gathers the number of extended reaches per testing phase, for 7 numbers per infant.

Prediction error for an axis node is the sum of differences between axes' predicted values, and their observed values in the executed reach. Both the predicted and observed values fall within the same numerical range, which is consistent over all joint axes. As such, we can sum the prediction errors for each axis to come to a total prediction error. For each trial, the prediction error can be defined as:

$$\sum_{a \in A} |a_d - a_o| \quad (14)$$

Where A is the collection of axis nodes, a_d is the desired value for node a , and a_o is the observed value for the axis in the executed reach. While the desired values for each axis node are restricted to the possible values the nodes can take, this does not hold for the observed values. Due to the nature of the kinematic calculations and the added muscle noise, the observed values for each axis can be any real number between 0 and 1. The possible range for the prediction error is any real number between 0 and 4. Prediction errors are gathered on a trial-by-trial basis, totalling 70 error results per testing phase, for 490 error results per infant.

The reach error is calculated as the difference between the end-effector position and the target position. Since both of the positions are vectors with shape $[x, y, z]$ for target t and hand h , the reach error can be defined as:

$$error = |[x_t, y_t, z_t] - [x_h, y_h, z_h]| \quad (15)$$

Since test phase targets are always out of reach, the lowest the reach error could be for an executed reach is 1, but on average it will be 5. Reach errors are gathered on a trial-by-trial basis, totalling 70 error results per testing phase, for 490 error results per infant.

Initial experimentation with node granularity up to 7 values per node showed little improvement in prediction- and reaching accuracy beyond a node granularity of 5. Run-times, however, increased significantly beyond this granularity. With a granularity of 5, the projected run-time per infant was approximately 30 minutes, taking approximately 3.5 seconds per reaching action. Granularities of 6 and 7 increased the time per action to approximately 9.8 and 27 seconds respectively, leading to projected run-times of 85 and 228 minutes respectively. Taking into consideration this increase in run-time compared to its low pay-off, we decided to run all simulations with node granularities of 5 across the causal models.

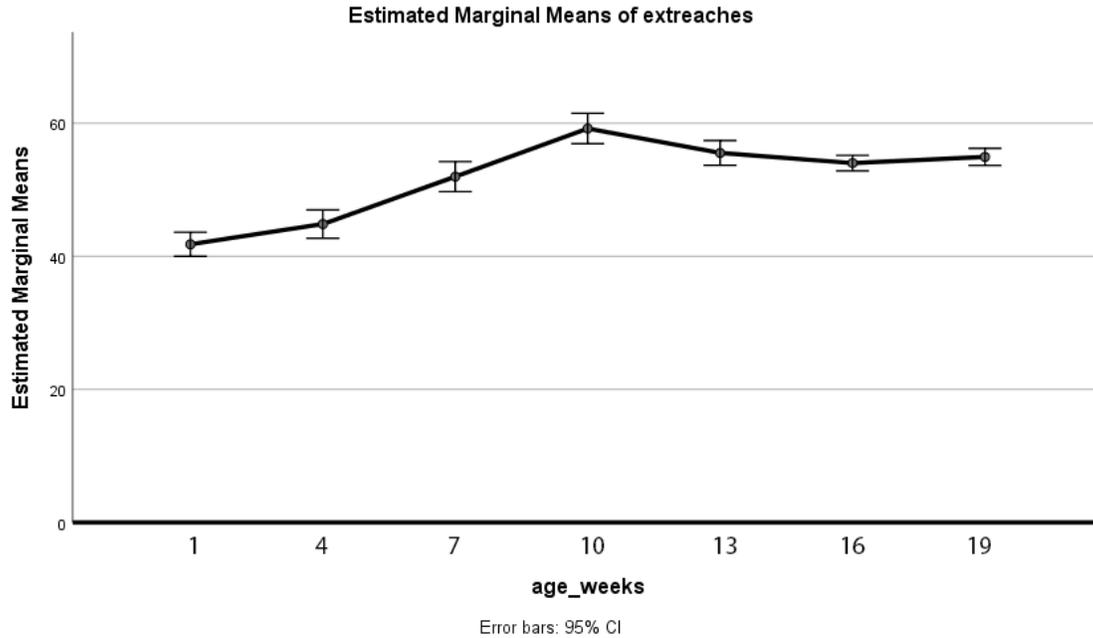


Figure 7: The mean number of extended reaches during the test phases for each age group.

3 Results

The results pertaining to the extended reaches can be found in figure 7. A one-way ANOVA was performed to test the effect of age as a within-subject factor on the number of extended reaches. The effect of age on the number of extended reaches was statistically significant ($F(6, 174) = 46.9, p < 0.001$). A post-hoc Bonferroni test revealed statistically significant differences only between consecutive weeks 3 and 7 ($p < 0.001$), and weeks 7 and 10 ($p = 0.01$). Furthermore, differences between non-consecutive age groups were statistically significant between weeks 1 and 4 and all other age groups ($p < 0.001$ for all differences). Additionally, the differences were statistically significant between week 7 and week 13 ($p < 0.05$) and week 10 and 16 ($p < 0.005$). Note that the statistically significant difference between weeks 3 and 7 and weeks 7 and 10 are caused by increases in reaching extension, rather than a (temporary) decrease comparable to the one found by von Hofsten.

To test for an effect of age on the prediction error, a one-way ANOVA was performed. A visualization of the tested data can be found in figure 8. The effect of age on prediction error was found to be statistically significant ($F(6, 174) = 1163.577, p < 0.001$). Further investigation using a post-hoc Bonferroni test revealed the following: Week 1 was statistically significantly different from all others ($p < 0.001$) except for week 13 ($p = 0.987$). Week 4 was found to differ statistically significantly from week 1 ($p < 0.001$), 7 ($p < 0.001$) and 10 ($p < 0.001$), but not from others. Week 7 and 10 were found to differ statistically significantly from all other age groups ($p < 0.001$ for all

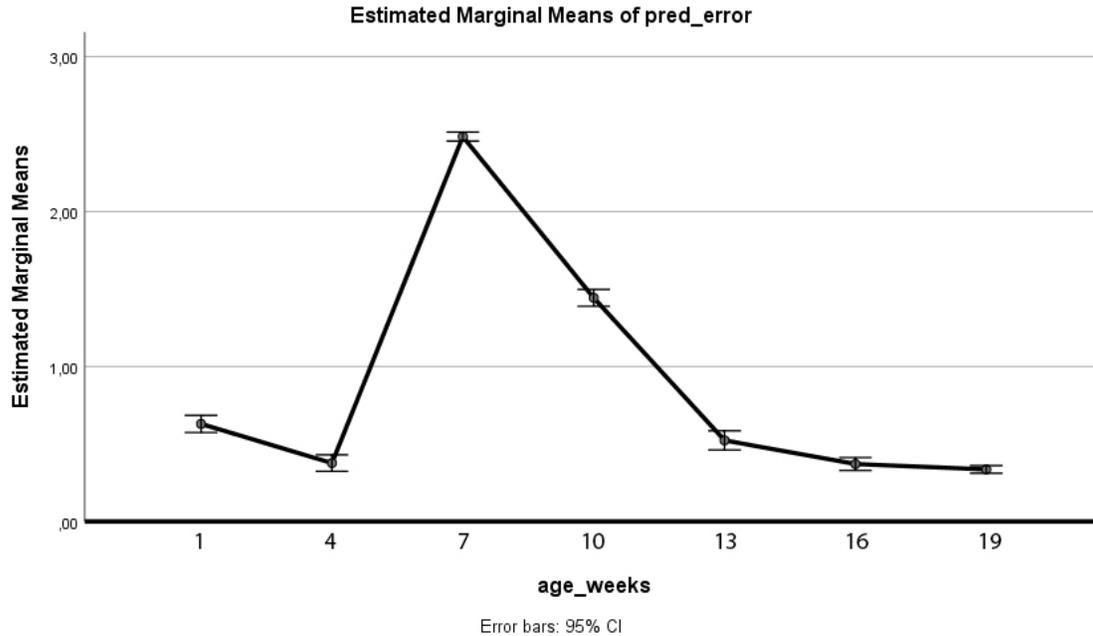


Figure 8: The mean summed prediction error during the test phases for each age group.

comparisons). Week 13 was found to not differ statistically significantly from weeks 1 ($p = 0.987$) and 4 ($p = 0.121$), but it did differ significantly from all other age groups ($p < 0.001$ for all comparisons). Weeks 16 and 19 were found to differ statistically significantly from all age groups ($p < 0.001$) except week 4 ($p = 1.0$), and from one another ($p = 0.155$).

Notable effects here are the statistically significant differences were found between all consecutive age groups ($p < 0.001$), except weeks 16 to 19, where no statistically significant difference was observed ($p = 0.155$).

To test for an effect of age on the reaching error, a one-way ANOVA was performed. A visualization of the tested data can be found in figure 9. We observed a statistically significant effect of age on reach error ($F(6, 174) = 27.095, p < 0.001$). Further investigation using a post-hoc Bonferroni test revealed significant differences between consecutive age groups at week 7 and 10 ($p < 0.005$), and between week 10 and 13 ($p < 0.01$). When looking at non-consecutive age group differences, weeks 7 and 10 differed statistically significantly from all other age groups ($p < 0.001$ for all comparisons to week 7, $P < 0.05$ for all comparisons to week 10). All other comparisons were not statistically significant. Note that, regardless of differences between age groups, none of the age groups present an average reach error below a unit distance of 24, and after outlier removal, standard deviations stay close to 1 unit distance.

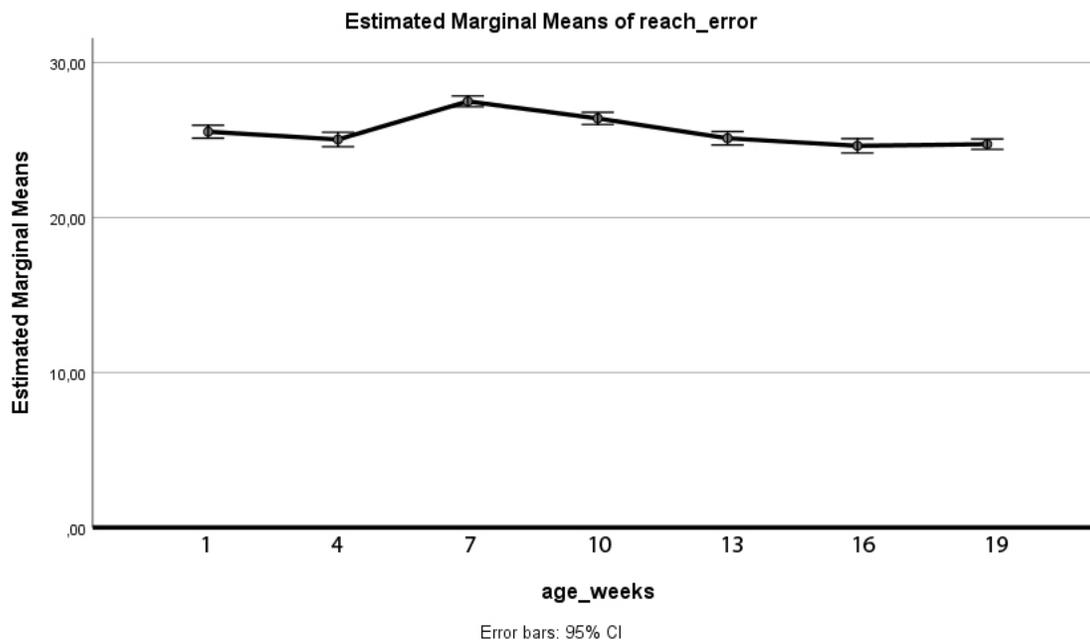


Figure 9: The mean reach error during the test phases for each age group.

4 Discussion of Results

Given the observed statistical results, we can follow up on our research questions and hypotheses. While these results can not prove or disprove hypotheses, they can provide support for or against them.

First and foremost, “*Can we employ imaging as counterfactual inference to model complex cognitive tasks such as motor-skill?*” can be answered with the most certainty. Complexity issues at the computational level alone pose a very serious issue for imaging as CI in cognitive modeling. Any intractable algorithm should raise concerns in this regard; after all, how are the operations inherent in such methods executed in the brain without requiring exorbitant amounts of time? We know that planning and execution of arm movements—which employ more muscles and degrees of freedom than our architecture—takes up only a few hundred milliseconds from when a motion is first detectable in the brain [43]. Present simulations took up over tenfold that amount of time per reaching motion, while both planning and execution were performed in highly simplified systems. Imaging as a method for CI deserves a commendation for minimizing its prediction errors given the unpredictability of its observations. Indeed, when we observe the results over prediction errors, we can see that, in relatively stable surroundings (little to no coactivation, little deviation from previous experience), the prediction error was minimized quite well. Although it does not approach 0 error, this is to be expected in a noisy surrounding. Furthermore, since the prediction error originates from the difference between an observed value (which can be any real number

between 0 and 1), and a restricted prediction value, one should expect some semblance of prediction error. Indeed, where two observations of 0.14 and 0.28 differ greatly, low-granularity causal models may represent both as 0.2. This *bin size* caused by discretization allows predictions to be done optimally given internal beliefs, and still incur prediction errors as large as half the bin size. Taking into account both movement noise and granularity-induced *minimum average error*, imaging seems to perform relatively well.

However, when we look at reaching error, we can observe that even minimizing prediction error as well as possible does not prove sufficient to learn the external task of motor control. While significant prediction error increases are indeed accompanied by significant reaching error increases (and likewise for decreases), reaching error simply does not approach the expected minimum average error of 5. The fact that such high reaching errors can arise from relatively low prediction errors shows the need for accurate prediction in motor action: The prediction error caused by node granularity and muscle noise seems small when expressed in node activations. However, translation of prediction error to joint angle differences shows an average of nearly 20° rotation error per joint. Such deviations in posture can lead to extensive differences in end-effector position, showing the necessity for low prediction error.

We can conclude that, although imaging can be used to make counterfactual predictions over nodes with good accuracy, it brings about various issues that hamper its usefulness in cognitive modeling. Furthermore, while taking complexity-reducing measures and abstracting physical systems barely allows imaging to be used as an inference method, the behavioural results show little improvement. Due to the complexity reductions, there is a limit to how accurate predictions can be, and in motor control, this is clearly an issue. This trade-off between the viability and the effectiveness of the methodology lead us to conclude that imaging in its current form does not allow for realistic modeling of the cognition underlying motor control. This shows factors that need to be taken into consideration when *formalizing* and *operationalizing* imaging in robo-havioural studies in the FOES method [44]: Imaging is a *formalization* of counterfactual imaging that proposes computational solutions to fundamental questions. The implementation, including the measures taken to *operationalize* imaging in a motor control setting highlights issues with the proposed formalization: Acquiring a combination of tractability and effectiveness requires further operationalization. Further improvements allow for *exploration* of the various implications of implementational details. *Studying* the computational and behavioural theories following this process should be the end goal.

Inspecting the statistical results in order to answer the question “*Can the onset of muscle coactivation cause an impetus motor-skill-acquisition?*” provides us with some insight, but adds considerations. The onset of muscle coactivation, in our simulations, is implemented with values that are speculative based on qualitative reports. The age that this onset coincides with correlates to an increase in prediction -and reaching error. When the coactivation coefficient approaches more mature levels the errors decrease. However, they do not decrease significantly past the level of pre-coactivation errors. Based on these results we can not claim support for the idea that the onset of muscle coactivation causes quick improvement in motor-skill. However, the low levels of

prediction error, and the low variance therein especially, do point to a benefit of low (but present) levels of muscle coactivation, which seems to be in line with behavioural studies.

Considering the restrictions placed upon reaching performance by the chosen generative model of cognition, one should keep in mind that the prediction errors may be artificially lower-bound by the uncertainty implicit in the architecture. As such, we cannot immediately conclude that our data does not support the idea that muscle coactivation or the onset thereof improves motor-skill acquisition. In fact, we should keep additional considerations in mind: The choice of generative cognitive models underlying motor action may (indirectly) affect how mechanistic constraints affect behavioural results. The mechanisms, restrictions and limits implicit in the choice of cognitive models are likely to affect the observable behaviour of the model. As such, one should take into consideration that the results observed in this study may be heavily influenced by choosing a counterfactual Bayesian model of cognition. After further operationalization of the computational theory of counterfactual inference, the results may be wholly different from those observed here. However, under the workings of imaging as a counterfactual inference theory, observed behavioural performance does not improve under muscle coactivation onset. As such, we should conclude that cognitive systems that function as counterfactual Bayesian imaging can not utilize the onset (or presence) of motor control sufficiently to experience strong improvement of motor-skill as a result of muscle coactivation onset.

Much like the previous question, *“Can the onset of muscle coactivation in 3-dimensional reaching behaviour cause observed developmental patterns?”* can be answered with additional considerations. Looking at the results for our simulations’ extended reaches, the differences should be clear immediately. at 13, 16 and 19 weeks of age, our simulated infants show significantly more extended reaches than ‘newborn’ infants, which is consistent with von Hofsten’s findings. However, around the onset of muscle coactivation, we find a statistically significant increase in extended reaches, followed by a significant decrease between week 10 and 13. It should be abundantly clear that this significant peak is the opposite of the behaviour that is observed in infants at the age of coactivation onset. The obvious conclusion one would make is that the onset of muscle coactivation in a 3D muscle based reaching model does not cause observed behavioural patterns, but it does rather the opposite. As such, these results do not support the coactivation hypothesis. However, one must again take into consideration the effect that the choice in generative causal models may have on observed behaviour. As with the previous question, one should view these results as a conditional rejection of the hypothesis that muscle coactivation causes observed developmental patterns. Indeed, we can conclude that, if cognition of motor control can be modeled as counterfactual Bayesian inference, muscle coactivation alone does not cause the behavioural patterns during human development.

A consistent theme among these conclusions should be clear: The choice of cognitive models can heavily influence behaviour exhibited by simulated agents. Although imaging as a method for making counterfactual predictions does seem to work given that measures are taken to reduce complexity, it does not suffice as a generative cognitive method if used to model complex mecha-

nisms. Due to the failings of imaging as a viable strategy for modeling cognitive functioning, one could go so far as to conclude that hypotheses regarding the effects of muscle coactivation onset cannot be confirmed nor discarded, as the used inference protocol simply does not mimic cognitive functioning sufficiently.

5 Conclusion and future work

Considering the nature of our current conclusions, further investigation into the themes of this study may provide additional insights, and may provide results that show different support than those presented here. Furthermore, the implementation of simulated experiments has forced certain decisions to be made, some of which could benefit from further inspection.

In terms of emulating von Hofsten’s experimental set-up, the choice of a reaching threshold, though an informed decision, remains an arbitrary one. The current study chose a threshold based on the simulation’s unit sizes of the arm, reasoning that reaches in the direction of a target would at least exceed a unit distance of $y = 12$. This threshold was kept consistent along all values on the x- and z-axis, as von Hofsten’s description of the threshold mentions a straight line drawn on a television screen. As exact distances are not named, we were forced to choose a threshold based on what could reasonably be called a ‘reach’. Furthermore, von Hofsten does not mention any displacement of the reaching threshold between weeks. Due to physical growth in infants, this may even have affected the relative effort needed for a reach to exceed the threshold. Since there was no mention of controlling for such an arguably difficult to balance factor, we chose to keep the threshold consistent at $y = 12$ over all simulations.

Furthermore, our simulations were one-shot trials, rather than repeated continuous operations. Von Hofsten measured extended reaches per minute, with any extended reach exceeding seconds counting as 2, and any reach exceeding 8 seconds counting as 3, et cetera. Additionally, the stimuli in the behavioural study moved in front of the infants at various speeds, allowing for various (significant) effects on the infant’s prereaching behaviour. Due to the trial-based nature of our simulations, such movement effects could not be emulated. Expanding on this architecture to be able to deal with time-streams of stimuli is a logical next step. Modeling motor behaviour as changes in motor states provides a more flexible, and indeed more natural approach to simulating motor-skill acquisition. Such an expansion can be expected to introduce new questions: What happens internally when no stimulus is present? What motivates an infant to reach for any stimuli, or to even explore and learn new skills at all? This idea of *intrinsic motivation* has been the subject of discussion in developmental robotics. The current study circumvents this issue by assuming intrinsic motivation to attempt a target-oriented reach. Predictive processing offers a working interpretation for this concept: When a stimulus enters an infant’s perceptive field, it changes observations, and causes prediction error. High rates of prediction error are unfavourable to any organism- the oscillation imply a lack of *homeostasis* [45]. The infant would attempt to minimize prediction error on both short- and long-term scales by observing as much about the stimulus as possible. The intrinsic motivation to learn stems from the prediction error minimization that implements homeostasis, thus giving the predictive infant a reason to seek out information on its stimuli. Such prediction error-motivated motor learning should follow from predictive methods if they are implemented such that they seek to minimize any prediction errors.

Presenting a moving stimulus in a time-locked series of positions would prove to be a problem for an imaging agent when one considers the required run-time for calculating a single desired posture (or change therein): The gathered data were restricted to causal models with low granularity structures due to the run-time required to calculate a single reach. With the chosen network granularity of 5, calculating a single reach took up approximately 3.5 seconds. If we were to time-lock the movement of a stimulus to emulate the stimulus presentation von Hofsten outlines, the agent would only perform one posture change every 3.5 seconds. Such low reaction speeds make it exceedingly difficult for any simulated infant to keep up with the stimulus, let alone follow its position with smooth hand movements.

Von Hofsten's findings include reports on visual fixation of the presented stimulus. These findings were important in favouring the coactivation hypothesis over the approach-and-withdrawal hypothesis, as they contradicted the latter. The current architecture does not include models of eye movement and binocular perception. However, in early stages there were plans to include a model of eye movement, which were discarded in favour of a streamlined, singular approach with respects to the scope of this thesis. After all, the main goals of the thesis were to test the effects of coactivation on reaching behaviour, and to test the viability of CI in cognitive models. The current sampling of spatial input from the target position was used as a simplified shorthand for such a model.

In essence, an integrated model of the eyes would work similarly to the arm model, but each eye would rotate along 2 DOFs. Each DOF would have an agonist-antagonist muscle pair to rotate the eye, which would be subjected to muscle coactivation and reciprocal innervation. Each eye would sample using the current method, and the combined sampled would be used to interpolate a 3-D location of the target. The observed positional information would be transformed to body-centered information, which would be usable in subsequent fixations and reaches. The resulting model would be able to use multiple subsequent eye movements to fixate the target, or look around randomly in search of certainty. In doing so, it would inform other motor systems (i.e. the arm, but the current architecture supports any multi-link limb system) about salient information.

Another point of discussion is the seemingly excessive simplification of internal representations of motor states. Representing muscle tension as a numeric value is most logically done using continuous values between some arbitrary lower and upper limits. The need to discretize internal representations came from the definition of causal worlds used in the imaging procedure. Representing all worlds in a system with continuous variables would require artificially discretizing these variables during construction of worlds, lest an infinite number of worlds are used in the imaging procedure. Keeping the causal model discretized is a simpler solution, both in belief updating and in consistency of values in the architecture. Naturally, when discretizing variables that should logically be continuous variables, one would prefer high granularity in these nodes. At first, we had aimed to use a causal network with consistent node granularities of 10 or higher. Even af-

ter implementing the additional complexity saving measures outlined in 2.2.1, it became quickly apparent that such high granularities were not viable.

These problems may not be exclusive to motor control, but can be generalized over cognitive functions that require fine-grained processing. One should note that the issue arising from the memory- and time-complexity of reasoning with too many different worlds is not exclusive to imaging. The notion of worlds was not introduced with imaging; it is a term that has been used in the discussion on CI since before imaging was coined. Any model of counterfactual inference based on worlds would cause the same problematic trade-off between model accuracy and complexity. Solutions to this trade-off would have to deal with the explosive complexity of dealing with continuous variables.

A solution one might think of is pre-selection of worlds based on some statistical measure such as informativeness, or the world's prior. However, this is a classical part of the frame problem [46, 47], where, in order to not consider a world, it implicitly *must* be considered, trapping the counterfactual agent in a infinite loop of rejecting worlds for consideration. Over the years, circumventions (and in some cases, claimed solutions) to the frame problem have been proposed [48, 49, 50, 51], most of which have different applications. Due to the specific nature of the current problem, we can not assume nor propose that any of these solutions or circumventions would provide improvement over the current issue surrounding CI without diving deeper into their workings. Most of these solutions are based in higher-order logic, and attempting to apply them to the current issue is far beyond the scope of this section. However, the idea stands that a solution in higher-order logic may provide an escape from the complexity issues surrounding CI. A point one might make in favour of one specific method of circumvention is that no reasonable agent that thinks ahead considers every possible event within the number of reasoning steps it makes. This is a valid argument- after all, any agent that reasons several steps forward will require strong boundaries on each set of subsequent steps. An example of this would be chess players, who should not take into account every possible move at all times when thinking ahead several turns, since not every move is viable in a competition. The bounding of inference steps within what can reasonably be expected is an example of the *sleeping dogs* strategy. Advocates of using this strategy within any AI application may find this a reasonable constraint, and it truly seems to be. However, when considering a potentially great number of worlds, attempting to apply the sleeping dogs strategy shows a problem: What worlds does the agent not even think of, and how does it determine those that are worth consideration? The problem here is that the sleeping dogs strategy is implicitly applied in the variables that are represented in the causal model, especially if subnetworks are separated in preparation for CI. Bounding the outcomes by what can reasonably be expected is the first step implicit in any workable implementation of CI, and attempting to apply any form of "*not waking sleeping dogs*" is a futile operation.

Critics may offer up the notion reasoning forward an arbitrary number of steps is bound to cause complexity issues in any scenario, be it in the real world or in simulations. As such, it is reasonable to argue that CI should not be considered too far ahead. The question that arises here is: How

far is *too far*? Since imaging and other world-based CI strategies are NP-hard, one could argue that every step is one too far, as every step introduces exponential growth in computational load. Regardless, one should take into consideration that the current study was performed with with CI over maximally four nodes with relatively low granularity at any given time. A possible response is that the non-binary nature of the causal model is what makes this specific use of the method so problematic, and that CI over binary causal models could work perfectly well. With sufficiently small inputs even NP-hard solutions can be executed quickly and— assuming such a small model accurately represent causal structures of the given problem— quite well. the ‘solutions’ to the complexity issues of CI presented in this study serve the same purpose of shrinking the input of the algorithm to a workable size. However, these pre-emptive steps merely circumvent high computational loads caused by the algorithm, rather than addressing its complexity.

One may think the current study may not have made optimal use of imaging’s strengths: Since imaging facilitates disjunctive counterfactual actions, it can be argued that its required computational load is wasted on processing single counterfactuals. However, it should not be forgotten that disjunctive actions lead to less exclusion of worlds. This means that each disjunctive action leaves more worlds to be considered during inference over following variables, increasing the run-time of the algorithm even further. Furthermore, while disjunctive action is indeed one of imaging’s strengths, Pearl forewarns its use for a reason. Reasoning over successive disjunctive actions assumes cognitive functions that we were not willing to make: Justifying the use of parallel sequential disjunctives in a model of cognition seemed like a fool’s errand. After all, how would the cortical regions responsible for controlling limbs process multiple disjunctive actions at once? This question can be applied to both reasoning with worlds and to disjunctive actions. Worlds (and the non-disjunctive reasoning over them) can generally be explained as follows: For each decision of a control parameter, a multitude of worlds disappears as though a set path has been chosen. This can be likened to lateral inhibition between layers of neurons in the cortex. In the motor pathways, neurons encoding the chosen activation for a physical action would be the most strongly activated, inhibiting others. As such, the notion of worlds is not necessarily unrealistic, if one were to imagine them as possible activations in pathways responsible for motor actions. For disjunctive action, this explanation would not hold. Several laterally inhibiting chains of neurons in a pathway cannot fire simultaneously, as they would inhibit each other, with one chain winning the ‘exchange’.

Various suggestions for further research can be considered given the questions and possible expansions this study provides. On the developmental and behavioural side, further experimentation could include extension of the experimental paradigm to include time-scale representations of stimuli and motor states will elucidate the effects the onset of muscle coactivation can have on reaching behaviour. Likewise, an inclusion of a model of eye movement and visual perception can be included, providing data regarding fixation behaviour and its interactions with coactivation onset. It should be clear that the underlying generative model should first be changed, and tested for possible changes in observed behaviour. Imaging in its current state does not seem to be the right

model for motor cognition.

Further development on imaging and other methods for CI should be considered with a focus on their ideal applications, as well as on improving their complexity. Although the basic form of imaging as presented by Pearl deals with complexity issues, further development of the method may introduce changes that allow for time-efficient imaging over possible world states. If such an improvement were to be developed, the bounds of what can be accomplished by counterfactual imaging will surely shift in favour of natural and realistic models of cognition.

6 References

- [1] C. von Hofsten, “Developmental changes in the organization of prereaching movements.,” *Developmental psychology*, vol. 20, no. 3, p. 378, 1984.
- [2] I. McGuire and G. Turkewitz, “Approach-withdrawal theory and the study of infant development,” *Cognitive growth and development*, pp. 57–84, 1979.
- [3] V. Gatev, “Role of inhibition in the development of motor co-ordination in early childhood,” *Developmental Medicine & Child Neurology*, vol. 14, no. 3, pp. 336–341, 1972.
- [4] K. J. Ciuffreda and L. Stark, “Descartes’ law of reciprocal innervation.,” *American journal of optometry and physiological optics*, vol. 52, no. 10, pp. 663–673, 1975.
- [5] L. A. Frey-Law and K. G. Avin, “Muscle coactivation: a generalized or localized motor control strategy?,” *Muscle & nerve*, vol. 48, no. 4, pp. 578–585, 2013.
- [6] G. Gottlieb, B. Myklebust, R. Penn, and G. Agarwal, “Reciprocal excitation of muscle antagonists by the primary afferent pathway,” *Experimental brain research*, vol. 46, no. 3, pp. 454–456, 1982.
- [7] C. Crone, “Reciprocal inhibition in man.,” *Danish medical bulletin*, vol. 40, no. 5, pp. 571–581, 1993.
- [8] M. Hoshiyama, Y. Kitamura, S. Koyama, S. Watanabe, M. Shimojo, and R. Kakigi, “Reciprocal change of motor-evoked potentials preceding voluntary movement in humans,” *Muscle & Nerve: Official Journal of the American Association of Electrodiagnostic Medicine*, vol. 19, no. 2, pp. 125–131, 1996.
- [9] P. L. Gribble, L. I. Mullin, N. Cothros, and A. Mattar, “Role of cocontraction in arm movement accuracy,” *Journal of neurophysiology*, vol. 89, no. 5, pp. 2396–2405, 2003.
- [10] P. Savastano and S. Nolfi, “A robotic model of reaching and grasping development,” *IEEE transactions on autonomous mental development*, vol. 5, no. 4, pp. 326–336, 2013.
- [11] N. E. Berthier, “Learning to reach: A mathematical model.,” *Developmental psychology*, vol. 32, no. 5, p. 811, 1996.
- [12] C. von Hofsten, “Structuring of early reaching movements: a longitudinal study,” *Journal of motor behavior*, vol. 23, no. 4, pp. 280–292, 1991.
- [13] P. Shaw, D. Lewkowicz, A. Giagkos, J. Law, S. Kumar, M. Lee, Q. Shen, and C. D. M. d’Autume, “Babybot challenge: Motor skills,” in *Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2015 Joint IEEE International Conference on*, pp. 47–54, IEEE, 2015.

- [14] J. Law, P. Shaw, and M. Lee, “A biologically constrained architecture for developmental learning of eye–head gaze control on a humanoid robot,” *Autonomous Robots*, vol. 35, no. 1, pp. 77–92, 2013.
- [15] S. K. Zibner, J. Tekülve, and G. Schöner, “The sequential organization of movement is critical to the development of reaching: A neural dynamics account,” in *Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2015 Joint IEEE International Conference on, pp. 39–46, IEEE, 2015.
- [16] J. Tekülve, S. K. Zibner, and G. Schöner, “A neural process model of learning to sequentially organize and activate pre-reaches,” in *Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2016 Joint IEEE International Conference on, pp. 318–325, IEEE, 2016.
- [17] W. Eijlander, L. Wijnen, A. van Maris, and B. J. Grzyb, “Muscle co-activation influences the development of prereaching movements,”
- [18] J. Pearl, *Causality*. Cambridge university press, 2009.
- [19] J. M. Kilner, K. J. Friston, and C. D. Frith, “Predictive coding: an account of the mirror neuron system,” *Cognitive processing*, vol. 8, no. 3, pp. 159–166, 2007.
- [20] A. Clark, “Whatever next? predictive brains, situated agents, and the future of cognitive science,” *Behavioral and brain sciences*, vol. 36, no. 3, pp. 181–204, 2013.
- [21] J. Kwisthout, H. Bekkering, and I. Van Rooij, “To be precise, the details don’t matter: on predictive processing, precision, and level of detail of predictions,” *Brain and cognition*, vol. 112, pp. 84–91, 2017.
- [22] R. P. Rao and D. H. Ballard, “Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects,” *Nature neuroscience*, vol. 2, no. 1, p. 79, 1999.
- [23] D. H. Hubel and T. N. Wiesel, “Receptive field and functional architecture in two non-striate visual areas (18 and 19) of the cat,” *J. neurophysiol*, vol. 28, no. 2, p. 229, 1965.
- [24] D. H. Hubel and T. N. Wiesel, “Receptive fields and functional architecture of monkey striate cortex,” *The Journal of physiology*, vol. 195, no. 1, pp. 215–243, 1968.
- [25] K. Friston, “The free-energy principle: a rough guide to the brain?,” *Trends in cognitive sciences*, vol. 13, no. 7, pp. 293–301, 2009.
- [26] R. A. Adams, S. Shipp, and K. J. Friston, “Predictions not commands: active inference in the motor system,” *Brain Structure and Function*, vol. 218, no. 3, pp. 611–643, 2013.
- [27] D. Lewis, “Counterfactuals and comparative possibility,” in *Ifs*, pp. 57–85, Springer, 1973.

- [28] M. Goldszmidt and J. Pearl, “Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions.,” *KR*, vol. 92, pp. 661–672, 1992.
- [29] J. Pearl, “Physical and metaphysical counterfactuals: Evaluating disjunctive actions,” *Journal of Causal Inference*, vol. 5, no. 2, 2010.
- [30] J. G. March, “Exploration and exploitation in organizational learning,” *Organization science*, vol. 2, no. 1, pp. 71–87, 1991.
- [31] J. Rosen, J. C. Perry, N. Manning, S. Burns, and B. Hannaford, “The human arm kinematics and dynamics during daily activities-toward a 7 dof upper limb powered exoskeleton,” in *Advanced Robotics, 2005. ICAR’05. Proceedings., 12th International Conference on*, pp. 532–539, IEEE, 2005.
- [32] R. P. Paul, *Robot manipulators: mathematics, programming, and control: the computer control of robot manipulators*. Richard Paul, 1981.
- [33] C. Jenkins and Z. Zeng, “Robot kinematics and dynamics.” <https://autorob.github.io/>. Course provided at University of Michigan, certain materials openly available online. The material in question consisted of slides from the topic *Forward Kinematics*.
- [34] G. E. MAYER, “Differential kinematic control equations for simple manipulators,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 1, no. 1, 1981.
- [35] D. Geiger, T. Verma, and J. Pearl, “d-separation: From theorems to algorithms,” in *Machine Intelligence and Pattern Recognition*, vol. 10, pp. 139–148, Elsevier, 1990.
- [36] S. Sinha, “A pedagogical walkthrough of computational modeling and simulation of wnt signaling pathway using static causal models in matlab,” *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2017, 12 2016.
- [37] S. R. Buss, “Introduction to inverse kinematics with jacobian transpose, pseudoinverse and damped least squares methods,” *IEEE Journal of Robotics and Automation*, vol. 17, no. 1-19, p. 16, 2004.
- [38] D. E. Orin and W. W. Schrader, “Efficient computation of the jacobian for robot manipulators,” *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 66–75, 1984.
- [39] Y. Nakamura and H. Hanafusa, “Inverse kinematic solutions with singularity robustness for robot manipulator control,” *Journal of dynamic systems, measurement, and control*, vol. 108, no. 3, pp. 163–171, 1986.
- [40] C. W. Wampler, “Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 16, no. 1, pp. 93–101, 1986.

- [41] J. P. Spencer and E. Thelen, “Spatially specific changes in infants’ muscle coactivity as they learn to reach,” *Infancy*, vol. 1, no. 3, pp. 275–302, 2000.
- [42] A. N. Meltzoff and M. K. Moore, “Explaining facial imitation: A theoretical model,” *Infant and child development*, vol. 6, no. 3-4, pp. 179–192, 1997.
- [43] B. Libet, C. A. Gleason, E. W. Wright, and D. K. Pearl, “Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential) the unconscious initiation of a freely voluntary act,” *Brain*, vol. 106, no. 3, pp. 623–642, 1983.
- [44] M. Otworowska, J. Riemens, C. Kamphuisa, P. Wolferta, L. Vuurpijla, and J. Kwisthout, “The robo-havioral methodology: Developing neuroscience theories with foes,” in *Proceedings of the 27th Benelux Conference on AI (BNAIC’15)*, 2015.
- [45] R. Carpenter, “Homeostasis: a plea for a unified approach,” *Advances in Physiology Education*, vol. 28, no. 4, pp. 180–187, 2004.
- [46] J. McCarthy, P. Hayes, B. Meltzer, and D. Michie, “Machine intelligence,” 1969.
- [47] D. Dennett, “The frame problem of ai,” *Philosophy of Psychology: Contemporary Readings*, vol. 433, pp. 67–83, 2006.
- [48] E. Sandewall, “An approach to the frame problem and its implementation,” *Machine intelligence*, vol. 7, no. 195-204, pp. 11–19, 1972.
- [49] M. Thielscher, “Introduction to the fluent calculus,” 1998.
- [50] M. Gelfond and V. Lifschitz, “Action languages,” 1998.
- [51] J. A. Fodor, “Modules, frames, fridgeons,” 1987.