Panta Rhei: Curiosity-Driven Exploration to Learn the Image-Schematic Affordances of Pouring Liquids

Mihai Pomarlan[1][0000−0002−1304−581X], Maria M. Hedblom[2][0000−0001−8308−8906], and Robert Porzel[3][0000−0002−7686−2921]

[1] Applied Linguistics Department, University of Bremen, Germany
[pomarlan]@uni-bremen.de
[2] Jönköping Artificial Intelligence Laboratory, Jönköping University, Sweden
[maria.hedblom]@ju.se
[3] Digital Media Lab, University of Bremen, Germany
[porzel]@tzi.de

Abstract. Despite rapid progress, cognitive robots have yet to match the facility with which humans acquire and find ways to reuse manipulation skills. An important component of human cognition seems to be our curiosity-driven exploration of our environments, which results in generalizable theories for action outcome prediction via analogical reasoning. In this paper, we implement a method to emulate this curiosity drive in simulations of the situation of pouring liquids between containers, and to use these simulations to construct a symbolic theory of pouring. The theory links qualitative descriptions of an initial state and manner of pouring with observed behaviors, and can be used to predict qualitative outcomes or select manners of pouring towards achieving a goal.

Keywords: Curiosity-driven Learning · Hybrid Approaches · Cognitive Robotics · Intrinsic Motivation · Commonsense Reasoning

1 Introduction

Lifelong open-ended learning is considered one of the hallmarks of intelligent behavior. Therefore, any formal system modelling intelligent behavior requires an autonomous component for unprompted learning. Looking at human behavior for inspiration, it is apparent that children engage in play and other learning activities with no other purpose than "for their own sake" [4]. Likewise, adults spend considerable time practicing hobbies that have little to do with evolutionary benefits, such as survival and reproduction. Instead, learning is intrinsically motivated, where curiosity drives cognitive processes to match external stimuli with internal representations [1][4]. Experiencing novel situations at all times, children quickly form and consistently update world assumptions regarding object

[4] In this particular research area, *intrinsic motivation* and *curiosity* are commonly treated as synonymous concepts.

properties, spatiotemporal relationships between objects, agent and environment and, in so doing, learn to anticipate outcomes of particular constellations [2, 23]. Through curious play, a developing agent learns to appreciate regularities of the world, e.g. when and how objects fall to the ground, that flat, solid objects can be stacked, and that contained objects will move with the container. As the agent's knowledge grows, increasingly sophisticated models of the world are formed, yet the conceptual skeleton remains as generalizable patterns that can be used for predicting outcomes of analogous situations.

Machine learning (ML) provides well-established techniques to teach artificial agents how to perform various tasks, e.g. image prediction or action selection [6, 30]. However, learning problems are often operationalized as feeding an existing dataset, perhaps augmented with random modifications, into a learning algorithm; autonomous learning is (most commonly) not included in such approaches, which negatively impacts generalizability of the learned models. To address these limitations, curiosity-driven learning methods are becoming increasingly common (e.g. [8, 11, 32, 33]). Usually starting with little or no symbolic knowledge about the world, these methods demonstrate how progressively more complicated predictive models can be built by agents via increasingly sophisticated play. Instead of tasking an artificial agent to solve a specific problem, the agent is encouraged to learn the rules of the environment and the affordances of objects without any particular goal state in mind. The idea is to simulate curiosity-induced autonomy and allow the agent to learn to predict outcomes and reason about scenes and activities it may be tasked with.

In this paper, we cut out a tiny slice of the full range of household situations robots may encounter and through curiosity-driven exploration, we build a theory about the interactions of liquids and solids, particularly the relationships involved in pouring a liquid from one container to another. We take inspiration from previous work by Oudeyer and Kaplan [27], arguing that curiosity and learnability should guide exploration in a play environment, in our case a simulation. However, we apply a symbolic flair to our models to anticipate the future, and have the robot learn symbolic rules in a logic of image schemas to describe how a situation might develop. These symbolic rules could then be used by an agent to construct perception queries and select actions [15] and we intend to pursue this line of research in future work.

2 Related Work

Cognitive robotics offers an interesting platform to investigate the benefits of curiosity-driven learning of everyday activities. On the one hand, robots are embodied, artificial agents that have the possibility to perceive and physically interact with their environment in a way similar to how humans would. On the other hand, they are subject to their programming and limited by the capacities of their hardware. Previous research demonstrates that while cognitive robots are proficient in well-defined tasks, they are usually unable to handle underspecified instructions and unfamiliar situations [3, 19, 25]. This means that their abilities for commonsense reasoning, prediction and adaptation require substantial improvements before they reach any level of human-like intelligence.

Our proposal is to combine the strength of curiosity-driven learning with symbolic description of functional relationships in the form of image-schematic micro-theories. Image schemas are spatiotemporal relationships between objects, agents and environments that have been argued to capture the meaning of affordances and the human conceptualisation of particular events [13, 17]. Following the reasoning behind embodied cognition[5], it is our hypothesis that this is a good foundation to formally model cognitive aspects of learning. This means that together with conceptual information about the environment, we add an ability for curiosity to, on its own terms, encourage the robots to engage with and extract meaningful information from its environment.

Before presenting the state of the art, we take a closer look at the cognitive processes of extracting generalizable information from unfamiliar scenarios.

2.1 Generalizable Patterns of Learning

Motivated by curiosity (and instinct for survival), humans learn novel things as they are exposed to unfamiliar scenarios. However, unlike artificial agents, they do not come empty-handed to most unfamiliar situations. Human cognition has mastered the skill of using generalized snippets of information learned from one situation, and in relation to other objects and concepts, and analogically apply it to different, unknown situations [16]. In this fashion, events are generalized into smaller temporal segments of motion events [34] and objects are thought to be conceptualized in relation to the affordances[6] they offer in the environment. Combined, these concepts formulate the foundation for image schemas [17, 21], described as *"dynamic analog representations of spatial relations and movements in space"* [22, p.591]. Formally they can be thought of as functional relationships over time and space that constitute the conceptual components that structure the understanding of the environment. For instance, a *cup* affords the image schema of Containment[7], and a *car* affords the motion-event of being transported from one point to another as the image-schematic pattern Source_Path_Goal. These conceptual patterns can then be combined into profiles [14, 26] to form increasingly complex understandings of the relationships between objects involved in a particular situation. For instance, if the knowledge that a plate can be Supported by a table is present, it can be analogously inferred that a book can be Supported by a desk. Originally stemming from cognitive linguistics to motivate the high amount of spatial metaphors and used to describe the shared structure that information is transferred on in analogy, they have been proposed to play the part of the most general components involved in commonsense reasoning of everyday activities for formal systems [13, 18]. Previous work

[5] The embodied cognition hypothesis argue that the body's interaction with its environment is the source for concept formation.

[6] By affordances we refer to the theory introduced by Gibson [12].

[7] Following convention, all image schemas are written in capitalized uppercase letters.

suggests it is motion-events [5] and changes in image-schematic states [14, 34] that constitute the principles for event segmentation in human cognition.

Image schemas have also been proposed as a way to link symbolic inference and numeric methods, and in so doing, impart a deeper understanding of functional relations in the physical world to an artificial agent. In [29], and the resulting hybrid reasoning is employed to discover which objects in a scene contribute towards maintaining a functional relation, beyond what is explicitly asserted at the symbolic level.

2.2 Intrinsic Motivation for Learning

Curiosity-driven approaches to learning are a specialized form of reinforcement learning (RL) algorithms that aim to maximize reward functions. The difference to traditional RL is that instead of rewards being offered from particular states of the environment, they are calculated in relation to the novelty of the learned information [32]. One account of intrinsic motivation is the work [31], in which prediction accuracy is enhanced based on the principle of maximizing intrinsic rewards. The theory encourages any method to explore event sequences unknown to the agent, but still regular enough to allow for learnable algorithmic structure. In [33], environment exploration is focused in areas where a classifier is uncertain. Information gain maximization is used in [11] to teach a robot motion planning.

In [28], the authors demonstrate how solving novel tasks in novel situations is significantly faster when using their curiosity-driven algorithm rather than classic reinforcement learning. Using this method on a large-scale, in [8] the authors show how the system learns to play several Atari games.

A different approach is taken by [27]. In their work they show how an AIBO robot learns to interact in increasingly complex ways with its environment by focusing its attention towards what it can learn best, as opposed to which interaction is expected to produce the most surprise. This results in behavior in which interactions that are too complex for the agent to comprehend are rarely attempted. Likewise, what is already well understood and predictable is also avoided. The result is an autonomously constructed curriculum for the robot, that strikes a balance between challenges and learning.

One particular area in which curiosity-driven learning could be applied, is learning object affordances. Through exploration, agents are able to identify the potential usages of particular objects and their environments. In similarity to how children learn image-schematic affordances by embodied interactions, so could artificial agents. The work on learning image-schematic notions and affordances is a large research area within cognitive robotics characterized by many different methodologies (e.g. [7, 9, 24]). For instance, in [20] the authors demonstrate how a robotic system can, through analogical reasoning, learn the object affordances of pouring coffee into different kinds of containers without having experience of that particular container.

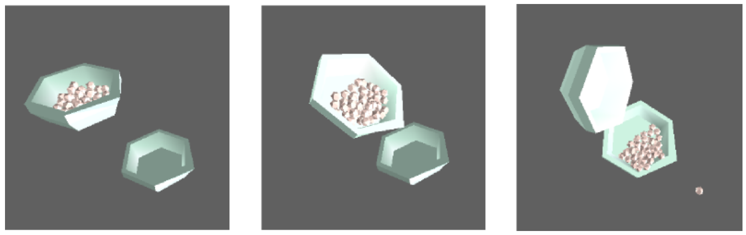

Fig. 1: A few frames from a pouring simulation.

3 Curiosity-Driven Exploration in a Symbolic Context

Previous work on curiosity-driven learning has focused on reinforcement learning methods that produce numeric models. Often, such models are difficult to interpret and difficult to reuse outside of the context in which they were learned. Inspired by the conceptual building blocks in image schemas, we will instead describe a method to construct symbolic rules based on a set of observations of simulated scenes. This set of simulations is constructed and expanded during the learning process, and each expansion aims to explore scene parameterizations expected to deliver not yet seen behaviors.

To showcase our method, we have selected the pouring liquids task. On the one hand, indirectly manipulating liquids via pouring is an important component of many everyday activities. On the other hand, the interactions of liquids with their surrounding objects are sensitive to many variables – pouring from too high may increase spillage, pouring too much to the side might tip an unstable container and so on. For tractability reasons, we approximate liquids in our work here by a collection of small rigid particles (Fig. 1). While not entirely realistic, this simplified model nonetheless allows interesting behaviors in simulation.

3.1 From Symbols to Scenes and Back – with a Curious Twist

Our goal is for an artificial agent to develop a "naïve physics" model of the world, at least in matters that pertain to transferring fluids between containers via pouring. This model will be a collection of rules stipulating that if some antecedent holds – a conjunction of propositions qualitatively describing a situation – then the observed behaviors of objects will obey a qualitative description – the consequent. The rules in the theory are defeasible, to capture the fact that such naïve rules are subject to exceptions that further experience may uncover.

Human playful exploration discovers what features of the world are interesting, or important, to look for. In this paper, we restrict our ambition by telling the agent what aspects of behavior will be interesting to monitor, and we are satisfied if the agent is merely able to adjust the parameters of its simulations such that it discovers behaviors not seen before.

We build on the approach from [29] where image schemas are represented at different levels of abstraction via theories for functional or spatial relations at the symbolic level, and methods to parameterize relations between geometric primitives at the numeric level. As such, the propositions we use to describe a scene relate to image-schematic aspects such as verticality, alignment, and relative movement (similarly to the method used in [10]). Further, these propositions are associated with generative models that allow the agent to either instantiate a spatial arrangement that obeys a qualitative description, or judge how well an arrangement or trajectory conforms to such. Through the generative models, antecedents – conjunctions of propositions – are associated to regions of the simulated scene parameter space describing an initial state and some action parameters. Conversely, the resulting trajectories from the simulation are used to decide which propositions qualitatively describe the observed behavior.

To discover the rules of a naïve physics model for pouring, one could use a dense grid of possible parameterizations of a simulation, and for each of these use the generative models to see which qualitative descriptions best fit the initial state and subsequent behavior. Such an approach is not feasible when there are many parameters – the curse of dimensionality. It follows that some heuristic to identify interesting parameterizations is necessary. A parameterization is interesting if it may result in behavior that has not been seen before, i.e. its qualitative description is different from that of previous simulations.

Our approach to tackle this problem consists of the following: for an antecedent, i.e. a qualitative description of an initial state, generate a set of parameterizations and for each of these, see how well the observed simulated behaviors obey each of a set of possible qualitative expectations. In effect, the simulator or physics engine is treated as a function mapping parameterizations to expectation fulfillment scores. This function is approximated by a set of hyperplanes, one for each expectation score, and each hyperplane gives a direction in which that expectation score is most sensitive to variation, at least in the region of parameter space covered by the considered antecedent. The antecedent is modified such that the parameterizations likely to be generated by its modified version are in the most interesting directions of variation, and the process is repeated.

It is useful to update the 'interestingness' of exploration directions in the region of parameter space covered by an antecedent as more simulation runs become available, so a nuance to the approach outlined above is to also compute and update weights that, intuitively, represent how strong the association is between a particular antecedent and parameterization. This allows the interesting directions of exploration around an antecedent to change as what was previously interesting gets explored and associated to other antecedents.

Thus, the "playful exploration" in simulation can be formalized. Let there be a list E of m propositions describing possible qualitative behaviors observed in simulation, a set D of propositions qualitatively describing either aspects of the initial state or of the pouring action, and a set of antecedents $A \subset 2^D$. A simulation is parameterized by some point p in $I = \mathbb{R}^n$ where each component gives the value of some parameter. As output, the simulation produces a point in

$O = ([0,1])^m$, where the kth component is a "score" describing how well the kth qualitative description from E applies to the simulated behavior. The physics engine running the simulation works according to some physical model but, for our purpose of learning a set of rules, we can treat it as a function $simulate : I \rightarrow O$ and we record samples of this function as they become available during the agent's exploration.

Because the propositions in an antecedent a have generative models attached to them to select parameter values, for a given simulation parameterization p there is a probability $P(p|a)$, estimating the likelihood of the parameterization if we want to instantiate a scene for description a. As our symbolic model grows and includes more rules, it may be the case that a parameterization is somewhat likely for several antecedents; however, it will be necessary to keep track of which simulations are particularly likely for only a few antecedents. To do this, we define a weight function w. Given a set of antecedents explored so far H, and an antecedent $a \in H$, the weight of a parameterization p for antecedent a is:

$$w(a,p) = \frac{P(p|a)}{\sum_{b \in H} P(p|b)} \tag{1}$$

$w(a,p)$ answers the following question: assuming a-priori that all antecedents in H are equally likely, and we know the antecedent we sampled for a parameterization produced p, how likely is it that the antecedent we sampled was a? As we will explain next, it is useful to allow the system to "forget" which antecedent actually generated a parameterization.

Our agent "explores" an antecedent a by generating several parameterizations, simulating them, and observing the results. However, the next step will be in deciding how to change the antecedent to (hopefully) get new outcomes. Therefore, it may be interesting to know which are the directions in which the *simulate* function is changing faster at the "point" corresponding to the antecedent. To estimate this, *simulate* is approximated by a hyperplane using weighted linear regression, where the weight of a parameterization p (whether it was actually sampled using a or not) is $w(a,p)$. This is because, as the number of recorded simulations and known antecedents grows, a previously recorded simulation may have a parameterization that fits several antecedents, and as a result is less informative about new directions to explore around any antecedent. "Learning" would then proceed according to the steps described in Algorithm 1. The algorithm uses the following functions:

- *sample* is a function which, given an antecedent and an integer n, produces n parameterizations sampled from $P(p|a)$.
- *simulate* is a function which, given a parameterization p, returns the resulting scores for the monitored expectations.
- *weight* is a function which maps an antecedent a, a parameterization p, and a set of antecedents H to which a belongs, to $w(a,p)$.
- *planeFit* performs weighted linear regression on a collection of parameterization/result pairs S with a vector of weights W.

Algorithm 1 Symbolic prediction model building via curious exploration

$Simulations, Rules, Antecedents \leftarrow \emptyset$
$A_{new} \leftarrow \{a_0\}$
while $A_{new} \neq \emptyset$ **do**
  $Planes, Weights \leftarrow \emptyset$
  $a_c \leftarrow A_{new}.pop()$
  $Antecedents \leftarrow Antecedents \cup \{a_c\}$
  for all $p \in sample(a_c, n+1)$ **do**
    $Simulations \leftarrow Simulations \cup \{(p, simulate(p))\}$
  end for
  for all $a \in Antecedents$ **do**
    for all $p \in Simulations$ **do**
      $Weights \leftarrow Weights \cup (a, p, weight(a, p, A))$
    end for
    $Planes \leftarrow Planes \cup \{(a, planefit(Weights(a), Simulations))\}$
  end for
  $dir, absslope, a_m \leftarrow max_{(a,h) \in Planes} bestabsslope(h)$
  if $absslope > threshold$ **then**
    $A_{new} \leftarrow A_{new} \cup modify(a_m, dir, Antecedents)$
  end if
end while
for all $a \in Antecedents$ **do**
  $c \leftarrow qualify(Simulations(argmax_{p \in Simulations} Weights(a, p)))$
  $Rules \leftarrow Rules \cup \{(a, c)\}$
end for

- $bestAbsSlope$ is a function which, given a hyperplane h, returns the direction of largest slope in this plane, and the absolute magnitude of this slope.
- $modify$ is a function which, given an antecedent a, a direction of variation dir, and a set of antecedents H, returns a set of antecedents obtained from modifying a such that these new antecedents are more likely to generate parameterizations that differ in direction dir compared to those sampled for a, and such that these new antecedents do not already appear in H.
- Finally, $qualify$ is a function which maps the result of a simulation s to the consequent that best describes it.

3.2 Exploring the Naive Physics of Pouring

We simulated scenarios in which a collection of small, rigid-body particles represent a liquid, and are poured from one bowl into another. The pouring motion is described via three poses of the source bowl and the duration of motions between those poses. Physical parameters for the source and destination bowls, and the particles representing the liquid, are also considered: mass, restitution, friction. The amount of particles used to represent the liquid in a simulation can vary. In total, 30 parameter values are needed for one simulation. To adjust these numeric parameters, several statements can be combined in antecedents. There are

86 propositions to select from and combine to create a qualitative description of a pouring scene. Not all of these combinations make sense however and the number of possible qualitative scenes is on the order of one billion.

The qualitative behaviors we monitor to construct our consequents relate to the motion of the liquid and the destination during the different stages of the source movement. We expect the destination to stay still throughout the pouring, and the liquid to leave the source, enter and eventually remain in the destination. In total, 9 propositions are used to check the simulation behavior.

We then ran the algorithm described in 1. Despite there being about one billion possible antecedents, the algorithm shows a clear tendency towards slowing down the growth of the A_new list: while the first antecedent considered produced 17 new ones to explore, the last explored antecedents added 1 or 0 new antecedents to explore. Nonetheless, we stopped exploration after 25 rules were produced, which include rules such as:

- pouring from the side will result in spillage
- too light a liquid will not entirely leave the source
- when pouring a lightweight liquid and then righting up the source while low above the destination will result in the liquid not remaining in the destination

4 Conclusion and future work

In this paper, we have described a method to use simulations to construct a symbolic model to predict qualitative outcomes of an action. The method focuses exploration on areas of the parameter space that are likely to yield new behaviors. The resulting model can then be used to predict action outcomes, or select action parameterizations. We showcase our method through a pouring liquids scenario.

As future work we intend to deploy such qualitative models for action selection for a simulated agent. Another direction is to adjust the generative models used to convert between qualitative and numeric descriptions, in an attempt to emulate the context- and task-dependent nature of human qualitative descriptions (e.g. "the weather is hot" vs. "the oven is hot" have different meanings in terms of expected temperature values).

Acknowledgements

The research reported in this paper has been supported by the FET-Open Project 951846 "MUHAI – Meaning and Understanding for Human-centric AI" (http://www.muhai.org/) funded by the EU Program Horizon 2020 as well as the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 "EASE - Everyday Activity Science and Engineering", University of Bremen (http://www.ease-crc.org/). The research was conducted in sub-projects *P01 Embodied semantics for the language of action and change: Combining analysis, reasoning and simulation.*

Bibliography

[1] Gianluca Baldassarre, Tom Stafford, Marco Mirolli, Peter Redgrave, Richard M. Ryan, and Andrew Barto. Intrinsic motivations and open-ended development in animals, humans, and robots: an overview. *Frontiers in Psychology*, 5:985, 2014.

[2] Lawrence W. Barsalou et al. Perceptual symbol systems. *Behavioral and brain sciences*, 22(4):577–660, 1999.

[3] Michael Beetz, Daniel Bessler, Andrei Haidu, Mihai Pomarlan, Asil Kaan Bozcuoglu, and Georg Bartels. Know Rob 2.0 - A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 512–519, 2018.

[4] Daniel E. Berlyne. Curiosity and exploration. *Science*, 153(3731):25–33, 1966.

[5] Jürgen Bohnemeyer, Nicholas J. Enfield, James Essegbey, Iraide Ibarretxe-Antuñano, Sotaro Kita, Friederike Lüpke, and Felix K. Ameka. Principles of event segmentation in language: The case of motion events. *Language*, pages 495–532, 2007.

[6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. 2016. Online: arXiv:1604.07316 [cs.CV].

[7] Jake Brawer, Meiying Qin, and Brian Scassellati. A causal approach to tool affordance learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8394–8399. IEEE, 2020.

[8] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. 2018. Online: arXiv:1808.04355 [cs.LG].

[9] Jefferson Coelho, Justus Piater, and Roderic Grupen. Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot. *Robotics and Autonomous Systems*, 37(2-3):195–218, 2001.

[10] Kaviya Dhanabalachandran, Vanessa Hassouna, Maria M. Hedblom, Michaela Küempel, Nils Leusmann, and Michael Beetz. Cutting events: Towards autonomous plan adaption by robotic agents through image-schematic event segmentation. In *Proceedings of the 11th on Knowledge Capture Conference*, pages 25–32, 2021.

[11] Mikhail Frank, Jürgen Leitner, Marijn Stollenga, Alexander Förster, and Jürgen Schmidhuber. Curiosity driven reinforcement learning for motion planning on humanoids. *Frontiers in neurorobotics*, 7:25, 2014.

[12] James J. Gibson. The theory of affordances. In Robert Shaw and John Bransford, editors, *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, pages 67–82. NJ: Lawrence Erlbaum, Hillsdale, 1977.

[13] Maria M. Hedblom. *Image Schemas and Concept Invention: Cognitive, Logical, and Linguistic Investigations.* Cognitive Technologies. Springer Computer Science, 2020.

[14] Maria M. Hedblom, Oliver Kutz, Reafael Peñaloza, and Giancarlo Guizzardi. Image schema combinations and complex events. *KI-Künstliche Intelligenz*, 33(3):279–291, 2019.

[15] Maria M. Hedblom, Mihai Pomarlan, Robert Porzel, Rainer Malaka, and Michael Beetz. Dynamic action selection using image schema-based reasoning for robots. In *Proceedings of the 7th Joint Ontology Workshops*, 2021.

[16] Douglas R. Hofstadter. Analogy as the core of cognition. *The analogical mind: Perspectives from cognitive science*, pages 499–538, 2001.

[17] Mark Johnson. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason.* The University of Chicago Press, Chicago and London, 1987.

[18] Werner Kuhn and Andrew U. Frank. A formalization of metaphors and image-schemas in user interfaces. In *Cognitive and linguistic aspects of geographic space*, pages 419–434. Springer, 1991.

[19] Lars Kunze, Moritz Tenorth, and Michael Beetz. Putting people's common sense into knowledge bases of household robots. In *Annual Conference on Artificial Intelligence*, pages 151–159. Springer, 2010.

[20] Safoura Rezapour Lakani, Antonio J. Rodríguez-Sánchez, and Justus Piater. Exercising affordances of objects: A part-based approach. *IEEE Robotics and Automation Letters*, 3(4):3465–3472, 2018.

[21] George Lakoff. *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind.* The University of Chicago Press, 1987.

[22] Jean M. Mandler. How to build a baby: II. Conceptual primitives. *Psychological review*, 99(4):587–604, 1992.

[23] Jean M. Mandler. *The Foundations of Mind: Origins of Conceptual Thought: Origins of Conceptual Though.* Oxford University Press, New York, 2004.

[24] Luis Montesano, Manuel Lopes, Alexandre Bernardino, and José Santos-Victor. Learning object affordances: from sensory–motor coordination to imitation. *IEEE Transactions on Robotics*, 24(1):15–26, 2008.

[25] Nils J Nilsson. Shakey the robot. Technical report, SRI INTERNATIONAL MENLO PARK CA, 1984.

[26] Todd Oakley. Image Schema. In Dirk Geeraerts and Hubert Cuyckens, editors, *The Oxford Handbook of Cognitive Linguistics*, pages 214–235. Oxford University Press, 2010.

[27] Pierre-Yves Oudeyer and Frédéric Kaplan. Discovering communication. *Connection Science*, 18(2):189–206, 2006.

[28] Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, pages 2778–2787. PMLR, 2017.

[29] Mihai Pomarlan and John A. Bateman. Embodied functional relations: A formal account combining abstract logical theory with grounding in simulation. In *Formal Ontology in Information Systems: Proceedings of the 11th*

International Conference (FOIS 2020), volume 330, page 155. IOS Press, 2020.

[30] Eder Santana and George Hotz. Learning a driving simulator. 2016. Online: arXiv:1608.01230v1 [cs.LG].

[31] Jürgen Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.

[32] Susanne Still and Doina Precup. An information-theoretic approach to curiosity-driven reinforcement learning. *Theory in Biosciences*, 131(3):139–148, 2012.

[33] Emre Ugur, Mehmet R. Dogar, Maya Cakmak, and Erol Sahin. Curiosity-driven learning of traversability affordance on a mobile robot. In *2007 IEEE 6th international conference on development and learning*, pages 13–18. IEEE, 2007.

[34] Jeffrey M. Zacks and Khena M. Swallow. Event segmentation. *Current directions in psychological science*, 16(2):80–84, 2007.