From imitation to emulation: A developmental framework for continual robot learning

Chaolan Lin, University of California, San Diego
chlin@ucsd.edu

***Abstract*—Imitation and emulation are foundational mechanisms through which humans, particularly children, acquire knowledge and adapt to complex social environments. Developmental research shows that children initially engage in high-fidelity imitation, including the reproduction of causally irrelevant actions, but gradually shift toward emulation, selectively reproducing goal-relevant aspects of behavior. This trajectory, from surface-level copying to intentional, goal-directed understanding, offers a powerful model for continual robot learning. In this position paper, we propose that robots can benefit from a similarly staged learning process, beginning with broad imitation and advancing toward flexible emulation guided by sustained human interaction and feedback. Drawing on theoretical and empirical insights from developmental psychology, we define key learning stages and introduce benchmark tasks that assess fidelity, causal reasoning, generalization, and social alignment. By aligning robot learning with human cognitive trajectories, we aim to develop systems that are not only robust and adaptive but also interpretable and capable of evolving within human environments.**

Robot learning; developmental robotics; imitation; emulation

I. Introduction

As robots become increasingly integrated into human environments, the demand for systems that can continually adapt to new social contexts, evolving human needs, and dynamic tasks becomes more pressing. Unlike static training paradigms, continual learning enables robots to update their knowledge and behavior incrementally over time. Human development, particularly in early childhood, offers a compelling model for this type of learning. One of the key trajectories in developmental psychology is the progression from imitation to emulation: From replicating observed actions to inferring and achieving goals through flexible strategies. In developmental psychology, this progression reflects not only growing cognitive sophistication but also increasing social attunement, including the ability to interpret communicative intent and infer causal structures [1-2].

In the context of human-robot interaction, continual learning should be similarly framed not only as a technical challenge but also as a process of social, situated, and cognitively interpretable adaptation. We propose a developmentally inspired framework in which robot learning mirrors children's progression from mimicry to norm-sensitive reasoning. By grounding robotic learning in this scaffolded trajectory, we argue that robots can better align with human expectations, adapt in socially appropriate ways, and generalize knowledge across tasks and users.

II. Imitation and Emulation in Human Development

Imitation refers to the replication of observed behavior, often performed without a full understanding of its purpose. From the earliest stages of life, infants demonstrate a capacity for imitation: Newborns can reproduce facial expressions, such as tongue protrusion, within the first few weeks of life [3], and by toddlerhood, children can replicate complex sequences of object-directed actions. Crucially, this imitation is often overly inclusive. Children often engage in overimitation, copying not only task-relevant actions but also causally unnecessary ones [2]

Overimitation is not a developmental error but a signal of social attunement. Research shows that 14-month-old infants are more likely to imitate an unconventional action, such as turning on a light with one's forehead, if it is presented in an ostensive-communicative manner [1]. This suggests that young learners are sensitive to cues indicating teaching intent and assume that demonstrated behaviors carry social or instructional significance, even when they appear inefficient.

As cognitive and social capabilities mature, children begin to transition from high-fidelity imitation to emulation, replicating the goal of an action rather than the specific means by which it is achieved. Emulation allows for flexibility, as children now use alternative methods to achieve the same outcomes, guided by emerging skills in causal reasoning [4], theory of mind [5], and social learning [6]. These developmental shifts support generalization across contexts, allowing children to abstract underlying principles from observed behavior—an essential capability for building robots that must operate in dynamic, human-centered environments.

III. Implications for Robot Learning

Despite advances in robot learning from demonstration and reinforcement learning, most systems remain limited in their ability to adapt flexibly and meaningfully over time, particularly in socially rich or uncertain environments. Current models are often trained in narrowly defined settings with fixed goals, resulting in behavior that is brittle, non-interpretable, and difficult to personalize. These systems typically lack mechanisms for building causal models, integrating long-term feedback, or interpreting ambiguous social signals, which are key capacities that underlie robust human learning.

In contrast, developmental research highlights how children's learning unfolds in a socially scaffolded, iterative fashion. Early overimitation facilitates the broad preservation of observed behaviors, whereas later stages involve causal inference, abstraction, and goal-based reasoning. This progression suggests that imitation should not be dismissed as a naïve strategy but embraced as an exploratory heuristic—a foundation from which more refined and abstract reasoning can emerge. We propose that robots adopt a similar trajectory, beginning with broad imitation and progressing toward emulation and norm inference, guided by interaction and feedback.

This developmental framing offers several implications for robot learning architectures. It encourages the design of systems that initially preserve high-fidelity demonstrations, even at the cost of redundancy, and then gradually prune behavior through causal learning. It also prioritizes human-aligned interpretability, where learning outcomes reflect not only task efficiency but also responsiveness to social norms, communicative cues, and user intent. Ultimately, this approach enables the development of robots that are not merely adaptable but socially embedded and capable of continual learning across time and context.

IV. Developmental stages for robot learning

To model human-like continual learning in robots, we outline a set of staged learning capabilities inspired by human development. These stages mirror how young children progressively acquire, refine, and generalize their knowledge through observation, action, and social interaction. Each stage reflects a qualitative shift in the learner's underlying cognitive representations and adaptive capabilities.

A. Stage 1: Mimicry

The first stage, mimicry, involves the learner reproducing all observable aspects of a demonstrated behavior, irrespective of their causal relevance. This corresponds to infants' earliest imitation behaviors, which typically reflect perceptual matching rather than intentional understanding [3]. For example, a child may copy an adult's unnecessary gestures or vocal tone while performing a simple task, not because those elements are functionally necessary, but because they are salient and socially meaningful within the interaction.

In robotic systems, mimicry can serve as a crucial bootstrapping mechanism, allowing the agent to collect rich action-perception data, form action mappings, and preserve behavioral structure that might later prove meaningful [7]. Although this strategy may include redundant or irrelevant components, such high-fidelity replication is not without value. It ensures that potentially essential but not yet understood actions are retained for later analysis and refinement as the robot's causal and goal inference capabilities develop. In this sense, mimicry functions not as a flawed imitation strategy but as a developmentally appropriate phase in a staged learning architecture.

B. Stage 2: Imitation

The second stage reflects selective imitation, where the learner begins to distinguish between essential and non-essential components of a demonstrated behavior. Children typically reach this stage in toddlerhood, developing the ability to selectively reproduce intentional, goal-directed actions while omitting inefficient or accidental components [8]. However, evidence from overimitation studies shows that young children often continue to copy causally irrelevant steps, particularly when those actions are presented in a pedagogical context, suggesting sensitivity to social and communicative cues [1-2]. For robotic systems, this stage involves segmenting action sequences, identifying elements that predict successful outcomes, and refining imitation through feedback or reinforcement learning mechanisms.

C. Stage 3: Emulation

The third stage involves emulation, where the learner achieves the same outcome as a demonstrator but through potentially novel means. In human development, this ability emerges as children begin to understand the causal structure of tasks and focus on reproducing results rather than specific actions [9]. Emulation indicates that the learner has formed an internal representation of the goal and can flexibly adapt behavior to achieve it, even when tools or environmental constraints vary. For robots, this stage depends on causal modeling, relational abstraction, and the capacity to operate across multiple strategies, enabling generalization beyond exact demonstrations.

D. Stage 4: Generalization

The fourth stage focuses on generalization, which involves applying learned behaviors or principles to novel but structurally similar tasks. In children, this capacity is supported by analogical reasoning and relational abstraction, allowing them to detect profound similarities across superficially different contexts [10]. Robots at this stage should demonstrate the ability to recognize underlying task patterns and transfer strategies across settings, tools, or goals. This stage requires integrating memory consolidation mechanisms and representation alignment frameworks, which are key components of continual learning systems that aim to minimize catastrophic forgetting while supporting adaptation over time [11].

E. Stage 5: Intent inference

The final stage is intent inference, where the learner can deduce the underlying goals or mental states of others, even when behaviors are ambiguous or incomplete. In humans, this capacity relies on the Theory of Mind (i.e., the ability to attribute beliefs, desires, and intentions to others) and plays a crucial role in flexible, socially guided learning [5]. By 12 months, infants engage in rational imitation, selectively reproducing actions they interpret as intentional or constrained [1]. In robotic systems, this maps onto the ability to interpret indirect feedback, ambiguous instructions, or evolving social norms. Robots operating at this stage should not only reproduce *what* was demonstrated but also infer *why*, adjusting behavior to better align with human expectations and unspoken goals.

V. A developmental benchmark for robot learning

To assess continual learning in robots, we propose a set of benchmark tasks modeled on paradigms from developmental psychology that reveal how children gradually shift from copying to understanding. Each task is designed to evaluate one or more dimensions of learning over time, such as imitative fidelity, causal flexibility, generalization, and social alignment. These tasks are grounded in well-established developmental theories and paired with evaluation metrics informed by research on imitation, pedagogy, and cognitive abstraction.

A. Redundant action task

In this task, a human demonstrator performs a multi-step behavior that includes both causally necessary and arbitrary or irrelevant actions. For example, before opening a box to retrieve an object, the demonstrator might first tap the lid three times—an action unrelated to task success. The robot observes several demonstrations of this sequence and is later tested to see if it replicates the entire sequence or selectively omits nonfunctional steps. This task assesses whether the robot begins with high-fidelity imitation, including irrelevant steps, but progressively refines its behavior as it gains causal insight and feedback.

This setup mirrors the developmental phenomenon of overimitation, in which children copy both relevant and irrelevant components of an action sequence. As mentioned earlier, overimitation is believed to reflect children's assumptions that all demonstrated actions are meaningful, particularly in pedagogical contexts. The redundant action task, therefore, provides insights into a robot's transition from mimicry to selective imitation, evaluating its ability to construct, revise, and generalize causal models through repeated exposure and interaction.

Performance should be measured across multiple sessions using metrics such as action efficiency (e.g., the number of steps versus the optimal steps), causal filtering score (e.g., the elimination of irrelevant actions), learning trajectory (e.g., changes over time), and responsiveness to feedback (e.g., behavioral adjustments after correction). A robot that initially imitates the full action sequence but eventually focuses on the causally relevant components would demonstrate developmental progression toward goal-sensitive emulation.

B. Causal distribution task

In this task, a method previously demonstrated to achieve a goal is rendered ineffective; for example, the tool used in the earlier demonstration is now broken or unavailable. To succeed, the robot must emulate the goal by discovering alternative means of achieving the same outcome. This task tests the robot's causal reasoning, behavioral flexibility, and its ability to decouple actions from goals, which are hallmarks of emulative rather than imitative learning.

Children begin to exhibit this capacity for emulation between the ages of four and five, transitioning from replicating exact behaviors to using novel strategies that fulfill the same function [9]. Robots modeled on this capability should develop and update internal causal models that allow them to predict outcomes based on environmental conditions, rather than simply memorizing successful action sequences. Evaluation metrics include success rate in altered context, solution diversity, and latency to adaptation.

C. Outcome reproduction task

The outcome reproduction task presents the robot with multiple demonstrations where different methods or tools are used to accomplish the same end-state goal. The robot is then asked to reproduce the goal independently, ideally using a novel or efficient approach. This task tests the ability to extract high-level, goal-invariant representations across varying perceptual or procedural conditions.

This developmental trajectory parallels children's growing capacity for analogical abstraction, in which learning shifts from superficial imitation to recognizing deeper relational similarities [10]. For robots, success in the task indicates the formation of internal representations that support goal generalization, a core requirement for lifelong learning. Key evaluation metrics include goal accuracy, novelty of the method used, and consistency across different environmental setups.

D. Analogical transfer task

The analogical transfer task introduces a novel but structurally analogous problem to one the robot has previously encountered. For example, if the robot previously learned to retrieve a toy using a stick, it is now expected to solve a similar task using a string to pull an object. Success depends on the robot's ability to recognize structural similarities, abstract prior knowledge, and transfer learned strategies to new domains.

In developmental psychology, analogical transfer marks a significant cognitive milestone, supported by structure-mapping theory, which posits that learners align relational correspondences between domains rather than surface features. For robots, this task assesses the modularity and abstraction of internal representations, as well as their applicability beyond the original learning context. Evaluation includes task success rate, degree of structural alignment, and representational similarity to prior learned solutions.

E. Novel goal task

The novel goal task challenges the robot to infer and execute a goal based on ambiguous or incomplete demonstrations. For example, a human may begin reaching toward an object but stop midway, or gesture vaguely in the direction of multiple items. Unlike tasks with explicit, full observed actions, this scenario requires the robot to rely on contextual cues, prior experiences, and social reasoning to determine the intended outcome.

This task parallels the development of Theory of Mind in children—the ability to attribute mental states such as beliefs, desires, and intentions to others, and to interpret behavior accordingly. Children typically begin to succeed at such intention-reading tasks around the age of four, as they develop the ability to infer goals from unobservable internal states [4]. A robot capable of similar inference demonstrates a deeper level of understanding: one driven not by behavior replication but by goal abstraction, predictive modeling, and social context interpretation.

Evaluation criteria for this task include goal inference accuracy (e.g., did the robot achieve the correct intended outcome?), action efficiency (e.g., did it reach the goal using minimal or optimal steps?), and feedback adaptation (e.g., can it revise its inference after receiving corrective signals?) Success in the novel task indicates a robot's capacity for socially intelligent learning, where behavior is guided not merely by demonstration but by inference, context, and intention modeling.

VI. Integrating Cognitive Evaluation Protocols

To strengthen the developmental benchmarking framework proposed in this paper, our current focus is on integrating structured cognitive evaluation protocols inspired by research in developmental psychology and recent AI benchmarks. A two-phase familiarization-test structure—where the robot first observes human behavior in repeated, structured demonstrations, followed by a generalization phase with altered constraints—can be used to probe deeper learning mechanisms. This methodology mirrors infant cognition studies, where researchers assess understanding by introducing subtle violations of previously observed patterns and measuring the learner's expectation regarding goal-directed behavior, action efficiency, and constraint sensitivity.

For instance, in the redundant action task, a robot could be familiarized with a demonstrator repeatedly performing a causally irrelevant step (e.g., tapping a lid) before opening a box. During the test phase, the context is altered (e.g., tapping action is no longer possible or visibly unnecessary), and the robot's behavior is observed. The key question becomes: does the robot continue to overimitate, or does it abstract the causal structure and omit the irrelevant action? This structured contrast allows researchers to pinpoint whether learning reflects surface mimicry or goal-based reasoning.

Additionally, future work should adapt naturalistic and collaborative tasks that require the robot to interpret evolving social expectations and feedback. Such tasks include role constraints, dynamic goals, and ambiguous social cues to test a robot's ability to act flexibly within human-centered environments. Evaluation metrics should go beyond task completion to include feedback-driven adaptation, norm generalization, and social alignment. Together, these protocols will support richer, developmentally inspired assessments of robot learning and better capture the long-term, interactive nature of continual learning in real-world settings.

However, implementing such cognitively inspired evaluation protocols presents several challenges and opens critical questions. First, designing tasks that balance ecological validity with experimental controls remains difficult, as naturalistic interactions often introduce variability that obscures fine-grained assessments of learning. Second, reliably measuring a robot's internal representations and distinguishing between surface-level mimicry and genuine causal inference requires interpretable models and multimodal behavioral analysis. Third, human feedback can be ambiguous or inconsistent, raising questions about how robots should integrate it over time to update their goals or strategies. These challenge points raise key open questions: How can we design benchmarks that adapt to individual robot learning trajectories? What types of errors are most informative about developmental progress? And how can feedback from human partners be structured to support both evaluation and improvement? Addressing these questions is critical for advancing robot learning systems that are not only technically robust but also socially responsive and cognitively grounded.

VII. Conclusion

Human cognitive development offers a powerful lens for reimagining how robots can learn in complex, interactive environments. In this paper, we have proposed a developmental framework for continual robot learning, grounded in the progression from imitation to emulation. This approach emphasizes the value of early-stage overimitation, the importance of causal abstraction, and the critical role of social inference in learning from humans over time.

Our proposed benchmark tasks reflect key dimensions of this developmental trajectory, that is, fidelity, flexibility, generalization, and social alignment. It also provides a scaffold for evaluating lifelong learning in robots. By adopting structured cognitive evaluation protocols and drawing on well-established principles from developmental science, we envision robots that not only perform tasks but also evolve alongside their human users. These systems will be more interpretable, more resilient, and more attuned to the norms and values that define effective human-robot collaboration.

References

[1] G. Gergely, H. Bekkering, and I. Király, "Rational imitation in preverbal infants," *Nature*, vol. 415, no. 6873, p. 755, Feb. 2002.

[2] D. E. Lyons, A. G. Young, and F. C. Keil, "The hidden structure of overimitation," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 104, no. 50, pp. 19751–19756, Dec. 2007.

[3] A. N. Meltzoff and M. K. Moore, "Imitation of facial and manual gestures by human neonates," *Science*, vol. 198, no. 4312, pp. 75–78, Oct. 1977.

[4] A. Gopnik, C. Glymour, D. M. Sobel, L. E. Schulz, T. Kushnir, and D. Danks, "A theory of causal learning in children: causal maps and Bayes nets," *Psychol. Rev.*, vol. 111, no. 1, pp. 3–32, Jan. 2004.

[5] H. M. Wellman, *The child's theory of mind*. in Bradford Books. London, England: MIT Press, 1992.

[6] M. Carpenter, J. Call, and M. Tomasello, "Twelve- and 18-month-olds copy actions in terms of goals," *Dev. Sci.*, vol. 8, no. 1, pp. F13–20, Jan. 2005.

[7] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Rob. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.

[8] A. L. Woodward, J. A. Sommerville, S. Gerson, A. M. E. Henderson, and J. Buresh, "The emergence of intention attribution in infancy," *Psychol. Learn. Motiv.*, vol. 51, pp. 187–222, 2009.

[9] A. Whiten, N. McGuigan, S. Marshall-Pescini, and L. M. Hopper, "Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee," *Philos. Trans. R. Soc. Lond. B Biol. Sci., vol.* 364, no. 1528, pp. 2417–2428, Aug. 2009.

[10] D. Gentner, "Structure-mapping: A theoretical framework for analogy," in *Readings in Cognitive Science*, Elsevier, 1988, pp. 303–310.

[11] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Netw.*, vol. 113, pp. 54–71, May 2019.