



HAL
open science

The SocialAI School: Insights from Developmental Psychology Towards Artificial Socio-Cultural Agents

Grgur Kovač, Rémy Portelas, Peter Ford Dominey, Pierre-Yves Oudeyer

► To cite this version:

Grgur Kovač, Rémy Portelas, Peter Ford Dominey, Pierre-Yves Oudeyer. The SocialAI School: Insights from Developmental Psychology Towards Artificial Socio-Cultural Agents. TOM 2023 - Theory-of-Mind Workshop at PMLR-202, May 2023, Honolulu (Hawaii), United States. hal-04365787v2

HAL Id: hal-04365787

<https://inria.hal.science/hal-04365787v2>

Submitted on 2 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

The SocialAI School: Insights from Developmental Psychology Towards Artificial Socio-Cultural Agents

Grgur Kovač¹ Rémy Portelas^{2,1} Peter Ford Dominey³ Pierre-Yves Oudeyer¹

Abstract

Developmental psychologists have long-established the importance of socio-cognitive abilities in human intelligence. These abilities enable us to enter, participate and benefit from human culture. AI research on social interactive agents mostly concerns the *emergence* of culture in a multi-agent setting (often without a strong grounding in developmental psychology). We argue that AI research should be informed by psychology and study socio-cognitive abilities enabling to *enter* a culture too. We discuss the theories of Michael Tomasello and Jerome Bruner to introduce some of their concepts to AI and outline key concepts and socio-cognitive abilities. We present The SocialAI school - a tool including a customizable parameterized suite of procedurally generated environments, which simplifies conducting experiments regarding those concepts. We show examples of such experiments with RL agents and Large Language Models. The main motivation of this work is to engage the AI community around the problem of social intelligence informed by developmental psychology, and to provide a tool to simplify first steps in this direction. Refer to the project website for code and additional information: <https://sites.google.com/view/socialai-school>.

1. Introduction

Our everyday life is immersed in a sociocultural world, which we navigate using a set of sophisticated socio-cognitive abilities. Although at first it might seem this sociocultural world is just another product of our cognition, decades of research in developmental psychology suggest the opposite. Our socio-cultural world, cultural knowledge, and our socio-cognitive abilities are the foundation of our development and both our social and asocial intelligence (Vygotsky & Cole, 1978; Bruner, 1990; Tomasello, 2019).

¹Inria, France ²Ubisoft, France ³CNRS, France. Correspondence to: Grgur Kovač <grgur.kovac@inria.fr>.

For Vygotsky, socio-cultural interactions are the main driver for “higher-level” cognition (Vygotsky & Cole, 1978). For him, many high-level cognitive functions first appear at the social level, and *then* develop at the individual level. This leap from interpersonal processes to intrapersonal processes is referred to as *internalization*. Vygotsky’s theories influenced multiple works within cognitive science (Clark, 1996; Hutchins, 1996), primatology (Tomasello, 1999) and the developmental robotics branch of AI (Billard & Dautenhahn, 1998; Cangelosi et al., 2010; Mirolli & Parisi, 2011).

Jerome Bruner emphasized the importance of culture in human development too. He presents a pragmatic view of how the practical use of referencing and requesting pushes language development through routinized social interactions (formats), in which those abilities are *necessary* to achieve various ends. He describes these interactions as scaffolded - the caretaker gradually helps less and demands more of the child to achieve those goals (Bruner, 1985).

Finally, Michael Tomasello’s work (Tomasello, 1999; 2019; 2020) constitutes a representative and contemporary assessment of the nature and central importance of culture in human development. Tomasello outlined core social abilities and motivations which (when combined with the relevant experience) enable us to participate in the cumulative cultural evolution (a powerful form of cultural transmission enabling the development and perpetuation of complex cultural artifacts and knowledge (Tomasello, 1999)).

Given the key role social cognition plays in human cognition and cultural evolution, it is natural that the field of AI studies those questions as well. A socially competent agent could learn our culture and participate in its cultural evolution, i.e. improve our concepts, theories, inventions, and create new ones. A system capable of out-of-the-box thinking creative solutions and discovering new relevant problems must learn our values and how we see and understand the world (it must learn our culture).

Enriching AI with those skills also has numerous more practical implications. Agents capable of online social learning could efficiently learn novel tasks and tools. A robot that can seamlessly infer the meaning of our gestures and utterances in new contexts could be easily *onboarded* into

human teams, without requiring humans to adopt new conventions. Furthermore, robots capable of learning human values and moral norms will be capable of performing tasks in the constraints defined by those values.

AI research on interactive agents is often focused on navigation and object manipulation problems, excised of any social dimension (Mnih et al., 2015; Lillicrap et al., 2016; Andrychowicz et al., 2017). The study of sociality is mostly studied in Multi-Agent settings, where the main focus is often on the *emergence* of culture (often with only a weak grounding in developmental psychology) (Jaques et al., 2019; Baker et al., 2019). While we agree that those directions are interesting and important, in this work we focus on *entering* an already existing complex culture. And we argue that it can be beneficial to be informed by developmental psychology theories.

We do not claim The SocialAI is sufficient to reach that far and complex goal. We only propose that being informed by the concepts discussed in this paper is useful, and we present SocialAI as a tool which could be used to start investigating such questions in more details.

Following the theories of Michael Tomasello and Jerome Bruner, this work identifies a richer set of socio-cognitive skills than those currently considered in most of the DRL literature. More precisely, we focus on three key aspects of social cognition as identified by Tomasello: 1) the ability to infer what others see and to engage in joint attention, 2) the development of referential communication through pointing and the beginnings of conventionalized communication through language, and 3) the use of imitation and role reversal imitation in social learning. We also outline two concepts from Jerome Bruner’s work: formats and scaffolding. Formats refer to the way in which social interactions are structured and presented, while scaffolding refers to the temporary support provided by a caretaker to help a learner achieve a task that would be otherwise too difficult.

We present The SocialAI school - a tool simplifying experiments studying the concepts outlined by those theories. It includes a parameterized procedural generation engine for social environments with various kinds of social interactions. This tool enables simple creation, modification, and extension of those environments. In our experiments, we show how various studies can be conducted using the SocialAI school. We present experiments regarding the following questions: generalization of social inferences (the pointing gesture) to new contexts, recreating an experiment from cognitive science (to study role reversal), and the impact of a scaffolded environment on the agent’s learning. We conduct those experiments with RL agents, and also present an additional case study with Large Language Models (LLMs). In the appendix, we explore additional questions regarding linguistic inferences, joint attention, imitation, inferring

others’ field of view, and formats. We hope to encourage future work extending and building on this first set to study various questions regarding social competence (ex. new sociocultural scenarios, architectures, training regimes).

We outline the following main contributions of this work:

- An introduction to Michael Tomasello’s and Jerome Bruner’s theories on child development and core socio-cognitive abilities
- An outline of a set of core socio-cognitive abilities important for current AI research
- The SocialAI school: a tool including a customizable procedural generation suite of environments aiming to simplify studies of socio-cognitive abilities of AI agents
- Examples of case studies demonstrating how SocialAI can be used to study various questions regarding socio-cognitive abilities in AI

2. Related work

This work aims to connect DRL with developmental robotics (Asada et al., 2009; Cangelosi & Schlesinger, 2014), a robotics field informed by developmental psychology. Developmental robotics has already argued for the importance of social intelligence in artificial agents (Billard & Dautenhahn, 1999; Lindblom & Ziemke, 2003; Mirolli & Parisi, 2011). We aim to expand this to the DL community.

Inside the deep learning community, multiple works studied disembodied social understanding from videos or text. These include classifying videos based on the nature of shown behavior (Shu et al., 2021; 2020), inferring agents’ goals, relationships, and predicting future trajectories (Nentanyahu et al., 2021). Machine learning models have also been used to model agents’ internal states, such as goals, beliefs, and desires (Rabinowitz et al., 2018; Baker et al., 2011). Regarding textual tasks, social reasoning abilities have been recently studied with LLMs. In extensive evaluation, LLMs were shown to struggle on two social benchmarks (Sap et al., 2022): SocialIQA (Sap et al., 2019), and especially on TOMi (Le et al., 2019). In other experiments, LLMs were evaluated on variations of false-belief tasks and exhibited promising performance (Trott et al., 2022; Kosinski, 2023). A gap with human performance was also demonstrated on implicatures - problems which heavily rely on inferring contextual information (Ruis et al., 2022). While our motivation is analogous to these works, we propose to focus on embodied and interactive agents.

DRL works studying embodied social interactions are mostly in multi-agent settings. Various intrinsic rewards were presented to foster cooperation between agents such as

influence over the others' actions (Jaques et al., 2019) and fostering joint attention (Lee et al., 2021). While multi-agent reinforcement learning is often used to study the emergence of culture or communication, here we study the process of *entering* an already existing culture.

Closer to our work, an RL agent was shown to adapt (through social learning) to a new environment with an expert (Ndousse et al., 2021). The independent RL agent was trained in a multi-agent environment with various environmental constraints and an auxiliary loss. Similar experiments were also conducted at a larger scale (Bhoopchand et al., 2022). The objective of the SocialAI school is to provide a tool simplifying similar studies, which could explore socio-cognitive abilities outlined by psychology.

3. Cognitive science background

This section introduces Michael Tomasello's and Jerome Bruner's theories and concepts.

3.1. M. Tomasello - The Shared Intentionality Theory

We are born into a culture filled with cultural artifacts, symbols and institutions like language, social norms, tool industries, or even governments (Richerson & Boyd, 2006; Tomasello, 2019). These artifacts are a product of a series of modifications over many generations. Tomasello calls this *cumulative cultural evolution*, and argues that it is behind our most impressive achievements (Tomasello, 1999).

Cumulative cultural evolution is grounded in our socio-cognitive abilities (e.g. social cognition, cultural learning, communication), which enable us to learn, improve, and teach our culture (Tomasello, 2019), i.e. *enter* a culture. Cultural artifacts inherited through this process become the core of our cognition. An example of this is language, which influences our cognition in many ways. For example, it defines how we categorize and construe the world, and enables a powerful form of social learning : instructed learning (Tomasello, 1999). This makes socio-cognitive abilities crucial, as their early development bootstraps both our social and asocial cognition (Herrmann et al., 2007).

Tomasello's *Shared intentionality theory* argues that human socio-cognitive abilities, such as communication and social learning, are transformed by two developmental steps : the emergence of *Joint intentionality* at around 9 months of age (the 9-month revolution), and the emergence of *Collective intentionality* at around 3 years of age (the objective/normative turn) (Tomasello, 2019).

Joint intentionality emerges at around 9 months of age (Tomasello, 2019). It enables children to form a *joint agent* (a dyadic "we") - they understand that they work with a partner towards the same joint goal. Children begin to view

dyadic social interactions through a "*dual-level structure*": a joint agent "we" on one level, and a personal "I" on another, i.e. we both understand that we both have separate roles ("I"), and that we work together towards the same joint goal ("we"). This enables them to take the perspective of others, which can also be done recursively - they are not only both attending to the same goal, they are also both attending to the partner's attention to the goal, and they both know that they both are doing so.

Collective intentionality emerges at around 3 years of age (Tomasello, 2019). It enables children to form a cultural *group-minded "we"*, which in comparison with a dyadic "we" represents an identity for a group. For example, a child might enforce a social norm because "this is how *we*, in this culture, do things". Consequently, children begin to participate in conventions and norms, and to view things from the "objective" perspective.

These two developmental steps transform countless abilities, motivations, and behaviors. For the purpose of this paper, we focus on the following three developmental pathways: social cognition (sec. 3.1.1), communication (sec. 3.1.2), and social learning (sec. 3.1.3), as we consider them the most relevant for AI at the moment.

3.1.1. SOCIAL COGNITION

In this section, we discuss the development of the ability to coordinate perspectives and view things from the *objective perspective* (a perspective independent from any individual) (Tomasello, 2019). The starting point is the ability to **infer what another sees or knows**. The earliest example of this is gaze following of six-month-olds (D'Entremont et al., 1997). Here, only one perspective is processed at the time. **Joint attention (JA)** emerges at around 9 months of age. Tomasello defines JA as consisting of two elements: *triangulation* (two participants attending to the same referent) and *recursiveness* (both participants being recursively aware that they are both sharing attention) (Tomasello, 2019). JA is characterized by the dual-level structure of shared attention on one level, and individual perspectives on another. Consequently, children start to align and exchange perspectives. Once children reach a sufficient level of linguistic competence, they start sharing attention to mental content in the form of linguistic discourse (at two to three years of age). The presence of conflicting perspectives in linguistic discourse pushes children to resolve those conflicts, which they do by forming the "objective" perspective, and **coordinating other perspectives** with it. Refer to appendix B.1.1 for details.

3.1.2. COMMUNICATION

Communication starts with **imperative gestures for self-serving purposes**. An example of such a gesture is the

child pulling the adult's hand, requesting them to pick them up. This gesture always has the same imperative meaning, and it never refers to an external object. The 9-month revolution brings forth **referential communication** - children start to communicate triadically to external referents through pointing and pantomiming. The pointing gesture is a powerful way of communicating, as the *same* gesture can be used to express many different meanings in many different scenarios, provided that the observer can correctly infer that meaning. The ability to infer this meaning is based on the emerging abilities of joint intentionality. Those of joint attention and, most notably, "recursive inference" - to interpret a pointing gesture, we make a recursive inference of what "you intend for me to think". For example, if we are looking for a ball together, and you point to a cupboard behind me. I should infer that you are drawing my attention to the cupboard to communicate that I should look for the ball in the cupboard. The next step is the appearance of **conventionalized linguistic communication**. The underlying principle stays the same: reference to an external entity combined with inferring the meaning through recursive inferences. The difference is that, now, the child also learns the conventional means of referring (for example, words and phrases). Tomasello argues that, at first, children don't understand language as conventional, and they use it as any other tool. The understanding of language as conventional follows the emergence of collective intentionality after the third birthday, and this gives rise to a myriad of different language uses, such as discourse or pedagogy. Refer to section B.1.2 for more details.

3.1.3. CULTURAL LEARNING

Human culture is characterized by a powerful form of cultural transmission called cumulative cultural evolution - inventions quickly spread and are improved by following generations (Tomasello, 1999). These inventions spread at such a pace that they are rarely forgotten or lost. This is referred to as the *ratchet* effect (Tomasello et al., 1993) - inventions are iteratively improved without *slippage* back. This effect is enabled by human social learning abilities (ex. imitation, instructed learning), and motivations (to learn from others, but also to affiliate and conform). The earliest form of cultural learning is the **mimicking of facial expressions** (observed even in neonates (Meltzoff & Moore, 1997)). Over the course of the first year, children begin to **imitate other's actions and goals**, and then, they begin doing so in ways which demonstrate their understanding of other's as intentional agents (Meltzoff, 1995). Joint intentionality brings forth a new form of cultural learning called **role reversal imitation**. Children can reverse the roles of a collaborative activity by learning about the partners role only from playing their own. For example, children respond to an adult tickling their arm, by tickling the adult's arm

(instead of its own) (Carpenter et al., 2005). This is enabled by the dual-level structure of joint intentionality through which children understand, at the same time, the joint goal of a dyadic interaction on one level, and the individuals' separate roles on another. The next big step in the development of cultural learning is learning from instructions - **instructed learning** (following the emergence of collective intentionality). It is based on the adults' motivation to teach children as well as on the children's ability to understand and learn from linguistic instructions. Children understand knowledge acquired through instructions as objective truth, and generalize it much better than knowledge acquired by other means (Butler & Tomasello, 2016). In this way we acquire the most complex knowledge and skills such as reading or algebra. Refer to section B.1.3 for more details.

3.2. Jerome Bruner

This work is also influenced by Jerone Bruner's theories, especially regarding the concepts of scaffolding (Wood et al., 1976) and formats (Bruner, 1985), which were recently reintroduced to AI as pragmatic frames (Vollmer et al., 2016).

Formats (Pragmatic frames) (Bruner, 1985) simplify learning by providing a stable structure to social interactions. They are regular patterns characterizing the unfolding of possible social interactions (equivalent to an interaction protocol or a grammar of social interactions). Formats consist of a deep structure (the static part) and a surface structure (the varying realizations managed by some rules). An example of a format is the common peek-a-boo game. The deep structure refers to the appearance and the reappearance of an object. The surface structure can be realized in different ways. For example, one might hide an object using a cloth, or hands; one might hide his face or a toy; one might do shorter or longer pauses before making the object reappear. We understand social interactions through such formats, and our social interactions are based on our ability to learn, negotiate, and use them.

Another relevant concept is scaffolding (Wood et al., 1976) (similar to Vygotsky's zone of proximal development (Vygotsky & Cole, 1978)). Scaffolding is a process through which an adult bootstraps the child's learning. The adult controls aspects of a task which are currently too hard for the child (scaffolds the interaction). The scaffold is gradually reduced as the child is ready to take on more aspects of the task, until they can solve the task alone (without scaffolding). An example is a child constructing a pyramid with the help of an adult (Wood et al., 1976). At first, the child is not even focusing on the task, and the adult tries to get its attention to the task by connecting blocks and building the pyramid in front of them. Once the child is able to focus on the task, the adult starts passing the blocks to the child to connect. In the next phase, the child is grabbing blocks by itself, and

the adult is helping through verbal suggestions. Then, only verbal confirmations are needed to guide the child. Finally, the child can construct the pyramid by itself. In summary, the adult observes the child and gradually transfers parts of the task (removes the scaffold) to the child. Through this process, the caretaker enables the child to master a task they would not be able to master alone.

Theory of Mind (ToM) can be defined as the understanding that others have intentions, desires, beliefs, perceptions, and emotions different from one’s own and that such intentions, desires, and so forth affect people’s actions and behaviors (APA, 2023). Theories discussed here (especially Tomasello’s) do not use this term much, rather they separate it into more fine-grained ones. For example, the ability to “imagine what another knows” require the processing of only one perspective, this is only the first step which later develops to the ability to coordinate multiple perspectives, and finally to form the objective perspective and address false-beliefs. Another example is the concept of recursiveness in JA: it is not sufficient to infer what another knows, but also that another knows about what you know, and so on. The socio-cognitive abilities discussed here are different aspects of ToM. Some of these those are a clear subset of the usual definition (ex. imagine what another knows), and some go beyond (coordinate multiple perspectives).

4. The SocialAI school

The SocialAI school is a tool for building interactive environments to study various questions regarding social competence, such as “What do concepts, such as social abilities and motivations, outlined by developmental psychology mean in the scope of AI?”, “How can we evaluate their presence in different agents?”, “What are their simplest forms and how can agents acquire them?”

To construct SocialAI, we rely on a set of key experiments and studies from developmental psychology, which were used to outline the most important abilities, motivations and developmental steps in humans. From the work of Tomasello, we focus on developments before and around the age of 9 months (we believe it is important to address those before more complex ones relating to development of 3-year-olds, see section 3.1). We study the following developmental pathways: Social cognition (inferring other’s perception and joint attention), Communication (referential communication through the pointing gesture and the beginning of conventionalized communication through simple language), and Cultural Learning (imitation and role reversal imitation). From the work of Bruner, we study the concepts of Formats and Scaffolding (see section 3.2).

SocialAI, which is built on top of Minigrad (Chevalier-Boisvert et al., 2018), includes a *customizable parameter-*

ized suite of *procedurally* generated environments. We implement this procedural generation with a tree-based structure (the parametric tree). This makes it simple to add and modify new environments, and control their sampling. All the current environments are single-agent and contain a scripted peer. The agent has to interact with the peer to reach an apple. This setup enables a controlled and minimal representation of social interactions. SocialAI also includes RL-based and LLM-based learners and evaluation protocols.

In our experiments, we present case studies as examples of how SocialAI could be easily used and modified to ask various questions and conduct diverse experiments. To facilitate future research, SocialAI was made to be very easy to use and modify. It will be completely open sourced, and we hope that it will be useful to the community to study the questions of social intelligence in AI.

We do not claim that the SocialAI school is sufficient to construct a socially competent agent as this is a very far-reaching and complex goal. However, we believe that in aiming for this goal, concepts from developmental psychology can serve as signposts for AI - give directions and enable us to define short term goals. Given that the outlined skills are at the very core of human social and cognitive competences, artificial agents aimed at participating in and learning from social interactions with humans are likely to require the same core competences. We present the SocialAI school merely as a first step towards this goal. Refer to section C in the appendix for technical details on the environment, agents, and the procedural generation mechanisms.

5. Experiments

We show how the SocialAI school can be used to conduct experiments inspired by theories and studies described in section 3. In all our case studies, except the study with language models (sec. 5.4), we use a PPO (Schulman et al., 2017) agents with different exploration bonuses (refer to appendix D for technical details). We study the pointing gesture, role reversal imitation and scaffolding, and present a study with large language models. In appendix G we present additional experiments regarding linguistic cues, joint attention, in-episode imitation learning, the ability to infer the peer’s field of view, and formats.

5.1. Understanding the pointing gesture

In this experiment, we study the ability of an RL agent to understand the pointing gesture. This experiment is motivated by a study of childrens’ ability to understand pointing gestures (Behne et al., 2005) (see section B.1.2 in the appendix). We study can an RL agent infer the meaning of a pointing gesture, and generalize this ability to new situations (infer




Figure 1. The Pointing experiments. The figure compares the success rate (mean \pm std over 8 seeds) on the training environments with the evaluation on the testing environment. The cross marks depict statistical significance ($p = 0.05$). The agent is able to infer the meaning of a pointing gesture on the training environments, but it is not able to generalize to a new social context.

the new meaning of a pointing gesture in a new context). It is interesting to study this kind of generalization as the power of inferring pointing gestures is based precisely on being able to infer its meaning to *new* referents based on *new* social contexts.

The environment consists of two objects (ex. boxes) and the peer that points to the correct object. The agent then has to interact with that object (ex. open the box) to get access to an apple. The agent is trained on five problems each with different objects (Boxes, Switches, Levers, Marble, Generators), and on the *asocial* version of the Doors problem (only one door and no peer). Training on the *asocial* version enables the agent to learn how to use a door, which is a prerequisite for generalization of the pointing gesture to an environment with two doors. The agent is evaluated on the Doors problem in the social setting (two doors and a peer pointing to the correct one) The agent needs to combine the knowledge of how to use a door (learned on the *asocial* version of that problem), with inferring the meaning of the pointing gesture (learned on the other five problems), and generalize that to a new scenario where the peer points to a door. Refer to section F.1 in the appendix for details.

Figure 1 shows the success rate of the agent on the training environments (“PPO_CB_train”) and its evaluation on the evaluation environment (PPO_CB(test)). We can see that while the agent easily solves the training environments (with the success rate of 95.2%), it fails to generalize. It reaches the success rate of 45.2%, which corresponds to randomly guessing the correct object. These results demonstrate that the agent can learn to infer the meaning of a pointing gesture in a familiar context, but cannot generalize to new social contexts. These results motivate future research on how an agent can be endowed with abilities for such combinatorial generalization, a potential solution could leverage LLMs.

Appendix G.1 presents experiments in which the peer, instead of pointing, provides linguistic cues for the color or the proximity of the correct object. As in the pointing experiments, we observe that while PPO agents master the training

environments, they fail to generalize to a new context.

5.2. Role reversal imitation

In this experiment, we study the role-reversal capabilities of an RL agent - to what extent can it learn about the partner’s role from playing its own. In doing so, we also show how a cognitive science experiment can be recreated in the scope of AI. In Fletcher (2012) apes and children were trained on one role (role B), and then tested on how long it took them to master the opposite role (role A). Results showed that children, but not apes, master role A faster than the control group (not pretrained). These results imply that children learn about the opposite role just from playing their own, i.e. they see the interaction from a bird’s eye perspective. We study the following two questions: 1) How much do RL agents learn about the partner’s role during a collaborative activity? 2) Does increasing diversity in the training (training on more tasks in both roles) enable the agent to learn more about the partner’s role?

We conduct this study on the MarblePass task. This task consists of two roles: one participant pushes the marble to the right side of the environment (role A), from where the other can push it to the a generator, which generates apples (role B). We aim to assess how much the agent learns about the opposite role (role A), from training in its own (role B). Following Fletcher (2012) we measure the sample efficiency of fine-tuning agents to the test role. Unlike in Fletcher (2012) it is not sufficient to compare an agent pretrained on the training role with an unpretrained agent. Even if the agent pretrained on the training role learns nothing about the testing role, it would still learn about environment dynamics and one would expect it to learn faster than the unpretrained agent. For this reason, we compare with an agent pretrained on the *asocial* version of the training role. In this version, the agent obtains reward in the same way as in the social version, but no peer is needed - the agent and the marble are placed on the right side of the environment and the agent has to push the marble towards the generator. Therefore, this agent learns all about the relevant environment dynamics, but not about the specific collaborative activity. This agent represents the control group in Fletcher (2012).

We conduct two experiments: *single* and *group*. In *single* experiments, the agents are trained only on one task : role B and the *asocial* version of the MarblePass problem. In *group* experiments, both agents are also trained both roles of all additional six collaborative problems (a total of 13 environments). In other words, we compare the agents pretrained in the four following ways: 1) experimental (*single*): pretrained only on role B of the MarblePass problem, 2) control (*single*): pretrained only on the *asocial* version of the MarblePass problem, 3) experimental (*group*): pretrained on role B of the MarblePass problem, and on both roles of all

other problems, 4) control (*group*): pretrained on the asocial version of the MarblePass problem, and on both roles of all other problems. Refer to appendix for additional details.

How much do RL agents learn about the partner’s role during a collaborative activity? Figure 2(a) shows the success rate of fine-tuning to role A of the MarblePass task. It compares the experimental and the control conditions of the *single* experiments. It is interesting to note that the agent pretrained on the asocial version (“asocial”) masters role A of the task slightly faster than the agent pretrained on role B of the task (“role.B”). This implies that, not only, the agent does not learn anything useful about the peer’s role, but pretraining on role B actually makes it harder for the agent to learn about role A. We believe that this is because, during training in role B, the agent learns to first wait for the peer, while in the asocial version it pushes the marble right away. As, in role A, the agent pushes the marble right away too, we believe this makes it slightly easier for the asocially pretrained agent to adapt to the new role. In other words, from an egocentric view the asocial version is closer (than role B) to role A. This shows that the RL agent, rather than understanding the interaction from a bird’s-eye perspective, finds the simplest way to solve the task.


Does training on additional problems enable the agent to learn more about the partner’s role? Figure 2(b) shows the success rate of fine-tuning to role A of the MarblePass task. It compares the experimental and the control conditions of the *group* experiments. Here we can see that there is no significant difference in sample efficiency. We can make two observations from this. First, as the socially pretrained agent was less sample efficient in the *single* experiments, we can conclude that pretraining on many tasks reduces overfitting on role B. And second, as this agent is not more sample efficient than the asocially pretrained baseline, we can conclude that this agent does not learn anything useful about the peer’s role too.

These results imply an interesting avenue of research into how agent’s attention can be directed to the partner’s role and the birds-eye-view of the activity.

5.3. Scaffolding

In this section, we study the concept of scaffolding (see sec. 3.2 for details): Can a scaffolded environment help an agent learn more complex interaction sequences (formats)? The environment is similar to the one in section 5.1. However, we evaluate on all six problems (instead of one) in the social version, and the peer does not point until the agent performs a more complex introductory sequence (establishes eye contact and utters “Help, please”).

We compare two types of scaffolding: “scaf_4” and “scaf_8”. The agent is trained in two phases. In the first phase, the



(a) *Single* experiment: learning role A given pretraining on role B (1 environment). (b) *Group* experiment: learning role A given pretraining on role B and 6 other two-roles tasks (13 environments).

Figure 2. Role reversal imitation experiments. We study to what extent is an RL agent able to transfer knowledge from one role of a collaborative activity to another. Figure shows the success rate of fine-tuning to role A (mean \pm std over 8 seeds), the cross marks depict statistical significance ($p = 0.05$). We compare a PPO agent pretrained on role B (“role.B”) to that pretrained on the asocial version of the environment (“asocial”), which learns only about the environment dynamics. Agents pretrained on role B do not master role A faster than asocially pretrained agents, implying that the RL agents do exhibit role reversal capabilities.

agent is trained on environments with varying complexity (defined by scaffolding type). After reaching a set success rate, the training goes to the second phase in which the agent is trained only on the six testing environments. The agent denoted by “scaf_4” is trained on four different introductory sequences (requiring or not requiring eye contact and the utterance). This agent is trained on 18 different environments (six problems, four sequences). The “scaf_8” agent is also trained with those four different options. In addition, the peer can help in two different ways: pointing to the object or interacting with it and leaving the apple for the agent to eat (36 environments). The easiest environments on which the “scaf_8” agent is trained do not require an introduction and the peer leaves the apple for the agent (the agent just goes to and eats the apple). The hardest ones require the introduction with both the utterance and eye contact and the peer points to the object. The agents are evaluated on those.

Figure 3 compares the success rate of the agents trained with the two scaffolding types (“scaf_4” and “scaf_8”) to that of an agent trained only on the six testing environments (“no_scaf”). We can see that only the scaffolded agents solve the testing environments, and that the agent with a more detailed scaffolding (“scaf_8”) solves the environment faster. These results show that scaffolding enables the agents to learn more complex formats, and that a more thorough scaffolding further improves the efficiency. In future work, more advanced scaffolding could be explored, ex. based on learning progress (Oudeyer & Kaplan, 2007) or other surrogate objectives (Portelas et al., 2020).




Figure 3. The scaffolding experiment. The comparison of agents trained on multiple environments of varying difficulty to that trained on an unscaffolded environment. The figure shows success rates on the testing environments (mean \pm std over 8 seeds) and the cross marks depict statistical significance ($p = 0.05$) with respect to the "no_scaf" baseline. Only the scaffolded agents ("scaf_4" and "scaf_8") solve the environment, and the scaffolding with eight difficulty levels is more sample efficient.

5.4. Large language models as interactive agents

Large language models (LLMs) are starting to be used in various tasks (Brown et al., 2020; Touvron et al., 2023), including to control interactive agents (Yao et al., 2022; Carta et al., 2023). In order to study LLMs as interactive agents, SocialAI school enables parsing of visual grid observations to text (this parsing can be easily modified and experimented with). This process is depicted in figure 24 in the Appendix.

We use two environments: AsocialBox and ColorBoxes. The AsocialBox environment contains a box, which the agent has to open to get the apple. The ColorBoxes environment contains two boxes and a peer. The peer says the color of the correct box, which the agent then has to open.

An LLM acts by continuing a prompt containing in-context examples, observations and actions from the current episode (last 3 steps for ColorBoxes and the full episode for AsocialBox as this gave the best performance), and the action query ("Act :"). We manually create expert trajectories for in-context examples - 7 episodes (545 words) for the AsocialBox environment, and 9 (687 words) for ColorBoxes (the full in context examples are given in appendix G.6). The model generates the textual continuation of this prompt (3 tokens for GPT models, and 3 words for bloom). If one of the available actions ("turn left", "turn right", "move forward", "toggle") is a substring of the generated text, that action is executed, otherwise, the "no_op" action is executed (the agent does not act this step). The executed action and the new observation are then added to the prompt.

We compare three LLMs: bloom-560m (Scao et al., 2022), and two GPT models (Brown et al., 2020), ada ("text-ada-001") and davinci-3 ("text-davinci-003"). We compare them to a baseline that samples a random action. We evaluate on a fixed test set of 20 environments, with 10 timesteps.

Table 1 shows that, on the AsocialBox environment, the GPT models achieve a success rate of 90%, despite only

Table 1. Comparison of large language models on two SocialAI environments (success rate on 20 environments). The best model ("davinci-3") reaches the success rates of 90% and 35%. While this performance is impressive given that the models observed only seven (for AsocialBox) and nine (for ColorBoxes) expert trajectories, it leaves much room for improvement.

	ada	davinci-3	bloom-560m	rand
AsocialBox	90%	90%	75%	5%
ColorBoxes	10%	35%	10%	0%

observing seven expert trajectories. On ColorBoxes, davinci-3 outperforms all other models with a success rate of 35%. As this is below 50%, this implies that none of the models are able to use cues from the peer. All models outperform the random values and both GPT models outperform bloom-560m. The environments used in this case study are much simpler than those in other case studies (only one problem, and no introductory sequence). Nonetheless, it is impressive that such a performance is achieved from only a few expert trajectories: seven for AsocialBox and nine for ColorBoxes.

We are optimistic that LLM-based agents could solve these and much more complex tasks with various improvements such as planning (Huang et al., 2022), chain-of-thought (Wei et al., 2022; Zhang et al., 2023), or fine-tuning (Carta et al., 2023). As the main motivation of this case study was to show that it is easy to study LLMs with the SocialAI school, we leave those experiments for future work.

6. Conclusion

We highlight the importance of socio-cognitive abilities in AI research and present an introduction to Michael Tomasello's and Jerome Bruner's theories of socio-cognitive development. Following these theories, we outlined a set of key socio-cognitive abilities and concepts for AI: social cognition (inferring other's perception and joint attention), communication (referential and early conventionalized communication), cultural learning (imitation and role reversal imitation), scaffolding, and formats.

We present the SocialAI school - a tool simplifying the research of core socio-cognitive abilities. We show how the SocialAI school can be used to easily create environments studying various questions inspired by developmental psychology. With RL agents, we conduct experiments regarding the pointing gesture, scaffolding, and role reversal (by recreating an experiment from developmental psychology). We demonstrate that, by using SocialAI to parse environments into text, Large Language Models can be easily studied too. In the appendix, we present additional studies concerning linguistic communication, joint attention, imitation learning, inferring others' field of view, and formats.

Our experiments demonstrated the diversity of studies that can be conducted with the SocialAI school, highlighted the limitations of standard RL agents, and showed that while large language models learn with high sample efficiency, additional methods such as fine-tuning or chain-of-thought might be needed.

References

- Abramson, J., Ahuja, A., Brussee, A., Carnevale, F., Cassin, M., Fischer, F., Georgiev, P., Goldin, A., Harley, T., Hill, F., Humphreys, P. C., Hung, A., Landon, J., Lillicrap, T. P., Merzic, H., Muldal, A., Santoro, A., Scully, G., von Glehn, T., Wayne, G., Wong, N., Yan, C., and Zhu, R. Creating multimodal interactive agents with imitation and self-supervised learning. *CoRR*, abs/2112.03763, 2021. URL <https://arxiv.org/abs/2112.03763>.
- Andreas, J. Language models as agent models. In *Conference on Empirical Methods in Natural Language Processing*, 2022.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, O. P., and Zaremba, W. Hindsight experience replay. In *NeurIPS*, 2017.
- APA, A. P. A. Theory of mind, 2023. URL <https://dictionary.apa.org/theory-of-mind>. Accessed: 2023-06-22.
- Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M., and Yoshida, C. Cognitive developmental robotics: A survey. *IEEE Transactions on Autonomous Mental Development*, 1(1):12–34, 2009. doi: 10.1109/TAMD.2009.2021702.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. *ArXiv*, abs/1909.07528, 2019.
- Baker, C. L., Saxe, R., and Tenenbaum, J. B. Bayesian theory of mind: Modeling joint belief-desire attribution. In Carlson, L. A., Hölscher, C., and Shipley, T. F. (eds.), *Proceedings of the 33th Annual Meeting of the Cognitive Science Society, CogSci 2011, Boston, Massachusetts, USA, July 20-23, 2011*. cognitivesciencesociety.org, 2011. URL <https://mindmodeling.org/cogsci2011/papers/0583/index.html>.
- Behne, T., Carpenter, M., and Tomasello, M. One-year-olds comprehend the communicative intentions behind gestures in a hiding game. *Developmental Science*, 8(6):492–499, 2005. doi: <https://doi.org/10.1111/j.1467-7687.2005.00440.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-7687.2005.00440.x>.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.
- Bhoopchand, A., Brownfield, B., Collister, A., Lago, A. D., Edwards, A. D., Everett, R., Frechette, A., Oliveira, Y. G., Hughes, E., Mathewson, K. W., Mendolicchio, P., Pawar, J., Pislari, M., Platonov, A. K., Senter, E., Singh, S., Zacherl, A., and Zhang, L. M. Learning robust real-time cultural transmission without human data. *ArXiv*, abs/2203.00715, 2022.
- Billard, A. and Dautenhahn, K. Grounding communication in autonomous robots: An experimental study. *Robotics and Autonomous Systems*, 24(1):71 – 79, 1998. ISSN 0921-8890. doi: [https://doi.org/10.1016/S0921-8890\(98\)00023-2](https://doi.org/10.1016/S0921-8890(98)00023-2). URL <http://www.sciencedirect.com/science/article/pii/S0921889098000232>. Scientific Methods in Mobile Robotics.
- Billard, A. and Dautenhahn, K. Experiments in learning by imitation - grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3-4):415–438, 1999. doi: 10.1177/105971239900700311. URL <https://doi.org/10.1177/105971239900700311>.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Bruner, J. Child’s talk: Learning to use language. *Child Language Teaching and Therapy*, 1(1):111–114, 1985. doi: 10.1177/026565908500100113. URL <https://doi.org/10.1177/026565908500100113>.
- Bruner, J. *Acts of meaning*. Harvard university press, 1990.
- Burda, Y., Edwards, H., Storkey, A. J., and Klimov, O. Exploration by random network distillation. *CoRR*, abs/1810.12894, 2018. URL <http://arxiv.org/abs/1810.12894>.
- Butler, L. P. and Tomasello, M. Two- and 3-year-olds integrate linguistic and pedagogical cues in guiding inductive generalization and exploration. *Journal of Experimental Child Psychology*, 145:64–78, 2016. ISSN 0022-0965. doi: <https://doi.org/10.1016/j.jecp.2015.12.001>. URL <https://www.sciencedirect.com/science/article/pii/S0022096515002945>.
- Cangelosi, A. and Schlesinger, M. *Developmental Robotics: From Babies to Robots*. The MIT Press, 2014. ISBN 0262028018.

- Cangelosi, A., Metta, G., Sagerer, G., Nolfi, S., Nehaniv, C., Fischer, K., Tani, J., Belpaeme, T., Sandini, G., Nori, F., et al. Integration of action and language knowledge: A roadmap for developmental robotics. *IEEE Transactions on Autonomous Mental Development*, 2(3):167–195, 2010.
- Carpenter, M., Akhtar, N., and Tomasello, M. Fourteen- to 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, 21: 315–330, 1998a. ISSN 0163-6383.
- Carpenter, M., Nagell, K., and Tomasello, M. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63 4:i–vi, 1–143, 1998b.
- Carpenter, M., Tomasello, M., and Striano, T. Role reversal imitation and language in typically developing infants and children with autism. *Infancy*, 8(3): 253–278, 2005. doi: 10.1207/s15327078in0803_4. URL https://www.tandfonline.com/doi/abs/10.1207/s15327078in0803_4.
- Carta, T., Romac, C., Wolf, T., Lamprier, S., Sigaud, O., and Oudeyer, P.-Y. Grounding large language models in interactive environments with online reinforcement learning. 2023.
- Carvalho, E. Social affordance. In Vonk, J. and Shackelford, T. (eds.), *Encyclopedia of Animal Cognition and Behavior*, pp. 1–4. 2020.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- Chevalier-Boisvert, M., Willems, L., and Pal, S. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Chevalier-Boisvert, M., Bahdanau, D., Lahlou, S., Willems, L., Saharia, C., Nguyen, T. H., and Bengio, Y. Babyai: A platform to study the sample efficiency of grounded language learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJeXC0cYX>.
- Chung, J., Gülçehre, Ç., Cho, K., and Bengio, Y. Gated feedback recurrent neural networks. In Bach, F. R. and Blei, D. M. (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 2067–2075. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/chung15.html>.
- Clark, A. *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA, USA, 1st edition, 1996. ISBN 0262032406.
- De Waal, F. *Are we smart enough to know how smart animals are?* WW Norton & Company, 2016.
- D’Entremont, B., Hains, S., and Muir, D. A demonstration of gaze following in 3- to 6-month-olds. *Infant Behavior and Development*, 20(4):569–572, 1997. ISSN 0163-6383. doi: [https://doi.org/10.1016/S0163-6383\(97\)90048-5](https://doi.org/10.1016/S0163-6383(97)90048-5). URL <https://www.sciencedirect.com/science/article/pii/S0163638397900485>.
- Erikson, E. H. *Childhood and society*. WW Norton & Company, 1993.
- Fletcher, G. E., Warneken, F., and Tomasello, M. Differences in cognitive processes underlying the collaborative activities of children and chimpanzees. *Cognitive Development*, 27(2):136–153, 2012. ISSN 0885-2014. doi: <https://doi.org/10.1016/j.cogdev.2012.02.003>. URL <https://www.sciencedirect.com/science/article/pii/S0885201412000093>.
- Gergely, G., Bekkering, H., and Király, I. Rational imitation in preverbal infants. *Nature*, 415(6873): 755–755, February 2002. doi: 10.1038/415755a. URL https://ideas.repec.org/a/nat/nature/v415y2002i6873d10.1038_415755a.html.
- Gopnik, A. and Meltzoff, A. N. *Words, thoughts, and theories*. Mit Press, 1997.
- Hare, B., Call, J., and Tomasello, M. Do chimpanzees know what conspecifics know? *Animal behaviour*, 61(1): 139–151, 2001.
- Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B., and Tomasello, M. Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, 317 (5843):1360–1366, 2007. doi: 10.1126/science.1146282. URL <https://www.science.org/doi/abs/10.1126/science.1146282>.
- Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Horgan, D., Quan, J., Sendonaris, A., Osband, I., et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

- Heyes, C. Précis of cognitive gadgets: The cultural evolution of thinking. *Behavioral and Brain Sciences*, 42:e169, 2019. doi: 10.1017/S0140525X18002145.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Huang, W., Abbeel, P., Pathak, D., and Mordatch, I. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *CoRR*, abs/2201.07207, 2022. URL <https://arxiv.org/abs/2201.07207>.
- Hui, D. Y.-T., Chevalier-Boisvert, M., Bahdanau, D., and Bengio, Y. Babyai 1.1, 2020.
- Hutchins, E. *Cognition in the Wild (Bradford Books)*. The MIT Press, September 1996. ISBN 0262581469. URL <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike04-20{%&}path=ASIN/0262581469>.
- Jaques, N., Lazaridou, A., Hughes, E., Gülçehre, Ç., Ortega, P. A., Strouse, D., Leibo, J. Z., and de Freitas, N. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97, pp. 3040–3049. PMLR, 2019. URL <http://proceedings.mlr.press/v97/jaques19a.html>.
- Keupp, S., Behne, T., and Rakoczy, H. Why do children overimitate? normativity is crucial. *Journal of Experimental Child Psychology*, 116(2):392–406, 2013. ISSN 0022-0965. doi: <https://doi.org/10.1016/j.jecp.2013.07.002>. URL <https://www.sciencedirect.com/science/article/pii/S0022096513001379>.
- Kosinski, M. Theory of mind may have spontaneously emerged in large language models. *arXiv preprint arXiv:2302.02083*, 2023.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- Kumar, S., Correa, C. G., Dasgupta, I., Marjeh, R., Hu, M., Hawkins, R. D., Daw, N. D., Cohen, J. D., Narasimhan, K., and Griffiths, T. L. Using natural language and program abstractions to instill human inductive biases in machines. *ArXiv*, abs/2205.11558, 2022.
- Le, M., Boureau, Y.-L., and Nickel, M. Revisiting the evaluation of theory of mind through question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 5872–5877, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1598. URL <https://aclanthology.org/D19-1598>.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Lee, D., Jaques, N., Kew, J. C., Eck, D., Schuurmans, D., and Faust, A. Joint attention for multi-agent coordination and social learning. *CoRR*, abs/2104.07750, 2021. URL <https://arxiv.org/abs/2104.07750>.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *ICLR*, 2016.
- Lindblom, J. and Ziemke, T. Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior*, 11(2):79–96, 2003. doi: 10.1177/10597123030112002. URL <https://doi.org/10.1177/10597123030112002>.
- Lyons, D. E., Young, A. G., and Keil, F. C. The hidden structure of overimitation. *Proceedings of the National Academy of Sciences*, 104:19751 – 19756, 2007.
- Meltzoff, A. N. Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental psychology*, 31 5:838–850, 1995.
- Meltzoff, A. N. and Moore, M. K. Explaining facial imitation: A theoretical model. *Infant and child development*, 6(3-4):179–192, 1997.
- Mirolli, M. and Parisi, D. Towards a vygotskian cognitive robotics: The role of language as a cognitive tool. *New Ideas in Psychology*, 29 (3):298–311, 2011. ISSN 0732-118X. doi: <https://doi.org/10.1016/j.newideapsych.2009.07.001>. URL <https://www.sciencedirect.com/science/article/pii/S0732118X09000348>. Special Issue: Cognitive Robotics and Reevaluation of Piaget Concept of Egocentrism.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.

- Moll, H. and Tomasello, M. Level 1 perspective-taking at 24 months of age. *British Journal of Developmental Psychology*, 24(3):603–613, 2006.
- Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., and Abbeel, P. Overcoming exploration in reinforcement learning with demonstrations. In *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 6292–6299. IEEE, 2018.
- Ndousse, K. K., Eck, D., Levine, S., and Jaques, N. Emergent social learning via multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 7991–8004. PMLR, 2021.
- Netanyahu, A., Shu, T., Katz, B., Barbu, A., and Tenenbaum, J. B. PHASE: physically-grounded abstract social events for machine social perception. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 845–853. AAAI Press, 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16167>.
- Oudeyer, P.-Y. and Kaplan, F. What is intrinsic motivation? a typology of computational approaches. *Frontiers in Neurobotics*, 1, 2007. ISSN 1662-5218. doi: 10.3389/neuro.12.006.2007. URL <https://www.frontiersin.org/articles/10.3389/neuro.12.006.2007>.
- Over, H. and Carpenter, M. The social side of imitation. *Child Development Perspectives*, 7(1): 6–11, 2013. doi: <https://doi.org/10.1111/cdep.12006>. URL <https://srcd.onlinelibrary.wiley.com/doi/abs/10.1111/cdep.12006>.
- Parker-Holder, J., Jiang, M., Dennis, M., Samvelyan, M., Forster, J. N., Grefenstette, E., and Rocktaschel, T. Evolving curricula with regret-based environment design. In *International Conference on Machine Learning*, 2022.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *ICML*, 2017.
- Perez, E., Strub, F., de Vries, H., Dumoulin, V., and Courville, A. C. Film: Visual reasoning with a general conditioning layer. *CoRR*, abs/1709.07871, 2017. URL <http://arxiv.org/abs/1709.07871>.
- Portelas, R., Colas, C., Weng, L., Hofmann, K., and Oudeyer, P. Automatic curriculum learning for deep RL: A short survey. *CoRR*, abs/2003.04664, 2020. URL <https://arxiv.org/abs/2003.04664>.
- Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Es-lami, S. M. A., and Botvinick, M. Machine theory of mind. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4215–4224. PMLR, 2018. URL <http://proceedings.mlr.press/v80/rabinowitz18a.html>.
- Raileanu, R. and Rocktäschel, T. RIDE: rewarding impact-driven exploration for procedurally-generated environments. *CoRR*, abs/2002.12292, 2020. URL <https://arxiv.org/abs/2002.12292>.
- Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Novikov, A., Barth-Maron, G., Gimenez, M., Sulsky, Y., Kay, J., Springenberg, J. T., Eccles, T., Bruce, J., Razavi, A., Edwards, A. D., Heess, N. M. O., Chen, Y., Hadsell, R., Vinyals, O., Bordbar, M., and de Freitas, N. A generalist agent. *ArXiv*, abs/2205.06175, 2022.
- Richerson, P. J. and Boyd, R. *Not by Genes Alone: How Culture Transformed Human Evolution*. University Of Chicago Press, June 2006. ISBN 0226712125. URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0226712125>.
- Ruis, L., Khan, A., Biderman, S. R., Hooker, S., Rocktaschel, T., and Grefenstette, E. Large language models are not zero-shot communicators. *ArXiv*, abs/2210.14986, 2022.
- Sap, M., Rashkin, H., Chen, D., Bras, R. L., and Choi, Y. Socialiqa: Commonsense reasoning about social interactions. *CoRR*, abs/1904.09728, 2019. URL <http://arxiv.org/abs/1904.09728>.
- Sap, M., Bras, R. L., Fried, D., and Choi, Y. Neural theory-of-mind? on the limits of social intelligence in large lms. *ArXiv*, abs/2210.13312, 2022.
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T. P., and Gelly, S. Episodic curiosity through reachability. *ArXiv*, abs/1810.02274, 2018. URL <http://arxiv.org/abs/1810.02274>.
- Scao, T. L., Fan, A., Akiki, C., Pavlick, E.-J., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A. S., Yvon, F., Gallé, M., Tow, J., Rush, A. M., Biderman, S. R., Webson, A., Ammanamanchi, P. S., Wang, T., Sagot, B., ..., and Wolf, T. Bloom: A 176b-parameter open-access multilingual language model. *ArXiv*, abs/2211.05100, 2022.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms.

- ArXiv*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- Shu, T., Kryven, M., Ullman, T. D., and Tenenbaum, J. Adventures in flatland: Perceiving social interactions under physical dynamics. In Denison, S., Mack, M., Xu, Y., and Armstrong, B. C. (eds.), *Proceedings of the 42th Annual Meeting of the Cognitive Science Society - Developing a Mind: Learning in Humans, Animals, and Machines, CogSci 2020, virtual, July 29 - August 1, 2020*. cognitivesciencesociety.org, 2020. URL <https://cogsci.mindmodeling.org/2020/papers/0723/index.html>.
- Shu, T., Bhandwaldar, A., Gan, C., Smith, K. A., Liu, S., Gutfreund, D., Spelke, E. S., Tenenbaum, J. B., and Ullman, T. D. AGENT: A benchmark for core psychological reasoning. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9614–9625. PMLR, 2021. URL <http://proceedings.mlr.press/v139/shu21a.html>.
- Siposova, B. and Carpenter, M. A new look at joint attention and common knowledge. *Cognition*, 189:260–274, 2019.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., Turck, F. D., and Abbeel, P. Exploration: A study of count-based exploration for deep reinforcement learning, 2017.
- Tennie, C., Call, J., and Tomasello, M. Push or pull: Imitation vs. emulation in great apes and human children. *Ethology*, 112(12):1159–1169, 2006.
- Tennie, C., Walter, V., Gampe, A., Carpenter, M., and Tomasello, M. Limitations to the cultural ratchet effect in young children. *Journal of experimental child psychology*, 126:152–160, 2014.
- Tomasello, M. *The Cultural Origins of Human Cognition*. Harvard University Press, 1999. ISBN 9780674000704. URL <http://www.jstor.org/stable/j.ctvjsf4jc>.
- Tomasello, M. Becoming human. In *Becoming Human*. Harvard University Press, 2019.
- Tomasello, M. The role of roles in uniquely human cognition and sociality. *Journal for the Theory of Social Behaviour*, 50(1):2–19, 2020. doi: 10.1111/jtsb.12223.
- Tomasello, M., Kruger, A. C., and Ratner, H. H. Cultural learning. *Behavioral and brain sciences*, 16(3):495–511, 1993.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Trott, S., Jones, C. J., Chang, T. A., Michaelov, J. A., and Bergen, B. K. Do large language models know what humans know? *ArXiv*, abs/2209.01515, 2022.
- Vollmer, A.-L., Wrede, B., Rohlfing, K. J., and Oudeyer, P.-Y. Pragmatic frames for teaching and learning in human–robot interaction: Review and challenges. *Frontiers in Neurobotics*, 10:10, 2016. ISSN 1662-5218. doi: 10.3389/fnbot.2016.00010. URL <https://www.frontiersin.org/article/10.3389/fnbot.2016.00010>.
- Vygotsky, L. S. and Cole, M. *Mind in society : the development of higher psychological processes*. Harvard University Press Cambridge, 1978. ISBN 0674576284 0674576292.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E. H., Le, Q., and Zhou, D. Chain of thought prompting elicits reasoning in large language models. *CoRR*, abs/2201.11903, 2022. URL <https://arxiv.org/abs/2201.11903>.
- Whiten, A., Horner, V., Litchfield, C. A., and Marshall-Pescini, S. How do apes ape? *Animal Learning & Behavior*, 32:36–52, 2004.
- Whiten, A., McGuigan, N., Marshall-Pescini, S., and Hopper, L. M. Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528):2417–2428, 2009.
- Wood, D., Bruner, J. S., and Ross, G. The role of tutoring in problem solving. *Journal of child psychology and psychiatry, and allied disciplines*, 17(2):89–100, 1976. doi: 10.1111/j.1469-7610.1976.tb00381.x.
- Wood, D., Bornstein, M., and Bruner, J. Interaction in human development, 1989.
- Wyman, E., Rakoczy, H., and Tomasello, M. Normativity and context in young children’s pretend play. *Cognitive development*, 24(2):146–155, 2009.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. React: Synergizing reasoning and acting in language models. *ArXiv*, abs/2210.03629, 2022.
- Zhang, T., Xu, H., Wang, X., Wu, Y., Keutzer, K., Gonzalez, J. E., and Tian, Y. Bebold: Exploration beyond the boundary of explored regions. *CoRR*, abs/2012.08621, 2020. URL <https://arxiv.org/abs/2012.08621>.

Zhang, Z., Zhang, A., Li, M., Zhao, H., Karypis, G., and Smola, A. J. Multimodal chain-of-thought reasoning in language models. *ArXiv*, abs/2302.00923, 2023.

A. Social agents are not objects.

Although social peers could be seen as merely complex interactive objects, we argue they are in essence quite different. Social agents (e.g. humans) can have very complex and changing internal states, including intents, moods, knowledge states, preferences, emotions, etc. The resulting set of possible interactions with peers (social affordances) is essentially different from those with objects (classical affordances). In cognitive science, an affordance refers to what things or events in the environment afford to an organism (Carvalho, 2020). A flat surface can afford "walking-on" to an agent, while a peer can afford "getting help from". The latter is a social affordance, which may require a social system and conventions (e.g. politeness), implying that social peers have complex internal states and the ability to reciprocate. Successful interaction might also be conditioned on the peer's mood, requiring communication adjustments. Training an agent for such social interactions most likely requires drastically different methods – e.g. different architectural biases – than classical object-manipulation training. In SocialAI we simulate such social peers using scripted peers. We argue that studying isolated social scenarios featuring scripted peers in tractable environments is a promising first step towards designing proficient social agents.

B. Extended Cognitive Science Background

The following section introduces additional detail to section 3.

B.1. Michael Tomasello - The Shared Intentionality Theory

Discussion on the uniquely human socio-cognitive abilities Michael Tomasello (2019) argues that the 9-month-revolution and the objective/normative turn are uniquely *human* developmental steps enabling uniquely *human* socio-cognitive abilities. There has been a lot of debate regarding this hypothesis (De Waal, 2016), and it still remains an open question. However, for the purpose of this work, the social proficiency of other great apes (or our last common ancestor with them) is not of primary importance. We find The Shared Intentionality Theory useful because it is systematic, extensive (covers a broad range of social abilities), and exact (is build upon a number of very clearly defined experiments). Furthermore, it is concerned with the questions regarding the development of core socio-cognitive abilities. We believe that this makes it a good basis to organize AI research on.

B.1.1. SOCIAL COGNITION

Imagining what others perceive In this paragraph, we discuss the ability to infer what another sees or knows. As compared to the later emerging ability to coordinate perspectives, this ability requires that only one perspective is processed at a time. Numerous studies have shown that both apes and children are capable of making such inferences. For example, in one experiment (Hare et al., 2001), a subordinate and a dominant chimpanzee were presented with a competitive scenario : competing for food. Results showed that the subordinate chimpanzee attempted to eat the food only if it was hidden from the dominant one. This experiment was then extended to children who were presented with two toys: one observed by an adult and one occluded from him (Moll & Tomasello, 2006). When asked to help the adult find a toy, 24-month-olds passed the occluded toy. These experiments, demonstrate that both children and apes are capable of inferring what a conspecific observes - i.e. they are able to infer another's perspective.

Joint Attention Joint attention is often defined in different ways (Siposova & Carpenter, 2019). To avoid confusion, we take the definition of joint attention from Michael Tomasello (Tomasello, 2019): joint attention consists of two aspects: *triangulation* and *recursiveness*. Triangulation refers to the child and the adult attending to the same external referent, and recursiveness refers to them both recursively being aware that they are both sharing attention. Joint attention enables children to process multiple perspectives at the same time, which can then be aligned and exchanged.

In a series of studies (Carpenter et al., 1998b), the amount of joint attention (number of joint attention episodes and their length) was measured in free play interactions between infants and their mothers. A steady rise in the amount of time spent in joint attention was observed in the period from 9 to 12 months.

Coordinating perspectives When sufficient linguistic competence is reached, children start to jointly attend to mental content through linguistic discourse. Through linguistic discourse children often encounter conflicting perspectives, which they are then pushed to resolve (e.g. one parent says it's raining outside, but another says it's not). They resolve those conflicts by learning to form an "objective" perspective - a bird's-eye-view perspective distinct from anyone's personal perspective - and coordinating the conflicting perspectives with it. (Tomasello, 2019) argues that this can only be achieved




Figure 4. Sketch of an experiment showing that apes can infer the conspecific's field of view (Hare et al., 2001). As the subordinate ape does not want to get into trouble, it will not steal the food from the dominant ape. In the experiment, the food was either occluded from the dominant ape or placed in plain sight. The subordinate ape ate the food only when it was occluded from the dominant ape. This shows that it was able to infer the dominant's field of view.

once a child has passed through the second developmental step, that of collective intentionality, which enables them to form such a "perspectiveless" bird's-eye view perspective.

B.1.2. COMMUNICATION




Figure 5. An experiment studying children's ability to infer the meaning of a pointing gesture (Behne et al., 2005). The child's attention is drawn to a toy. This toy is then hidden in one of the two boxes (the child does not know which one). The experimenter then points to one of the two boxes, and the child is able to infer this to mean that the toy is in that box.

Referential communication - The Pointing gesture Following the 9-month revolution, children start to communicate *referentially* - to an external referent (Tomasello, 2019). This is primarily achieved through *pointing* and pantomiming.

Figure 5 depicts an experiment studying the children's understanding of the pointing gesture (Behne et al., 2005). First, the child's attention is drawn to the toy, which is then hidden in one of the two boxes. The experimenter then points to a box, and the child infers this to mean that the toy is in that box. 14-month-old children were able to successfully follow a pointing gesture to find the toy. In this scenario, the child makes the following recursive inference: the adult is helping by directing the child's attention to the box, and she wants the child to infer that the toy is in the box.

Linguistic communication Linguistic communication is based on the same principle as gestural referential communication: sharing attention to a referent and recursively inferring the intended meaning. However, linguistic communication in addition requires learning conventionalized means of reference, such as words or phrases. Where once was a single pointing gesture,

now there is a complex grammar of gestures, with specific conventions assigned to each gesture. In a series of studies, Malinda Carpenter (1998b) observed children’s understanding of words steadily increase in the period after 9 months. This was measured by questioners given to their caretakers on regular intervals.

Michael Tomasello argues that when language use first appears, children do not yet understand it as conventional, rather they use it as any other artifact or tool. It is only after the emergence of collective intentionality, when children start to understand and use conventions and norms, that they also begin to perceive language as such. This is evidenced by specific new ways in which they come to use and understand language. For example, when others break the rules of a game they protest by normative statements such as "No! It does not go like this!" (Wyman et al., 2009). It is needless to say that language plays many important roles in children’s development. Here we will outline just a few of countless possible examples. Language provides children with abstract constructions which gives them a new organized format for cognitive representation. Through discourse, children encounter many conflicting perspectives, which brings them to resolve those conflicts by forming the "objective" perspective. Finally, language opens up a new way of cultural learning - instructed learning - in which adults directly teach children "objective" truths about the world. Knowledge learned in that manner is easier to generalize (Butler & Tomasello, 2016).

B.1.3. CULTURAL LEARNING

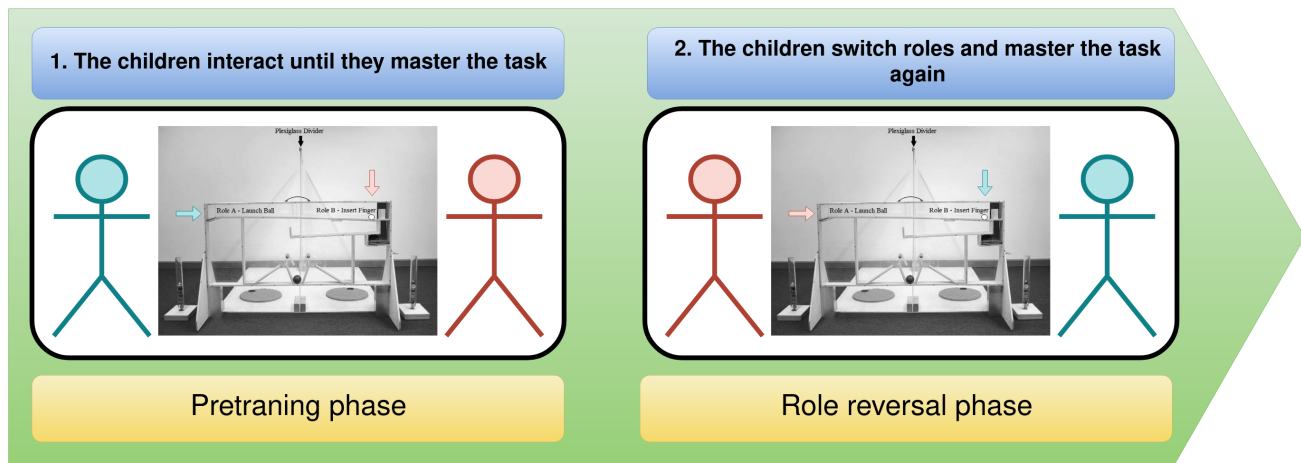


Figure 6. Depiction of an experiment on role reversal (Fletcher et al., 2012). The task consists of two roles: one participant pushes a ball into the apparatus, and the other redirects it with their finger. The ball then pushes two marbles toward each of the participants. In the pretraining phase, children collaborate until they master the task (three consecutive successful trials). Then, in the role reversal phase, their roles are reversed and they master the task again. Total number of trials to master the task is compared between the two phases. Children, but not apes, needed less trials to master the task in the role reversal phase than in the pretraining phase.




Figure 7. Depiction of an experiment(Carpenter et al., 1998b). The experimenter activates the party favor (sound) by pushing the spring, and the child imitates and waits for the sound. The sketch was taken and modified from Carpenter (1998b).

Imitation, Emulation, and Overimitation Imitation and emulation learning both refer to observing a demonstration and learning from it. Imitation learning refers to the learning of means (actions), while emulation to the learning of ends (goals) of a demonstration (Whiten et al., 2004; 2009; Tennie et al., 2006). Refer to Whiten (2009) for a discussion and taxonomy of imitative and emulative learning processes.

Figure 7 shows an experiment studying children’s imitation abilities (Carpenter et al., 1998b). In this experiment, the experimenter demonstrates an instrumental action (e.g. pressing a spring attached to a box) which activates the light on top of the box. The children repeated the instrumental action and looked expectedly at the light. This kind of learning emerges over the course of the first year - children reconstruct the outcome of others’ actions. However, soon after this, children begin imitating in a way which demonstrates the understanding of other’s goals. Children perform an action that an adult attempted, but failed to perform (Meltzoff, 1995), and do not imitate accidental actions (Carpenter et al., 1998a). Similarly, rational imitation appears. If the action was forced upon the demonstrator, the children recreate the result through more rational means (Gergely et al., 2002). For example, in an experiment a demonstrator pressed a button with its head while having tied hands and 14-month-olds responded by pressing the button with their hand (Gergely et al., 2002).

Emulation is a type of social learning where the focus is on the outcome, and not on the actions performed (Wood et al., 1989). In other words, the learning is about some property of the environment rather than the social interaction. For example, by observing someone turn on the lights we learn that lights can be turned on by flicking a light switch. The learner tries to recreate some observed outcome, in doing so they can, but do not have to, recreate the actions.

On the other side of this spectrum is *overimitation* - children repeat actions that are not relevant for the outcome. Children often prefer to not only recreate the outcome (as in emulation), but also do it in the same way as the adults (even if this requires doing additional unnecessary actions). For example, in one experiment (Tennie et al., 2014) children were presented with a demonstration of a rice-pouring task. The experimenter performed a useless preliminary action before grabbing the rice. 4-year-old children responded by repeating both the useless and the necessary actions. It has been proposed that children overimitate to affiliate and conform for the purpose of in group bonding (Over & Carpenter, 2013), but this remains an open question (Keupp et al., 2013; Lyons et al., 2007)

Role reversal Imitation Role reversal imitation appears following the 9-month revolution. Figure 6 depicts an experiment with children and apes from Fletcher (2012). An apparatus is used where one participant pushes the marble, and the other inserts a finger to redirect the ball so that it falls to the correct location. Then, both participants get a reward. Children who previously played role A mastered role B in less trials than children who never played role A. In another experiment (Carpenter et al., 2005), children were asked to immediately reverse the role. An experimenter did some action on the child (e.g. poke the child and say “your turn”) and the child responded with the same action on the experimenter (poke the experimenter back). These experiments show that children understand the separate roles and how each is relevant for the activity.

C. Technical details on the SocialAI school

In this section, we discuss in detail the parameterized environments and their procedural generation. Then, we discuss the architecture of the RL agent and the exploration bonuses used in our experiments.

C.1. Parameterized Social Environments

The SocialAI school is built on top of the MiniGrid codebase (Chevalier-Boisvert et al., 2018), which provides an efficient and easily extensible implementation of grid world environments. SocialAI environments are grid worlds consisting of a room. In all of our environments, the task of the agent is to eat the apple, at which point it is rewarded. The reward is diminished according to the number of steps it took the agent to complete the episode. The episode ends when the agent eats the apple, uses the *done* action, or after a timeout of 80 steps.

The agent interacts with the environment as shown in figure 8. This multimodal observation space consists of the full dialogue history, and a 7x7x8 tensor corresponding to the 7x7 grid in front of the agent. Each cell is encoded by six integers representing the object type, color, and some additional object-dependent information (e.g. is the door open, point direction, gaze direction, etc.). Refer to figure 22 for a list of all objects.

The agent acts in the environment through a multimodal action space, which consists of 6 primitive actions (*no*, movement actions, *toggle*, and *done*) and a 4x16 templated language. The agent also has the option not to speak, which is implemented

with an additional binary output from the agent. Refer to section C.4 for further details about the architecture of the agent.

All environments contain a scripted social peer, and the task can only be solved by interacting with this peer (for which socio-cognitive abilities are needed). A social peer observes the world in the same way as the agent does (as a grid in front of it), and it also observes the agent’s utterances. Their action space consists of primitive actions for movement, pointing, and the *toggle* action. The peer can also communicate with words and sentences. As the peer is scripted, there are no constraints on the language it can utter (it is not constrained to a templated language). The language it uses depends on the environment, which defines which sentence the peer will utter at which point. The peer is represented in the agent’s observation by 7 integers depicting their: object type, position, color, type (cooperative or competitive), gaze direction, point direction, and the last executed primitive action. The peer’s gaze and point directions are represented relatively the agent (e.g. 1 - to the left of the agent). The pointing direction can also be set to 0, which signifies that the peer is not pointing. Figure 9 shows an example of an environment with the corresponding encoding of the peer. The agent (red) and the scripted peer (purple) are making eye contact - the peer and the agent are in the same row or column and their gazes meet frontally. In this example, the scripted peer is also pointing to the blue box.

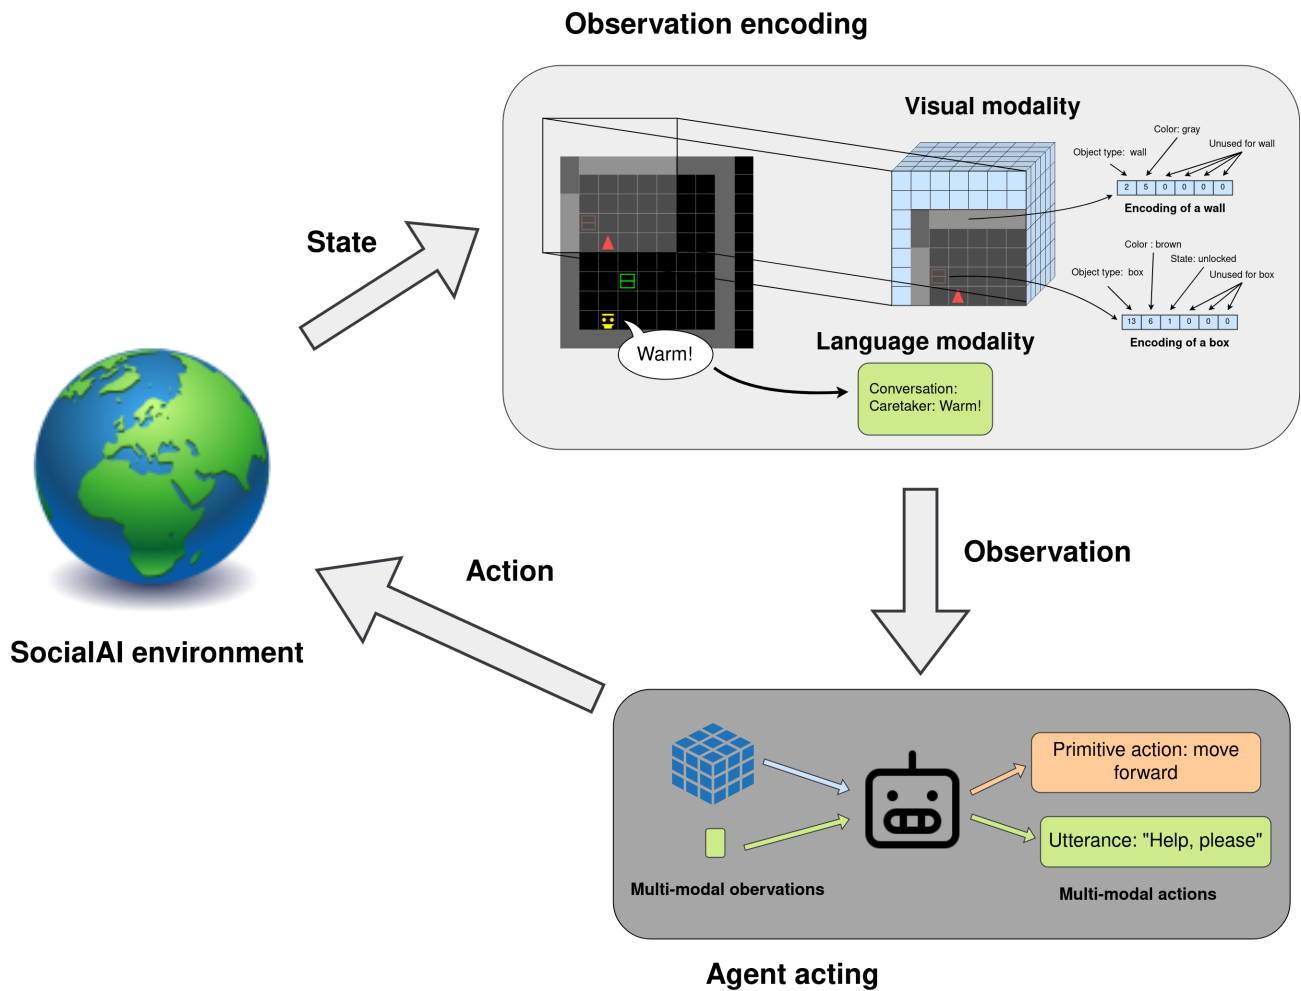


Figure 8. Workflow of an agent acting in the SocialAI school. The environment generates a state, which is represented as multimodal observations: a 7x7x6 tensor and the full dialogue history. The agent acts through a multi-modal action space consisting of primitive actions and utterances.

The SocialAI environments are parameterized, and those parameters define the social dimensions of the task. In other words, parameters define which socio-cognitive abilities are needed to solve the task. For example, depending on the ENVIRONMENT




Figure 9. A depiction of a peer and its encoding. The agent and a peer are in eye contact, and the peer is pointing to the blue box. To the right is an encoding of the peer. The encoding contains information about the peer, e.g. the gaze and point direction.

TYPE parameter, the peer can give information, collaborate with the agent, or be adversarial. Then, in the case of the peer giving information, additional parameters define what is the form of this information (linguistic or pointing).

C.2. The parameter tree

SocialAI enables the creation of many parameterized environments, and those parameters are implemented as nodes in a parameter tree. A parameter tree is a structure through which the experimenter can easily define which parameters (and their values) can be sampled. An example of such a tree can be seen in figure 10. The standard procedure is that an experimenter defines such a parameter tree and each episode begins with the sampling of a new parameter set from this tree. Once a parameter set has been sampled, an environment is created, and the agent placed inside.

A parameter tree is used to sample parameter sets from it, three examples of such sampling are shown in figure 10. There are two kinds of nodes: parameter nodes (rectangles) and value nodes (ovals). Parameter nodes correspond to parameters, and value nodes corresponds to possible values for those parameters. Sampling proceeds in a top-down fashion, starting from the root node. In all our experiments, the root node is the ENV_TYPE parameter node. Sampling from a parameter node selects one of its children (a value node), i.e. sets a value for this parameter. This can be done by uniform sampling over the node’s children, or by prioritized sampling with a curriculum. Once a value node has been chosen, the sampling continues through it to all of its children (parameter nodes). In other words, setting a value for one parameter, defines which other parameters (the value node’s children) need to be set. In our codebase, it is simple to create such trees, and to add additional parameters and environments. In the remained of this section, we explain some of the most relevant parameters.

C.3. Environment types

The most important parameter in the parameter tree is the environment type (ENV_TYPE). This parameter node is the root node in all our experiments. We implemented three different values: INFORMATIONSEEKING, COLLABORATION, and ADVERSARIALPEER. A parameter tree doesn’t have to contain all of those values, this choice entirely depends on the type of experiment one wants to conduct, most often only one of type will be present in a tree. For example, figure 10 shows the tree with only the INFORMATIONSEEKING environment type. This tree was used to study understanding of the pointing gesture in section 5.1.

Information Seeking type

We used this environment type in experiments studying communication, joint attention, and imitation learning. In figure 11 we can see examples of INFORMATIONSEEKING type environments.

The general principle of this environment type is as follows. The agent is rewarded upon eating the apple, which is hidden at the beginning of the episode and can be accessed through interacting with an object.

The PROBLEM parameter defines which objects will be present in the environment, and different objects make the apple

accessible in different ways (ex. pushing, toggling, ...). We implement six different problems: BOXES, SWITCHES, MARBLE, GENERATORS, DOORS, OR LEVERS. For example, opening the box will make the apple appear at the location of the box, while pulling the lever will open the door in front of the apple.

The `N` parameter defines the presence of a distractor (present if set to 2) A distractor is an object of the same type as the correct object, but if used both objects are blocked and the apple cannot be obtained in this episode.

To find out which object is the correct one, the agent must interact with the scripted peer. This interaction starts with the agent introducing itself. The way in which the agent should introduce itself is defined by the `INTRODUCTORY_SEQUENCE` parameter. We define the following four values: `NO`, `EYE_CONTACT`, `ASK`, `ASK-EYE_CONTACT`. For the value `NO`, no introduction is needed. The peer will give information at the beginning of the episode. In most of our experiments, we will use the value `EYE_CONTACT` where the peer turns to look at the agent and waits for the agent to look back (to establish eye contact) - the agent must direct its gaze directly towards the scripted peer. An example of an established eye contact can be seen in figure 9. For the value `ASK`, the agent needs to utter "Help, please". The agent does so using templated language, by selecting the "Help, X" template and the word "please". A full grammar of the language is given in table 2 in the appendix. Finally, the `ASK-EYE_CONTACT` value is a combination of the previous two. It requires that the agent utters "Help, please" during eye contact.

Once the agent introduces itself, the `HELP` parameter defines the peer's behavior. If it is set to `Y` the peer will obtain the apple, and leave it for the agent to eat. Alternatively, it will give cues to the agent about which object to use. The nature of this cue is defined by the `CUE_TYPE` parameter. We define four different values: `POINTING`, `LANGUAGE_COLOR`, `LANGUAGE_FEEDBACK`, and `IMITATION`. For the `POINTING` type, the peer will point to the correct object. It will move to a location from which it can unambiguously point (e.g. the same row) and point to the object (see figure 11(a)). For the `LANGUAGE_COLOR` type, the peer will say the color of the correct object (see figure 11(b)). For the `LANGUAGE_FEEDBACK` type, the peer will hint to the proximity of the agent to the correct object. Every step, the peer will utter "Cold", "Medium", "Warm" or "Hot", depending on how close the agent is to the correct object. For example, "Cold" means that the agent is far from the object, and "Hot" that it is right next to it (see figure 11(c)). For the `IMITATION` type, the peer will demonstrate the use of the correct object. The peer will use the correct object, obtain the apple, and eat it, and then reset the environment to its initial state.

For the purpose of analyzing the agent's behavior in more detail, information seeking environments can also be created without the peer (by setting the `PEER` parameter to `N`). The asocial version refers to an environment contains no distractor and no peer, i.e. the agent just needs to use the object in the environment.

Collaboration type

We used this environment type to study the ability of the agent to reverse the roles of a collaborative activity. Environments are separated into two halves (corresponding to different roles) by a fence over which the agent can see, but which it cannot cross (see figure 12). If both roles are fulfilled correctly, two apples will become accessible (one on each side of the fence).

The most important parameters are `ROLE` and `PROBLEM`. The parameter `ROLE` designates the role to be assigned to the agent. The `PROBLEM` parameter defines the collaborative activity, of which we implemented seven: `DOORLEVER`, `MARBLEPUSH`, `MARBLEPASS`, `BOXES`, `SWITCHES`, `GENERATORS`, `MARBLE`. In `DOORLEVER` one participant opens the door by pulling the lever and the other passes through them, and activates the generator (generating two apples). In `MARBLEPUSH` one participant opens the door by pulling the lever, and the other pushes a marble through them. This marble activates the *marble generator* upon contact with it. In `MARBLEPASS` one participant pushed the marble to the right side of the room, and then the other pushes it towards the *marble generator*. In the remaining four problems, one participant is presented with two boxes of different colors. The other participant is presented with two objects of the same colors as the two boxes and of the type defined by the `PROBLEM` parameter (e.g. two generators). First, the participant that was presented with boxes open one box (an apple will be in both). After this, to obtain its apple, the other participant must use the object of the same color as the opened box. In figure 12 we can see examples of `COLLABORATION` type environments.

Like the information seeking environments, collaboration environments can also be instantiated in their asocial versions. The peer is not present in the environment, and the environment is initialized so that the task can be solved alone. For example, in `MARBLEPASS` the marble is already on the right side of the room, so the agent just has to push it towards the *marble generator*.

Adversarial environment type This environment type is used to study the ability of the agent to infer the peer's field of view. An apple will be present in the environment right away. However, the agent will get rewarded only if it eats it while

not being observed by the peer (the peer is adversarial). Therefore, the agent needs to infer the right moment to eat the apple. There is one important parameter in this environment type. It refers to the amount of obstacles present in the environment. Figure 19 shows this environment type without any obstacles (figure 19(a)) and with obstacles present (figure 19(b)).

C.4. Architecture of the RL agent

The multimodal observation space consists of a 7x7x6 tensor (vision) and the full dialogue history (language). The multimodal action space consists of 6 primitive actions (no_op, movement actions, *toggle*, and *done*), and a 4x16 templated language.

In this work we use a PPO-trained (Schulman et al., 2017) DRL architecture initially designed for the BabyAI benchmark (Chevalier-Boisvert et al., 2019). The policy design was improved in a follow-up paper (Hui et al., 2020) (more precisely, we extend their *original_endpool_res* model). See figure 13 for a visualization of the complete architecture. First, symbolic pixel grid observations are fed into two convolutional layers (LeCun et al., 1989; Krizhevsky et al., 2012) (3x3 filter, stride and padding set to 1), while dialogue inputs are processed using a Gated Recurrent Unit layer (Chung et al., 2015). The resulting image and language embeddings are combined using two FiLM attention layers (Perez et al., 2017). Max pooling is performed on the resulting combined embedding before being fed into an LSTM (Hochreiter & Schmidhuber, 1997) with a 128D memory vector. The LSTM embedding is then used as input for the navigation action head, which is a two-layered fully-connected network with tanh activations and has an 6D output (i.e. 5 navigation actions and no_op action).

In order for our agent to be able to both move and talk, we add to this architecture a talking action head, which is composed of three subheads. All of them consist of two fully-connected layers with tanh activations, and take the LSTM’s embedding as input. The first one is used as a switch: it has a one-dimensional output to choose whether the agent talks (output >0.5) or not (output <0.5). If the agent talks, the two other networks are used to sample the template and the word. Grammar of the templated language is depicted in table 2 and examples of multi-modal actions in table 3.

C.5. Exploration bonuses

The exploration bonuses we use are inspired by recent works in intrinsically motivated exploration (Pathak et al., 2017; Savinov et al., 2018; Tang et al., 2017). These intrinsic rewards estimate the novelty of the currently observed state and add the novelty based bonus to the extrinsic reward.

In this work we present two techniques for computing the count-based exploration bonus. Both of our count-based exploration bonuses are episodic - they estimate the diversity of states observed within an episode, and assume that beneficial episodes are those with more diverse observations. We compare those two exploration bonuses to RIDE (Raileanu & Rocktäschel, 2020), RND (Burda et al., 2018), and to the agent without any exploration bonus.¹

Language-based exploration bonus (CBL) For some utterance s_{lang} observed at state s , we count how many times was this utterance observed during the episode. We compute the bonus for this step using the following equation:

$$r_{intr} = T * \tanh \left(\frac{C}{(N(s_{lang}) + 1)^M} \right) \quad (1)$$

, where M , C , and T are hyperparameters and $N(s_{lang})$ is the number of times the utterance s_{lang} was observed during this episode so far.

Vision-based intrinsic reward (CB) We reward the agent for observing diverse encodings. An encoding is the 6D representation of a cell (see figure 8 for more details). A visual observation consists of 47 (7x7) encodings representing cells in front of the agent. For some visual observation s_{viz} at step s , a set of encountered unique encodings is created (duplicates are removed) $U(s_{viz})$, and then the reward computed using the following equation:

$$r_{intr} = T * \tanh \left(\sum_{e \in U(s_{viz})} \frac{C}{(N(e) + 1)^M} \right) \quad (2)$$

¹We verify our implementation of RIDE and RND by recreating the results of those baselines on environments from RIDE (Raileanu & Rocktäschel, 2020).

, where M , C , and T are hyperparameters, $U(s)$ is a set of unique encodings visible in state s , and $N(e)$ is the number of times an encoding e was encountered in the current episode.

D. Detail on the baselines and exploration bonuses used

In all of our case studies, except the study with language models (sec. 5.4), we use a PPO (Schulman et al., 2017) reinforcement learning agent (as described in section C.4) with different exploration bonuses.

In a set of pilot experiments (appendix E), we compared various PPO agents (Schulman et al., 2017) with different exploration bonuses (for details see section C.5), and outlined the most promising ones to use in the following case studies. Visual count-based exploration bonus ("PPO-CB") performed best on the tasks in which language is not used, and its linguistic variant "PPO-CBL" performed best in environments with the peer giving linguistic cues. In our experiments, these exploration bonuses outperformed RND (Burda et al., 2018) and RIDE (Raileanu & Rocktäschel, 2020). Both of those two exploration bonuses are episodic. They estimate the diversity of observations in an episode and give reward proportional to that diversity. The linguistic exploration bonus uses the number of different words, and the vision-based exploration bonus the number of different encodings observed (refer to C.5 for further details).

In the case studies in sections 5.1 and 5.2 we use the "PPO-CB" exploration bonus. The case study in section 5.3 requires raw PPO for the purposes of the study, and one in section 5.4 uses LLMs as agents. In additional experiments (section G), we use PPO-CB and PPO-CBL (in those case-studies in which the peer provides linguistic feedback).

E. Pilot experiments

In our pilot studies, we encoded the peer in a way which used the mix of egocentric and allocentric vision - the peer's gaze and pointing direction were encoded in terms of absolute direction ("NSEW"). We decided to change this to fully egocentric as we found it more natural with regards to the question of socio-cognitive artificial intelligence. We believe that the best performing baselines would also perform best with purely egocentric encodings (the one we use in the rest of the paper). For that reason, and to avoid unnecessary energy spending, we do not compare with other baselines on the purely egocentric encoding.

Figure 14 compares PPO agents trained with different exploration bonuses discussed in section C.5 on two different INFORMATIONSEEKING type environments. The first environment involves the peer pointing to the correct object. Figure 14(b) shows that the best performing agent is the one leveraging the visual count-based exploration bonus (PPO_CB). The second environment involves the peer uttering the color of the correct object. Figure 14(a) shows that the best performing agent is the one leveraging the linguistic count-based exploration bonus (PPO_CBL). We conclude that PPO_CBL is the most suitable baseline for environments involving linguistic cues, and PPO_CB for the other environments.

F. Details on the environment generation parameters used in the main experiments

F.1. The Pointing experiments

The parameter trees used in this experiment are depicted in figure 27. We used the INFORMATIONSEEKING environment type described in section C.3. The INTRODUCTORY_SEQUENCE is set to EYE_CONTACT, and the CUE_TYPE to POINTING - the peer will point to the correct object after eye contact. The agent is trained on the following five problems: BOXES, SWITCHES, LEVERS, MARBLE, GENERATORS, and on the asocial version of the DOORS problem (a version without the distractor or peer). Training on this asocial version is important as it enables the agent to learn how to use a door, which is needed to evaluate generalization.

We use the PPO agent with visual count-based exploration bonus (PPO-CB).

F.2. Role reversal imitation experiments

The parameter trees used in this experiment are depicted in figure 28. We used the COLLABORATION type environments described in section C.3. We evaluate agents on role A of the MARBLEPASS task - the agent has to push the marble to the right side of the environment, from where the peer can push it to the *marble generator*.

We train the agent with the visual count-based exploration bonus (PPO_CB)

F.3. Scaffolding experiments

The parameter trees used in this experiment are depicted in figure 29. In this experiment, we use the INFORMATION SEEKING environment type with the LANGUAGE FEEDBACK cue type. We train agents on all six problems, using different values of the INTRODUCTORY_SEQUENCE and HELP parameters. We evaluate the agents on all six problems, with the most complex introductory sequence - ASK_EYE_CONTACT.

The agent denoted by "scaf_4" is trained on four different values of the INTRODUCTORY_SEQUENCE parameter, and with the HELP parameter set to N (the peer will provide cues). This agent will be trained on a total of 18 different environments: six problems, and four introductory sequences. The second agent (denoted by "scaf_8") is also trained on all values of the INTRODUCTORY_SEQUENCE parameter, but it is in addition trained on both values of the HELP parameter (N and Y) - a total of 36 environments. In half of those environments (with HELP set to Y) the peer will provide the apple to the agent after the introduction (e.g. it will go to the correct box, and open it). In the other half (with HELP set to N), the peer will only provide linguistic feedback cues.

In this experiment, we use the PPO agent without an exploration bonus.

G. Additional case studies

G.1. Inferring the meaning of linguistic cues

In this section, we study the ability of the agent to infer the meaning of simple words. We follow the same procedure as in section 5.1. This case study is motivated by the experiments from cognitive science discussed in section 3.1.2. In (Carpenter et al., 1998b) infants' word understanding steadily increased in the period between 9 and 15 months after birth. We study the following questions:

- Can an RL agent learn to interpret simple utterances?
- Can the agent generalize to new situations, and infer the meaning of those utterances for objects in a new context?

The best performing agent on the linguistic environments in the pilot study was the one using the linguistic count-based exploration bonus (PPO-CBL) (see appendix E). We use this agent to address both questions.

Environments The environments are the same as those in section 5.1: the INFORMATION_SEEKING environment type, with the INTRODUCTORY_SEQUENCE set to EYE_CONTACT. The only difference is that the peer will give linguistic cues instead of pointing. We run two experiments with two different types of linguistic cues: *Color* and *Feedback*. In *Color* the peer will utter name color of the correct object. In *Feedback* the peer will utter a description of how close the agent is to the correct object: "Cold", "Medium", "Warm", and "Hot" meaning, respectively, "far", "medium", "close" and "right next to". The experimental procedure is the same as the one in section 5.1. The agent is trained on the same five problems and the asocial version of the DOORS problem.

Can RL agents learn to interpret simple utterances?

Figures 15(a) and 15(b) show the performance of the agent with the linguistic count-based exploration bonus (denoted PPO_CBL_train). We can see that the agent (PPO-CBL) solves these environments efficiently, reaching a final performance of 95.9% and 71.4% for COLOR and FEEDBACK cue types, respectively. We further analyse the performance of each separate seed for the agent trained on the FEEDBACK cue type. This is shown in figure 16 where it is visible that the agent is normally able to achieve high performance, but that there are two seeds which, due to their instability, reach a success rate of 0. This experiment shows that the agent is capable of learning to infer the meaning of simple utterances in familiar contexts.

Can the agent generalize to new situations? A more interesting question is whether that agent can infer the meaning of the same word based on a new context. Therefore, we evaluate the agent's generalization abilities in a new scenario - the DOOR problem - following the same procedure as in section 5.1.

This kind of generalization is particularly interesting as communication depends on our ability to ground words in *new* social contexts: inferring meaning by combining the convention associated to a word with the recursively inferred intention of the speaker. For example, while "red" can mean "open the red box" in one context, it can mean "push the marble towards the red generator" in another.

Figures 15(b) and 15(a) show the performance of the same agent evaluated on the DOORS problem (denoted "PPO_CBL_test") They show that neither of the agents is capable of such generalization, which is consistent with the experiments with the pointing gesture in section 5.1.

These results motivate future research on what kind of biases could be built into the agents (and in what way) so that they could infer the meaning of familiar words in new contexts. For example, an interesting avenue of future work is to try to combine an agent with a large language models, and see if the knowledge contained in it could make the agent generalize better.

G.2. Joint Attention

Tomasello describes joint attention as consisting of two parts: triangulation and recursiveness (Tomasello, 2019). He argues that joint attention plays a key role in the 9-month revolution by transforming dyadic interactions (e.g. mimicking facial expressions) to triadic (e.g. imitating an action on an object). Joint attention was also required in the previous experiments (sections 5.1 and G.1). The agent and the peer triangulated on an external referent, however, the agent could assume that the peer was participating in the interaction.

In this experiment, we aim to conduct a more thorough test of the second aspect of joint attention - *recursiveness* (both participants being aware that they are both sharing attention). To solve the task, the agent needs to infer if the peer is participating and is aware that the agent is participating too. We create environments where the peer, in addition to giving regular cues inside joint attention, gives *misleading* cues outside joint attention. These cues are implemented uttering a random cue, and are given before the agent completes the introductory sequence. In other words, the agent should learn to discriminate between cues given for the agent during joint attention (after the introduction) and cues given regardless of the agent outside joint attention (before the introduction).

We study the following question:

- Can RL agents learn to differentiate between cues given inside and outside joint attention, i.e. can they learn to infer whether the peer is participating in the interaction?

Environments In this section, we extend the environment from section G.1 studying the COLOR cue type. The environment is extended so that the agent must also recursively infer whether the cue is intended for the agent. This misleading cue is given before the introductory sequence is completed, and takes the form of the peer uttering a color of a random one of the two objects. Apart from that, the experiments are conducted in the same way as in section G.1 (we train on the same problems for and use the same baseline).

Results Figure 17 compares the performance (success rate) of the agent trained on this extended environment (denoted by JA) with the agent trained on the regular environment (from the experiment in section G.1). These results show that the agent is not able to differentiate between cues given inside and outside of joint attention. We believe that this is due to the cues being highly misleading in this environment. As the peer utters the color of a random object present in the environment, there is a 50% possibility of the misleading cue being the same as the helpful one.

These results open many avenues for future research. One might study which kinds of biases can be integrated into the agent to make such cues less misleading. The generalization abilities of those agents should also be investigated. For instance, we could study if an agent that learned to ignore misleading linguistic cues would ignore misleading pointing cues.

G.3. Imitation learning

In the following section, we study the ability of the agent to learn how to obtain the apple by imitating the peer. This experiment is motivated by an experiment from (Carpenter et al., 1998b) discussed in section 3.1.3. In it, infants showed a steady increase in imitation learning abilities in the period between 9 and 15 months after birth. We want to test the agent's ability to imitate an instrumental action on an object.

From an AI perspective, this can be seen as meta-imitation learning. We want to see if an agent can obtain (through gradients) the imitation learning mechanism, which it could then use (during the episode) to learn how to use a new object. In this section, we study the following questions:

- Can RL agents learn (through gradients) an imitation mechanism?

In these experiment, we use the agent with the visual count-based exploration bonus (CB), as we found it worked best in our pilot study (see appendix E). We compare three agents trained with the same exploration bonus scaled by different weights: 0.25, 0.5, and 1.

Environment The Environment is an INFORMATION SEEKING type environment without a distractor. After the introductory sequence (EYE_CONTACT), the peer will demonstrate using an object to obtain the apple. For example, it will toggle a box or push a generator. Then the peer will then eat the apple, and revert the environment to its initial state. The agent should then imitate the peer - use the same action on the object - to obtain the apple for itself. If the agent uses the object it in the wrong way (e.g. pushes the box instead of toggling it) it will be blocked, and the apple will not be obtainable in this episode. The agents are evaluated on a new problem in which the agent encounters a new object for the first time. This means that the agent must pay attention to how the peer uses the object, and use it in the same way.²

The agents are trained on five problems (all expect DOORS). Most importantly, compared to the experiments in sections 5.1 and G.1, these agents will not be trained on the asocial version of the DOORS problem. That is because, in the generalization testing, we want to see if the agent can learn to use a completely new object.

Results Figure 18 shows the performance (success rate) of the agents on the training environments, the percentage of succesful introduction with the peer, and the evaluation on the (unseen) DOORS problem.

On figures 18(a) and 18(b) we can see that the agent with a lot of exploration bonus (PPO_CB_1) is too focused on the peer and, and is unable to solve the task. This is implied by the high percentage of the successful introductory sequence, and low success rate on the training environments. On the other hand, the agent with smaller exploration bonus weight (PPO_CB_0.25) solves these environments without problems, however it does use the peer. As such, the agent can solve the training environments by ignoring the peer and discovering how to use each object by itself. However, this agent is not able to generalize to a new object as the only way to know how to use that object is to observe the peer’s demonstration (see figure 18(c)). Figure 18(c) shows the performance of those agents on the testing environment. The figure shows that neither of the three agents is capable of acquiring a meta-imitation learning mechanism that can generalize to a novel object (door).

These results are not surprising, as current exploration bonuses are not well suited to enable RL agents to meta-learn mechanisms. These results imply that an interesting avenue of research is to study how to endow agents with such meta-imitation learning mechanisms that would enable them to learn a behavior in a new scenario. An promising solution to this problem are large language models and other large transformer-based networks pretrained on many other tasks. It would be interesting to study if such agents already have an imitation learning mechanism which would enable such online imitation. This would open up countless avenues of research into various forms of online imitation and emulation learning.

G.4. Inferring another’s field of view

In this section, we study the ability of the agent to infer what the other observes. This experiment is modeling the one in Hare (2001). In it, apes were shown to be able to infer what another sees, as they only took the food the alpha male could not see.

In this section, we want to study the following questions:

- Can agents learn to infer the other’s field of view?

Environment In the following experiment, we are using the ADVERSARIALPEER environment type, in which the agent has to eat the apple while not being seen by the peer. We study two versions of this environment: with and without obstacles. Obstacles make the problem of inferring the peer’s field of view harder.

Experiment We study how the agent infers the peer’s field of view by training the agent on the AdversarialPeer task.

It is important to note that this agent can sometimes use other (asocial) information to achieve performance. For example, if the object is surrounded by occlusions the agent could guess that it is not observed by the peer, which is not necessarily the case. To better understand the performance of the agent we compare the agent with two baselines. First, we assess to what extent the agent is making inferences based on the peer’s location and gaze direction. We train an agent (“invisible_peer”) that has the peer filtered from its observations (it cannot observe the peer). This baseline estimates the maximum possible performance. If the standard agent outperforms this baseline this implies that it is leveraging the social information in the

²The encoding of the peer includes the peer’s previous timestep action.

environment. Second, to estimate the upper bound on the performance we train an agent in the environment without the peer present (this agent is reward every time it eats the apple).

Results Figure 20 shows the performances of those three agents. It shows that the agent outperforms the agent with the peer filtered from its observations (“invisible_peer”), which implies that the agent is using the peer’s location and gaze direction to infer whether to eat the apple or not. Furthermore, the agent is not able to match performance of the agent trained without the peer present in the environment (“no_peer”). This results imply that, while the agent is able to leverage some social information in the environment, there still remains room for improvement. Future research could focus on constructing novel types of exploration bonuses to bridge this gap.

G.5. Formats

In the following experiment, we study the ability of the agent to learn formats (also referred to as pragmatic frames in (Vollmer et al., 2016)). Formats are a concept introduced by Jerome Bruner, which we discussed in more detail in section 3.2. They can be regarded as protocols of social interactions. We study the following question:

- To what extent can an exploration bonus help with the acquisition of a complex format.

We address this question by training two agents (one with an exploration bonus, and one without it).

Environment We use the INFORMATION SEEKING environment type with the LANGUAGE FEEDBACK cue type. We train all agents on all six problems. In contrast to section G.1, where the introductory sequence was always set to EYE CONTACT, here it is set to ASK EYE CONTACT - the peer will give cues after the agent utters “Help, please” during eye contact.

Results Figure 21 compares the performance of an agent that does not use any exploration bonus (“PPO_no_bonus”) to an agent that uses the visual count-based exploration bonus (“PPO_CBL”). The agent with the exploration bonus achieves high performance (97.9% success rate) and greatly outperforms the agent without the exploration bonus. These experiments show that, as expected, learning complex formats can be made easier with exploration bonuses.

This experiment can be interpreted in tandem with the experiment in section 5.3 where we show how more complex formats can be learned by weaker agents (without an exploration bonus) when learning in a scaffolded environment. Future work could explore how these two different approaches - modifying the agent and modifying the environment - can be used in tandem to learn even more complex formats. Furthermore, one interesting research direction is to study which kinds of problems are better addressed by modifying the agent and which by modifying the environment.

G.6. Additional information on the case study with large language models as interactive agents

Figures 25 and 26 show the in context examples given to large language models in section 5.4. The in context examples were created by hand.

H. Extended discussion

In this work, we outline and discuss several concepts from developmental psychology – mostly regarding the development before and around 9 months of age – which we found to be most relevant for AI at the moment. Even among this restricted set it is not reasonable to aim for an exhaustive introduction. As such, several socio-cognitive concepts are either discussed very briefly (e.g. conformity, social norms, instructed learning) and a lot of others are not mentioned (e.g. morality, fairness, sense of self). We leave their analysis for future work. Furthermore, while we argue that the work of Tomasello and Bruner provides an interesting framework to guide AI research in social skill acquisition, many other perspectives could have been considered, e.g. Erik Erikson (Erikson, 1993), Alison Gopnik (Gopnik & Meltzoff, 1997), or Cecilia Heyes (Heyes, 2019).

Similarly, as the present work merely represents a first step towards socially proficient artificial learners, many technical dimensions were simplified. In particular, we refrain from free form language dialogues and consider simple templated language. Likewise, we do not use human or trained peers, but scripted peers (which enables to isolate social abilities). Rather than implementing rich 3D visual worlds with continuous actions, we use grid-worlds with discrete primitive actions. We argue that such simplifying assumptions affords tractable studies while maintaining enough social complexity to model and isolate various social challenges. Assuming progress is made over these social scenarios, an interesting avenue for future work will be to extend the parametric generation towards environments with more complex sensorimotor challenges.

Table 2. Template-based grammar used in all of the SocialAI environments. If the agent decided to speak it chooses a template and a noun to insert into the template.

Nouns		
Action	Template	Noun
0	Where is <noun>	please
1	Help <noun>	the exit
2	Close <noun>	the wall
3	How are <noun>	you
4		the ceiling
5		the window
6		the entrance
7		the closet
8		the drawer
9		the fridge
10		the floor
11		the lamp
12		the trash can
13		the chair
14		the bed
15		the sofa

Table 3. Examples of actions in the environment. Second and third dimension must both either be underfined or not. In practice, there is an additional binary output which defines if the agent will speak.

Action	description
(1, -, -)	moves left without speaking
(1, 1, 5)	moves left and utters "Help the window"
(-, 1, 5)	doesn't move but utters "Help the window"
(-, -, -)	nothing happens

Given recent works showcasing the importance of Automatic Curriculum Learning in "asocial" DRL (Parker-Holder et al., 2022; Portelas et al., 2020), an interesting direction for future work would be to study whether this can also be observed in SocialAI. Our short case study on the importance of scaffolding (sec. 5.3) suggests a positive impact, although we restricted our analysis to simple expert curricula. An important challenge will be to design curriculum methods able to leverage the hierarchical structure of SocialAI's parametric tree, rather than the usual low-dimensional flat spaces of task-encoding parameters (predominant in the literature).

Large language models (LLMs) are present in many branches of artificial intelligence. A promising avenue of future research is the application of language models to interactive agents (Andreas, 2022). In this paper, we studied LLMs only on simple environments with a simple method - prompting the model with a few expert trajectories. While this approach showed impressive sample efficiency, it is very limited due to the constraints on the prompt size. These experiments should be revisited with more powerful methods such as fine-tuning or chain-of-thought prompting. Such methods could potentially make more complex social inferences, leading to better performance on many case studies in this paper, especially the ones related to generalization to new scenarios.

An important factor for the observed learning failures of our PPO agents in our case-studies might be linked to the simple forms of exploration bonuses that we used (variants of count-based bonuses (Bellemare et al., 2016)). Finding efficient exploration bonuses for social settings is a challenging task. In appendix E we show that RIDE and RND (Raileanu & Rocktäschel, 2020; Burda et al., 2018), two state-of-the-art exploration bonuses from classical DRL underperformed compared to our simple methods. An interesting avenue would be to study the robustness of recent exploration bonus methods designed for social scenarios (e.g. Zhang (2020)).

A limitation of this work is that, in our experiments with reinforcement learning, we only consider agents learning through pure extrinsic or intrinsic reward. An interesting avenue for future work is to study how existing imitation learning methods such as GATO (Reed et al., 2022) or the decision transformer (Chen et al., 2021) could acquire social skills better if provided with expert demonstrations and trained jointly on other tasks which require social competence (e.g. Sap (2019), Netanyahu (2021), Shu (2021)). Another promising line of works are methods combining reinforcement learning signals with imitation learning (Hester et al., 2018; Nair et al., 2018).

Another important limitation of the present work is the lack of human-AI interactions during learning. In light of human social learning, such interactivity appears like an important step towards socially proficient learners. While we argue that scripted peers constitute a good basis for initial progress, it is crucial for long-term progress that more interactive learning paradigms are considered inside SocialAI. (Abramson et al., 2021) constitutes a noteworthy step towards this goal. On a more general note, featuring human inputs inside the learning pipeline of autonomous learners could constitute an efficient way to instill human-like inductive biases inside agents, which is an important feature to align agents' cognition to ours. For instance, in Kumar (2022) the authors show in a 2D-tile revealing task that autonomous agents learning from human-generated language description reached performances more similar to humans compared to agents trained with synthetic language.

From a high level perspective, the SocialAI school presents many relevant concepts (e.g. skills, motivation) which are at the core of human cognition. We discussed in more details those we considered most relevant for AI at the moment. We do not claim to have covered every relevant aspect of early social cognition: many other dimensions remain to be studied, both from the work of Tomasello and Bruner and from other authors. The SocialAI school is a tool which can be easily extended and modified to study many of these aspects. We hope this paper will motivate the AI community around the questions of social cognition. This avenue of research, on a longer timeline, promises to endow AI with the abilities and motivations enabling it to enter and push our culture forward, i.e. participate in our cumulative cultural evolution.

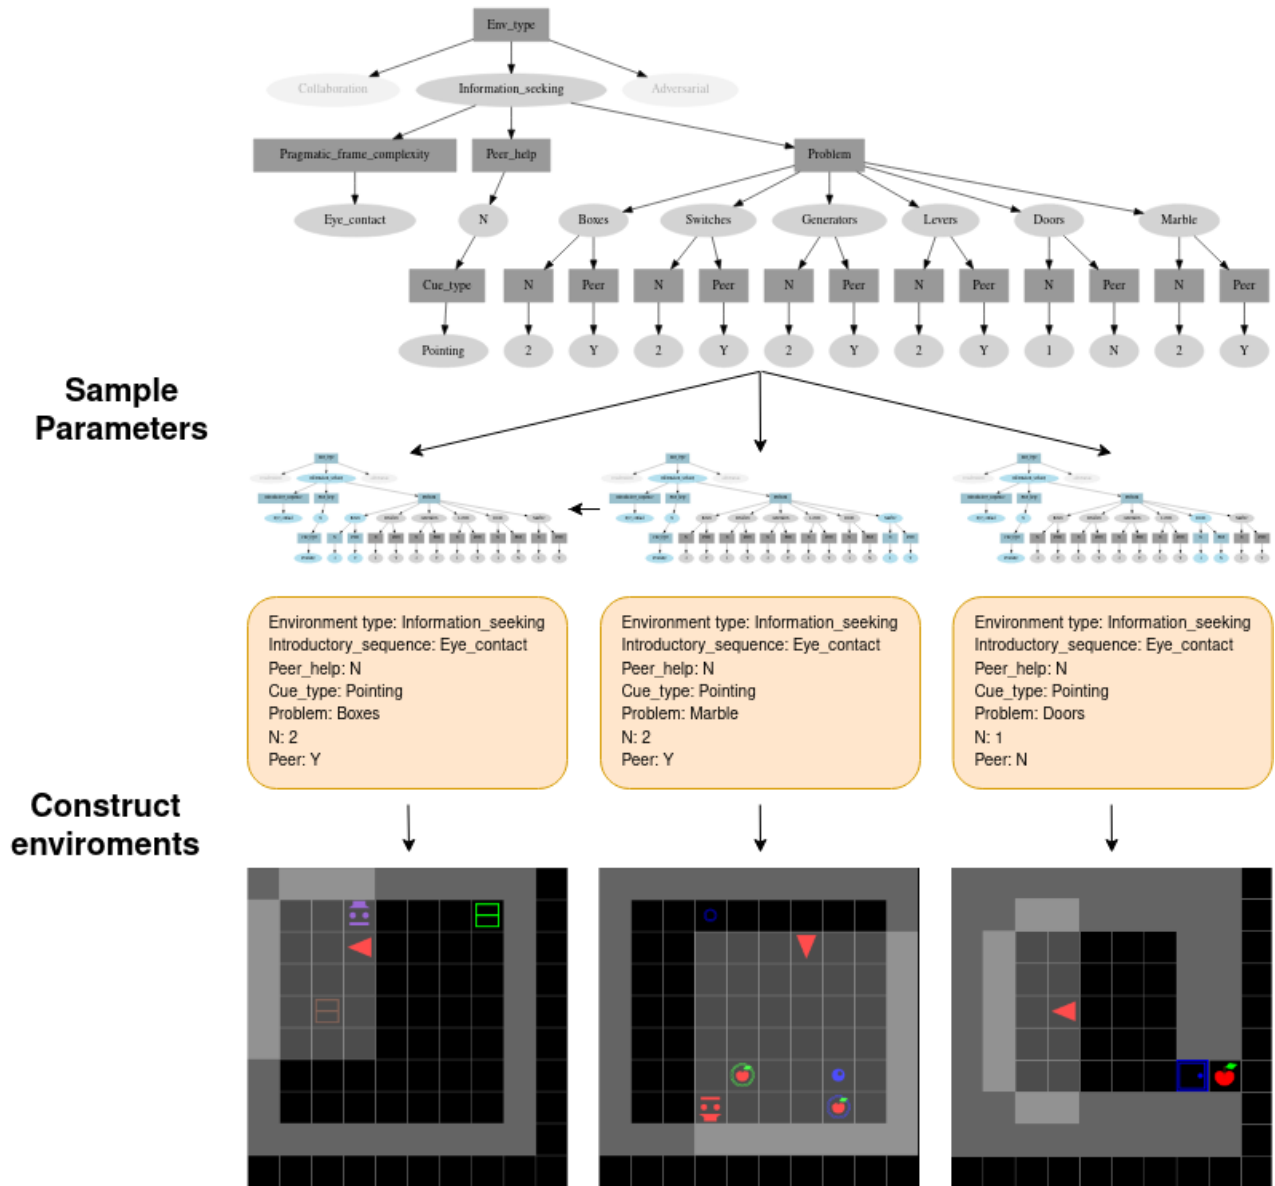

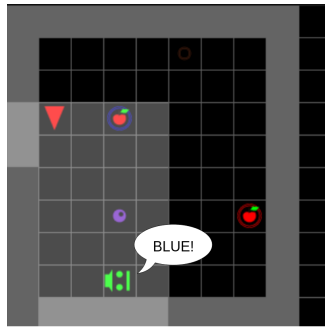



Figure 10. An example of procedural environment generation using tree-based parametric sampling. There are two kinds of nodes: parameter nodes (rectangles) and value nodes (ovals). Parameter nodes require that one of its children (a value node) is selected. Value nodes require that sampling progresses through all of its children (parameter nodes). In this tree, all parameter nodes except "Problem" have only one child. This means that only the Problem parameter can be set in different ways. We show three examples of parameter sampling, and the three environments constructed from those parameters.



(a) A scripted peer pointing to a box. The agent needs to open the red box.




(b) A scripted peer uttering the color of the correct generator. The agent needs to push the marble onto the blue generator.




(c) A scripted peer hinting the distance to the correct lever ("Hot" means very close). The agent needs to pull the purple lever to open the door.


Figure 11. Examples of INFORMATIONSEEKING type environments, in which agents learn to find hidden apples using textual or non-verbal communication with social peers.



(a) The MARBLEPASS problem with the agent in role B. The peer pushes the marble to the right and then the agent pushes it further to the purple marble generator. This makes two apples appear on the blue and red platforms.



(b) The LEVERDOOR problem with the agent in role B. The peer opens the red door by pulling on the green lever. This enables the agent to go through the door and activate the purple generator. This makes two apples appear on the gray and yellow platforms.



(c) The MARBLEPUSH problem with the agent in role A. The peer opens the yellow door using the green lever. Then the agent pushes the marble through the door to the purple marble generator. This makes two apples appear on the purple and green platforms.

Figure 12. Examples of COLLABORATION type environments, in which agents must learn cooperative strategies with a (scripted) peer to achieve two-player puzzles.

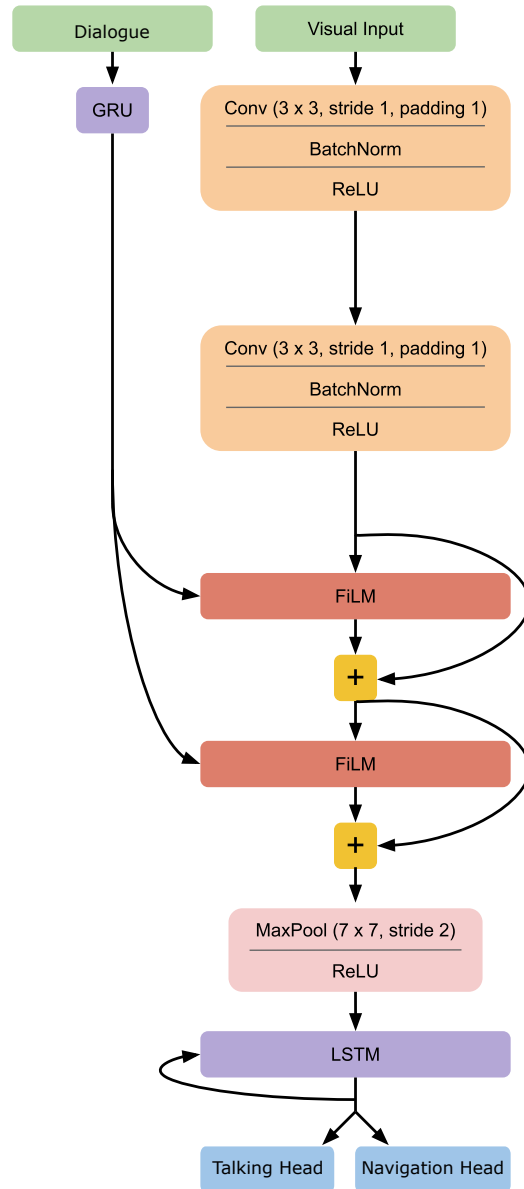

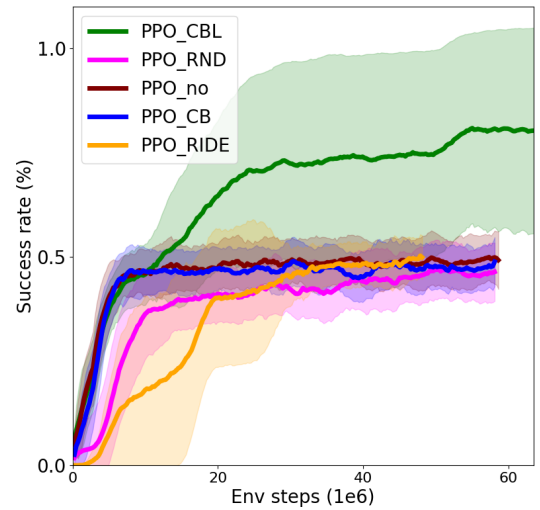


Figure 13. Our Multi-Headed PPO baseline DRL agent. Architecture visualization is a modified version of the one made by Hui (2020). We perform two modifications: 1) Instead of fixed instruction inputs our model is fed with NPC’s language outputs (if the agent is near an NPC), and 2) We add a language action head, as our agent can both navigate and talk.




(a) Pilot experiments with the peer pointing to the correct object.




(b) Pilot experiments with the peer uttering the color of the correct object.

Figure 14. Pilot experiments showing that our count-based exploration bonuses outperform other baselines. On the environments with the pointing gesture, visual count-based (“CB”) exploration bonus is the best performing condition. On the environments with utterances, linguistic count-based (“CBL”) exploration bonus is the best performing condition.



(a) Language Feedback cue type experiments: the peer gives cues regarding the proximity of the agent is to the correct object (e.g. Hot, Warm, Cold).



(b) Language Color cue type experiments: the peer utters the color of the correct object.

Figure 15. **The linguistic cues experiments.** We study if an RL agent is able to infer the meaning of linguistic cues in order to use the correct object. We consider two types of cues: *language feedback* and *color*. In both settings, the agent was trained on five different problems, and on the asocial version of the Doors problem (only one door and no peer present in the environment) - denoted by “train”. Agents were periodically evaluated on the social version of the Doors problem (two doors and a peer giving cues) - denoted by “test”. The figure compares the success rate (mean +/- std over 8 seeds) on the training environments with the evaluation on the testing environment. The cross marks depict statistical significance ($p = 0.05$). In both cases the agents achieve much better performance on the training problems, but fail to generalize to a new problem - the agent is not able to infer the meaning of an utterance in a new context.

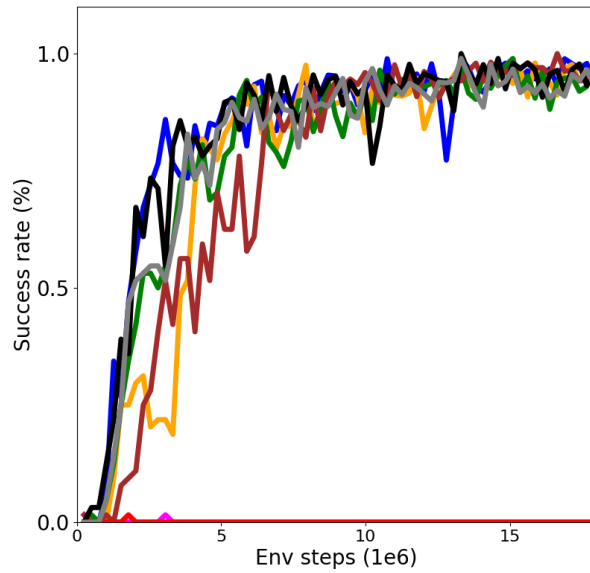


Figure 16. Per-seed performance on the training environments of the agent from figure 15(a) ("PPO_CBL(train)"). The figure shows that the agent is able to solve the training tasks efficiently, but that there are two unstable seeds which result in the success rate of 0%.

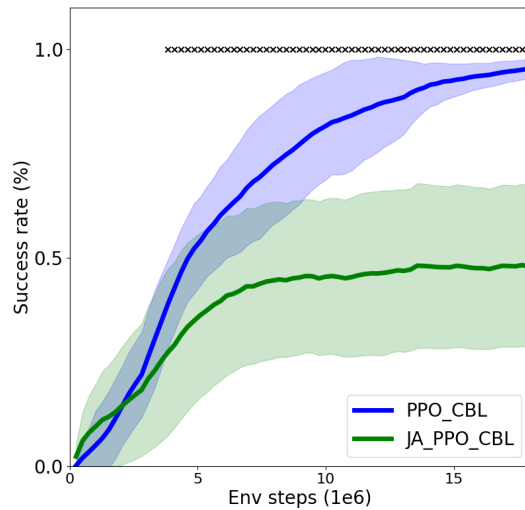
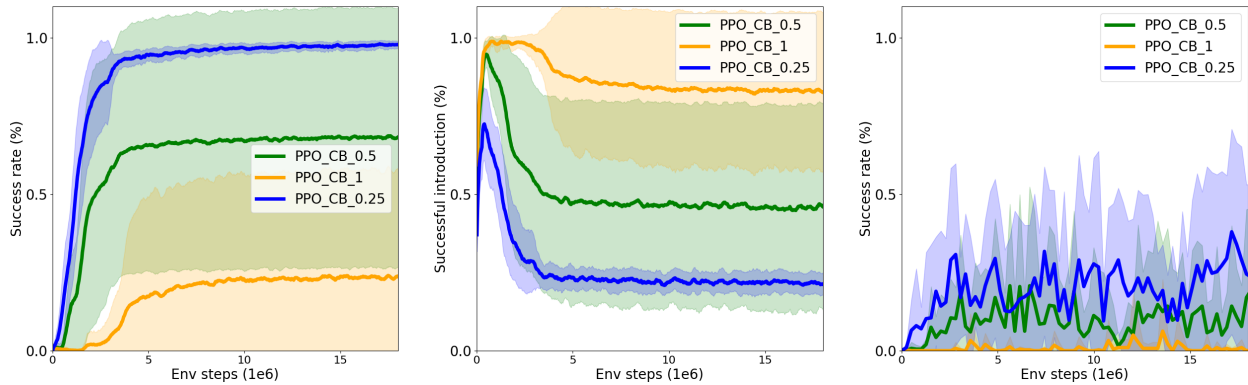
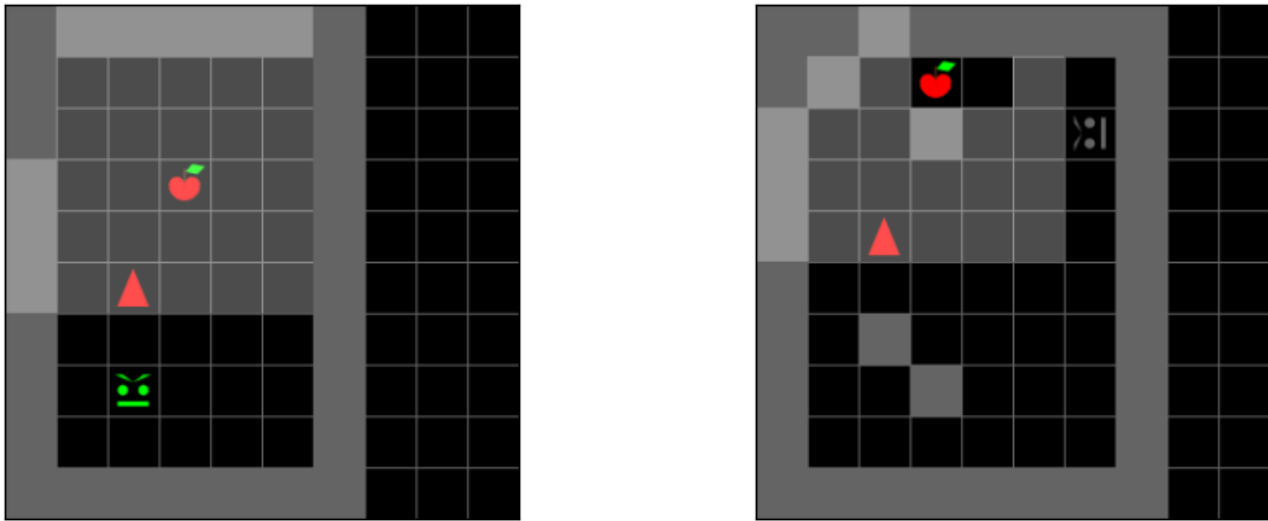


Figure 17. **The joint attention experiment.** The environments feature a test for recursiveness - infer if the peer knows that they are working together. The environments are same as the ones from figure 15(b), but with the addition of misleading cues - random cues given regardless of the agent (a random color). The peer gives misleading cues outside of joint attention (before the introductory sequence). The agent should ignore these cues, and use only cues given inside joint attention. The figure compares the success rate (mean +/- std over 8 seeds) of the agent trained on the environments with both regular and misleading cues ("JA_PPO_CBL"), to the agent trained on the environments with only regular cues ("PPO_CBL(train)" from figure 15(b)). The figure shows that the agent unable to master the Joint Attention variant.



(a) Imitation experiments performance (success rate) on the training environments. (b) The percentage of successful introductory sequences on the training environments. (c) Imitation experiments performance (success rate) on the testing environment.

Figure 18. Imitation learning experiments. The peer demonstrates how to use an object (after the agent successfully introduces itself). The agent is trained on five different problems and evaluated on a new problem with a previously unobserved object (a door). A socially proficient agent should be able to learn (by observing the demonstration) which action (toggle or push) to use on the new object. The curves compare three agents trained with a different scaling factor for the visual count-based exploration bonus. One can see that the agent with high exploration bonus (“PPO_CB_1”) focuses too much on the peer, which results in ignoring the task. This is evidenced by high success in completing the introductory sequence (fig. 18(b)), but low success rate on the task (fig. 18(a)). On the other hand, using low exploration bonus (“PPO_CB_0.25”) pushes the agent to solve the training task whilst ignoring the peer. Rather than observing the peer’s demonstration, this agent learns how to use objects by themselves. This results in perfect performance on the training object, but it makes it impossible to generalize to a new object. Neither of the agents is able to achieve high performance on the held out testing environment. This implies that they are not able to learn (online) through imitation which action to use with a new object.



(a) Adversarial peer environment without occlusions. (b) Adversarial peer environment with occlusions.

Figure 19. Environments from the Adversarial peer experiments in which the agent has to infer the peer’s field of view. The agent is rewarded upon eating the apple on the condition that it was not in the field of view of the peer while doing so. We run the experiments with two different settings: with and without occlusions (depicted in figures 19(b) and 19(a)). Occlusions make it harder to infer the peer’s field of view as it is no longer rectangular.




Figure 20. Adversarial peer experiments. We compare three agents on two environments (depicted on figures 19(b) and 19(a)). The "PPO_CB" agent is trained on the regular environment (rewarded upon eating the apple while not being observed by the peer). The "PPO_CB_no_peer" agent is trained in the environment without the peer (the agent is rewarded every time it eats the apple). This represents the upper bound of the performance. The "PPO_CB_invisible_peer" agent is trained on the regular environment with the peer filtered from the agent's observations. This represents the performance of a completely asocial agent which ignores the peer. Figures 20(a) and 20(b) compare the performance of these three agents (8 seeds +- std), the crosses depict a statistically significant difference ($p_1 < 0.05$) compared to the "PPO_CB" agent. The results show that the "PPO_CB" agent is able to partially infer the peer's field of view (as it outperforms the "invisible_peer" baseline), but is not able to reach perfect performance (as defined by the "PPO_CB_no_peer" baseline).




Figure 21. Comparison of an agent with and without the exploration bonus on an environment with a more complex introductory sequence (format). The task consists of the agent doing the introductory sequence by making eye contact and uttering "Help, please". The peer will then give linguistic cues regarding the proximity of the agent to the target object (e.g. Hot, Warm, Cold). Based on these cues, the agent should use the target object, instead of the distractor, to obtain the apple. The figure shows that using the visual count-based exploration bonus enables the agent to learn a more complex introductory sequence and solve the task.




Figure 22. Visualizations and descriptions for all objects featured in SocialAI environments.




Figure 23. Two environments used in the experiments with large language models. The observations are parsed into pure text.

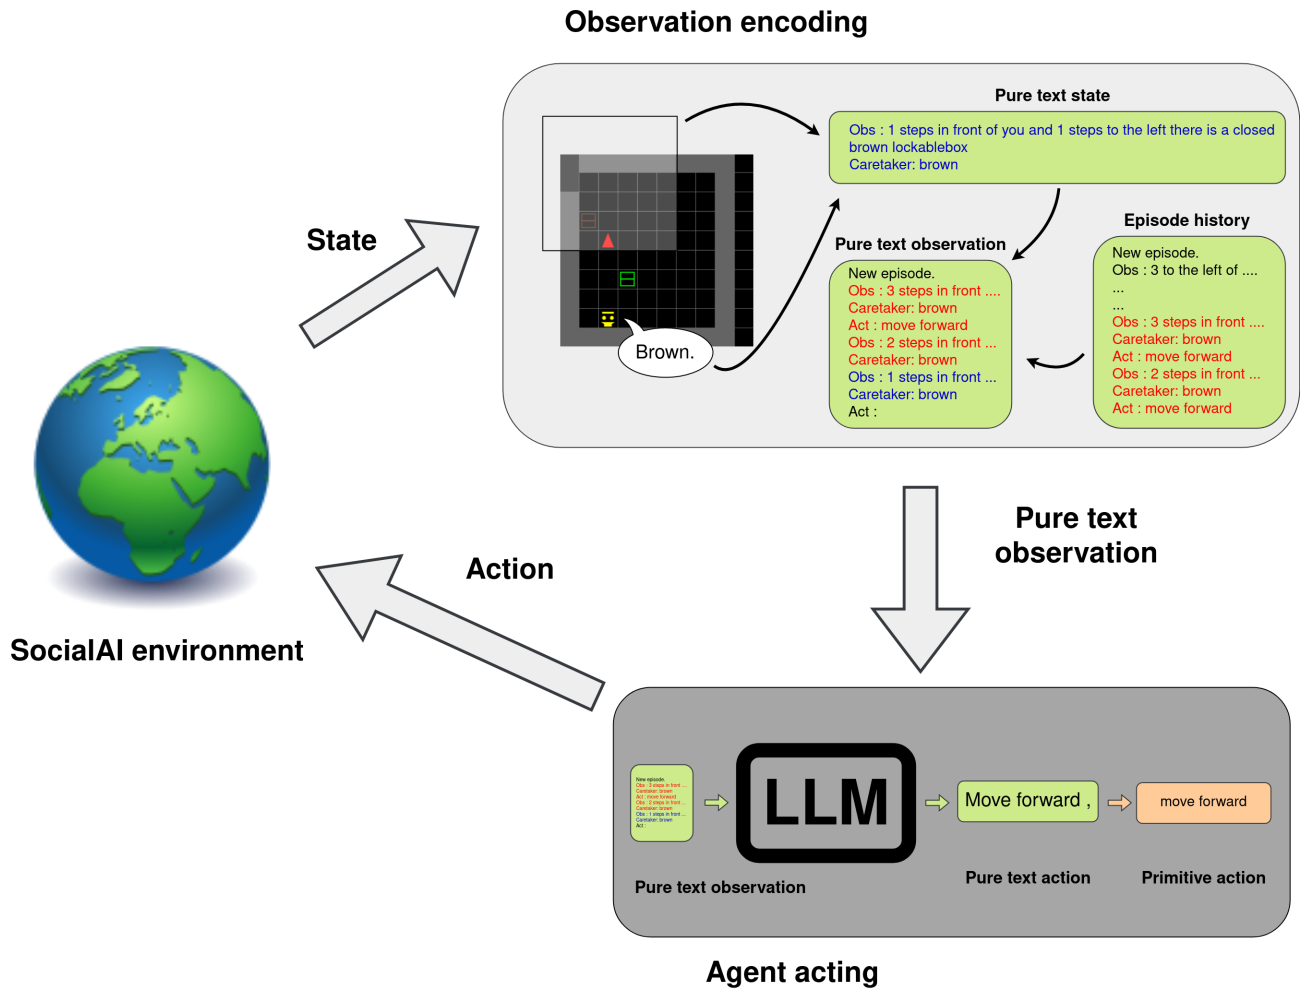


Figure 24. An example of how a language model can be used as an interactive agent in SocialAI. A state is parsed into a pure text observation and combined with previous two observations and actions. This is, appended to the in context examples, is used as prompt for the LLM. The agent generates the text which is then matched (as case insensitive substring) with the list of possible actions. The matched action is executed in the environment.

New episode.
Obs : 3 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : 2 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : Right in front of you there is a closed green lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!
New episode.
Obs : 2 steps to the right there is a closed green lockablebox
Act : turn right
Obs : 2 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : Right in front of you there is a closed green lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!
New episode.
Obs : 2 steps to the left there is a closed green lockablebox
Act : turn left
Obs : 2 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : Right in front of you there is a closed green lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!
New episode.
Obs : 2 steps in front of you and 2 steps to the right there is a closed blue lockablebox
Act : move forward
Obs : 1 steps in front of you and 2 steps to the right there is a closed blue lockablebox
Act : move forward
Obs : 2 steps to the right there is a closed blue lockablebox
Act : turn right
Obs : 2 steps in front of you there is a closed blue lockablebox
Act : move forward
Obs : Right in front of you there is a closed blue lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!

New episode.
Obs : 2 steps in front of you and 1 steps to the left there is a closed green lockablebox
Act : move forward
Obs : 1 steps in front of you and 1 steps to the left there is a closed green lockablebox
Act : move forward
Obs : Just to the left of you there is a closed green lockablebox
Act : turn left
Obs : Right in front of you there is a closed green lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!
New episode.
Obs :
Act : turn right
Obs : 1 steps in front of you and 1 steps to the right there is a closed blue lockablebox
Act : move forward
Obs : Just to the right of you there is a closed blue lockablebox
Act : turn right
Obs : Right in front of you there is a closed blue lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!
New episode.
turn right
Obs :
Act : turn right
Obs :
Act : turn right
Obs :
Act : move forward
Obs :
Act : turn right
Obs : 4 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : 3 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : 2 steps in front of you there is a closed green lockablebox
Act : move forward
Obs : Right in front of you there is a closed green lockablebox
Act : toggle
Obs : Right in front of you there is a red apple
Act : toggle
Success!

Figure 25. In context examples for the SocialAIBox environment.

<p>New episode. Obs : 3 steps in front of you there is a closed green lockablebox Caretaker: green Act : move forward Obs : 2 steps in front of you there is a closed green lockablebox Caretaker: green Act : move forward Obs : Right in front of you there is a closed green lockablebox Caretaker: green Act : toggle Obs : Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : 2 steps to the right there is a closed green lockablebox Caretaker: green Act : turn right Obs : 2 steps in front of you there is a closed green lockablebox Caretaker: green Act : move forward Obs : Right in front of you there is a closed green lockablebox Caretaker: green Act : toggle Obs : Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : 2 steps to the left there is a closed green lockablebox Caretaker: green Act : turn left Obs : 2 steps in front of you there is a closed green lockablebox Caretaker: green Act : move forward Obs : Right in front of you there is a closed green lockablebox Caretaker: green Act : toggle Obs : Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : 2 steps in front of you and 2 steps to the right there is a closed blue lockablebox Caretaker: blue Act : move forward Obs : 1 steps in front of you and 2 steps to the right there is a closed blue lockablebox Caretaker: blue Act : move forward Obs : 2 steps to the right there is a closed blue lockablebox Caretaker: blue Act : turn right Obs : 2 steps in front of you there is a closed blue lockablebox Caretaker: blue Act : move forward Obs : Right in front of you there is a closed blue lockablebox Caretaker: blue Act : toggle Obs : Right in front of you there is a red apple Caretaker: blue Act : toggle Success! New episode. Obs : 2 steps in front of you and 1 steps to the left there is a closed green lockablebox Caretaker: green</p>	<p>Act : move forward Obs : 1 steps in front of you and 1 steps to the left there is a closed green lockablebox Caretaker: green Act : move forward Obs : Just to the left of you there is a closed green lockablebox Caretaker: green Act : turn left Obs : Right in front of you there is a closed green lockablebox Caretaker: green Act : toggle Obs : Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : Caretaker: blue Act : turn right Obs : 1 steps in front of you and 1 steps to the right there is a closed blue lockablebox Caretaker: blue Act : move forward Obs : Just to the right of you there is a closed blue lockablebox Caretaker: blue Act : turn right Obs : Right in front of you there is a closed blue lockablebox Caretaker: blue Act : toggle Obs : Right in front of you there is a red apple Caretaker: blue Act : toggle Success! New episode. Obs : 3 steps in front of you and 3 steps to the left there is a friendly yellow peer. It is looking towards you. 1 steps to the right there is a closed brown lockablebox Caretaker: brown Act : turn right Obs : Right in front of you there is a closed brown lockablebox Caretaker: brown Act : toggle Obs : Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : 1 steps to the right there is a closed blue lockablebox Right in front of you there is a closed brown lockablebox Caretaker: brown Act : toggle Obs : 1 steps to the right there is a closed blue lockablebox Right in front of you there is a red apple Caretaker: green Act : toggle Success! New episode. Obs : 1 steps to the right there is a closed blue lockablebox Right in front of you there is a closed brown lockablebox Caretaker: brown Act : toggle Obs : 1 steps to the right there is a closed blue lockablebox Right in front of you there is a red apple Caretaker: green Act : toggle Success!</p>
--	---

Figure 26. In context examples for the ColorBoxes environment.

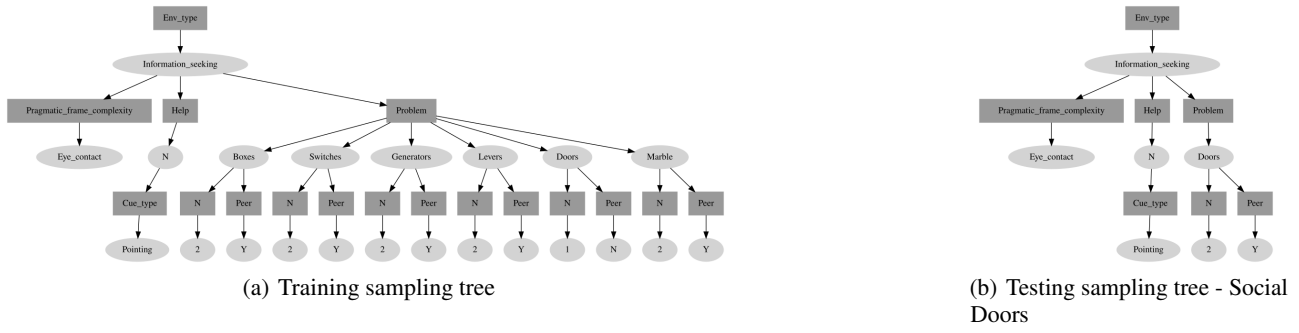


Figure 27. Sampling trees used in the pointing case study in section 5.1



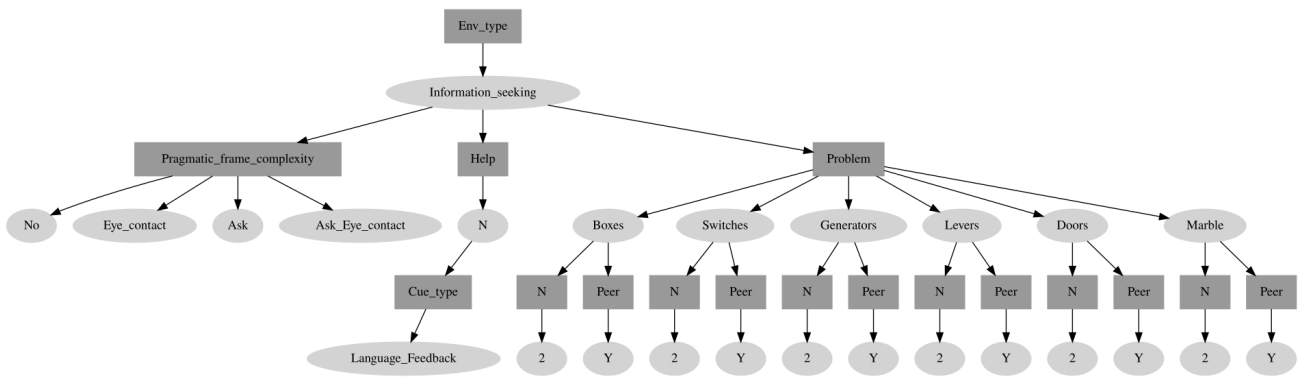



Figure 28. Role reversal sampling trees from the case study in section 5.2



(a) Testing tree.



(b) Scaf_4 tree.



(c) Scaf_8 tree.

Figure 29. Sampling trees used in the first phase of the scaffolding case study in section 5.3

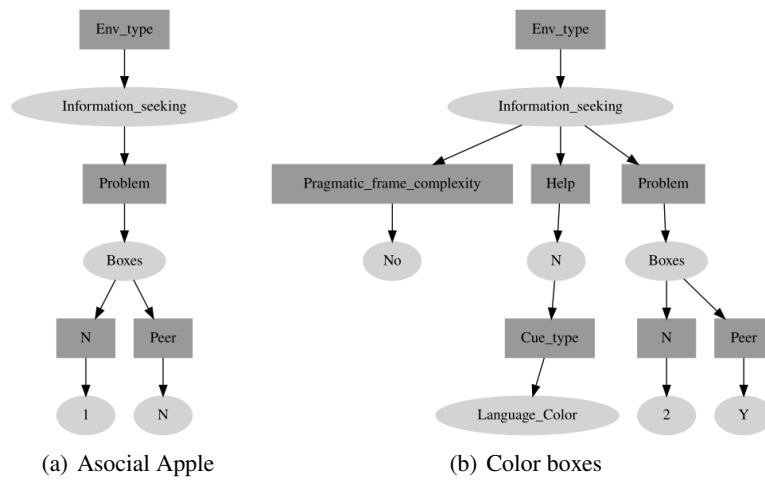


Figure 30. Sampling trees used in the case study with LLMs as interactive agents (section 5.4)