

Project-team proposal, Jan 2025

Flowers AI & CogSci

(AI: Studying Autotelic and Aligned Intelligence, using Artificial Intelligence modelling, with applications in Assisted and Augmented human Intelligence)

FLOW in Exploration, leaRning, and diScovery

Inria and University of Bordeaux

Domain: Perception, Cognition, Interaction

Theme: Robotics and intelligent environments

Table of Contents

Abstract.....	3
Team Composition.....	4
1. Scientific context, challenges and general objectives.....	5
2. Research Program: origins, evolution and new directions.....	10
2.1 Origins of research program: Theories of curiosity-driven learning, sensorimotor development and developmental robotics (2011-2017).....	10
2.2 Current research program (2017-ongoing): testing theories of curiosity-driven learning in humans, language augmented autotelic deep RL and generative models, and new application domains.....	13
Objective 1: Improve understanding of human autotelic and aligned intelligence.....	14
Objective 2: Building curiosity-driven autotelic and aligned machines.....	14
Objective 2.1: Connecting and disseminating autotelic learning to the machine learning community.....	14
Objective 2.2: Curiosity-driven autotelic AI agents leveraging large foundational models.....	16
Objective 2.3: Cultural evolution, innovation and curiosity at the group level.....	17
Objective 3: Applications in education and assisted scientific discovery.....	18
Objective 3.1: Translational educational technology research: training curiosity and metacognition.....	18
Objective 3.2: Autotelic algorithms for assisted scientific discovery.....	21
Evolution of team structure.....	22
2.3 The horizon: research and applications in 10 years.....	22

Autotelic learning in human individuals and social groups: development, cross-cultural, neurodiversity (Objective 1).....	23
Autotelic and aligned machine learning: program synthesis for abstract, grounded, generative and frugal system 2 intelligence (Objective 2).....	24
Generalized training of curiosity, metacognition and creativity across the lifespan (Objective 3.1).....	26
Autotelic and aligned assisted scientific discovery: towards artificial labs for assisting the scientific community (Objective 3.2).....	27
3. Software.....	28
4. Impact in science, public policy, industry and society.....	29
4.1 Scientific impact.....	29
4.2 Societal and industrial impact.....	31
4.3 Support to public policies.....	32
4.4 Key industrial collaborations.....	33
5. Positioning and collaborations.....	34
6. References.....	37
7. Appendix: Detailed research program and applications.....	44
7.1 Objective 1: Improve understanding of human autotelic and aligned intelligence in humans.....	44
7.2 Objective 2: Building curiosity-driven autotelic and aligned machines.....	48
7.2.1 Objective 2.1: Connecting and disseminating autotelic learning to the machine learning community.....	48
7.2.2 Objective 2.2: Curiosity-driven autotelic AI agents leveraging large foundational models.....	49
7.2.3 Objective 2.3: Cultural evolution, innovation and curiosity at the group level.....	54
7.3 Objective 3: Applications in education and assisted scientific discovery.....	57
7.3.1: Objective 3.1: Train curiosity-driven autotelic learning in humans across the lifespan.....	57
7.4 Objective 3.2: Assisted scientific discovery with autotelic exploration algorithms.....	62

10-line abstract:

This project-team aims to study the mechanisms enabling **open-ended learning in humans and machines**, with a focus on **curiosity-driven autotelic learning** and on **language** (both as a cognitive tool and as a tool enabling alignment and **collective intelligence**).

We study these mechanisms along 3 axes:

- 1) **computational modelling** of autotelic learning in humans (individuals and groups), and testing the predictions of these models in novel human experimental paradigms using **methods from experimental psychology**;
- 2) building language-augmented autotelic **deep reinforcement learning** systems, leveraging, grounding and improving **foundational models**, enabling creative exploration, sample efficient learning and robust generalization, and studying how autotelic agents can align with objectives of a group/culture;
- 3) applications in **educational technologies** fostering curiosity-driven learning in humans, and

in **assisted scientific discovery**, using curiosity-driven exploration algorithms to help scientists discover new complex systems.

Abstract

This project-team aims to study the fundamental mechanisms that can enable **open-ended learning and development in humans and machines**, i.e. how individuals, or groups of individuals, can continuously discover and learn novel skills of increasing complexity. We also aim to leverage this fundamental understanding for human-centered real-world applications in education and in assisted scientific discovery.

In particular, we focus on studying mechanisms enabling **Autotelic and Aligned Intelligence** in humans and machines. A first key ingredient of open-ended learning is curiosity-driven **autotelic learning**, which is the ability of individuals to set and pursue their own goals (from the greek ‘telos’/goal, and ‘auto’/self), a form of intrinsic motivation pushing organisms to continuously seek new knowledge and skills. self-organizing their own learning curriculum, using meta-cognition and leading to creative exploration.

To enable abstraction, collective intelligence, and alignment of autotelic systems on human cultures (values, preferences), we also aim to study how **language and social interaction**, both as a communication system and as a cognitive tool, can guide autotelic exploration. Symmetrically, using multi-scale models, we aim to study how curiosity-driven autotelic exploration could self-organize at the group level. We also aim to study what are the ecosystemic and evolutionary origins of autotelic systems.

We study these mechanisms from three complementary scientific perspectives:

1) **Improving our understanding of human curiosity-driven autotelic and aligned intelligence**: here we use methods from psychology to design new experimental protocols, and computational theories and models to design and test hypotheses about human autotelic learning and how it interacts with metacognition, language and social interaction;

2) **Building curiosity-driven autotelic and aligned machines**: we aim to design and build open-ended learning machines that are autotelic, data frugal, with strong generalisation skills, and leverage language and social interaction to integrate within human cultures (values, ethics, preferences); Here we leverage **generative AI** systems, encoding forms of human cultural knowledge, as cognitive tools, and we aim to improve the frugality and grounding of generative AI systems using autotelic deep reinforcement learning;

3) **Applications**: in **Educational technologies**, we aim to stimulate curiosity-driven autotelic learning, meta-cognition and creativity in humans across the lifespan and across neurodiversity (leveraging both models of human autotelic learning and frugal generative edTech tools); In the domain of **Assisted scientific discovery**, we aim to study how curiosity-driven autotelic exploration algorithms can help scientists (physicists/chemists/mathematicians, ...) make discoveries in complex systems.

Beyond core scientific questions across disciplines, this project addresses two key societal challenges: 1) How can we **build AI systems that serve humans and human societies in their diversity**, helping their curiosity and cultures to bloom? 2) How can we provide **educational opportunities for all children, and adults across the lifespan**, in a world with many challenges, to become intrinsically motivated learners, critical thinkers, autotelic explorers?

Team Composition

C: computer science background; **H**: human sciences background

Senior Research Scientists and Faculty Members

Pierre-Yves Oudeyer (Team leader, Inria, DR, HDR, **C, H**) ; Hélène Sauzéron (INRIA, Inria, DR, HDR, **H**) ; Clément Moulin-Frier (Inria, Researcher, **C**) ; Cécile Mazon (Univ. Bordeaux, Associate Professor, **H**)

Post-Doctoral Fellows

Cedric Colas (Inria, Post-Doctoral Fellow, **C**) ; Marion Pech (Inria, Post-Doctoral Fellow, **H**) ; Maria Teodorescu (Inria, Post-Doctoral Fellow), **C** ; Olivier Clerc (Inria, arri. Nov. 2024, **H**)

PhD Students

Rania Abdelghani (EvidenceB, **H**) ; Maxime Adolphe (OnePoint, **H**) ; Thomas Carta (UNIV BORDEAUX, **C**) ; Marie-Sarah Desvaux (Univ. Bordeaux, **H**) ; Gautier Hamon (Inria, **C**) ; Grgur Kovac (Inria, **C**) ; Jeremy Perez (Univ. Bordeaux, from Oct 2023, **C,H**) ; Matisse Poupard (Catie, **H**) ; Julien Pourcel (Inria, **C**) ; Clément Romac (HuggingFace, **C**) ; Isabeau Saint-Supery (Univ. Bordeaux, **H**) ; Nicolas Yax (ENS Paris, **C**)

Research Engineers

Zaccharie Bugaud, **C** ; Corentin Léger, **C**

Administrative Assistant

Nathalie Robin [INRIA]

External Collaborator

Didier Roy **C,H** ; Eleni Nisioti **C** ; Nathanaël Fijalkow **C**

Extended team (people with whom we have/plan to have major collaborations over several years): Olivier Sigaud (ISIR, Sorbonne Université, **C**) ; Stefano Palminteri (ENS Paris, **H**) ; Nathanaël Fijalkow (Labri, **C**) ; Jacqueline Gottlieb (Columbia Univ., **H**) ; Matthias Gruver (Univ. Cardiff, **H**) ; Yana Fandakova (Univ. Trier, **H**) ; Gaia Molinaro (Berkeley, **H**) ; Peter Ford Dominey (Inserm, **C**) ; Sylvain Lamprier (Univ. Angers, **C**) ; and researchers from companies that collaborate with us on a personal basis: Bert Chan (Google Deepmind, **C**) ; Ida Mommenejab (Microsoft, **C, H**) ; Benjamin Clément (EvidenceB, **C**) ; Rémy Portelas (Ubisoft, **C**).

1. Scientific context, challenges and general objectives

Humans explore, learn and discover continuously novel skills and knowledge, through **open-ended processes**. In fact, humans, and some other life forms, are equipped with intrinsic motivation systems (“curiosity”) pushing them to spontaneously explore and actively seek new knowledge (Gottlieb and Oudeyer, 2018), setting and pursuing their own goals (they are **autotelic**, Colas et al., 2022b), ranging from the most concrete (e.g. stack cubes) to the most abstract (e.g. invent new maths problems).

This happens at the level of **individuals**, starting with children who eagerly and spontaneously explore their bodies and their environment as they **develop**, up to adults of all ages and all backgrounds. This autotelic exploration process also benefits from social and collective dynamics, leveraging past discoveries and being guided to align with the culture (values, preferences, ethics) of a given group (Henrich et al., 2016). The iteration of this collective intelligence process, accumulating and transmitting discoveries over generations, gives rise to open-ended **cultural evolution, and to autotelic exploration at the level of collectives**.

Understanding the mechanisms that enable the origins and functionalities of autotelic learning in interaction with social groups and culture, giving rise to open-endedness, is still a major mystery for science.

Objective 1: Improve understanding of human autotelic and aligned intelligence

Thus, a first objective of this project is to advance our fundamental understanding of the origins, mechanisms and functionalities of autotelic learning and exploration, and how this interacts and is aligned with collective dynamics. As described below, this will involve a combination of computational models for developing new theories and hypotheses, as well as design and analysis of new human experimental paradigms analysed with these computational models. Particular scientific questions we target include studying links between autotelic learning, metacognition (one’s own ability to know and control one’s own knowledge and cognitive functions) and creativity in humans, e.g. how do these skills develop in childhood and which internal and external factors influence them ? How do they link to language and processes of social interaction ?



An intriguing feature of the human species is our ability to continuously invent new problems (autotelic exploration) and to proactively acquire new skills in order to solve them, while leveraging and being guided and aligned on the culture of the groups in which they develop (adapted from Moulin-Frier, 2022)

Artificial systems. The richness and efficiency of **open-ended processes in humans**, and more generally in nature, strongly **contrasts with most current artificial systems** (such as those developed in the machine learning communities). Let's consider child development: in just a few years, most children learn a wide diversity of skills ranging from locomotion, visual navigation, dexterous manipulation, building lego structures, drawing, riding bicycles, playing collective sports, understanding the mind of others and joint intentionality, emotion communication, language, or invent strategies to communicate non verbally when visiting countries in a foreign language. And contrary to most classical AI systems, they do this without the intervention of an engineer: no-one opens their brain to tune parameters, or provide a carefully curated dataset, for acquiring new skills. And they are most often well integrated in the culture of their groups as a result of social learning and education. Furthermore, they **learn all this diversity of skills within very constrained resources of time, energy and data**. They are not given a pre-made dataset to learn from, rather **they need to collect their own learning data by exploring the physical world and interaction with social peers**: this is very costly in time and energy (even reading books!), and yet they learn smoothly.

On the contrary, even the most powerful and impressive recent AI systems, such as **large foundational models** (e.g. GPT-4, Claude, Llama3, Gemini or Mistral), are very far from such open-endedness and efficiency. First, their mode of learning is fundamentally different: they learn from **very large datasets pre-collected by engineers**, orders of magnitude larger than data self-collected by children. Even though it was shown that diverse "skills" self-organise as a result of learning regularities in such datasets (Brown et al., 2020; Bubeck et al., 2024), these generative models still require engineers to provide more data, or specific programs or prompts to generate domain specific synthetic data, to extend their skill sets. Crucially, **they do not include mechanisms for exploring autonomously novel environments, they do not learn to achieve goals in these environments, in particular they are not autotelic**: they do not

self-generate goals as children (any learner) or human scientists (knowledge builders) would do. Finally, because they are not trained to solve problems in external environments (they are just trained to fill the gaps in texts, images and videos, or to maximise human preferences in RLHF), they **lack grounding and causal understanding** (Bender and Koller, 2020; Mahowald et al., 2023), and methods to align them with human values are severely limited due to the lack of rich social interaction and feedback with and from humans. In brief, generative AI systems, as they are developed today, are not autonomous, not autotelic, not grounded, not well aligned and not open-ended. This also limits the ability to use them in the service of humans.

Objective 2: Building curiosity-driven autotelic and aligned machines

Thus, our second major objective will be to address these limits by building and studying **curiosity-driven autotelic artificial agents that learn by interacting with external environments** and within socio-cultural collectives. To do this, we leverage and extend state-of-the-art deep reinforcement learning algorithms and transform them into **autotelic RL systems** (Objective 2.1). Also and crucially, we will study how algorithms for **autotelic learning can be made better aligned (teachable, driveable), more robust and more creative using language** both as a tool for social interaction, and as a cognitive tool to make abstractions and leveraging knowledge acquired by others (Objective 2.2). To achieve this, we will use pre-trained **generative AI models** as a cognitive tool bootstrap, enabling us to address the poor sample efficiency (i.e. require large amounts of environment interaction) and **poor generalisation** of classical (autotelic) RL algorithms. In addition, autotelic architectures will enable us to achieve incremental **grounding and alignment of generative AI models** with external physical and social dynamics.

To improve further abstraction, generalisation, we aim to establish links between autotelic learning and program synthesis techniques, whereby 1) autotelic generative models will self-improve their coding abilities by setting learnable coding problems of increasing complexity; 2) we will use code for autotelic procedural self-generation of environments, tasks and policies. Because of its expressivity and abstractness, we believe working in code spaces will open new perspective on open-endedness.

We also aim to study how autotelic algorithms, using the learning progress theory, will enable **frugal** adaptation of generative models thanks to **automatic curriculum learning**. Finally, we aim to study how groups of autotelic language-augmented agents can work in group and give rise to higher-order forms of **autotelic collective intelligence**, using multi-agent autotelic reinforcement learning techniques and measurement tools from the field of cultural evolution to track collective innovations (Objective 2.3).

Applications. Given the importance of curiosity in human learning and cultural evolution, fundamental progress in modelling and understanding human curiosity-driven autotelic learning on the one hand, and progress on building efficient artificial autotelic learning systems aligned

with human cultures, offer a high potential for applications of high human and societal values. We aim to exploit this potential and achieve impact in two particular domains of application.

Objective 3: Applications in education and assisted scientific discovery

Objective 3.1: Train curiosity-driven autotelic learning in humans across the lifespan

We aim to develop educational technologies and interventions **that help children and adults across the lifespan to learn in ways that are more motivating and more efficient**, for example by **stimulating curiosity and meta-cognition and using both models of curiosity and generative AI**. Our approach combines 1) an **interdisciplinary** perspective using both cognitive science, educational sciences and machine learning; 2) a **user-centric** approach with **real-world field studies**, in particular with real classrooms in the French educational system, or with field studies with adult of ageing populations; 3) consideration of both **neurotypical and neurodiverse populations**.

Beyond directly training curiosity and meta-cognition, and given their transversal role, we also aim to study how personalised training techniques (e.g. using adaptive curriculum with algorithms that maximize learning progress measures) can enable more efficient and more motivating training of disciplinary skills (e.g. maths, languages) and other cognitive dimensions (attention, working memory, etc).

Beyond showing effective impact in RCTs (randomized control trials), our objective is that our techniques and interventions be **used large-scale in the real world**. To achieve this, we will combine focused collaboration with the educational institutions and edTech industry, with user-centered design of open licence pedagogical material that will aim to be directly and easily reusable by teachers. We also aim to **help public action** in this domain, through interaction and advising with national and European public institutions.

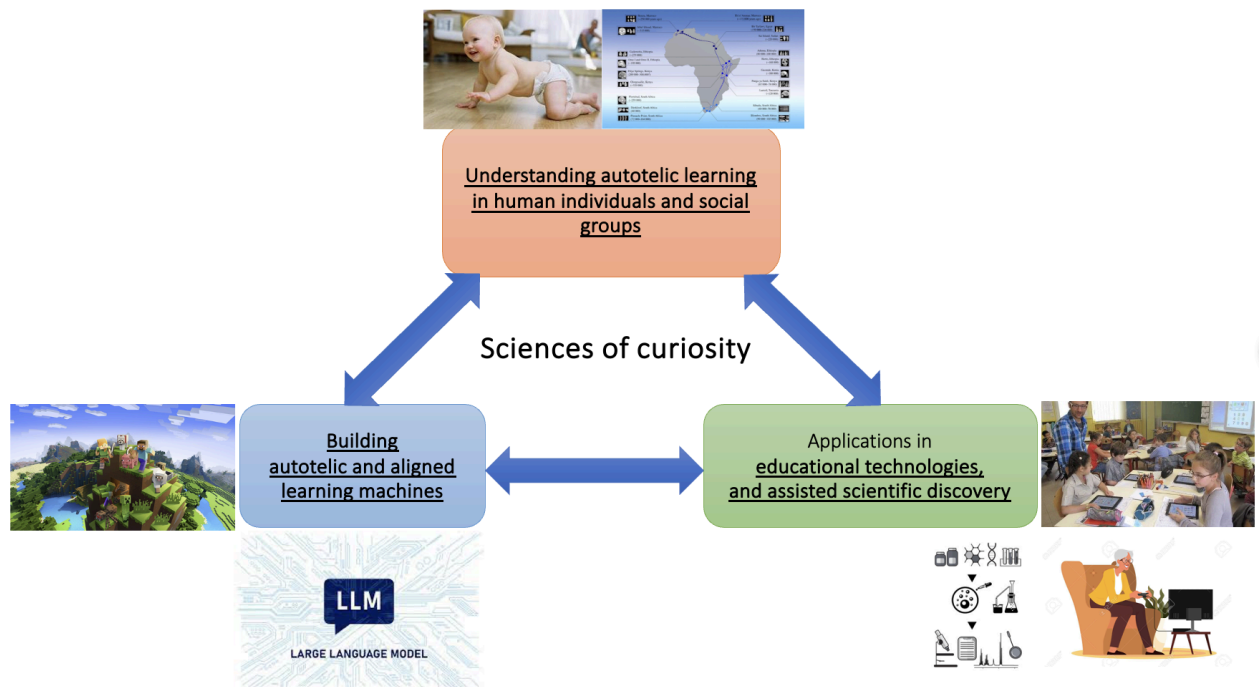
Objective 3.2: Assisted scientific discovery with autotelic exploration algorithms

Some of the greatest scientific challenges include the study and **design of novel materials, molecules or networks with complex dynamics**, where the space of possible self-organized behaviour is often initially mostly unknown, and the space of parameters very large, making exploration and discoveries **very costly and difficult for physicists, chemists or biologists**. We aim to study and show how **autotelic aligned exploration algorithms** can be used as powerful **discovery assistants** in these contexts. We believe they have specific capabilities making them highly relevant for this application: they are made to explore and discover in a sample efficient manner a high diversity of behaviours in complex systems (autotelic), while being driveable so that scientists can drive them in directions of interest

(aligned). To maximize diversity, we aim to develop methods learning a diversity of goal representations (autotelic and quality-diversity exploration using meta-diversity search). To enable abstractness and high-level guidance from human scientists (e.g. to provide feedback on measures of interestingness), we aim to leverage language and multimodal generative models.

To make **fast progress** in this direction, we first aim to use **artificial life environments**, such as continuous cellular automata, as an experimental domain, aiming to use autotelic exploration algorithms to help discover the **origins of autopoietic systems (and even autotelic systems self-organised from the ground up) as well as study how evolutionary processes themselves could self-organise.**

For further real world impact, we aim to develop collaborations with **physics/chemistry/biology academic labs**, as well as various **industrial companies** working on the design of new physical or biomolecular systems.

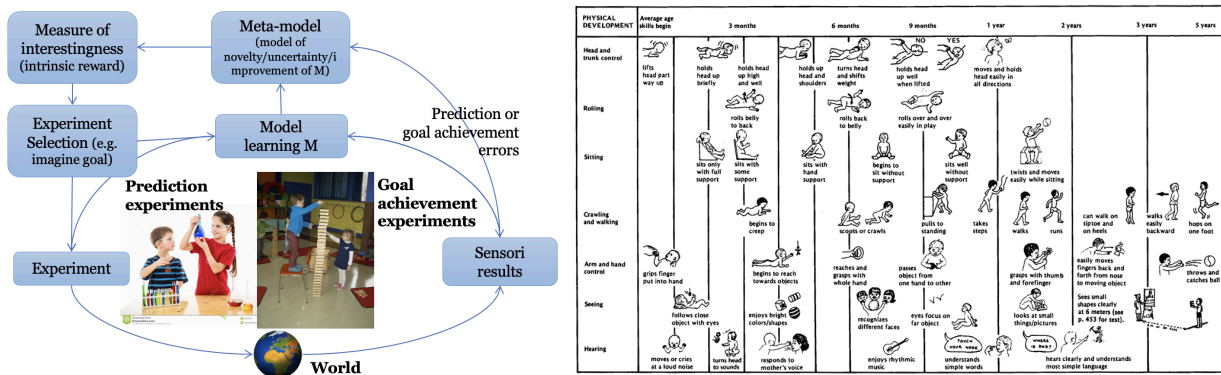


Overview of the main dimensions of our research program

2. Research Program: origins, evolution and new directions

The Flowers project-team was created in 2011, with an initial research program focused on computational modelling of curiosity-driven learning of sensorimotor skills in humans and machines, with experimentations and applications in developmental robotics. While modelling curiosity-driven learning in humans and machines has remained at the core of our research program across the years, and will remain so, the team's scientific strategy to study this topic, both in terms research focus and applications, significantly evolved, in particular since 2017 with several key new dimensions: 1) studying the role of language as a cognitive tool in curiosity-driven autotelic exploration; 2) studying the interaction and alignment between individual curiosity-driven learning and collective preferences and innovation (alignment); 3) testing theoretical predictions of our models in new human experiment paradigms using methods from experimental psychology; 4) connecting and integrating with modern deep learning methods, from deep RL to foundational models, and with the associated research communities; 5) shifting our application focus to educational technologies and assisted scientific discovery. This evolution was associated with a progressive change of the team structure and of our national and international collaborations. Our current work and research program is thus framed and explained below within this perspective, and our current activities are focused and organised along these novel directions started around 2017 (and fully in place around 2019-2020).

2.1 Origins of research program: Theories of curiosity-driven learning, sensorimotor development and developmental robotics (2011-2017)



Computational models of curiosity-driven learning account for the self-organisation of developmental trajectories in humans (Smith and Oudeyer, 2016).

Around the mid-20th century, psychologists started studying the hypothesis that humans, and some other animals, are endowed with mechanisms of **intrinsic motivation**, also called “curiosity” in everyday language, leading them to spontaneously explore novel activities for their

own sake. Such curiosity-driven exploration processes were hypothesised to play important roles in learning, both in cognitive and educational sciences: however, until the start of the 21st century, research for understanding of the underlying mechanisms was still very scarce. This also explains why **such mechanisms were overlooked in machine learning and robotics**.

In the first years of the 2000s, several labs in the world began studying these mechanisms through proposing various computational theories and hypotheses. Among these groups, Pierre-Yves Oudeyer and his colleagues, first at Sony CSL Paris and then at Inria Bordeaux, proposed several theoretical ideas and techniques to build some of the foundations of a new emerging field studying curiosity at the cross-roads of AI, machine learning, cognitive sciences, psychology and neuroscience. In particular, one major contribution has been the development of the **Learning Progress Hypothesis (LPH)**, proposing that human brains are intrinsically motivated to explore activities with high learning progress, leveraging meta-cognitive processes and leading to the self-organisation of efficient learning curricula (Kaplan and Oudeyer, 2007; Oudeyer et al., 2016). A second major contribution has been the development of a theoretical framework to account for **autotelic learning**, a form of learning where individuals learn to represent, sample and pursue their own goals (Oudeyer and Kaplan, 2007a; Colas et al., 2022b).

Based on several proof-of-concept studies of these computational theories (e.g. Oudeyer et al. 2007b), the Flowers team was founded in 2011 by Pierre-Yves Oudeyer (joining Inria) and David Filliat (Ensta ParisTech), with a research program aiming at scaling up these theories along two main dimensions: 1) showing how the LPH could account for key properties of sensorimotor in human infants; 2) showing how it was possible to develop curiosity-driven autotelic learning algorithms that would enable high-dimensional real world robots to acquire complex sensorimotor skills in a human-like way. Several major results were achieved in the 2011-2016 period along these lines.

Modelling curiosity-driven sensorimotor development in humans. First, the team developed a series of computational models implementing the LPH and studied how associated curiosity-driven exploration and learning mechanisms could self-organize developmental trajectories that reproduce structural regularity and diversity that are well known in humans, and yet lacking so far principled causal explanations. For example, the team developed a model of vocal development accounting for the main developmental vocal stages in the first year of life (Moulin-Frier et al., 2014). Then, the team developed extensions of these models that accounted for known yet previously unexplained developmental phenomena in tool use development during infancy (Forestier and Oudeyer, 2016). Also, the team showed how these same models could also account for how infants discover that vocalizations can be used as tools to influence social peers (Forestier and Oudeyer, 2017). This allowed us to propose a synthetic theory of how curiosity-driven learning self-organized developmental trajectories on the long term, and could explain the discovery and acquisition of language through mechanisms that are initially purely sensorimotor (Oudeyer and Smith, 2016).

A major goal and challenge for this work has been to have impact in the human sciences, and in particular in psychology and neuroscience. Leveraging the original and powerful theories and computational experiments we contributed, we also started to build strong collaborations with researchers in these domains, in particular with with Jacqueline Gottlieb (Cognitive Neuroscience lab, Columbia). This allowed us to publish foundational scientific papers in key journals of these domains, such as Trends in Cognitive Sciences (Gottlieb et al., 2014), Topics in Cognitive Science (Oudeyer and Smith, 2016) or Nature Reviews Neuroscience (Gottlieb and Oudeyer, 2018). These papers, presenting our synthetic theoretical computational perspective, played a significant role in structure the new emerging domains of the Sciences of Curiosity. Associated to this, we started organizing a series of international interdisciplinary conferences, starting with the Neurocuriosity symposium in Inria Bordeaux in 2014, which we have then held multiple times (2016 London, 2018 Penn Univ, 2023 Columbia Univ), together with many other workshops. This has played a key role at the international level for building and structuring this scientific field.

Application to developmental robotics. As these theories of curiosity-driven learning modeled how human infants explore and learn efficiently in the real world, a key target test for these theories had to be experimenting their efficiency in enabling to self-organize learning of skill repertoires in high-dimensional real world robots. This is what the team achieved in parallel during this period, associated to the design and refinement of algorithms of autotelic curiosity-driven learning, using sample efficient memory based architectures. This led to several key successes. First, we showed how such autotelic exploration algorithms could enable multi-articulated robots to learn inverse models of their own bodies, enabling efficient arm control or quadruped locomotion (Baranes et al., 2013). Then, the team showed how these algorithms could enable a robot to efficiently learn to manipulate complex and soft objects that were beyond reach of traditional control methods in robotics (Nguyen and Oudeyer, 2014). Finally, we showed how these algorithms could also be used to actively select learning sources, such as actively asking example demonstrations from several humans having diverse skill levels (Nguyen and Oudeyer, 2012). Finally, a major result was obtained using modular autotelic exploration algorithms, enabling humanoid robots to learn complex tool use skills from scratch in a couple dozen hours, leading us to obtain the 2nd best demo prize at Neurips 2016, several key conference papers in robotics venues (e.g. IROS 2016, IEEE ICDL-Epirob 2016), and a journal paper in JMLR (Forestier et al., 2022).

Starting to explore a new application domain: education. Beyond robotics, two different projects led us start exploring education as a new application area. First, after we developed the Poppy humanoid platform, which was the first world wide open-source humanoid robot (Lapeyre et al., 2013a,b), first as a tool to study the role of varying morphologies in skill acquisition, we discovered such a platform matched the needs of the growing educational robotics field, using robots as a micro-world to teach the basis of computer science, AI, but also arts at levels ranging from primary to high schools and universities. Driven by our complementary mission for popularizing science and interacting with the society at large, we built large scale educational robotics projects, such as the [Poppy Education](#) project. This led us to design and share both open-source hardware for educational robotics, but also educational content co-designed with a

large community of teachers and artists (> 1000 people on the forum we animated). This led to many projects both in the educational system, and in the associated world (e.g. several dance companies used our tools to build large audience shows), see (Lapeyre et al., 2014). While this community grew fast and had multiple forms of societal uses, we found continuing to develop it ourselves was beyond the mission of a research team. For this reason, this dynamics was transferred outside the team for further development: on the one hand, the [Pollen Robotics](#) company was founded as a spin-off (now still existing and having great international success), and on the other hand the [Poppy Station](#) NGO was founded as a vehicle for carrying associative educational and artistic projects.

On the other hand, we started to study, in a very exploratory manner, the use of our LPH computational theory of human curiosity-driven learning to design algorithms for personalizing learning curricula in intelligent tutoring systems, i.e. for personalizing sequences of exercises in educational software. This led us to develop the ZPDES algorithm within the KidLearn project (Clement et al., 2016), and start experimenting it for teaching mathematics in primary schools. We conducted a Random Clinical Trial study with > 1000 primary school children, showing the ZPDES algorithm enables both better efficiency in learning and in intrinsic motivation compared to a teaching curriculum designed by an expert in the didactics of mathematics. This work and experiment was foundational in our whole applicative research program in the domain of education (see below).

Team structure: During this period, the team structure followed the logic of our activities. The team was co-founded with the robotics team of David Filliat, and several robot learning researchers joined the team both at Inria and Ensta (Manuel Lopes, Freek Stulp, Alexander Gepperth, Anna-Lisa Vollmer). Work towards psychology and neuroscience was mostly achieved through the interdisciplinary work of Pierre-Yves Oudeyer, and through key external collaborations (e.g. J. Gottlieb). Work in education largely benefitted from the key recruitment of Didier Roy, who had previously been developing a career in educational institutions and took a crucial role both in the Poppy Education project, and in collaboration with M. Lopes and PY Oudeyer, in the KidLearn project.

2.2 Current research program (2017-ongoing): testing theories of curiosity-driven learning in humans, language augmented autotelic deep RL and generative models, and new application domains

Note: a detailed description of this research program is given in Appendix

While major successful outcomes were achieved in the 2011-2017 period, we identified several limits in our projects, leading us to set new objectives. These included in particular:

- **Testing the predictions of our computational theories of curiosity in human experiments**, and improving their articulation with other psychological constructs such as metacognition;

- Broaden the impact outside the developmental robotics community, and in particular **disseminate the ideas and algorithms of autotelic learning in the machine learning community**
- Understanding how to scale up computational theories of curiosity-driven learning towards more abstraction, more creativity, more robustness: this led us to develop a new framework of **autotelic learning articulated with language acquisition and collective/cultural dynamics**; This also led us to build **bridges with the emerging field of generative AI**, as well to the societally key question of alignment of AI technologies with human culture(s);
- Achieving larger-scale educational impact through **translational research in educational technologies**, and training directly curiosity and metacognition as key educational enablers;
- Exploring a **new application domain** for autotelic exploration algorithms: **automated scientific discovery**

Objective 1: Improve understanding of human autotelic and aligned intelligence

While our computational theory (LPH) enabled us to account for already known developmental structures, a limit was the lack of experimental tests of associated predictions on human learning and exploration. Yet, at the time we developed the LPH theory, no experimental paradigm in psychology and neuroscience existed to test its predictions. Through key collaborations with Jacqueline Gottlieb, then with H  l  ne Sauz  on as she joined the team, we started to work on the design of various novel experimental paradigms. This led to several key experimental results confirming major dimensions of the LP theory. For example, we showed that when human adults are free to explore several learning activities of varied learnability levels, their exploratory behaviour is best explained by utility models that include a LP component (Ten et al., 2021). This was confirmed in related experiments conducted in several other labs in the world, (e.g. Poli et al., 2022; Sayali et al., 2023). In further psychology experiments with humans, we characterised several links between curiosity and cognitive load (Poupard et al., 2022), and between curiosity and spatial memory (Sivashankar et al., 2024).

Crucially, we are now developing a conceptual articulation between our theories of autotelic curiosity-driven learning and meta-cognition, seen as a process supporting various stages of the curiosity cycle, as well as collect a first series of experimental results supporting this framework (Abdelghani et al., 2023c).

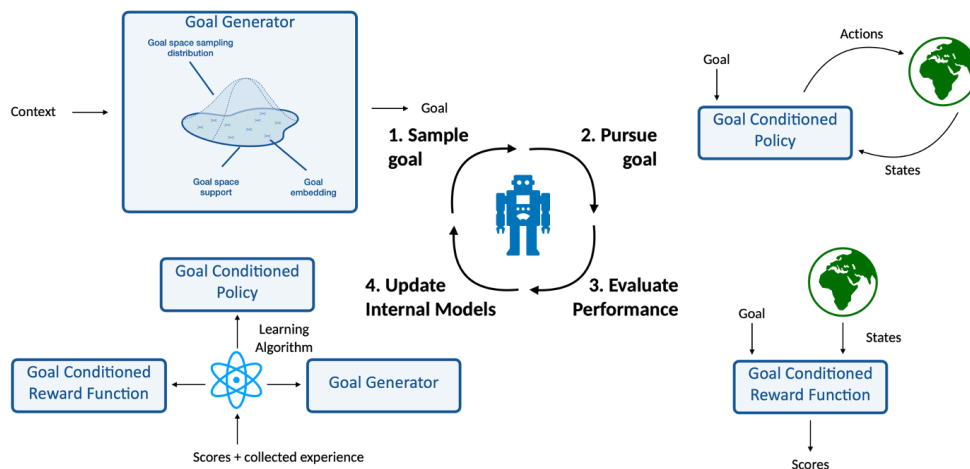
Objective 2: Building curiosity-driven autotelic and aligned machines

Objective 2.1: Connecting and disseminating autotelic learning to the machine learning community

While the team's work had great impact and influence in the developmental robotics community (e.g. in some years, around 30% of papers in the IEEE ICDL-Epirob conference discussed

curiosity-driven learning in robots and cited our works), the ideas and algorithms were still little known in the machine learning community up to 2016-17. However, our experimental results showed great potential of intrinsically-motivated learning algorithms, and in particular autotelic to address key unsolved problems in machine learning, in particular both RL problems with very sparse rewards, multi-goal RL problems. Further, some fundamental scientific questions we were exploring about open-ended learning were little considered at that time by the machine learning community, and we believed they were important for the future of that field. In reverse, our works did not use at that time strong technical advances in machine learning, in particular deep reinforcement learning methods for which impressive results were shown in 2016 (e.g. Mnih et al., 2015).

For all these reasons, the team made a strategic pivot: we started to focus our research and dissemination efforts towards ML communities gathered in conferences such as Neurips, ICML and ICLR. This was a challenge as at this point, we were both not experts in Deep RL methods, and the epistemology of developmental robotics was different from the habits of the ML community in terms of writing scientific papers. To address these challenges, we benefitted from welcoming Olivier Sigaud, one of the best Deep RL researcher in France, in 2017 as a detached senior researcher in the team to help us make this pivot successful. We quickly made substantial contributions, in particular developing algorithmic architectures first mixing autotelic memory-based and deep RL techniques to solve sparse reward problems (Colas et al., ICML 2018), then the UGL algorithm for learning goal representations (Pere et al., ICLR 2018) then the CURIOUS algorithm which was the first autotelic deep reinforcement learning system (Colas et al., 2019), as well as autotelic learning algorithms used for automatic curriculum training of Deep RL agents (Romac et al., ICML 2021), all published at ICML. These works were soon followed by a whole series of publications at Neurips, ICLR and ICML (e.g. Colas et al., Neurips 2020; Akakzia et al., ICLR 2021; Karch et al., Neurips 2021; Liu et al., ICML 2022; Lemesle et al., ICLR 2022; Barde et al., ICLR 2022; Carta et al., Neurips 2022; Carta et al., ICML 2023), some of them obtaining very selective oral accept (best 1% papers, e.g. Reinke et al, ICLR 2020; Etcheverry et al., Neurips 2020). This series of work was identified and recognized by a growing set of academic machine learning labs in the world, and strongly contributed to grow the importance of the research topic of intrinsically motivated machine learning: we recently wrote a review paper surveying this emerging field of curiosity-driven autotelic deep reinforcement learning (Colas et al., JAIR 2022), as well as published one of our key autotelic learning algorithm in JMLR (Forestier et al. , 2022). Several major AI labs in the world started to use this family of methods to solve sparse reward problems that were previously unsolvable (including labs at Stanford, UCL, Google Brain, Meta, Uber AI labs, see e.g. Bellemare et al. 2016; Ecoffet et al. 2019; LeCun 2022).



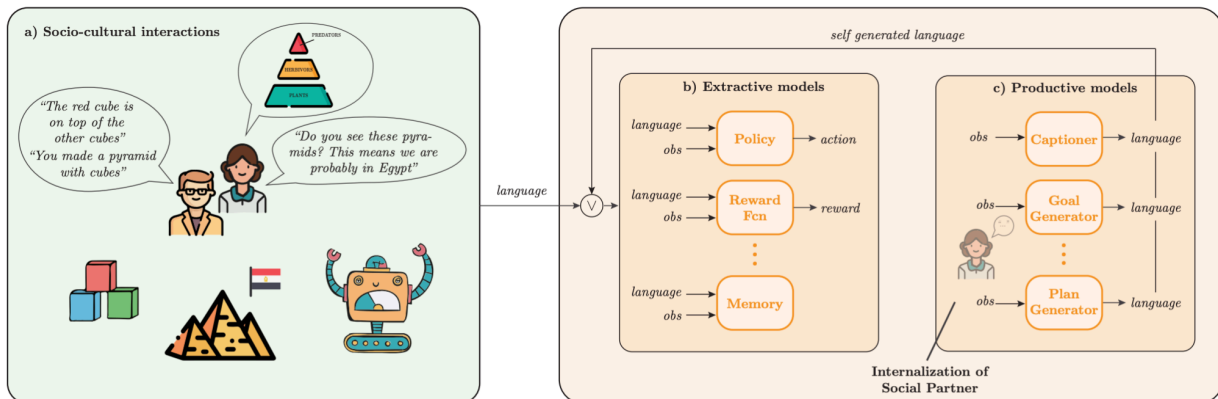
The autotelic exploration and learning loop (adapted from Karch, 2023)

Objective 2.2: Curiosity-driven autotelic AI agents leveraging large foundational models

One major limit of the generation of models of curiosity-driven autotelic learning, described above, is that they were mostly operating in low-level sensorimotor spaces, **limiting capacities to abstract and generalize**, thus limiting capabilities for **creative exploration**. Second, this generation of models of curiosity-driven exploration led to developmental **trajectories that were little influenced by social and cultural environments**: on the contrary, children’s development is highly influenced by their social environment, and from an application point of view, **it is desirable that open-ended AI systems develop in a way that is steerable by humans so that they become aligned and integrated within human culture**. This led us to study the role of language, and more generally collective and cultural intelligence, in sculpting open-ended learning.

In particular, we started exploring and modeling the idea that, beyond its communicative nature, **language is internalised by children to become a cognitive tool empowering abstract reasoning, planning and imagination**, such as argued by Vygotsky (1934) and Bruner (1990). This idea is associated to the question of how **social peers scaffold learning environments of children into a zone of proximal development**, where activities are just beyond the current capacities of the child, enabling efficient curriculum learning trajectories, and guidance of autotelic exploration by the social (and educational) environment. Beyond words and syntax, Jerome Bruner also showed how **language conveys knowledge and thought heuristics**, for example embedded in stories children are exposed to, also internalised to boost individual cognition, as well as sculpt internal systems of motivation and values.

Concretely, one foundational work in the team on this topic has been the **IMAGINE language-grounded autotelic architecture** (Colas et al., Neurips 2020). In this work, an autotelic deep reinforcement learning agent could first learn language as a communication tool through interaction with a simulated social peer, and then **reuse the learned language-behaviour compositional model to self-generate creative new goals** and self-compute goal achievement rewards (Colas et al., 2020). This effectively led to an agent capable of expanding autonomously a diversity of abstract skills that are relevant to and steerable by a social peer. A limit of this proof-of-concept work was that the social peer, and in particular its **language system, were simplistic, far from the complexity of human language**, both in terms of language complexity and in terms of the culture it conveyed. To address this challenge, we started **leveraging large language models and use them as culture models** to build simulate social and linguistic environments with human-like complexity (Colas et al., 2022b). We also started studying how **LLMs can be used directly as cognitive tools** in autotelic agents, helping them to imagine new goals and plan towards them (Colas et al., 2022b, 2023; Pourcel et al., 2024). We also identified an approach where such use of LLMs in autotelic agents could enable agents to **align and ground LLMs in the dynamics of external environments, thus addressing one of the main current limitations of LLMs** (and more generally of multimodal foundational models).



After being acquired as a tool for communication, language can be internalized and become a cognitive tool for autotelic agents: it enables to imagine creative and abstract goals, generate high-level plan hypotheses, re-describe observed behaviours and self-generate internal goal-achievement rewards (adapted from Colas et al., 2022; see also Colas, 2021)

Objective 2.3: Cultural evolution, innovation and curiosity at the group level

Beyond this work studying how language and social dynamics can guide and align autotelic agents, we also started studying how groups of autotelic agents could interact and self-organize forms of collective intelligence. To do this, we considered the larger scientific picture of **cultural evolution** (Henrich, 2016; Derex and Mesoudi, 2020), showing how groups of individuals can display **collective abilities to explore, discover**, vary and transmit new skills and new knowledge, in effect showing forms of **curiosity-driven exploration and innovation at the**

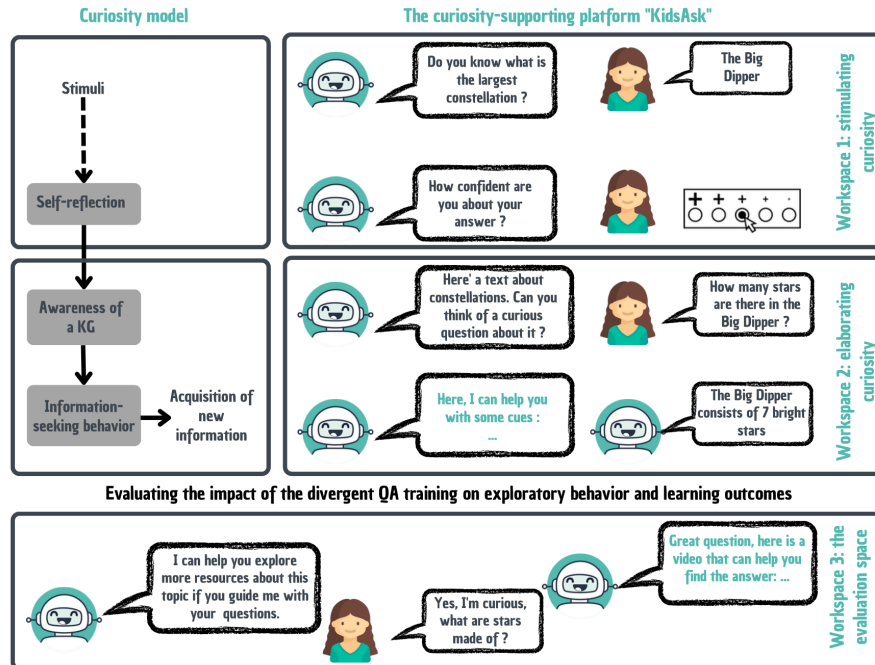
group level and over cultural evolution scales. In this perspective, and building on the expertise of Clément Moulin-Frier and Pierre-Yves Oudeyer in computational modelling of language evolution (both did their PhDs on this topic, Oudeyer, 2003; Moulin-Frier, 2011), the team started developing **models of the dynamics of exploration and innovation in groups of culturally interacting autotelic agents**. This includes understanding how the topology of social networks impact innovation (Nisioti et al., 2022), the formation of language communication protocols in groups of agents (Barde et al., 2022), and how groups of autotelic agents can learn to coordinate in their exploration of the world (Nisioti et al., 2023b). **A key open question here is understanding how curiosity and autotelic learning at the level of individuals interacts with exploration and innovation processes at the group level**. We also connected this line of work on modeling collective dynamics with recent innovations in generative AI. Given the growing impact of generative AI in our everyday life, and in particular their role in helping humans produce new cultural artefacts, it has also become crucial to **understand how the transmission dynamics happen in systems of generative AI agents, or in hybrid AI-human groups**: our research aims to contribute to this new scientific field that has been called “**Machine Culture**” (Brinkmann et al., 2023). In this direction, we started investigating how LLMs can encode superpositions of socio-cultural perspectives, for example value systems, and developed tools to analyze the robustness and controllability for the expression of these socio-cultural abilities (Kovac et al., 2024).

This recent line of approach of the team forms a backbone of our current research program, both to build artificial systems and to model language-augmented autotelic learning in humans and its interaction with collective dynamics.

Objective 3: Applications in education and assisted scientific discovery

Objective 3.1: Translational educational technology research: training curiosity and metacognition

Based on the very encouraging results of our exploratory projects in educational technologies in the 2014-2017 period, and in particular the KidLearn project and the ZPDES algorithms (see above), we aimed to scale up this work both scientifically and in terms of **societal, institutional and industrial impact**, and targeting to **help children and adults across the lifespan to learn in ways that are more motivating and more efficient**. A key property of projects we developed in this period, and we aim to continue in this new proposal, is the combination of 1) an **interdisciplinary** perspective using both cognitive science, educational sciences and AI; 2) a **user-centric** approach with **real-world field studies**, in particular with real classrooms in the French educational system, or with field studies with adult of ageing populations; 3) consideration of both **neurotypical and neurodiverse populations**. Part of this work is covered by the 2022-2025 research program of the associated "Curious tech" team with the University of Waterloo in Canada (E. Law from CHI Lab and M. Fernandes from The Cognitive neuroscience lab) (<https://flowers.inria.fr/curioustech-associate-team/>)



In the KidsAsk platform, we use conversational agents to train children to ask curiosity-driven questions, as well as related meta-cognitive skills. We also study how generative AI techniques can be used to generate automatically efficient cues in this training process. Such educational interventions are evaluated in the field, in primary schools (Abdelghani et al., 2022, 2023b)

From the scientific point of view, we identified the need and challenge to develop educational interventions that directly train curiosity, and its associated metacognitive skills, as a converging set of works shows they are both central for academic success and well-being of diverse populations of learners, and yet are currently outside the scope of educational methods in the school context. As first steps in this direction, we leveraged our theoretical perspective on autotelic learning to design and evaluate educational interventions that foster curious question (Abdelghani et al., 2022), as well as meta-cognitive skills associated to the curiosity cycle (Abdelghani et al., 2023c) in primary schools. Thanks to strategic partnership with Académie de Nouvelle Aquitaine, we run a series of field studies showing significant impact of our interventions on the targeted skills. We also started to study how the ZPDES personalization algorithm can be used to train maths skills in children with ASD (Mazon et al., 2023), and to train attention and other related transverse cognitive skills such as working memory and executive functions (Adolphe et al., 2023), with encouraging results for both young and aging populations.

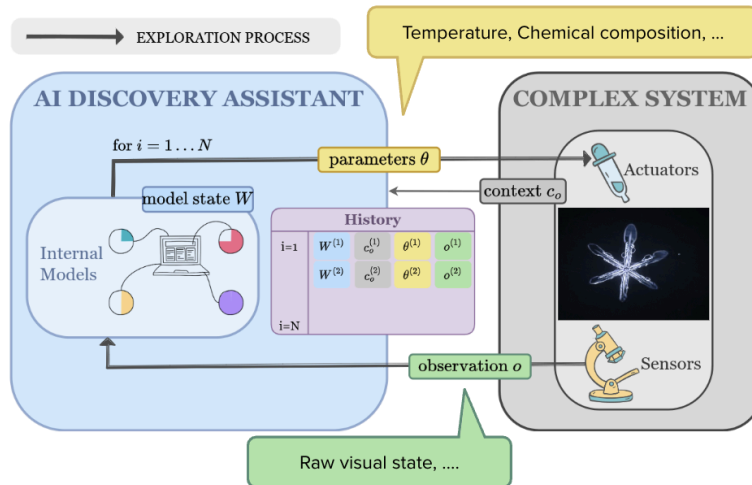
Given the rising impact of generative AI in education, with high societal and ethical stakes, we also started to study the limits, dangers and opportunities provided by these new tools. First, we published one of the first international field studies investigating the pedagogical use of LLMs (here GPT-3) in real classrooms: in Abdelghani et al. (2023) (collab. with E. Yuan and T. Wang from MSR Montreal, and with Y-H. Wang). We showed that when used appropriately, GPT-3 can be used to build conversational agents that train efficiently curious question asking in primary

school children, enabling to scale up a pedagogical approach for training curiosity skills we developed previously (Abdelghani et al., 2022). We also developed a conceptual framework to think about the opportunities and challenges associated with using generative AI in the classroom, and in particular asking how this could be done by enabling children to keep and develop active learning skills (Abdelghani, Sauz on & Oudeyer, 2023). We claim that one key challenge is to **improve the AI literacy of both children and their teachers**: with this aim in mind, one objective of our project is to make contributions aiming to boost AI literacy for diverse audiences, and children in particular, for example by building pedagogical materials (videos, intelligent tutoring systems aimed at AI literacy).

From the industrial point of view, we developed a strategic collaboration with the EvidenceB edTech company: with the support from the French Ministry of Education, the ZPDES personalization system was deployed in the large-scale AdaptivMaths system now available in all French primary schools (> 68k classrooms). EvidenceB further used ZPDES in the “MIA seconde” system aimed for training high-school students in maths and french. One challenge of this line of work is the empirical adoption of these technologies by teachers, on which we work through co-design with teachers and through pedagogical material to explain teachers the how and why of these technologies. We also characterised various limits of emerging pure LLM-based artificial tutors (e.g. the Khanmigo system from Khan Academy), both in terms of pedagogical relevance and in terms of energetic cost. We are currently working with the EvidenceB company to study how generative AI could rather be used in a frugal and pedagogically relevant manner to pre-generate pedagogical content (exercices, hints, explanations, feedback) to feed the graph content in intelligent tutoring systems on which algorithms like ZPDES operate. We are also cooperating with the CATIE company to develop XR (extended realities) based educational technologies with a sustainable cognitive load (Poupard et al., 2024, collab with the Catie compagnie), as well as exploring the use of personalization algorithms to onboard video game players with the Ubisoft company.

Societal, institutional and public policies. Given the very high societal importance of these topics, we have started a number of actions where we directly discuss and interact with public policy representative in order to explain them recent scientific insights on these educational issues, and advice them on the analysis of the situation in France and other countries, as well as on which actions shall be prioritized (e.g. interactions with Conseil Scientifique de l’Education Nationale, with various DRANEs department in Academies, with Senat or OPCST commissions, etc). Associated to these actions, we conduct also several projects to improve AI literacy in middle and high schools (e.g. the design of the “[ChatGPT expliqu  en 5mn](#)” video series), and participate to many popular science actions (e.g. writing popular science books for adolescents).

Objective 3.2: Autotelic algorithms for assisted scientific discovery



Curiosity-driven AI discovery assistants can enable chemists, physicists or biologists to explore efficiently the landscape of self-organized structures in complex systems (figure adapted from Etcheverry, 2023)

Curiosity-driven exploration for assisted scientific discovery is another major application of our fundamental research which we started to explore in this recent period. We showed how curiosity-driven exploration algorithms can enable scientists (physicists, chemists, biologists, etc) **to explore and map the space of self-organised behaviours in diverse complex systems** (Etcheverry et al., 2020; Etcheverry, 2024), especially in cases where scientists do not already have a model of the studied complex system, when they do not know the space of possible behaviour, and when it is costly to make experiments. For example, through a collaboration with M. Levin at Tufts University, we studied how autotelic AI systems (IMGEP algorithms) can enable cost effective discovery of diverse sophisticated and robust behaviours in gene regulatory networks (Etcheverry et al., 2024). We also studied how these exploration algorithms can enable novel discoveries of **self-organised structures in continuous CAs like Lenia**, such as forms of primitive agency (Hamon et al., 2024), making contributions to the domain of artificial life, such as our Flow Lenia system which received the Best Paper Award at Alife 2023 (Plantec et al., 2023). Notably, our computational framework for curiosity-driven assisted scientific discovery has been independently reused by external researchers in physics at Univ. Washington (Falk et al., 2022), with whom we are exploring new collaborations.

Summary: In the 2017-present period, we have operated a strategic scientific and applicative pivot: while keeping curiosity-driven autotelic learning in humans and machines as our core research activity, we 1) started projects testing our theoretical predictions in human psychology experiments, and articulated links between curiosity and metacognition; 2) Integrated modern Deep RL techniques with autotelic algorithms, and shifted from the developmental robotics to the machine learning community as target of our contributions to the design of more open, flexible and robust learning machines; 3) Shifted from sensorimotor autotelic learning to language-based abstract yet grounded autotelic learning, and built synergetic bridges with

recent advances in generative AI; 4) Scaled up our research in educational technology by taking a translational approach and developing industrial collaborations, with actions to support public policies; 5) Started the new application domain of automated scientific discovery. These constitute the pillars of our current research program (and we have already funded projects on these dimensions already in place for the next 4-5 years).

A **detailed description** of this ongoing **research program** is provided in the **Appendix**.

Evolution of team structure

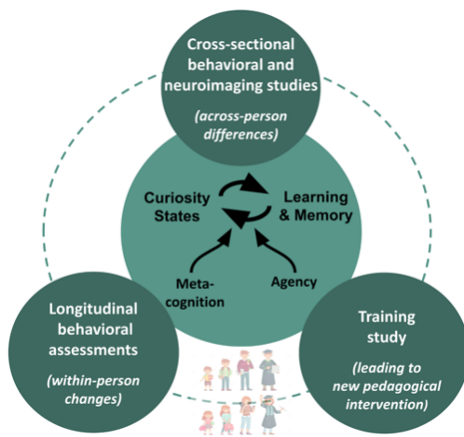
Given this evolution of techniques and topics in this period, the team structure evolved significantly. First, H  l  ne Sauz  on and C  cile Mazon, Professor and Assistant Professor in psychology/cognitive sciences at University (and responsible of the Cognitive Science master in Univ. Bordeaux), joined the team (2020) as permanent members, providing great **expertise in methodologies from experimental psychology** (both for fundamental and educational technology research). H  l  ne Sauz  on is now detached at Inria (DRCE). In parallel, Cl  ment Moulin-Frier joined the team as an Inria permanent researcher (CRCN, 2019), with a strong background in computational modelling of collective and social dynamics in groups of autotelic agents, and in particular on language and cultural evolution. In machine learning, Olivier Sigaud, a key figure of Deep RL research in France, joined the team for 2 years in 2017 (detachment). In parallel, during this period, the collaborations with the Paris part of the team at ENSTA Paris, focused on robotics, progressively decreased.

Beyond permanent members, we have also developed very strong collaborations with several external individual researchers. First, in the context of Pierre-Yves Oudeyer's sabbatical year at Microsoft Research Montreal, which allowed acquiring state-of-the-art knowledge in generative AI and developing a network of collaboration with the Montreal AI ecosystem, several collaborations with individual Microsoft Research scientists were further developed, including Ida Mommenejab (MSR NY, computational neuroscience and computational models of collective dynamics), Marc-Alexandre C  t   and Eric Yuan (LLMs and Deep RL). In France, we also developed strong collaborations with Deep NLP specialists Sylvain Lamprier and Laure Soulier at Sorbonne Universit  . In the domain of automated scientific discovery, we develop strong collaborations with Bert Chan (external member of the team and Google Deepmind scientist, on artificial life projects), Michael Levin (Head of the synthetic biology lab at Univ. Tufts/Harvard). In educational technologies, we have a strong collaboration with Daphne Bavellier (Univ. Geneva) and Benjamin Cl  ment at the EvidenceB company.

2.3 The horizon: research and applications in 10 years

While the previous sections presented an overview of our ongoing research program (see also our detailed research program in Appendix), there are already several further long term objectives and topics on the horizon. Some of these projects are already being operationalized (e.g. in grant proposals and in starting collaborations), and some others are more speculative

and exploratory. On the path towards these long term (autotelic at the team level :) goals, we may discover new directions through serendipity and curiosity.



Centre circle: We hypothesise that both metacognition and agency impact the development of elicitation of curiosity states, how curiosity-driven information search is organised, and the efficiency of curiosity to enhance learning and memory. **Outer circles:** To study various dimensions of this hypothesis, we will conduct cross-sectional behavioural and neuroimaging studies using complementary experimental approaches, a longitudinal behavioural study with multiple assessments to understand within-person changes, and a training study which has the potential to lead to new pedagogical interventions to stimulate curiosity in the classroom.

Autotelic learning in human individuals and social groups: development, cross-cultural, neurodiversity (Objective 1)

Our recent and ongoing projects include the design of novel human experimental paradigms to study various dimensions of autotelic learning and exploration, either in individual or in groups. We believe these will already address several very important yet open scientific questions. However, further in the horizon, we identified three related objectives leveraging these new experimental paradigms. First, one key open scientific questions is how properties of curiosity-driven exploration (e.g. such as measured in the “Monster Task”, Ten et al., 2021), divergent thinking and metacognition (e.g. such as measured in tools used in Abdelghani et al., 2023c), develop and correlate from childhood, to adolescence, to adulthood and in ageing populations. No existing study exists addressing such fundamental and important questions. This is the core of a large-scale project proposal H  l  ne Sauz  on and Pierre-Yves Oudeyer recently wrote in collaboration with Matthias Gruber (Univ. Cardiff) and Yana Fandakova (Univ. Trier), combining psychology and neuroscience imaging techniques with computational modelling tools. In a related line, we also aim to address a related open question: how curiosity and metacognition, as measured with these tools, develop and correlate in neurodiverse populations (e.g. along the autistic spectrum or along attention disorders). Finally, most existing studies of curiosity in the world have been conducted in western or asiatic academic labs, and associated human subject populations: it is still a largely open question whether results observed would also apply in very different cultural populations. Leveraging a starting collaboration with Maxime Derex (IAST Toulouse), who is a specialist of the study of cultural evolution and innovation, we are setting up projects to do field studies in populations like Tsimane in south america. We plan to scale up such cross-cultural studies in the coming years, leveraging collaborations with labs studying diverse populations from diverse cultural backgrounds.

Methods from psychology and neuroscience to study human curiosity and metacognition

To achieve these aims, we will combine different **methodological approaches from psychology, neuroscience and computational modeling**, including behavioral experiments and functional neuroimaging (done in our collaborators labs, e.g. at Mathias Gruber and Yana Fandakova's labs, see Fandakova and Gruber, 2021), training studies, and longitudinal assessments to investigate how across-person differences and within-person changes in metacognition and agency contribute to curiosity-based learning across late childhood and adolescence, and across neurodiversity and cultures. We will **leverage established and newly developed paradigms**, e.g. the Trivia paradigm (Fandakova and Gruber, 2021), the Monster paradigm (Ten et al., 2021), a new **Virtual Reality (VR) paradigm** (Sivashankar et al., 2024), the curious question generation paradigm (Abdelghani et al., 2022); the curious question training paradigm (Abdelghani et al., 2023), for in-depth analysis of how exploratory behaviors lead to curiosity-based learning across age. Furthermore, we will use **large language models (LLMs) for qualitative analyses of think-aloud components** (Xiao et al., 2023) in the experiments to characterize verbal responses associated with metacognitive processes in curiosity-based learning.

In addition to behavioural and neuroimaging measures used to assess state curiosity and intrinsic motivation, we will validate questionnaires from experimental psychology: **Vallerand's types of motivation scale** (Vallerand et al. 1989) that probes intrinsic and extrinsic motivational mechanisms in educational settings. Trait curiosity across experiments will be tested using the **Interest and Deprivation Type Epistemic Curiosity Scale** (Piotrowski et al., 2014) which distinguishes between a desire for new information anticipated to increase pleasurable feelings of situational interest (I-type) and a motive to reduce unpleasant experiences of feeling deprived of knowledge (D-type). Finally, metacognitive skills related to curiosity will be examined with the **Metacognitive Strategies Questionnaire for Learning** (MSQL, Pintrich, 1991), which taps into broad areas of goal orientation, control beliefs and affect, learning, metacognitive and resource management strategies. We will use the Children's Images of and Attitudes towards Curiosity (CIAC, Post et al., 2015), which assesses individuals' perceptions of curiosity. Across studies, pubertal status will be assessed using the Puberty Developmental Scale (PDS, Petersen et al., 1988).

Autotelic and aligned machine learning: program synthesis for abstract, grounded, generative and frugal system 2 intelligence (Objective 2)

There are currently several major limits in machine learning systems, in particular large generative models: they lack abstraction capabilities enabling to adapt fast to very new situations (e.g. create hypotheses or cognitive tools on the fly); they lack metacognitive capabilities (related to System 2 cognition, Kahneman, 2018); they lack grounding and they are

very costly in terms of energy and data. Our ongoing research project started to address grounding and frugality, as autotelic agents self-organize a learning curriculum driving them to progressively learn interaction skills of increasing complexity, with self-collection of high quality grounded data. One first step towards **abstraction and generalization** is the use of language as a cognitive tool in autotelic agents. However, to go further, we are already starting to explore the use of formal languages, such as **programs or maths formalisms**, as cognitive tools to build verifiable hypotheses or theories, and build policies with strong generalization properties. In short, we aim to build autotelic system that will represent goals as programs, and search for policy solutions also expressed as programs. Our approach will consist in combining autotelic exploration strategies with generative AI techniques tailored for automated **program synthesis**. A first step in this direction is our ACES project (Pourcel et al., 2024), whereby we are working on the design of autotelic generative models that self-improve their coding skills through iterative self-generation of diverse coding problems and associated code solutions. Further in our plans, we will explore how such approaches could help make progress on the ARC benchmark, which consists of visually simple yet abstract puzzles that are easy for humans but still very difficult for even the most powerful LLMs (Chollet, 2019). Further down this road, we also aim to **explore how autotelic algorithms could enable self-improvement in formal maths**, through iterative generation and solution search of maths problems, checked using automated proof systems like Coq or Lean. This work will involve strong collaboration within the new **LLM4Code Inria Défi**, in particular with **Nathanel Fijalkow (Labri)**, expert in program synthesis, with **Xavier Hinault (Mnemosyne)**, and with **Guillaume Baudard and Mathieu Acher (coordinators of LLM4Code)** and their teams. We will also leverage collaborations with the **MIT lab of Josh Tenenbaum** on probabilistic program synthesis and its synergies with generative models (**Cédric Colas** is currently doing a Marie Curie postdoc prior to coming back to the team).

In complement, we will also further develop our approaches to **improve metacognitive skills** in autotelic generative agents. This will include first a systematic study of various dimensions of metacognition in generative models using measurement methods from **experimental psychology** (Fleming and Lau, 2014). Then, we will focus on building human-cognition inspired cognitive architectures that combine generative AI components within agentic architectures that dynamically learn to represent and predict their own capabilities (e.g. estimating learning progress about tasks expressed in natural language or in programming languages). This will involve a collaboration with the **human reinforcement learning lab of Stefano Palminteri (ENS Ulm)**.

As these techniques reach a level that enable both abstract and frugal learning and generalisation, one major target will be studying how they can be systematically used in contexts that require rich real-time interaction and learning with humans. This will be highly relevant in general for education and alignment of machines towards human preferences/values/ethics, but also more specifically useful to enable our systems to learn powerful individualized tutoring strategies in educational technologies, and to enable fast and abstract interactive exploration and discovery in the sciences (biology, chemistry, physics, maths, computer science, ...).

Generalized training of curiosity, metacognition and creativity across the lifespan (Objective 3.1)

Our long term objective is to design educational technologies for training curiosity, meta-cognition and creativity across the lifespan, leveraging our advances in fundamental understanding of the development of these skills in humans (Objective 1) and our computational AI models of these mechanisms (Objective 2). One key objective is to address the societal challenges inherent to the School of the 21st Century, in particular the inclusive school (learners with disabilities, especially cognitive); cross-disciplinary skills such as attention, curiosity, learning to learn, perseverance.

Building on our current projects, we will aim for these targets along several dimensions. First, we will study how to **extend the ZPDES approach** (Clément et al., 2015, 2024) **to new domains, in particular cognitive training of attention and meta-cognitive control**. Also, we will aim to extend the evaluation of this approach when targeted to diverse populations in terms of age (including ageing populations thanks the collaboration with the national network of free-time universities and with the newly created Vascular Brain Health Institute - VBHI - at Bordeaux), cognitive capabilities (e.g. neurotypical or neurodiverse thanks collaboration with Bordeaux hospital but also with the [atypiefriendly](#) network for inclusive universities in France), or occupations (e.g. training attentional skills in professional sport players; or training new video game players to enter more easily complex multi-player games, through collaboration with Ubisoft). From a technical point of view, we also aim to automate a large part of the manual work currently needed to design pedagogical graphs of exercices on which ZPDES-like algorithms are functioning: we aim to **use large foundational models for pre-generation of pedagogical content** that complies with didactics specification (we just wrote a large project proposal on this topic with the EvidenceB company). We will also explore the use of small language models (for frugality) to provide personalized feedback, hints and explanations for exercices, leveraging meta-cognitive strategies.

We will also scale up work we began to train metacognitive skills and curious question asking in children, going further the proof-of-concept by framing it within the LP theory and experimentally studying the role of language in these meta-cognitive training interventions (in synergy with the cognitive science study described above). Then, we will extend the study of the impact of such metacognitive training to a set of creativity abilities, as measured by standard psychological tools in psychology (e.g. Said-Metwaly et al. 2017). We will also consider using such training, and measuring impact on creativity, for various adult populations.

Teacher training to instil curiosity in classrooms, and then in schools: we will also consider how to train teachers (as proximal social support) so that they learn adequate strategies to foster curiosity, metacognition and creativity in the classroom, through an appropriation of our methods, particularly in the context of STEM (science, technology, engineering, and mathematics) subjects. This challenge is particularly critical as it has been shown that the teacher has a critical role to play in shaping and spreading curiosity-based learning (Peterson, 2020). So we're planning studies to evaluate curiosity training for teachers, so that they can use

our educational technologies (KidsAsk, KidsReflect), as well as unplugged versions of these activities for wider accessibility in schools, and then compare the results in terms of the impact on children's learning and curiosity.

Stimulating social curiosity in educational contexts with neurodiverse populations.

Neurodiverse populations face societal barriers due to their atypical cognitive and behavioural profiles. In edTech, traditional approaches for inclusion focus on matching technologies to individual needs, ignoring the broader (social) environment as a support for inclusion. However, family-professional partnerships are essential for supporting and improving children's academic performance and social adaptation. Thus, in an original way, we are developing educational programs aimed at improving the culture and social perceptions of neurodiversity of the neurodiverse student's entourage, in order to create conditions for social interactions that are favourable for everyone. (Mazon et al., 2023; Saint-supery, et al., 2023). We aim to design and to assess new technology promoting social curiosity across stakeholders, fostering dialog as a lever of success for a shared, enriched and powerful social representation of the neurodiverse student, with as a result an improved well-being for all stakeholders (student, parents, teachers and caregivers).

Generative AI to evaluate edTech systems and human learning/training outcomes.

A key dimension of societal application like education is that they require strong and systematic evaluations. One dimension of these evaluations consists in analysing verbal productions from students and other stakeholders, e.g. think aloud reports during experimentations, responses to psychology questionnaires, and interviews. To analyse such kind of data, qualitative analysis, and deducting coding methods in particular, are the reference methods. However, they require a large amount of expert time, which limits the possibility to do qualitative analysis and deductive coding at a large scale, thus limiting evaluation. To address this limit, we will build on our preliminary work exploring the use of generative AI tools to automate deductive coding and assess this automation to ensure reliability in the interpretation of results (Xiao et al., 2023). Because of the high stakes of such methods in psychology and the human sciences at large, we will aim to release open-source tutorials enabling researchers in these disciplines to reuse and leverage these methods.

This applicative research program leverages various collaborations with stake-holders, ranging from public educational institutions (e.g. Académie de Bordeaux) to industry (e.g. EvidenceB, OnePoint, Ubisoft, Catie) and NGOs (e.g. France IOI).

Autotelic and aligned assisted scientific discovery: towards artificial labs for assisting the scientific community (Objective 3.2)

Building on our recent work, a major mid and long term objective is to adapt and experiment our autotelic exploration algorithms for assisted scientific discovery in real world applications related to both the design of real new physical materials, and new biological systems. Indeed, we believe this line of research, beyond helping biologists/physicists/chemists in academic labs, has high potential for industrial application: we have for example started discussion with the Solvay/Sysenqo company, and the LOF in Bordeaux, to establish collaborations on the use of

autotelic exploration algorithms in material and chemical design processes. We are also considering new collaboration with companies at the cross-roads of biology and machine learning/generative AI, to use autotelic exploration algorithms to help map the space of biomolecular systems and design new molecules with targeted properties.

Towards teams of autotelic agents assisting teams of human scientists. Beyond these applications, we also aim to integrate the collective intelligence perspective in autotelic exploration for assisted scientific discovery. More precisely, we aim to study how a group of autotelic agents, with particular social topologies for fostering innovation (Nisioti et al., 2022), or particular specialities to enable meta-diversity search (Etcheverry et al., 2020) at the group level, could form a kind of artificial scientific team, with the aim to help a team of human scientists. One challenge of particular interest will be studying how to balance creative exploration of this group of autotelic agents, and alignment with the preferences and interest of human scientists. To achieve this, we will work on extending simple alignment method we introduced in (Etcheverry et al., 2020), to methods leveraging language and vision using large multimodal foundational models such as Llava (Liu et al., 2024). For example, we will study how human scientists could express their intuitions, interests and preferences using natural language over samples of discoveries, and how this may be transposed into multimodal transformer representations that can drive automatically a group of autotelic agents to explore corresponding phenomena in the targeted physical/chemical/biological system.

3. Software

We have been approaching software development with the aim to facilitate dissemination, impact and reproducibility of our scientific and technological contributions, and will continue to do so. Thus, the vast majority of the software we develop to conduct our scientific projects are shared open-source within MIT-style licences on our github web site: <https://github.com/flowersteam>

There are two broad categories of software we are working on. First, most of our papers are accompanied with source code implementing both the core algorithmic contributions as well as scripts enabling others to reproduce the main experimental results in the papers. These pieces of code leverage very standard machine learning libraries (e.g. Jax, Pytorch or BabyAI) and consist of small self-contained new pieces of code. This fosters reusability and reproducibility on the mid and long term (avoiding to rely on more complex but non mainstream libraries that may not be well maintained).

There are also several projects for which we are and will continue developing more involved software libraries that aim to be reused both within and outside the team to facilitate construction and development of novel large scale experiments. Three examples of such libraries being currently developed are:

1. **The Lamorel library** (<https://github.com/flowersteam/lamorel>): Lamorel is a Python library designed for people eager to use Large Language Models (LLMs) in interactive environments (e.g. RL setups). It has been developed to enable various researchers in

the team to implement and conduct large-scale experiments combining LLMs and Deep RL on the Jean Zay platform, solving a number of non trivial engineering challenges. It is also available for the general research community where researchers in either LLMs or Deep RL could benefit from it to design large scale experiments in these domains.

2. **The adtool library** (https://github.com/flowersteam/adtool_legacy) : adtool is a software package developed in the Inria FLOWERS research team which provides an integrated solution for studying complex systems through curiosity-search methods, consisting of a user-friendly Web UI and an extensible Python library for user-defined experimentation systems and search algorithms. We aim to use this library as a tool for disseminating our exploration algorithms in non computer scientists population, such as chemistry, physics and biology labs.
3. A new python library aiming to facilitate the **use of foundational models in the human sciences for automatic or assisted evaluation of human verbal and behavioural data**, i.e. to enable large scale qualitative analysis (e.g. annotations) following precise codebooks. This relates to the LLM-as-Judge approach used in machine learning, but here with a specific target of enabling the scientifically solid and easy use of the approach in human sciences (and fostering good methodological practices). We already made proofs-of-concept of this approach in (Xiao et al., 2023; Abdelghani et al., 2023).

4. Impact in science, public policy, industry and society

4.1 Scientific impact

Beyond interdisciplinarity: Building a new science of curiosity. A central and unifying topic of our research and application activities is the study of curiosity. While our past and future projects involve conceptual and methodological tools and contributions from multiple disciplines, we view our project as going beyond traditional interdisciplinarity. Indeed, just like climate science had become a focused research domain that has crossed the traditional boundaries of university departments, our aim has been to build and grow internationally an integrated science of curiosity in natural and artificial systems, crossing the traditional boundaries of AI, neuroscience, psychology, education, ethology, evolution, biology, physics. We are not primarily experts in AI or psychology or education working on a project about curiosity: rather we are primarily experts in curiosity and its links with open-ended learning, and we leverage and integrate concepts and techniques also often used in AI, psychology or education. **We aim to reinforce our existing contributions in this direction, ranging from building theories and experiments that add to the corpus of scientific knowledge on curiosity, to leading the organisation of international events dedicated to this integrated science.** As an example, PY Oudeyer is currently co-leading the organisation of a Gordon Research Conference series entitled “The New Science of Curiosity” (see <https://www.grc.org/>), and with H. Sauz on has also co-led the preparation of a large-scale European project with colleagues in Germany and UK on studying the developmental links between curiosity and meta-cognition.

Development and open-endedness in generative AI. There has been revolutionary advances in AI in the last few years, especially around generative systems such as multi-modal foundational models. However, as described above, these systems are still strongly limited in several key dimensions: they are not pro-active agents interacting with external environments, they lack grounding, meta-cognition and curiosity. **We believe our project can make fundamental scientific and technological contributions to adapt and extend current generative AI systems by integrating forms of curiosity, meta-cognition and grounding,** for which we recently made proofs of concepts, and vice-versa take advantage of powerful capabilities of foundational models to build new kinds of curiosity-driven learning systems capable of creative and abstract exploration learning and discovery. Beyond the new fundamental understanding it may provide, we believe this will in turn improve the scalability and impact of **real-world applications with high societal stakes**, such as helping scientists make new discoveries that may for e.g. help build more sustainable materials, generate cleaner energy or save energy, find molecules with medical applications, design accessible and efficient educational tools, or help design more sustainable forms of plant growing in agriculture.

Machine culture. Beyond technological advances, generative AI is also starting to have a major influence on human cultural evolution. They are now massively used as intermediation platforms between individuals and existing corpuses of knowledge and culture, conveying multiple forms of biased cultural perspectives that they can amplify. This phenomenon has recently become massive as social networks are pervaded by bots powered by generated AI systems, playing the roles of humans with particular opinion or backgrounds, and increasingly interacting directly among each other, beyond interaction with humans. While generative AI offers unique potential in enabling humans make discoveries and know and understand better each other's cultures, these properties have also been leveraged by diverse organisations to influence in unfair and dangerous manners what populations think and do. Even though this poses major societal issues, this evolution has been so rapid that basic scientific understanding of cultural evolution in hybrid human-machine groups is strongly lacking. Thus, we believe the parts of our project which aim at modelling cultural evolution in groups of generative AI agents, or hybrid groups, as well as its links with properties of curiosity-driven learning at the level of individuals, has a potential to make very useful contributions to these high stake issues.

Building self-organising AI from the ground up. When using continuous cellular automata as a playground for designing and evaluating our algorithms for curiosity-driven automated discovery for the sciences, we are also actually making direct contributions to the domain of Artificial Life. In particular, we believe the tools and approach we are taking, in particular exploring the self-organisation of sensorimotor agency and open-ended evolutionary processes, has the potential to have significant impact in this domain. This has been attested recently by our Best Paper award at the Alife 2023 conference (and also wider impact, e.g. through > 2 millions views of the popular science videos of Sciences Etonnantes and Ego presenting - in part- our work on this topic). As we are aiming to study the self-organisation of basic forms of memory, learning, and even autotelic learning, in such environments, this may also constitute a

foundational approach to build AI systems from the ground up, possibly opening new possibilities in terms of robustness, adaptivity and generalisation.

4.2 Societal and industrial impact

Translational educational technologies to foster curiosity-driven learning and critical mind in a fast evolving world. We live in a world that is evolving fast. For example, climate change will require human societies to renew their modes of living. Geopolitical processes cause large populations, with millions of children living in fragile and high poverty contexts, to migrate. New technologies, such as generative AI, are profoundly impacting economic dynamics, democracy, cultural evolution, and yet their advances were not predictable by researchers in AI five years ago. Yet, in most educational contexts, including in Europe, what is taught to children in classrooms is very similar to what was taught 50 years ago. And even for the so-called “fundamental knowledge” of mathematics and language, studies like PISA show a worrying decrease of skills and motivation in children. **How could we renew children’s education to help them thrive, be happy, and become actors of a world they will be proud of?** As mentioned in a recent report from OECD (Vincent-Lancrin et al., 2020), we believe it is essential to train children to become autonomous lifelong learner, through fostering and training their curiosity and their critical minds, their ability to go search by themselves new information, and to question the validity of information they collect, as well as question their own knowledge and opinions. Thus, our research program aimed to train curiosity and the associated metacognitive skills that underlie the critical mindset, has the potential to contribute in this perspective. In particular, **we aim to leverage our fundamental research in translational projects where we will work directly with major educational stakeholders from the start** (e.g. students, teachers, parents, educational institutions like individual schools, Académie de Bordeaux, edTech companies like EvidenceB, government and in particular ministry of education) to build educational interventions that will be efficient, adapted to the needs and constraints of real world educational contexts, and with the aim of large scale adoption and use (a first step in this direction are the AdaptivMaths and MIA Seconde educational software now deployed in all French primary schools and supported by the French ministry of education).

Generative AI and education: scientific understanding of stakes and opportunities. One particular topic we will focus on is the study of the opportunities and challenges of generative AI in education. While very recent (ChatGPT was introduced only 1.5 years ago), generative AI has already very importantly impacted the educational world in the last few months. More than 50% of children in the 12-18 age range have already used generative AI systems for their homework, and this tendency is quickly rising, including in Europe. This is associated with major challenges and opportunities. Challenges include forms of uses of generative AI by students that may harm their abilities to learn, understand, and be motivated to put effort and be actively engaged in these processes. Also, this impacts profoundly the way teachers shall design homeworks as students are now massively using these tools for helping them in homeworks. On the other hand, generative AI opens unique opportunities for rich personalised tutoring, ranging from opportunities to obtain tailored explanations and feedback, to getting the opportunity to discuss and train in foreign languages. Such opportunities may be particularly magnified for countries

where the educational system is underdeveloped (Nie et al., 2024). **Key aspects of our research program are geared towards studying these opportunities and challenges, for example running field studies in middle and high schools to understand how students currently (mis)understand and (mis)use generative AI tools.** In complement, we will continue working on **outreach**, especially developing educational tools enabling to improve generative AI literacy in students, teachers and parents, for example by further developing and disseminating our pedagogical video series “ChatGPT explained in 5mn” (https://developmentalsystems.org/chatgpt_5_minutes/en/), as well as through participating to popular science events, visiting middle and high schools and welcoming students in the lab, writing [popular science books](#), and participating to discussions on these topics in wide audience media.

Empowering discovery in science and engineering. Some major societal challenges, such as sustainable development and health, require designing new forms of materials, new techniques for producing energy, or new molecules that may be used in medicine. In many cases, the complexity of self-organising materials or biological systems involves significant scientific and engineering challenges for understanding, controlling and inventing. Following several of our recent proof-of-concept projects (Etcheverry, 2023 ; Etcheverry et al., 2024), we aim to do translational research also in this domain, **enabling chemists, physicists and biologists, in both academia and industry**, to efficiently use our tools for curiosity-driven exploration to help them make new discoveries. In particular, we are now starting exploring several new collaborations in these fields (see below).

4.3 Support to public policies

Given the high societal challenges associated to our work, associated on the one hand on the impact of generative AI at large in society and industry, and on the other hand on the specific challenges related to **education and how it may be impacted by AI technologies**, we also aim to **strongly develop our activities aimed at informing and supporting public policies**, which is one major objective of Inria. This can take several forms, but one key vector for such public support is actively participating in **interactions and discussions with public bodies that are both working on analyzing the current stakes and proposing new actions and new laws**. As examples, we recently supported Inria direction in writing notes on generative AI and its societal dimensions for the cabinet of E. Macron, we participated to interviews from **senators** preparing a report on AI and education, we made presentations of the stakes associated to training curiosity and metacognition using AI technologies at **Conseil Scientifique de l'Éducation Nationale**, and at an annual scientific event organized by **DNE** (Direction du Numérique Educatif is the department of the ministry of education dedicated to digital technologies), and we were invited by BPI to participate to evaluation and monitoring of projects related to education/edTech by this institution. We are also working to develop collaborations between Inria and the UK AI safety institute, towards building a French institution similar to the UK one. This includes developing a collaboration with Chris Summerfield on doing field studies to assess the current state of use of generative AI in middle and high schools to inform public policies on this topic.

Another form of support consists in developing **pedagogical materials** that can be readily reused and included in pedagogical platforms from public bodies, e.g. such as our [ChatGPT explained in 5mn video series](#) that has been integrated in various tools from DNE and in the European mooc made for introducing teachers to AI (the [AI4Teacher mooc](#)). We are also writing general audience books aimed at introducing AI to primary, middle and high school students, such as the book [“C’est pas moi, c’est l’IA”](#) that will be published by Nathan in september 2024.

Another form of support of public policies is our activities aimed at **sharing AI models and data that are fully open-source** (open weights and open data) and trained on data associated with appropriate rights (we are here also collaborating with the Hugging Face company to distribute these open models and data on their platform). For example, we recently built a project with the EvidenceB company, in collaboration with Région Ile-de-France, to build an open model trained on data from free manuals, for which authors will be retribution in an appropriate manner: this kind of model will enable wider and legally compliant access to AI models by the edTech ecosystem in France.

4.4 Key industrial collaborations

As described in sections above, we have diverse industrial collaborations for all dimensions of our research programme. Here, we summarise our key current industrial collaborations, as well as key target collaborations to be further developed.

Key current industrial collaborations.

EvidenceB (<https://evidenceb.com/>) In the domain of educational technology applications, we have had several major collaborations with EvidenceB company, ranging from P2IA/AdaptivMaths project to a CIFRE PhD thesis. We have an ongoing licencing agreement between Inria and EvidenceB enabling EvidenceB to use our ZPDES algorithm and software as a core component of all their intelligent tutoring systems (AdaptivMaths, AdaptivLang, MIA Seconde, etc). We recently built two large scale project proposals, one for P2IA and one for BPI call on generative AI applications.

Hugging Face (<https://huggingface.co/>) In the domain of AI, we have ongoing key collaborations with Hugging Face, for example through co-supervision of a PhD thesis with Thomas Wolf (co-founder and chief science officer at Hugging Face), both for developing new machine learning algorithms, e.g. for grounding LLMs in the dynamics of external environments, but also for disseminating our scientific and technological contributions in an open-source manner on their platforms.

IRT Saint Exupery (<https://www.irt-saintexupery.com/fr/>) In the domain of automated scientific discovery, we have an ongoing consortium collaboration around the project AIxIA, where we

study the use of autotelic exploration algorithms to map the space of interference behaviours on embedded software and hardware.

Key industrial collaborations to be further developed.

Solvay/Syensqo (<https://www.syensqo.com/fr/france/bordeaux>) In the domain of automated scientific discovery, we have started several discussions to develop collaborations on using autotelic exploration algorithms to efficiently explore and map the space of material design and properties, with the aim to help scientists at Syensqo to discover new materials with high environmental and functionality properties.

Ubisoft La Forge (<https://www.ubisoft.com/fr-fr/studio/laforge>) We started collaborations in the context of a Region Aquitaine project, and have several new ongoing discussions with Ubisoft La Forge France to develop collaborations around two topics: 1) grounding LLMs in the dynamics of particular video games and using them to play the role of human-like NPCs; 2) use automatic curriculum learning techniques to help onboarding new players in games that currently require high initial skills for playful integration in normal game modes.

Pollen Robotics (<https://www.pollen-robotics.com/>) This startup company was created in 2016 as a spin off the Flowers team, to develop open-source robotics hardware and software. While robotics is not currently a priority of the team, the Pollen Robotics company is now exploring the use of multimodal foundational models for enabling fast learning and grounded language interaction between humans and robots: this opens various possibilities to collaborate given our research on grounding foundational models.

The OpenLLM consortium (ClassCode et Linagora, <https://www.openllm-france.fr/>) This consortium gathering various public institutions/associations (e.g. ClassCode) and companies, such as Linagora, aims to develop open LLMs trained in a fully transparent and law-compliant way, to be shared with the French generative AI ecosystem, in particular for educational applications. We started setting up collaborations with this consortium both to use their models and to contribute improving these models and the associated data based on our research program.

Some other discussions have started with companies including Safran, Scaleway, The H company on topics related to autotelic and grounded generative AI.

5. Positioning and collaborations

Broader positioning in AI, cognitive science and educational technologies. As explained in the Scientific Impact section above, our main original positioning has been to develop an **integrated scientific approach to curiosity and open-ended learning, enabling us to build bridges and open new perspectives in AI, cognitive sciences and applications in educational technologies and assistive scientific discovery.** The ideas we developed across

the years on curiosity-driven learning have for example been reused and adapted by multiple academic and private AI research labs in the world. Our participation in mainstream machine learning conferences has been instrumental in achieving this impact, and enabled us to develop multiple collaborations that will be further expanded. In particular, while simple forms of curiosity-driven learning are now commonly used in the **machine learning community** (e.g. using learning progress to drive exploration), **the concept of autotelic exploration, which we have shown to be extremely powerful and general, has yet to be better disseminated and developed in this community**, especially in relation to building generative AI systems that can self-improve and verify their own knowledge. This will be one of our major objectives towards the AI community.

Similarly, while the **cognitive science community** is now significantly studying simple forms of curiosity, **the concept of autotelic learning and exploration is still under explored**: for example, existing experimental paradigms to study the self-generation of goals are still to be invented and tested. We aim to contribute to achieving the objective to disseminate further this concept, in particular through designing experimental protocols suited to study these advanced forms of curiosity-driven exploration. In our two main application domains, education and assisted scientific discovery, there are two dimensions of positioning. First, we will aim to increase our visibility in the targeted scientific communities (e.g. EDM or AIED conferences in education). Second, we will aim to work directly with actual stakeholders to maximise real world use of the technologies we develop. Our priority will be on this second dimension.

International labs working on similar topics. The topic of open-ended learning and its links with generative AI has become recently the focus of several international teams, and in particular: Sebastian Risi's team at the IT University of Copenhagen, Jeff Clune's team at University of British Columbia, and Tim Roscktaschel's team at Google Deepmind. These three teams have indeed recently developed projects studying how diversity search algorithms (e.g. quality-diversity algorithms) can enable generating continuously novel tasks and environments, in particular leveraging foundational models, to train open-ended agents. We have started collaborations with Sebastian Risi's team (two previous members of the team are now working in his team) as well as with Jeff Clune's team. While there is similarity, and even synergy, between our work and these team's projects, our project is unique in that it goes beyond studying diversity search in machine learning and artificial life: we have indeed strong research dimensions in cognitives science and in educational technology and assisted scientific discovery that are not at all present in these other teams.

In cognitive sciences, neurosciences and education, other teams in the world also study human curiosity-driven learning, such as Goren Gordon's lab (Tel Aviv University), Mathias Gruber (Univ. Cardiff), Yana Fandakova (Univ. Trier), Azzurra Ruggeri (Max Planck institute for human development), Jamie Jirout (Univ. Virginia), Kou Murayama (Univ. Tübingen) and many other participants of the Neurocuriosity symposia we have been organizing over the years. We have strong interactions with this community, with particularly strategic collaborations with Jacqueline Gottlieb (Univ. Columbia), Mathias Gruber (Univ. Cardiff), Yana Fandakova (Univ. Trier), Celeste Kidd (Univ. Berkeley) and Kou Murayama (Univ. Tübingen). Again, our project is unique as in

contrast to these teams, we also have research dimensions on curiosity-driven exploration and learning in machine learning, with applications beyond education, such as assisted scientific discovery.

Main active collaborations. In **academia**, our **key collaborators in AI** include Olivier Sigaud (Sorbonne Univ.; Deep RL), Laure Soulier (Sorbonne Université; LLMs) and Sylvain Lamprier (Univ. Angers; Deep RL and LLMs), Peter Ford Dominey (CNRS, cognitive modelling); Nathanael Fijalkow (Labri, program synthesis), Guillaume Baudard (PiCube Inria team, program synthesis), Xavier Hinault (Mnemosyne, neuroinspired machine learning), Pierre Dragicevic (Inria Bivwac); Marti Sanchez-Fibla (Univ. Pompeu Fabra, multi-agent systems); Josh Tenenbaum (MIT; cognitive modelling); Edith Law (Univ. Waterloo; HCI). Our key **academic collaborations in cognitive, culture and biological sciences** include Maxime Derox (Institute of Advanced Studies in Toulouse, cultural evolution), Andre Tricot (Univ. Montpellier, learning sciences); Matthias Gruber (Univ. Cardiff, cognitive neuroscience), Yana Fandakova (Univ. Trier, cognitive neuroscience), Daphné Bavelier (Univ. Geneva, learning and brain science) J. Gottlieb (Univ. Columbia, cognitive neuroscience), Laurianne Rat-Fisher (Nanterre Univ, developmental psychology), Celeste Kidd (Univ. Berkeley, developmental psychology), Anne Collins (Univ. Berkeley, human reinforcement learning); Myra Fernandes (Univ. Waterloo; cognitive neuroscience), Michael Levin (Tufts University, biology), Daphne Bavelier (Univ. Geneva; cognitive training), Ricard Solé (Univ. Pompeu Fabra, complex systems).

Some of these collaborations are associated with **large scale collaborative research programs**, such as our **CuriousTech associate team** with Univ. Waterloo focusing on cognitive science research on curiosity and applications in edTech (with E. Law and M. Fernandes), the LLM4Code Inria “Défi” focusing on studying how generative AI can foster research in computer science (with N. Fijalkow, X. Hinault and G. Baudard), a European ORA project on the cognitive science study of curiosity and metacognition (with M. Gruber and Y. Fandakova), the AlxIA project on assisted scientific discovery (with E. Jenn), the InnovCare project (with S. Lechevalier) within the PPR Autonomie-France 2030 (and with Fondation France-Japan of EHESS) to study the role of intrinsic motivation in adoption of technologies fostering autonomy in ageing populations, the project VBHI - France 2030 (IHU, S. Debette) to study the use of AI-driven personalization of cognitive training curricula in ageing populations (with F. Lotte and F. Wagner from Inria), and project CAP Santé Numérique at University of Bordeaux (with R. Thibault) where we study the role of curiosity in edTech for training neuroanatomy (with D. Liguoro, A. Tricot and M. Bertrand).

We are currently exploring the development of new collaborations including: Arvin Murugan (Univ. Chicago, physics); Cyril Aymonier (ICMCB, Bordeaux, chemistry); Christopher Summerfield (Univ. Oxford; UK Institute of AI Safety).

In the industry, the key organizations with whom we collaborate are described in the previous section. In addition to these organizations, we have strong informal collaborations with individuals from various industrial companies, including Benjamin Clément (EvidenceB, edTech), Bert Chan (Google Deepmind, artificial life), Ida Mommenejab

(Microsoft Research, models of cultural evolution), Marc-Alexandre Côté and Eric Yuan (Microsoft Research, Deep RL and generative AI).

We also have major collaborations with government organisations in the domain of education, such as Académie de Bordeaux and networks of primary, middle and high schools.

6. References

- Abdelghani, R., Oudeyer, P. Y., Law, E., de Vulpillières, C., & Sauzéon, H. (2022). Conversational agents for fostering curiosity-driven learning in children. *International Journal of Human-Computer Studies*, 167, 102887.
- Abdelghani, R., Sauzéon, H., & Oudeyer, P. Y. (2023a). Generative AI in the Classroom: Can Students Remain Active Learners? *NeurIPS 2023 - GAIED Workshop - Conference on Neural Information Processing Systems*, Dec 2023, New orleans, USA, United States. *arXiv preprint arXiv:2310.03192*.
- Abdelghani, R., Wang, Y. H., Yuan, X., Wang, T., Lucas, P., Sauzéon, H., & Oudeyer, P. Y. (2023b). Gpt-3-driven pedagogical agents to train children's curious question-asking skills. *International Journal of Artificial Intelligence in Education*, 1-36.
- Abdelghani, R., Law, E., Desvaux, C., Oudeyer, P. Y., & Sauzéon, H. (2023c). Interactive environments for training children's curiosity through the practice of metacognitive skills: a pilot study. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference* (pp. 495-501).
- Adolphe, M., Pech, M., Sawayama, M., Maurel, D., Delmas, A., Oudeyer, P. Y., & Sauzéon, H. (2023). Exploring the Potential of Artificial Intelligence in Individualized Cognitive Training: a Systematic Review. *OSF (submitted to PlosOne)*.
- Akakzia, A., Colas, C., Oudeyer, P. Y., Chetouani, M., & Sigaud, O. (2021). Grounding Language to Autonomously-Acquired Skills via Goal Generation. In *ICLR 2021*
- Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6), 26-38.
- Baranes, A., & Oudeyer, P. Y. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1), 49-73.
- Barde, P., Karch, T., Nowrouzezahrai, D., Moulin-Frier, C., Pal, C., & Oudeyer, P. Y. (2022). Learning to Guide and to Be Guided in the Architect-Builder Problem. In *International Conference on Learning Representations*.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., & Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29
- Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 5185-5198).
- Borgerhoff Mulder, M., & Mulder, M. B. (2003). Human behavioural ecology. *Encyclopedia of Life Sciences:(Human Behavior Ecology, A3671)*.
- Bornemann, R., Hamon, G., Nisioti, E., & Moulin-Frier, C. (2023). Emergence of collective open-ended exploration from Decentralized Meta-Reinforcement learning. In *Second Agent Learning in Open-Endedness Workshop*.

- Brinkmann, L., Baumann, F., Bonnefon, J. F., Derex, M., Müller, T. F., Nussberger, A. M., ... & Rahwan, I. (2023). Machine culture. *Nature Human Behaviour*, 7(11), 1855-1868.
- Brown, G. R., Dickins, T. E., Sear, R., & Laland, K. N. (2011). Evolutionary accounts of human behavioural diversity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1563), 313-324.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
- Bruner, J. (1990). *Acts of meaning: Four lectures on mind and culture* (Vol. 3). Harvard university press.
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., ... & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Carta, T., Romac, C., Wolf, T., Lamprier, S., Sigaud, O., & Oudeyer, P. Y. (2023). Grounding large language models in interactive environments with online reinforcement learning. *ICML 2023*.
- Carta, T., Oudeyer, P. Y., Sigaud, O., & Lamprier, S. (2022). Eager: Asking and answering questions for automatic reward shaping in language-guided rl. *Advances in Neural Information Processing Systems*, 35, 12478-12490.
- Chan, B. W. C. (2020). Lenia and expanded universe. In *Artificial Life Conference Proceedings 32* (pp. 221-229). One Rogers Street, Cambridge, MA 02142-1209, USA journals-info@ mit. edu: MIT Press.
- Chollet, F. (2019). On the measure of intelligence. *arXiv preprint arXiv:1911.01547*.
- Chu, J., & Schulz, L. E. (2020). Play, curiosity, and cognition. *Annual Review of Developmental Psychology*, 2, 317-343.
- Chu, J., Tenenbaum, J. B., & Schulz, L. E. (2023). In praise of folly: flexible goals and human cognition. *Trends in Cognitive Sciences*
- Clément, B., Roy, D. Pierre--Yves Oudeyer, Manuel Lopes (2016) Multi--Armed Bandits for Intelligent Tutoring Systems. *Journal of Educational Data Mining (JEDM)*, 7(2), 1310-13174.
- Clément, B., Sauzéon, H., Roy, D., & Oudeyer, P. Y. (2024). Improved Performances and Motivation in Intelligent Tutoring Systems: Combining Machine Learning and Learner Choice. *arXiv preprint arXiv:2402.01669*.
- Colas, C., Sigaud, O., & Oudeyer, P. Y. (2018) GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms. In *Proceedings of International Conference on Machine Learning (ICML 2018)*
- Colas, C., Fournier, P., & Chetouani, M., Sigaud, O., Oudeyer, P. Y., (2019). CURIOS: Intrinsically Motivated Modular Multi-Goal Reinforcement Learning. In *International Conference on Machine Learning (ICML 2019)*.
- Colas, C., Karch, T., Lair, N., Dussoux, J. M., Moulin-Frier, C., Dominey, P., & Oudeyer, P. Y. (2020) Language as a cognitive tool to imagine goals in curiosity driven exploration. *Advances in Neural Information Processing Systems*, 33, 3761-3774.
- Colas, C. (2021). *Towards Vygotskian Autotelic Agents: Learning Skills with Goals, Language and Intrinsically Motivated Deep Reinforcement Learning* (Doctoral dissertation, Université de Bordeaux).
- Colas, C., Karch, T., Moulin-Frier, C., & Oudeyer, P. Y. (2022). Language and culture internalization for human-like autotelic AI. *Nature Machine Intelligence*, 4(12), 1068-1076.
- Colas, C., Karch, T., Sigaud, O., & Oudeyer, P. Y. (2022b). Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey. *Journal of Artificial Intelligence Research*, 74, 1159-1199

- Colas, C., Teodorescu, L., Oudeyer, P. Y., Yuan, X., & Côté, M. A. (2023). Augmenting autotelic agents with large language models. In *Conference on Lifelong Learning Agents* (pp. 205-226). PMLR.
- Davidson, G., Todd, G., Togelius, J., Gureckis, T. M., & Lake, B. M. (2024). Goals as Reward-Producing Programs. arXiv preprint arXiv:2405.13242
- Delmas, A., Clément, B., Oudeyer, P. Y., & Sauzéon, H. (2018, November). Fostering health education with a serious game in children with asthma: pilot studies for assessing learning efficacy and automatized learning personalization. In *Frontiers in Education* (Vol. 3, p. 99). Frontiers Media SA.
- Derex, M., & Mesoudi, A. (2020). Cumulative cultural evolution within evolving population structures. *Trends in Cognitive Sciences*, 24(8),
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2019). Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*
- Etcheverry, M., Moulin-Frier, C., & Oudeyer, P. Y. (2020). Hierarchically organized latent modules for exploratory search in morphogenetic systems. *Advances in Neural Information Processing Systems*, 33, 4846-4859.
- Etcheverry, M. (2023). *Curiosity-driven AI for Science: Automated Discovery of Self-Organized Structures* (Doctoral dissertation, Inria & Labri, Université Bordeaux).
- Etcheverry, M., Moulin-Frier, C., Oudeyer, P. Y., & Levin, M. (2024). Ai-driven automated discovery tools reveal diverse behavioral competencies of biological networks, *Elife*.
- Falk, M. J., Roach, F. D., Gilpin, W., & Murugan, A. (2022). Curiosity-driven search for novel non-equilibrium behaviors. *arXiv preprint arXiv:2211.02589*.
- Fandakova, Y., & Gruber, M. J. (2021). States of curiosity and interest enhance memory differently in adolescents and in children. *Developmental Science*, 24(1), e13005.
- Flavell, J. H., Beach, D. R., & Chinsky, J. M. (1966). Spontaneous verbal rehearsal in a memory task as a function of age. *Child development*, 283-299
- Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in human neuroscience*, 8, 443.
- Forestier, S., & Oudeyer, P. Y. (2017). A unified model of speech and tool use early development. In *39th Annual Conference of the Cognitive Science Society (CogSci 2017)*
- Forestier, S., & Oudeyer, P. Y. (2016, August). Curiosity-driven development of tool use precursors: a computational model. In *38th annual conference of the cognitive science society (cogsci 2016)* (pp. 1859-1864).
- Forestier, S., Portelas, R., Mollard, Y., & Oudeyer, P. Y. (2022). Intrinsically motivated goal exploration processes with automatic curriculum learning. *Journal of Machine Learning Research*, 23(152), 1-41.
- Gasse, M., Grasset, D., Gaudron, G., & Oudeyer, P. Y. (2023). Using Confounded Data in Latent Model-Based Reinforcement Learning. *Transactions on Machine Learning Research*
- Gottlieb, J., & Oudeyer, P. Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19(12), 758-770
- Goupil, L., & Proust, J. (2023). Curiosity as a metacognitive feeling. *Cognition*, 231, 105325.
- Guilford, J. P. (1967). Creativity: Yesterday, today and tomorrow. *The Journal of Creative Behavior*, 1(1), 3-14.
- Guo, Y., & Ayoun, B. (2023). What's in it for them? The role of social curiosity and social needs in motivating and retaining hospitality employees. *International Journal of Hospitality Management*, 115, 103596.
- Hagtvedt, L. P., Dossinger, K., Harrison, S. H., & Huang, L. (2019). Curiosity made the cat more creative: Specific curiosity as a driver of creativity. *Organizational Behavior and Human Decision Processes*, 150, 1-13.

- Haluptzok, P., Bowers, M., & Kalai, A. T. (2022). Language models can teach themselves to program better. *arXiv preprint arXiv:2207.14502*.
- Hamon, G., Nisioti, E., & Moulin-Frier, C. (2023). Eco-evolutionary Dynamics of Non-episodic Neuroevolution in Large Multi-agent Environments. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation* (pp. 143-146).
- Hamon, G., Etcheverry, M., Chan, B. W. C., Moulin-Frier, C., & Oudeyer, P. Y. (2024). Discovering Sensorimotor Agency in Cellular Automata using Diversity Search. *arXiv preprint arXiv:2402.10236*.
- Hargrove, R. A., & Nietfeld, J. L. (2015). The impact of metacognitive instruction on creative problem solving. *The Journal of Experimental Education*, 83(3), 291-318.
- Hartung, F. M., & Renner, B. (2011). Social curiosity and interpersonal perception: A judge× trait interaction. *Personality and Social Psychology Bulletin*, 37(6), 796-814.
- Hartung, F. M., & Renner, B. (2013). Social curiosity and gossip: Related but different drives of social functioning. *PLoS One*, 8(7), e69996.
- Henrich, J. (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. princeton University press.
- Henrich, J., Boyd, R., Derex, M., Kline, M. A., Mesoudi, A., Muthukrishna, M., ... & Thomas, M. G. (2016). Understanding cumulative cultural evolution. *Proceedings of the National Academy of Sciences*, 113(44), E6724-E6725.
- Jia, N., Luo, X., Fang, Z., & Liao, C. (2024). When and how artificial intelligence augments employee creativity. *Academy of Management Journal*, 67(1), 5-32.
- Kahneman, D. (2017). *Thinking, fast and slow*, Penguin books
- Kaplan, F., & Oudeyer, P. Y. (2007). In search of the neural circuits of intrinsic motivation. *Frontiers in neuroscience*, 1, 9.
- Karch, T., Teodorescu, L., Hofmann, K., Moulin-Frier, C., & Oudeyer, P. Y. (2021). Grounding spatio-temporal language with transformers. *Advances in Neural Information Processing Systems*, 34, 5236-5249
- Karch, T. (2023). *Towards Social Autotelic Artificial Agents: Formation and Exploitation of Cultural Conventions in Autonomous Embodied Artificial Agents* (Doctoral dissertation, Bordeaux).
- Kashdan, T. B., Disabato, D. J., Goodman, F. R., & McKnight, P. E. (2020). The Five-Dimensional Curiosity Scale Revised (5DCR): Briefer subscales while separating overt and covert social curiosity. *Personality and individual differences*, 157, 109836.
- Koutstaal, W., Kedrick, K., & Gonzalez-Brito, J. (2022). Capturing, clarifying, and consolidating the curiosity-creativity connection. *Scientific reports*, 12(1), 15300.
- Kovač, G., Portelas, R., Hofmann, K., & Oudeyer, P. Y. (2021). SocialAI 0.1: Towards a Benchmark to Stimulate Research on Socio-Cognitive Abilities in Deep Reinforcement Learning Agents.
- Kovač, G., Portelas, R., Sawayama, M., Dominey, P. F., & Oudeyer, P. Y. (2024). Stick to your Role! Stability of Personal Values Expressed in Large Language Models. *Proceedings of the annual meeting of the cognitive science society*
- Lapeyre, M., Rouanet, P., & Oudeyer, P. Y. (2013a). The poppy humanoid robot: Leg design for biped locomotion. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 349-356). IEEE.
- Lapeyre, M., Rouanet, P., & Oudeyer, P. Y. (2013b). Poppy humanoid platform: Experimental evaluation of the role of a bio-inspired thigh shape. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (pp. 376-383). IEEE.
- Lapeyre, M., Rouanet, P., Grizou, J., Nguyen, S., Depraetre, F., Le Falher, A., & Oudeyer, P. Y. (2014, September). Poppy project: open-source fabrication of 3D printed humanoid robot for science, education and art. In *Digital Intelligence 2014* (p. 6).

- LeCun, Y. (2022). A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1).
- Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., & Hutter, M. (2020). Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47).
- Lemesle, Y., Sawayama, M., Pérez, G. V., Adolphe, M., Sauzéon, H., & Oudeyer, P. Y. (2022). Evaluating language-biased image classification based on semantic representations. In International Conference on Learning Representations (ICLR 2022).
- Levine, S., Finn, C., Darrell, T., & Abbeel, P. (2016). End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39), 1-40
- Liu, I. J., Yuan, X., Côté, M. A., Oudeyer, P. Y., & Schwing, A. (2022, June). Asking for knowledge (afk): Training rl agents to query external knowledge using language. In International Conference on Machine Learning (pp. 14073-14093).
- Liu, H., Li, C., Wu, Q., & Lee, Y. J. (2024). Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Li, M. Y., Fox, E. B., & Goodman, N. D. (2024). Automated Statistical Model Discovery with Language Models. arXiv preprint arXiv:2402.17879
- Mahowald, K., Ivanova, A. A., Blank, I. A., Kanwisher, N., Tenenbaum, J. B., & Fedorenko, E. (2024). Dissociating language and thought in large language models. *Trends in Cognitive Sciences*.
- Maslin, M. A., Shultz, S., & Trauth, M. H. (2015). A synthesis of the theories and concepts of early human evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1663), 20140064.
- Mazon, C., Fage, C., Consel, C., Amestoy, A., Hesling, I., Bouvard, M., ... & Sauzéon, H. (2019). Cognitive mediators of school-related socio-adaptive behaviors in ASD and intellectual disability pre-and adolescents: A pilot-study in french special education classrooms. *Brain sciences*, 9(12), 334.
- Mazon, C., Etchegoyhen, K., Saint-Supery, I., Amestoy, A., Bouvard, M., Consel, C., & Sauzéon, H. (2022). Fostering parents-professional collaboration for facilitating the school inclusion of students with ASD: Design of the "ToGather" web-based prototype. *Educational technology research and development*, 70(1), 231-262.
- Mazon, C., Clément, B., Roy, D., Oudeyer, P. Y., & Sauzéon, H. (2023). Pilot study of an intervention based on an intelligent tutoring system (ITS) for instructing mathematical skills of students with ASD and/or ID. *Education and Information Technologies*, 28(8), 9325-9354.
- Meyer, E., Sauzéon, H., Saint-Supery, I., & Mazon, C. (2023, August). Systematic review of technologies to collaborate and co-educate students with special educational needs and supporting their schooling. In *IHIET 2023-10th International Conference on Human Interaction and Emerging Technologies* (Vol. 111, pp. 1-12). AHFE International.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540), 529-533.
- Moulin-Frier, C. (2011). Rôle des relations perception-action dans la communication parlée et l'émergence des systèmes phonologiques: étude, modélisation computationnelle et simulations (Doctoral dissertation, Université de Grenoble).
- Moulin-Frier, C., Nguyen, S. M., & Oudeyer, P. Y. (2014). Self-organization of early vocal development in infants and machines: the role of intrinsic motivation. *Frontiers in psychology*, 4, 1006.
- Moulin-Frier, C. (2022). The Ecology of Open-Ended Skill Acquisition (HdR, Université de Bordeaux (UB)).
- Murayama, K. (2022). A reward-learning framework of knowledge acquisition: An integrated account of curiosity, interest, and intrinsic-extrinsic rewards. *Psychological Review*, 129(1), 175.

- Nettle, D., Gibson, M. A., Lawson, D. W., & Sear, R. (2013). Human behavioral ecology: current research and future prospects. *Behavioral Ecology*, 24(5), 1031-1040.
- Nguyen, S. M., & Oudeyer, P. Y. (2014). Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots*, 36, 273-294
- Nguyen, S. M., & Oudeyer, P. Y. (2012). Active choice of teachers, learning strategies and goals for a socially guided intrinsic motivation learner. *Paladyn*, 3(3), 136-146.
- Nie, A., Chandak, Y., Suzara, M., Malik, A., Woodrow, J., Peng, M., ... & Piech, C. (2024). The GPT Surprise: Offering Large Language Model Chat in a Massive Coding Class Reduced Engagement but Increased Adopters' Exam Performances (No. qy8zd). Center for Open Science
- Nisioti, E., Mahaut, M., Oudeyer, P. Y., Momennejad, I., & Moulin-Frier, C. (2022). Social network structure shapes innovation: experience-sharing in RL with SAPIENS. arXiv preprint arXiv:2206.05060.
- Nisioti, E., & Moulin-Frier, C. (2023a). Dynamics of niche construction in adaptable populations evolving in diverse environments. In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press.
- Nisioti, E., Masquil, E., Hamon, G., & Moulin-Frier, C. (2023b). Autotelic Reinforcement Learning in Multi-Agent Environments. In *Conference on Lifelong Learning Agents* (pp. 137-161). PMLR.
- Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, Nat McAleese, Nathalie Bradley-Schmieg, Nathaniel Wong, Nicolas Porcel, Roberta Raileanu, Steph Hughes-Fitt, Valentin Dalibard, and Wojciech Marian Czarnecki. Open-Ended Learning Leads to Generally Capable Agents. 2021. arXiv: 2107. 12808 [cs.LG].
- Oudeyer, P. Y. (2003). *L'auto-organisation de la parole*. (Doctoral dissertation, University Paris VI, Paris)
- Oudeyer, P. Y., & Smith, L. B. (2016). How evolution may work through curiosity-driven developmental process. *Topics in Cognitive Science*, 8(2), 492-502.
- Oudeyer, P. Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. *Progress in brain research*, 229, 257-284.
- Oudeyer, P. Y., & Kaplan, F. (2007a). What is intrinsic motivation? A typology of computational approaches. *Frontiers in neurorobotics*, 1, 108.
- Oudeyer, P. Y., Kaplan, F., & Hafner, V. V. (2007b). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2), 265-286.
- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1-22).
- Perez, J., Léger, C., Ovando-Tellez, M., Foulon, C., Dussaud, J., Oudeyer, P. Y., & Moulin-Frier, C. (2024). Cultural evolution in populations of Large Language Models. arXiv preprint arXiv:2403.08882.
- Peterson, E. G. (2020). Supporting curiosity in schools and classrooms. *Current Opinion in Behavioral Sciences*, 35, 7-13.
- Piaget, J. *The Language and Thought of the Child* (Routledge, 1926).
- Phillips, R. (2016). Curious about others: Relational and empathetic curiosity for diverse societies. *New Formations*, 88(88), 123-142.
- Péré, A., Forestier, S., Sigaud, O., & Oudeyer, P. Y. (2018). Unsupervised Learning of Goal Spaces for Intrinsically Motivated Goal Exploration. In *Proceedings of International Conference on Learning Representations (ICLR 2018)*, arXiv preprint arXiv:1803.00781.
- Plantec, E., Hamon, G., Etcheverry, M., Oudeyer, P. Y., Moulin-Frier, C., & Chan, B. W. C. (2023). Flow-Lenia: Towards open-ended evolution in cellular automata through mass conservation and

- parameter localization. In *Artificial Life Conference Proceedings 35* (Vol. 2023, No. 1, p. 131). One Rogers Street, Cambridge, MA 02142-1209, USA journals-info@mit.edu: MIT Press.
- Poupard, M., Larrue, F., Sauzéon, H., & Tricot, A. (2024, accepted). A systematic review of immersive technologies for education: effects of cognitive load and curiosity state on learning performance. *British Journal of Educational Technology*.
- Pourcel, J., Colas, C., Oudeyer, P. Y., & Teodorescu, L. (2024). *Generating a Diversity of Challenging Programming Puzzles with Autotelic Generative Models*. *arXiv preprint arXiv:2310.10692*. Submitted to NeurIPS.
- Puente-Diaz, R., & Cavazos-Arroyo, J. (2020). Creative metacognitive feelings as a source of information for creative self-efficacy, creativity potential, intrapersonal idea selection, and task enjoyment. *The Journal of Creative Behavior*, 54(3), 499-507.
- Reinke, C., Etcheverry, M., & Oudeyer, P. Y. (2020). Intrinsically Motivated Discovery of Diverse Patterns in Self-Organizing Systems. In Proceedings of International Conference on Learning Representations (ICLR).
- Romac, C., Portelas, R., Hofmann, K., & Oudeyer, P. Y. (2021). TeachMyAgent: a Benchmark for Automatic Curriculum Learning in Deep RL. ICML 2021.
- Rule, J., Goddu, M., Chu, J., Pinter, V., Reagan, E. R., Bonawitz, E., & Ullman, T. (2023). Fun isn't easy: Children choose more difficult options when "playing for fun" vs. "trying to win"
- Saint-Supery, I., Sauzéon, H., Maillart, C., Neu, N., Meyer, E., & Mazon, C. (2023, August). Cross-cultural evaluation of a web application to support communication and collaboration among stakeholders of the school inclusion of children with ASD. In *AAATE 2023-The 17th International Conference of the Association for the Advancement of Assistive Technology in Europe*.
- Sigaud, O., Baldassarre, G., Colas, C., Doncieux, S., Duro, R., Perrin-Gilbert, N., & Santucci, V. G. (2023). A Definition of Open-Ended Learning Problems for Goal-Conditioned Agents. *arXiv preprint arXiv:2311.00344*.
- Sivashankar, Y., Fernandes, M., Oudeyer, P. Y., & Sauzéon, H. (2024). The beneficial role of curiosity on route memory in children. *Frontiers in Cognition*, 3, 1346280.
- Poupard, M., Larrue, F., Sauzéon, H., & Tricot, A. (2022). A systematic review of immersive technologies for education: effects of cognitive load and curiosity state on learning performance.
- Poli, F., Meyer, M., Mars, R. B., & Hunnius, S. (2022). Contributions of expected learning progress and perceptual novelty to curiosity-driven exploration. *Cognition*, 225, 105119.
- Said-Metwaly, S., Van den Noortgate, W., & Kyndt, E. (2017). Approaches to measuring creativity: A systematic literature review. *Creativity. Theories–Research–Applications*, 4(2), 238-275.
- Sayalı, C., Heling, E., & Cools, R. (2023). Learning progress mediates the link between cognitive effort and task engagement. *Cognition*, 236, 105418.
- Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Hambro, E., ... & Scialom, T. (2024). Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36.
- Schutte, N. S., & Malouff, J. M. (2020). A meta-analysis of the relationship between curiosity and creativity. *The Journal of Creative Behavior*, 54(4), 940-947.
- Sear, R., Lawson, D. W., & Dickins, T. E. (2007). Synthesis in the human evolutionary behavioural sciences. *Journal of Evolutionary Psychology*, 5(1), 3-28
- Shulman, M. (2024). Strange new universes: Proof assistants and synthetic foundations. *Bulletin of the American Mathematical Society*.
- Sigaud, O., Baldassarre, G., Colas, C., Doncieux, S., Duro, R., Perrin-Gilbert, N., & Santucci, V. G. (2023). A Definition of Open-Ended Learning Problems for Goal-Conditioned Agents. *arXiv preprint arXiv:2311.00344*.

- Ten, A., Kaushik, P., Oudeyer, P. Y., & Gottlieb, J. (2021). Humans monitor learning progress in curiosity-driven exploration. *Nature communications*, 12(1), 5972.
- Ten, A., Oudeyer, P. Y., & Moulin-Frier, C. (2022). Curiosity-driven exploration. *The Drive for Knowledge: The Science of Human Information Seeking*, 53.
- Tomasello, M. (2019). *Becoming human: A theory of ontogeny*. Harvard University Press.
- Vincent-Lancrin, S., et al. (2020), *Développer la créativité et l'esprit critique des élèves : Des actions concrètes pour l'école*, La recherche et l'innovation dans l'enseignement, Éditions OCDE, Paris,
- Vygotsky, L. S. Thought and Language (MIT press, 1934)
- Wang, G., Xie, Y., Jiang, Y., Mandlkar, A., Xiao, C., Zhu, Y., ... & Anandkumar, A. (2023). *Voyager: An Open-Ended Embodied Agent with Large Language Models*. In *Intrinsically-Motivated and Open-Ended Learning Workshop@ NeurIPS2023*.
- Xiao, Z., Yuan, X., Liao, Q. V., Abdelghani, R., & Oudeyer, P. Y. (2023). *Supporting qualitative analysis with large language models: Combining codebook with GPT-3 for deductive coding*. In *Companion proceedings of the 28th international conference on intelligent user interfaces*
- Zhang, Y., Li, Y., Cui, L., Cai, D., Liu, L., Fu, T., ... & Shi, S. (2023). *Siren's song in the AI ocean: a survey on hallucination in large language models*. arXiv preprint arXiv:2309.01219.

7. Appendix: Detailed research program and applications

7.1 Objective 1: Improve understanding of human autotelic and aligned intelligence in humans

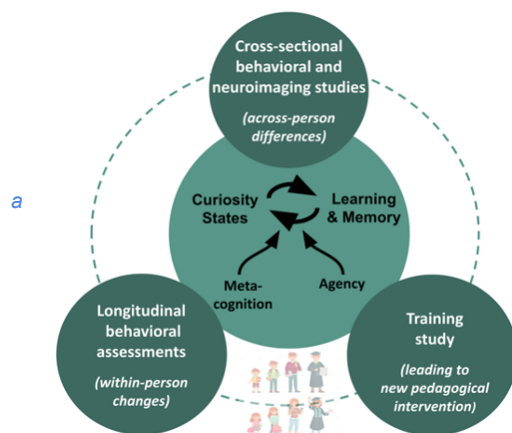
The team's overall approach to curiosity and cultural evolution in AI has relied on our cognitive science research towards understanding these processes in humans. In the last two decades, we both made fundamental contributions in theoretical and experimental understanding of human curiosity, for example through introducing the Learning Progress Hypothesis (Gottlieb & Oudeyer, 2018.) and testing some of its predictions in novel experimental paradigms with human subjects (Ten et al., 2021) . We also made key contributions in modelling human language evolution, in particular through our works on the evolution and self-organisation of speech sound systems (Oudeyer 2003, Moulin-Frier 2011). In recent years, we started building theoretical bridges for joint understanding of curiosity and language origins in humans (Oudeyer and Smith, 2016). This work has been strongly fuelled by interdisciplinary approaches (within the team where H. Sauzéon and C. Mazon are cognitive scientists, and through collaborations with various labs) and scientific community building activities, e.g. through our co-organizing role in the series of Neurocuriosity symposia in collaboration with Jacqueline Gottlieb (Columbia Univ. US). We aim to continue working on several key open scientific questions in this line of work.

Curiosity, meta-cognition and agency across the lifespan. The Learning Progress hypothesis, as well as other theories of curiosity-driven learning, all assume meta-cognitive competencies (e.g. ability to evaluate one's own uncertainty, knowledge gaps or learning progress) as well as forms of agency. However, experimental studies of human curiosity have so far mostly overlooked studying the influence of meta-cognition and agency, let alone simply

measure them together with various dimensions of curiosity (e.g., Murayam, 2022)). Another major limit of current models and experimental studies of human curiosity has been that they have not studied how curiosity develops across the lifespan. Actually, the scientific community knows very little on how various forms of curiosity change across childhood, adolescence, and up to ageing populations.

We will aim to address some of these limitations by collaborating with various international groups, including M. Gruber (Univ. Cardiff) and Y. Fandakova (Max Planck Institute for Human Development) with whom we just submitted a major ANR/DFG/ESRC project on this topic. In particular, we propose an interdisciplinary approach to make new breakthroughs in understanding how metacognition contributes to the development of curiosity-based learning, and set the stage for educational interventions that could help children develop their curiosity. Given the links between curiosity and metacognition, and the fact that metacognition continues to improve across childhood and adolescence, we formulate the hypothesis that the efficiency of curiosity-based learning, i.e. the ability to inquire about and prioritise learning of information associated with high curiosity, improves across child and adolescent development.

In this collaborative project, we will answer this question with a series of experiments (Figure FIG). We will combine different methodological approaches, including behavioural and functional neuroimaging experiments, training studies and longitudinal assessments to investigate how across-person differences and within-person changes in metacognition and agency contribute to curiosity-based learning. We will leverage established and newly developed paradigms, using digital technologies not previously used in the field of curiosity-based learning, which will enable us to better understand how metacognition and agency affect curiosity-based learning across development. For instance, we will use Virtual Reality (VR) for in-depth analysis of how exploratory behaviours lead to curiosity-based learning across age. Furthermore, we will use large language models (LLMs) for qualitative analyses of think-aloud components in the experiments to characterise verbal responses associated with metacognitive processes in curiosity-based learning. Finally, we will translate the lab-based findings from our proposed experiments to design pedagogical interventions that stimulate curiosity and metacognition in the classroom, and conduct an associated field study (described in the Educational Technology application section below). Building on our previous pioneering research in the field across our labs, this timely project will provide a unique opportunity for a step change in our understanding of curiosity development.



Centre circle: We hypothesise that both metacognition and agency impact the development of elicitation of curiosity states, how curiosity-driven information search is organised, and the efficiency of curiosity to enhance learning and memory. **Outer circles:** To study various dimensions of this hypothesis, we will conduct cross-sectional behavioural and neuroimaging studies using complementary experimental approaches, longitudinal behavioural study with multiple assessments to understand within-person changes, and a training study which has the potential to lead to new pedagogical interventions to stimulate curiosity in the classroom.

Experimental paradigms for studying autotelic learning in humans. Another limit of existing experimental studies of human curiosity, including the ones mentioned above, has been that most of them focused so far on studying how humans prefer exploring one of several pre-existing stimuli or learning activities (Ten, Oudeyer & Moulin-Frier, 2022). However, as shown in our theoretical and AI work described above, and as argued in complementary arguments from Laura Schulz and Junyi Chu (Chu and Schultz, 2020), exploration of self-generated goals, including arbitrary goals or games, may be key in accounting for human development, and further in accounting for human innovation and cultural evolution (Chu et al., 2023). Only very few exploratory experimental protocols have started to be investigated in the literature (e.g. Davidson et al., 2024; Rulet et al., 2023): we aim to further develop this form of experimental protocol, informed by predictions made by our theoretical models, in collaboration with researchers such as G. Molinaro and A. Collins (Univ. Berkeley, both in a 6 months research visit at Inria Flowers and Mnemosyne in 2024) J. Chu (Harvard Univ, US), L. Rat-Fischer (Univ. Nanterre) and A. Ruggeri (TU Munich).

Links between curiosity and creativity for autotelic learning in children: The ability to imagine abstract and new goals is essential for creative discovery and open-ended learning throughout life. Children achieve this by using the compositionality of language as a tool to imagine situations they have never experienced before, targeting them as goals during play (Piaget, 1926; Vygotsky, 1938). Echoing the IMAGINE architecture (Colas et al., 2020), an intrinsically motivated deep reinforcement learning architecture modelling compositional imagination (the creation of new linguistic associations for new goals), we aim to investigate the links between curiosity and creativity in humans, focusing on the metacognitive role of language in guiding autonomous learning behaviours. Although the nature of the links between curiosity and creativity is currently not well defined, a recent meta-analysis shows that higher levels of curiosity are significantly associated with higher levels of creativity (Schutte & Malouff, 2020). Divergent thinking mechanisms are said to be the cognitive resource common to both skills (Guilford, 1967; Koutstaal et al., 2022), and some authors even identify curiosity as a facilitator, a trigger for creativity (Hagtvedt et al. 2019): high curiosity states induce better ideation and greater idea associations conducive to problem solving. Also, both creativity and curiosity are governed by metacognitive processes of self-regulation of learning (Puente-Diaz & Cavazos-Arroyo, 2020) enabling the identification of information gaps, problem situations or uncertainties, the generation of ideas, paths to resolution, and monitoring and evaluating the value of ideas as creative output or as majoring knowledge (Hargrove & Nietfeld, 2015; Jia et al., 2019). We aim to investigate developmental differences on curiosity-based learning and problem-solving tasks while studying their relationships and their dependency to intrapersonal factors (especially metacognitive skills and personality dimensions such as epistemic curiosity, creativity or intellectual humility traits) in late childhood (from 6 to 11 yo). To achieve experiments needed to address these topics, we will leverage an educational Léa-Ifé collaboration network established with 10 primary schools around Bordeaux. As a whole, in this part of the project, we aim to demonstrate that curiosity as a process to seek knowledge in the face of self-generated goals of knowledge gaps, or as a metacognitive feeling (Goupil & Proust, 2023), leads to better initiation of the creative process.

Curiosity to learn about others and social interaction:

Social curiosity is defined as the desire to acquire knowledge about others in society, encompassing an interest in their emotions, thoughts, and behaviours. This type of curiosity can be divided into two forms (Philips, 2016): 1) empathetic curiosity (the desire to acquire knowledge about others), and 2) relational curiosity (the desire to interact with others). Like other types of curiosity, social curiosity motivates people to engage in exploratory behaviours directed toward the social world, seeking novel information about how people think, behave, and feel. Hartung and Renner (2013) proposes three functions of social curiosity: 1) acquiring information useful for learning and development, 2) establishing interpersonal relationships and increasing a sense of social belonging, and 3) controlling the social world by making it more predictable and manageable. Thus, social curiosity enhances social functioning and has been linked to improved social behaviour adaptation, the ability to establish and maintain social relationships, and better social judgement abilities (Hartung and Renner, 2011, 2013). Recently, another distinction has been proposed in social curiosity (Kashdan et al., 2020): 1) overt social curiosity, an explicit interest in understanding other people, which motivates direct communication with others; 2) covert social curiosity, an “hidden” interest that motivates more indirect and furtive behaviours to understand others, such as discreetly observing people, listening to others’ conversations, and reading tabloids and human-interest stories. Covert curiosity is often associated with negative outcomes like gossiping or spying (Kashdan et al., 2020), but it can also drive the understanding of the social world through observation and finally motive interactions with others (Guo & Ayoun, 2023). On the other hand, overt social curiosity has been linked with open-mindedness, extraversion, and sociability (Kashdan, et al. 2020), and was associated with better job performance (Guo & Ayoun, 2023).

Autotelic game invention and cultural transmission. Leveraging the theoretical ideas on the interaction between autotelic learning and cultural evolution as described in the previous section, we also aim to study experimentally these interactions in chains of humans incentivized to transmit to each other games or artefacts of their own intrinsically motivated invention (either physical or video games, e.g. using experimental setups like Davidson et al., 2024). We aim to design new experimental protocols and run them both in various age ranges in European populations, as well as in populations in non western culture (e.g. Tsimane indian populations in south america), leveraging associated collaborations with Maxime Derex at IAST, Toulouse.

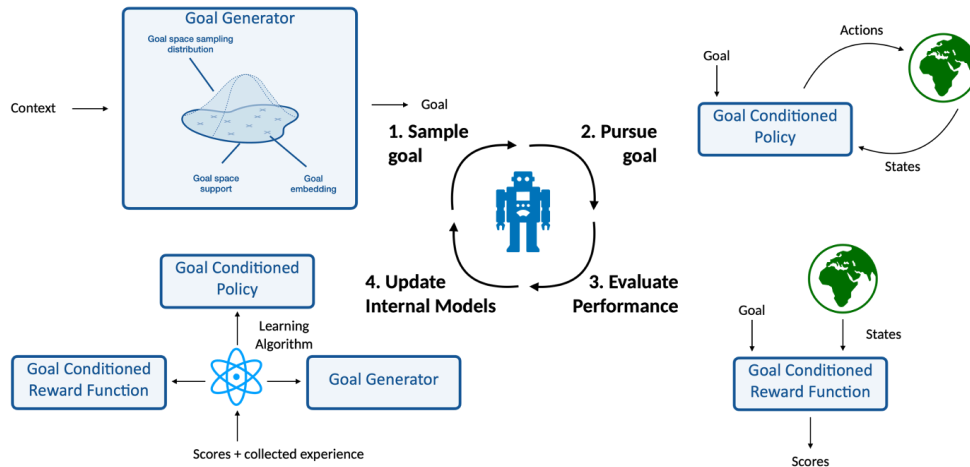
7.2 Objective 2: Building curiosity-driven autotelic and aligned machines

7.2.1 Objective 2.1: Connecting and disseminating autotelic learning to the machine learning community

While the team's work had great impact and influence in the developmental robotics community (e.g. in some years, around 30% of papers in the IEEE ICDL-Epirob conference discussed curiosity-driven learning in robots and cited our works), the ideas and algorithms were still little known in the machine learning community up to 2016-17. However, our experimental results showed great potential of intrinsically-motivated learning algorithms, and in particular autotelic to address key unsolved problems in machine learning, in particular both RL problems with very sparse rewards, multi-goal RL problems. Further, some fundamental scientific questions we were exploring about open-ended learning were little considered at that time by the machine learning community, and we believed they were important for the future of that field. In reverse, our works did not use at that time strong technical advances in machine learning, in particular deep reinforcement learning methods for which impressive results were shown in 2016 (e.g. Mnih et al., 2015).

For all these reasons, the team made a strategic pivot: we started to focus our research and dissemination efforts towards ML communities gathered in conferences such as Neurips, ICML and ICLR. This was a challenge as at this point, we were both not experts in Deep RL methods, and the epistemology of developmental robotics was different from the habits of the ML community in terms of writing scientific papers. To address these challenges, we benefitted from welcoming Olivier Sigaud, one of the best Deep RL researcher in France, in 2017 as a detached senior researcher in the team to help us make this pivot successful. We quickly made substantial contributions, in particular developing algorithmic architectures first mixing autotelic memory-based and deep RL techniques to solve sparse reward problems (Colas et al., ICML 2018), then the UGL algorithm for learning goal representations (Pere et al., ICLR 2018) then the CURIOUS algorithm which was the first autotelic deep reinforcement learning system (Colas et al., 2019), as well as autotelic learning algorithms used for automatic curriculum training of Deep RL agents (Romac et al., ICML 2021), all published at ICML. These works were soon followed by a whole series of publications at Neurips, ICLR and ICML (e.g. Colas et al., Neurips 2020; Akakzia et al., ICLR 2021; Karch et al., Neurips 2021; Liu et al., ICML 2022; Lemesle et al., ICLR 2022; Barde et al., ICLR 2022; Carta et al., Neurips 2022; Carta et al., ICML 2023), some of them obtaining very selective oral accept (best 1% papers, e.g. Reinke et al, ICLR 2020; Etcheverry et al., Neurips 2020). This series of work was identified and recognized by a growing set of academic machine learning labs in the world, and strongly contributed to grow the importance of the research topic of intrinsically motivated machine learning: we recently wrote a review paper surveying this emerging field of curiosity-driven autotelic deep reinforcement

learning (Colas et al., JAIR 2022), as well as published one of our key autotelic learning algorithm in JMLR (Forestier et al. , 2022). Several major AI labs in the world started to use this family of methods to solve sparse reward problems that were previously unsolvable (including labs at Stanford, UCL, Google Brain, Meta, Uber AI labs, see e.g. Bellemare et al. 2016; Ecoffet et al. 2019; LeCun 2022).



The autotelic exploration and learning loop (adapted from Karch, 2023)

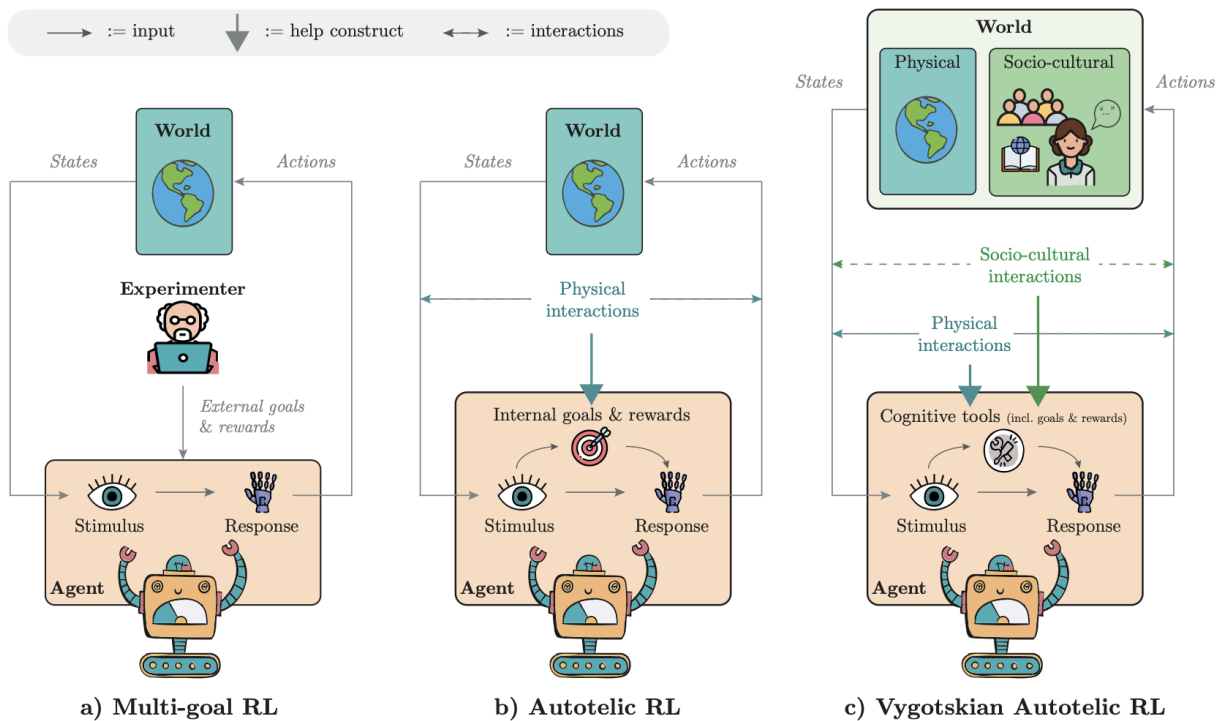
7.2.2 Objective 2.2: Curiosity-driven autotelic AI agents leveraging large foundational models

Here we aim to jointly address current limits and challenges of curiosity-driven autotelic AI systems on the one hand, and of large foundational models on the other hand. As argued above, the first generation of curiosity-driven autotelic AI systems, developed in the team and in a few other teams in the world, opened new perspectives for open-ended learning in interactive physical environments but were limited in terms of abstraction, sample efficiency, generalisation, and alignment on human cultures. On the other hand, large foundational models, and language models in particular, opened extraordinary perspectives for fluid handling of diverse language and visual tasks, encoding rich cultural models, but are strongly limited in terms of autonomy, requirement of very large datasets and compute resources, and lack grounding and causal understanding of external interactive environments.

Interestingly, some major limits of autotelic learning systems are strengths of language models, e.g. LLMs as culture models can help align on human cultures, bootstrap efficient exploration by leveraging prior knowledge, enable few-shot learning through in-context processes, and enable generalisation and abstraction through language combinatoriality. Reversely, some major limits of LLMs are strengths of autotelic learning systems, e.g. autotelic architectures can provide

forms of agency to LLMs, enable them to collect or generate autonomously new data for self-improvement, and align their internal knowledge on the external dynamics of environment.

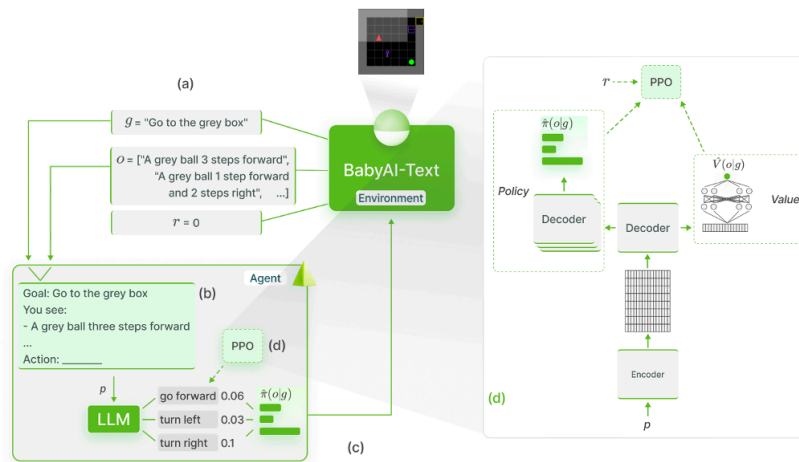
This project will leverage these complementarities to **build autotelic agents leveraging large foundational models**, which we believe could set the ground for a new generation of machines that would be both more open-ended, more flexible and robust, more sample efficient, with higher capabilities for abstraction and generalisation, and aligned on human cultures (as we started to outline in Colas et al., 2022, 2022b). Here are examples of key concrete dimensions we will explore, most of which are already started to be explored in current PhD and postdoc projects in the team:



Language and culture can be leveraged as cognitive tools to augment autotelic reinforcement learning agents, e.g. enabling abstract goal generation or imagination of high-level plans. We study how generative models (e.g. LLMs) can be used to bootstrap such architectures (figure adapted from Colas et al., 2022).

Large foundational models for representing, sampling and solving goals in autotelic architectures. We recently started to explore how LLMs (e.g. ChatGPT) can be used as cognitive tools supporting exploration and planning of reinforcement learning agents, such as in the LMA3 architecture (Colas et al., 2023). Here, agents are architectures combining a “traditional” autotelic RL algorithm with LLMs used as interactive tools that can be probed to 1) sample new relevant goals given the list of already achieved goals and state of the environment (this leverages in-context learning); 2) evaluate goal-achievement given the realized behavioural trajectory; 3) achieve hindsight relabelling for efficient learning. However, goals remained quite concrete even though they were expressed in language. We aim to study how LLMs can be leveraged to generate new kinds of goals, including more creative goals (e.g. leveraging

prompting strategies that stimulate creativity, inspired from human techniques), time-extended goals, and learning goals. Learning goals may be expressed as self-generated questions about the environment, but could also consist in generating hypotheses to explain observed phenomena, or to account for gaps in theories that a language model could read in books: autotelic exploration will aim to test these hypotheses and answer these questions. Thus, this will pave the way towards curiosity-driven agents that explore to build scientific and causal theories of their environments (Gasse et al., 2023; Li et al., 2024). We will also leverage LLMs as tools that can provide relevant suggestions to generate plans to achieve self-generated goals, including through building and handling repertoires of hierarchical skill libraries, potentially expressed as code like in (Wang et al., 2023).



Large language models can be grounded in interactive environments when they are used as agent policies: in the GLAM architecture, feedback from the environment enables to align their internal representations to the external dynamics using online RL (Carta et al., 2023).

Large foundational models as agents: grounding and alignment on external physical and social dynamics. In LMA3-like architectures, LLMs are fixed tools: they do not evolve. Yet, letting autotelic agents rely on LLMs also poses challenges. LLMs are known to convey false information and harmful biases that may come directly from the data sources or be introduced through inadequate training processes (Bender and Koller, 2020; Zang et al, 2023). Autotelic agents relying on these models could be misled in the exploration of the physical world, and also demonstrate harmful behaviours and contribute to reinforcing stereotypes and inequalities when interacting with humans. Ideally, autotelic agents equipped with LLMs should align with the dynamics of external environments and with the objectives and needs of a particular family of users, tasks, and contexts. To address this challenge, we will study how LLMs can themselves learn incrementally through environment interaction, leveraging an online model-free reinforcement learning approach we introduced recently (Carta et al., 2023), and which we will extend to autotelic exploring agents as well as to rich social-cultural environments (see below).

Learning internal (abstract, language level) world models with curiosity-driven LLMs. While pre-trained LLMs encode cultural knowledge about the world (which may be inaccurate, see previous paragraph), they do not typically encode world models, i.e. they do not have the ability to simulate and predict faithfully what could happen if certain states and actions were

encountered. This is associated with their lack of causal understanding and counterfactual capabilities. We will combine autotelic architectures manipulating learning goals (e.g. generating and testing hypotheses about future world states) and reinforcement learning (in particular model-based) to study how LLMs could explicitly learn causal models of the world through iterating well-chosen interventions and achieving corresponding representations updates.

Learning and coordinating diverse models of interestingness at multiple time scales. One key dimension of models of curiosity-driven autotelic learning is how measures of goal interestingness are computed: this enables to prioritize goals, generating a learning curriculum that is essential to enable actual learning and discovery in vast environment where many goals are either trivial or unlearnable, and where time is too limited to enable learning all learnable goals. The learning progress hypothesis suggests that goals shall be sampled according to expected learning progress: however, models considered so far have applied learning progress over concrete and short term goal types for which it was relatively straightforward to compute learning progress. But how can one compute learning progress, or even learnability, for more abstract or vaguely defined goals (e.g. “building a beautiful wooden tower”, “becoming an expert in juggling”, “programming a novel fun video game”)? Progress towards many of such goals may not be measured objectively, but rather be relative to a subjective culturally-grounded conceptual system. To address this challenge, we will study how various cultural perspectives can be elicited in language models, and leveraged by autotelic agents to build subjectives and culturally grounded measures of achievement and progress. More generally, novelty is also a key measure used to drive curiosity-driven exploration, and as argued in a recent paper to which we contributed (Sigaud et al., 2023), this is also a notion that shall be conceptualised relative to observers and their conceptual systems. Thus, multiple kinds of novelty may exist, and exploring them requires form of meta-diversity search, an approach we introduced in Etcheverry et al. (2020): we will here study how LLMs can be used to generate novel conceptual dimensions as new discoveries are made by autotelic agents, leading to LLM-based meta-diversity search. Beyond the notion of objective versus subjective learning progress (or novelty), the cultural environment may also encode preferences over goals types that are akin to social norms, which may also guide autotelic agents.

Overall, this entails that multiple kinds of measures of interestingness may co-exist in agent’s motivational systems: we will study how this can be formalised and modelled using both heuristic and normative approaches.

Training meta-cognitive skills in autotelic LLMs. One major limit of current foundational models, and LLMs in particular, is that they do not have capabilities for measuring and expressing uncertainty about knowledge or predictions they express. Yet, such meta-cognitive abilities are at the core of autotelic learning, as estimating learning progress, learnability, novelty or surprise are key measures of interestingness (see above). In a context where LLMs can be used and trained themselves through online RL to generate goals and plans, we will also investigate how they can be trained in the same manner to assess quantitatively learnability and uncertainty measures about their behaviour and predictions, de facto equipping them with meta-cognitive abilities. We will aim to study how learning such meta-cognitive skills in a particular domain may then transfer in other domains.

Teachable autotelic agents that align with the culture of human users. The interaction between autotelic LLMs and social peers is of key importance both from a fundamental and application/societal points of views: autotelic AI systems can leverage human guidance to learn more efficiently, and they also need to be steerable and alignable on the culture of human users so that they remain in service of these users in application scenarios. As argued by Vygotsky, social peers help learners acquire complex knowledge and skills by scaffolding learning in the zone of proximal development, achieving curriculum learning. Such guidance is likely to be essential to enable efficient learning in the real world and across modalities, and to enable alignment with human needs and values, especially when learning social skills (e.g. theory of mind, cooperative thinking, ethics or social norms – Tomasello, 2019). This will require LLMs to have the ability to leverage diverse forms of guidance from real human users, ranging from instructions, corrections, advice, plan suggestions, explanations, or pointers to cultural artefacts (e.g. textbooks). As we are still a long way from achieving efficient learning from real time interactions with real humans, a first step in this direction will consist in immersing autotelic agents in rich simulated socio-cultural environment, where video game like worlds will be populated by complex social peers simulated using LLMs (e.g. Park et al., 2023), and involving representative social challenges, leveraging our pilot work on the Social AI school (Kovač et al., 2021). Importantly, this socio-cultural environments will need to be multimodal, as human culture is also embedded in rich and diverse non-verbal modalities: this will involve leveraging audio-visual foundational models and applying the learning methodologies mentioned in previous paragraphs.

Learning to use cognitive tools (from chain of thoughts to external APIs): We will leverage recent work showing how LLMs can learn to use external tools or APIs (Schick et al., 2024) to extend the capabilities of autotelic LLM-based agents. This will on the one hand consist in computational tools enabling agents to generate and execute code (e.g. Python code), for example when autotelic agents might generate and explore behaviours that are encoded directly as a computer program (and use interpreters or simulators as tools to optimise these programs before running and testing them in the environment). Another kind of tool will consist in search and retrieval tools used by LLM-based autotelic agents to find relevant information in external knowledge sources, ranging from domain specific knowledge bases to directly asking humans (thus, these tools will be useful for points 1 and 3 above). We will in particular study how autotelic agents can learn to use these tools using online RL (= when to make calls to these APIs? Which calls? How to interpret the results?).

Self-improving small LLMs that learn how to invent and solve code and maths problem

Leveraging recent work showing the potential of LLMs to generate code and interact with code interpreters (Haluptzok et al., 2022), we will develop and study autotelic LLMs that learn how to code. More precisely, we will study how autotelic exploration algorithms can be used to drive LLMs to self-generate new coding problems, guided by linguistic descriptor spaces, and using interaction with code interpreters to self-evaluate and self-improve on self-generated problems, while grounding their skills in these coding external environments. This will aim to address limits of traditional environments used to study autotelic learning: most experimental environments, including rich video games like Minecraft, are actually not open-ended, while code generation and interpretation is open-ended. This will also constitute a domain of application for autotelic

agents, as an alternative approach within the growing literature aiming to train small models with advanced coding capabilities, through the automatic generation of high quality synthetic data. Then, in a second step, we will consider extending this approach to the invention and solving of maths problems, leveraging formal correspondences between code and proof systems. In the context of the Inria “Défi” LLM4Code (see below), this will involve having autotelic LLMs learn to interact with automatic proof systems like Coq or Lean, an building on recent attempts to use LLMs as interfaces between mathematicians and proof assistants (Shulman, 2024), considering Coq as an external cognitive tool learned using techniques described in previous paragraphs.

Pursuing long-term goals. Current autotelic agents mostly pursue goals at the timescale of an episode. Humans, on the other hand, can pursue goals they can barely hope to achieve within their lifetime (e.g. building an efficient fusion reactor). Because there is an infinity of potential goals and little time to explore them alone, autotelic agents may need cultural models to bias their selection of long-term goals towards more feasible, interesting, or valuable options—turning an individual exploration into a population-based exploration. Keeping long-term goals in mind will require improvements in architecture’s memory systems, but might also benefit from language and culture. Indeed, verbalisation is known to increase humans’ memory span (Flavell et al., 1966) and writing lets us set our goals in stone. Young children progressively become future-oriented as they are taught to project themselves into the future through education, social interactions (what do you want to do when you grow up?) and cultural metaphors (e.g. the self-made man). If autotelic agents will need better hierarchical RL algorithms to achieve long-term goals, they could also leverage cultural artefacts evolved for improved collaboration and long-term planning (Colas et al., 2022) — think of roadmaps, organisation systems and project management tools. Because long-term goals are not immediately rewarding, human cultures supply shorter-term reward systems to track progress on these goals (good grades in the educational system, money and social recognition in professional careers)— a form of reward shaping. Studying this perspective will be a long term goal of this project proposal.

Evaluation approaches: As reviewed in Colas et al. (2022b), progress towards the scientific objectives mentioned above will be evaluated using diverse sets of measures quantifying exploration structure, the (meta-)diversity of discovered skills, the capability to generalise on out-of-distribution tests, fast adaptation and transfer learning to initially unknown downstream tasks, robustness to unknown evolutions of environments, and using explainability techniques when opening the black-box of the learned systems.

7.2.3 Objective 2.3: Cultural evolution, innovation and curiosity at the group level

The focus of the planned research described in previous paragraphs was on how language could empower autotelic agents to achieve long-term creative exploration and learning. Beyond language, the social structure in which autotelic agents are embedded introduces both opportunities and challenges, associated with further open scientific questions, for curiosity-driven exploration: this is the focus of this second dimension of our research program.

Interaction between curiosity in individuals and collective innovation dynamics. First, we will study how autotelic agents can learn to coordinate to achieve goals they are curious about. A challenge here is that various autotelic agents in a shared environment may have conflicting

self-generated goals, making their exploration more difficult. On the other hand, some goals individual agents may imagine and set for themselves may not be reachable if they try alone: in that case, coordinating with others to achieve joint self-generated goals may enable to learn and discover more complex goals (e.g. if some goals require different prior competencies in a population of diverse agents). Thus, we will study which mechanisms can enable autotelic agents to negotiate joint goals that serve their individual curiosity (we started first steps in recent work, Nisioti et al., 2023b), for example by playing coordination games akin to language games. Further, we hypothesize that a population of autotelic agents may increase several kinds of diversity as compared to single autotelic agents: diversity of kinds of goals explored (meta-diversity), but also diversity of ways to solve the goals, leading to a diversity of exploration and learning curricula. This may help address some of the exploration difficulties in complex and/or deceptive environments, e.g. with many local minima, as hinted in our recent study studying the role of social network topology in collective innovation (Nisioti et al., 2022). Finally, we aim integrating this perspective on collective innovation dynamics in groups of autotelic agents with our application domain on automated scientific discovery, by studying how groups of autotelic agents could learn to coordinate to help human scientists make discoveries in physics, chemistry, biology or computer sciences, i.e. by building **hybrid science labs mixing artificial agents and human scientists**.

Curiosity to learn about others and social interaction: In several modeling works we achieved in the past, we proposed the hypothesis that general curiosity-driven learning mechanisms could be forces that push individuals to explore vocal and language interactions with others, and thus be a key mechanism at the origins of language in individuals and societies (Oudeyer and Smith, 2026 ; Moulin-Frier et al., 2014 ; Forestier and Oudeyer, 2017). However, many open questions remain in this direction. First, beyond language, an open-question is to understand what are the mechanisms needed to enable efficient learning of diverse and complex social skills beyond language, involving for example theory of mind, joint intentionality, morality and social norms in general. We recently developed an experimental framework, called SocialAI School (Kovač et al., 2021), where learning agents are situated in environments where they need to learn such skills. Our preliminary experiments show that traditional RL approaches fail to learn such skills efficiently. We aim to study whether forms of curiosity and autotelic learning could help in learning these skills. In particular, beyond general forms of curiosity, we aim to consider and model computationally a special form of curiosity towards understanding others and social interaction, i.e. social curiosity (Philips, 2016)). We also aim to study how such mechanisms for social curiosity, beyond fostering learning social skills at the level of individuals, may drive groups of agents to self-organize macro-structures at the social level, such as special interaction topologies, which may in turn influence the general dynamics of collective innovation at the group level, as well as the potential regularities within the space of cultural artefacts, stories, conventions, formed at the group level.

Collective exploration and innovation: curiosity at the group level. As we discussed in previous paragraphs, groups of agents can display forms of collective intelligence to solve certain kinds of tasks. We are interested in exploring the possibility to consider the group as an integrated higher-level entity with its own form of curiosity, and study how groups may display structured and coordinated forms of exploration and diversity search. In brief, can curiosity

operate at the group level? Here, we aim to reuse the IMGEP framework (Forestier et al., 2022) to study how simple mechanisms at the individual level (including simple imitation with variation of goals considered by agents in the neighbourhood, and simple forms of knowledge sharing) can self-organise structured goal exploration at the group level (e.g. expanding the diversity of explored and learned goals while avoiding goals that are either trivial or unlearnable)?

Cultural evolution in groups of generative AI agents, and in hybrid human-machine groups. As generative AI techniques play an increasing role in the daily access, learning and creation of cultural content, their influence on human cultural evolution is likely to become a very important societal issue. Thus the fundamental study of the dynamics of cultural evolution in groups of generative AI machines, and in hybrid human-AI groups, a domain recently called Machine Culture (Brinkmann et al., 2023), has become highly relevant, and yet has been overlooked so far by both the AI and the human cultural evolution communities. We aim to contribute to the development of this field by broadening our studies of cultural evolution beyond their links to curiosity, in particular through systematically studying how various interaction protocols, various topological structures, various combinations of simulated personas, various combinations of foundational models with diverse capabilities, can undergo cultural evolution. A first step in this direction will be achieved by modelling dynamics of story formation and evolution in groups of artificial agents, and in hybrid groups of artificial agents and humans, and compare these results to experiments achieved in humans from various cultural backgrounds (this will be done in collaboration with Maxime Derex, a researcher in human cultural evolution in IAST, Toulouse). Given we here use LLMs as both models and simulated individuals in models of cultural evolution, we also aim to advance understanding of how culture is encoded in foundational models, and what is its level of robustness and controllability. We recently started steps in this direction by studying the robustness and controllability of how LLMs express values in various kinds of contexts (Kovač et al., 2024), as well as cultural evolution in populations of LLMs, allowing the manipulation of variables such as network structure, personality, and the way social information is aggregated and transformed (Perez et al., 2024).

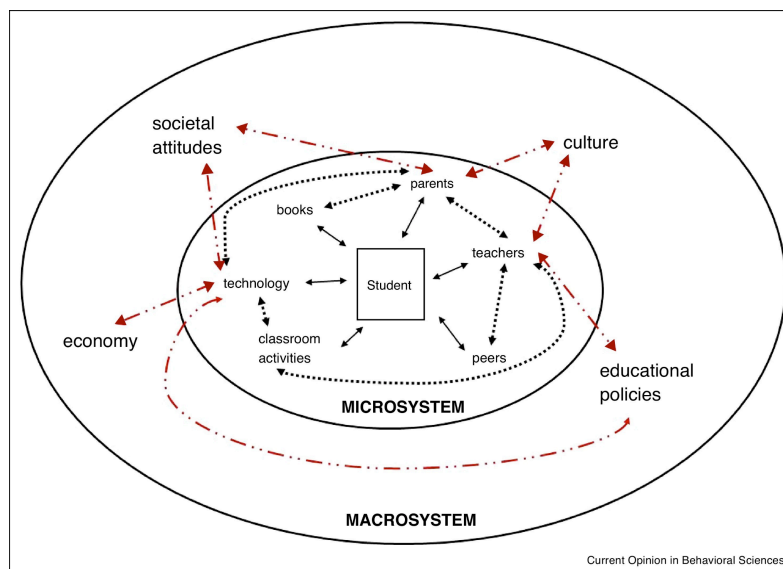
Evolutionary origins of autotelic learning. Although the ability to acquire an open-ended repertoire of skills is fundamentally related to the characteristics of human intelligence, research in AI rarely considers the mechanisms that may have guided the emergence of this ability during the evolution of the human species. Human behavioural ecology (HBE) is a field studying how ecological constraints and opportunities could have shaped human behaviour throughout its evolutionary history (Maslin et al., 2015; Brown et al., 2011; Sear et al., 2007). In recent works we have started to study how theories and hypotheses from HBE could inform the design of novel environments and training paradigms in AI and in Artificial life. Our central hypothesis here is that evolving artificial agents in highly-dynamic simulated environments, inspired by theories in HBE on the role of environmental factors in early human evolution, can result in the emergence of proto-autotelic learning abilities from the bottom-up. In this context, autotelic learning can be conceived as an evolutionary adaptation to highly unpredictable environments, where evolving an ability to self-generate and pursue intrinsic goals could be an advantage, as a way to be better prepared for extreme environmental changes. Recent contributions on this topic concerns the proposition of a conceptual framework aiming at cross-fertilizing a dialogue between the HBE and AI communities (Moulin-Frier et al., 2022), as well simulation experiments

where agent's populations can adapt both evolutionarily and developmentally to environments varying at multiple spatio-temporal scales (using methods such as neuroevolution, Hamon et al., 2023, and meta reinforcement learning, Bornemann et al., 2023) and potentially construct their own ecological niches (Nisioti et al., 2023a). We aim to continue developing this new research direction along several dimensions: (1) proposing massively multi-agent simulation environments with complex dynamics inspired by theories on early human evolution ; (2) proposing novel bi-level adaptation algorithms inspired by the interplay between evolution and development in biology ; (3) exploring societal applications of the proposed framework to better understand the trajectory of the human species in relation to the ongoing climate change.

7.3 Objective 3: Applications in education and assisted scientific discovery

7.3.1: Objective 3.1: Train curiosity-driven autotelic learning in humans across the lifespan

Leveraging our fundamental research work at the crossroads of AI and cognitive sciences, we will continue developing application projects in the domain of educational technologies according to a systemic approach of education (e.g., Peterson, 2020)




Supporting curiosity in schools and classrooms according to Peterson (2020)

The general objective is to design educational technologies leveraging AI models for training curiosity, meta-cognition and creativity across the lifespan. One key objective is to address the societal challenges inherent to the School of the 21st Century, in particular the inclusive school (learners with disabilities, especially cognitive); cross-disciplinary skills such as attention, curiosity, learning to learn, perseverance; digital/programming education and the digital transformation of the school. Part of this work is covered by the 2022-2025 research program of

the associated "Curious tech" team with the University of Waterloo in Canada (E. Law from CHI Lab and M. Fernandes from The Cognitive neuroscience lab) (<https://flowers.inria.fr/curioustech-associate-team/>)

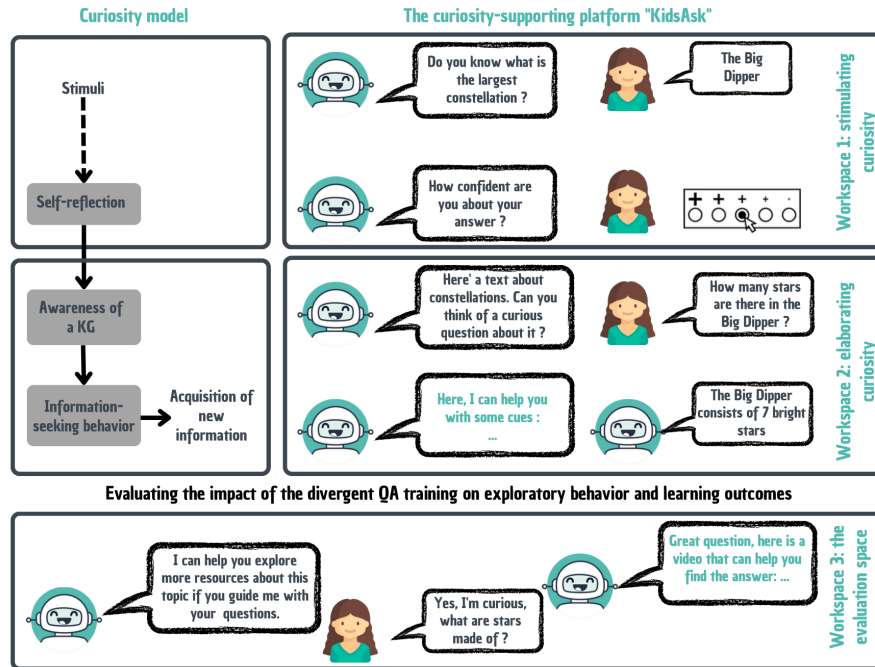
The LP theory for personalising sequences of exercises in cognitive training in neurodiverse populations and across the lifespan. First, we will study how to extend the ZPDES approach (leveraging the LP theory and multi-armed bandit to personalise sequences of exercises in intelligent tutoring systems, Clément et al., 2015, 2024; Delmas et al., 2018) to new domains, in particular cognitive training of attention and meta-cognitive control, but also learning languages. Also, we will aim to extend the evaluation of this approach when targeted to diverse populations in terms of age (lifespan, including ageing populations thanks the collaboration with the national network of free-time universities and with the newly created Vascular Brain Health Institute - VBHI - at Bordeaux), cognitive capabilities (e.g. neurotypical or neurodiverse thanks collaboration with Bordeaux hospital but also with the [atypiefriendly](#) network for inclusive universities in France), or occupations (e.g. training attentional skills in professional sport players; or training new video game players to enter more easily complex multi-player games).



We design algorithms (e.g. ZPDES, Clément et al., 2015; 2024) that personalise learning curricula in intelligent tutoring systems, and conduct field experiments in classrooms evaluating their impact on learning efficiency within the domain (e.g. maths), motivation and metacognition.

Generative AI for educational content creation (exercises, explanations, feedback). Second, as a limit of the ZPDES approach is that it requires human experts to format pedagogical content in a graph and generate many variations of exercises of a given type, we will also study how generative AI techniques (and in particular structured hierarchical prompting techniques) can be used to automate the generation and automatic correction of exercises, e.g. given traditional textbooks. This approach will also be leveraged to train smaller LLMs with

high-quality datasets in order to incorporate such automatic exercise generation in educational systems that are lightweight and do not rely on using LLMs hosted in foreign countries. This approach will finally be used more specifically for automatically generating programming exercises for computer science education, leveraging the work on autotelic LLMs that learn how to code as described above.



In the KidsAsk platform, we use conversational agents to train children to ask curiosity-driven questions, as well as related meta-cognitive skills. We also study how generative AI techniques can be used to generate automatically efficient cues in this training process. Such educational interventions are evaluated in the field, in primary schools (Abdelghani et al., 2022, 2023b)

Training meta-cognition, curiosity and stimulating agency. Third, we will scale up work we began to train metacognitive skills and curious question asking in children, going further the proof-of-concept by framing it within the LP theory and experimentally studying the role of language in these meta-cognitive training interventions (in synergy with the cognitive science study described above). Then, we will extend the study of the impact of such metacognitive training to a set of creativity abilities, as measured by standard psychological tools in psychology. We will also consider using such training, and measuring impact on creativity, for various adult populations. More precisely, we will work on the following dimensions:

- **Metacognitive training of curiosity for fostering agency in learning:** We will aim to develop pedagogical intervention aiming to help children develop their meta-cognitive perception and understanding of their own learning progress, so that they could be empowered to become efficient in self-generating a learning curriculum when faced with freedom to organise themselves (thus, one aim would be to train them in such a way that they could adequately personalise their own sequences of exercises without needing an EdTech system like ZPDES above, which is currently an efficient training tool also helping

considerably to progress in our understanding of curiosity-driven learning and personalization, but requires the use of technology, thus limiting its use in a number of contexts).

- ***Training creativity through curiosity and meta-cognition intervention.*** Related to our fundamental study of the links between curiosity and creativity, we aim to demonstrate that curiosity training as a process of knowledge maximisation - or as metacognitive feeling (Goupil & Proust, 2023), leads to better initiation of the creative process. An educational technology that reinforces divergent thinking in curiosity-driven processes, is expected to both improve learning without the discomfort of cognitive load, and facilitate creative thinking processes mobilised in a large set of problem-solving tasks (deduction, inference, association or analogy). We will design experimental protocols comparing new conversational agents providing verbal incentives to foster curiosity or creativity, in learning and problem-solving tasks, and conduct them in various age ranges (from 6 to 11 yo) thanks to the Léa-ifié Network established with 10 primary schools around Bordeaux (project will be submitted in next spring). We aim to study how curiosity's effectiveness may vary (or may be invariant) during this critical period during which metacognitive mechanisms develop.
- ***Teacher training to instil curiosity in classrooms, and then in schools:*** we will also consider how to train teachers (as primary social support) so that they learn adequate strategies to foster curiosity, metacognition and creativity in the classroom. Here, we aim to empower teachers to implement training of metacognition and curiosity-driven learning in the classroom, through an appropriation of our methods, particularly in STEM (science, technology, engineering, and mathematics) subjects. This challenge is particularly critical as it has been shown that the teacher has a critical role to play in shaping and spreading curiosity-based learning (Peterson, 2020). So we're planning studies to evaluate curiosity training for teachers, so that they can use our educational technologies (KidsAsk, KidsReflect), as well as unplugged versions of these activities for wider accessibility in schools, and then compare the results in terms of the impact on children's learning and curiosity. We will also compare with our results obtained by seasoned experimenters convinced of the value of curiosity-driven learning. We also plan to conduct qualitative methods (focus groups, questionnaires on acceptance and adoption) to assess how teachers perceive our curiosity-based learning technologies and related pedagogical methods, how they might improve their practices, and how they might apply them concretely to teaching sequences they would have to create.

Stimulating social curiosity in educational contexts with neurodiverse populations.

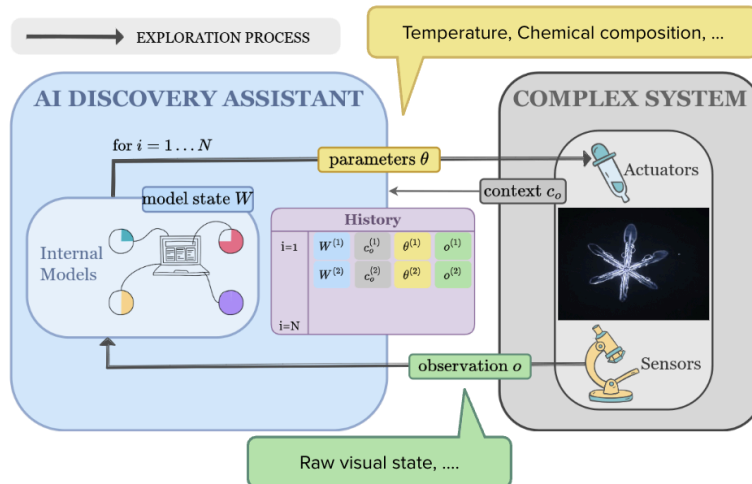
Neurodiverse populations face societal barriers due to their atypical cognitive and behavioural profiles. In educational technologies, traditional approaches for inclusion focus on matching technologies to individual needs, ignoring the broader (social) environment as a support for inclusion. This places the burden of overcoming disabilities on the individual rather than creating inclusive and disability-friendly environments. Family-professional partnerships are essential for supporting and improving children's academic performance and social adaptation. Thus, in an original way, we are developing educational programs aimed at improving the culture and social perceptions of neurodiversity of the neurodiverse student's entourage, in order to create

conditions for social interactions that are favourable for everyone. (Mazon et al., 2023; Saint-supery, et al., 2023). Promoting collaboration can be approached through social curiosity, with interventions fostering the desire of each stakeholder to understand others' feelings and thoughts and increasing their motivation to interact with each other (Meyer et al., 2024). Indeed, social curiosity has been positively associated with improved social functioning, enhanced behavioural adaptation, and the ability to establish and maintain relationships (Philips, 2016), which are also essential ingredients in building the relationships of trust needed to promote collaborative practices. We aim to design and to assess new technology promoting social curiosity across stakeholders, fostering dialog as a lever of success for a shared, enriched and powerful social representation of the neurodiverse student, with as a result an improved well-being for all stakeholders (student, parents, teachers and caregivers).

Generative AI to evaluate edTech systems and human learning/training outcomes. A key dimension of societal application like education is that they require strong and systematic evaluations. One dimension of these evaluations consists in analysing verbal productions from students and other stakeholders, e.g. think aloud reports during experimentations, responses to psychology questionnaires, and interviews. To analyse such kind of data, qualitative analysis, and deducting coding methods in particular, are the reference methods. However, they require a large amount of expert time, which limits the possibility to do qualitative analysis and deductive coding at a large scale, thus limiting evaluation. To address this limit, we will build on our preliminary work exploring the use of generative AI tools to automate deductive coding and assess this automation to ensure reliability in the interpretation of results (Xiao et al., 2023). Because of the high stakes of such methods in psychology and the human sciences at large, we will aim to release open-source tutorials enabling researchers in these disciplines to reuse and leverage these methods.

This applicative research program leverages various collaborations with stake-holders, ranging from public educational institutions (e.g. Académie de Bordeaux) to industry (e.g. EvidenceB, OnePoint, Ubisoft, Catie) and NGOs (e.g. France IOI).

7.4 Objective 3.2: Assisted scientific discovery with autotelic exploration algorithms




Curiosity-driven AI discovery assistants can enable chemists, physicists or biologists to explore efficiently the landscape of self-organized structures in complex systems (figure adapted from Etcheverry, 2023)

In many areas of natural and engineering sciences, researchers aim to study complex systems that can self-organize diverse structures and properties (e.g. new materials, biological networks such as GRNs, etc). **One major challenge they face at the early stage of research is that it is very hard to explore and build maps of the possible behaviours of such systems: they are high-dimensional, and at the same time there are often scarce resources of time and energy to explore them. Furthermore, scientists often do not initially know how to represent structures and behaviours of a new complex system.** We have recently showed proof-of-concepts that **curiosity-driven AI exploration algorithms**(especially autotelic learning systems that incrementally learn goal representations) could be very useful to help scientist **explore and learn models of these systems and learn representations of the self-organized behaviours** (e.g. [Blog post](#)). These proof-of-concept were mostly focused so far on abstract models of morphogenetic processes such as continuous cellular automata (e.g. [Etcheverry et al., 2020](#)). Recently, we collaborated with [Michael Levin's biology lab in Harvard](#) and showed how this approach can be used to map the space and characterize new behaviours in models of **gene regulatory networks**, setting the stage for new approaches to design or leverage self-organized functionalities in biological networks (see [preprint](#)).

In the coming years, our major objective will be to extend these methods and show how they can be used to study actual **physical, chemical and biological systems in the real world**. We will leverage the ongoing collaboration with M. Levin, as well as another developing collaboration with **bio-physicists at University of Chicago (A. Murugan's team)** to explore the space of behaviours of soft condensed matter, and finally a developing collaboration with [Cyril Aymonier](#) and the ICMCB research institute (Chemistry) in Bordeaux. We have also started

a collaborative project with IRT Saint Exupery in Toulouse for applying our automated discovery framework to better understand interference phenomena in multi-core hardware architecture (which can be viewed as an instance of a particularly complex system). Other potential applications of this framework concern the management of decentralized power network (contact with T. Dragicevic, Aalborg Univeristy, Danemark), as well as the exploration of synergies in neuro-muscular (contact with INCIA) and speech production (contact with Dan Dediu from the University of Barcelona) models. Besides science, we also plan to explore applications in the artistic domain, for example in the context of 3D visual animations (contact with Andy Lomas) or sound synthesis.



Artificial Intelligence (AI) is interested in implementing artificial agents able to learn an open-ended repertoire of skills (Left, Stooke et al. 2023). Artificial Life (Alife) is interested in discovering open-ended dynamics from the local interaction of simple elements (Right, Chan, 2020). We propose to integrate methods from both field with the objective of: A. discovering agents with behavioral functionalities self-organizing from the dynamics of a continuous cellular automata (CA). B. discovering update rules of a CA resulting in the open-ended evolution of diverse species (see Hamon et al., 2024; Plantec et al, 2023).

Exploring the role of environment dynamics in the origins of sensorimotor agency, memory, learning, autotelic learning in Lenia. We are also actively contributing to the domain of Artificial Life, in relation to one of the most fundamental question in Science: the origins of life and of Darwinian evolution. In this context, we have applied our automated discovery algorithms to state-of-the-art continuous Cellular Automata, which are considered as a test bed for the computational study of the origins of life (we are actively collaborating with Bert Chan from Google Deepmind Tokyo, creator of the Lenia cellular automata). In Hamon et al. (2024), we show how our approach can enable the discovery of self-organizing patterns displaying forms of sensorimotor agency, i.e. “artificial creatures” with emergent self-constitution, self-maintenance and behavioral functionality. In Plantec et al. (2023), we extend the Lenia cellular automata to a mass-conservation constraint, allowing to embed update rules locally in the system and enabling the simulation of an artificial “primitive soup”; showing how diverse creatures can emerge and interact from an environment containing originally no agent. This last contribution obtained the Best Paper Award at the ALife conference in 2023.

Finally, we will also explore how autotelic exploration algorithms can be used to help mathematicians explore a space of new problems and their proofs, in the context of the **LLM4Code Inria défi**, and leveraging work described above on autotelic LLMs learning diverse coding skills (thanks to similarities between computer programs and maths proofs), as well as a developing collaboration with the [PiCube](#) Inria team.