

Cédric COLAS, PhD

Website: ccolas.github.io
GScholar: [VBz8gZ4AAAAJ](https://scholar.google.com/citations?user=VBz8gZ4AAAAJ)
Github: [ccolas](https://github.com/ccolas)
Twitter: [@cedcolas](https://twitter.com/cedcolas)

Education

- Nov 2017 - **PhD in Artificial Intelligence**, *INRIA - Flowers Lab*, Bordeaux, FR.
- June 2021 Title: *Towards Vygotskian Autotelic Agents: Learning Skills with Goals, Language and Intrinsically Motivated Deep Reinforcement Learning*. See list of publications below. Defended on June 30, 2021. I obtained the *Best PhD Thesis Award 2022* in France, from the French Artificial Intelligence Association (AFIA) with a cash prize of 1000 EUR.
- 2016-2017 **Master in Cognitive Science**, *École Normale Supérieure*, Paris, FR.
Top-tier French university. Main subjects: Cognitive Neuroscience of the Prefrontal Cortex, Human Reasoning, Human Psychology Experimentations, Neuroscience of Consciousness. Results: 16.98/20, ranked 2/48.
- 2015-2016 **MSc in Biomedical Engineering**, *Imperial College*, London, UK.
Top-tier English university. Neurotechnology Stream. Main subjects: Biomedical Imaging, Speech Processing, Image Processing, Computational Neurosciences, Brain-Machine Interfaces. Results: 78/100, with distinctions.
- 2013-2016 **BSc and MSc in Electrical Engineering, Computer Science and Telecommunications**, *Supelec*, Gif-sur-Yvette, FR.
Now CentraleSupélec, 2nd best French engineering school (Usine Nouvelle ranking). The biomedical engineering MSc degree was an exchange year unlocking Supélec's MSc. Main subjects: Algorithmic, Signal Processing, Optimization, Statistics, Probability Theory. GPA: 3.7/4.

Research Projects

- Nov 2017 - **PhD Projects**, *INRIA - Flowers Lab*, Bordeaux, FR.
- June 2021 See list of publications below. Highlights:
- Supervised by Dr. Pierre-Yves Oudeyer and Pr. Olivier Sigaud;
 - Technical skills: self-taught in Python, Pytorch, Tensorflow, computing infrastructures, LaTeX;
 - Doctoral training: Ethics, Scientific Integrity, Answer Set Programming, Random Forest, Reproducibility, Ecological Impact of Research, Epistemology of Science;
 - Supervision experience: 2 master students, 2 bachelor students;
 - Developed collaborations with UberAI labs, Microsoft Research Montreal/Cambridge, OpenAI, SISTM team (Inria) and CNRS.
- Jun-Sept 2019 **Research Internship**, *Uber AI Labs*, San Francisco, US.
Project: scaling the Quality-Diversity algorithm *Map-Elites* to Deep Neuroevolution via the use of Evolution Strategies. Highlights:
- Supervised by Dr. Vashisht Madhavan and Dr. Jeff Clune;
 - Research experience: International collaboration, in a private research lab in San Francisco;
 - Technical skills: use of massively parallel computing (1000 cpus);
 - Complementary research skills: Deep NeuroEvolution (deep Genetic Algorithms, Evolution Strategies), Quality-Diversity (Map-Elites, Novelty-Search (NS), NS-ES).
 - Led to a publication at GECCO 2020 (see below).
- Jan-Jun 2017 **Master Project**, *Brain and Spine Institute - Motivation, Brain and Behavior Lab*, Paris, FR.
Project: computational models of the exploration-exploitation dilemma in a two-armed bandit task using variational Bayesian inference. Highlights:
- Supervised by Dr. Jean Daunizeau;
 - Complementary research skills: human cognitive modeling, models of human curiosity, variational Bayesian inference, model selection, navigating cognitive science literature.

- May-Sept 2016 **MSc Project**, *Imperial College - Brain and Behaviour Lab*, London, UK.
 Project: design of a brain-machine interface using EEG and convolutional neural networks to control an avatar in a video game for the international Cybathlon competition. Highlights:
 - Supervised by Dr. Aldo Faisal;
 - Technical skills: EEG setup, EEG data processing, Python, Theano, CNNs;
 - Research experience: collaborative project with colleague Dr. Pablo Ortega; deployment in an international competition (Cybathlon 2016, Zurich).
- Apr 2016 **MSc Project**, *Imperial College - Brain-Machine Interfaces Class*, London, UK.
 Project: offline decoding of a monkey's hand trajectories from 98 neuronal spike trains. Highlights:
 - Research experience: team project (team of 4);
 - Technical skills: feature selection, random forest, linear regression, nearest neighbors;
 - Ranked 2nd.
- Jul-Aug 2015 **Internship**, *Center of Psychiatry and Neuroscience*, Paris, FR.
 Project: I assisted a PhD student in the development of a fear renewal protocol in rats exploring wide environments. I setup the controlled experiment (rat conditioning, camera for movement detection, automatic protocol for stimuli).

Publications (peer-reviewed, main author)

- RL **Colas, C.**, Karch, T., Sigaud, O. & Oudeyer, P-Y. (2021). *Intrinsically Motivated Goal-Conditioned Reinforcement Learning: a Short Survey*. **Accepted at JAIR**. A useful review and framework to think about goal-conditioned RL approaches and their relation to intrinsically motivated learning and the quest for autonomous open-ended agents. Link: arxiv.org/pdf/2012.09830.pdf.
- Epidemiology and RL **Colas, C.**, Hejblum, B., Rouillon, S., Thiébaud, R., Oudeyer, P-Y., Moulin-Frier, C. & Prague, M. (2020). *EpidemiOptim: A Toolbox for the Optimization of Control Policies in Epidemiological Models*. Accepted to **JAIR**. General software library developed in collaboration with epidemiologists to facilitate the design and comparisons of epidemiological models and automatic intervention strategies to fight epidemics. Link: arxiv.org/pdf/2010.04452.pdf.
- RL Akakzia A., **Colas, C.**, Oudeyer, P-Y., Chetouani, M. & Sigaud, O. (2020). *Grounding Language to Autonomously-Acquired Skills via Goal Generation*. Accepted at **ICLR 2021**. An intrinsically motivated agents that builds its own curriculum to solve complex stacking tasks autonomously. Link: arxiv.org/pdf/2006.07185.pdf. Equal contributions of the two first authors.
- RL **Colas, C.**, Akakzia A., Oudeyer, P-Y., Chetouani, M. & Sigaud, O. (2020). *Language-Conditioned Goal Generation: a New Approach to Language Grounding for RL*. Accepted at **LaReL workshop, ICLR 2020**, 2nd best paper. A new model that turns linguistic descriptions into concrete goal embeddings for agents to target. This enables efficient generalization, and behavioral diversity, as several goals can be sampled for a given linguistic description. Link: arxiv.org/pdf/2006.07043.pdf.
- Evolutionary Algorithms **Colas, C.**, Huizinga, J., Madhavan, V., & Clune, J. (2020). *Scaling MAP-Elites to Deep Neuroevolution*. Accepted at **GECCO 2020**. Scaling of the most popular Quality-Diversity algorithm to the era of deep neural networks. New problems can be solved (hard exploration problems, high-dimensional behavioral adaptation). Link: arxiv.org/pdf/2003.01825.pdf.
- RL **Colas, C.**, Karch, T., Lair, N., Dussoux, J. M., Moulin-Frier, C., Dominey, PF. & Oudeyer, P-Y. (2020). *Language as a Cognitive Tool to Imagine Goals in Curiosity-Driven Exploration*. Accepted at **NeurIPS 2020**. The first autotelic agent using language as a cognitive tool to support the imagination of goals. Is currently leading to several PhD and postdoc projects in the lab. Link: arxiv.org/pdf/2002.09253.pdf.
- RL **Colas, C.**, Sigaud, O. & Oudeyer, P-Y. (2018). *CURIIOUS: Intrinsically Motivated Modular Multi-Goal Reinforcement Learning*. Accepted at **ICML 2019**. The first autotelic agents based on state-of-the-art deep reinforcement learning methods. It uses automatic curriculum learning to craft its own learning trajectory learn to target multiple goals involving multiple types of affordances. Link: arxiv.org/pdf/1810.06284.pdf.

- RL Colas, C., Sigaud, O. & Oudeyer, P-Y. (2018). *GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms*. Accepted at **ICML 2018**. First steps to merge developmental robotics autotelic methods with deep reinforcement learning. An intrinsically motivated exploration helps bootstrap a standard Deep RL algorithm. This led to multiple works on decoupling exploration and exploitation in RL. Link: arxiv.org/pdf/1802.05054.pdf.

Publications (peer-reviewed, second author)

- RL Portelas, R., Colas, C., Weng, L., Hofmann, K. & Oudeyer, P-Y. (2020) *Automatic Curriculum Learning For Deep RL: A Short Survey*. Accepted at **IJCAI 2020**. A useful survey and framework to think about automatic curriculum learning methods in deep RL. Link: arxiv.org/pdf/2003.04664.pdf.
- RL Portelas, R., Colas, C., Hofmann, K. & Oudeyer, P-Y. (2019). *Teacher Algorithms for Curriculum Learning of Deep RL in Continuously Parameterized Environments*. Accepted at **CoRL 2019**. A novel automatic curriculum learning method to shape agents learning trajectory in complex domains. Link: arxiv.org/pdf/1910.07224.pdf.
- RL Fournier P., Colas, C., Chetouani, M. & Sigaud, O. (2019). *CLIC: Curriculum Learning and Imitation for feature Control in non-rewarding environments*. Accepted at **IEEE Transactions on Cognitive and Developmental Systems**. A discrete variant of the CURIIOUS algorithm that integrates natural demonstrations. Link: arxiv.org/pdf/1901.09720.pdf.
- BCI Ortega, P., Colas, C. & Faisal, AA. (2018). *Compact Convolutional Neural Networks for Multi-Class, Personalised, Closed-Loop EEG-BCI*. In 2018 7th **IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)** (pp. 136-141). The first BCI to consider CNN implementations to allow disabled people to control machines. Link: ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8487644

Publications (preprints or in review)

- RL Sigaud, O., Caselles-Dupré, H., Colas, C., Akazia, A., Oudeyer, P-Y. & Chetouani, M. (2021). *Towards Teachable Autonomous Agents*. **In review**. An opinion paper on the need for agents that are both intrinsically motivated and interactive (teachable). Link: arxiv.org/pdf/2105.11977.pdf.
- Stats for RL Colas, C., Sigaud, O. & Oudeyer, P-Y. (2019). *A Hitchhiker's Guide to Statistical Comparisons of Reinforcement Learning Algorithms*. A useful comparative study of statistical methods to meaningfully compare RL algorithms, and associated guidelines. Link: arxiv.org/pdf/1904.06979.pdf.
- Stats for RL Colas, C., Sigaud, O. & Oudeyer, P-Y. (2018). *How Many Random Seeds? Statistical Power Analysis in Deep Reinforcement Learning Experiments*. A discussion of the lack of statistical testing in the comparison of RL algorithms. Link: arxiv.org/pdf/1806.08295.pdf.

Non-scientific publications

- Blog post *Language as a Cognitive Tool: Dall-E, Humans and Vygotskian RL Agents*. March 2021. Link: developmentalsystems.org/language_as_cognitive_tool_vygotskian_rl.
- Blog post *Intrinsically Motivated Modular Multi-Goal RL*. March 2020. Link: developmentalsystems.org/curious_intrinsically_motivated_multi_modular_goal_rl.
- Blog post *How Many Random Seeds?* February 2020. Link: developmentalsystems.org/how_many_random_seeds.
- Blog post *Bootstrapping Deep RL with Population-Based Diversity Search*. February 2020. Link: developmentalsystems.org/bootstrapping_rl_with_diversity.

Technological Development

- EpidemiOptim Online demonstration of a pre-trained lockdown policy in the context of simulated COVID19 Demo epidemics. epidemioptim.bordeaux.inria.fr.
- EpidemiOptim Library interfacing optimization algorithms with models of epidemic propagation. Epidemiological library models are wrapped in OpenAI Gym interfaces, making them readily compatible with state-of-the-art optimization algorithms. github.com/flowersteam/epidemioptim.

- ME-ES Software of the ME-ES architecture. github.com/uber-research/Map-Elites-Evolutionary.
- Playground RL learning benchmark. Designed to facilitate the study of systematic generalization in language-conditioned RL architectures. Easy to extend to more objects, more interactions and more complex language. github.com/flowersteam/playground_env.
- Imagine Software of the IMAGINE architecture. github.com/flowersteam/imagine.
- Curious Software of the CURIOUS architecture. github.com/flowersteam/curious.
- GEP-PG Software of the GEP-PG architecture. github.com/flowersteam/geppg.
- RL stats Library providing easy-to-use tools to compare the performance of RL algorithms. Allows the replication of the experiments from *A Hitchhiker's Guide to Statistical Comparisons of Reinforcement Learning Algorithms*. github.com/flowersteam/rl-difference-testing.

Communications

- PhD Defense *Towards Vygotskian Autotelic Learning: Learning Skills with Language, Goals and Intrinsically Motivated Deep Reinforcement Learning*. June 2021 (online). Link: youtube.com/watch?v=x4vS557rhAM.
- Invited Talk EpidemiOptim: A Toolbox for the Optimization of Control Policies in Epidemiological Models. DeepMind Seminar, January 2020 (online).
- Conference Talk *Scaling MAP-Elites to Deep Neuroevolution*. Accepted at GECCO 2020. GECCO conference, July 2020 (online). Link: youtube.com/watch?v=m2peevXlgKY.
- Conference Talk *Language as a Cognitive Tool to Imagine Goals in Curiosity-Driven Exploration*. NeurIPS conference, December 2020 (online). Link: slideslive.com/38937386/language-as-a-cognitive-tool-to-imagine-goals-in-curiosity-driven-exploration?ref=speaker-28788-popular.
- Conference Talk *CURIOUS: Intrinsically Motivated Modular Multi-Goal Reinforcement Learning*. ICML conference, Long Beach CA (US), June 2019. Link: youtube.videoken.com/embed/v-W4JSWUX28?tocitem=24.
- Conference Talk *GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms*. ICML conference; Stockholm (SW), June 2018. Link: youtu.be/MK-oAqHjdmg?t=2080.

Grants & Awards

- PhD Fellowship 3 years (2018-2021). Co-funded by INRIA and DGA (Direction Général de l'Armement, FR). 59k€ each.
- Thesis Award PhD Thesis Award 2022 from the French AI Association (AfIA).
- MSCA Postdoctoral Fellowship Marie Skłodowska-Curie Postdoctoral Fellowship for 3 years (2022-2025). Obtained for the project *Help Me Grow: Artificial Cognitive Development via Human-Agent Interactions Supported by New Interactive, Intrinsically Motivated Program Synthesis Methods*. 273k€.
- Phd Thesis Award Award for the Best PhD Thesis in France (2022) from the French Association for AI (AFIA). Prize: 1,000€.

Supervision Experience

- Feb-Aug 2021 I supervised two 2nd year master student from Univ Paris-Sorbonne (Paris, FR). With one of them, we worked on a new benchmark and dataset to train scene captioners to enhance cutting-edge autotelic reinforcement learning algorithms (to be released soon). With the other, we worked on the addition of minimal social interventions to help autotelic agents explore in a block manipulation domain (publication submitted for review).
- Feb-June 2020 I supervised two bachelor students from Ecole de Cognitive (Bordeaux, FR). I mentored them in their project on the parallels between cognitive science and recent developments in autonomous AI systems.
- Feb-Aug, 2018 I supervised a 2nd year master student from Ecole Normale Supérieure (Rennes, FR). We worked on coupling curiosity mechanisms with model-based reinforcement learning.

— Misc Training & Qualifications

Programming Languages Python (proficient), Matlab (proficient), LaTeX (proficient), Processing (intermediary), Arduino-C++ (beginner).

Languages French (native), English (fluent, 1 year in the UK, 4 months in the US), Spanish (beginner).