

Curiosity and Intrinsic Motivation for Autonomous Machine Learning

Pierre-Yves Oudeyer, Manuel Lopes, Celeste Kidd, Jacqueline Gottlieb

► **To cite this version:**

Pierre-Yves Oudeyer, Manuel Lopes, Celeste Kidd, Jacqueline Gottlieb. Curiosity and Intrinsic Motivation for Autonomous Machine Learning. ERCIM News, ERCIM, 2016, 107, pp.2. <<http://ercim-news.ercim.eu/en107/special/curiosity-and-intrinsic-motivation-for-autonomous-machine-learning>>. <hal-01404304>

HAL Id: hal-01404304

<https://hal.inria.fr/hal-01404304>

Submitted on 28 Nov 2016

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

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

Curiosity and Intrinsic Motivation for Autonomous Machine Learning

by Pierre-Yves Oudeyer, Manuel Lopes (Inria), Celeste Kidd (Univ. of Rochester) and Jacqueline Gottlieb (Univ. of Columbia)

Autonomous lifelong multitask learning is a grand challenge of artificial intelligence and robotics. Recent interdisciplinary research has been investigating a key ingredient to reach this goal: curiosity-driven exploration and intrinsic motivation.

A major difference between human learning and most current machine learning systems is that humans are capable of autonomously learning an open-ended repertoire of skills, often from very little data that they actively

for every new task the machines are given.

One of the major components that enables autonomous, open learning in humans is curiosity, a form of intrinsic

aims at pushing the frontiers of what we know about human active learning and how it can be built into machines.

Various strands of work in developmental robotics, AI and machine



Figure 1: Curiosity-driven learning in humans and robots (left: photo by Adam Fenster/Univ. Rochester; right: Milo Keller/ECAL).

collect themselves. Humans show an extraordinary capacity to adapt incrementally to new situations and new tasks. They proactively seek, select, and explore new information to develop skills before they are actually needed.

On the contrary, typical machine learning systems—including those associated with recent advances in deep (reinforcement) learning—learn to solve finite sets of tasks that are predefined by the engineer, and only by access to very large databases of examples. As a consequence, such machines require a new dedicated reward/cost function to be programmed by an engineer and time to reprocess millions of learning examples

motivation that pushes us to actively seek out information and practice new skills for the mere pleasure of learning and mastering them (as opposed to practicing them for extrinsic rewards such as money or social recognition). In the context of the interdisciplinary HFSP project “Curiosity”, Flowers team [L1] at Inria (France), Gottlieb Lab [L2] at Columbia University (US) and Kidd Lab [L2] at University of Rochester (US) are joining forces to study the mechanisms of curiosity-driven active learning in children, adults and monkeys and how they can be modelled and applied with machine learning systems. Mixing artificial intelligence, machine learning, psychology and neuroscience, this project

learning have begun to explore formal models of curiosity and intrinsic motivation (see [1] for a review), providing theoretical tools used in this project. In these models, curiosity is typically operationalised as a mechanism that selects which action to experiment or which (sub-)goals to pursue, based on various information-theoretic measures of their “interestingness”. Many such measures have already been studied with machines and robots—e.g., Bayesian surprise, uncertainty, information gain, learning progress or empowerment—and are often optimised within the reinforcement learning framework, where they are used as intrinsic rewards.

Such algorithmic systems were recently shown to allow machines to learn how to solve efficiently difficult tasks in which extrinsic rewards are rare or deceptive, precluding an easy solution through traditional reinforcement learning methods [1]. These systems were shown to allow robots to efficiently learn multiple fields of parameterised high-dimensional continuous action policies [1]. They also allow robots to self-organise their own learning curriculum, self-generating and self-selecting their own goals, showing a progressive development of new skills with stages that reproduce fundamental properties of human development, for example, in vocal development or tool use [2].

However, many open questions remain. For example, what are the features of interestingness that stimulate the curiosity of human brains? Can current computational models account for them, or be improved by taking inspiration from the heuristics used by humans? Are these mechanisms of

curiosity hardwired or adapted during lifelong learning? As curiosity is a form of guidance for exploration and data collection for autonomous machines, it is also possible to investigate how it can be combined with other forms of guidance used by human-like imitation. For example, in recent robotics experiments, curiosity-driven robots learn repertoires of skills by actively seeking help from human teachers [2].

Finally, curiosity has long been known to be key in fostering efficient education. Computational models of these mechanisms open the possibility for new kinds of educational technologies that could foster intrinsically motivated learning. In recent work, Clement et al. [3] showed one way by presenting active teaching algorithms that were capable of personalising sequences of pedagogical exercises (e.g., math exercises for primary school children), through the dynamic selection of exercises that maximise informational quantities such as learning progress.

Links:

- [L1] <https://flowers.inria.fr>
- [L2] <http://www.gottliebblab.com>
- [L3] <http://www.bcs.rochester.edu/people/ckidd/>

References:

- [1] P-Y. Oudeyer, J. Gottlieb, M. Lopes: “Intrinsic motivation, curiosity and learning: theory and applications in educational technologies”, *Progress in Brain Research*, 2016.
- [2] P-Y. Oudeyer, L. Smith: “How Evolution may work through Curiosity-driven Developmental Process”, *Topics in Cognitive Science*, 1-11, 2016.
- [3] B. Clement, D. Roy, P-Y. Oudeyer, M. Lopes: “Multi-Armed Bandits for Intelligent Tutoring Systems”, *Journal of Educational Data Mining (JEDM)*, Vol 7, No 2, 2015.

Please contact:

Pierre-Yves Oudeyer
Inria and Ensta ParisTech
pierre-yves.oudeyer@inria.fr

Applied Data Science: Using Machine Learning for Alarm Verification

by Jan Stampfli and Kurt Stockinger (Zurich University of Applied Sciences)

A novel alarm verification service applying various machine learning algorithms can identify false alarms.

False alarms triggered by sensors of alarm systems are a frequent and costly inconvenience for the emergency services and owners of alarm systems. Around 90% of false alarms are caused by either technical failures such as network downtimes or human error.

To remedy this problem, we develop a novel alarm verification service by leveraging the power of an alarm data warehouse. In addition, we apply various machine learning algorithms to identify false alarms. The goal of our system is to help human responders in their decision about whether or not to trigger costly intervention forces.

Approach

We are working with a security company that is a major player in secure alarm transmission with the aim of sig-

nificantly reducing the number of false alarms. Alarms can be triggered by devices installed at banks, jewellery stores, private homes, etc.

The problem of alarm prediction is conceptually similar to anomaly detection [1] or prediction of failures [2]. Hence, we can borrow some ideas from these fields and take advantage of the latest progress in deep learning [3]. We have chosen four machine learning approaches to predict false alarms:

- Random forests
- Support vector machines
- Logistic regression
- Deep neural networks.

We based our experiments on a set of more than 300,000 alarms. For each alarm we had multiple features such as device ID, device location, type of

alarm, alarm trigger time, etc. We used these features as input for our machine learning approaches. An overview of the system architecture is given in Figure 1.

For security reasons we do not have direct information about whether an alarm is a false alarm or not. However, we have indirect information that we can use as labels for our machine learning approach. In particular, for each alarm we know when it was triggered and when it was reset again. Our hypothesis is, that if the time difference Δt between triggering and resetting an alarm is small, there is a high chance that the alarm is false.

Consider the case of an alarm system deployed at a private home. Further assume that kids or a pet triggered an