CREATIVITY: A DRIVER FOR RESEARCH ON ROBOTICS IN OPEN ENVIRONMENTS

STEPHANE DONCIEUX

Abstract. Many robots solve complex tasks in closed environments that are fully known by the robot designer. Robots are much rarer in our every day environment. The main reasons are its complexity and openness that frequently result in unpredictable situations created by new objects or new dynamics of interaction. Numerous applications of robotics would benefit from robot abilities to deal with open environments. A lot of research is focused on this topic, but a question remains unanswered: how to evaluate the ability of a robot to face new contexts? The performance of a robot in a particular environment says little about what this performance will be in a different environment, even if it seems similar. A robot programmed to manipulate boxes may not be able to manipulate balls of the same size. Likewise, a robot perception system may be fooled by a change of luminosity. Consequently, the performance to expect in a new situation cannot be systematically deduced from what has been observed in known situations. There is then a need for a criterion that allows to compare approaches with respect to their ability to endow robots with the robustness or adaptive abilities to deal with new situations. To this end, it is proposed to define criteria based on the notion of creativity. A definition is given in the context of robotics and examples of use for learning and developmental processes are given.

1. Introduction

Robots are used on a daily basis in factories, where their environment is fully controlled and tuned to make their work fast and efficient. These successes can be attributed to adapted mechanical devices associated with efficient control laws. Artificial Intelligence has also obtained successes with the victory of a computer program against a professional human player in the game of Go [Silver et al., 2016], a notoriously hard game [Müller, 2002]. Researches in all fields related to robotics from mechanics to artificial intelligence have then reached a maturity that can be observed in these noticeable successes where most humans, if not all, are clearly outperformed by machines. Meanwhile, activities performed on a daily basis by humans, like washing dishes, folding laundry or even just emptying a dishwasher, remain challenges for robotics whereas every human can do it after training. What makes the difference between factory problems like 'pick and place', games like Go or Chess for which machines are now extremely efficient and emptying any kind of dishwasher containing any kind of tableware that is easy for humans but completely out of reach for robots? This gap mainly stems from the way the problem is defined. For factory problems like 'pick and place' or for the game of Go, the world is closed and static: all the concepts and rules are known beforehand and do not change along time. In a real-life context of a dishwasher emptying task, the tableware or the place where to put it may change. The dishwasher itself may also change. Many different models actually exist that have different features and are thus not used *exactly* in the same way. Changing features of the problem to solve requires the robot, its mechanical features, perception system and control system to be able to cope with it. In some applications, a fixed but robust controller associated to an adapted morphology can deal with new situations. Autonomous vacuum cleaners and other lawnmowers can be deployed in very different environments without any change. They have thus been able to invade the consumer market [Jones, 2006]. In other applications, like the dishwasher example, no robust robot has been found yet.

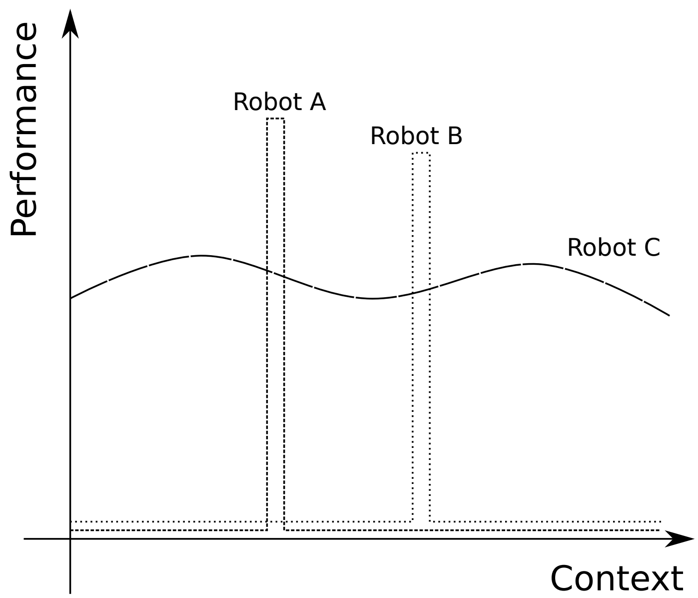

Figure 1. The performance of a robot is not necessarily constant over different contexts. Robots A and B are specialists, very efficient in a specific context, but poorly performing in others. Robot C generalizes better, but has a lower performance with respect to A and B in their specific contexts.

Demonstrations have been made of robots cooking [Bollini et al., 2013], folding towels [Maitin-Shepard et al., 2010] or emptying a dishwasher [Srinivasa et al., 2008]. These demonstrations prove that it is possible for a robot to perform such tasks, at least in the conditions used for these experiments, so why does it remain a challenge? A demonstration of a robot's behavior shows its ability to deal with a *particular* task in a *particular* environment. It does not demonstrate *at all* what its performance would look like for the same task, but in a different environment, even if the difference is limited, in particular if the robot designer is not here to adapt its software or hardware (Figure 1). The gap is between what is within reach for a robot in *known* situations and what is possible in situations whose precise details are unknown to the robot designer. These environments

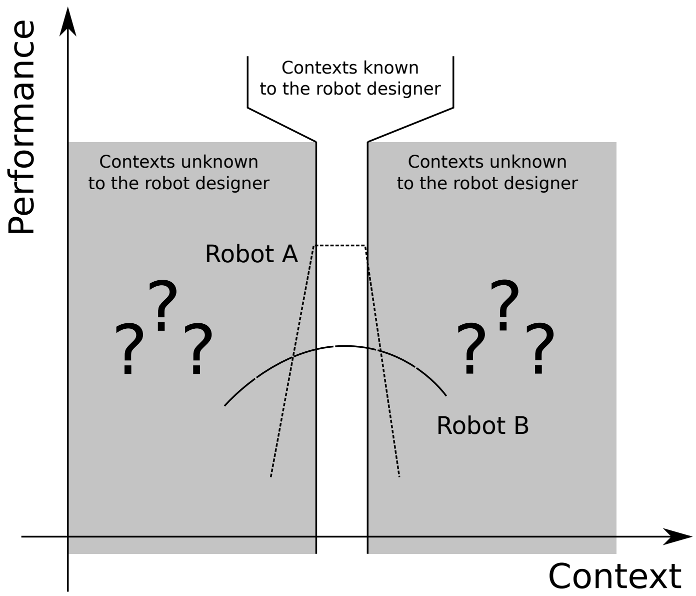

FIGURE 2. Dealing with an open environment implies to deal with contexts that may be unknown to the robot designer. Robot A performs better than B in known contexts, but B outperforms A in other contexts. This difference is not measurable during the design phase. The challenge is to find measures that can reflect this difference. Creativity is proposed to address this issue.

are called open environments. If a demonstration is not enough to estimate progresses on robotics applications in open environments, what is missing? It is argued here that there is a need for a new and specific criterion to better report the ability of a robot to deal with such open environments (Figure 2). I propose to take into account the knowledge available to the robot when it is switched on. Does it know the objects it will manipulate? Does it know their shape, size, weight, color, texture? Does it already know how to manipulate them? How to discriminate them from the background? Such choices are aimed at simplifying robot programming, but in the same time every choice made limits the range of environments and objects the robot can adapt to. A robot that can empty a dishwasher without precisely knowing the tableware or the dishwasher model, is expected to easily adapt to new tableware or dishwasher models, whereas a robot programmed for specific ones will hardly adapt to new situations. The new criterion should then take into account an originality or novelty criterion with respect to what is known beforehand. This originality criterion alone is not enough as it does not take into account the efficiency of the robot to fulfill the task. A random behavior can easily be reported as original, but it does not mean that it is interesting. Originality needs to be associated with performance. Creativity is defined as a process that can create both original and effective products [Stein, 1953, Runco and Jaeger, 2012]. I propose to use this notion of creativity to estimate the ability of robots to deal with open environments.

In the following, I propose a definition of creativity adapted to robotics in open environments and propose some measures that could allow Kuhn's normal science in this domain [Kuhn, 1970]. Learning and development are two kinds of approaches that are expected to generate the adaptive properties a robot would need to be able to face an open environment. I briefly propose how to take creativity into account in both approaches.

2. Creativity in robotics: a definition

Creativity is an important feature of human activity. As such, it has drawn a lot of attention to describe and characterize it [Stein, 1953, Guilford, 1967, Runco and Jaeger, 2012], find its neuroanatomical substrate [Dietrich, 2004] and try build computer programs that can generate creative products like paintings or drawings [Boden, 1998, Kowaliw et al., 2009]. It has been defined as the ability to design a product that is both original and efficient [Stein, 1953, Runco and Jaeger, 2012]. Originality corresponds to the core of this concept: a solution that is not new nor original will not be qualified as creative. Efficiency is also important to distinguish a creative solution from a randomly generated one. It is proposed here to start from this general definition and to refine it in the context of robot adaptation to new situations.

The product considered here is a policy and its observable result: the behavior, i.e. the sequence of observed robot states and actions in a particular environment together with their impact on this environment, including any available performance measure. The expression 'policy design process', noted $\mathcal{P}$, will be used to refer to any kind of decision, learning or developmental process that generates a policy in a particular context.

Assessing creativity is an important issue when considering artificial systems expected to generate creative products. It is required to compare competing approaches and drive research work towards the most creative systems. In the case of artistic creations, the evaluation can be attributed to the way humans react to the generated product [Colton et al., 2012]: can they make the difference between art works made by the computer generated and by a human artist? Would people buy it or vote for it? Relying on human evaluations to assess the creativity of robot behaviors is possible, but it is proposed here to define a quantitative criterion. It is a difficult task [Boden, 1998], but the special case of Robotics may provide some guidelines about how to proceed.

Evaluating efficiency requires specific criteria, if it may be difficult to define in general [Boden, 1998], such criteria are often already available in Robotics. Most learning algorithms are driven by a function estimating the performance of the behavior policy under consideration, may it be called value, reward or fitness function.

Originality is not so straightforward to get. It is not an absolute notion. It makes creativity difficult to assess in general as it depends on a cultural context [Boden, 1998] and on prior knowledge [Oxman, 1990]. Design is considered to be, in fact, *"a dynamic process of adaptation and transformation of the knowledge of prior experiences in order to accommodate them to the contingencies of the present"* [Oxman, 1990]. This general definition perfectly fits the robotics application considered here. The behaviors are built by the behavior design process that relies on the knowledge used to build and train it. The originality of a behavior can then be considered as relative to a database K of the knowledge available to the observer that estimates it and that will typically be the knowledge base used to build and train the behavior design process . The originality of a behavior policy π is expected to significantly change when it is measured with respect to knowledge bases $K - \{\pi\}$ and $K + \{\pi\}$.

Dealing with open environments requires to adapt to new and unforeseen *future* situations. As the future is, by definition in an open environment, not fully predictable, the proposed originality criterion shifts future to the present, and present to the past: the originality of a behavior found to solve a problem in the present time given what the system knew in the past, is expected to be an estimation of how the system will react to future situations given what it knows in the present.

Given a knowledge base and a behavior observed, how to estimate its originality? A behavior that is already present in the knowledge base is not expected to be original, but what about other behaviors? A behavior that can be deduced in a straightforward manner from the knowledge base is not expected to be original either. A robot that knows how to walk or run at a constant speed can actually reach any intermediate speed by alternating between the two behaviors. To what extent is it original with respect to its initial knowledge base? The definition of originality requires to define a process that uses the reference knowledge base. Given a knowledge base K, a robot policy design process $\mathcal{P}$ chooses a policy π in a particular context c, defined by a task and a domain[1]: $\pi = \mathcal{P}(K, c)$. The originality criterion needs to be defined with respect to a given reference policy design process $\mathcal{P}_{ref}$ and a reference knowledge base K_{ref}. The originality criterion estimates the ability of another policy design process $\mathcal{P}$ to go beyond $\mathcal{P}_{ref}$ capability with the *same* knowledge base K_{ref} in a particular context.

At least two different definitions have been proposed in the literature to compute novelty or originality. The first one relies on the notion of surprise and on the ability to predict (or not) observed patterns [Schmidhuber, 1991, Schmidhuber, 2010]. A behavior is considered as surprising if the robot cannot accurately predict the results of the corresponding motor commands. A robot with the ability to make predictions and observe how they change, can use this information as an artificial curiosity driving the predictive model learning [Schmidhuber, 1991, Oudeyer et al., 2007, Schmidhuber, 2010]. A policy with a behavior that differs from those predictions is then original. Described in the proposed formalism, K_{ref} is here the set of trained predictors at a particular moment and $\mathcal{P}_{ref}$ relies on these predictors to determine what to do. The second definition is similar, but does not rely on a prediction system. It considers originality as the novelty with respect to what has been observed so far [Lehman and Stanley, 2011]. A policy is considered as novel if its behavior is different enough from what was observed up to now. This definition allows to design a novelty metric as an average distance between the behavior of the policy to test and the behaviors in a set of observations [Lehman and Stanley, 2011]. K_{ref} in this case is the set of observed behaviors and corresponding policies, and $\mathcal{P}_{ref}$ is a process that chooses the

[1]The task is described by a reward function and the domain is the configuration of the environment.

policy to apply in K_{ref} with no modification. A policy generating a behavior that differs from those in K_{ref} is then original.

These different analyses lead to the following definition for creativity in a robotics context:

Definition: Given a policy π generated by a design process $\mathcal{P}$ in a context c with the help of a knowledge base K_{ref} ($\pi = \mathcal{P}(K_{ref}, c)$), the creativity of π is a $(v(\pi), \rho(\mathcal{P}_{ref}, K_{ref}, \pi))$ pair where:

- $v(\pi)$ is a value function evaluating the performance of a policy π ($v(.)$ defines a task);
- $\mathcal{P}_{ref}$ is a reference policy design process;
- $\rho(\mathcal{P}_{ref}, K_{ref}, \pi)$ is an originality metric that evaluates the originality of a policy π with respect to the use of knowledge base K_{ref} by a policy design process $\mathcal{P}_{ref}$.

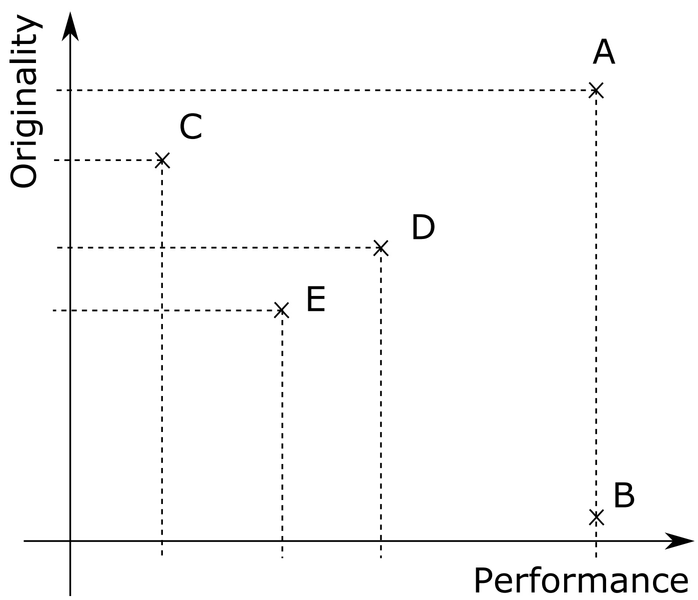

FIGURE 3. Comparison of the creativity of different behaviors. See text for details.

This definition associates both efficiency and originality criteria. Comparing two different problem solving approaches requires to take both into account. They can be aggregated with a weighted average, but this approach has a drawback: two studies using different weights can reach opposite conclusions. The Pareto dominance relations used in multi-objective optimization [Deb, 2001] can thus be more appropriate to reach a parameter-independent conclusion. In this relation, a solution y is said to dominate another solution x, if both conditions 1 and 2 are true:

(1) the solution y is not worse than x with respect to all objectives;
(2) the solution y is strictly better than x with respect to at least one objective.

With this definition, nothing can be said about C and D (Figure 3): C is more original but D is better performing. C neither dominates nor is dominated by D. E is dominated by D, as D is better on both objectives. B has the same performance than A, but less originality. It is thus less creative. A is then the most creative behavior in this example as it dominates B, C, D and E.

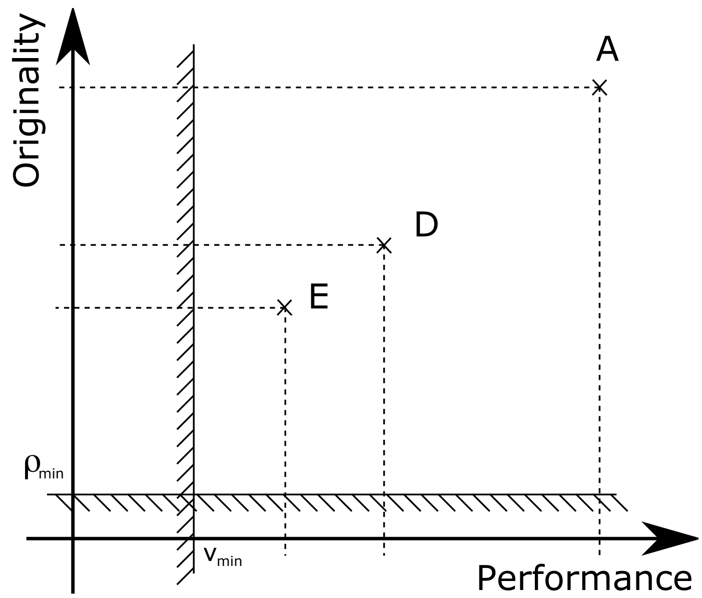

Figure 4. Comparison of creativity values with low boundaries for originality and performance. With these values, both B and C (Figure 3) are discarded.

Extreme solutions may not be interesting: a random solution may have the largest possible originality, as it can be generated with an empty knowledge base, but at the same time, it will have a poor performance. Likewise, a behavior generated with a system that has a full knowledge of the environment may have a high performance, but it has no originality. With the Pareto dominance relation, these two particular behaviors are equally creative, or more precisely, they are both non dominated. As both are of no real interest with respect to creativity as they are either not performing the task at all or completely task-specific, I suggest to define minimal values for both performance and originality and discard the behaviors that are below these limits (Figure 4). As the robot is expected to solve the task, at least partly, the lower bound on performance can reflect this minimal expected task resolution. The lower limit for originality may be chosen to reject all policies that can be deduced by $\mathcal{P}_{ref}$ from K_{ref}. This results in $\rho_{min} = \epsilon > 0$ (no matter how small ϵ is).

3. A simple example

To illustrate how this criterion could be used, the example of a robot that has to grasp a basket will be considered. This object is complex enough so that it can be grasped in several different ways.

The proposed reference behavioral design process $\mathcal{P}_{ref}$ consists in a program able to make the robot arm follow a particular trajectory, including a particular grasp strategy – grasping the basket at the handle, for instance. This is a typical low-level controller generated with methods from control theory. Given a particular trajectory – including gripper movement – it makes the robot follow it with as much stability and accuracy as possible. This behavioral design process is clearly not adaptive. If the basket is moved and if the given trajectory is unchanged, a robot arm following the behavior proposed by $\mathcal{P}_{ref}$ will still try to grasp it at the former position and will thus fail.

A second behavior design process $\mathcal{P}_{vis}$ may be a vision-based approach that determines the trajectory based on visual inputs [Hutchinson et al., 1996]. $\mathcal{P}_{vis}$ will typically rely on object visual features to track the object and determine robot movements to reach and grasp it[2]. It will rely on the knowledge of object visual features like its shape, color, etc. The grasping strategy is imposed by the robot programmer. It will be considered that it is the same as the one above: grasping by the basket handle. The main difference is that the trajectory will be adapted to different basket positions thanks to the visual feedback.

A third behavior design process $\mathcal{P}_{dev}$ may rely on a developmental approach, as in [Kraft et al., 2010], to learn how to grasp any kind of objects. $\mathcal{P}_{dev}$ relies on the exploration of grasping hypotheses with a dedicated algorithm and may thus find a lot of different grasping strategies for a basket [Kraft et al., 2010].

To be fair, K_{ref} should contain all the trajectories used to design both behavior design processes, but of course not the trajectories that are automatically generated by the exploration process of $\mathcal{P}_{dev}$. $\mathcal{P}_{ref}$ consists in following the most appropriate trajectory in K_{ref}, if any.

$\mathcal{P}_{vis}$ and $\mathcal{P}_{dev}$ should be able to grasp the basket at positions for which no trajectory exists in K_{ref}. The behaviors generated by both approaches in these cases, will be different from trajectories in K_{ref} and thus have a significant originality. The performance should also be above that of $\mathcal{P}_{ref}$, that will fail here. $\mathcal{P}_{vis}$ will always use the same grasping strategy. $\mathcal{P}_{dev}$ may find different grasping strategies. Some may have similar efficiency or even a better one, but others may be less efficient. This situation will be frequent, in particular if $\mathcal{P}_{dev}$ did not have enough time to explore the relevance of all possible grasping strategies. The ones that differ from $\mathcal{P}_{vis}$ grasping strategy should also be more different from K_{ref} and thus more original than any trajectory generated by $\mathcal{P}_{vis}$, as K_{ref} contains trajectories that are representative of $\mathcal{P}_{vis}$ grasping strategy. The originality of the behaviors generated by $\mathcal{P}_{dev}$ will then counterbalance their eventual and relative inefficiency. $\mathcal{P}_{dev}$ will then not be dominated by $\mathcal{P}_{vis}$. Both approaches would be considered as equally creative when compared with the Pareto dominance relation. This ability to innovate by finding different grasping strategies would have been probably completely dominated with a comparison

[2]It can actually compute the trajectory and rely on $\mathcal{P}_{ref}$ to make the robot follow it.

relying on performance only, as usually done in statistical learning when a generalization ability is tested on new experimental setups. The proposed creativity criterion allows to balance performance with the ability to go beyond the boundaries of the knowledge available when the behavior design process has been implemented. This is an expected feature of the creativity criterion, as this ability, even if it can result in less performing solutions in some cases, should help to face unforeseen situations.

The creativity of $\mathcal{P}_{dev}$ could also be measured relative to $\mathcal{P}_{vis}$ and K_{ref}. It should lead to the same conclusion.

4. Creativity to study learning and development in robotics: an example

Learning is a concept with many different facets [Wilson and Keil, 2001]. The focus here will be on learning in the presence of a performance measure or reward. Given (1) a task, (2) a training experience and (3) a performance measure, a computer program is said to learn if its performance at the task improves with experience [Mitchell, 1997]. Reinforcement learning in particular aims at learning what to do so as to maximize a numerical reward signal [Sutton and Barto, 1998]. A robot with such a feature should be able to deal with new and unforeseen situations by discovering the adapted behavior through a trial and error process. Robotic features are anyway still challenging for current learning methods. First, the environment is continuous, as are most perceptions and motor commands whereas reinforcement learning, in its most standard definition, is discrete [Sutton and Barto, 1998]. Discrete states and actions can be defined, but their numbers needs to be minimized because of the curse of dimensionality. Furthermore, this definition has a critical impact on learning efficiency [Kober et al., 2013] and requires an expertise on the task that may not be available and impedes robot adaptation abilities. Reinforcement learning has been adapted to continuous domains [Doya, 2000], but the exploration of possible actions remain an open issue. Modifying the behavior policy so as to follow the gradient of increasing performance is a simple strategy to improve robot efficiency. These strategies make anyway a strong assumption: the gradient of increasing reward leads to the searched policy. This assumption is reasonable if the starting policy is close enough to the searched one. Demonstrations made by human experts are expected to be close to efficient solutions. Using them to bootstrap the learning process has led to successful experiments in challenging scenarios like helicopter aerobatics [Ng et al., 2006] or robots playing table tennis [Mülling et al., 2013], but following a gradient does not always lead to optimal solutions [Stanley and Lehman, 2015]. The reward may be constant over large areas of the policy space: no gradient is available there to drive the learning process. Furthermore, the gradient of increasing performance can be misleading and drive to a dead-end [Lehman and Stanley, 2011]. The creativity criterion should help to balance the ability to find new behaviors through learning with their potential inefficiency with respect to what dedicated methods can do. It is then a quantitative criterion that will give a chance to adaptive methods with respect to non adaptive ones.

A performance criterion is already available with these learning algorithms, but what could the originality criterion look like for learning? It could focus on the limits of gradient following approaches and could thus compare the policy $\pi = \mathcal{P}(K_{ref}, c)$ to the policies obtained by performing local gradient ascent on the policies in K_{ref}.

Learning, as defined above, requires the definition of fixed representations for the reward, the policy and any other concept used by the learning algorithm. The performance of the learning process highly depends on these representations [Kober et al., 2013]. In an open-ended environment, this representation may need to be updated along time, for instance if new objects come into play, or if the environment changes.This ability seem to be one of the main feature of humans cognitive abilities [Karmiloff-Smith, 1995]. Developmental robotics proposes to go further than mere learning and to build a system not aimed at the resolution of one particular task, but built to develop along time and switch from one task to another [Weng, 2001], while drawing inspiration from human development [Guerin et al., 2013]. In a developing robot, the path towards getting an appropriate behavior may be long and require that it builds its own intrinsic motivations [Oudeyer and Kaplan, 2007, Baldassarre and Mirolli, 2013] as well as representations [Guerin et al., 2013, Doncieux, 2015].

What could the creativity criterion look like for developmental robotics and what could it be used for ? A performance criterion needs to be defined. It can either be provided by a human or by an external device and it should estimate the performance of the robot on the task chosen to test the creativity of the system. The knowledge base used for the originality criterion should include the behavior primitives available at the beginning of the developmental process as well as the preprogrammed motivations and representations. If an external caregiver helps the robot by showing it what to do, these demonstrations may also be taken into account. As for learning, the reference policy design process may be a gradient following approach. Here, it would try to improve the provided primitives and demonstrations on the basis of available motivations. The creativity criterion would highlight the possibility to go beyond this initial knowledge, without the need to necessarily perform better than other alternative approaches.

5. Discussion

The proposed creativity criterion is a tool aimed at encouraging progress on the road towards robots that are adaptive to open environments. It should enable Kuhn's normal science by allowing quantitative comparisons between approaches. The creativity of a robot, as defined here, is relative to a given knowledge base and a reference policy design process. It has two interesting consequences. First, it encourages researchers to highlight the knowledge they provide to their robot. Providing knowledge is not a problem *per se*, but it may significantly reduce robot adaptive abilities. A typical example consists in simplifying the robot vision process by defining objects as blobs of a given color that are on top of a flat surface. Such assumptions clearly limit the adaptive abilities of the robot. Knowledge like the effect of gravity on objects, for instance, will not limit robot adaptive abilities that much, except if the robot is sent to space. Secondly, it requires to define a reference policy design process and thus to make explicit the adaptive abilities that are searched for: the ones that go beyond the reference policy design process.

It is usual to test the generalization ability of algorithms relying statistical methods on a test set that hasn't been used during learning [Pinville et al., 2011]. The drawback of this method with respect to the creativity criterion proposed here is that is only takes into account the performance of the method. The creativity criterion also takes into account the ability to go beyond the boundaries of the knowledge that has been initially provided. A method that generates original behaviors that are performing a bit below alternative approaches would be eliminated with such a comparison criterion. But is the ability to perform better a good sign of adaptive abilities in new situations? The ability to generate original behaviors may be much more promising, this is the hypothesis that has motivated this work.

Does a robot need to be creative to face open environments? The creativity defined here is relative and aims at driving research work. Each time some knowledge or policy design process is identified to limit robot adaptive ability in a well identified context, a dedicated creativity criterion can be built to test the ability of proposed robots to go beyond these limits. Once the expected adaptivity has been reached, a new creativity criterion can be defined for which what has been found is not creative anymore. The creativity of a policy design process is then of interest only from a transient point of view to drive scientific research and has neither absolute nor intrinsic meaning.

How could a robot be creative? Learning in continuous domains can be formalized as an optimization problem [Stulp and Sigaud, 2013]. Local optimization methods alone converge to local optima and their performance depends on the initial policies. Global optimization methods are less dependent on this knowledge base. Evolutionary algorithms allow global optimization by relying on the principles of variation and selection [Eiben and Smith, 2008]. They can optimize parameters, but also allow to design structures [Floreano et al., 2008]. It has been proposed that human creativity may actually rely on similar mechanisms [Campbell, 1960, Simonton, 2010, Dietrich and Haider, 2015] that are also the basis of robot learning and adaptation algorithms [Doncieux et al., 2015]. An interesting feature with respect to creativity is their ability to define a divergent search process driven by the diversity of obtained behaviors [Mouret and Doncieux, 2012] or their novelty with respect to the behaviors observed so far [Lehman and Stanley, 2011] and also to combine it with convergent search objectives [Doncieux and Mouret, 2014]. Such a divergent search process has been used to discover many different ways of walking for an hexapod robot. It has been associated to a behavior selection process to allow the robot to deal with unforeseen damages, one of the situations a robot may face in an open environment [Cully et al., 2015]. As no initial walking behaviors were provided, the behaviors discovered by the whole process are then very creative with respect to the proposed measure.

6. Conclusion

Performance criteria do not evaluate the ability of a robot to deal with new situations. A robot may be programmed to efficiently fulfill a task in a well identified and closed context, but it does not mean that the same robot will still be efficient in environments that are even slightly different. This is not a major issue for some applications of robotics, notably in factories, where the environment can be controlled, but it critically limits their potential application in open contexts like our every day environment. The notion of creativity has been defined to design new performance criteria more adapted to this context. Creativity is a relative notion depending on a given knowledge base and on a reference policy design process. Its definition aims at providing quantitative measures of progress in the field of adaptive robotics.

Acknowledgments

This work has been supported by the FET project DREAM[3], that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 640891. The author would like to thank Olivier Sigaud, Benoît Girard and Mehdi Khamassi for their fruitful comments on the first draft of this article.

References

[Baldassarre and Mirolli, 2013] Baldassarre, G. and Mirolli, M. (2013). *Intrinsically motivated learning systems: an overview*. Springer.

[Boden, 1998] Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103(1):347–356.

[Bollini et al., 2013] Bollini, M., Tellex, S., Thompson, T., Roy, N., and Rus, D. (2013). Interpreting and executing recipes with a cooking robot. In *Experimental Robotics*, pages 481–495. Springer.

[Campbell, 1960] Campbell, D. T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review*, 67:380–400.

[Colton et al., 2012] Colton, S., Wiggins, G. A., et al. (2012). Computational creativity: The final frontier? In *ECAI*, volume 12, pages 21–26.

[Cully et al., 2015] Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, 521(7553):503–507.

[Deb, 2001] Deb, K. (2001). *Multi-objectives optimization using evolutionnary algorithms*. Wiley.

[Dietrich, 2004] Dietrich, A. (2004). The cognitive neuroscience of creativity. *Psychonomic bulletin & review*, 11(6):1011–1026.

[Dietrich and Haider, 2015] Dietrich, A. and Haider, H. (2015). Human creativity, evolutionary algorithms, and predictive representations: The mechanics of thought trials. *Psychonomic bulletin & review*, 22(4):897–915.

[Doncieux, 2015] Doncieux, S. (2015). Representational redescription: the next challenge? *IEEE-CIS CDS Newsletter*, 12(2).

[Doncieux et al., 2015] Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. G. (2015). Evolutionary robotics: What, why, and where to. *Frontiers in Robotics and AI*, 2(4).

[Doncieux and Mouret, 2014] Doncieux, S. and Mouret, J.-B. (2014). Beyond Black-Box Optimization: a Review of Selective Pressures for Evolutionary Robotics. *Evolutionary Intelligence*, 7(2):71–93.

[Doya, 2000] Doya, K. (2000). Reinforcement learning in continuous time and space. *Neural computation*, 12(1):219–245.

[Eiben and Smith, 2008] Eiben, A. E. and Smith, J. E. (2008). *Introduction to Evolutionary Computing (Natural Computing Series)*. Springer.

[Floreano et al., 2008] Floreano, D., Dürr, P., and Mattiussi, C. (2008). Neuroevolution: from architectures to learning. *Evolutionary Intelligence*, 1(1):47–62.

[3] http://www.robotsthatdream.eu/

[Guerin et al., 2013] Guerin, F., Kruger, N., and Kraft, D. (2013). A survey of the ontogeny of tool use: from sensorimotor experience to planning. *Autonomous Mental Development, IEEE Transactions on*, 5(1):18–45.

[Guilford, 1967] Guilford, J. P. (1967). Creativity: Yesterday, today and tomorrow. *The Journal of Creative Behavior*, 1(1):3–14.

[Hutchinson et al., 1996] Hutchinson, S., Hager, G. D., and Corke, P. I. (1996). A tutorial on visual servo control. *Robotics and Automation, IEEE Transactions on*, 12(5):651–670.

[Jones, 2006] Jones, J. (2006). Robots at the tipping point: the road to irobot roomba. *Robotics Automation Magazine, IEEE*, 13(1):76–78.

[Karmiloff-Smith, 1995] Karmiloff-Smith, A. (1995). *Beyond Modularity: A Developmental Perspective on Cognitive Science*.

[Kober et al., 2013] Kober, J., Bagnell, J. a., and Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274.

[Kowaliw et al., 2009] Kowaliw, T., Dorin, A., and McCormack, J. (2009). An empirical exploration of a definition of creative novelty for generative art. In *Artificial Life: Borrowing from Biology*, pages 1–10. Springer.

[Kraft et al., 2010] Kraft, D., Detry, R., and Pugeault, N. (2010). Development of object and grasping knowledge by robot exploration. *IEEE Transactions on Autonomous Mental Development*, 2(4):1–17.

[Kuhn, 1970] Kuhn, T. S. (1970). *The Structure of Scientific Revolutions, 2nd enl. ed.* University of Chicago Press.

[Lehman and Stanley, 2011] Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.

[Maitin-Shepard et al., 2010] Maitin-Shepard, J., Cusumano-Towner, M., Lei, J., and Abbeel, P. (2010). Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 2308–2315. IEEE.

[Mitchell, 1997] Mitchell, T. (1997). *Machine Learning*. McGrax Hill.

[Mouret and Doncieux, 2012] Mouret, J.-B. and Doncieux, S. (2012). Encouraging Behavioral Diversity in Evolutionary Robotics: An Empirical Study. *Evolutionary computation*, 20(1):91–133.

[Müller, 2002] Müller, M. (2002). Computer go. *Artificial Intelligence*, 134(1):145 – 179.

[Mülling et al., 2013] Mülling, K., Kober, J., Kroemer, O., and Peters, J. (2013). Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research*, 32(3):263–279.

[Ng et al., 2006] Ng, A. Y., Coates, A., Diel, M., Ganapathi, V., Schulte, J., Tse, B., Berger, E., and Liang, E. (2006). Autonomous inverted helicopter flight via reinforcement learning. In *Experimental Robotics IX*, pages 363–372. Springer.

[Oudeyer and Kaplan, 2007] Oudeyer, P.-Y. and Kaplan, F. (2007). What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6.

[Oudeyer et al., 2007] Oudeyer, P.-y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265–286.

[Oxman, 1990] Oxman, R. (1990). Prior knowledge in design: a dynamic knowledge-based model of design and creativity. *Design Studies*, 11(1):17–28.

[Pinville et al., 2011] Pinville, T., Koos, S., Mouret, J.-B., and Doncieux, S. (2011). How to Promote Generalisation in Evolutionary Robotics : the ProGAb Approach Formalising the Generalisation Ability. In *GECCO '11: Proceedings of the 13 th annual conference on Genetic and Evolutionary Computation*, pages 259–266.

[Runco and Jaeger, 2012] Runco, M. A. and Jaeger, G. J. (2012). The Standard Definition of Creativity. *Creativity Research Journal*, 24(1):92–96.

[Schmidhuber, 1991] Schmidhuber, J. (1991). A possibility for implementing curiosity and boredom in model-building neural controllers. In *From animals to animats: proceedings of the first international conference on simulation of adaptive behavior (SAB90)*.

[Schmidhuber, 2010] Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990-2010). *Autonomous Mental Development, IEEE Transaction on*, 2(3):230–247.

[Silver et al., 2016] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489.

[Simonton, 2010] Simonton, D. K. (2010). Creative thought as blind-variation and selective-retention: combinatorial models of exceptional creativity. *Physics of Life Reviews*, 7(2):156–179.

[Srinivasa et al., 2008] Srinivasa, S., Ferguson, D. I., Vande Weghe, M., Diankov, R., Berenson, D., Helfrich, C., and Strasdat, H. (2008). The robotic busboy: Steps towards developing a mobile robotic home assistant. In Burgard, W., Dillman, R., Plagemann, C., and Vahrenkamp, N., editors, *Intelligent Autonomous Systems 10*, pages 129–136. IOS Press.

[Stanley and Lehman, 2015] Stanley, K. O. and Lehman, J. (2015). *Why Greatness Cannot Be Planned: The Myth of the Objective*. Springer Publishing Company, Incorporated.

[Stein, 1953] Stein, M. I. (1953). Creativity and Culture. *The Journal of Psychology*, 36(2):311–322.

[Stulp and Sigaud, 2013] Stulp, F. and Sigaud, O. (2013). Robot skill learning: From reinforcement learning to evolution strategies. *Paladyn, Journal of Behavioral Robotics*, 4(1):49–61.

[Sutton and Barto, 1998] Sutton, R. S. and Barto, A. G. (1998). *Introduction to Reinforcement Learning*. MIT Press.

[Weng, 2001] Weng, J. (2001). ARTIFICIAL INTELLIGENCE: Autonomous Mental Development by Robots and Animals. *Science*, 291(5504):599–600.

[Wilson and Keil, 2001] Wilson, R. A. and Keil, F. C. (2001). *The MIT encyclopedia of the cognitive sciences*.