Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Towards Game-based Metrics for Computational Co-creativity

Rodrigo Canaan, Stefan Menzel, Julian Togelius, Andy Nealen

2018

Preprint:

This is an accepted article published in IEEE Conference on Computational Intelligence and Games (CIG). The final authenticated version is available online at: https://doi.org/10.1109/CIG.2018.8490429

Towards Game-based Metrics for Computational Co-Creativity

Rodrigo Canaan
Tandon School of Engineering
New York University
Brooklyn, NY, USA
rodrigo.canaan@nyu.edu

Stefan Menzel
Honda Research Institute Europe GmbH
Offenbach/Main, Germany
stefan.menzel@honda-ri.de

Julian Togelius
Tandon School of Engineering
New York University
Brooklyn, NY, USA
julian.togelius@nyu.edu

Andy Nealen
Tandon School of Engineering
New York University
Brooklyn, NY, USA
nealen@nyu.edu

The article has been accepted for publication in the IEEE Conference on Computational Intelligence and Games

Copyright notice:

©2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

DOI: 10.1109/CIG.2018.8490429

Towards Game-based Metrics for Computational Co-Creativity

Rodrigo Canaan
Tandon School of Engineering
New York University
Brooklyn, NY, USA
rodrigo.canaan@nyu.edu

Stefan Menzel
Honda Research Institute Europe GmbH
Offenbach/Main, Germany
stefan.menzel@honda-ri.de

Julian Togelius
Tandon School of Engineering
New York University
Brooklyn, NY, USA
julian.togelius@nyu.edu

Andy Nealen
Tandon School of Engineering
New York University
Brooklyn, NY, USA
nealen@nyu.edu

***Abstract*—We propose the following question: what game-like interactive system would provide a good environment for measuring the impact and success of a co-creative, cooperative agent? Creativity is often formulated in terms of novelty, value, surprise and interestingness. We review how these concepts are measured in current computational intelligence research and provide a mapping from modern electronic and tabletop games to open research problems in mixed-initiative systems and computational co-creativity. We propose application scenarios for future research, and a number of metrics under which the performance of cooperative agents in these environments will be evaluated.**

***Index Terms*—artificial intelligence, cooperative systems, games**

I. INTRODUCTION

Designing intelligent agents characterized by a co-creative, cooperative behavior would mark a major breakthrough in the age of industrial man-machine interaction. Exchanging relevant information with suitable time frequency and enriching the partner (human or machine) with novel perspectives and solution strategies on the problem are key factors for desirable results (considering the value of the output and the effort required). Cooperative games offer the valuable opportunity to realize an interactive environment for developing and evaluating computational methods used by these agents.

In this paper we review concepts and implementations of cooperative games in the light of their capability to impact development processes in (industrial) environments with co-evolution and co-creativity as important expressions for cooperation. Having a working definition of computational creativity, and how creative systems and their outputs are judged in terms of their value, novelty, interestingness, and surprise, will help us understand cooperatively creative agents and might help us build them as well. Computational creativity and AI-assisted design are important application areas for computational intelligence techniques such as neural networks, reinforcement learning and evolutionary computation; further, the conceptualization of creativity as search in a design space fits well with design applications of evolutionary computation.

Essentially, this paper tries to answer the following question: what game-like interactive system would provide a good environment for measuring the impact–and success– of a co-creative, fully cooperative agent?

We begin with a survey of the definition of computational creativity-related terms in the literature, how they relate to each other and how they apply to future work on our own co-creative agents in Section II. When considering cooperation between multiple actors (be they human or machine), in addition to the abilities and characteristics of each individual, the attributes of the relationships between individuals and the surrounding environment also impact the success of the endeavor. Section III explores some of these relational or environmental attributes of creative efforts, such as the exchange of information and the share of responsibility. In section IV we propose a set of metrics under which to evaluate cooperative agents in game-like environment, and section V gives our vision of how cooperative agents integrating all discussed techniques should operate in the long term.

II. COMPUTATIONAL CREATIVITY

Creativity is often understood as the production of novel and valuable concepts [1]. Computational creativity is a subfield of Artificial Intelligence (AI) that focuses on computational systems whose behavior can be deemed creative [2]. While much theoretical and practical work exists on systems that aim to be creative in their own right, with little or no human intervention [3]–[6], there are also many systems designed to cooperate with humans to achieve better results than either can presently do alone [7]–[9]. We focus on concepts of computational creativity and how they map to game-based tasks to further propose a number of concrete game-based metrics for co-creativity in a computational setting.

A. Novelty, Interestingness, Surprise

In his CSF framework [10], Wiggins says an artifact produced by a system is novel if there are no previously existing similar or identical artifacts in the context in which the artifact is produced. Ritchie [11] builds upon Wiggins' work and introduces the notion of the inspiring set as the "knowledge base of known examples which drives the computation within the program". Ritche calls an artifact generated by the program novel if it is not part of the inspiring set (or not too similar to its members). Both authors admit the possibilities of Novelty being either an absolute assessment (based on the existence of

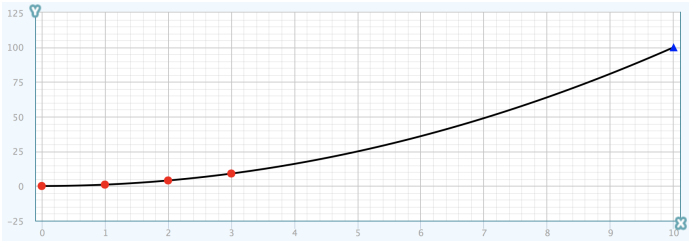


Fig. 1. A simple polynomial regression model trained on the dataset $I = [(0,0),(1,1),(2,4),(3,9)]$ (in red) would perfectly predict the point $x = (10,100)$ in blue, even though the Euclidean distance to the inspiring set is large

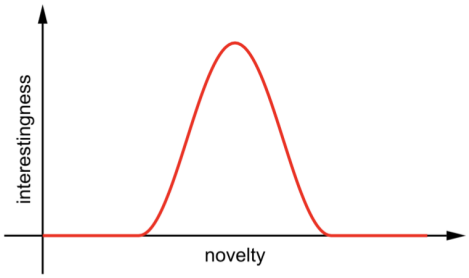


Fig. 2. A Wundt Curve, as shown in [13]

identical artifacts) or, more flexibly, to depend on some metric that establishes degrees of similarity between objects.

Reehuis *et al.* [12] provide an overview of Novelty metrics used by researchers, and propose dividing them between distance-based metrics and learning-based metrics. Distance-based metrics depend only on the distance, in a specified metric space, between a candidate solution and the archive of earlier solutions (what Ritchie would call the inspiring set). They define uniqueness as the minimum distance between a solution and a member of the archive, as used by [13] and [14]. Sparseness is defined as the average distance from a candidate solution to its k nearest neighbors in the archive, as used by Lehman and Stanley [15]. Reehuis *et al.* note that uniqueness is equivalent to sparseness with a value of k = 1.

Learning-based metrics take the agents expectations into account. Formally, let an agent (or an external viewer) be imbued with a model of the world, which ascribes probabilities to certain events. High novelty, or surprise occurs when the agent comes into contact with examples which contradict the model. Reehuis *et al.* provide the prediction error, dispersion in predictions and predictive variance of the model as examples of learning-based novelty. Itti and Baldi [16] provide a bayesian denition of surprise using the relative entropy, or Kullback-Leibler (KL) divergence [17]. Since the KL divergence depends on a prior probability distribution, we could also classify it as learning-based novelty.

B. Analysis of distance- and learning-based novelty metrics

We provide a simple example of the distinction between the two kinds of novelty in figure 1. The points in red are part of the inspiring set $I = [(0,0),(1,1),(2,4),(3,9)]$ and a candidate solution $x = (10,100)$ is shown in blue. A naive Euclidean distance-based metric would ascribe high novelty to x, while a simple learning model based on polynomial regression could might ascribe zero novelty to x, since it is a perfect fit to the parabola $y = x^2$. Thus, under learning-based novelty, what is novel to one observer might not be to another.

It is clear to us that the distinction between distance and learning-based novelty is didactic only. A high novelty value in a distance-based metric such as sparseness or uniqueness is equivalent to a low probability in a simple model that takes only the Euclidean distance from the points in the inspiring set into account (with more distant points being less probable). On the other hand, a more complex learning model can be abstracted as a distance metric in a sufficiently high dimension.

Thus, the choice of novelty metric to use depends on the problem. If one must describe a model being refined over time, or multiple agents with individual models making different predictions, a learning-based metric might be ideal. If there is no explicit model, or a single static model and a distance metric is readily available, it might be preferred. Richter [18] defines a "neighborhood structure" as an integral part of a fitness landscape, so we believe evolutionary computation is a good environment for distance-based metrics.

Whatever the kind of metric used, is important to note that a higher value of novelty is not necessarily desirable. As a simple example, consider a set of observations consisting entirely of random noise, such as a "poem-generator" that simply generates long strings of random characters. These would have high novelty (either in the distance or learning sense), but it could hardly be called a poem generator. It is clear that both low-novelty and extreme-novelty can be undesirable to a system, which is why some authors define the *interestingness* of an object as a function relating its novelty to some desirability metric. A Wundt curve [19] is a hedonistic function commonly used to express this relationship [13] [12] [20]. In this sense, interestingness might be characterized by just the right amount of novelty - not too much, not too little.

Learning-based interestingness is also defined in a way to avoid excessively high novelty (unpredictability). Schmidhuber, as part of his theory of artificial curiosity, provides a comprehensive framework [21] for characterizing the learning progress of an agent by noting the intimate relationship between prediction and compression. An observation is termed interesting if it enables the agent to learn some previously unknown irregularity, that is, further compress the available data. Rehuis *et al.* [12] discuss a number of different learning-based interestingness metrics, which attempt to maximize the learning progress induced by including a new observation in the model: Actual Learning Progress, Previous Learning Progress, Previous Competence Change and Reducible Error. These are all based on the difference between the prediction error in a region of the problem space at two points in time.

The use of these terms (novelty, interestingness, value) is not entirely consistent across all literature. For this reason, we find it convenient to settle on some definitions for our

purposes, which lean closer to the way the terms are used in [12]. These definitions are:

Novelty: any measure of dissimilarity between a sample concept and a collection of concepts (distance-based novelty) or an expression of the prediction error of a surrogate model (learning-based novelty).

Interestingness: a function of how desirable a solution is based on its novelty. This will typically assign a low score both to low-novelty and excessive-novelty solutions.

Surprise: synonymous with learning-based novelty, that is, a measure of how much a candidate solution deviates from a model's expectation.

C. Value

Wiggins defines Value as "The property of an artifact (abstract or concrete) output by a creative system which renders it desirable in the context in which it is produced". Given that we also defined interestingness with regards to desirability, a closer look at the relationship between interestingness and value is necessary.

We define value as any measure of desirability, possibly domain-specific, while interestingness will be used solely as a more domain-agnostic measure of desirability that depends only on the underlying novelty metric and possibly the agent's internal state, but not on any externally assigned goals. To make the distinction clear, we propose an example inspired the space probe described in [20]. Imagine a space probe designed for mining some kind of ore in a distant planet. It has a number of sensors to measure some features of the world and is able of movement in four different directions. Via reinforcement learning, it uses data from its sensors to build a model that predicts the concentration of ore in parts of the world.

Consider now two regions of the world $R1$ and $R2$. At some point in time the model predicts a high concentration of ore in $R1$ and low concentration in $R2$. After exploring both regions, $R1$ is found to have low concentration of ore, $R2$ is found to have a high concentration and the model is updated. From a pure learning perspective (that is, in terms of learning progress), both observations can be equally useful. From a value perspective, it is clear that $R2$ has more value. $R1$ is only useful to the extent that, by exploring similar regions, the probe might eventually learn a new pattern that enables it to avoid such low-value regions in the future.

As Graziano puts it, a reinforcement learning agent can be given an "internal or curiosity reward", which directs its learning, and an "external reward", defined to achieve some pre-defined goal. These must be balanced against each other, as, unless the agent is provided with an accurate model from the start, it first needs to learn where the high-value regions are by exploring unknown (possibly low-value) regions. This is known as the exploration and exploitation problem.

A more classic formulation of the exploration and exploitation problem is given by the Multi-Armed Bandit (MAB) problem, in which a gambler is faced with N slot machines (also known as "one-armed bandits") with unknown reward distributions and must decide which machine to play at each point in time. An in-depth study of the MAB problem is outside the scope of this article. For more information, see [22]. In a Reinforcement Learning context, we will take novelty or interestingness (depending on the formulation of the problem) to be related to an agent's internal reward, encouraging exploration, and value to be related to an agent's external objective, encouraging exploitation. For a pure learning agent, an external definition of value might not be necessary.

Another interesting application of the relationship between novelty and value is seen in Lehman and Stanley's novelty-based evolution [15]. They implement novelty search as an extension of the NEAT method [23], using sparseness as metric for novelty, where distance is a domain-dependent measure of behavioral difference. Sparseness is, in turn, measured against the current population plus an archived set of high-novelty solutions. The novelty of a solution is used as selection factor for the evolving population, and the external objective is only used as a stopping condition test. By not using a fitness function based on the external objective, they outperform traditional methods in some deceptive environments, that is, where the fitness function leads too often to local optima. This indicates that when a good heuristic for the desired objective is unavailable, search through novelty alone can still lead to good results. Another possibility is a combined approach, where both novelty and traditional fitness are rewarded concurrently in a multi-objective formulation of the problem [24].

III. Games as Mixed-Initiative Research Platforms

The recent years have seen advancements both in systems that facilitate human creation and systems able of autonomous creation. However, researchers have noted a gap in systems that can work in tandem with one or more human agents, and achieve similar levels of initiative and responsibility as would be expected from a human partner. These are known as mixed-initiative systems. Some authors also use the term human-computer co-creativity, or mixed-initiative co-creativity, when emphasizing the creative nature of the output of such systems.

Carbonel [25] defines mixed-initiative systems as "one in which both humans and machines can make contributions to a problem solution, often without being asked explicitly". This notion is developed by Burstein and Mcdermott [9], who investigate how humans and machines can "best share information about and control of plan development" in a mixed-initiative system so that each agent works in areas where they perform best, use the appropriate representation for the communication of plans and have the means of acquiring and transferring authority over tasks. They identify six areas of AI research that needed to be addressed to enable their proposed model of mixed-initiative planning systems: plan-space search control management, representations and sharing of plans, plan revision management, planning and reasoning under uncertainty, learning and inter-agent communications and coordination.

Yannakakis *et al.* [7] identify "mixed-initiative co-creation (MI-CC) as the task of creating artifacts via the interaction of a

human initiative and a computational initiative", emphasizing the proactivity of the contributors, and differentiating it from "non-proactive computer support tools (e.g. spell-checkers or image editors)". They also argue that, if such a system is able to foster human creativity, then it can be called mixed-initiative co-creativity.

Krüger *et al.* [8] classify interaction between human and machine system in three levels of cooperation complexity: tools, adaptive tools and cooperative assistants. With a "tool" the human user has complete responsibility for the success of the operation and adaptation to different tasks. An "adaptive tool" has a model the environment to adapt to different situations, but has no capability to resolve possible mismatch between its goals and the humans goals. "Cooperative assistants" have a model not only of the environment, but of the human user, and are equipped with a transparent interface enabling the negotiation of responsibilities and goals. Although they do not use the term mixed-initiative, it is our view that such a cooperative assistant would qualify as mixed-initiative.

A similar distinction is drawn by Davis *et al.* [26], between what they call Creativity Support Tools (which support a creative person), Computational Creativity systems (which autonomously create products) and Computer Colleagues, which are "Co-creative agents (that) collaborate with humans in real-time improvisation to enrich the creative process". Davis [27] previously defined human-computer co-creativity as "a situation in which the human and computer improvise in real time to generate a creative product", where "the contributions of human and computer are mutually influential" and that "introduces a computer into this collaborative environment as an equal in the creative process". (Though one can of course think of useful co-creative processes where the computer is not an equal.)

Games have been considered the "killer app" for computational creativity [28], due to being multifaceted , content intensive, benefiting greatly from procedual generation techniques and rich (highly interactive and engaging). Games have also traditionally been used as benchmarks for AI. Of particular interest are general game-playing algorithms, which can in principle be applied to any games and better generalize to other real-world problems. For example, the GVGAI competition offers a set of 2D arcade-like games [29]. The use of games as AI benchmarks has received recent media attention due to the success of DeepMind's success at the game of Go with AlphaGo [30], AlphaGo Zero [31] by combining reinforcement learning and Monte Carlo tree search. This paradigm has also yielded success in other games by Anthony *et al.* [32] and by DeepMind's AlphaZero [33]. Games are also fun. Perez *et al.* [29] suggest that this leads to higher interest in AI research by the general public, and a 2014 review of gamification studies by Hamari *et al.* [34] concludes that, although some methodological issues were found, most studies yielded positive effects of gamification. We would like to investigate whether the use of game-like techniques can lead to the design of better co-creativity tools for real world problems.

Finally, we have identified several modern games where we believe a good AI controller, especially one designed for co-operative play with humans, would benefit from addressing many specific issues listed by mixed-initiative and co-creativity researchers as research topics for the development of the field. Tables I and II illustrate a mapping between these research topics and games that would serve as interesting problems for those research topics. We further detail the correspondence between research topics and games below:

Agent modeling: A lot of research in mixed-initiative systems and co-creativity is concerned with building a good model of the other agent's behavior and goals. For Burstein and Mcdermott [9], intent recognition (e.g. filling in the gaps of a plan that is not specified to the degree of atomic actions) and learning user preferences are important tasks of mixed-initiative planing systems. The ability to build a model of the user is one of the factors that distinguish a cooperative assistant from an adaptive tool for Krüger *et al.* [8].

Hadfield-Menel *et al.* [35] introduce Cooperative Inverse Reinforcement Learning (CIRL), a framework of cooperation between a Human H and a robot R, where both players are rewarded by the same reward function, which is known only by H. R tries to infer the reward function from H's actions. They show that when H tries to greedily maximize its own rewards, R might learn a poor approximation to the real reward function and achieve suboptimal results, so optimal solutions may involve active instruction by the human. The use of Generative Adversarial Networks (GANs) [36] to generate novel artifacts based on the design objectives of a user [37] or emulating a specific art style [38] is also a recent and promising approach to this problem.

The amount of time or data available for learning can also impose constraints on the techniques used. If a behavior must be learned over the course of a single game session, for example (rather than over a large number of games), one approach used by Barret *et al.* [39] is to pre-compute a set of strategies S and assume the other player is using strategy S_i with probability p_i, using Bayesian reasoning to update the probability of each strategy whenever the other player makes an action. The value of a prospective action with each possible paired strategy is weighted by their probability to determine the best action. They show this can lead to better results than simply mirroring the other player, even when the actual strategy is not one of the strategies contained in S.

Another useful technique is empowerment maximization [40], [41]. Empowerment is an information-theoretic, intrinsic motivation metric that formalizes how much potential causal influence an agent has upon the world it can perceive. An artificial agent motivated to maximize its human partner's empowerment could sidestep the issue of creating a complex model of the other agent's intentions by simply acting to leave their partner's options open.

In games, the need to predict the other player's actions and objectives arises naturally in competitive environments, especially those involving simultaneous action selection (like *Race for the Galaxy* (Tomas Lehmann, 2007) and other forms

TABLE I
RESEARCH TOPICS AND GAMES

Agent Modelling	Changing Environment	Nontrivial Goals		
		Emerging goals	Hidden Goals	Dynamic Goals
Poker Race for the Galaxy *Hanabi* *Magic Maze*	*Pandemic* *Flashpoint* *Overcooked*	Minecraft Roleplaying Games	Mafia Werewolf The Resistance *Shadows over Camelot** *Dead of Winter**	Ticket to Ride Terra Mystica *Pandemic Legacy: Season 1*

Mapping of research topic to games. Games in italics are cooperative. Games with an asterisk are cooperative with an optional traitor mechanic. Underlined games are electronic games

TABLE II
RESEARCH TOPICS AND GAMES (CONT.)

Asymmetric responsibilities	Communication	
	Unconstrained	Constrained
Pandemic *Magic Maze* *Can't Drive This*	*Pandemic* *Flashpoint*	*Hanabi* *Magic Maze* Real-Time Games (in general) Competitive Games (in general)

of bluffing (like Poker). In cooperative games, the need for agent modeling is alleviated if players are allowed to freely coordinate their actions. However, some cooperative games like *Hanabi* (Antoine Bauza, 2010) and *Magic Maze* (Kasper Lapp, 2017) enforce communication restrictions, which makes agent agent modeling a key factor for success.

Changing environment: Referring to traditional AI planning systems, Burstein and Mcdermott [9] state "the worlds in which these planners worked tended not to change much, fight back at all", and regards plan revisions and reasoning under uncertainty as two major areas of necessary research. For Krüger *et al.* [8], the ability to"change one or more of its own parameters in response to environmental variations" separates regular tools from adaptive tools, and is one of the requirements for cooperative assistants.

Many modern tabletop games excel in thematically representing environment changes inspired in real world uncertainties. In *Pandemic* (Matt Leacock, 2008), in during the Infection phase, cards are drawn from and infection deck to randomly add disease cube to cities in the board. If not treated timely by the players, these might induce chain reactions and defeat the players. *Flashpoint: Fire Rescue* (Kevin Lanzing, 2011) has the Advance Fire phase, where smoke and fire can be added to the board, which can cause explosions, structural damage to a collapsing building and knock down player-controlled Firefighter units. These phases usually occur in between player action phases to randomly provide either resources or obstacles to the players, and we term them environment phases for generality. In *Overcooked*, an electronic cooperative cooking game, the ingredients each player has access to changes with shifts in the map layout.

In some games, the goal of the game itself (that is, the scoring function) may change unpredictably during the course of the game (for example, limited-time scoring opportunities). We investigate these and other goal-related features below.

Nontrivial goals: Real-life goals are often nontrivial. They might be unknown to some of the agents, as in [35]. The goal might change during the execution of a project or parts of it may be implicitly specified [9]. The goal might be complex and broken into subtasks, and the responsibility for each subtask must be properly assigned, which could involve negotiation [8], [9]. In short, Davis *et al.* [26] characterize goals as "socially negotiated, dynamic and emergent".

In some "games"[1], such as *Minecraft* (Mojang, 2008) and roleplaying games, there is no overall objective stated by the rules, although the players might still define objectives for themselves based on what is fun for them, negotiate it with other players and attempt to achieve them via cooperation or competition. We term these games with regards to their goal as Emerging, due to their emergent nature as a product of the interaction between players and the environment.

Modern tabletop also employ many variations of secret objectives. Although we could not find a fully cooperative game with hidden goals, social deduction games such as *Werewolf* (Davidoff, Plotkin, 1986) and *The Resistance* (Don Eskridge, 2009) feature competition between two or more factions (whose members cooperate among themselves), where each player typically only know the allegiance of a small fraction of the other players (and thus, their objectives). *Shadows Over Camelot* (Cathala, Laget, 2005) and *Dead of Winter: A Crossroads Game* (Gilmour, Vega, 2014) are semi-cooperative games with a random probability of one player being assigned a traitor role. The mere possibility of a traitor encourages players to second-guess other player's reasons. Dead of Winter features a fairly unique mechanic where, even if no traitor is present, each player's goal is composed of a public objective, shared by all non-traitor players, and a secret objective, where a player only wins if the group fulfills the public objective and they personally fulfill their secret objective (so that one or more players might still lose even

[1]At this point, we want to acknowledge the controversy in calling these activities games. In *Rules of Play* [42], Salen and Zimmerman's definition of game involves there being a quantifiable outcome. We sidestep this discussion and call them games for simplicity and consistency with common usage.

if the group achieves success). This adds another layer of complexity where seemingly strange behavior by a player can be justified either by their secret objective or by a traitor role, and a non-traitor player's need to fulfill their secret objective might lead to the failure of the entire group.

Dynamic goals (where the scoring function itself changes over the course of a game) are also common: in *Ticket to Ride* (Alan Moon, 2004), players have the option of drawing extra objective cards, achieving extra score if they manage to fulfill these new objectives, at the risk of score penalties if they fail. In *Terra Mystica* (Drögemüller, Ostertag, 2012), a unique scoring tile is randomly drawn for each turn, enabling limited-time scoring opportunities for all players. A cooperative example is *Pandemic Legacy: Season 1* (Daviau, Leacock, 2015), a variation of *Pandemic* where players play missions in a persistent and evolving world, and a mission's objective may be altered mid-course by specific storyline events.

Asymmetric responsibilities and areas of expertise: For Burstein and Mcdermott [9], two of the high-level goal are to enable proper communication between agents with different areas of expertise, and that each agent works in areas where they perform best. Krüger *et al.* [8] gives an example of a cleaning robot that is able to identify areas where it cannot access (e.g. due to being blocked by an object) and proactively request assistant from the human user (who has a different set of skills and is able to e.g. move the object away). Different responsibilities (such as teaching and learning) can also be a result of asymmetric information, such as in [35].

In *Pandemic* and many other games, each player controls a unique character with special abilities, such as performing one specific type of action more efficiently. Some games are more radical in the variability between player powers. In *Magic Maze*, players share control of a group of character pawns, but each player can only move a pawn in one specific cardinal direction. In *Can't Drive This* (Pixel Maniacs, 2016), one player takes the role of a driver while the other dynamically builds the road on which the first player must drive.

Communication: Researchers highlight the need for a shared representation [9] or interface [8] in which communication can happen. Burstein and Mcdermott [9] also implicitly acknowledge a cost to associated with communication when stating "it is almost necessarily the case that details will be left out, if the communication is to be succinct enough to make it worth defining the task for another to carry out". Lu *et al.* [43] use a cooperative co-evolutionary approach to demonstrate how the frequency at which communication occurs impacts cooperative performance under different communication costs. Finally, the problem definition itself might disallow certain forms of communication, or allow no communication at all, in which case agents still can gain information by reasoning about other agents' actions [35].

Games offer an avenue for exploring all of these problems. In games that allow unrestricted communication, such as *Pandemic* and *Flash Point*, complex communication involving conditional logic and algorithm building can emerge, as shown by Berland [44]. Designing a communication scheme with comparable expressive power for effective human-AI and AI-AI cooperation is an open problem. In the Tiny Coop environment [45] communication actions are available, allowing each agent to signal the direction it would like its partner to move in the immediate future. However, human communication often happens not at the level of individual actions, but in terms of higher level goals and their dependencies. The need for communication can also be triggered by specific events, such as the completion of a goal or a change in environment. A recent example of development in this direction are by Schrodt *et al.* [46], whose agent is able to establish cooperative goals in a variant of *Super Mario Bros.* while thinking out loud its current intentions and state.

In other games, the communication is restricted by the game rules. In *Hanabi*, players can only communicate by expending a limited number of hints, which can only state the color or value of cards in another player's hand. In *Magic Maze*, players can freely communicate, but only at specific points in time. As a real-time game, time spent elaborating the plan comes at the cost of time for execution of the plan.

In some competitive games, the rules allow full communication, including negotiation and partnerships, but it must occur in full view of other players. In this scenario, the cost of communication is the information that is leaked to antagonist players, and so communication is a strategic decision.

IV. Metrics for Co-creative Agents

We propose the following types of metrics for co-creative agents in game environments:

- **Value:** For any game with fixed objectives stated by the rules, a natural way to measure value is the game's scoring function. For games with emerging or hidden goals, explicit feedback from the user, if available, can also be used as a value metric. For procedurally generated content, value could be measured by a pre-determined fitness function of the generated artifact's features (as seen in [7]), by results of simulation [47] or by subjective evaluation of the human player, who selects their preferred generated artifact [7].
- **Learning-based metrics:** An agent might attempt to build a model of the other agents over the course of a game session or multiple sessions. A model of the user's behavior can be used to predict their action in a tree search algorithm. A model of the user's preferences can be used to predict the probability of acceptance of an artifact by the user. The accuracy of these predictions is a metric of learning-based novelty, and the higher the confidence of the model in a result, the higher the surprise if the prediction fails.

  For an agent attempting to gradually build a model of a player, care must be taken to isolate gains in performance (either in terms of value or in terms of accuracy of predictions) due to improvements on the part of the agent and improvements on the part of the human. After all, a human player could play with a simple, non-learning agent, and the agent could still report an increase in

performance due to the human player better learning to play the game or play matching the agent's expectations. A statistical analysis of player improvement over time is given in [47].

To avoid this confounding factor, it is important implement baseline agents, with statistically unchanging behavior, who would serve as proper control groups when paired with learning agents and humans, so that the impact of an agent's learning on performance cannot be overestimated.

- **Distance-based metrics:** In some scenarios, the product of each decision by the agent will not be a single, atomic action, but a number of options for the other agents to choose from, such as a number of action plans or a number of in-game artifacts for use of the human player. In these cases, distance-based metrics of novelty and interestingness can be used to make sure the suggestions offer a varied sample of the decision space, rather than small variations of a single idea. That way, the user is most likely to find a suggestion they identify it, and tweak some finer details to their own preference.
- **Empowerment metrics:** Empowerment "grows when different actions lead to different perceivable outcomes", and is a form of intrinsic motivation [41]. As such, it can be used in the absence of explicitly stated goals. We believe empowerment can also be used to maximize chance of acceptance of a suggestion, similar to distance-based metrics, by providing many relevant choices to the user.
- **Communication metrics:** The most direct way to measure the effectiveness of a communication scheme is simply to measure the difference in value (or in accuracy of predictions) achieved by cooperating under different schemes (or with no communication), as is done in [45]. The frequency of communication [43] can also be an important metric in scenarios where there is a communication cost or where player experience could be negatively impacted by a high-frequency stream of low-level communication actions.

The proposed metrics pose an initial approach to quantify the success of co-creative agents in cooperative games and similar environments.

V. The way forward

In tables I and II, we listed some characteristics of games that provide interesting research topics for human-computer cooperation. In our view, the most promising application scenarios are those focused on **agent modeling** and **communications** and are the biggest gap in current cooperative systems. They are at the core of co-creative cooperative activities, while the remaining entries of the table serve as challenges to be addressed by better cooperative systems (which include agent modeling and communication as core components): how will the other player react by a change in the environment? How to infer an unknown goal from a player's actions? How to communicate a change of plans due to a change in the environments? How to communicate (or predict) which activities are to be performed by each agent, especially under time constraints?

Going forward, we believe communication and agent modeling can also feed off each other. On one hand, an agent can use its ability to communicate to build more accurate models, either by directly asking for missing information or by picking up on cues from information provided by its peers. On the other hand, having an accurate model can help determine what information to share or ask for. A very clear example of this dynamic is in the game *Hanabi*, where different players are comfortable operating under different levels of implicit information (e.g. how willing are they to risk playing a card with incomplete information?). Observing the hints given by a player can help us infer how much information they need for their own actions, while knowing how they act under uncertainty helps us determine what hints to give.

While section III provides many examples of application scenarios to achieve progress in human-machine cooperation in the short term, our long term view is that this research can lead to applications where high-level goals and plans can be negotiated between human and artificial agents, taking into account their specific abilities and knowledge. The artificial agents will then be able to fill in small gaps in the plan by reasoning about a model of the world and of the other agents. Alternatively, the agent can proactively request any information it is missing if the gaps are too large to be filled.

It will be able to detect events that require a change of plans (such as a change in environment, available resources or goals of the group) and once again communicate and negotiate the new plan. All along the process, novel and valuable artifacts will be produced through computational creativity techniques, where novelty and value are judged in regards to a model of the knowledge and preferences of the target audience.

CONCLUSION

We started this paper by asking what game-like environments would be ideal for measuring the impact and success of co-creative cooperative agents. We answer that question by proposing several types of metrics, based on a thorough research on computational creativity and metrics used in the computational intelligence community for the related concepts of novelty, value, interestingness and surprise. We have shown how research in these scenarios, and similar games, can help shed light on open questions of the field and provided a vision of how these systems could operate in the long term.

We hope that this can lead to the development of better mixed-initiative, co-creative systems for a variety of domains, including industrial applications, where human and machine can cooperate working in areas where they perform best, communicating efficiently to achieve nontrivial goals under a changing, uncertain environment.

ACKNOWLEDGMENT

Rodrigo Canaan gratefully acknowledges the financial support from Honda Research Institute Europe (HRI-EU). We

would also thank Nikolas Dahn (TU Ilmenau) and Dr. Thomas Weisswange (HRI-EU) for their valuable input to this research.

REFERENCES

[1] M. A. Boden, "Creativity and artificial intelligence," *Artificial Intelligence*, vol. 103, no. 1-2, pp. 347–356, 1998.

[2] S. Colton, G. A. Wiggins *et al.*, "Computational creativity: The final frontier?" in *ECAI*, vol. 12, 2012, pp. 21–26.

[3] P. McCorduck, *Aaron's code: meta-art, artificial intelligence, and the work of Harold Cohen*. Macmillan, 1991.

[4] S. Colton, "Creativity versus the perception of creativity in computational systems." in *AAAI spring symposium: creative intelligent systems*, vol. 8, 2008.

[5] ——, "The painting fool: Stories from building an automated painter," in *Computers and creativity*. Springer, 2012, pp. 3–38.

[6] C. Guckelsberger, C. Salge, and S. Colton, "Addressing the why? in computational creativity: A non-anthropocentric, minimal model of intentional creative agency," in *Proceedings of the Eight International Conference on Computational Creativity*, 2017.

[7] G. N. Yannakakis, A. Liapis, and C. Alexopoulos, "Mixed-initiative co-creativity." in *FDG*, 2014.

[8] M. Krüger, C. B. Wiebel, and H. Wersing, "From tools towards cooperative assistants," in *Proceedings of the 5th International Conference on Human Agent Interaction*. ACM, 2017, pp. 287–294.

[9] M. H. Burstein and D. V. McDermott, "Issues in the development of human-computer mixed-initiative planning," in *Advances in Psychology*. Elsevier, 1996, vol. 113, pp. 285–303.

[10] G. A. Wiggins, "A preliminary framework for description, analysis and comparison of creative systems," *Knowledge-Based Systems*, vol. 19, no. 7, pp. 449–458, 2006.

[11] G. Ritchie, "Some empirical criteria for attributing creativity to a computer program," *Minds and Machines*, vol. 17, no. 1, pp. 67–99, 2007.

[12] E. Reehuis, M. Olhofer, M. Emmerich, B. Sendhoff, and T. Bäck, "Novelty and interestingness measures for design-space exploration," in *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 2013, pp. 1541–1548.

[13] R. Saunders, P. Gemeinboeck, A. Lombard, D. Bourke, and A. B. Kocaballi, "Curious whispers: An embodied artificial creative system." in *ICCC*, 2010, pp. 100–109.

[14] T. Hester and P. Stone, "Intrinsically motivated model learning for a developing curious agent," in *Development and Learning and Epigenetic Robotics (ICDL), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1–6.

[15] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evolutionary computation*, vol. 19, no. 2, pp. 189–223, 2011.

[16] L. Itti and P. Baldi, "Bayesian surprise attracts human attention," *Vision research*, vol. 49, no. 10, pp. 1295–1306, 2009.

[17] S. Kullback, *Information theory and statistics*. Courier Corporation, 1997.

[18] H. Richter, "Analyzing coevolutionary games with dynamic fitness landscapes," in *Evolutionary Computation (CEC), 2016 IEEE Congress on*. IEEE, 2016, pp. 609–616.

[19] W. M. Wundt, *Grundzüge de physiologischen Psychologie*. W. Engelman, 1874, vol. 1.

[20] V. Graziano, T. Glasmachers, T. Schaul, L. Pape, G. Cuccu, J. Leitner, and J. Schmidhuber, "Artificial curiosity for autonomous space exploration," *Acta Futura*, vol. 4, pp. 41–51, 2011.

[21] J. Schmidhuber, "Formal theory of creativity, fun, and intrinsic motivation (1990–2010)," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 230–247, 2010.

[22] S. Agrawal and N. Goyal, "Analysis of thompson sampling for the multi-armed bandit problem," in *Conference on Learning Theory*, 2012, pp. 39–1.

[23] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.

[24] J.-B. Mouret, "Novelty-based multiobjectivization," in *New horizons in evolutionary robotics*. Springer, 2011, pp. 139–154.

[25] J. R. Carbonell, "Mixed-initiative man-computer instructional dialogues." BOLT BERANEK AND NEWMAN INC CAMBRIDGE MASS, Tech. Rep., 1970.

[26] N. M. Davis, Y. Popova, I. Sysoev, C.-P. Hsiao, D. Zhang, and B. Magerko, "Building artistic computer colleagues with an enactive model of creativity." in *ICCC*, 2014, pp. 38–45.

[27] N. Davis, "Human-computer co-creativity: Blending human and computational creativity," in *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2013.

[28] A. Liapis, G. N. Yannakakis, and J. Togelius, "Computational game creativity." in *ICCC*. Citeseer, 2014, pp. 46–53.

[29] D. Perez-Liebana, S. Samothrakis, J. Togelius, T. Schaul, S. M. Lucas, A. Couëtoux, J. Lee, C.-U. Lim, and T. Thompson, "The 2014 general video game playing competition," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 8, no. 3, pp. 229–243, 2016.

[30] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.

[31] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.

[32] T. Anthony, Z. Tian, and D. Barber, "Thinking fast and slow with deep learning and tree search," in *Advances in Neural Information Processing Systems*, 2017, pp. 5366–5376.

[33] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel *et al.*, "Mastering chess and shogi by self-play with a general reinforcement learning algorithm," *arXiv preprint arXiv:1712.01815*, 2017.

[34] J. Hamari, J. Koivisto, and H. Sarsa, "Does gamification work?–a literature review of empirical studies on gamification," in *System Sciences (HICSS), 2014 47th Hawaii International Conference on*. IEEE, 2014, pp. 3025–3034.

[35] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, "Cooperative inverse reinforcement learning," in *Advances in neural information processing systems*, 2016, pp. 3909–3917.

[36] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.

[37] P. Bontrager, W. Lin, J. Togelius, and S. Risi, "Deep interactive evolution," *arXiv preprint arXiv:1801.08230*, 2018.

[38] A. Elgammal, B. Liu, M. Elhoseiny, and M. Mazzone, "Can: Creative adversarial networks generating art by learning about styles and deviating from style norms, june 2017," *arXiv preprint arXiv:1706.07068*.

[39] S. Barrett, P. Stone, and S. Kraus, "Empirical evaluation of ad hoc teamwork in the pursuit domain," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 567–574.

[40] C. Salge, C. Glackin, and D. Polani, "Empowerment–an introduction," in *Guided Self-Organization: Inception*. Springer, 2014, pp. 67–114.

[41] C. Guckelsberger, C. Salge, R. Saunders, and S. Colton, "Supportive and antagonistic behaviour in distributed computational creativity via coupled empowerment maximisation," in *Proceedings of the Seventh International Conference on Computational Creativity*, 2016.

[42] K. Salen and E. Zimmerman, *Rules of play: Game design fundamentals*. MIT press, 2004.

[43] X. Lu, S. Menzel, K. Tang, and X. Yao, "Cooperative co-evolution based design optimisation: A concurrent engineering perspective," *IEEE Transactions on Evolutionary Computation*, 2017.

[44] M. Berland and V. R. Lee, "Collaborative strategic board games as a site for distributed computational thinking," *Developments in current game-based learning design and deployment*, vol. 285, 2012.

[45] J. Walton-Rivers, "Controlling co-incidental non-player characters," http://www.iggi.org.uk/assets/IGGI-2016-Piers.pdf access on 02/01/2018.

[46] F. Schrodt, Y. Röhm, and M. V. Butz, "An event-schematic, cooperative, cognitive architecture plays super mario," *Cognitive Robot Architectures*, vol. 10, pp. 10–15, 2017.

[47] A. Isaksen and A. Nealen, "A statistical analysis of player improvement and single-player high scores." in *DiGRA/FDG*, 2016.