The Thirty-Third International
FLAIRS Conference (FLAIRS-33)

Modeling Novel Solutions to Creative Problem Solving Tasks with Subjective Observers

Chris Miller, Arnav Jhala
Department of Computer Science
North Carolina State University

Abstract

We propose a categorization of solution-centric evaluation metrics for a class of domain-independent AI challenge tasks known as MacGyver problems. Our definitions formally describe different classes of novel solutions for general creative problem solving tasks described in the MacGyver framework. Furthermore, inspired by existing theories of creativity, we extend the MacGyver problem formalism to incorporate subjective observers of problem solving tasks. By doing this, we explicitly model solutions to creative problem solving tasks as subjective evaluations based on the varying domain knowledge of observing agents. As an application of our extended formalism, we then illustrate how previous work on goal-driven conceptual blending represents a powerful form of human creativity whose creative solutions can be more formally described through our classes of novel solutions. Additionally, we conclude by highlighting strong connections between observer-oriented creative problem solving as described here and personalized procedural content generation in games.

Introduction

In this paper we explicitly model and classify novel solutions to creative problem solving tasks defined as classical planning problems. Additionally, we model these novel solutions with respect to observing agents which function as critics of proposed problem solutions. We propose these classifications in order to provide a formal vocabulary for describing solution-centric measures for creative problem solving tasks in terms of their novelty in a way that also allows for these solutions to be considered and evaluated outside of the sole context of the solver. We present these models through a class of domain-independent AI challenge tasks framed as classical planning problems referred to as MacGyver problems (Sarathy and Scheutz 2018). The introduction of MacGyver problems to the cognitive systems research community is designed to elucidate how humans exhibit significant degrees of creativity and flexibility when improvising solutions to seemingly unsolvable problems with limited resources. These problems are introduced as the first step towards formalizing creative improvisation and problem solving for the design of creative cognitive systems. However, we posit that concepts such as *novelty* and *creativity* are subjective evaluation criteria. Thus, we maintain that perceptions of novel solutions are particular to an observing individual's experience, knowledge, and familiarity with a given application domain.

By further building upon the formalization of MacGyver problems, we show how one can more precisely model what it means for a solution to be novel in the context of a set of observers. We believe that this explicit model of novel solutions can inform the design and evaluation of systems that aim to produce new or creative content. As a more concrete example of how our definitions of novel solutions can be seen in existing computational systems, we also reinterpret previous work on goal-driven conceptual concept blending tasks as creative problem solving tasks for finding novel solutions (Li et al. 2012). Overall, in this paper, we reintroduce MacGyver problems, establish a vocabulary for describing novel solutions, extend the formalism to incorporate observers of problem solving processes and solutions, and apply our formulation to the task of creative conceptual blending. We then conclude by highlighting connections between the formalizations presented here and personalized procedural content generation for game and level design.

MacGyver Problems

The motivation for introducing MacGyver problems begins with one of the most persistent questions in AI: how should machine intelligence be evaluated (Sarathy and Scheutz 2018)? Sarathy and Scheutz address this question by first exploring features of general intelligence. They posit that resourcefulness, improvisation, and creative problem solving can encapsulate the idea of general intelligence and use these notions to explore how machines can learn to improvise when they encounter problems that seem unsolvable. It is here that MacGyver problems are introduced using the language of classical planning as a means of formalizing intuitions of creative problem solving. MacGyver problems are defined as a class of classical planning problems that are initially unsolvable for an agent given their current knowledge. By definition, these problems require an agent to sufficiently expand its domain representation or understanding of its world within the context of a larger universe to solve the problem.

Formal Definition of MacGyver Problems

We now summarize the formal definition of MacGyver problems given that having an understanding of this formalization is imperative to our proposed extensions.

Definition 1 (Universe). A universe $\mathbb{U} = (S, A, \gamma)$ represents the entirety of the perceivable and actionable world that agents exist in–regardless of whether or not the entire universe is accessible to agents. In terms of classical planning, the universe is the classical planning domain encompassing all allowable states S, actions A, and transitions γ.

Definition 2 (World). A world $\mathbb{W}^t = (S^t, A^t, \gamma^t)$ represents a subspace of the Universe $\mathbb{U}$ that is perceivable and actionable by a particular species t of agent.

Definition 3 (Agent Subdomain). An agent subdomain $\sum_i^t = (S_i^t, A_i^t, \gamma_i^t)$ of type t represents a planning subdomain for an agent's perceptions and actions within its particular world $\mathbb{W}^t$. A key aspect of this definition is that it encapsulates the agent not being aware of all of its entire perceivable world at all times with $\sum_i^t$ representing the subspace of the world it perceives and can act upon at time i.

Definition 4 (MacGyver Problem). A MacGyver problem for an agent t is a planning problem in the agent's world $\mathbb{W}^t$ with the goal state g that is unreachable by the agent from its current initial state s_0. Hence, a MacGyver problem is defined as $\mathcal{P}_M = (\mathbb{W}^t, s_0, g)$ where $s_0 \in S_i^t$ and g is a set of ground atoms.

Problem Solving

Following the formal definition of MacGyver problems, it is now more clear how a MacGyver problem $\mathcal{P}_M$ is framed as a classical planning problem that is initially unsolvable from an agent's perspective. It also follows that solving $\mathcal{P}_M$ necessitates a transformation of the agent's initial subdomain $\sum_i^t$ to reach the goal state g. The transformation(s) may include, for example, the addition of a state, the addition of an action, or even the addition of a transition function to the agent's domain. Determining the series of domain modifications that can transform the agent's subdomain to a domain encompassing a solution is the crux of MacGyver problems and is also what makes them an interesting class of AI challenges.

Definition 5 (Agent Domain Modification). $\sum_j^{t*}$ represents a domain modification of $\sum_i^t$ in which $\sum_i^t$ has either been extended or contracted. $\sum_j^{t+}$ defines an agent-subdomain extension at time j that exists in $\mathbb{W}^t$ but not in $\sum_i^t$ at the previous time i; namely $\sum_i^t \leq \sum_j^t$. In the context of domain extension, $\sum_j^t = \sum_i^t \cup \sum_j^{t+}$.

Additionally, Sarathy and Scheutz state that an agent can extend its subdomain by sensing and perceiving its environment or itself (e.g. making observations, introspection, or receiving advice). They also show that agents will necessarily require heuristics to explore search spaces to solve MacGyver problems given that MacGyver problem decidability is intractable.

Evaluating Agent Performance

Sarathy and Scheutz (2018) also pose the important question: How should we evaluate creative problem solving and novel solutions? They briefly address this question by proposing three subclasses of evaluation criteria for MacGyver problems which vary in their objectivity and subjectivity: problem-centric, agent-centric, and solution-centric measures. Problem-centric measures are evaluated based on the inherent difficulty of a MacGyver problem via qualitative criteria such as the reachability of goal states, domain size, etc. Agent-centric measures primarily focus on inherent limitations of the agent and its environment. Lastly, solution-centric measures evaluate an agent's proposed solution for solving a MacGyver problem. In order to start exploring criteria for solution-centric measures in more detail, we build upon this prior work by proposing a classification of novel solutions as a solution-centric measure for creative problem solving tasks.

Characterizing Novel Solutions

Here we describe our three classifications of novel solutions: generally defined novel solutions, novel alternatives, and novel discoveries. As we introduce these definitions, we make the simplifying assumption that the problem solver is only being evaluated within its own context (i.e. the problem solver is its own and only observer). This is the default perspective of evaluation that is used in MacGyver problems given that evaluation is initially proposed only considering the context of the problem solver and not other entities. We will revisit these definitions in the context of larger problems with multiple observers. *We retain t as a variable describing an agent, not as a variable describing time in accordance with the MacGyver framework.* The exact criteria for a goal depend on the application domain–a generative system for example might describe its goals as the desired set of attributes a generated artifact should have.

Definition (General Novel Solution). A novel solution $\mathcal{N}_s = (\mathcal{P}_M, t, \lambda_s, O)$ is a 4-tuple where,

- $\mathcal{P}_M$ is the given MacGyver problem with initial state s_0 and goal state g at time i
- t is the problem solving agent
- $\lambda_s = (a_1, ..., a_n)$ for $a \in A^t$ representing a solution action sequence for reaching g from s_0
- O is a set of observers (t can be its own observer)

such that $\{\forall o : O | (g \in W^t \cap W^o), (g \in \sum_i^t \cap \sum_i^o), \lambda_s$ is not a known solution at time i for neither t nor o, λ_s is the

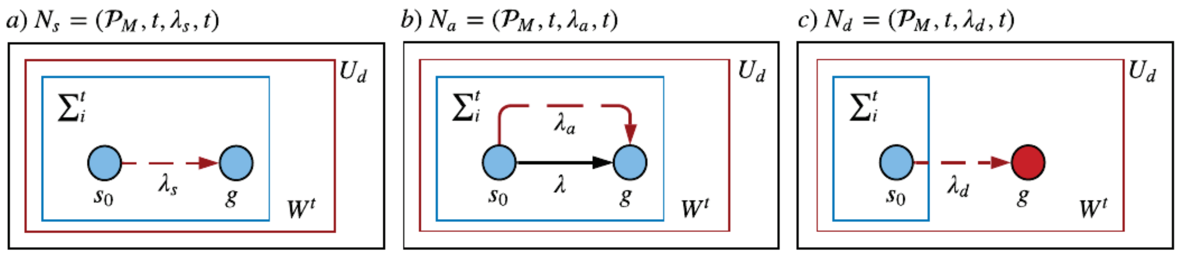

Figure 1: Three proposed classifications of novel solutions: a) general novel solution, b) novel alternative, c) novel discovery. Known states (i.e. in an agent's subdomain) are colored blue and known solutions are solid black lines. Unknown states (i.e. outside of an agent's subdomain) are colored red and unknown solutions are drawn with red dashed lines. In these simple examples, the problem solving agent is cast as its own observer.

first derived solution from s_0 to $g\}$.

As defined here, a general novel solution is essentially the prevailing notion of a "novel solution". Namely, this is a solution to a problem with an explicitly known, desired goal state that has not been solved yet. As shown in Figure 1a, λ_s defines the first solution from s_0 to g in the context of a problem solver and some set of observers (which is currently just t). Of course, often times, there exist several different ways of solving previously solved problems. Often times, these alternative solutions are more efficient, more insightful, or just more interesting ways of solving the same problem. To express these types of solutions we introduce another definition using the same notation from the definition of general novel solutions.

Definition (Novel Alternative). A novel alternative $\mathcal{N}_a = (\mathcal{P}_M, t, \lambda_a, O)$ is a 4-tuple such that $\{\forall o : O | (g \in W^t \cap W^o), (g \in \sum_i^t \cap \sum_i^o)$, λ_a is not a known solution at time i for neither t nor o, once λ_a is known it is *not the only known solution* from s_0 to $g\}$.

More informally, as shown in Figure 1b, a novel alternative $\mathcal{N}_a$ describes the creation of a new solution λ_a for reaching an *explicitly known* (see the definition of $\sum$) goal state to which a different, *known* solution λ already exists. We frame this as *novel alternative* due to the observation that λ_a represents a new method of reaching a known, desired state (i.e. a new means of reaching some state). Lastly, we formalize novel discovery.

Definition (Novel Discovery). A novel discovery $\mathcal{N}_d = (\mathcal{P}_M, t, \lambda_d, O)$ is a 4-tuple such that $\{\forall o : O | (g \in (W^t \cap W^o), (g \notin \sum_i^t \cup \sum_i^o)$, λ_d is not a known solution at time i for neither t nor $o\}$.

By its definition, novel discovery can be considered as a harder to achieve form of novelty. As shown in Figure 1c, unlike generally defined novel solutions and novel alternative, novel discovery *necessitates* a subdomain transformation that not only brings g into the known range of states in $\sum$ but also affords a solution λ_d for reaching that goal from an initially inadequate subdomain. By doing this, the agent is able to discover a novel solution for reaching a goal state that was not even under explicit consideration (at least in that agent's context).

Observers of Creative Problem Solving Tasks

Observers as Subjective Critics

Now we can postulate if an agent's solution is novel to any number of observers by modeling this evaluation as the subjective perception of an observer with varying amounts of domain knowledge. Inspired by Boden's (2004) seminal work on defining creativity, we maintain that creativity is a largely subjective phenomenon. In particular, we focus on Boden's definition of psychological creativity which is described as a creative process through which an idea is created that is new to the person who came up with the idea. Two core aspects of Boden's informal definitions of creativity are that creativity is inherently subjective and that novelty and value are essential criteria for creativity. Similarly and intuitively, we can model novelty as also being a context-specific judgement. We present this perspective not to dismiss claims of creativity in existing system but rather as a way to better understand and communicate what it means for a computational system to be creative or novel and to better understand how this is framed.

Considering creativity as a fundamentally subjective phenomenon, we postulate that perceptions of creativity are particular to an observing individual's experience, knowledge, and familiarity with a given application domain. Given that observers are able to represent artifacts in an application domain, observers can be cast as agents in $\mathbb{U}_d$ whose respective subdomains $\sum^o$ represent the extent of their explicitly held knowledge of that application domain. With this addition to the MacGyver framework, we can make interesting comparisons between an agent solving a MacGyver problem and observing agents within that application domain.

Observers as Subjective Critics of Problem Solvers

Figure 2 depicts a sample domain with a wide variety of MacGyver problems for a problem solving agent t with two observer agents o_1, o_2. The left subdomain belongs to the

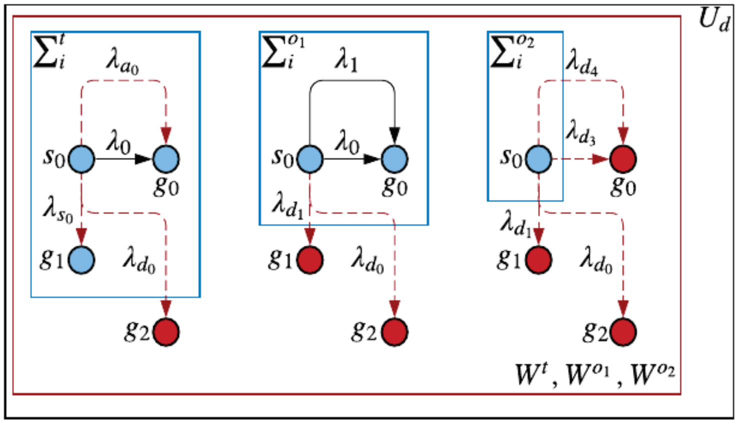

Figure 2: A sample domain with a problem solving agent t and two observers o_1, o_2. The left subdomain is the problem solving agent t and the middle and right subdomains are respectively those of o_1 and o_2. The diagram illustrates the different types of potentially novel solutions as perceived by each agent through a wide variety of MacGyver problems at time i. Known states are colored blue. Known solutions are solid black lines. Unknown states are colored red. Unknown solutions are drawn with red dashed lines.

problem solving agent t and the middle and right subdomains are respectively those of o_1 and o_2. For simplicity, we establish the problem solving agent and observer(s) worlds as all coinciding (i.e. these agents share the same physical and sensory capabilities). The distinction of these last two agents as observers is primarily to illustrate that they will evaluate the problem solving agent's solution(s) based on their own respective knowledge of potential solutions. As agents, observers are still capable of representing solutions in the application domain. In Figure 2, states and solutions that are not in each respective agent's initial domain are colored red or drawn with dashed lines representing currently unknown plans for reaching a particular state from another. Conversely, solid lines represent known plans for reaching a state from another state.

Let us focus on the problem solving agent t in Figure 2 as if it is proposing solutions to reaching goal states to its observers. Let us suppose that t derives the solution λ_{a_0} for reaching goal g_0. From the perspective of t, this represents a novel alternative given that it suggests a new solution for solving a previously solved problem (previously solved by λ_0). But what if t proposed this solution to o_1? Observer o_1 is already aware of t's proposed solution (formerly λ_{a_0}) and has already listed it as the known solution λ_1. Thus, from the perspective of o_1, t's proposed solution λ_{a_0} would not be considered a novel alternative, let alone a novel solution.

Observer o_2 however, has a significantly smaller domain representation and has yet to discover any solutions to any of its goal states. Let us now suppose that agent t proposes λ_{a_0} to o_2. Unlike agents t and o_1, g_0 lies outside of observer o_2's known subdomain. Accordingly, agent t's novel alternative through λ_{a_0}–which was considered not to be novel by observer o_1–would be considered a novel discovery to o_2.

Naturally, we also intend for these observers to be dynamic learners. Once an observer becomes aware of a solution, they then incorporate it into their knowledge and subsequently expand their domain knowledge. This captures the common expectation that once one is explicitly aware of a solution after having been shown it, they will no longer consider that solution as creative, let alone novel (at least without a retrospective analysis). Having dynamic learners also helps capture what we believe to be one of the key challenges in producing novel concepts–how to produce novel concepts for observers that constantly expand their threshold of novelty as they are exposed to novel concepts. In the context of generative systems, we consider this formalization to be particularly insightful. For a system that generates content, having a model of observers that evaluate generated artifacts based on their own understandings of a domain can provide an evaluation framework for the system to adapt its generation process to the needs of its observers (i.e. critics).

Through Figure 2, we illustrate how different types of novel solutions to classical planning problems can be categorized and how they may be perceived differently by different observers. It quickly becomes apparent that even slight differences in each agent's subdomain can affect their evaluations of the type of novel solutions that proposed solutions belong to. This allows for interesting, varied interpretations of the same solution based on the context of the evaluator. We have presented this classification of novel solutions as one of possibly many expansions of solution-centric measures for creative problem solving tasks as initially proposed by Sarathy and Scheutz (2018). Collectively, these proposed classifications provide a formal vocabulary for describing solutions to creative problem solving tasks in terms of their novelty in a way that also allows for these solutions to be considered and evaluated *outside of the sole context of the solver*.

On Conceptual Blending

Conceptual blending is a great example of a general cognitive operation with clear connections to creative problem solving processes and producing novel solutions. Fauconnier and Turner's widely disseminated work on conceptual integration networks describes the process of conceptual blending as the merging of two or more input mental spaces to produce a blended mental space (Fauconnier and Turner 1998). Thus, the key result of conceptual blending is that one is able to produce a blended space with new relations that did not exist in the individual input spaces. In this way, conceptual blending represents one of the most powerful and recognizable forms of human creativity (Fauconnier and Turner 2008). The concept of a mythological chimera for example, results from the conceptual blending of a lion, a goat, and a snake. As a cognitive process, conceptual blending is a clear example of a creative generation strategy which can be suggested as a type of heuristic for the creative problem solving tasks that MacGyver problems pose.

An even more direct interpretation of conceptual blending as a goal-oriented creative problem solving task is that of goal-driven conceptual blending (Li et al. 2012). Goal-driven conceptual blending is suggested as a more algorithmic approach to creative conceptual blending following the belief that theories of creativity should be computable (Johnson-Laird 2002). An example conceptual blending system that Li et al describe is one that uses goal-driven conceptual blending to generate fictional gadgets in computer-generated stories. As Li et al point out, "the aim of the gadget generation algorithm is to break out of these static world configurations and create new types of objects previously unknown to the system" (Li and Riedl 2011a; 2011b). The gadget generation algorithm iteratively constructs gadget behaviors by working backwards from goals derived from a story. Goals are represented as first-order logic predicates which instruct the system on how to identify the input spaces it can blend to achieve these goals. In particular, the behaviors of iteratively generated gadgets are constructed by projecting actions–which consist of preconditions and effects following a partial-order planning representation–from input spaces. Projections continue until all of the design goals are satisfied; hence, this selective projection process is goal-driven.

We can also consider goal-driven conceptual blending as a direct analogy to MacGyver problems. It is apparent that goal-driven conceptual blending is a creative problem solving task designed around producing creative solutions to initially unsolvable problems. In the vein of MacGyver problems, solutions to these problems are initially unknown to the agent and require clever manipulations of domain spaces to be solved. In the case of Li et al's goal-driven conceptual blending, these clever solver heuristics can be regarded as the efficient backward-chaining process used for selective projection (Li et al. 2012). Additionally, with our expanded definition of novel solutions for MacGyver-style problems, we can describe the specific type of solutions that goal-driven conceptual blending produces as novel discoveries. Here, the desired goal state g of the conceptual blending problem is a state containing all of the desired properties of a gadget. By the definition of goal-driven conceptual blending problems, a solution for reaching g is not already known and more importantly, *although the characteristics of a desired solution are known, a specific instance of an acceptable gadget is not known to the blending agent*. This last detail allows us to conclude that g exists outside of the blending agent's known subdomain (i.e. $g \notin \sum_i^t$). In this way, we consider the specific type of novel solutions that goal-driven conceptual blending is geared towards is that of novel discovery.

On Personalized Procedural Content Generation

In the context of interactive digital media, procedural content generation (PCG) generally refers to computational systems that aid or automate the design and creation of content for games, interactive narratives, etc (Hendrikx et al. 2013). PCG approaches include but are by no means limited to evolutionary algorithms and machine learning (Liapis, Yannakakis, and Togelius 2013; Summerville et al. 2018).

Personalized PCG represents a branch of PCG research that focuses on designing PCG systems that generate content specifically tailored towards particular player's preferences as a means of enhancing the player's experience of fun with the game. One such example of personalized PCG includes the use of evolutionary algorithms for the online generation of weapons for a space shooter game based on inferred models of the player's weapon preferences (Hastings, Guha, and Stanley 2009). As Sorenson and Pasquier point out however, this notion of fun is incredibly broad and quantitatively ill-defined (this mirrors similar challenges with defining creativity) (Sorenson and Pasquier 2010; Jordanous 2012). Regardless, successful progress towards automated personalized PCG necessitates some form of player modeling. Shaker et al (2010) describes a method of player modeling for automated level generation in platform games and personalized PCG for level design has also been explored through the Mario AI Championship (Shaker et al. 2011). One of the major conclusions from the Mario AI Championship Level Generation Track is that the relationship between level characteristics and fun is highly subjective and likely not a simple linear function.

To us, personalized PCG represents yet another creative problem solving task. This time however, the 'goal state' is the more loosely defined concept of fun from the perspective of an specific observer. In the PCG systems and studies described above, the creative problem solving process is now being guided by a PCG system. However, our addition of observers to the creative problem solving task framework also allows us to more formally express how these problem solving tasks work in the additional context of subjective observers. We believe that personalized PCG in games is a rich area of further exploration for our extended theory of creative problem solving tasks in which progress on either front will greatly benefit the other. *In particular, we hypothesize that fostering specific types of novel solutions for creative problem solving as defined in this paper can greatly impact player experiences in games utilizing personalized PCG*. In the near future, we plan to operationalize the theories of creativity and novelty presented here in a creative board game generator inspired by Cameron Browne's work on board game generation (Browne 2011).

Conclusion

In this paper we set out to accomplish the following:

1. Reintroduce the MacGyver problem framework as an insightful interpretation of creative problem solving tasks
2. Propose a categorization solution-centric evaluation metrics through a classification of novel solutions for creative problem solving tasks
3. Extend the MacGyver problem formalism to incorporate subjective observers of problem solving tasks
4. Illustrate how goal-driven conceptual blending represents a creative problem solving methodology for deriving novel discoveries
5. Highlight connections between observer-oriented creative problem solving and personalized procedural content generation in games

After introducing MacGyver problems, we defined three types of novel solutions: general novel solutions, novel alternative, and novel discovery. Each novel solution type is formally described in terms of the relationships between a desired goal state described in classical planning and the known subdomain of a creative problem solving agent. We presented these definitions as a means of more rigorously defining one of potentially many forms of solution-centric measures for MacGyver problems (Sarathy and Scheutz 2018).

Inspired by Boden's formulation of creativity as an inherently subjective criterion, we also extended the MacGyver framework to include subjective observer agents (Boden 2004). These observers are represented as additional agents with their own respective domain knowledge and representations that do not explicitly strive to solve creative problem solving tasks but rather critique and analyze proposed solutions from the problem solving agents they observe.

We then concluded by exploring how our extended formalisms are related to different problems that can be regarded as creative problem solving tasks–namely goal-driven conceptual blending and personalized PCG in games. We showed that goal-driven conceptual blending is directly analogous to MacGyver problems and that it strives for a specific type of novel solutions: novel discoveries. Afterwards, we highlighted personalized PCG as being closely related to creative problem solving tasks with observer agents as defined in this paper. We end by hypothesizing that fostering specific types of novel solutions for creative problem tasks with respect to subjective observers can greatly impact player experiences in games utilizing personalized PCG. We suggest that continued progress on either front can greatly benefit the other and subsequently future work in computational creativity, cognitive systems, and artificial intelligence.

References

Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Routledge.

Browne, C. 2011. *Evolutionary game design*. Springer Science & Business Media.

Fauconnier, G., and Turner, M. 1998. Conceptual integration networks. *Cognitive science* 22(2):133–187.

Fauconnier, G., and Turner, M. 2008. *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.

Hastings, E. J.; Guha, R. K.; and Stanley, K. O. 2009. Automatic content generation in the galactic arms race video game. *IEEE Transactions on Computational Intelligence and AI in Games* 1(4):245–263.

Hendrikx, M.; Meijer, S.; Van Der Velden, J.; and Iosup, A. 2013. Procedural content generation for games: A survey. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9(1):1.

Johnson-Laird, P. N. 2002. How jazz musicians improvise. *Music Perception: An Interdisciplinary Journal* 19(3):415–442.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3):246–279.

Li, B., and Riedl, M. O. 2011a. Creative gadget design in fictions: generalized planning in analogical spaces. In *Proceedings of the 8th ACM Conference on Creativity and Cognition*, 41–50. ACM.

Li, B., and Riedl, M. O. 2011b. A phone that cures your flu: Generating imaginary gadgets in fictions with planning and analogies. In *Workshops at the Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

Li, B.; Zook, A.; Davis, N.; and Riedl, M. O. 2012. Goal-driven conceptual blending: A computational approach for creativity. In *Proceedings of the 2012 International Conference on Computational Creativity, Dublin, Ireland*, 3–16.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013. Sentient sketchbook: Computer-aided game level authoring. In *FDG*, 213–220.

Sarathy, V., and Scheutz, M. 2018. Macgyver problems: Ai challenges for testing resourcefulness and creativity. *Advances in Cognitive Systems* 6:31–44.

Shaker, N.; Togelius, J.; Yannakakis, G. N.; Weber, B.; Shimizu, T.; Hashiyama, T.; Sorenson, N.; Pasquier, P.; Mawhorter, P.; Takahashi, G.; et al. 2011. The 2010 mario ai championship: Level generation track. *IEEE Transactions on Computational Intelligence and AI in Games* 3(4):332–347.

Shaker, N.; Yannakakis, G.; and Togelius, J. 2010. Towards automatic personalized content generation for platform games. In *Sixth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Sorenson, N., and Pasquier, P. 2010. The evolution of fun: Automatic level design through challenge modeling. In *ICCC*, 258–267.

Summerville, A.; Snodgrass, S.; Guzdial, M.; Holmgård, C.; Hoover, A. K.; Isaksen, A.; Nealen, A.; and Togelius, J. 2018. Procedural content generation via machine learning (pcgml). *IEEE Transactions on Games* 10(3):257–270.