Limits Theorems for Creativity with Intentionality

Lav R. Varshney[†,‡]
[†]University of Illinois at Urbana-Champaign, Urbana, IL, USA
[‡]Salesforce Research, Palo Alto, CA, USA
varshney@illinois.edu

Abstract

Creativity is the generation of an artifact that is judged to be novel and high-quality. Several computational creativity systems with diverse algorithmic foundations are now meeting standards of novelty and quality, as judged by experts in creative fields. The existence of these myriad design approaches suggest a natural question analogous to the one addressed in information theory: Are there fundamental limits to creativity? Here, we first review a recent mathematical formalism that captures key aspects of combinatorial creativity and yields fundamental tradeoffs between novelty and quality. The fundamental limit resembles Shannon's capacity-cost function. Then we extend the theory to capture intentionality in creativity, treating it as a communication problem, where a creative artifact must not only be novel and high-quality, but also reliably convey a message of given information rate. The resulting fundamental limit resembles the information bottleneck optimization in machine learning.

Introduction

Computational creativity systems are now able to produce ideas and artifacts that are judged to meet standards of novelty and utility by experts in creative domains, see e.g. (Boden 2004; 2015; Colton and Wiggins 2012; Varshney et al. 2019; Keskar et al. 2019), according to a variety of assessment criteria (Jordanous 2012; Colton et al. 2014; Lamb, Brown, and Clarke 2018; Riedl 2015; Hashimoto, Zhang, and Liang 2019). Yet, it has been unclear whether there are upper bounds on how creative any system can be, whether human, machine, or hybrid. This suggests the need for a general theory of creativity that would yield fundamental limits, analogous to the Shannon limit for reliable communication in the presence of noise (Shannon 1948) or the Carnot limit for efficiency of engines (Carnot 1824). Note that such limit theorems are prevalent in mathematical systems theories (Auyang 2004) and determined within closed deductive systems that require abstraction to establish.

Fundamental limit theorems serve several different purposes. First, they establish which resources and performance criteria are fundamental and which are largely unimportant. Second, they demarcate what is possible from what is impossible, providing design insights into operating at the boundary, that is, principles for optimal designs. Third, they define fundamental benchmarks that allow an evaluation of new creativity algorithms on an absolute scale, rather than only compared to people or previous technologies. Finally, they state ideals for pushing people to build technologies that approach/achieve these absolute limits. Note that the kinds of informational fundamental limits we consider here do not take computational complexity into account.

A basis for such a general mathematical systems theory was given in prior work that formalized the structure of conceptual spaces and computational creativity (Wiggins 2006; Ritchie 2007; 2012; Hung and Choy 2013; Velardo and Vallati 2016). We recently extended this theory to include a statistical dimension, which further enabled a characterization of what is possible and what is impossible in designing computational creativity systems. The result detailed in (Varshney 2019a) provided a limit theorem on the tradeoff between novelty and quality for a given creative domain. The result focused on combinatorial creativity but also captured transformational creativity as a wholly different kind of creativity. This previous work, however, only considered the technical problem of creativity and excluded consideration of intentionality—a (human) intent, inspiration, or desire to express something (Collingwood 1938; Hertzmann 2018; 2020). As described there, "These aspects of intent in creativity are irrelevant to the engineering problem". Consideration of intentionality is also absent in previous theories of creativity (Wiggins 2006; Ritchie 2007; 2012; Hung and Choy 2013; Velardo and Vallati 2016).

Yet, it is said—especially in the Western tradition following Romanticism (but see criticisms, e.g. (Wimsatt and Beardsley 1946))—that communication of meaning in art is necessary for eliciting an aesthetic experience (Csikszentmihalyi and Robinson 1990; Cilliers 1998; Ritchie 2007). For example considering narration or poetry, (linguistic) *meaning* is the relation between a linguistic form and communicative intent, where communicative intents are about things that are outside of language. Communicative intent is distinct from standing meaning, which is constant across all of its possible contexts of use (Bender and Koller 2020).

Recent surveys further indicate that people want not just novelty/quality, but also intentionality and autonomy, to attribute creativity to an artificial system (Ventura 2019). These further layers of desiderata for creative systems are redolent of the three layers of communication put forth by

Proceedings of the 11th International
Conference on Computational Creativity (ICCC'20)
ISBN: 978-989-54160-2-8

Warren Weaver (Shannon and Weaver 1949):

- the technical problem (How accurately can the symbols of communication be transmitted?),
- the semantic problem (How precisely do transmitted symbols convey the desired meaning?), and
- the effectiveness problem (How effectively does received meaning affect conduct in the desired way?).

Beyond the technical problem of creativity from prior work (Varshney 2019a), here we are concerned with incorporating intentionality to consider the semantic problem of creativity.

Meaning and understanding have long been described as a key to intelligence (Bender and Koller 2020). Intention is realized when the system produces an artifact with the goal of communicating a particular message and the observer reliably understands that message from the artifact (Ventura 2019). As such, here, we aim to extend (Varshney 2019a) to include an intentionality layer, by requiring a creative artifact to not only be novel and high-quality, but also reliably convey a message.

The remainder of the paper is organized as follows. In the next section, we summarize our past theoretical framework. Next we extend that framework by considering the communicative intent of intentionality. Finally, we conclude.

Review of Mathematical Formalisms and Limits of Creativity

We first review the existing formalism and resultant limit theorems for the technical problem of combinatorial creativity. For brevity, here we restrict to finite artifacts, but see (Varshney 2019a) for extensions. To facilitate exposition of the mathematical formalism, we use culinary creativity as a running example.

The basic idea is to formalize the conceptual space for combinatorial creativity, to formalize the two dimensions of merit (novelty and quality), and to formalize a general description of a computational creativity algorithm that operates in the conceptual space to optimize the tradeoff between novelty and quality.

Definition 1. *A* component *is an atomic unit in the creative domain, drawn from the set* Ω*, from which artifacts are constructed.*

In culinary, this would be the list of possible ingredients that people eat.

Definition 2. *A* discrete artifact *is an unordered combinatorial object* α *selected from the power set* 2^Ω *of possible components,* Ω*, that define the creative domain.*

In culinary, this would be a specific recipe, expressed in terms of its ingredients (and not amounts or instructions).[1]

Definition 3. *The* known set *is a set of artifacts that are already known in the creative domain,* Θ*, also called the* inspiration set*. In the discrete case,* $\Theta \subseteq 2^\Omega \in 2^{2^\Omega}$.

[1]For brevity, we focus on ingredient lists, but extensions to structured objects like sequential recipes (which may be represented as directed acyclic graphs) or music compositions with non-commutative novelty and quality functions follow directly.

In culinary, this may be all recipes in published cookbooks, or recipes a given person has cooked/eaten.

Definition 4. Novelty *is determined using a non-negative function, in the discrete case,* $s : 2^\Omega \times 2^{2^\Omega} \mapsto \mathbb{R}_+$ *that measures the surprise of a given artifact* α_0 *in the presence of a known set* Θ.

A particular novelty measure that was considered is the empirical *Bayesian surprise* (Itti and Baldi 2006; 2009; Baldi and Itti 2010; Varshney 2019b). If we let the prior probability distribution of artifacts in the known set Θ be P_θ, the creation of a new artifact α will update it to a posterior distribution $P_{\theta|\alpha}$. Then the Bayesian surprise $s(\alpha, \Theta)$ for a given artifact with respect to the inspiration set Θ is

$$s(\alpha, \Theta) = \int_\Theta P_{\theta|\alpha} \log \frac{P_{\theta|alpha}}{P_\theta} d\theta.$$

Notice that such a Bayesian notion of surprise is an expectation-based novelty measure (Grace and Maher 2019) with respect to the prior P_θ. When there is lack of absolute continuity, we have an infinite value that indicates transformational creativity in a kind of hierarchy (Wiggins 2019).

Definition 5. Quality *is determined using a non-negative function* $q : 2^\Omega \mapsto \mathbb{R}_+$ *that measures the utility of a given discrete artifact* α_0.

In culinary, quality may be a measure of flavor derived from the hedonic psychophysics of olfactory perception.

Definition 6. *A* creativity algorithm $\mathcal{G}$ *is a probabilistic process* $P_A(\alpha)$ *that produces a set of* n *artifacts* $\{\alpha_i\}_{i=1}^n$.

This view of creativity as stochastic sampling in the conceptual space is very general. In degenerate forms, such a definition encompasses generative algorithms that enumerate the entire space or optimization algorithms that directly and deterministically generate just $n = 1$ possibility.

With this formalism in place, the fundamental tradeoff between the average quality and average surprise produced by sampling algorithm $P_A(\alpha)$ is cast as follows.

$$S(Q) = \max_{P_A(\alpha):\mathbb{E}[q(A)]\geq Q} \mathbb{E}[s(A, \Theta)].$$

This is not only a limit theorem, but also implies an optimal creativity algorithm, the extremal $P_A^*(\alpha)$,

$$P_A^*(\alpha) = \arg \max_{P_A(\alpha):\mathbb{E}[q(A)]\geq Q} \mathbb{E}[s(A, \Theta)].$$

Quite unexpectedly, when taking $s(\cdot)$ as Bayesian surprise, using techniques from information geometry, the result is a flipped version of Shannon's capacity-cost function (Shannon 1948; 1959; Varshney 2008).

Theorem 1 (Varshney (2019))**.** *The fundamental tradeoff between novelty and quality in combinatorial creativity is given by the following expression:*

$$S(Q) = \max_{P_A(\alpha):\mathbb{E}[q(A)]\geq Q} I(A, \Theta). \quad (1)$$

Intriguingly, creativity has an equivalence to information transmission, with quality playing the role of energy and novelty playing the role of information rate. All of this, however, without communicative intent from the creator.

Introducing Intentionality

Now we extend the formalism to bring intentionality into the picture. In particular, we formalize intentionality as the need to reliably communicate a message m from the creator to an audience member using creative artifacts α, where perception of the message-bearing part of the creative artifact (the signal) is modeled as a noisy channel with transition probability assignment $p_{\hat{A}|A}$. Here $\hat{\alpha}$ are perceived signals, decoded as messages $\hat{m}$. Now the goal is to communicate so error probability, $\Pr[m \neq \hat{m}]$, via creative artifacts is arbitrarily small.

For reliable communication alone, the limiting information rate, R, is the channel capacity C. If there are constraints on the signaling strategy, this may be reduced. Due to the noisy channel coding theorem, this fundamental limit of channel capacity is given as follows.

Theorem 2 (Shannon (1948))**.** *The fundamental limit of reliable communication in the presence of noise under the input constraint requiring the input distribution to be in the family $\mathcal{P}$ is the channel capacity*

$$C(\mathcal{P}) = \max_{P_A(\alpha)\in\mathcal{P}} I(A;\hat{A}). \tag{2}$$

Now we essentially combine Theorem 1 (information geometry argument) with Theorem 2 (random coding argument) to get a limit theorem for creativity with intentionality. The detailed proof is omitted for brevity; a sketch is given.

Theorem 3. *For a given perception channel $p_{\hat{A}|A}$ and known set Θ, we require a minimal amount of average quality Q and minimal novelty S, then the maximum information rate of communicative intent R that can be reliably transmitted is:*

$$C(Q,S) = \max_{P_A(\alpha):\mathbb{E}[q(A)]\geq Q, I(A,\Theta)\geq S} I(\hat{A};A).$$

Proof. Notice that Eq. (2) in Theorem 2 holds for any constrained family of input distributions $\mathcal{P}$. Here we choose $\mathcal{P} = \{P_A(\alpha) : \mathbb{E}[q(A)] \geq Q, I(A,\Theta) \geq S\}$ to satisfy the novelty and quality constraints and the result follows. Detailed arguments are needed to ensure that asymptotic arguments have appropriate interaction. □

Again details omitted for brevity, but the information-theoretic optimization problem in Theorem 3 can be reformulated in a dual formulation as an optimization of novelty under quality and communicative intent constraints as:

$$S(Q,R) = \max_{P_A(\alpha):\mathbb{E}[q(A)]\geq Q, I(\hat{A};A)\geq R} I(A,\Theta).$$

to become more along the lines of the form of Eq. (1) in Theorem 1, which is a constrained optimization for novelty. As we can see in this form, the requirement for positive communicative intent rate may decrease the novelty that can be achieved, when the mutual information constraint is active.

Note that there is a natural Markov relationship as $\hat{A} \leftrightarrow A \leftrightarrow \Theta$, based on the order of how variates are chosen. Hence, the constraint will indeed often be active.

In a third alternative form as a constrained optimization for quality, one can likewise observe that the communicative intent requirement may reduce the achievable quality.

As far as we know, this is the first characterization of how introducing intentionality into combinatorial creativity may reduce achievable novelty and quality.

With this limit expression in hand, we can also notice a formal connection to the information bottleneck functional (Tishby, Pereira, and Bialek 1999; Gilad-Bachrach, Navot, and Tishby 2003) which has become prevalent in the machine learning community and can be thought of as expressing the idea of minimal sufficient statistics with respect to a relevance variable. In particular, both settings are mutual information optimization under mutual information constraints but with the inequality going the other way. For creativity, one term is concerned with communication and the other with novelty, whereas for sufficient statistics, one term is still communication but the other is relevance. In this sense, it is interesting to think about novelty as a kind of irrelevance to the inspiration set (irrelevance rather than relevance due to the reversed inequality).

Conclusion

In this paper, we have summarized our recent work in engineering systems theory that uses information geometry to establish the fundamental limits of creativity and that may therefore be of interest to researchers in computational creativity. Importantly, (Varshney 2019a) had dismissed the role of intentionality in creativity when establishing limit theorems for the technical problem of creativity.

Here we have therefore brought intentionality back into the picture and established a further limit theorem for intentional creativity, connecting our previous limits of creativity with Shannon's previous limits of reliable communication. This investigation of semantic creativity shows that requiring communicative intent may reduce the quality and/or novelty of creative artifacts that are generated. Moreover, connections to the information bottleneck in machine learning essentially show that novelty in creativity can be thought of as a kind of irrelevance to the inspiration set.

Going forward, it will be of interest to compute fundamental limits with and without intentionality in given creative domains, and to compare them with the performance of existing computational creativity algorithms.

Acknowledgment

Discussions with D. Oppenheim are appreciated.

References

Auyang, S. Y. 2004. *Engineering—An Endless Frontier*. Cambridge, MA: Harvard University Press.

Baldi, P., and Itti, L. 2010. Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Netw.* 23(5):649–666.

Bender, E. M., and Koller, A. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *2020 Annual Conference of the Association for Computational Linguistics (ACL)*.

Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. London: Routledge, 2nd edition.

Boden, M. A. 2015. Foreword. In Besold, T. R.; Schorlemmer, M.; and Smaill, A., eds., *Computational Creativity Research: Towards Creative Machines*. Springer. v–xiii.

Carnot, S. 1824. *Réflexions sur la Puissance Motrice du Feu et Sur Les Machines Propres a Développer Cette Puissance*. Paris: Chez Bachelier. Libraire.

Cilliers, P. 1998. *Complexity and Postmodernism: Understanding Complex Systems*. Routledge.

Collingwood, R. G. 1938. *The Principles of Art*. Clarendon Press.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In De Raedt, L.; Bessiere, C.; and Dubois, D., eds., *ECAI 2012: 20th European Conference on Artificial Intelligence*. Amsterdam: IOS Press BV. 21–26.

Colton, S.; Pease, A.; Corneli, J.; Cook, M.; and Llano, T. 2014. Assessing progress in building autonomously creative systems. In *Proc. Int. Conf. Comput. Creativity (ICCC 2014)*, 137–145.

Csikszentmihalyi, M., and Robinson, R. E. 1990. *The Art of Seeing: An Interpretation of the Aesthetic Encounter*. Los Angeles, CA, USA: Getty Publications.

Gilad-Bachrach, R.; Navot, A.; and Tishby, N. 2003. An information theoretic tradeoff between complexity and accuracy. In Schölkopf, B., and Warmuth, M. K., eds., *Learning Theory and Kernel Machines*, volume 2777 of *Lecture Notes in Computer Science*. Berlin: Springer. 595–609.

Grace, K., and Maher, M. L. 2019. Expectation-based models of novelty for evaluating computational creativity. In Veale, T., and Cardoso, F. A., eds., *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*. Springer. 195–209.

Hashimoto, T.; Zhang, H.; and Liang, P. 2019. Unifying human and statistical evaluation for natural language generation. In *Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Hertzmann, A. 2018. Can computers create art? *Arts* 7(2):18.

Hertzmann, A. 2020. Computers do not make art, people do. *Communications of the ACM* 63(5):45–48.

Hung, E. C. K., and Choy, C. S. T. 2013. Conceptual recombination: A method for producing exploratory and transformational creativity in creative works. *Knowledge-Based Systems* 53:1–12.

Itti, L., and Baldi, P. 2006. Bayesian surprise attracts human attention. In Weiss, Y.; Schölkopf, B.; and Platt, J., eds., *Advances in Neural Information Processing Systems 18*. Cambridge, MA: MIT Press. 547–554.

Itti, L., and Baldi, P. 2009. Bayesian surprise attracts human attention. *Vis. Res.* 49(10):1295–1306.

Jordanous, A. K. 2012. *Evaluating Computational Creativity: A Standardised Procedure for Evaluating Creative Systems and its Application*. Ph.D. Dissertation, University of Sussex.

Keskar, N. S.; McCann, B.; Varshney, L. R.; Xiong, C.; and Socher, R. 2019. CTRL: A conditional transformer language model for controllable generation. arXiv:1909.05858 [cs.CL].

Lamb, C.; Brown, D. G.; and Clarke, C. L. A. 2018. Evaluating computational creativity: An interdisciplinary tutorial. *ACM Comput. Surv.* 51(2):28.

Riedl, M. O. 2015. The Lovelace 2.0 test of artificial intelligence and creativity. In *Proc. 29th AAAI Conf. Artif. Intell. Workshops*.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds & Mach.* 17(1):67–99.

Ritchie, G. 2012. A closer look at creativity as search. In *Proc. Int. Conf. Comput. Creativity (ICCC 2012)*, 41–48.

Shannon, C. E., and Weaver, W. 1949. *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.

Shannon, C. E. 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27:379–423, 623–656.

Shannon, C. E. 1959. Coding theorems for a discrete source with a fidelity criterion. In *IRE Nat. Conv. Rec., Part 4*, 142–163.

Tishby, N.; Pereira, F. C.; and Bialek, W. 1999. The information bottleneck method. In *Proc. 37th Annu. Allerton Conf. Commun. Control Comput.*, 368–377.

Varshney, L. R.; Pinel, F.; Varshney, K. R.; Bhattacharjya, D.; Schörgendorfer, A.; and Chee, Y.-M. 2019. A big data approach to computational creativity: The curious case of Chef Watson. *IBM J. Res. Develop.* 63(1):7:1–7:18.

Varshney, L. R. 2008. Transporting information and energy simultaneously. In *Proc. 2008 IEEE Int. Symp. Inf. Theory*, 1612–1616.

Varshney, L. R. 2019a. Mathematical limit theorems for computational creativity. *IBM J. Res. Develop.* 63(1):2:1–2:12.

Varshney, L. R. 2019b. Must surprise trump information? *IEEE Technology and Society Magazine* 38(1):81–87.

Velardo, V., and Vallati, M. 2016. A general framework for describing creative agents. arXiv:1604.04096 [cs.AI].

Ventura, D. 2019. Autonomous intentionality in computationally creative systems. In Veale, T., and Cardoso, F. A., eds., *Computational Creativity*. Cham, Switzerland: Springer. 49–69.

Wiggins, G. A. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.

Wiggins, G. A. 2019. A framework for description, analysis and comparison of creative systems. In Veale, T., and Cardoso, F. A., eds., *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*. Springer. 21–47.

Wimsatt, Jr., W. K., and Beardsley, M. C. 1946. The intentional fallacy. *The Sewanee Review* 54(3):468–488.

Proceedings of the 11th International Conference on Computational Creativity (ICCC'20)
ISBN: 978-989-54160-2-8