Exploring Application Domains for Computational Creativity

Ashish Jagmohan, Ying Li, Nan Shao, Anshul Sheopuri, Dashun Wang
IBM TJ Watson Research Center, Yorktown Heights, NY

Lav R. Varshney
University of Illinois, Urbana-Champaign, IL

Pu Huang
Marvelous Face, Inc., NY

Abstract

We are motivated by the recent application of computational creativity in the culinary domain. Given the increasing commercial importance of data-driven computation, we explore and provide unified framework in three new domains to which computational creativity can be applied and yield business value. The three domains are travel, fashion, and science. Reflecting on the framework characterization, we identified two properties common across these domains, related to the creative space and codified domain knowledge. We believe that these properties may have value as sufficient, but not necessary, conditions to identify domains suitable for industrializing computational creativity. We are working towards finding tight properties common across different domains as well as ones that exclude domains.

1 Introduction

Computational creativity is the study of how computers can create, or help create artifacts that humans perceive to be creative. The field attempts to better understand human creativity and design programs that can enhance human creativity, bringing together ideas from artificial intelligence, cognitive psychology, design, philosophy and the arts. An overview of this field can be found in (Colton and Wiggins 2012).

A recent attempt in culinary computational creativity ((Varshney et al. 2013; Pinel and Varshney 2014)) motivates the work presented in this paper. The system described in that study can create novel and flavorful recipes as perceived by people. It gathers data from culinary science as well as other domains such as hedonic psychophysics, establishes evaluation metrics based on codified expert knowledge in recipe design and human flavor perception, and can create quintillions of recipes which are far beyond the number of existing recipes. We describe this work in Section 2.

Given the increasing commercial importance of data-driven computation, we explore whether a similar design framework can be extended and bring business value to other domains which are analogous to the food domain in terms of compositional models for artifacts. Focusing on fashion, travel, and science, we describe how a computational creativity framework can be developed in Section 3.

In Section 4, we reflect upon the framework, and argue that there are two properties which appear to be common across the aforementioned domains, and they are related to the combinatorial complexity of the domain creation space, and the state of the codified domain knowledge respectively.

In this preliminary position paper, we argue that while these two properties may not be necessary conditions to identify domains wherein computational creativity yields business value, they are sufficient to provide a general framework for computational creativity in industrial deployment. We are working towards finding tight properties as well as ones that exclude unsuitable domains.

2 Culinary Creativity: Case Study

Culinary design has long been seen as a creative domain in the history of human creativity research. "Made up a recipe" is one of the 100 creative activities listed on the first human creativity rating questionnaire developed by Torrance in 1962 (Sawyer 2012). But can a computer be creative for culinary recipes? With the availability of large-scale online recipe repositories in recent years, some recipe design principles have been validated using a data-driven approach, such as the food pairing hypothesis (Ahn et al. 2011). In addition, human flavor perception is gradually being uncovered by advanced scientific study of food chemistry, hedonic psychophysics and neurogastronomy (Haddad et al. 2010; Shepherd 2006). These efforts have made computational creativity possible for generating novel and flavorful recipes.

In fact, (Morris et al. 2012) discussed a recipe creation system which was restricted to soups, stews and chili. The more recent culinary computational creativity system (Varshney et al. 2013; Pinel and Varshney 2014) is more general and has a cognitive flavor assessment component motivated by the scientific study of human flavor perception. The recipes created by this system have been served in multiple venues and have been well received. An independent assessment done by *Wired* (Davis 2013) of a recipe created by the system concluded that "while the IBM dessert tasted better, it was also insanely elaborate, so we'll call it a draw."

We briefly describe the system here and characterize the culinary domain to understand why and how computational creativity brings value to this domain. The system used Wikia recipes as an inspiration set (around 25,000 recipes), and it can produce quadrillions or even more newly created recipes. Generally speaking, the volume of existing recipes from recipe repositories is usually around tens of thousands, possibly up to millions. So the inspiration set is large enough for us to draw prior information. The dimensionality of the culinary domain is captured by the number of possible ingredients which is in the range of hundreds or thousands. Looking at the combinatorial complexity of recipes which may contain 10-20 ingredients, the creativity space can be in the scale of quintillions, a much larger number than the size of the existing recipes. There is plenty of room for creativity in generating novel artifacts.

As mentioned before, there exists codified knowledge in the scientific study of human flavor perception and culinary design principles. The system gathers data and organizes such knowledge in a structured model and therefore can provide quality and novelty assessment (pairing, pleasantness and surprise) of the creative artifacts in such a large design space that people would not be able to do so. Indeed, tracking so many ingredients and reasoning their combinations to build quintillions of ideas are only feasible in a computational creativity system.

3 Application Domains

Having developed and deployed a system in the trillion-dollar food industry, are there other domains where computational creativity can bring business value? There should be enough room for creativity in a domain so that creating new artifacts is more valuable than searching among existing ones. There should also be codified knowledge for a computer to learn in a structured way in order to establish formalized predictors for creativity in terms of quality and novelty, so that a human expert can gain support from the computational system rather than relying only on the intrinsic human expertise. Following these thoughts, we explore three domains in this section: travel, fashion and scientific discovery. The results are summarized in table 1.

3.1 Travel

The advent of large-scale online networks over the last two decades has affected a fundamental transformation of how the travel sector interacts with and sells to customers. Although online travel sales now account for nearly $100 billion, there is a high dissatisfaction rate (53%) among customers, and while most focus is on price competition, little concern is given to the added value that digital channels can bring to customers (Carey, Kang, and Zea 2012; Peterson 2011).

While no comprehensive personalized travel planning solution exists, there are websites and apps which leverage social and mobile modalities to facilitate travel planning. These include informational websites such as TripAdvisor and Fodors, niche websites such as Flextrip, and itinerary planning and organization websites and mobile apps such as TripIt, Plannr, and mTrip. There is also some prior research on using collaborative filtering to recommend travel packages to tourists (Liu et al. 2014).

It is our view that the travel domain offers a promising opportunity for computational creativity to drive business value. The expected artifact produced by the system is a travel "experience", which is a sequence of activity/time-range pairs. Here, an activity may denote visiting a specific destination such as a cultural artifact, or a physical activity such as taking a specific tour. A time-range is the time period over which that activity is to be carried out. Designing a travel experience is a search in a high-dimensional space defined by the Cartesian product of possible activities and time-ranges. The inspiration set can be existing travel packages, or itineraries culled from social media as suggested, for example, in (Lempel et al. 2014).

The high dimensionality of the travel experience space, and the many possibilities of extracting inspiration itineraries, combine to make ample room for creativity. As witnessed by the relatively low satisfaction rates of current travelers, the design of a personalized and comprehensive travel experience is a creative endeavor with non-trivial difficulty. Further, there exists rich domain literature studying the issue of travel satisfaction from psychological and sociological perspectives (del Bosque and Martin 2008). Thus the gap between codified knowledge and intrinsic customer expertise is significant. We believe that computational creativity is uniquely suited to designing satisfying travel experiences, by marrying computationally-intensive learning from big data with expert insights.

We propose two chief metrics of goodness. Personalized novelty measures how different an experience is from prior itineraries and the user's prior travel experiences. Travel satisfaction measures how likely the user is to be satisfied at the end of the travel experience, and requires codified domain knowledge to compute. An example of such domain knowledge is the cognitive-affective model for travel satisfaction derived and validated in (del Bosque and Martin 2008). The key finding is that overall travel satisfaction (and travel loyalty) is driven by an interplay of cognitive and emotional aspects, including destination image, trip expectations, and positive and negative emotions accumulated during the trip. Specifically, high travel satisfaction is driven by positive expectations which are disconfirmed positively during the trip. This understanding points to how a computational creative system can design travel experiences so as to maximize personalized satisfaction, by leveraging a user's personal notions of destination image and expectations.

3.2 Fashion

Creating fashion artifacts is challenging, both due to the fact that there are many factors to weigh in (such as fashion style, color and fabric) and many design options (such as pockets and belts) to play with. In addition, even without taking the design aspect into account, creating good and tasteful outfits from a set of given clothing articles and subsequently ranking them based on certain criteria is a challenging problem. In this section, we discuss and formulate this problem with a focus on outfit creation based on individuals' wardrobes.

Consider a wardrobe containing clothing articles: we need to find an outfit with a combination of clothing articles that meets particular requirements; the goal is to create an outfit that is both aesthetically pleasing and satisfying. Equivalently, we can generate a list of outfits then rank them based on certain metrics. There has been some prior work on this front. For example, Lin *et al.* (Lin et al. 2012) described a personalized clothing recommendation system

Table 1: Characterized Domains

	Culinary	Travel	Fashion	Scientific Discovery
Output Artifacts	**Recipe**: Mixture of ingredients.	**Travel experience**: Sequence of (activity, time-range) pairs.	**Dress**: A set of outfits that are aesthetically pleasing.	**Hypotheses**: A set of existing literatures.
Volume of Inspiration Set	Existing recipes from recipe repositories.	Existing itineraries, from travel packages or social media.	Existing examples of aesthetic/stylish dress examples.	Pool of concepts and relations (published connections).
Dimensionality	High: Ingredients.	High: Activity$\times$ Time.	High: Top $\times$ Bottom $\times$ Any additional layers.	High: trivial and non-trivial combinations.
Metrics of Goodness	**Surprise**: Difference from inspiration recipes. **Pleasantness**: Likely pleasantness of recipe.	**Novelty**: Against inspiration itineraries, user experience. **Satisfaction**: Likely user satisfaction.	**Surprise**: Style difference against personal inspirations. **Aesthetics**: Color and pattern matching.	**Impact**: how many citations a certain combination of concepts may receive?
Codified Expert Knowledge	Principles of flavor pairing. Principles of pleasantness.	Psycho-social principles of travel satisfaction.	Fashion design, color science, psycho-social dress principles.	Metaknowledge, Swanson hypotheses.

based on a modified Bayesian network. Specifically, the system constructs the outfit by first selecting a top, then a bottom which matches the selected top. Another related work is proposed by Shen *et al.* (Shen, Lieberman, and Lam 2007), where each clothing item is first labeled with brand, type and a sentence to describe its style; and then the user tells the system about a particular occasion in her mind; finally, based on commonsense reasoning, the system matches the clothes' styles and functions with the concepts needed for the context, and returns suggestions for complete outfits.

In the outfit creation problem, the inspiration set contains existing dress examples that are aesthetically pleasing, and one possible data source is the individual's photo album. In this case, the inspiration set is large enough for a meaningful outfit creation, while not so large compared to the combinatorial space of dress artifacts. Moreover, we can leverage specialized principles in fashion design, color science, and even psychology and sociology. There is significant prior literature that codifies such knowledge, and a significant gap exists between expert knowledge and individual knowledge.

A complete personalized outfit creation system consists of five components: 1) *catalog of personal wardrobe*, which records an individual's wardrobe including both the clothing articles and their features (*e.g.*, the garment type, color, pattern, fabric, brand, *etc.*); 2) *personal needs or requirements collection*, examples of which are specific dressing occasions (e.g. evening party or daily work), context information (weather, season), and user profession and age; 3) *outfit creation strategy*, which determines how to generate the list of outfits. Both sequential and integrated approaches can be applied here. Specifically, the sequential approach creates an outfit by selecting the needed clothing articles piece by piece based on certain criteria. In contrast, the integrated approach learns "good" outfit examples from existing knowledge and creates an outfit as a single artifact that it deems "good"; 4) *ranking metrics*, which will be applied to rank the generated outfits based on formal design principles such as color and pattern matching, and novelty value; and 5) *system evaluation*, probably best conducted through a user study.

3.3 Genesis of Scientific Hypotheses

Over the past years, the exponential expansion in knowledge is changing the landscape of science, representing both pressing challenges and exciting opportunities. Indeed, the volume of scientific papers has increased to the extent that no individual can read all papers within a field.

We take one recent study in the context of biomedical chemistry (Foster, Rzhetsky, and Evans 2013) as an exemplary case to illustrate the process of applying computational creativity to generating scientific hypothesis. The first challenge is to define the artifacts and the items within each artifact. This corresponds to defining the underlying space of possible search paths and conceptual entities within the space. The network of scientific knowledge proposed by (Girvan and Newman 2002) and taxonomy of research strategies building on top of this network (Foster, Rzhetsky, and Evans 2013) provides a promising direction for constructing such a conceptual space with semantic entities. For example, Foster *et al*. analyzed 6.5 million abstracts in biomedicine and biomedical chemistry to construct a network of relations between chemicals. One can use this network as a representation of knowledge, hence each artifact corresponds to a study into the relationship between chemicals, with items being chemicals involved in the study (Evans and Foster 2011). On another coarse-grained level, applying community detection algorithms to this network yields knowledge clusters, corresponding to tightly related concepts. In this view, items within each artifact are represented by knowledge clusters. A key insight in this process comes from citations. As citations are often taken as proxies of impact, one can study how and why certain combination of conceptual entities within the knowledge representation would generate artifacts with higher impact.

Taken together, computational genesis of high-quality scientific hypotheses is an active and promising line of inquiry, mainly following two directions. On one hand, there have been a number of fascinating studies into clever mechanisms of combining existing knowledge. Besides biomedical chemistry, there are also literature-based discovery methods pioneered by (Swanson 1987) and more recently combination of novelty and conventionality through co-citation pairs by (Uzzi et al. 2013).

4 Discussion

Reflecting on the computational creativity framework developed in Section 3, we find that there are two common properties across these domains. The first property is related to the combinatorial complexity of the creation space and its relation to the number of extant inspiration artifacts. On one hand, the size of the inspiration set is suitably large for a data-driven approach to learn basic cultural principles of the domain. On the other hand, the full combinatorial creation space is significantly larger than the inspiration set, so that creating new artifacts is more valuable than searching among existing ones. The second property is about the cognitive difficulty of evaluating artifacts. Codified knowledge exists and can be learned by computer, and therefore data-driven predictors of novelty and domain-appropriateness can be deployed for evaluation and selection of ideas. In this case, there is a significant computable knowledge asymmetry in favor of a computational creativity system than a human expert with intrinsic expertise, that computers can quickly access more knowledge than human creators.

A foundation of creativity is knowledge, and codified knowledge exists for many domains. A computationally creative system needs to effectively and efficiently represent, manipulate, and reason with such codified knowledge in application domains. Organizing such knowledge into a well-structured scheme or model may not, however, be easy. Identifying domains of industrial importance where there is an ability to learn about parts and combining rules from examples is therefore crucial.

Exploring the whole creation space is possible for some application domains but for many others, this space is combinatorially large. For such cases, we need creativity metrics to carve out the space for efficient selection, though finding good heuristic metrics can be a process of trial-and-error, and as much art as science. Principles from psychology, however, provide a good starting point.

We believe the two properties discussed here to be sufficient conditions, but not necessarily necessary conditions, to identify domains suitable for computational creativity in industrial deployment. We are working towards finding tight properties common across different domains, as well as ones that exclude domains.

References

Ahn, Y.-Y.; Ahnert, S. E.; Bagrow, J. P.; and Barabasi, A.-L. 2011. Flavor network and the principles of food pairing. *Sci. Reports* 1:196.

Carey, R.; Kang, D.; and Zea, M. 2012. The trouble with travel distribution. *Mckinsey and Company, Insights and Publications*.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier. *Proc. Euro. Conf. A. I.*

Davis, A. P. 2013. Digital gastronomy: When an IBM algorithm cooks, things get complicated–and tasty. *Wired* 20(10).

del Bosque, I. R., and Martin, H. S. 2008. Tourist satisfaction a cognitive-affective model. *Ann. Tour. Res.* 551–573.

Evans, J. A., and Foster, J. G. 2011. Metaknowledge. *Science* 331(6018):721–725.

Foster, J. G.; Rzhetsky, A.; and Evans, J. A. 2013. Tradition and innovation in scientists' research strategies. *arXiv:1302.6906*.

Girvan, M., and Newman, M. E. 2002. Community structure in social and biological networks. *Proc. Nat. Acad. Sci.* 99(12):7821–7826.

Haddad, R.; Medhanie, A.; Roth, Y.; Harel, D.; and Sobel, N. 2010. Predicting odor pleasantness with an electronic nose. *PLoS Comput. Biol.* 6(4):e1000740.

Lempel, R.; Golbandi, N.; Choudhury, M.; Feldman, M.; Amer-Yahia, S.; and Yu, C. 2014. Automatic construction of travel itineraries using social breadcrumbs. *ACM Hypertext*.

Lin, Y.; Kawakita, Y.; Suzuki, E.; and Ichikawa, H. 2012. Personalized clothing recommendation system based on a modified Bayesian network. *IEEE Int. Symp. Appl. Internet* 414–417.

Liu, Q.; Chen, E.; Xiong, H.; Ge, Y.; Li, Z.; and Wu, X. 2014. A cocktail approach for travel package recommendation. *IEEE Trans. Knowl. Data Eng.* 278–293.

Morris, R. G.; Burton, S. H.; Bodily, P. M.; and Ventura, D. 2012. Soup over beans of pure joy: Culinary ruminations of an artificial chef. *Proc. Int. Conf. Comput. Creati.* 119–125.

Peterson, S. 2011. Travel 2020: The distribution dilemma. *IBM Institute for Business Value*.

Pinel, F., and Varshney, L. R. 2014. Computational creativity for culinary recipes. *Proc. ACM CHI Conf. on Human Factors in Computing Systems*. To appear.

Sawyer, R. 2012. *Explaining Creativity: The Science of Human Innovation*. Oxford: Oxford University Press.

Shen, E.; Lieberman, H.; and Lam, F. 2007. What am i gonna wear?: Scenario-oriented recommendation. *International Conference on Intelligent User Interfaces* 365–368.

Shepherd, G. M. 2006. Smell images and the flavor system in the human brain. *Nature* 444(7117):316–321.

Swanson, D. R. 1987. Migraine and magnesium: eleven neglected connections. *Perspectives in biology and medicine* 31(4):526–557.

Uzzi, B.; Mukherjee, S.; Stringer, M.; and Jones, B. 2013. Atypical combinations and scientific impact. *Science* 342(6157):468–472.

Varshney, L. R.; Pinel, F.; Varshney, K. R.; Bhattacharjya, D.; Schoergendorfer, A.; and Chee, Y.-M. 2013. A big data approach to computational creativity. arXiv:1311.1213.