

LAV R. VARSHNEY
JUN WANG
KUSH R. VARSHNEY

Special Issue on Creativity and Food

Associative Algorithms for Computational Creativity

ABSTRACT

Computational creativity, the generation of new, unimagined ideas or artifacts by a machine that are deemed creative by people, can be applied in the culinary domain to create novel and flavorful dishes. In fact, we have done so successfully using a combinatorial algorithm for recipe generation combined with statistical models for recipe ranking and selection. However, the algorithm of creation in our prior work may be difficult for people to interpret, understand, and ultimately adopt because the process differs from the process of human creativity theorized in the psychology and neuroscience literatures. In this paper, to address this issue, we discuss how human creativity, including in the food arena, may be built on associations and how an algorithm also built on associations can be more relatable to people so they can interact with the tool more easily. We propose a computational creativity approach that extends the data mining technique of association rule mining to generate new food recipes. We illustrate this associative algorithm on real-world culinary data.

Keywords: association mining, computational creativity, food.

Computational creativity is the use of computers to generate and select ideas or artifacts that would be regarded as creative if produced by people alone. This area of study has a long history with workers motivated by disparate goals. Some researchers are after scientific understanding of the mind, whereas others have the technological goal of producing useful tools; some practitioners may try to illuminate the psychological processes underlying creativity, whereas others try to produce esthetically interesting results (Boden, 2015). Herein, we take the position of engineering researchers interested in producing tools that are useful and easy to use by people. We do not take the position of

This paper is part of a special issue of JCB devoted to work on creativity and food, guest edited by Dr Paul Sowden, Dr Anita Eves, and Professor Monique Raats, that follows on from the 2014 International Workshop on Understanding and Fostering Creativity in the Kitchen, held at the Institute of Advanced Studies, University of Surrey, UK. All articles have been peer reviewed by two reviewers.

scientific researchers or of practitioners of either kind. We are therefore concerned with the psychological plausibility of the algorithm only to the extent it enables people to interact easily with it.

Boden (2004) had put forth four questions that computational creativity may try to answer: (1) Can computational ideas help us understand human creativity? (2) Can computers do things that at least appear to be creative? (3) Can a computer appear to recognize creativity? (4) Can computers themselves ever *really* be creative (as opposed to merely producing apparently creative performance whose originality is wholly due to the human programmer)? To be clear, this paper is neither concerned with the psychological question (1), nor with the deep question (4). It is simply concerned with the engineering question (2) and a little bit with (3) insofar as it informs (2).

As such, we adopt the social constructivist definition of creativity that has been well-developed in studying human creativity (Sawyer, 2012): Creativity is the generation of an idea or artifact that is judged to be novel and also to be appropriate, useful, or valuable by a suitably knowledgeable social group. This definition can be operationalized and tested using the consensual assessment technique (CAT; Amabile, 1982), where the creativity of an idea or artifact is rated by several experts in the creative domain, and the measured creativity is the average rating of these judges. The CAT displays strong inter-rater reliability (Baer, 1993; Kaufman, Baer, & Cole, 2009; Kaufman, Lee, Baer, & Lee, 2007).

Within the broader field of artificial intelligence (AI), an initial stated goal was to make machines that are creative,¹ but this fell out of the main focus. Recently, tests for creativity similar to CAT have been proposed to augment or replace the Turing test in assessing machines' ability to exhibit intelligent behavior equivalent to that of a human (Bringsjord, Bello, & Ferrucci, 2001; Riedl, 2015). Computational creativity is again becoming a central topic in AI (Buchanan, 2001; Cardoso, Veale, & Wiggins, 2009; Colton, López de Mántaras, & Stock, 2009; Colton & Wiggins, 2012; Sundararajan, 2014).

Creativity is often described as having divergent thinking to generate ideas and convergent thinking to select the best ones. Association is a generative process typical in human creativity (Finke, Ward, & Smith, 1992) and how human memory retrieval is primarily thought to function. Classical studies of human creative breakthroughs have found that insights typically result from associations with closely related material, rather than extremely remote associations (Gough, 1976; Perkins, 1983). Further studies point to the value of remote (but not extremely remote) associations (Benedek, Könen, & Neubauer, 2012; Benedek & Neubauer, 2013). Our algorithm will have associative strength as a tunable parameter, cf. Acar and Runco (2014) for discussion of close, remote, and extremely remote associations.

When a computationally creative computer operates as a colleague or partner to people, such that the creative act is accomplished through integrated human-computer cooperation rather than autonomously (Lubart, 2005), there is often value in having the computer act more like a person. Motivated by such a mixed-initiative approach to human-computer interaction for creativity where human and computer enter into a creation conversation where each contributes ideas (Smith, Whitehead, & Mateas, 2011), herein we investigate using generative algorithms for creativity based on association.

¹ As noted in the proposal for the Dartmouth Summer School of 1956 which initiated the study of artificial intelligence (McCarthy, Minsky, Rochester, & Shannon, 2006).

Within the domain of culinary creativity, the associative generative algorithm we develop greedily pursues strong associations to build up a culinary artifact composed of parts: an ingredient list comprising ingredients. The goal is to create completely new artifacts that have never been imagined, rather than just proposing minor modifications of existing products (Shahbaz, Srinivas, Harding, & Turner, 2006). The types of associations among culinary ingredients when constructing association networks to traverse include: co-occurrence in recipes, grown in same season, similar function, similar topic, or similar chemical composition. Association mining is further used to find interesting patterns in the relationships between ingredients (Agrawal & Srikant, 1994).

New recipe creation proceeds as follows. Perform structural decomposition of a repository of extant artifacts to find the right level of granularity of components. Start with a seed for inspiration and design objectives, such as novelty, flavor, balance, or cost. Build up the ingredient list by iteratively adding components that are interestingly associated to already-included components and also optimize the design objective. Consider adding novel artifacts into the repository of extant artifacts.

We have recently developed and deployed a data-driven computational creativity system for culinary recipes that has been judged creative along dimensions of recipe novelty and flavor, measured using both impact on professional chefs and with the CAT (Varshney et al., 2013a). As detailed elsewhere (Pinel & Varshney, 2014; Pinel, Varshney, & Bhattacharjya, 2015; Pinel, Varshney, & Tounsi, 2013; Varshney, Pinel, Varshney, Schörgendorfer, & Chee, 2013b; Varshney et al., 2013a) and reviewed herein, our system design encapsulates the divergent thinking to conceive of ideas and the convergent thinking to select the best ones into distinct computational models. The selective component involves predictive analytics using chemical, psychophysical, attentional, and culinary principles, whereas the generative component uses a combinatorial stochastic sampling approach drawing on an inspiration set of extant culinary recipes. In contrast to the proposed associative algorithm, this previous system is not as relatable to people.

Note that the predictive model within the selective component does, in a sense, allow the overall system to perceive the flavors of the proposed dish. It computes flavor pleasantness and flavor pairing synergy (Varshney et al., 2013a).

To validate the new associative algorithm for generating recipe ideas, we replace the combinatorial generative algorithm in the larger system architecture, create culinary recipes, and examine these outputs.

The remainder of the paper is organized as follows. First we review some of the literature on the associative nature of human creativity, which serves as inspiration for the algorithmic development to follow. Included in this review is a discussion of associative memories, which shows the central role of stored data in creativity. Then, after giving background on the data mining technique of association rule mining, we propose a data-driven approach to create new culinary recipes based on association rules. We show example empirical results created by the proposed associative generative algorithm. Finally, we summarize and conclude.

ASSOCIATION IN HUMAN CREATIVITY

Association is the connection between two or more concepts. Mednick (1962) defined the creative process as “the forming of associative elements into new combinations which either meet specified requirements or are in some way useful.” Mednick has further

described ways in which fruitful associations may form in the mind through contiguous evocation via serendipity, similarity, or mediation. Associative elements may be evoked *serendipitously* by the contiguous appearance of requisite stimuli in the environment, usually accidentally. Associative elements may be evoked in contiguity as a result of the *similarity* of associative elements or stimuli eliciting these associative elements. Finally, associative elements may be evoked in contiguity through the *mediation* of common elements.

To understand human creativity in the culinary domain, we interviewed chef James Briscione, Director of Culinary Development at the Institute of Culinary Education, about how he designs new recipes. He noted, “I always start with one ingredient, I have to identify that main item first and, you know, what are we going to build. You want to start with the really good looking zucchini that was at the market that day, you want to start with these cool little mushrooms you’ve never seen before. You want to start with that ingredient, begin pulling in all the flavors.” He went on to describe the way he proceeds with other flavors: “For me it’s memory, it’s all on this taste memory, you don’t have to be a chef: anyone with a lot of experience and who focuses on those kinds of things—you start to build out that memory and you start accessing and grabbing from these things that you’ve seen or tasted or smelled before and start putting them into little pairs and it’s one of those things that evolves.” Finally noting that “This ingredient grabs that one, and that ingredient grabs another one that you wouldn’t necessarily have thought of with the first one but you start building this chain and that’s where the really interesting things start to happen.”

Evidently Briscione’s introspective view of culinary recipe construction follows exactly the associative paradigm described classically, with associations built largely on memory and information retrieval through similarity and mediation.

Associative thought processes require links, which we can label with the relationships they express. Although not strictly necessary, such labeling of links has been found quite useful in classical experiments for simulating human memory with an associative model (Anderson, 1972; Quillian, 1968; Schvartz, Burger, & Simmons, 1970). For culinary creativity, associative relationships among ingredients could be due to common usage in regional cuisines, shared flavor compounds, agricultural similarities in terms of harvesting season, or even linguistic similarity in names. This is depicted in Figure 1. In later sections, we will develop implementations of these associations and resultant associative creativity with algorithmic precision.

ASSOCIATIVE MEMORY

As chef Briscione notes, association is all about memory. In fact, neuroscientists have long described various brain regions involved in memory as working associatively. A notable example is the piriform cortex, which is involved in olfactory sensing and memory (Wilson & Sullivan, 2011), and therefore flavor perception (Shepherd, 2012). A characteristic of odor memory in humans is the unique ability of odors to vividly trigger evocation of emotional experiences. This is perhaps because of the direct connection between the piriform cortex and the limbic system (amygdala and hippocampus) involved in emotion and general memory. Unlike other sensory modalities that are mediated through several cortical layers, there is also a direct connection to the orbitofrontal cortex involved in decision-making (Mouly & Sullivan, 2010).

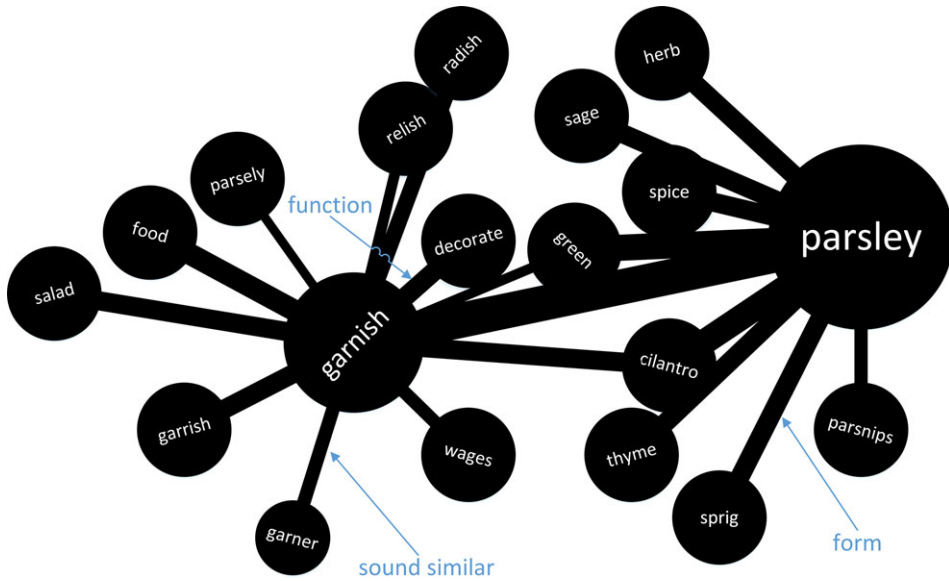


FIGURE 1. Items connected to other items for many possible different reasons via many possible relationships.

Broadly, our memory content retrieval is done by stimulating with external or internal inputs: one piece of information gives rise to related information that was stored earlier. So we see a strong similarity between Mednick's description of associative creativity and associative recall.

Moving to modern neuroscience, recent work argues that the distributed, content-addressable architecture of human memory is critically important for creativity. Without a distributed character, there would be no overlap between items that share microfeatures and so no way of forging associations between them. Without content-addressing, associations would not be meaningful. Representations sharing features are encoded in overlapping distributions of neurons, so activation can spread from one to another (Gabora & Ranjan, 2013). In fact, certain neuroactive drugs can increase or decrease the associative range of conceptual thinking (Boden, 2013).

Associative memories are used in a variety of information retrieval settings (Chen, 1995) and in computing applications including memory caches for processing units (Chisvin & Duckworth, 1989), relational databases (Lin, Smith, & Smith, 1976), among many others. Typical connectionist structures of associative memories include Hopfield neural networks (Hopfield, 1982), Kohonen maps (Kohonen, 1988), Boltzmann machines (Ackley, Hinton, & Sejnowski, 1985), and constructions based on sparse graph codes (Karbasi, Salavati, Shokrollahi, & Varshney, 2014); connectionist approaches have been described for creativity systems. Herein, however, we use symbolic algorithms for association. As far as we know, specific associative recall techniques (whether connectionist or symbolic) have not previously been deployed as part of culinary computational creativity systems.

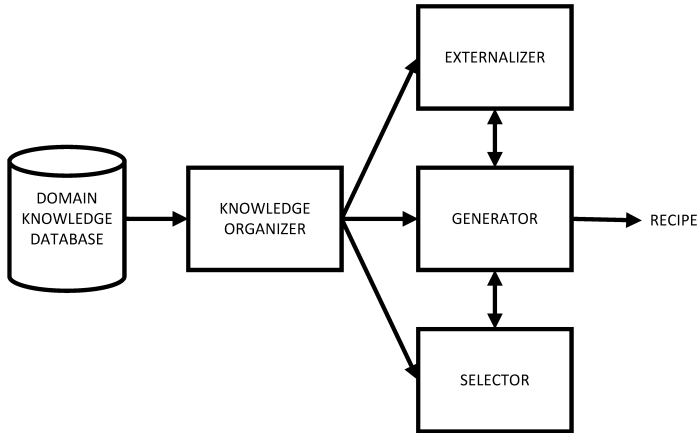


FIGURE 2. Block diagram of computational creativity system that produces a culinary recipe.

A COMPUTATIONAL CREATIVITY SYSTEM ARCHITECTURE

In this section, we review a general, modular system architecture for data-driven computational creativity that has been described previously (Pinel et al., 2015; Varshney et al., 2013a) and has structural similarity to other culinary computationally creative systems (Morris, Burton, Bodily, & Ventura, 2012). Note that this is different from certain optimization-based approaches to computational creativity which do not separate the generative and the selective parts of the system into separate modules, cf. Shao, Murali, and Sheopuri (2014). Modularity promotes flexibility in the design of cognitive tools.

A block diagram for the computational creativity system architecture is given in Figure 2, with three main algorithmic components: a generator, a selector, and an externalizer, which interact to produce a culinary recipe. These components are fed by a database of domain knowledge and a way to organize this knowledge. It is important to note that in the system architecture, the externalizer and the selector do not directly interact, but only do so through the idea generator.

The domain knowledge database represents information collected on the culinary field, including information on styles, tastes, constituents, combinations, evolution, regionality, culture, and methods of preparation. It also includes a repository of existing recipes that have been deemed creative by human audiences. These raw data are resolved and organized, to serve as a knowledge source that the generator, selector, and externalizer components draw from. Information from related, but distinct fields to the culinary domain are also kept in the database. Much data engineering and natural language processing is required for creating and using this knowledge store.

New ideas for ingredient lists are produced by the generator. These potential design ideas are evaluated for creativity and ranked/chosen by the selector. Ingredient proportions and recipe steps for the methods are determined by the externalizer. All three components take input from the categorized database: the generator to draw inspiration for new ideas, the externalizer to learn from extant methods of preparation, and the selector

to evaluate a design idea for creativity against the repository of existing artifacts as well as against properties of ingredient constituents and combinations.

A creative computer is limited if it cannot evaluate proposed recipes for creativity. Boden (1998) emphasizes the importance of evaluation in creativity, and says this is perhaps more difficult than generating new ideas. The selector component models human perception, taste, and culture using data analytics, drawing in data from the chemistry of food and its sensory perception, cf. Lawless and Heymann (2010); Dunkel et al. (2014). It examines generated ideas along two main dimensions: novelty and quality. These metrics are defined on the basis of datasets within the culinary domain, information related to culinary science, and experimental data from hedonic psychophysics. Intriguingly, perceived novelty, familiarity, and complexity all affect flavor perception itself (Giacalone, Duerlund, Bøegh-Petersen, Bredie, & Bom Frøst, 2014). Note that there is limited training data available for complete recipes (which lie in a high-dimensional space). Hence, unlike many other data analytics problems and systems, our data-driven approach to computational creativity is not based on traditional supervised learning. We decompose recipes into ingredients and have assessment methods for these parts and for their recombination rules. Novelty can be assessed via information-theoretic or other similar quantifications of innovation within the context of all other existing ideas or artifacts in the domain of interest. One implementation of this component is detailed elsewhere (Varshney et al., 2013a).

The final component, the externalizer, determines steps and ingredient proportions needed to take the ingredient list to a complete recipe. The recipe steps provide constraints on what designs are possible and they can be optimized for efficient production, for example, using techniques from planning and operations research. Ingredient proportions may be determined by matching the nutritional and ingredient type distributions of similar extant recipes. One implementation of this component is detailed elsewhere (Pinel et al., 2015).

The generator is the lead component in the system architecture, and data-driven algorithms within it are the focus of this paper. In the next section, we start to build up algorithmic techniques for data-driven association mining.

ASSOCIATION RULES THROUGH DATA MINING

In this section, we briefly review a simple yet powerful data mining tool, *association rule mining*, covering its basic concepts, solutions, as well as the related applications. As a popular and well-investigated method, association rule mining aims at discovering meaningful relations between variables in large databases. The earliest idea for association analysis roots back to 1960s, when a method called GUHA was developed for the general purpose of data mining (Hájek, Havel, & Chytil, 1966). However, this technique was popularized more recently by Agrawal, Imielinski, and Swami (1993). A popular example of association rule mining comes from market basket analysis in retail. By analyzing transaction records, one aims to identify correlations and associations between different products that shoppers tend to buy together. Stores can use such knowledge for promotions, shelf placement, catalog design, store layout, customer segmentation, etc.

Given a series of transactions T where each transaction includes a set of products, the objective of association analysis is to generate rules like $x \Rightarrow y$, which suggests a strong relationship exists between the sale of x and y as the shoppers who buy x also buy y .

Here, both x and y can represent a collection of products, namely an *itemset*. There are two important quantities related to an association rule. We first use *support* to measure the frequency of an itemset, which is equivalent to the fraction of transactions that contains a specific itemset. For instance, the support for an itemset x can be computed as: $s(x) = \frac{\sigma(x)}{|T|}$, where the function $\sigma(\cdot)$ counts the frequency of an itemset. Similarly, we can derive the support for a union set $x \cup y$ as $s(x \cup y) = \frac{\sigma(x \cup y)}{|T|}$. Accordingly, the *confidence* of the rule $x \Rightarrow y$ is computed as:

$$\alpha = \frac{s(x \cup y)}{s(x)} = \frac{\sigma(x \cup y)}{\sigma(x)}$$

Note that *support* measures the occurring frequency of a rule and *confidence* describes the strength of the association.

Depending on the setting, there are several variants of association rule mining, including binary, continuous, and fuzzy association rules (Tan, Steinbach, & Kumar, 2006). The mining process for the association rules could be very computationally costly. Given N transactions with a total of unique products, the number of possible association rules is $O(M2^M - 1)$ and the corresponding computational complexity $O(NM2^M)$ if using exhaustive search.² However, there are some efficient algorithms developed recently to cope with data with many transactions. The Apriori algorithm is a well-known solution for mining association rules (Agrawal & Srikant, 1994). It iterates between two key steps: estimating the support using a breadth-first search strategy and then continuing exploration using the downward closure property of support. Other popular algorithms include a lattice-theoretic approach called Equivalence Class Transformation (Zaki, 2000) and the pattern-matching approach called frequent pattern-growth (Witten & Frank, 2005). Due to the empirical efficiency and success, association rule mining has been employed in a wide range of applications, ranging from marketing research (Agrawal et al., 1993) to bioinformatics (Creighton & Hanash, 2003; Oyama, Kitano, Satou, & Ito, 2002) to web

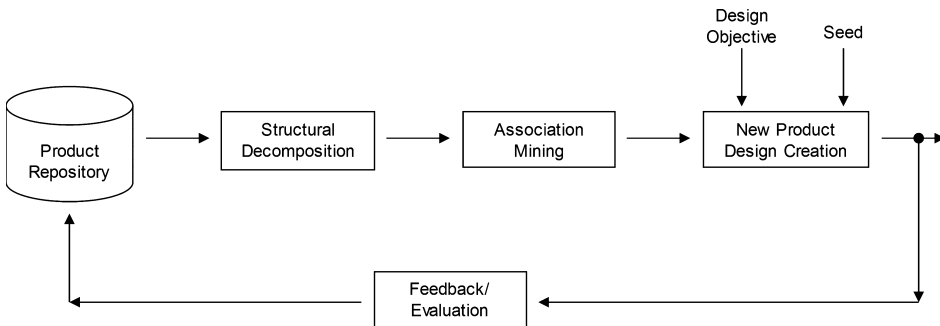


FIGURE 3. Schematic diagram of an associative algorithm for computational creativity idea generation.

² The so-called big O notation for computational complexity indicates how the algorithm processing time changes as a function of the size of the input asymptotically. The presence of exponential terms such as 2^M implies that an algorithm is not efficient.

TABLE 1. Groupings of Crops That Share Seasonality in Central Illinois

Grouping	Crops
Earliest Spring Perennials	Chives; Sorrel; Wild dandelions, Nettles; Rhubarb; Asparagus
Early Spring Greens	Spinach; Lettuce; Arugula; Cold-weather herbs (dill and cilantro); Pea Shoots; Peas (regular, snow, sugar snap); Kohlrabi; Broccoli; Chard; Spring Radishes, Spring turnips; Baby onions, Beets, Carrots
Summer's Fruits (and more)	Strawberries; Raspberries; Currants, Gooseberries; Blueberries; Aronia berries; Kale; Beans (green, yellow, purple, edamame); Potatoes; Cabbage; Cucumber; Eggplant; Basil; Tomatoes; Tomatillos, Ground cherries; Peppers (sweet, hot); Full-size beets, Carrots; Okra; Spaghetti squash; Melons
Autumn's Roots and Shoots	Bok Choi, other chois; Radicchio, Endive, Escarole; Greens (turnip, mustard); Acorn, Red Kuri, Hubbard, Delicata, other squashes; Pumpkins; Leeks, Scallions; Brussels sprouts
Winter Storage Crops	Sweet potatoes; Parsnips, Turnips, Rutabagas, Potatoes, Beets; Celery root, Parsley roots, Carrots, Winter Radishes; Daikon, Burdock, Salsify; Butternut Squash, Acorn squash, all sorts; Popcorn, Wheat berries, Dry beans; Corn meal, Flour, Oatmeal

analytics (Pei, Han, Mortazavi-Asl, & Zhu, 2000). In the next section we use association rule mining in the generative portion of the computational creativity system.

ASSOCIATIVE GENERATIVE ALGORITHM

In this section, we develop an algorithm to create novel recipes having components associated with one another by optimizing a design objective. It proceeds as follows (Figure 3 shows the operation of the algorithm as a block diagram).

The first step is to decompose recipes structurally into the appropriate level of modularity of components. So far we have talked about ingredients as this right level, but perhaps larger components such as complete sauces, spice mixtures, or other ingredient groups would make sense. This choice involves a tradeoff between flexibility and coherence, and can be inspired by the nature of the existing recipes in a “Product Repository” that is present in the knowledge database. This choice also involves determining the relationships by which components may possibly be associated.

The second step uses the association mining technique described above to find interesting patterns in the relationships between components, such as co-occurrence. For other kinds of relationships, we can use certain data sources directly. For example, to apply associations through the principle of seasonality, we can use data from growing guides that are produced by various state agriculture departments. As a typical example, for Central Illinois, Table 1 shows crops that are grown in the same season (The Land Connection (Illinois Department of Agriculture), 2015). Likewise shared flavor compounds among ingredients can be determined from appropriate databases (Burdock,

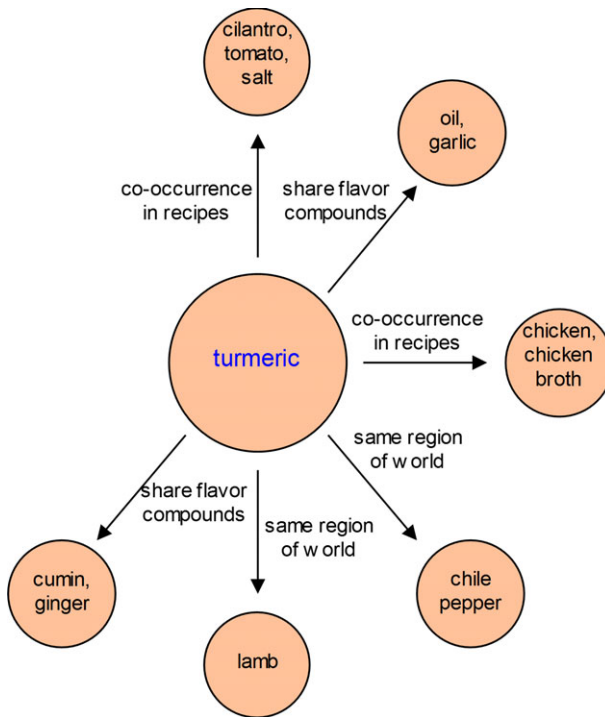


FIGURE 4. Decompose food into ingredients and consider relationships like co-occurrence in recipes, shared flavor compounds, being from same region of world, and being grown in same season of year.

2009). Figure 4 shows an example of a food ingredient, turmeric, and several other ingredients with which it has associations of various kinds.

To demonstrate association rule mining so as to learn co-occurrence relationships, we apply the Apriori algorithm to a large corpus of culinary recipes. A subset of association rules extracted from recipe data is shown in Figure 5.

The next step is new recipe creation. We start with a seed for inspiration and design objectives, such as novelty, quality, balance, or cost. The recipe is built up iteratively by adding components that are interestingly associated to already included components and also optimize the design. Iteratively adding components to the product is performed by traversing interesting associations while fulfilling a design objective. At the end, will have a complete recipe composed of several components. Feedback and evaluation steps may later add newly designed recipes to the Product Repository for future runs of the algorithm. Continuing with the turmeric-seeded example, in Figure 6, we can see how the iterative construction of the idea proceeds.

The complete flow of the iterative algorithm is depicted in Figure 7, using a visual schema taken from the SPSS Modeler software package.

By running the algorithm we have developed herein, we can generate ingredient list ideas that may be novel and flavorful, that is, creative. Three examples, with seeds of pumpkin pie spice, turmeric, and tomato, respectively are shown in Table 2.


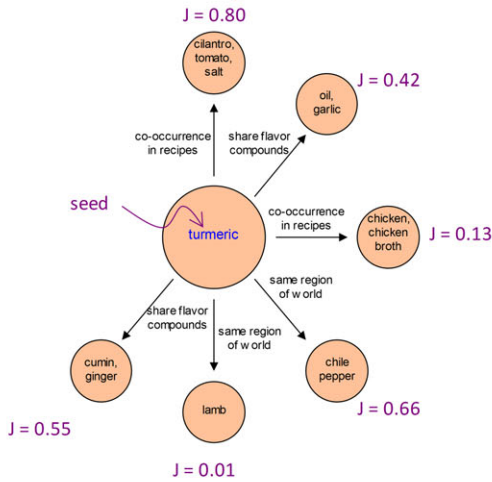
Consequent	Antecedent	Support %	Confidence %
chicken broth	chicken salt	3.819	20.270
black pepper	water onion salt	5.815	20.266
tomato	cheddar	3.333	20.258
soy sauce	ginger water	2.017	20.256
flour	vegetable oil salt	4.568	20.245
cinnamon	vanilla egg	2.594	20.232
carrot	parsley pepper	3.359	20.230
oil	tomato onion	8.615	20.220
cinnamon	flour granulated sugar	3.596	20.215
rice	chicken	6.426	20.214

FIGURE 5. Subset of learned co-occurrence associations from association mining.

design objective:

$$J = \lambda * \text{novelty} + (1 - \lambda) * \text{pleasantness}$$

The ingredients that maximize the design objective are chosen as the next ingredients to include.


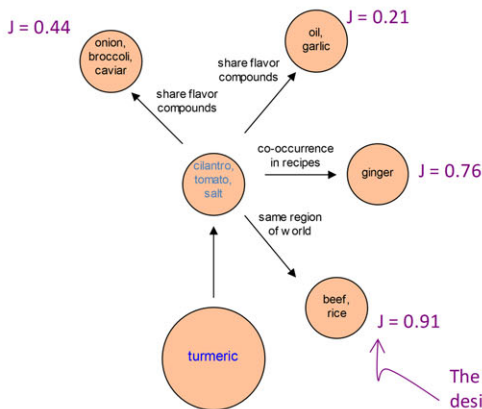


The feedback helps to refine the design objective by updating the definitions of novelty and pleasantness.

(a)

design objective:

$$J = \lambda * \text{novelty} + (1 - \lambda) * \text{pleasantness}$$



The ingredients that maximize the design objective are chosen as the next ingredients to include.

(b)

FIGURE 6. Iterative traversal of associations by optimizing a design objective, so as to build up a complete recipe. The first step is (a), and the second step is (b). The algorithm proceeds similarly for further steps.

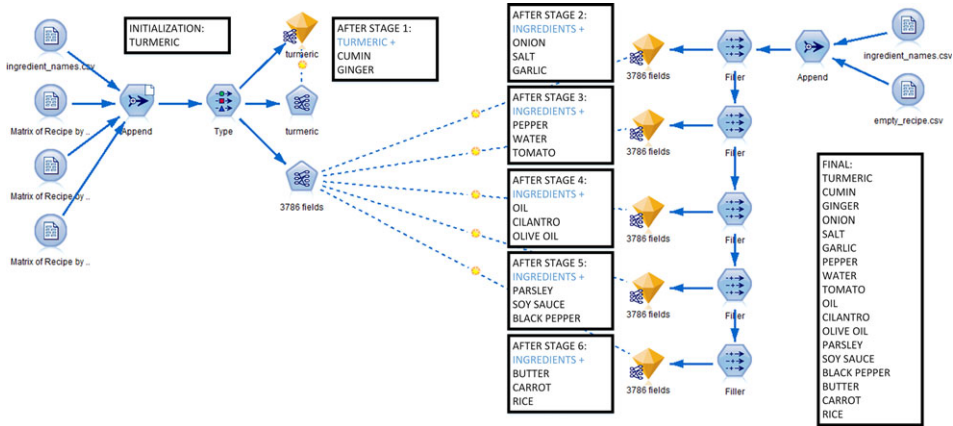


FIGURE 7. System flow diagram.

TABLE 2. Examples of Ingredient Lists Generated by Our Associative Algorithm

Seed	Ordered List of Ingredients Generated
Pumpkin pie spice	Pumpkin pie spice, pumpkin, salt, cinnamon, egg, flour, granulated sugar, butter, baking soda, baking powder, milk
Turmeric	Turmeric, cumin, ginger, onion, salt, garlic, pepper, water, tomato, oil, cilantro, olive oil, parsley, soy sauce, black pepper, butter, carrot, rice
Tomato	Tomato, green bell pepper, garlic, onion, salt, olive oil, pepper, water, black pepper, parsley, rice, red bell pepper

CONCLUSION

In the psychology literature, combinational creativity is said to involve generating unfamiliar combinations of familiar ideas (Boden, 2013). However, associative pathways are not all there is to combinational creativity; there is also the issue of relevance. Novel combinations we value, and thus regard as creative, invariably involve relevance, even if not immediately apparent (Boden, 2013).

Taking this view to design a computational creativity algorithm whose process people can relate to and collaborate with, we have proposed an algorithm that finds associations using data mining to generate combinations, but importantly also optimizes for relevance on top of the mined associations. We have illustrated its workings in the culinary domain, in which ingredients are natural components that can have various associations, including co-occurrence of flavor compounds, co-occurrence in existing recipes, similarity in color, similarity in growing season, and similarity in growing region. This approach to computational creativity, which parallels human creativity, is the first in our knowledge to utilize association rule mining.

REFERENCES

- Acar, S., & Runco, M.A. (2014). Assessing associative distance among ideas elicited by tests of divergent thinking. *Creativity Research Journal*, 26, 229–238; doi:10.1080/10400419.2014.901095.
- Ackley, D.H., Hinton, G.E., & Sejnowski, T.J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147–169; doi:10.1207/s15516709cog0901_7.
- Agrawal, R., Imielinski, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, 22, 207–216. doi:10.1145/170036.170072.
- Agrawal, R., & Srikant, R. (1994). *Fast algorithms for mining association rules*, pp. 487–499. In Proceedings of the 20th International Conference on Very Large Data Bases (VLDB'94).
- Amabile, T.M. (1982). Social psychology of creativity: A consensual assessment technique. *Journal of Personality and Social Psychology*, 43, 997–1013. doi:10.1037/0022-3514.43.5.997.
- Anderson, J.M. (1972). FRAN: A simulation model of free recall. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5, pp. 315–378). New York: Academic Press.
- Baer, J. (1993). *Creativity and divergent thinking: A task-specific approach*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Benedek, M., Könen, T., & Neubauer, A.C. (2012). Associative abilities underlying creativity. *Psychology of Aesthetics, Creativity, and the Arts*, 6, 273–281. doi:10.1037/a0027059.
- Benedek, M., & Neubauer, A.C. (2013). Revisiting Mednick's model on creativity-related differences in associative hierarchies. Evidence for a common path to uncommon thought. *Journal of Creative Behavior*, 47, 273–289. doi:10.1002/job.35.
- Boden, M.A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103, 347–356. doi:10.1016/S0004-3702(98)00055-1.
- Boden, M.A. (2004). *The creative mind: Myths and mechanisms*. London: Routledge.
- Boden, M.A. (2013). Creativity as a neuroscientific mystery. In O. Vartanian, A.S. Bristol & J.C. Kaufman (Eds.), *Neuroscience of creativity* (pp. 3–18). Cambridge, MA: MIT Press.
- Boden, M.A. (2015). Foreword. In T.R. Besold, M. Schorlemmer & A. Smail (Eds.), *Computational creativity research: Towards creative machines* (pp. v–xiii). Paris: Atlantis Press; doi: 10.2991/978-94-6239-085-0.
- Bringsjord, S., Bello, P., & Ferrucci, D. (2001). Creativity, the Turing test, and the (better) Lovelace test. *Minds and Machines*, 11, 3–27. doi:10.1023/A:1011206622741.
- Buchanan, B.G. (2001). Creativity at the metalevel. *A. I. Magazine*, 22, 13–28.
- Burdock, G.A. (2009). *Fenaroli's handbook of flavor ingredients*. Boca Raton, FL: CRC Press.
- Cardoso, A., Veale, T., & Wiggins, G.A. (2009). Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *A. I. Magazine*, 30, 15–22.
- Chen, H. (1995). Machine learning for information retrieval: Neural networks, symbolic learning, and genetic algorithms. *Journal of the American Society for Information Science*, 46, 194–216; doi: 10.1002/(SICI)1097-4571(199504)46:3<194::AID-ASIA4>3.0.CO;2-S
- Chisvin, L., & Duckworth, R.J. (1989). Content-addressable and associative memory: Alternatives to the ubiquitous RAM. *IEEE Computer*, 22, 51–64. doi:10.1109/2.30732.
- Colton, S., López de Mántaras, R., & Stock, O. (2009). Computational creativity: Coming of age. *A. I. Magazine*, 30, 11–14.
- Colton, S., & Wiggins, G.A. (2012). Computational creativity: The final frontier? In L. De Raedt, C. Bessiere & D. Dubois (Eds.), *ECAI 2012: 20th European Conference on Artificial Intelligence* (pp. 21–26). Amsterdam: IOS Press BV; doi: 10.3233/978-1-61499-098-7-21.
- Creighton, C., & Hanash, S. (2003). Mining gene expression databases for association rules. *Bioinformatics*, 19, 79–86. doi:10.1093/bioinformatics/19.1.79.
- Dunkel, A., Steinhaus, M., Kotthoff, M., Nowak, B., Krautwurst, D., Schieberle, P., & Hofmann, T. (2014). Nature's chemical signatures in human olfaction: A foodborne perspective for future biotechnology. *Angewandte Chemie International Edition*, 53, 7124–7143. doi:10.1002/anie.201309508.
- Finke, R.A., Ward, T.B., & Smith, S.M. (1992). *Creative cognition: Theory, research, and applications*. Cambridge, MA: MIT Press.
- Gabora, L., & Ranjan, A. (2013). How insight emerges in a distributed, content-addressable memory. In O. Vartanian, A.S. Bristol & J.C. Kaufman (Eds.), *Neuroscience of creativity* (pp. 19–44). Cambridge, MA: MIT Press.
- Giacalone, D., Duerlund, M., Bøegh-Petersen, J., Bredie, W.L.P., & Bom Frøst, M. (2014). Stimulus collative properties and consumers' flavor preferences. *Appetite*, 77, 20–30.

- Gough, H.G. (1976). Studying creativity by means of word association tests. *Journal of Applied Psychology*, *61*, 348–353.
- Hájek, P., Havel, I., & Chytil, M. (1966). The GUHA method of automatic hypotheses determination. *Computing*, *1*, 293–308. doi:10.1007/BF02345483.
- Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, *79*, 2554–2558.
- Karbasi, A., Salavati, A.H., Shokrollahi, A., & Varshney, L.R. (2014). Noise facilitation in associative memories of exponential capacity. *Neural Computation*, *26*, 2493–2526. doi:10.1162/NECO_a_00655.
- Kaufman, J.C., Baer, J., & Cole, J.C. (2009). Expertise, domains, and the consensual assessment technique. *The Journal of Creative Behavior*, *43*, 223–233. doi:10.1002/j.2162-6057.2009.tb01316.x.
- Kaufman, J.C., Lee, J., Baer, J., & Lee, S. (2007). Captions, consistency, creativity, and the consensual assessment technique: New evidence of reliability. *Thinking Skills and Creativity*, *2*, 96–106. doi:10.1016/j.tsc.2007.04.002.
- Kohonen, T. (1988). *Self-organization and associative memory*. New York: Springer-Verlag.
- Lawless, H.T., & Heymann, H. (2010). *Sensory evaluation of food: Principles and practices* (2nd edn). New York: Springer.
- Lin, C.S., Smith, D.C.P., & Smith, J.M. (1976). The design of a rotating associative memory for relational database applications. *ACM Transactions on Database Systems*, *1*, 53–65. doi:10.1145/320434.320447.
- Lubart, T. (2005). How can computers be partners in the creative process. *International Journal of Human-Computer Studies*, *63*, 365–369. doi:10.1016/j.ijhcs.2005.04.002.
- McCarthy, J., Minsky, M.L., Rochester, N., & Shannon, C.E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence. *A. I. Magazine*, *27*, 12–14. doi:10.1609/aimag.v27i4.1904.
- Mednick, S.A. (1962). The associative basis of the creative process. *Psychological Review*, *69*, 220–232. doi:10.1037/h0048850.
- Morris, R.G., Burton, S.H., Bodily, P.M., & Ventura, D. (2012). *Soup over beans of pure joy: Culinary ruminations of an artificial chef*, pp. 119–125. Proceedings of the International Conference on Computational Creativity (ICCC 2012).
- Mouly, A.-M., & Sullivan, R. (2010). Memory and plasticity in the olfactory system: From infancy to adulthood. In A. Menini (Ed.), *The neurobiology of olfaction* (pp. 119–125). Boca Raton, FL: CRC Press.
- Oyama, T., Kitano, K., Satou, K., & Ito, T. (2002). Extraction of knowledge on protein–protein interaction by association rule discovery. *Bioinformatics*, *18*, 705–714. doi:10.1093/bioinformatics/18.5.705.
- Pei, J., Han, J., Mortazavi-Asl, B., & Zhu, H. (2000). Mining access patterns efficiently from web logs. In T. Terano, H. Liu & A.L.P. Chen (Eds.), *Knowledge discovery and data mining. Current issues and new applications, Vol. 1805 of lecture notes in computer science* (pp. 396–407). Berlin: Springer; doi: 10.1007/3-540-45571-X_47.
- Perkins, D.N. (1983). Novel remote analogies seldom contribute to discovery. *The Journal of Creative Behavior*, *17*, 223–239. doi:10.1002/j.2162-6057.1983.tb00358.x.
- Pinel, F., & Varshney, L.R. (2014). Computational creativity for culinary recipes. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2014)* (pp. 439–442). doi: 10.1145/2559206.2574794.
- Pinel, F., Varshney, L.R., & Bhattacharjya, D. (2015). A culinary computational creativity system. In T.R. Besold, M. Schorlemmer & A. Smalld (Eds.), *Computational creativity research: Towards creative machines* (pp. 327–346). Paris: Atlantis Press; doi: 10.2991/978-94-6239-085-0_16.
- Pinel, F., Varshney, L.R., & Tounsi, L. (2013). *Information in networks as inspiration: Value for culinary computational creativity*. New York: Workshop on Information in Networks (WIN).
- Quillian, M.R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 216–270). Cambridge, MA: MIT Press.
- Riedl, M.O. (2015). *The Lovelace 2.0 test of artificial intelligence and creativity*. Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence Workshops.
- Sawyer, R.K. (2012). *Explaining creativity: The science of human innovation*. Oxford: Oxford University Press.
- Schwarcz, R.M., Burger, J.F., & Simmons, R.F. (1970). A deductive question-answerer for natural language inference. *Communications of the ACM*, *13*, 167–183. doi:10.1145/362052.362058.
- Shahbaz, M., Srinivas, M., Harding, J.A., & Turner, M. (2006). Product design and manufacturing process improvement using association rules. *Proceedings of the Institution of Mechanical Engineers. Part B: Journal of Engineering Manufacture*, *220*, 243–254. doi:10.1243/095440506X78183.
- Shao, N., Murali, P., & Sheopuri, A. (2014). *New developments in culinary computational creativity*. Proceedings of the International Conference on Computational Creativity (ICCC 2014).

- Shepherd, G.M. (2012). *Neurogastronomy: How the brain creates flavor and why it matters*. New York: Columbia University Press.
- Smith, G., Whitehead, J., & Mateas, M. (2011). *Computers as design collaborators: Interacting with mixed-initiative tools*. Proceedings of the ACM Creativity and Cognition Workshop on Semi-Automated Creativity (SAC 2011).
- Sundararajan, L. (2014). Mind, machine, and creativity: An artist's perspective. *The Journal of Creative Behavior*, 48, 136–151. doi:10.1002/jocb.44.
- Tan, P.-N., Steinbach, M., & Kumar, V. (2006). *Introduction to data mining*. Boston: Addison-Wesley.
- The Land Connection (Illinois Department of Agriculture) (2015). Farm fresh now! Local food time ticker. Available from: <http://thelandconnection.org/community/farm-fresh-now/> [last accessed February 2015].
- Varshney, L.R., Pinel, F., Varshney, K.R., Bhattacharjya, D., Schörgendorfer, A., & Chee, Y.-M. (2013a). *A big data approach to computational creativity*. arXiv:1311.1213v1 [cs.CY].
- Varshney, L.R., Pinel, F., Varshney, K.R., Schörgendorfer, A., & Chee, Y.-M. (2013b). *Cognition as a part of computational creativity*, pp. 36–43. In Proceedings of the 12th IEEE International Conference on Cognitive Informatics and Cognitive Computing (ICCI*CC 2013); doi: 10.1109/ICCI-CC.2013.6622223.
- Wilson, D.A., & Sullivan, R.M. (2011). Cortical processing of odor objects. *Neuron*, 72, 506–519. doi:10.1016/j.neuron.2011.10.027.
- Witten, I.H., & Frank, E. (2005). *Data mining: Practical machine learning tools and techniques*. San Francisco: Morgan Kaufmann.
- Zaki, M.J. (2000). Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering*, 12, 372–390. doi:10.1109/69.846291.

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

Jun Wang, Alibaba Group

Kush R. Varshney, IBM Thomas J. Watson Research Center

Correspondence concerning this article should be addressed to Lav R. Varshney, 123 Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801. E-mail: varshney@illinois.edu