Title
The importance of chance and interactivity in creativity

Journal
Pragmatics & Cognition: 5 –26, 22(1)

Author
Kirsh, David

Publication Date
2014

Peer reviewed

The importance of chance and interactivity in creativity

David Kirsh
University of California at San Diego

Individual creativity is standardly treated as an 'internalist' process occurring solely in the head. An alternative, more interactionist view is presented here, where working with objects, media and other external things is seen as a fundamental component of creative thought. The value of chance interaction and chance cueing — practices widely used in the creative arts — is explored briefly in an account of the creative method of choreographer Wayne McGregor and then more narrowly in an experimental study that compared performance on a Scrabble-like word discovery problem. Subjects were presented with seven letters and given two minutes to call out three-to-seven-letter English words. There were three conditions: The tiles were fixed in place, subjects were free to move the tiles manually or the tiles could be randomly shuffled. Results showed that random shuffling was best, with manual movement second. Three reasons are provided: Shuffling is faster and cheaper than mentally thinking of candidates; randomizing strings covers the search space better than a deterministic method based on past successes; and randomizing is equivalent to adding diversity to a team, which is known empirically to lead to more creative solutions.

Keywords: Creativity, problem solving, interactivity, thinking, randomness, choreography, Scrabble

Chance has a privileged role in creativity. It can be used to thwart bias, overcome the drive to imitate past solutions, and stimulate new ideas. Nature uses it; many of the most efficient algorithms use it; and so do performers, choreographers, designers and mathematicians — but not quite in the same way. In what follows, I explore how chance and interactivity operate in creative thinking.

Pragmatics & Cognition 22:1 (2014), **5–26.** DOI 10.1075/pc.22.1.01kir
ISSN 0929–0907 / E-ISSN 1569–9943

1. Creative thinking: A more interactionist perspective

My starting assumption is that creative thinking, whether in science, art or sense-making, is not something that occurs solely in the head — the internalist view. In most cases, creativity depends on an interactive cycle of working with artifacts, reacting to interim changes in the environment and then interacting again. Humans think interactively (Kirsh and Maglio 1994; Hollan et al. 2000; Clark 2006; Cowley and Vallée-Tourangeau 2010). Mozart, for example, has been held as an exemplar of internalist creativity because he once allegedly wrote, "[P]rovided I am not disturbed, my subject enlarges itself, becomes methodized and defined, and the whole, though it be long, stands almost finished and complete in my mind" (published in Rochlitz 1815). He is now thought to have worked interactively, typically needing a piano or harpsichord, the letter now discredited as forged (Keefe 2006).

Internalist accounts of thinking tend to treat thought as if it were inner speech or some other form of internal imagery (visual, musical or kinesthetic). This makes it seem that thought is exclusively under the control of the person doing the thinking, much like our choice of words seems to be under our control when we speak, a view that downplays the mutuality holding between thinkers and the environment they are embedded in. Yet what happens outside a brain often affects what happens inside. What we perceive and what we attend to inevitably prime and trigger associations that bias what we think next. Creativity does not occur in a situational vacuum. While we're busy thinking one idea, the world moves on. Depending on what happens, that next-world state may capture our attention and take a causal lead in what we subsequently think. People do not always uniquely determine their thought trajectory, so it is good to surround ourselves with provocative material — even if self-generated. Indeed, so tightly linked are we to our surrounds that nature would have missed a trick if it had not evolved us to make the most of our dynamic partnership. Just as our bodies co-evolved with our biome, so too have our brains co-evolved with the materials we create. Humans are nature's finest niche constructors, and we keep changing things moment by moment.

When we examine creative acts — making a humorous comment or gesture, solving a challenging problem, interpreting a musical sequence, seeing unexpected possibilities in a statement — these and a thousand other creative things are processes that typically involve agent–environment interaction.

Start with the philosopher's favorite: writing. A person must interact with external tokens of words through inscriptions made by pen or keyboard. Despite a writer knowing what s/he has penned, something special happens as the words accumulate. The externalized text can be read and reread, reinterpreted and critically evaluated. The text becomes a local environment for creative reflection and seeing

— an opportunity for a reflective dialogue (cf. Schön 1992; Suwa and Tversky 1997 on sketching). This leads to revisions that can be surprisingly non-local — major edits and rewrites. It is hard to imagine how writers relying on memory alone could create so many sizable and precise alterations if the text they were revising was hidden from view.

Speech is another interactive process that might feel internalist, exclusively under the control of a single agent — the speaker — but is not. Spoken words leave no persistent trace to review. It might seem that nothing is there to shape our utterances other than our own thoughts. And yet, conversation is a joint activity. Both participants co-determine conversational topic and focus. Together they create a complex space of shared meanings, intents and possibilities that inevitably affects word choices, facial expressions and hand gestures. Even when speech is private, the externalization of sounds alters a solitary person's environment through intonation and prosody. Subsequent thought is affected.

Such ideas are widely appreciated by now, as is the idea that interaction helps us to become better thinkers. Interactive thinkers *reach solutions faster* than individuals thinking on their own because external processes perform some of the work that otherwise would be performed internally (Norman 1993; Kirsh and Maglio 1994). That's true and the way the story used to be told. Of equal importance, though, is the *way* thinking changes. When people think interactively, they work differently than when thinking in isolation. The functional properties of the elements they think with drive them toward different and often *novel modes of thought*.

For instance, when an architect rotates a physical model, the external rotation is usually said to enable off-loading computation. The analog computation implicit in physical rotation complements internal thinking, and this usually saves time and mental effort, especially if the rotation is complex and otherwise would have to be attempted in the head. But examine the interaction closely, and it can be seen that more is at work; the way thought unfolds is different. Because physical models have mass and form they have different manipulation profiles than mental models and images. They can be studied, touched and observed; they can be arbitrarily complex and detailed; and dynamic properties such as deformation under pressure can be implemented with greater precision than in their mental counterparts. They can be moved faster and put in *surprising* positions. And importantly, what can be seen in the visual display of models, especially physical ones, is unceasingly ambiguous. These enhanced visual and kinesthetic possibilities partly explain why physical and computer models are so important in technology and science. Our interaction with external things becomes a driving force of what is thought next. Interaction extends our cognitive reach.

Owing to interaction with external things, it is best to see processes inside an agent as just one component in a larger network of parts extending far beyond brain and body. With external tools and resources, a person and his or her environment form a distributed system — an unequal partnership but a partnership nonetheless — that spreads state and control across brain, body and environment (Clark 2008). We are bound in complex ways to the outside and have evolved to count on this binding. For instance, because our brains are tuned to the location and behavior of things, we can time when and where to look next for relevant information. Our posture and head position hold state: They orient us to what we should attend to, given what we have just done and what is likely to happen next. Jointly, these three things — brains, bodies and environment — have a profound effect on the microstructure of what we will subsequently do. They hold the distributed state that explains the flow of thought and activity far better than internalist theories can.

The importance of interaction is even more apparent when we examine how different media idiosyncratically shape the time course of thought. Consider how people think through drawing vs. thinking through writing, or through talking with colleagues, or how a violin phrase is thought through by the activity of practicing. Each medium — talk, drawing, music making — with its distinctive tool set has its own governing principles, inevitably influencing how thought progresses. We think and perform differently in different media.

Take drawing. When people draw, their cognition is partly driven by the representational properties intrinsic to the formalism of drawing and partly by the physics of drawing, which is derived from the way hands, arms and pencils work. On the formalism side, drawing is a representational system with certain expressive and stylistic biases. It is good at making certain attributes explicit — the effect of perspective, the relative position of parts, relative length, size — but it is not good at exposing certain others like the relative weight, strength or flexibility of the objects being portrayed. These are not *on display* explicitly (Kirsh 1990). Drawing is even worse at displaying negation, conditionals and disjunctions. Thinking through drawing is an embedded activity with certain moves and biases, certain patterns and strategies of thought. Some inferences are facilitated; others are inhibited or limited. Our thought process is biased by the representation.

Videoing tool use and eye tracking can empirically deepen the story. A drawing is brought into being piecemeal, using tools like pencil and paper or stylus and tablet. Each tool set inevitably constrains how fast agents can work and where their attention will be directed. You can't jump a stylus from one part of a sketch to another if you want a continuous line. Nor can you direct attention to some other region without stopping work on the first. Distraction and dropped attention affect performance. These physical realities of eye, hand and attention constrain what is attended to and thought at each moment. They also relentlessly generate the internal state related to tracking what one is doing and must do next. The overall effect of thinking while drawing is that the choice of micro-topics and features to think about are partially forced on the drawer as s(he) moves step by step through the image-making process. *Thought unfolds through constrained interaction.* The same applies to other forms of interactive reasoning: logical proof, choreographic creation, music composition, essay writing, sculpture and so on. This follows directly from accepting that thought is not uniquely a logical or abstract process; it is a physical process spanning brain, body and environment. Creating thought and structure is an interactive process.

2. Opportunism, interactivity and creativity

Although the physical and embodied nature of cognition is widely discussed, the role of interactivity has not been a central focus in the psychological study of thought and creativity. (Kaufman and Sternberg 2010) Traditional experimental studies have taken an internalist perspective and treated thought as something happening entirely within an agent. The major components of creativity on this view are those that enable a thinker to break from the pack: divergent thinking and problem reframing. These have been studied primarily in insight problems, and semantic spread tasks. How should this research be rethought once interactivity is made a central component?

Insight problems and reframing. Here are two examples: The first is a language-based example, the second a visual problem.

> A1: A hospital correctly records that two children, who are not twins, have the same birth date (day, month, year) and are born to the same biological mother and father. How is that possible?

As the problem is written, most people choose a way of framing the problem that makes finding a solution impossible. To discover the solution, the initial problem representation must be restructured. This is thought to be the moment of creativity.

In this case, subjects rarely consider that the children might be triplets because the problem is approached assuming that no key information has been intentionally suppressed — in accordance with Gricean maxims of relevance (Grice 1975). But it has. Consequently, their search space does not contain the solution. Rational searching "leads to an *impasse* (Ohlsson 1992) in which progress halts; to break the impasse, the problem representation must be changed or *restructured* to allow fresh directions of search" (Gilhooly and Murphy 2005: 279–280).

Interactivity helps re-framing. For instance, the answer, for some people, will pop out if they discuss the problem with others or if they explicitly write down the givens, assumptions and constraints, as suggested by Polya (1957), and then distinguish those that come from the problem from those that are their own presuppositions. In typical laboratory studies, there are no opportunities to talk or write. In prohibiting these types of interactive thinking, psychologists miss the chance to observe problem-solving methods pervasive in everyday life. It is precisely because people do talk and work things out on paper that we value tools and collaboration.

The same truth holds for visual insight problems. In the triangle problem (Katona 1940), subjects are told to:

> A2: Arrange six sticks, all of equal length, to make four identical equilateral triangles with sides one stick long. See Figure 1.

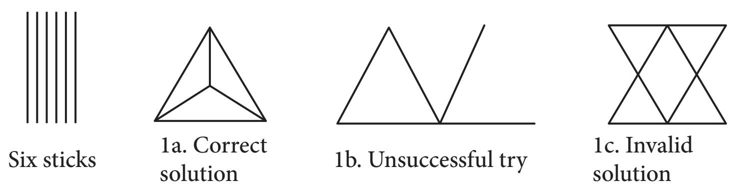

Figure 1. The figure in 1a is a solution only if the sticks are seen as assembled in three dimensions. Figure 1b represents a typical intermediate effort, while Figure 1c shows four equilateral triangles whose sides are not full sticks.

In a typical experiment, participants work alone and rarely with real sticks (though Katona's [1940] subjects did). Instead, they work interactively on paper, on a computer or non-interactively in their head, just calling out a solution (e.g., Knobloch et al. 2001). Perhaps that is why they nearly always assume the solution will be a two-dimensional pattern.[1] However, there is no two-dimensional solution. To solve the problem, subjects must think three dimensionally. This step is hard because someone drawing on paper naturally assumes that to represent sticks of equal length, lines of equal length must be used. They think flat. In a perspective-based drawing, the same stick can be represented by different length lines (see Figure 1a). Sometimes interactivity helps because a subject may notice that a shape s/he is making can be interpreted in a three-dimensional manner, encouraging a reconsideration of the problem givens. Duncker called this "intelligent fumbling" (1945: 26).

1. For an example of a study that explores the effect of putting interactivity and ecological naturalness back into insight problems, see Vallée-Tourangeau et al. (2011).

Figure 2. Children's magnetized bars. On the left, the bars are unassembled; on the right, they have been assembled into a pyramid.

When subjects work with actual 3D sticks, especially those suited to 3D construction, intelligent fumbling is far more productive. The solution invariably leaps out during manipulations because the physical possibilities of sticks (rather than lines) completely change the problem landscape, as anyone will know who has played with children's magnetized bars (see Figure 2). Experimentally, Vallée-Tourangeau et al. (2011) found subjects performed significantly better on matchstick algebra when they could move three-dimensional matches compared to working with a static illustration.

The inference to make is that experimental versions of insight problems have been good at showing that we are easily stymied, but rarely do they show *how* we normally find our way out via natural interaction. Seldom do they show how tuned we are to working with things physically. Like it or not, the world has a say in how we work, and this results in differences in the way we think interactively vs. internally. Interaction is richer and potentially more powerful. This matters for a theory of thinking.[2] But equally, it matters for a theory of creativity, since working interactively with things may be our most powerful engine of creativity.

Semantic spread problems, divergent thinking. The challenge in these problems is less about finding novel ways of looking at a problem than it is about thinking repeatedly of non-standard, non-obvious associations. For instance, in the Alternative Uses Task (Guilford 1967), subjects have two minutes to:

---

2. Vallée-Tourangeau and colleagues have argued even more strongly that "thinking is the product of an interactive assemblage of resources internal and external to the agent" (Weller et al. 2011: 429) and that "[p]roblem solving from a distributed cognitive system perspective is an emergent product of the strategic and opportunistic manipulation of artefacts populating a physical space" (Vallée-Tourangeau et al. 2011: 273).

> B. Enumerate all the ways you can imagine using a common household item, like a brick or a newspaper.

Some uses are obvious; others are more creative. For a newspaper, examples of alternative uses might include to hit flies, to be a door stop, wrap fish, shield the sun from one's eyes, light a fire, make into spitballs, pack fragile dishes, use as insulation under clothes, advertise in …

Semantic spread problems are sometimes explained as tasks requiring defocused attention (Vartanian 2009). Solutions gradually develop without a single 'aha' experience that rejigs or redefines the whole search space. The more defocused a subject's attention, the more widely dispersed are the connections they will make and the more low frequency are the associations they will find. It's like accumulating many little aha's each time a new way of seeing occurs.

The critical issue with all these laboratory tasks, whether insight or semantic spread, is that the problem is always posed in abstraction from the resources a person usually has. This is necessary for controlled experimentation, but a consequence is that the affordances and systems of cues that subjects have learned to exploit in everyday life are unavailable. We have been tool-using social animals for thousands of generations. Studies that abstract from our ecologically natural setting — our artifact-rich one — are in danger of encouraging conclusions that don't generalize well. The pulse and pattern of thinking in everyday life differs from that found in experimental cases. Whether in ateliers, offices or computer environments, hints are thrown up everywhere. Indeed, the most innovative environments are designed to make it easy to encounter and proliferate hints by offering aids like material sheds or idea carts, and supplying whiteboards, paper and working surfaces.[3] Yet insight and semantic spread problems, as studied experimentally, are hint-free. Can we be certain that experiments designed to remove these cues, hints and supports reliably expose true human capabilities and limitations? Perhaps unearthing hints is precisely what some people are good at. Interactive creatures develop interactive strategies.

Over 70 years ago, gestalt psychologists noticed the value of hints. In early work on insight problem solving by Duncker (1945) and others (Wertheimer 1961), subjects were presented with problems like the candle problem that involved working with the actual objects required for a solution rather than pictures or drawings of those objects.

In Duncker's original candle problem, the trick — as is the case with many insight problems — is to realize that an object (in this case a small box) can be used

3. "… a dedicated space with walls — a war room — always helps us do better work. The walls of a war room can extend a team's memory, provide a canvas for shared note-taking, and act as long-term storage for works in progress" (Knapp 2014).

in multiple ways: as a container, but also as a shelf or support. On a table there are three boxes presented in one of three conditions: 1) Relevant materials are boxed: candles, tacks and matches are each stored in their own box; irrelevant materials such as buttons, paper clips, etc. are scattered on the table; 2) irrelevant materials are boxed, and the relevant objects are scattered on the table; 3) empty — all task relevant materials and task-irrelevant materials are scattered on the table. The objective is to figure out how to mount three candles on a nearby door at eye height, as if for use in an old-fashioned vision experiment.

In such an environment, where subjects are picking up and putting down objects, it is easy for hints to arise at any moment. Duncker's (1945) main inquiry was not specifically about hints but rather functional fixedness, a conjecture supported by the better solution rate he found for empty boxes — the condition where boxes were not already functioning as containers. But he also observed that among the minority of subjects who succeeded in other conditions, most of them said that thinking or working with tacks was a key factor in their solutions (Duncker 1945: 92) — another instance of intelligent fumbling. Just having to work with a box to remove the candles seemed similarly to facilitate insight, as if having more "contact" with the boxes led to thinking of them more flexibly (Duncker 1945: 90; also see Glucksberg 1964). It is surprising when and why something serves as a hint to overcome functional fixedness.

The same sort of opportunistic hints were observed to be important for monkeys in solving the monkeys and banana problem. In Kohler's (1925) classic work on ape insight, a small bunch of bananas was suspended well above the reach of a monkey. In the cage, a stick and some boxes were placed at some distance from the bananas. The solution required the monkey to move the box, stand on it and then use the stick to hit the bunch.

Typically, monkeys 'pace around' waving the stick (Kohler 1925; Schiller 1957). Moving the box requires two hands. Thus, monkeys must solve a difficult problem: Give up one sub-goal (hold stick) to achieve another (move block), even though their experience in this setting is that standing on a box (by itself) will not get them to the bananas. They rarely think far enough ahead to imagine letting go of the stick, then moving the box, and then picking up the stick again before climbing the box.

What often happens, though, is that in the course of walking around, they eventually see the box and bananas line up visually. At that moment, or shortly thereafter, they seem to have a simian aha! and immediately drag the box under the bananas, recover the stick and use it to hit the bunch.

Would we be wrong to say that the monkeys were opportunistically exploiting a hint? Although the concept of a hint is not a well-defined concept in the theory of problem solving (Kirsh 2009), a hint often functions as a cue that triggers a reassessment of a move's goodness — a thing that was previously valueless is suddenly revalued. Sometimes hints are cues that trigger a rethinking of what is possible — things that were seen as prohibited or even impossible before are reconsidered. To a hapless monkey jumping for bananas, a person standing on a box waving a stick qualifies as a big hint — a strong cue — to take a second look at nearby sticks and boxes. Merely resting a stick on a box is a hint, though a considerably weaker one.

When hints are effective, they realign thinking. Goals are changed; objects of thought are replaced. Thus, if a monkey or person were initially thinking about how a stick might strike a banana, they must shift focus to think about a box and how it can extend reach. Until that moment, boxes are irrelevant — out of mind. In AI, when an external stimulus suddenly forces a subject to change goals or plans in order to take advantage of unanticipated opportunities, it is called opportunism. The classic example is shopping for dinner. If one buys ingredients according to a list based on a fixed idea about dinner (e.g., preparing a chicken stew), it precludes taking advantage of unrelated sales (sale on salmon) that would require rethinking the main course. A more opportunistic system, however, would reassess things because the top-level goal of a satisfying dinner (at a good price) can be reached in a different and better way. The system is more open to new thoughts.

The idea that chance events can serve as hints or opportunities for creative discovery in human problem-solving contexts is often referred to as an evolutionary or Darwinian approach to creativity (Campbell 1960; Simonton 2011) — the view that random variation in the form of chance mutation or arbitrary recombination is the generative mechanism producing the candidates that selection must operate on. In this Darwinian theory, though, chance does not provide hints or inspiration or empirical insight — our view; it is the literal source of variation.

The position supported here is that randomness and chance can serve as a stimulus to broaden thinking and to overcome cognitive set and functional fixedness — an approach underlying many modern creativity techniques. But again, chance is not literally the source of variation; it is the prompt — the seed — that facilitates novel invention. The chance event must be interpreted and worked with. It provides *inspiration*. This means that almost any chance idea might be made to work by a skilled and creative person. On evolutionary accounts, there is no room for follow-through to reshape a variation. There is only selection (or rejection).

This hint- or 'seed'-based approach is practiced by choreographers like Cunningham, Forsythe and McGregor who have their dancers invent movements by using exercises that rely on whatever objects and structures are nearby — found objects. The stimuli are random in that anything in the room might serve — a chair, handbag or light fixture, the edge of a brick or curtain. But dancers can show better or worse taste in the attributes they initially choose as the hints or guiding ideas from which to fashion movement. Using those stimuli as seeds, dancers are supposed to create novel movements that are not random but skilled and interesting.

In related fashion, architects often look for inspiration by crumpling or folding paper or other materials (Bergdoll et al. 2009; Vyzoviti and Dimitriou 2012) or through free association with clients (Neustein 2010). Mathematicians, analogously, explore the behavior of poorly understood functions by using computers to calculate the outcomes of randomly selected input values — a quasi-experimental approach (Bunt et al. 2009). In design and business, there are workshops devoted to teaching how to use randomness as a method for increasing generativity in the ideation phase of thinking (Reis 2013). Although the things people try out often lack immediate relevance to their problems as originally conceived, they help to break habitual thinking and lead to making new connections or to developing new ideas. Chance events trigger creativity in the actual candidate-generation process.

3. An experiment on chance in Scrabble

To explore the value of chance as a creative stimulant, we carried out a small study to test how well subjects would play a computerized Scrabble-like game in three conditions: 1) Shuffle: Hitting a 'shuffle' button randomly shuffled seven Scrabble tiles; 2) Interactive: Subjects could manually move tiles using a mouse; and 3) Static: Scrabble tiles were fixed in place.[4] Scrabble resembles insight problems in being plagued by mental set and functional fixedness as shown by the little aha people often feel when finding a word that was invisible before; it also somewhat resembles semantic spread problems because performance often benefits from defocused attention (Martingdale 1999). Using the shuffle button rarely produces directly usable Scrabble words — as with true mutation — but new words are reliably stimulated.

3.1 Method

Participants

We tested 26 subjects (11 females, 15 males, average age 25) recruited from students at the University of California, San Diego and the local population. All participants were native English speakers, defined as having learned English before the age of five.

4. This is a revision of an experiment we ran many years ago (Maglio et al. 1999) that now includes the random reset condition and a new set of analytical techniques.

Materials

There were four stimuli consisting of strings of seven letters such as 'ETCUSFA' in each of three conditions, each capable of making about 80–110 three-to-seven-letter words. Each string supported about the same number of high- (>100 per 10^6 words) and medium-frequency (10–100 per 10^6) words as calculated from the word frequency table in Google Ngram. Subjects sat in front of a computer with a Web application that we built. Tiles could be moved only in the shuffle and interactive conditions (randomly in shuffle and manually in the interactive condition) by dragging images of the Scrabble tiles.

The main dependent measures were the number of valid words produced (not counting repetitions) and their word frequency (High, Medium, Low and Very Low). A video camera was used to observe gestures and make it easier to transcribe subject speech.

Procedure

Participants were given three minutes to call out all the English three-, four-, five-, six- and seven-letter words they could compose from the seven-letter stimulus. They were told that non-English words, repeat words and made-up words would be ignored. They were then given a practice run in each of the three conditions, and their questions were answered.

The experiment was run using a Latin Square design in two blocks of six trials (two trials per condition) with a break of three minutes between the blocks. The order of the conditions was counterbalanced within and between blocks. Before the experiment, subjects filled in a brief questionnaire about their experience with word games, education and handedness.

3.2 Results and discussion

Results

Our main prediction was that movement — both the interaction and shuffling conditions — would be better for word discovery than viewing motionless (static) tiles, and specifically that shuffling tiles would be best for creativity, followed by standard interaction. Shuffling, moreover, should serve as a safety net saving subjects from doing excessively badly. We also predicted that less common words (i.e., lower-frequency words) would be more numerous in the shuffle condition than in either of the other two conditions, suggesting that shuffling enables search of less-probable regions of the search space.

Data were analyzed using a repeated measures analysis of variance (ANOVA) incorporating Greenhouse-Geiser correction. We report the original degrees of freedom along with the corrected p values.

Subjects produced the most words in Shuffle ($M = 18.88$, $SD = 4.14$), followed by Interactive ($M = 17.67$, $SD = 4.31$) and then Static ($M = 16.56$, $SD = 5.35$), with the main effect being significant, $F(2,50) = 11.28$, $p < .001$. Post-hoc t tests revealed that Shuffle was reliably better than Interactive ($t = 2.49$, $p = .02$), Interactive was better than Static ($t = 2.20$, $p = .038$), and Shuffle was robustly better than Static ($t = 4.90$, $p < .001$).

Moreover, a significant majority of participants performed better in the Shuffle condition than in the Static condition (92%, $p < .001$, binomial test), showing that the effect was nearly universal. We predicted that random reshuffling would raise the floor on performance — it would act as a safety net. This prediction was borne out with 85% of participants performing better on their worst Shuffle trial than on their worst Static trial ($p < .01$, binomial test), again showing how general the effect was.

Our second conjecture, however, was not confirmed: Notably, no more low-frequency words were found by moving tiles around than when they were static. It did not help subjects find harder words. Nor did movement help most at the beginning or at the end, when subjects might be assumed to have run out of ideas.

Discussion

In Scrabble, it is easier to evaluate ideas (decide if a letter sequence is a legal work made from the right letters) than to generate them. So any process that helps generate good candidates ought to improve performance. The same often applies to problem solving and creative thought more broadly.

Why would shuffling and interaction facilitate inventiveness? An obvious conjecture is that letter sequences created through movement contain high-value raw materials — useful trigrams like 'ing', 'pre' or 'ile'. It is obviously easier to find words in trigram or morpheme space than single-letter space. This may be true for interaction, where subjects move tiles into trigrams and bigrams as part of a strategy. But it cannot be right for shuffling. In shuffling, the overwhelming majority of letter sequences are jumbles. Rarely are they useful trigrams. The overwhelming majority of letter sequences are jumbles. A better view is that shuffling surfaces hints, letter sequences that stimulate word discovery. This is not an aleatoric process because actual candidates are not generated; they are only made easier to find.

There are a few reasons to think that shuffling is about hinting and that its success is generalizable. First, random generators lower the cost of hypothesis production. Pressing a shuffle or randomize button is fast, so fast that a button can be pressed many times in the hope that a good search region for closer exploration will eventually surface. Internal and ordinary interactive generation is slower, more complex and more effortful. People hoping to produce more ideas might want to find a hint-maker that relies on chance.

Second, by using randomness, an entire search space is probabilistically guaranteed to be covered. Randomized algorithms break patterns that we don't know we are following. Random generators produce candidates with larger variance than subjects can produce on their own because there are no biases or filters that limit candidates based on prior expectations of what is good. Every element and every region in a domain has equal probability of being chosen. As long as agents can decide quickly whether a candidate is interesting — once they see it — they can cash in on the variance, taking up good opportunities when they arise and ignoring the rest. If taken literally, this method might warrant seeing the success of shuffling as supporting a 'natural selection' account — random-generation selective retention. But because random candidates are rarely viable, the process cannot be natural selection: It is their parts, or even more indirectly, the associations induced in subjects by reflecting on those parts, that are valuable. It is more like random stimulation-hinting.

On closer examination, the hinting process is complicated. Given that most shuffling generates bad candidates and primarily low-value material, shuffling ought to be a bad source of hints. Mostly it is, as shown by how often the biggest users of shuffle immediately reject their shuffle. Several subjects, for example, seemingly bereft of ideas shuffled more than a hundred times in the third minute to discover only a few valid words. As might be expected, there is a happy medium. Some shuffling is helpful, but too much shuffling is not. Better players — those in the top third — shuffled about 50% less than those in the bottom third (means 48 vs. 94, $p < .04$). That is, the best players shuffled every 3.7 seconds vs. 1.9 seconds for the worst ones. This difference may suggest that better players take longer to judge the likely 'hint quality' of a shuffle, or it may reflect the extra time it takes to extract ideas from each shuffle. Regardless, all subjects shuffle so much that the cost in time and mental effort must be sufficiently low that it pays to keep fishing for hints.

The third and perhaps most interesting reason why random resetting facilitates performance better than other conditions is that it duplicates the effect of adding additional people to a team. Adding cognitive diversity to a team is known to facilitate creativity (Kurtzberg 2005). Each new person operates with a different cognitive outlook and method; hence, biases are partly washed out. Where this is not viable, adding chance may be the easiest way to diversify. In fact, incorporating chance may be even more facilitative: There are no group dynamics to hinder participants from suggesting truly wild ideas.

4. Is it really random?

Most creativity methods incorporating chance use generators that are less random than the quasi-random shuffle function in Scrabble. Suppose a person has a trick to help her crossword play. She watches a stock market ticker on TV, hoping the acronyms will trigger an association. Market tickers cycle; the same symbols come up at regular intervals. So the process is not random. But it is exogenous to the crossword domain. If this trick works, then maybe it is not randomness that is required for stimulation but independence. It may be enough if the process is driven by a function unconnected to crosswords. To the crossword player, the feeling of randomness is strong because whether ticker acronyms cycle or not, the acronym set has no greater relevance to one crossword question than another. There is no prior way of deciding when in the ticker stream an acronym will serve as a hint, if ever. But it is not randomness per se that is needed.

If we work with chance as something that is exogenous and unpredictable relative to a problem rather than requiring true randomness, then any process that is chaotic, strongly non-linear or uncorrelated with problem components may provide hints that, though determinate, are unpredictable to a subject with normal information and computational power. To harness these sorts of chance events to stimulate creativity, subjects must find a simple method of making the chance events relevant. For instance, ticker acronyms might provoke looking for solutions that incorporate ticker letters for the unknowns. This will cause a departure from one's normal problem-solving method. In the case of the Alternate Uses Task, the ticker letters might be used as a constraint: For example, name possible uses beginning with each letter. Using Apples' ticker symbol AAPL as a cue, a subject might guess that a newspaper could be used for A, attracting attention; A, advertising a used car; P, prodding cattle; or L, letters to paste on a ransom note. The simple act of constraining association to a single letter can be helpful.

I now turn to applications of this sort of randomness in art.

5. The use of randomness in choreography

As has been noted by others, "performance artists, painters, musicians, writers, and poets (particularly those from the early to mid-20th century) relied on their own indeterminate actions, free association, and accidental movements to generate (aspects of) randomness" (Leong et al. 2008: 16). In the arts, these random elements were often explicitly woven into the artistic performance, as for instance, John Cage famously used randomness to make musical sound, and Merce Cunningham used random placement and positioning during live performances of dance. These processes were called aleatoric, but a lot followed that was not random, since choosing a starting position and placement is just part of the composition.

The process I report here is based on a rather different version of this method as it plays out in the choreography of Wayne McGregor (henceforth WM), the resident choreographer of the Royal Ballet and director of Random Dance, an experimental company. In his work, chance elements serve very clearly as no more than stimulation or hints used in the early formation of ideas rather than in the performance phase.

For several years, I have been studying how WM creates a new dance (Kirsh et al. 2009; Kirsh 2013). Like Cunningham and Bill Forsythe before him, WM relies on giving dancers problems or tasks that involve using arbitrarily selected 'seeds', typically derived from materials found in the local environment as stimuli to generate movement. For example, dancers might be told to look around the dance studio and pick out six arbitrary[5] curves, surfaces or structures. They are then asked to choose six arbitrary parts of their body — for instance, their left hip, navel, right ear, left shoulder, mid-thigh and tailbone — and then sequentially 'sketch' their chosen curves and structures using a random assignment of body part to curve. A dancer might sketch the elliptical curve of a ceiling lamp using her tailbone, and then using her left shoulder she might sketch the joint of a folding chair, pick up the end of that chair's curve with her right ear and continue on to sketch the pedals of the piano with her navel before using her hip, and so on. The result is an odd mixture of smooth and irregular movements that may contain forms that rarely arise. No one expects the dancer to stop there. These movements themselves are ideas that dancers can build on to make more integrated, interesting phrases.

Sketching tasks are for beginners. WM has extended this method in many ways for his super-expert dancers. Here is a brief transcript derived from a video we took of an actual task he gave his dancers while making his 2009 piece, *Dyad 1909*. The group was working in a large modern studio with a 30-foot-high ceiling from which diverse electronic equipment was suspended. WM asked the dancers to lie down and after a few moments said:

> Look up. Pick boundaries in your field of vision to pay attention to. Describe the boundary in which to work in detail (while on the floor). Pick one detail or aspect within the boundary and describe it. Where is it most black? Assign that blackest thing to your body as a weight. Use that concept of weight as an anchor for what

5. This is not quite accurate. Dancers invariably have some prior idea of what counts as an interesting curve or structure, so their choice is not arbitrary. But of the many 'interesting' curves or structures present around them, we may assume their choice is arbitrary.

> you are exploring. Describe something that is outside your boundary. Move to a new location and repeat these steps. Do the task about five times total.

In dance, the term 'describe' often means to depict through body motion. In this case, WM wanted the dancers to characterize the detail or aspect using internal language or imagery before using that description as inspiration for movement. They spent the next 55 minutes on the floor, moving their bodies in original and often striking ways as they worked on the task. They came back to the task a few times in later sessions.

The output from the task, selected and refined by the dancers and then selected and modified by WM, goes through many more transformations. Two dancers, each with their own independent movements in pace and shape, might be told to work together to make a duet. Dancer B might be told to take part of A's movements and incorporate them somehow, or use them as counterpoint movements to help B modify his or her own. The seed ideas that the task generated are edited, cherry-picked by WM, added to, altered, combined, stretched, reworked and used as seeds for new random functions. At no time is the process mechanical or out of human intervention, especially by WM, who personally appropriates the movements himself to ensure consistency in their feel. Nonetheless, over the next three or four weeks, the dance remains open to novel influences while at the same time being refined and thematically shaped. In choreography, as in other areas of design, a germ idea is subject to hundreds of constraints and adaptive forces. Steve Jobs put it like this:

> … there is a tremendous amount of craftsmanship in between a great idea and a great product. And as you evolve that great idea, it changes and grows. It never comes out like it starts because you learn a lot more as you get into the subtleties of it and you also find that there are tremendous trade-offs that you have to make (quoted in Tideman 1995).

But without that great idea, there would be no great product.

6. Final discussion

In the search for methods to enhance creativity, artists, scientists, designers and businesses often rely on techniques that incorporate chance elements as a source of unpredictable stimulation. Ideation wants diversity. The power of chance lies in its departure from tradition. It releases us from predetermined ideas about the good that constrain our vision of the possible. Too often predictions about what will be a good idea are unreliable. The same holds for bad ideas. Jazz musicians regularly err on purpose, just to paint themselves into a position whereby to escape they must do something surprising — something that departs from the foreseeable (Barret 1998).

Robert Sutton (2001) described an anecdote he once heard:

> I got this idea from Karl Weick of the University of Michigan, who has described the ritual used by Naskapi Indians to determine where to hunt game. They placed the shoulder bone of a caribou over a fire until it cracked — then read the cracks as a map. Weick asserts that the ritual was effective because plans for future hunts were not shaped by the results of past hunts. It kept the Naskapis from mindlessly returning to — and depleting — territory they had covered before (Sutton 2001: 102).

Past success is the enemy. Familiarity is the enemy. Bayes is the enemy. In design companies a materials shed, sometimes called an ideas cart, is often relied on for stimulation during the ideation phase of design. A materials shed is a collection of artifacts like mechanisms, fabrics, patterns, iconic designs and toys or gadgets that a designer, in the early part of research, is encouraged to visit. Viewing the thousands of items filling the shed, someone in search of inspiration might respond to an odd metallic texture, an unusual gear configuration, a suction device or a mechanism exploiting a physical principle in an uncommon way. The aspect that triggers interest need not be an element that would ever work in the final design or a part of something central to the design. Its role is to provoke an idea that had not been considered before, an idea that helps the designer realign how he or she thinks or that arouses consideration of candidates outside the norm. This operates quite differently than another familiar strategy: using pictures, written descriptions and physical examples of *available solutions*, where the reverse effect often happens. Past ideas, or idea analogs, become fixated. The reviewed solutions bias thinkers to look for variants on existing solutions or to use existing search spaces, rather than invent new ones.

We found that chance in a Scrabble-like word discovery task facilitated performance better than carefully controlled thought alone. Like an ideas cart, it generates alternatives from left field. The reasons are that 1) hints are provided at low cost — the shuffle button can be hit quickly many times until a 'juicy' letter sequence emerges; 2) using chance to stimulate candidate generation increases the probability of searching more of the search space (hint-stimulated candidates are on average more different from each other than non-hint candidates); and 3) adding randomness to a search is like adding diverse team members to a panel. Diversity helps overcome bias. Empirically, this has been shown to lead to better decisions, presumably because more of the search space is considered.

To computationalists, none of this is surprising. Since the late 1970s, when Rabin (1976) first introduced chance into programming and computation theory, randomized algorithms have been appreciated for their speed and simplicity. For many applications, a randomized algorithm is the simplest algorithm available (e.g., Quicksort), or the fastest, finding a solution sooner than any deterministic algorithm; or the best in the sense that if a solution exists, a randomized algorithm has $1-\varepsilon$ probability of finding it, whereas no deterministic algorithm can offer such a guarantee (Motwani and Raghavan 1995).

Such findings support using evolutionary methods where candidates are directly generated. These methods may work in large populations where parallel processing of candidates can take place. But single humans try out ideas serially. They haven't time to evaluate the majority of candidates that would be generated randomly, so they must be more selective. Providing them with hints promises to strike the right balance between stimulating novelty and exploiting their prior knowledge — a method we find widely used in the arts.

The bottom line for creatives is that incorporating chance into one's creative method is not an admission of inadequacy; it is often theoretically the best that can be done. As the saying goes, if you can't walk in the right direction, sometimes it's enough to walk randomly. This holds true in Scrabble, in computer science, in evolution and, with a twist, likely in most areas of intellectual endeavor. Oddly, having access to randomness is one way we are better off than the angels: Because we are embedded in the material world, we can use physical things and chance to overcome our natural bias for preconceived order. Chance elevates the soul by grounding it in matter.

Acknowledgements

I have benefited from valuable discussions and the support of Wayne McGregor and his Studio, Sheldon Brown and the Arthur C. Clarke Center for Human Imagination, and from the help of Seana Coulson, Tyler Marghetis, the excellent commentary of an anonymous reviewer and by the computational skills of Joseph Stacy.

References

Barrett, F.J. 1998. "Coda—Creativity and improvisation in jazz and organizations: Implications for organizational learning". *Organization Science* 9(5):605–622. DOI: 1047-7039/98/0905/0605/$05.00

Bergdoll, B., Dickerman, L., Buchloh, B. and Doherty, B. 2009. *Bauhaus 1919–1933*. New York: The Museum of Modern Art.

Bunt, A., Lank, E. and Terry, M. 2009. "Understanding the role of computers in mathematical problem solving". Mathematical User-Interfaces Workshop 2009. Retrieved from http://

www.activemath.org/workshops/MathUI/09/proc/Bunt-Lank-Terry-Role-Computers-Math-Solving-MathUI09.pdf

Campbell, D.T. 1960. Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review* 67(6): 380–400. DOI: 10.1037/h0040373

Clark, A. 2006. "Soft selves and ecological control". In Ross, D., Spurrett, D., Kincaid, H. and Stephens, G.L. (eds), *Distributed Cognition and the Will: Individual Volition and Social Context*. Cambridge, MA: The MIT Press, 101–122.

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Cowley, S. and Vallée-Tourangeau, F. 2010. "Thinking in action". *AI & Society* 25(4): 469–475. DOI: 10.1007/s00146-010-0269-3

Duncker, K. 1945. *On Problem Solving*. Less, L.S. (trans.). Westport, CT: Greenwood Press.

Gilhooly K.J. and Murphy, P. 2005. "Differentiating insight from non-insight problems". *Thinking & Reasoning* 11(3):279–302. DOI: 10.1080/13546780442000187

Glucksberg, S. 1964. "Functional fixedness: Problem solution as a function of observing responses". *Psychonomic Science* 1(1–12): 117–118. DOI: 10.3758/BF03342819

Grice, P. 1975. "Logic and conversation". In Cole, P. and Morgan, J.L. (eds), *Syntax and Semantics. 3: Speech Acts*. New York: Academic Press, 41–58.

Guilford, J.P. 1967. *The Nature of Human Intelligence*. New York: McGraw-Hill.

Hollan, J., Hutchins, E. and Kirsh, D. 2000. "Distributed cognition: Toward a new foundation for human-computer interaction research". *ACM Transactions on Computer-Human Interaction (TOCHI)* 7(2): 174–196. DOI: 10.1145/353485.353487

Katona, G. 1940. *Organizing and Memorizing*. New York: Columbia University Press.

Kaufman, J.C. and Sternberg, R.J. (eds). 2010. *The Cambridge Handbook of Creativity*. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9780511763205

Keefe, S.P. 2006. "Rochlitz, (Johann) Friedrich". In Eisen, C. and Keefe, S.P., (eds), *The Cambridge Mozart Encyclopedia*. Cambridge: Cambridge University Press, 427–428. DOI: 10.1017/CBO9780511481383

Kirsh, D. 1990. "When is information explicitly represented?" *The Vancouver Studies in Cognitive Science*. Reissued by Oxford University Press, 1992, 340–365.

Kirsh, D. 2009. "Problem solving and situated cognition". In Robbins, P. and Aydede, M. (eds), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press, 264–306.

Kirsh, D., Muntanyola, D., Jao, J., Lew, A., and Sugihara, M. 2009. "Choreographic methods for creating novel, high quality dance". In Chen, L.-L., Feijs, L., Hessler, M., Kyffin, S., Liu, P.-L., Overbeeke, K. and Young, B. (eds), *Proceedings of Design and Semantics of Form and Movement DeSForM 2009 October 26–27, 2009, Taipei, Taiwan*. Amsterdam: Koninklijke Philips Electronics N.V., 188–195.

Kirsh, D. 2013. "Embodied cognition and the magical future of interaction design". *ACM Transactions on Computer-Human Interaction (TOCHI)* 20(1): 1–30. DOI: 10.1145/2442106.2442109

Kirsh, D. and Maglio, P.P. 1994. "On distinguishing epistemic from pragmatic actions". *Cognitive Science* 18(4): 513–549. DOI: 10.1207/s15516709cog1804_1

Knapp, J. 2014. "Google ventures: Your design team needs a war room. Here's how to set one up". Retrieved from http://www.fastcodesign.com/3028471/google-ventures-your-design-team-needs-a-war-room-heres-how-to-set-one-up

Knoblich, G., Ohlsson, S. and Raney, G.E. 2001. "An eye movement study of insight problem solving". *Memory & Cognition* 29(7): 1000–1009. DOI: 10.3758/BF03195762

Kohler, W. 1925. *The Mentality of Apes. Winter, E. (trans.)*. London: Kegan Paul.

Kurtzberg, T. 2005. "Feeling creative, being creative: An empirical study of diversity and creativity in teams". *Creativity Research Journal* 17(1): 51–65. DOI: 10.1207/s15326934crj1701_5

Leong, T., Howard, S. and Vetere, F. 2008. "Take a chance on me: Using randomness for the design of digital devices". *Interactions* 15(3): 16–19. DOI: 10.1145/1353782.1353787

Maglio, P.P., Matlock, T., Raphaely, D., Chernicky, B. and Kirsh, D. 1999. "Interactive skills in Scrabble". In Hahn, M. and Stoness, S.C. (eds), *Proceedings of the Twenty-First Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates, 326–330.

Martindale, C. 1999. "Biological bases of creativity". In Sternberg, R.J. (ed.), *Handbook of Creativity*. Cambridge: Cambridge University Press, 137–152.

Motwani, R. and Raghavan, P. 1996. "Randomized algorithms". *ACM Computing Surveys* 28(1): 31–37. DOI: 10.5860/CHOICE.33-5148 DOI: 10.1145/234313.234327

Neustein, D. 2010. "Frank Gehry's crumpled vision for Sydney". Retrieved from http://www.australiandesignreview.com/opinion/1842-frank-gehrys-crumpled-vision-for-sydney

Norman, D.A. 1993. *Things That Make Us Smart*. Reading, MA: Addison-Wesley.

Ohlsson, S. 1992. "Information-processing explanations of insight and related phenomena". In Keane, M.T. and Gilhooly, K.J. (eds), *Advances in the Psychology of Thinking*. London: Harvester-Wheatsheaf, 1–44.

Polya, G. 1957. *How to Solve It: A New Aspect of Mathematical Method*. Garden City, NY: Doubleday.

Rabin, M.O. 1976. "Probabilistic algorithms". In Traub, J.F. (ed.), *Algorithms and Complexity*. New York: Academic Press, 21–40.

Reis, D. 2013. "Let randomness kick-start your creativity". *Thinkergy Innovation Company*. Retrieved from http://thinkergy.com/let-randomness-kick-start-your-creativity-part-1/.

Rochlitz, F. (ed.). 1815. *Allgemeine Musikalische Zeitung* 17: 561–566.

Schiller, P.H. 1957. "Innate motor action as a basis of learning". In Schiller, C.H. (ed. and trans.), *Instinctive Behavior*. New York: International Universities Press, 264–287.

Schön, D.A. 1992. "Designing as reflective conversation with the materials of a design situation". *Research in Engineering Design* 3(3): 131–147. DOI: 10.1016/0950-7051(92)90020-G

Simonton, D.K. 2011. "Creativity and discovery as blind variation: Campbell's (1960) BVSR model after the half-century mark". *Review of General Psychology* 15(2): 158–174. DOI: 10.1037/a0022912

Sutton, R.I. 2001. "The weird rules of creativity". *Harvard Business Review* 79(8): 94–103.

Suwa, M. and Tversky, B. 1997. "What do architects and students perceive in their design sketches? A protocol analysis". *Design Studies* 18(4): 385–403. DOI: 10.1016/S0142-694X(97)00008-2

Tideman, K. 1995. "Steve Jobs on the difference between a great idea and a great product". Retrieved from http://www.kaspertidemann.com/steve-jobs-on-the-difference-between-a-great-idea-and-a-great-product

Vallée-Tourangeau, F., Weller, A. and Villejoubert, G. 2011. "Distributed cognition and insight problem solving". In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society, 273–278.

Vartanian, O. 2009. "Variable attention facilitates creative problem solving". *Psychology of Aesthetics, Creativity, and the Arts* 3(1): 57–59. DOI: 10.1037/a0014781

Vyzoviti, S. and Dimitriou, M. 2012. "The grasping hand as form generator: Generative modelling in physical and digital media". In Liveneau, P. and Marin, P. (eds), *MC2012 Symposium: Materiality in Its Contemporary Forms: Architectural Perception, Fabrication and Conception*. Villefontaine: Les Grands Ateliers de l'Isle D'Abeau, 91–99.

Weller, A., Villejoubert, G. and Vallée-Tourangeau, F. 2011. "Interactive insight problem solving". *Thinking & Reasoning* 17(4): 424–439. DOI: 10.1080/13546783.2011.629081

Wertheimer, M. 1961. *Productive Thinking* (rev. ed.). London: Tavistock.

Author's address

University of California, San Diego
Cognitive Science 0515
La Jolla, CA 92093–0515
USA

About the author

David Kirsh is Professor and past Chair of the Department of Cognitive Science at the University of California at San Diego. He was educated at Oxford University (DPhil), did postdoctoral research at MIT in the Artificial Intelligence Lab and has held research or visiting professor positions at MIT and Stanford University. Some of his recent projects focus on ways humans use their bodies as things to think with, specifically in dance making and choreographic cognition.