A Quantitative Study of Creative Leaps

Lior Noy*, Yuval Hart*, Natalie Andrew*, Omer Ramote, Avi Mayo and Uri Alon
Molecular Cell Biology
Weizmann Institute of Science
Rehovot, Isreal
lior.noy@weizmann.ac.il

Abstract

We present a novel quantitative approach for studying creative leaps. Participants explored the space of shapes composed of ten adjacent squares, searching for 'interesting and beautiful' shapes. By recording players' actions we were able to quantitatively study aspects of their exploration process. In particular our goal is to identify populated sub-regions in the shape space and study the dynamics of 'creative leaps': a jump from one such area to another. We present here the experimental system, our methods of analysis and some preliminary results. We show that the network of shapes created by human participants is different from the class of networks created by applying a simple random-walk algorithm. Chosen shapes show an interesting negative correlation between their abundance and the probability to be chosen as beautiful. We further analyzed the human network unique signature using its *network motifs* profile. Intriguingly, this signature shows similarity to words-adjacency networks extracted from texts. Lastly, we find preliminary evidence that human players exhibit two types of exploration: 'scavenging', where shapes similar in their visual-iconic meaning are quickly accumulated, and 'creative leaps', where players shift to a new region in the shape space after a prolonged search. We plan to build upon this result to quantitatively study creative processes in general and creative leaps in particular.

Introduction

In his book "the Act of Creation" the author Arthur Koestler describes the similarities between three types of creative acts: the pun of the joker, the discovery of the scientist and the lyric expression of the poet (Koestler 1964). The crux of the creative act is the creative leap, the momentary intersection of two different matrices of association (Fig. 1, left). Consider a search resulting in a creative solution for a given problem. Before the creative leap the search is confined to some familiar sub-space (the horizontal plane in Fig. 1, left). Using chance or intuition the solver has managed somehow to reach a point on the plane which also belongs to another plane, a totally different class of solutions (the vertical plane in Fig. 1, left). The creative leap is the ability to recognize this transition point and to jump from one class of solutions to another.

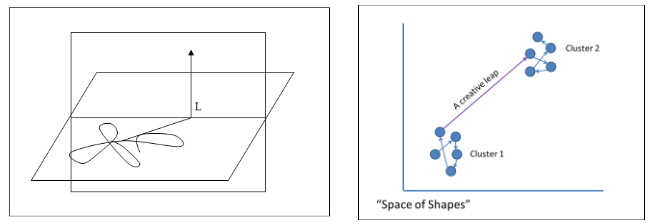

Figure 1. A Symbolic representation of creative leaps. Left: according to Koestler the heart of any creative act is a creative leap between two intersecting domains. Right: a hypothetical creative space. Solutions are grouped into two clusters. Searching within a cluster requires short moves and creates similar solutions. In order to move to a different cluster of solutions the agent needs to perform a creative leap.

Little is known about the dynamics of creative leaps. Previous work has described creative leaps of exceptional creators (Miller 1996) while empirical work has focused mainly on moments of insight in problem solving, such as the Remote Association Test, using both behavioral (Dominowski and Dallob 1995) and brain studies (Sandkühler 2008). It is difficult to capture creative leaps in a laboratory setting. Moreover, many solution spaces might be high-dimensional and complex, with no clear metric defining the similarity between points. For example, consider the space of all answers to the following question used in a group creativity test: "how can the number of tourists visiting your city be increased" (Nijstad and Stroebe 2006). While this problem has solutions that belong to different classes (for example 'increase advertisement' vs. 'improve infrastructure') it is not clear how to define and construct the space of all such ideas.

Our goal is to study a creative task with an underlying solution space that is (a) simple and well defined to enable a quantitative investigation of the search dynamic (b) that contains clusters of solutions, with the possibility of performing creative leaps between them (see Fig. 1, right). Our approach resembles recent work by Jennings that similarly studied people's search trajectories in a visual domain (Jennings 2010; Jennings et al. 2011).

We searched for a parameterized space that will be complex enough to allow for possible creative leaps, but not too complex to allow a computational description of human search in this space. We suggest using the set of all Nsize *polyominoes* – the set of two dimensional shapes composed of N adjacent squares (Golomb 1994).

Besides its well defined structure which allows for establishing a metric on the search space, the polyominoes space provides a crucial advantage: the shape space exploration complexity is tunable by changing the parameter N. We can thus aim to have an exploration process which is on one hand not too trivial and on the other hand not too complex to quantify. In that we hope to capture the gist of what Boden describes as 'an exploratory frame of mind' (Boden, 2004). Since this exploration process resembles a creative process undertaken by, say, a graphic designer designing a new icon in a limited space, we hope to gain insights in the growing field of computational models for design processes (Gero, 2000).

We analyzed the network of shapes and moves created by human participants and compared the human exploration with a simple random-walk algorithm that transverses the network of shapes discovered by the human participants. This comparison shows that the human search behavior is not simply the results of a random travel between the shapes. Our results suggest that humans perform two types of searches: 'scavenging', a simple search in an area of shapes, which can be explained by an algorithmic search, and 'insight' moves, or leaps, that cannot be explained by simple algorithm. The first type of moves corresponds to the within cluster exploration in Fig. 1, while the second type contains, we hope, the creative leaps.

We next describe our experimental setup, the methods of analysis we employed and some initial findings supporting the notion that creative leaps can be quantitatively studied using the suggested approach.

Experimental Setup

System

We developed a system to experimentally test human trajectories in the shape space of polyominoes. We are currently experimenting with *decominoes,* 10-size polyominoes (consisting of 4655 unique shapes and 36,446 shapes if rotations and mirror images are counted).

We tested several variants of the creative task and report here results from the 'journey in shape space': exploring the space by moving one square at a time, transforming one legitimate shape to another. The starting point shape is always the horizontal line. We ask people to "explore the space of 'shifting shapes' and to discover shapes that you find interesting and beautiful".

We developed an experimental setup using Processing, an open source, cross-platform, programming language used for visualization (see Fig. 2).

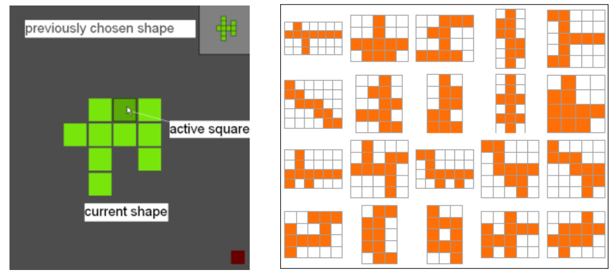

Figure 2. Exploring the space of shapes. Left: a screen shot of the 'Shape Shifter' game. At each step players move one square to create a new polyomino. Shapes can be stored in the 'shape gallery' by pressing the gray rectangle at the top-right corner. Right: examples of different shapes created by human players.

Procedure

123 participants (58 females and 65 males, ages 12-75 years, mean = 34.3), recruited through emails and social networks, were invited to participate in a short experiment in creativity. At any point players could store the current shape to a 'shape gallery'. The players moved freely between shapes, within a time limit of 25 minutes (no participant reached this limit). When choosing to finish the exploration they continued to the 'rating stage'. In this last stage players observed the 'shapes gallery' and were asked to choose 'the five most creative shapes you discovered'. We recorded square moves between shapes and their timing, as well as each player chosen gallery shapes and the final five shapes.

Analysis

A random-walk algorithm over the entire shape network

We used a network representation (a graph) of the shape space in the following way. Each shape is a node in the graph. Shape A and B are connected by an edge if shape A can be reached from shape B by moving a single square in a valid way. This structure is a directed graph representing all possible valid moves

The algorithm explores the network by first randomly removing one square from the current shape. The next decomino in the path is then generated by placing the 10th square in a new random location (self-loops are not excluded). This extends the path by one step. The path is further extended by repeating these steps up to a predetermined steps number.

This algorithm was used to establish both the entire shape space network and a random walk generated network to compare with the human generated network of travelled shapes. For the entire shape space the algorithm was run until all possible 36,446 decominoes were generated (with mean path length of 150,000 steps). For comparison with the human network, the algorithm was run 123 times (the number of human participants) with a number of steps which is sampled from the number of steps distribution of the human players.

A random-walk algorithm over the human generated network

In order to create computer generated networks which are more closely related to the human networks we restricted the algorithm to travel only on edges which were travelled by at least one human player.

First the human generated decominoes network is generated and the allowed steps are listed. Although the network is naturally directed, the computerized walker is allowed to move on the undirected network (that is, the computer can also move backward on any human edge).

The algorithm is seeded and a new shape is chosen randomly from the set of shapes which are connected by allowed edges. The length of the path is sampled from the distribution of lengths of paths traversed by the human players. This process is repeated 123 times.

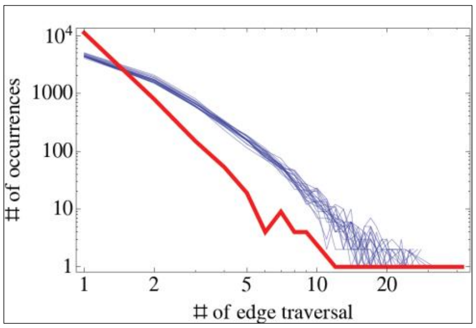

Figure 3. Comparing human and computational exploration networks. The number of occurrences of edges where edges are grouped by the number of times they were traversed. Shown are the values for human players' network (red) and the random walk network restricted to the human network shapes (mean of 10 simulation in dark blue, each specific simulation in light blue).

Our current goal is to compare the features of the human generated network to a network generated by a random-walk algorithm and to study if there is a noticeable difference between the two, in order to show that the human behavior cannot be explained as a result of a random-walk in the shape-space.

Triad Significance Profile Calculation

The 13 network motif frequencies of the human and random generated networks were calculated. The normalized Z score of each of the 13 possible triads was then calculated. Z score is computed by the difference of the triad frequency to the mean frequency of the same triad in a computerized agents' network, measured in STD units. Frequency mean and STD were calculated from 10 simulations of the computational networks.

Results

Human and Computational Networks

We first asked whether the exploration network created by human players is different from the network created by a random-walk algorithm traveling the entire shape networks. We find that the exploration network created by human players is much more compact. Furthermore, the players' network obeys a power-law distribution of node degree frequencies (how many edges go in or out from a specific node), while the computational algorithm produces a Gaussian-like distribution of node degree frequencies. In addition, human exploration on the network of all allowed edges is very constrained and compact relative to a random exploration process of the whole shapes space.

We next asked whether the type of exploration players perform is dictated only by some constraint on shapes available to people's perception. We thus compared the human exploration network with an ensemble of networks created by allowing a random-walk algorithm to choose shapes randomly, but restricting it to shapes that were selected by the human players. We find that the algorithm travels much less than the human players and so create a much smaller network than the players' network. Furthermore, the properties of the computational exploration networks, such as the distribution of nodes degrees is markedly different from the human exploration network (Fig. 3).

Consensus in Participants' Choices

A possible concern regarding our creative task is whether there is some consensus among different participants regarding their aesthetic choices. While we do not expect to have total agreement – for example some players preferred iconic shapes, while other preferred more abstract ones, a total lack of consensus could raise doubts on the validity of this task to measure human creativity.

To assess the consensus in participants' choices we plotted the *selection ratio*, the percentage of times a shape was chosen (number of times chosen divided by number of times traversed) against the number of times a shape was traversed (Fig. 4). We differentiated between shapes ranked as interesting shapes in the last stage of the game (in blue) and those that were only chosen to the gallery (in red).

We note that there is a large number of shapes with high (>50%) selection ratio, with few shapes exhibiting selection ratio of more than 90%. At least for these shapes there seems to be a consensus among the different human participants. In addition, shapes that were ranked in the last stage had a statistically significant higher selection ratio (ranked: centered around (23.34, 50) with STD (19.41, 20); not-ranked: centered around (15.6, 20) with STD (6.7, 13); non-paired t-test, $p<10^{-7}$).

We also note the negative correlation (Pearson correlation = -0.25, $p<0.05$) between the prevalence of a shape (how many times it was traversed) and its selection ratio. Intriguingly, this might suggest that shapes 'less traveled by' are appreciated more by the people who have reached them.

A Network Motifs Signature

In order to further characterize the human exploration network we measured its network motifs signature, termed triad significance profile (TSP). This network signature is calculated by taking the frequencies of all three node subgroups of a network and normalizing each frequency by the triad frequency in a network created by a similar random process (Milo 2002). In our case, we compared triad frequencies of human network with triad frequencies created by the random walk algorithm on the human network (see Analysis). Previous studies in our lab showed that networks with similar structure and function have a similar TSP signature. Thus, this method offers another quantitative classification to networks.

This preliminary calculation (Fig. 5) indicates that the network motifs significance profile shares a similar frequency signature of text networks (Milo 2004), suggesting that the human visual exploration process in shape space consists of visual rules similar to those of language networks, having categories of words with a certain formulated way of combining between different categories. Future work should check the dependency of the calculated triad significance profile on the randomization process used to create the base-line random network.

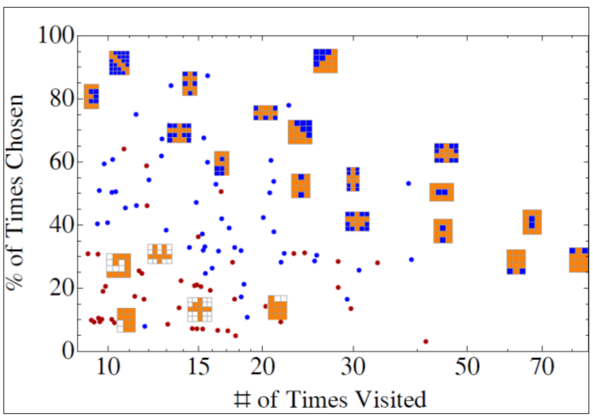

Figure 4. Consensus in participants' choice of shapes. Y-axis: the number of times a shape was admitted into the gallery out of the number of times it was visited. X-axis: the number of times a shape was visited. Only shapes that were visited at least 10 times are presented. Dots in blue represent shapes that were also ranked in the final stage while red dots represent shapes that were chosen to the gallery but were not ranked. Correspondingly, shapes shaded in blue are representative of the set of finally chosen shapes.

Initial Evidence for Creative Leaps

In order to more closely examine the exploration process of individual players, we focused on the 'chosen to the gallery' shapes (Fig. 6), enumerating both the number of steps between two sequential shapes (the number above each shape) and the time interval between selections of the two shapes (the y axis). For several players we observe an interesting pattern: the time and number of steps between two sequential chosen shapes is declining at the beginning, usually creating similar content shapes. Then, a long traversal exploration process is commenced, usually leading to shapes belonging to a new cluster of similar shapes. As exemplified in Fig. 6, the player moves from "Animals" shapes to "Space invaders" shape to "Symbolic male/female" shapes. One can interpret this saw-tooth pattern as consisting of scavenger explorations connected by a creative leap, which serves to reach a new iconographic domain.

We hope to utilize these processes to cluster the shapes automatically into different domains and thus create a semi-metric on the shape space. Another utility to aid building the metric comes from the use of the rating process at the end of the game. Subjects are requested to choose the five most creative shapes. Our assumption is that subjects will choose shapes that they see as most distinct from one another, thus providing another metric measure on the shape space.

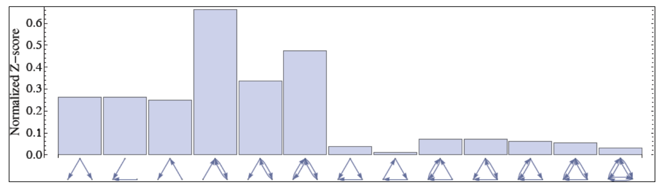

Figure 5. The triad significance profile (TSP) of the human players' network suggests a similarity to word-adjacency networks of texts. The main feature of the TSP is the under-representation of triangle-shaped triads 7 to 13.

Conclusions and Future Work

We presented a novel quantitative approach for studying creative leaps. Our goal is to study a creative task using computational tools. Specifically we aim to define the space of products of the creative task, to detect clusters of similar products and to study creative leaps between them.

Working toward this goal we developed a web-based game in which players explored a visual space composed of 10-size polyominoes, while searching for interesting and beautiful shapes. As a first step we tested whether human behavior in this task can be explained as a result of a random-walk algorithm. We therefore compared the exploration network created by human players to two computational exploration networks. The first network was created by random walks on all possible shapes, and the second one was created by random walks restricted to shapes chosen by human players. We compared general properties of these networks, such as in/out degree, and found a significant difference between the human and the computational networks. Compared to a network made by a random walk on shapes chosen by players, the computer's random walk is much smaller, suggesting that the trajectories of human exploration contain also segments of directed motion toward interesting regions of the space. Following the *fogginess* metaphor of Jennings (2011) these segments might correspond to the areas of the landscape with have 'good visibility'.

We also used the concept of network motifs to characterize the human search network. We identified which of a known super-families of networks (e.g. social, transcription networks, and language originated), matches the human exploration network. We find that the human network is similar to language-originated network, and are planning to further study the connection between these two networks.

We further find preliminary evidence of players' paradigm shift while playing the game. Players show periods of 'scavenging', where they exploit shapes similar in iconic meaning (e.g. animals, letter, symmetric shapes) accompanied by long walks on the grid of possible shapes, which leads to a different region in the shapes space. The 'saw-tooth' pattern we have found in the time between chosen shapes (Fig. 6) might be the first clue for the existence of clusters in our shape space. We plan to corroborate these finding by different methods that can be used to detect clusters of shapes in this visual domain. In particular, we plan to use the human choices embedded in our task at multiple levels (which shape to move to; which shapes to insert to the gallery; which shapes to choose in the final stage) as a different probe into the structure of the shape space.

This paper presents work-in-progress aiming to develop a computational platform for studying human search in creative tasks, and in particular to study creative leaps. We are currently performing a large-scale human experiment with this platform and plan to apply a host of quantitative methods to further test the preliminary results presented here. Using these methods we hope to be able to measure and study the dynamics of creative leaps.

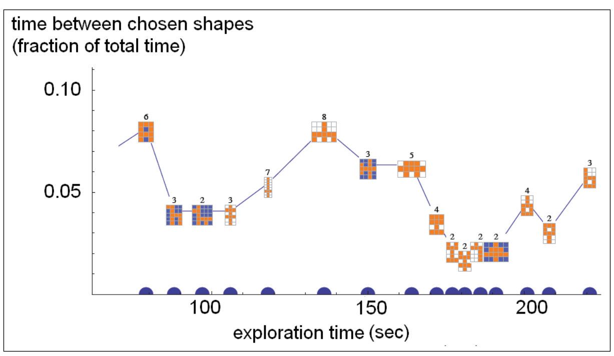

Fig 6. Preliminary evidence for clusters in the shape space. Looking at the time differences between chose shapes we often see 'saw-tooth' patterns. Humans seem to reach a fruitful region, 'scavenge' it, that is, to quickly pick a few similar shapes, and then to move to another region, a move that takes much more time. Notice for example the two clusters of similar shapes around 100 and 180 seconds. Only chosen shapes are shown, and shapes in the 'top five' (chosen between all gallery shapes) appear with a blue background. The number above each shape is the number of moves from the previous shape.

References

Boden, M.A. 2007. Creativity in a nutshell. *Think* 5 (15): 83.

Dominowski, R.L. and Dallob, P. 1995. Insight and problem solving. In: Sternberg, R. J.; Davidson J. E. (Eyds). The nature of insight, pp. 33–62. MIT Press, Cambridge, MA.

Gero, J.S. 2000. Computational models of innovative and creative design processes. *Technological forecasting and social change* 64 (2): 183.

Golomb, S. W. 1994. *Polyominoes* (2nd ed.). Princeton University Press, Princeton, NJ.

Jennings, K.E. 2010.Search strategies and the creative process, Proceedings First International Conference on Computational Creativity, 130

Jennings, K.E.; Simonton, D.K. and Palmer, S.E. 2011. Understanding exploratory creativity in a visual domain, *Proceedings of the 8th ACM conference on Creativity and cognition*, 223.

Koestler, A. 1964. *The Act of Creation*. Penguin Books, NY.

Miller, A. I. 2000. Insights of genius: Imagery and creativity in science and art. MIT Press, Cambridge, MA.

Milo, R.; Shen-Orr, S.; Itzkovitz, S.; Kashtan, N.; Chklovskii, D.; and Alon, U. 2002. Network motifs: simple building blocks of complex networks. *Science* 298 (5594): 824.

Milo, R.; Itzkovitz, S.; Kashtan, N.; Levitt, R.; Shen-Orr, S.; Ayzenshtat, I.; Sheffer, M.; and Alon, U. 2004. Superfamilies of evolved and designed networks. *Science* 303 (5663): 1538.

Sandkühler, S. and Bhattacharya, J. 2008. Deconstructing insight: EEG correlates of insightful problem solving. *PLoS One* 3 (1): e1459.