

COMPUTER IMAGINATION USING PATTERN RECOGNIZERS WITH RECURRENT ACTIVITY CONTROL

TAPIO TAKALA
Helsinki University of Technology
02150 Espoo, Finland

Abstract. The paper discusses free creative imagination. Though it often seems out of control, it usually follows association chains, where previously learned patterns recurrently appear, but always in a new context. A computational model producing this kind of activity on a visual screen is developed. It is based on a set of simultaneous pattern recognizers, working on an array of cells corresponding to the retina. When a recognizer finds partial fitness, it will feed its preferred image back, boosting itself to full activity. It gives the original image an "interpretation", like is done by humans when recognizing an emergent pattern. A higher level searchlight controller allows only one recognizer to dominate at a time. After a while the dominating pattern will saturate and give way to other pending interpretations. If there are multiple possible interpretations (recognizers responding) to a constant image input to the retina, the controller makes them to alternate in random order. Without an input, a chain of images is produced by free associations, when a fading image activates another recognizer. This continuous reinterpretation of emerging situations is typical to creative thought processes. A preliminary implementation using a variation of layered neural networks has created promising results. Further experiments with more sophisticated pattern recognizers are going on.

1. Background

It is commonly accepted that creativity is the ability to produce something which is recognizable as new and useful. Most often this is taken as the ability to solve problems: given a setting of constraints and (functional) requirements, find a solution (object) that satisfies them and perhaps is even optimal in some sense. This approach emphasizes the usefulness of creativity, but neglects other important properties like playfulness and ability to enjoy the results. Relaxing the pragmatic aspect, free imagination becomes an interesting issue – not as much *what* is imagined but *how* it proceeds.

Here I concentrate on the dynamics of imagination or other creative action: how does a mind get interested in things, recognize them, make

expectations on what is recognized, actually produce mental images, and eventually get tired of the creation and start looking for other interests? What are the mechanisms behind the fluctuations of interest of a creative mind?

For the first Heron Island conference I made an overview of creativity and presented a theoretical model based on neuropsychology (Takala, 1993), which partly has its roots in design theory (Takala, 1987a, 1987b). The present paper is an attempt to demonstrate that model with an implemented case of computer imagination, a mechanism continuously producing associative sequences of simple visual images.

This work, as well as my previous one, has strong coherence with Mitchell's (1993) paper on structured recognition of emergent shapes. As he points out, "design is not description of what *is*, it is exploration of what *might be*", and "the meaning of a drawing is not adequately captured by imposing *one* structure on it". Reinterpretation of a drawing means imposing other structures on it, which are possible (supported by the drawing) but may be very different from the first one and from each other. In this sense "creativity is the art of misunderstanding (a drawing) in a meaningful way" (Takala, 1993). The sense of humor is largely based on such reinterpretations. A lot of delight follows the finding of a new interpretation that perfectly fits the given information (e.g. a joke), but is against expectations, in some aspect opposite to the most obvious interpretation.

I think that implementation of such mechanisms in computers is the key to make them actively creative, instead of just straight-forward acting (but not so amusing) problem-solvers.

In the experiment that follows, the computer explores an input image by continuously giving it new interpretations, modifying it according to the current interpretation, and making the result subject to reinterpretation. Shortly, the action of computer in the experiment goes like follows: A small portion of an input image's features at a time are compared with previously learned potential interpretations. When enough correlation appears, that interpretation is brought into attention, and it becomes the expectation of the viewer. This means that those features supporting the interpretation are boosted and others suppressed or neglected, and even new features are created to support the expected figure. This positive feedback makes the current interpretation even stronger, until it perfectly fits the produced image¹.

However, a single interpretation would be too boring. Instead, there are multiple interpretation processes going on in parallel. They compete for the attention of the viewer (the consciousness) and for taking over the expectations and control of input image. When one interpretation has been

¹ *Sic*, what you see is not necessarily what you get, but what you expect to see.

on for a while, it becomes saturated and gives way to the next strongest one. This results in a continuous chain of interpretations, each influenced by the previous one in chain.

2. A neural network system for imagination

Analogously to a design theory describing the design action as a loop where design synthesis and evaluation iterate between requirements and product proposal (figure 1a), the system described here consists of pattern synthesis (PS) and pattern recognition (PR) modules working on an image and its interpretation (figure 1b). In design the requirements are given and a product is to be developed, whereas in image understanding the image is given and an interpretation is to be formed. Creative imagination is both combined: starting from any image, an interpretation is formed, then the image is modified according to the interpretation, a new interpretation is developed, etc.

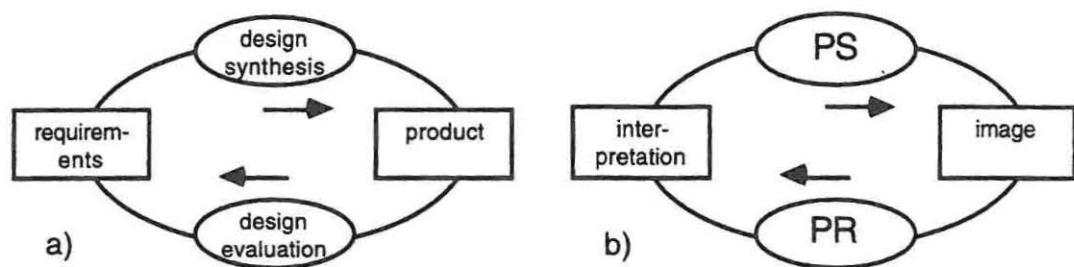


Figure 1. Design (a) and imagination (b) processes.

Many pattern synthesis methods producing very appealing images have been developed in computer graphics. Such are for example fractals produced by iterative mathematical functions (Mandelbrot, 1977), plant models produced by symbolic grammar-like L-systems (Prusinkiewicz and Lindenmayer, 1990) and textures created by band-limited random noise (Perlin, 1985) or reaction-diffusion models (Turk, 1991, Witkin and Kass, 1991). However, none of these devices can appreciate its own products. A human evaluator is needed to select those products worth of publishing. The importance of this fact is emphasized in the work by Sims (1991), where images are generated by elementary mathematical functions combined by genetic programming. The high artistic quality of these images is fully due to his taste as a recognizer and selector during the evolution.

Most pattern synthesis methods do not provide an inverse, a pattern recognition method working with corresponding principles. Syntactic methods (e.g. L-systems) could do that, but they only work on a highly

symbolic level – another method is required to extract the basic syntactic elements from an image.

Pattern recognition with artificial neural networks (ANN) works directly on the pixel level of an image. An approach similar to the present is that by Liu (1994). Based on a model of visual perception (Palmer et al., 1993), he used a two-level architecture of ANNs to recognize emergent subshapes of an image. Unfortunately, not much has been done with ANNs for image synthesis. The present work attempts to go in that direction.

2.1. OVERVIEW OF THE SYSTEM

The system described here consists of layers of neural cells connected to each other across the levels. The overall organization is depicted in figure 2. At the bottom is a layer corresponding to the retina of the eye, where sensory input is fed in. At the next layer, small features like line segments or corners are recognized. A third layer represents higher level pattern recognizers, the activity of which is scheduled by yet another unit.

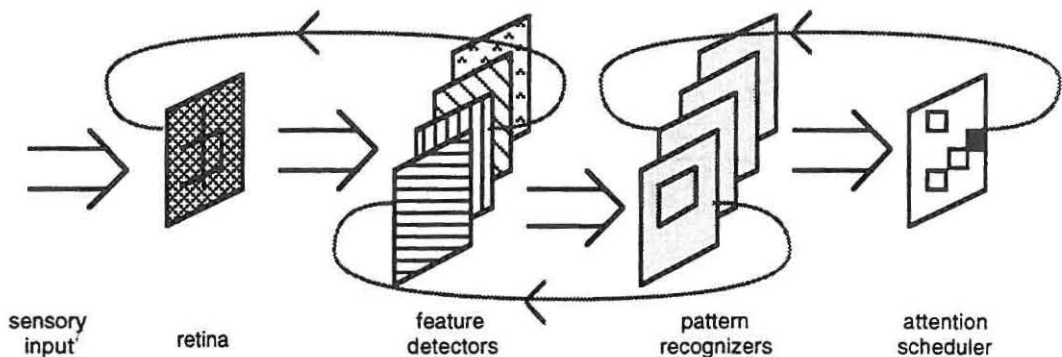


Figure 2. The neural network with recurrent feedback between layers.

Though similar in overall architecture, the implementation does not use any standard type of ANN (perceptron, backpropagation, etc.). Like in most ANNs, a neural cell here combines synaptic input from a number of cells on the previous layer, and provides a single output (figure 3). The main difference is that for all forward synaptic connections there are corresponding backward connections also, providing positive feedback and that way reverberation of cell activity.

An active cell will boost its input cells towards the values that keep the cell itself active. Such recurrent processing corresponds to the idea of self-activating neural loops in the brain (Hebb, 1958, Takala, 1993). Liu (1994) used this "recurrent attention" only as one step in the process, to intensify a partially recognized subshape in order to make its detection clearer. Here

the feedback is applied across all subsequent layers of cells, and becomes an essential mechanism for attention-directed image synthesis.




Figure 3. A neuron cell (a), and a typical laterally organized arrangement of cells (b).

2.2. THE IMAGINARY RETINA

The first layer in our system is a two-dimensional array of cells, where the imagined shapes actually form. Since the input image is directly fed in to this layer, it will be called the retina². These cells take as input only the sensor values (input image) and feedback from the next upper level. If both input signals are weak, random noise is added to the cell output value.

2.3. FEATURE DETECTORS

The next layer consists of a number of parallel arrays, each performing a simple feature detection task. The features used in this experiment are patterns within a small 3x3 area, like short vertical and horizontal line segments or elementary corners (figure 4). For each cell of a detector's array, the template pattern is positioned at corresponding place on the retina, and its correlation with the cell values there is calculated. Thus each feature detector performs digital filtering using the feature pattern as filter kernel.




Figure 4. The 3 x 3 feature detector patterns.

² It is doubtful if imagined shapes really exist as activity of neurons at the eye. However, measurable sound has been observed to come out of the ear due to activity of efferent nerves.

Each detector cell will feed the corresponding feature pattern, multiplied by the cell's output activity, back to the retinal cells. Unless disturbed by other input or feedback from other detectors, a recognized feature subimage will thus remain stable.

2.4. PATTERN RECOGNIZERS

The task of the higher levels is to recognize more complicated patterns formed by the features detected at the previous level. There are a number of such recognizers working in parallel. Each of them calculates the fitness of its template pattern against the features detected below, and will boost the template back if the fitness is high. Thus a once well-recognized pattern will stabilize itself - just like the feature detectors alone would do.

In this experiment the purpose was to be able to recognize arbitrary rectangles on the retina. Due to the huge number of all possible rectangles, only a few recognizers were implemented, but their template patterns are assigned dynamically. Though ideally a recognizer should be built out of similar neural cells as the other layers consist of, a more heuristic approach was taken here.

The only type of pattern to be recognized is a rectangle. It can be described with a template with slots for the rectangle's left, right, bottom and top coordinates. A rectangle's recognition is based on counting the number of vertical features at each column and horizontal features at each row of the retina (i.e. calculating the marginal densities of the two feature detector cell arrays). From these, a set of largest values are selected, whose column and row numbers serve as potential x and y coordinates. From those the (L,R,B,T) parameters of the rectangle are selected at random, but checking that no two recognizers have the same template.

The rectangle's fitness is the relative number of features matching with its boundary, i.e. the horizontal and vertical features used above, plus the four corner features.

An active recognizer will feed back its pattern by boosting all its boundary features and black (empty) features elsewhere. Through the competition and feedback of the feature detectors, this information will gradually appear on the retina also.

2.5. ATTENTION CONTROL

The positive feedback from level to level down would alone cause a stable pattern (a single rectangle) to appear on the retina. In biological neural networks, however, such would not happen due to adaptation of cells. A constant stimulus will saturate the cells and its effect will fade out after a

while, only changes are taken into account. Total sensory deprivation causes the brain to generate endogenous stimuli, which are perceived as imagination or hallucination.

Besides adaptation, there seems to be another control mechanism in brain, called searchlight (Crick, 1984). It takes care that attention is paid to one focus at a time, and for a short time only.

In this experiment adaptation and searchlight are simulated on the pattern recognizer level only. While all recognizers work in parallel upwards and calculate continuously their output values from the feature level activities, an attention scheduler selects only one of them at a time for feedback down to the feature detectors. Let us assume the scheduler compares the activities of the recognizers and selects one with a high value. From that on the cumulated activity of the selected recognizer is calculated and serves as a measure of saturation. When this measure becomes high enough, attention is released and another recognizer selected. The new recognizer will probably have some similarity in its pattern, in order to have high fitness value and become selected.

3. Results

Figure 5 shows what the first two layers alone can do. If only the empty and full patterns (figure 4, a and h, respectively) are used, they tend to form larger homogeneous black and white areas. At the borders the two features compete, and gradually either one will dominate and fill the whole image. Meanwhile, however, very interesting patterns form (5a) resembling the reaction-diffusion patterns used to simulate e.g. animal fur colorings (Turk, 1991). The other feature patterns (figure 4, b-g) contain boundaries between black and white and can be assembled to fit exactly together like dominoes. Such formation is shown in (5b). If an input image (5c) is added, it will act as a constraint, to which the other features attach in a process similar to annealing (5d).

With higher layers in the system much more dynamic behavior appears. Although only one interpretation of the image appears at a time, the dominating interpretation will continuously change due to the searchlight mechanism. Figures 6a-6t show a sampling of consecutive situations (further apart than consecutive time steps of the simulation). At each snapshot, the retinal image is shown to the left and its current interpretation (the pattern selected to attention) to the right. Between them is a small indicator showing the current fitness between the image and the interpretation.

The fluctuation of interest and the gradual formation of patterns can be clearly seen. For example, after an interpretation is selected in (6f), its large rectangle shape becomes boosted and clearer on the retina, and the indicator

shows better fitness (6g). Usually consecutive interpretations are similar in the sense that they share at least some common features (corner or partial side). Interestingly, this similarity may jump over a stage - for example, (6k) resembles more (6i) than (6j). This is because the common features (vertical line to the very left) did not quite fade during the intermediate stage (6j).

The constant input (5c) kept on in this experiment guides all interpretations, though occasionally something totally unexpected may appear with low fitness (6d). Among others there are some "natural" emergent interpretations of the input (the squares in 6s and 6t) showing up in the sequence.




Figure 5. Patterns produced by feature level feedback alone: (a) non-stable diffuse pattern formed by competition of full and empty features alone, (b) stable featured pattern developed from random input, (c) the input image used in subsequent experiments, (d) stable featured pattern formed around the input image. (Note: the resolution of (a) is 300x300 cells, others have 30x30 cells on the retina)

Without any input the system will rather randomly wander from one rectangle to another, following associations formed by their common features.




Figure 6. a-j: A sequence of responses to a constant input (each pair has the retinal image to the left and its interpretation to the right)

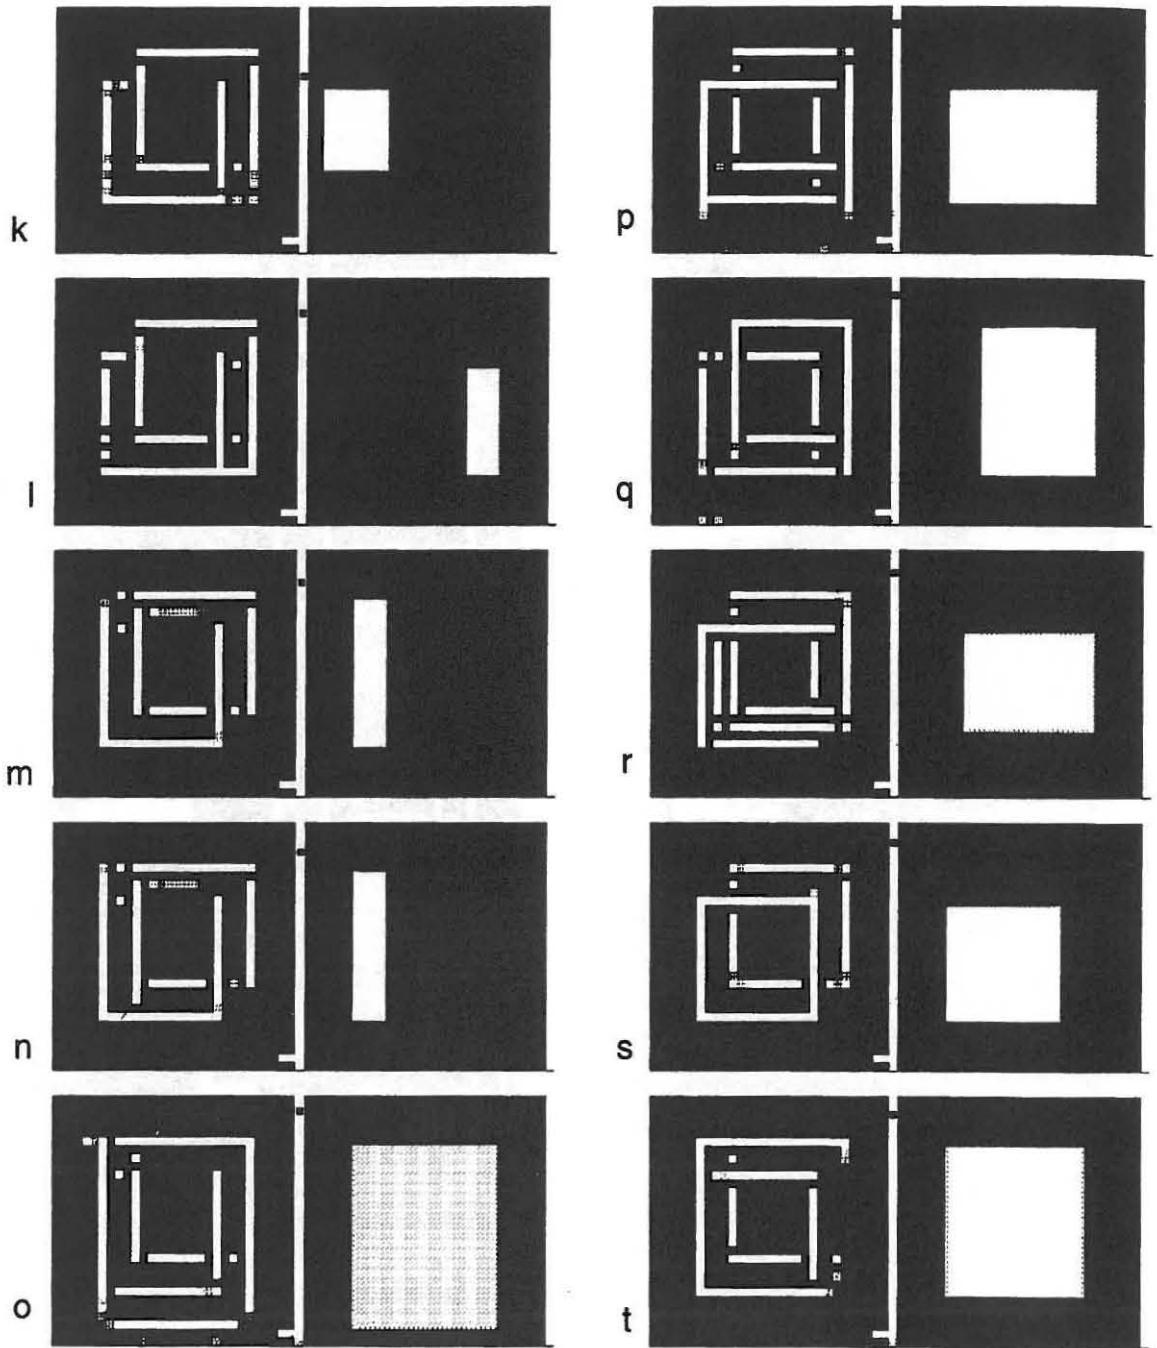


Figure 6. k-t: A sequence of responses to a constant input (cont.)
 (each pair has the retinal image to the left and its interpretation to the right)

4. Conclusions

The main theme in this paper has been virtual imagination, or creative image synthesis. Though this is opposite to pattern recognition research, its methods are utilized here. The novel idea is to combine recognition with positive feedback that will resynthesize a recognized pattern. In this respect the present work extends e.g. that by Liu (1994), where recognition was the only concern. Another difference from his work is the interpretation of Crick's searchlight hypothesis. While he took a rather concrete approach by dynamically delimiting the visual focus on the level of the retina, here the searchlight works on the highest level of abstraction by dynamically selecting which pattern recognizers are active.

The more layers there are in the network, the more sophisticated are the results. The pattern recognizers alone, working directly on the retina, could well produce sequences of images that appear and fade away each in turn. However, the feature layer placed between the retina and the pattern recognizers lets partial images also formate and grow spatially like in an annealing process. Adding more intermediate layers in the future will facilitate more complicated ways of images to build up and to make more subtle associations (connotations) between each other.

In the initial experiment described in this paper the images are not very artistic, since the only concept that the system understands is rectangles. However, more interesting fantasy can be expected with the same mechanism by using a richer collection of recognizable shapes. An important future extension will be a learning or self-organizing memory (Kohonen, 1984) that lets the system form the recognizable patterns itself, based on its visual experience through a video camera. Another extension will be a mechanism to simulate positive and negative attitudes to recognized signals, implementing a kind of pleasure indicator (Campbell, 1973, Takala, 1993).

References

- Campbell, H.J. (1973). *The Pleasure Areas - A New Theory of Behaviour*, Delacorte Press, New York.
- Crick, F.:1984, Function of the thalamic reticular complex: The searchlight hypothesis, *Proc. National Academy of Sciences of USA*, Vol.81, Neurobiology, pp. 4586-4590.
- Hebb, D.O. (1958, 1972). *A Textbook of Psychology*, W.B.Saunders, Toronto.
- Kohonen, T. (1984). *Self-Organization and Associative Memory*, Springer-Verlag, Berlin.
- Liu, Y-T.:1994, Encoding explicit and implicit emergent subshapes based on empirical findings about human vision, in Gero, J. and Sudweeks, F. (eds), *Artificial Intelligence in Design '94*, Kluwer Academic Publishers, pp. 401-418.

- Mandelbrot, B.B.:1977, *The Fractal Geometry of Nature*, W.H.Freeman & Co, New York.
- Mitchell, W.J.: 1993, A computational view of design creativity, in Gero, J. and Maher, M-L. (eds.). *Modeling Creativity and Knowledge-Based Creative Design*, Lawrence Erlbaum, 1993, pp.25-42.
- Palmer, J., Ames, C.T., Lindsay, D.T.:1993, Measuring the effect of attention on simple visual search, *Journal of Experimental Psychology: Human Perception and Performance*, **19**(1), 108-130.
- Perlin, K.:1985, An image synthesizer, *Computer Graphics*, **19**(3) (Proc. SIGGRAPH'85), pp. 287-296.
- Prusinkiewicz, P. and Lindenmayer, A.:1990, *The Algorithmic Beauty of Plants*, Springer-Verlag.
- Sims, K.:1991, Artificial Evolution for Computer Graphics, *Computer Graphics*, **25**(4) (Proc. SIGGRAPH'91), pp. 319-328.
- Takala, T.:1987a, Theoretical framework for computer aided innovative design, in H.Yoshikawa and E.A.Warman (eds.), *Design Theory for CAD*, pp. 323-334.
- Takala, T.: 1987b, Intelligence beyond expert systems: a physiological model with applications in design, in P.J.W.ten Hagen and T.Tomiyama (eds.), *Intelligent CAD Systems I*, Springer-Verlag, Berlin, p.286-294.
- Takala, T.: 1993, A neuropsychologically-based approach to model creativity, in Gero, J. and Maher, M-L. (eds.). *Modeling Creativity and Knowledge-Based Creative Design*, Lawrence Erlbaum, pp.91-108.
- Turk, G.:1991, Generating textures on arbitrary surfaces using reaction-diffusion, *Computer Graphics*, **25**(4) (Proc. SIGGRAPH'91), pp. 289-298.
- Witkin, A. and Kass, M.:1991, Reaction-diffusion textures, *Computer Graphics*, **25**(4) (Proc. SIGGRAPH'91), pp. 299-308.