

Learning to Colour Greyscale Images

Penousal Machado*, André Dias[†] and Amílcar Cardoso[†]

* Instituto Superior de Engenharia de Coimbra, Quinta da Nora, 3030 Coimbra, Portugal, machado@dei.uc.pt

[†] CISUC - Centre for Informatics and Systems, University of Coimbra, Píhal de Marrocos, 3030 Coimbra, Portugal, adias@student.dei.uc.pt ; amilcar@dei.uc.pt

Abstract

This paper is about the colouring of greyscale images. More specifically, we address the problem of learning to colour greyscale images from a set of examples of true colour ones. We employ Genetic Programming to evolve computer programs that take as input the Lightness channel of the training images and output the Hue channel. The best programs evolved can then be used to give colour to greyscale images. Due to the computational complexity of the learning task, we use a genome compiler system, GenCo, specially suited to image processing tasks.

1 Introduction

The work presented here is part of a wider research project, NEvAr, whose aim is to build a constructed artist (i.e. a program that generates artworks autonomously).

NEvAr is an Evolutionary Art Tool inspired on the work of K. Sims (Sims, 1991). It relies on Genetic Programming (GP) (Koza, 1992) to evolve populations of images, based on aesthetic principles. Fitness assignment plays, like in most Evolutionary Computation systems, a key role since it guides the evolutionary process.

NEvAr can be used as an Interactive Evolution tool. In this mode of operation the user supplies the fitness values to the evolved images. It can also be used as a fully autonomous system. In this case fitness is assigned through an explicit fitness function, which takes into consideration several complexity estimates of the images (Machado and Cardoso, 2002).

However, the fitness assignment procedure only takes into account the lightness information of the images, discarding the hue and saturation information. Therefore, in this mode of execution, we are limited to greyscale images. A full description of NEvAr and of the automatic fitness assignment can be found in (Machado and Cardoso, 2002).

There are good theoretical and artistic reasons to deem colour less important than lightness. The development of the colouring procedures of systems like AARON (Cohen, 1995) is, to some extent, based on this notion. However, this collides, at least apparently, with the importance given to colour by some of the most prominent painters (e.g. (Kandinsky, 1911)). Moreover, NEvAr's limitation to greyscale images, in its autonomous version, was frustrating, to say the least. In this paper we address the problem of giving colour to greyscale images.

An analysis of the role of colour and the way colour is assigned, particularly in abstract art, leads to the conclusion that artists (certainly not all, but at least a significant

proportion) usually work with a limited colour palette, and that the spatial relation between colours usually follows some rules. This is consistent with the view that each artist constructs its own artistic language, which complies with an implicit grammar. It is also consistent with the approach used in AARON to colour its drawings (Cohen, 1995; Cohen, 1999).

The idea of creating a program to give colour to the greyscale images created by NEvAr emerged naturally. Unfortunately this poses several problems. Our system is based on a non-symbolic approach and produces bitmap images, hence there is no clear definition of closed forms, shapes, etc.

Therefore, although we could define a palette to work with, assigning the colours of that palette to specific forms would be difficult since we have no forms to begin with. Even assuming that the forms could be properly identified by some sort of pre-processing method, assigning the right colour to each shape and keeping a proper spatial relation among colours would still be a problem, due to the unstructured nature of the output.

Additionally, the creation of a colouring system, by itself, doesn't appear to be an easy task, involving the choice of an adequate set of palettes, establishing a consistent colouring grammar, etc.

Taking these facts into consideration, and also the fact that the generality of this type of approach would be limited, we decided to abandon this idea. Instead, we are trying to create a system that learns to colour images from a set of training examples.

This approach has, potentially, several advantages over a built-in colouring procedure, namely:

- We don't need to code by hand a set of colouring rules.
- The results of the system are less predictable.
- We can use paintings made by well-known artists as training set, hence learning to colour images according to their style.

Furthermore, it's also an indirect way of testing if the colouring procedures followed by some artists can be formally expressed.

In the remainder of the paper we will describe our efforts to build a system to colour greyscale images. The paper is structured as follows: in Section 2 we describe our approach, giving the necessary implementation details; in Section 3, we present some experimental results attained by this approach and make a brief analysis; finally, in Section 4, we will draw some conclusions and present our ideas for future research.

2 Our Approach

GP is one of the most recent Evolutionary Computation techniques. Its goal is to evolve populations of computer programs, which improve automatically as evolution progresses (Banzhaf et al., 1998).

Due to the outstanding influence of the work of Koza (Koza, 1992) it is common, within the Machine Learning community, to associate the term GP to the evolution of tree structures. In this paper we follow this "classical" definition. Therefore, when we talk about GP we are talking about the evolution of tree structures, which are built from a set of functions (f-set) and terminals (t-set). The internal nodes of the tree are members of the f-set, and the leafs are members of the t-set.

In our approach we use GP to evolve populations of programs that give colour to greyscale images. We start by selecting a true-colour training image (or set of images), which is split in its Lightness, Hue and Saturation channels.

The evolved programs take the greyscale image corresponding to the Lightness channel as input, and output a greyscale image. The output is compared with the Hue channel of the training image, the closer the output is to the desired one, the higher the fitness. The same procedure can be applied using the Saturation channel, to evolve programs that generate the saturation information.

As evolution progresses, the quality of the individuals increases and, eventually, we find programs which generate colourings very close to the original ones.

It is important to notice, however, that this is not enough - the idea is to use these programs to give colour to different images. The fact that a program produces an output that exactly matches the training one does not guarantee that it will produce an interesting colouring on different images. In other words, the evolved programs must generalise well. To promote their generalising capabilities, we took some precautions in the selection of the function set and also in the construction of the fitness function. In the remainder of this section, we describe the options taken and give justification for these options.

2.1 Implementation details

A first word goes to how the output of each program is calculated: given a particular individual, it is run for each of the pixels belonging to the training image (or images). Therefore, assuming a training image of 100*100 pixels, each individual must be run 10000 times. Each execution of an individual implies the transversal of its tree and calling, for each node, the corresponding function.

When we take into account that the individuals can easily reach sizes of several thousand nodes, and that GP populations usually contain several hundred individuals, the conclusion that the execution step is of considerable computational weight clearly follows. In order to minimize this problem we implemented our system with GenCo, a Genome Compiler system specially suited for image processing tasks (Dias et al., 2002).

A Genome Compiler is a GP system that makes online compilation of the evolved programs. In situations like the one previously described, i.e. in which each individual must be executed several times, this type of system can provide significant speed improvements, since each individual is compiled once and the resulting machine code executed several times (10000 considering a training image of 100*100 pixels). In this scenario, genome compiler systems are, typically, 50 to 100 times faster than standard C based GP implementations (Fukunaga et al., 1998; Nordin, 1994; Nordin and Banzhaf, 1995; Dias et al., 2002).

Next we describe the design options made in the selection of the function and terminal set, and the reasons that justified them.

2.1.1 Function and Terminal Sets

Our problem has some similarities with the symbolic regression of functions or images (Koza, 1992; Nordin and Banzhaf, 1995). There are, however, some important differences. In a symbolic regression task, one would usually resort to a function set composed by the arithmetic operations and the if statement, and use as terminal set the variables X, Y . This type of set-up unavoidably results in programs whose output depends exclusively on the coordinates of the pixel being calculated.

In an attempt to solve this problem, we added to the terminal set the lightness value of the pixel being evaluated, thus giving more information to the program. This allows the output to vary in accordance to changes on the lightness channel. Additionally, we also added the lightness values of the adjacent pixels, so that the programs have access to the surrounding context of each pixel.

In what concerns the function set, we decided to keep the "traditional" one. The reason for this choice is threefold: it was deemed sufficient for a first approach; the inclusion of more complex functions (e.g. for-next loops, or high level constructs) would severely increase the computational weight of the evaluation step; the inclusion of this type of functions doesn't necessarily benefit the evolutionary process, in fact, the opposite can, and frequently happens (see, e.g., (Banzhaf et al., 1998)).

Taking all these factors in consideration we used the following function and terminal sets:

- F-set = $\{+, -, *, \%, \sin, \cos, if\}$, where % stands for the protected division operator, and if for the if-less-then-else statement (Koza, 1992).
- T-set = $\{X, Y, A..I\}$, where X, Y are variables corresponding to the coordinates of the pixel, and $A..I$ are the lightness values of the pixel being calculated and of the surrounding ones.

The results achieved were very disappointing, basically because the behaviour of programs continued to be determined, almost exclusively, by the values of the X and Y variables. This is clearly undesirable, since it means that we are not evolving programs that give colour to images according to their lightness channel, we are merely evolving a function that, when applied over an interval of X, Y values, generates the hue channel of the training image. Thus, we are only memorizing the training instance(s).

Taking the X, Y variables from the terminal set resulted in having poor evolution, and convergence to trivial and uninteresting colourings (e.g. having as a result the predominant colour of the training image, or always assigning to a specific lightness value the same hue or saturation).

It was clear that the variables X, Y were not producing any desirable impact on the programs. Therefore, we decided to delete them from the terminal set. It was also clear that without these variables the programs hadn't enough information to determine the appropriate colour for each pixel, since the programs only have a local view of the lightness channel (the pixel being evaluated and the nine surrounding ones). To compensate this lack of information, we introduced a new function, $get(\Delta x, \Delta y)$, whose behaviour can be described as follows: assuming that the current pixel has the coordinates a, b , it returns the lightness value of the pixel situated on $(a + \Delta x, b + \Delta y)$.

Since the get function provides a way to access the lightness values of the surrounding pixels, there was no reason to make these values as part of the terminal set. Accordingly, our function and terminal sets became:

- F-set = $\{+, -, *, \%, \sin, \cos, if, get\}$
- T-set = $\{E\}$, where E stands for the lightness value of the pixel being evaluated.

This set-up proved to be adequate, allowing the evolution of suitable colouring programs with good generalization capabilities, as will be shown in section 3. In the following section we focus on the used fitness functions, and on the grounds for using them.

2.1.2 Fitness Functions

We start by describing the fitness function used when we are evolving the saturation channel of an image from the lightness one. In this situation we use the root mean square error between the desired output and the real one to assign fitness, i.e. considering that I holds the desired saturation values and that P is the output of a program the fitness is given by the following formula:

$$s = \frac{1}{1 + \frac{\sqrt{\sum_{x=1}^{\max_x} \sum_{y=1}^{\max_y} (I(x,y) - P(x,y))^2}}{\max_x \times \max_y}} \quad (1)$$

When we are trying to evolve the Hue information, this formula must be slightly altered due to the circular nature of the Hue channel. Assuming that the images take values in the $[0, 1]$ interval, the above formula yields a maximum distance when $I(x, y) = 1$ and $P(x, y) = 0$ (or vice-versa). However, in this situation the difference in hue would be quite small and probably unnoticeable to a human observer. Therefore, we use the following formula:

$$h_a = \frac{1}{1 + \frac{\sqrt{\sum_{x=1}^{\max_x} \sum_{y=1}^{\max_y} [\min_{\theta}(I(x,y), P(x,y))]^2}}{\max_x \times \max_y}} \quad (2)$$

where I holds the desired Hue values, P the program's output, and $\min_{\theta}(x, y)$ returns the minimum angle distance between x and y . This fitness function can be further improved if we take into consideration that the perceived difference in hue depends on the lightness of the pixel (e.g. if the pixel as lightness equal to zero it will always be black, no matter what hue we assign to it; even when the lightness is only close to zero, giving to that pixel a hue different than the desired one will hardly be noticeable to a human viewer). Taking this into account, and considering that L holds the lightness information we obtain the following formula:

$$h_{a'} = \frac{1}{1 + \frac{\sqrt{\sum_{x=1}^{\max_x} \sum_{y=1}^{\max_y} [\min_{\theta}(I(x,y), P(x,y)) \times (1 - 2 \times |0.5 - L(x,y)|)]^2}}{\max_x \times \max_y}} \quad (3)$$

The tests conducted using formulas 2 and 3 as fitness functions yield deceptively good results. In the early steps of evolution fitness increases steadily and swiftly, giving the impression that an adequate colouring program will be easily found. However, after a few hundred generations, the improvements in fitness drop suddenly and evolution seems to halt. This wouldn't be a cause for concern if the evolved programs solved the problem at hand satisfactorily. To some extent they do, since they usually have high fitness values, relatively close to the maximum attainable fitness. The problem is that a qualitative analysis of the results reveals that the colourings are usually uninteresting from an aesthetic perspective - typically, only a small subset of the original colours is used, resulting in the loss of nuances that make the original colouring work. This situation becomes more severe when the training images have a large area filled with a particular hue value, and other areas filled with hue values close to that one.

To better explain the problem we will use an example: consider, for instance, a landscape painting mostly filled with green (for the grass and trees) and with a blue sky (green

and blue are relatively close in the colour spectrum); a program that outputs the hue corresponding to the green colour will have a relatively good fitness, since it is not penalised in the dominant green area of the image and only suffers a little penalisation on the blue area. To make things even worse, such a program is easy to evolve, so it will be found in a small number of generations. Changing it by mutation or crossover will tend to decrease its fitness, meaning that the fittest descendants will tend to be similar to it. Therefore, in a few generations the population will be dominated by similar programs, which will begin to increase in their size to protect themselves from destructive crossover and mutation; soon the increase in size becomes exponential and evolution becomes impossible (the exponential growth of program size is usually called bloat problem; a more detailed explanation of why bloat occurs can be found in (Banzhaf et al., 1997).

In an attempt to force our system to evolve more interesting colourings, covering a wider colour spectrum, we decided to add another factor to our fitness function. A description of the procedure used to calculate that factor follows.

We start by taking the I image (that holds the desired hue values) and decrease its colour depth, using the optimised median cut algorithm, obtaining an image I' composed by only 16 different hue values ($v_1..v_{16}$). Then we count how many pixels exist of each different hue and store the pixel count values in an array, $A_{I'}$. For each pixel of image P , which holds the output of the program, we determine the closest hue value, v_i , and the distance Δv between the pixel value and v_i . We add to $A_{P'}[i]$ the value $1 - \Delta v$, thus if the pixel's hue matches exactly one of the sixteen hues present in the training image we add one, when it doesn't match exactly we add slightly less. After performing this procedure for all the pixels of the output image we compute the following formula:

$$h_b = \frac{1}{1 + \frac{\sqrt{\sum_{x=0}^{15} (A_{I'}[i] - A_{P'}[i])^2}}{16}} \quad (4)$$

Thus, we are basically comparing the number of pixels of each hue value of the output and target image. Returning to our previous example of the green and blue landscape, if the output image has the same amount of blue pixels and green pixels than the original (and also assuming that it is the exact blue and green tone), formula 4 will give a value close to one. However, if the output image is dominated by the green colour there will be a huge discrepancy between the amount of green and blue of the two images, and therefore H_b will be close to zero. Notice that, in what H_b is concerned, the placement of the colours has no real influence on the resulting value. To further force a wider coverage of the colour spectrum, we also used the following modification to this formula:

$$h_{b'} = \frac{1}{1 + \frac{\sqrt{\frac{\sum_{x=0}^{15} (A_{I'}[i] - A_{P'}[i])^2}{\max(A_{I'}[i], A_{P'}[i])}}}{16}} \quad (5)$$

which ensures that all sixteen different hues have the same weight in the calculation. Resorting again to our example, and considering that we have a small yellow area (for the sun), when we use formula 5 having the correct number of yellow pixels is as important as having the correct number of blue or green ones. In the next section we will present some of the experimental results achieved.

3 Experimental Results

In order to test our approach to the colouring of greyscale images we conducted a series of experiments. As training images we selected some of the early works of Wassily Kandinsky. These images were reduced to a size of 96*96 pixels in order to allow a reasonably fast evolution. In its current form our system is only capable of handling square images, so the aspect ratio of the original pictures was altered in order to comply to this constraint. The criteria for selection of the training images was, mainly, individual taste, but we also took into consideration the complexity of the image and the original aspect ratio, to ensure that the images would still be recognizable after being submitted to the previously described operations.

Although our system is prepared to use several training images at the same time, we decided to use, in each experiment, just one training image, since this allowed faster testing and still gives a good indication of the potential of the system. The results presented in this section concern the evolution of programs that take the lightness channel of an image as input and generate the hue channel.

The experimental settings were the following:

- Population size = 100; Tournament selection, Tournament size = 5; Initialisation method = Ramp half and half; Initial tree depth = 2-6, Maximum tree depth = 20; Crossover rate = 0.7, Mutation rate = 0.25, Reproduction rate = 0.05;

The runs were stopped after 1000 generations. We used $h_{a'} + h_{b'}$ as fitness function, since preliminary testing indicated that this type of fitness yield better results.

The best way of presenting the results of our system would be to show the colourings that it generates. Due to the impossibility of presenting true-colour images, we decided to make the figures that should be included in the paper available at the following address:

- <http://www.aisb.org.uk/aisbj/extra/1/2/learningtocolour.pdf>

In the remainder of this section we will make a qualitative analysis of the results achieved. To get a clear view of the system's strengths and weakness the reader should visit the indicated address.

In the appraisal of the results we must take two different aspects into consideration, namely: 1) How well do the evolved programs colour the training image; 2) How do they behave when applied to different images.

In what concerns to the first criterion, an analysis of the results shows that the generated colourings match closely, although not exactly, the colourings of the training images.

Not having exact matches isn't cause for concern. In fact, we were not expecting to, nor particularly interested in, achieving them, since our primary goal is to obtain programs that generalise well.

In the experimental results present here the maximum number of generations was set to 1000. We did, however, conduct some experiments in which this number was significantly higher. These tests showed that fitness continued to increase steadily (though slower). However, after 50000 generations the exact match was still to be found, and at this point, the size of the programs had increased enormously (the average size was close to 15000 nodes) and the rate of fitness increase was extremely slow.

A careful analysis of the generated colourings will also reveal that, in some cases, there are small areas of noise. These noise areas considerably affect the quality of colourings, since it is possible for a human observer to identify these spots, without resorting to a comparison to the original colouring.

Learning to Colour

The fact that we are applying the programs to images significantly larger than 96*96, contributes to the appearance of noise. There are, however, other factors involved.

We tried to identify in which situations evolution may lead to this kind of results. Apparently, they occur when, in the early stages of the evolutionary process, an individual is found that generates a reasonable colouring with some noise areas. Since the evolutionary process is in the beginning, the average fitness is low, and the individual will have a fitness value higher than average, and hence great chances of generating descendants. When the areas of noise are relatively small, they will not hinder significantly the individual's fitness. In these circumstances evolution will tend to improve other areas of the drawing (e.g. try to match exactly the hue of a certain area). The problem is that, after a few generations, the portions of code generating noise will be, not only hidden in the middle of the pieces of code that generate a positive effect on fitness, but, probably, also mixed with them. Thus, it becomes extremely difficult to eliminate the undesirable pieces of code.

This type of problem is not particular to our approach. Both natural and artificial evolutionary systems show this type of shortcoming. In our particular case, the best way to avoid it would be to introduce a penalty for the individuals that generate noise areas. This involves, off course, implementing a procedure to identify these noise areas. Although, we still haven't tried it, we feel that the application of a median filter to the target and generated hue channels, followed by a comparison of the changes introduced by it, should be enough to identify most of the noise areas.

More important than the performance of the programs when applied to the training, is their performance when used to colour images that were not use in their training. After all, this was the original goal of our research.

We weren't, obviously, expecting the programs to reproduce the colourings of images that weren't used in their training. Considering that in each experiment only one training image was used, that could only occur by sheer luck.

Like we stated before, the original motivation for our work was to evolve programs that could be later used to colour the greyscale images generated by NEvAr. Therefore, what we wanted to verify was if the evolved programs could generate consistent, appealing and believable colourings.

The experimental results indicate that this is the case. Some of the colourings generated by the programs are quite striking. Moreover, they present a good and (at least in our opinion) aesthetically sound combination of colours.

It's also interesting to notice that, in some cases, the generated images are surprisingly close to the original ones. Like we stated before we were not expecting for this to happen.

We consider these results to be extremely promising, especially if we take into consideration that there is still lot's of room for improvement. To allow a more equitable assessment of the results, avoiding the bias induced by already knowing the original image, we decided to present three alternative colourings of Wassily Kandinsky painting "Interior (My Dining Room)", without presenting the original painting.

A final word goes to the fact that, in some cases, the generated colourings exhibit some of the already identified illnesses, the presence of noise areas. To some extent, this was expected, since the problem lies somewhere in the individuals' code. The application of the programs to images other than the training one tends to make these problems more visible, which was also foreseeable.

4 Conclusions and further work

In this paper we presented our ongoing research whose goal is to learn to colour greyscale images from a set of training instances. This effort is part of a wider research project that aims at building a fully autonomous constructed artist.

The results achieved so far concern, mostly, the evolution of programs that generate the hue channel from the lightness one. The experiments performed on the evolution of programs to generate the saturation channel seem to indicate that this task is quite simpler. Nevertheless, the replacement of the original saturation channel by one generated through a program will unavoidably imply the introduction of further noise. At the time of writing we can't access exactly how much this will affect the quality of the generated colourings (although we foresee little impact).

There is also the need to conduct additional experiments, especially to use several training images simultaneously, which will hopefully increase the generalization capabilities of the evolved programs. Another aspect that can be improved is the fitness function, we are currently altering it so that it also takes into account the vicinity relations between areas of colour, once again the idea is to promote the evolution of programs that assign colour based in more general concepts.

A final word goes to other types of application of the proposed techniques, for instance, the colouring of black and white movies. Assuming that one of the frames is coloured (either by hand or by some other method) it should be possible to evolve a program that gives correct colours to the surrounding frames.

Acknowledgements

This research project was approved by FCT (Fundação para a Ciência e Tecnologia) and POSI (Programa Operacional "Sociedade da Informação"), and is partially funded by FEDER, project n. POSI/34756/SRI/2000.

References

- Banzhaf, W., Nordin, P., and Francone, F. D. (1997). Why introns in genetic programming grow exponentially. In *Workshop on Exploring Non-coding Segments and Genetics-based Encodings, ICGA*, East Lansing, MI, USA.
- Banzhaf, W., Nordin, P., Keller, E., and Francone, F. D. (1998). *Genetic Programming - An Introduction*. Morgan Kaufman.
- Cohen, H. (1995). The further exploits of aaron, painter. *Constructions of the Mind*, 4(2).
- Cohen, H. (1999). Colouring without seeing: a problem in machine creativity. *AISB Quarterly*, (102):26–35.
- Dias, A., Machado, P., and Cardoso, A. (2002). Improving genome compiler's performance. In *Genetic and Evolutionary Computation Conference, Graduate Student Workshop, GECCO'02*, New York, USA.
- Fukunaga, A., Stechert, A., and Mutz, D. (1998). A genome compiler for high performance genetic programming. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 86–94.

Learning to Colour

- Kandinsky, W. (1911). *On the Spiritual in Art*. Republished by Dover Publications in 1997.
- Koza, J. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press.
- Machado, P. and Cardoso, A. (2002). All the truth about nevar. *Applied Intelligence, Special issue on Creative Systems*, 16(2):101–119.
- Nordin, P. (1994). A compiling genetic programming system that directly manipulates the machine-code. In Kenneth, E., editor, *Advances in Genetic Programming*, pages 311–331. MIT Press.
- Nordin, P. and Banzhaf, W. (1995). Complexity compression and evolution. In *Sixth International Conference on Genetic Algorithms, ICGA95*, pages 310–317. Morgan Kaufmann.
- Sims, K. (1991). Artificial evolution for computer graphics. *ACM Computer Graphics*, 25:319–328.

Figures




Figure 1: On the left column the training images, on right the images resulting from the application of the best individual of each run to the lightness channel of the corresponding training image.

Learning to Colour




Figure 2: The images on the first column result from the application of the program that generated the image a' from Fig. 1; on the second column, images resulting from the application of b'; on the third column images resulting from the application of c'. In the bottom row we present three alternative colourings of W. Kandinsky painting "Interior (My Dining Room)", 1909; The original colouring isn't presented to allow an unbiased assessment of the results.