Evolving Image Enhancement Pipelines

João Correia[1(✉)], Leonardo Vieira[1], Nereida Rodriguez-Fernandez[2], Juan Romero[2], and Penousal Machado[1]

[1] CISUC, Department of Informatics Engineering, University of Coimbra, 3030 Coimbra, Portugal
{jncor,machado}@dei.uc.pt, lmavieira@student.dei.uc.pt
[2] Faculty of Computer Science, University of A Coruña, Coruña, Spain
{nereida.rodriguezf,jj}@udc.es

Abstract. Image enhancement is an image processing procedure in which the original information of the image is improved. It alters an image in several different ways, for instance, by highlighting a specific feature in order to ease post-processing analyses by a human or machine. In this work, we show our approach to image enhancement for digital real-estate-marketing. The aesthetic quality of the images for real-estate marketing is critical since it is the only input clients have once browsing for options. Thus, improving and ensuring the aesthetic quality of the images is crucial for marketing success. The problem is that each set of images, even for the same real-estate item, is often taken under diverse conditions making it hard to find one solution that fits all. State of the art image enhancement pipelines applies a set of filters that solve specific issues, so it is still hard to generalise that solves all types of issues encountered. With this in mind, we propose a Genetic Programming approach for the evolution of image enhancement pipelines, based on image filters from the literature. We report a set of experiments in image enhancement of real state images and analysed the results. The overall results suggest that it is possible to attain suitable pipelines that visually enhance the image and according to a set of image quality assessment metrics. The evolved pipelines show improvements across the validation metrics, showing that it is possible to create image enhancement pipelines automatically. Moreover, during the experiments, some of the created pipelines create non-photorealistic rendering effects in a moment of computational serendipity. Thus, we further analysed the different evolved non-photorealistic solutions, showing the potential of applying the evolved pipelines in other types of images.

Keywords: Image enhancement · Image processing · Computer vision · Evolutionary computation · Genetic programming

1 Introduction

Digital images are, now more than ever, an essential element in our daily lives, considering that almost all online activities depend, in one way or another, on

J. Romero et al. (Eds.): EvoMUSART 2021, LNCS 12693, pp. 82–97, 2021.
https://doi.org/10.1007/978-3-030-72914-1_6

this type of resource. Every image we see in our daily life has a set of attributes that define the way it is perceived. Often, these attributes are not well balanced or optimised for the context of the image and affect image quality.

Image Enhancement (IE) is an image processing approach that aims to improve the perception of a feature or the overall quality of the image, by a person or a computer. Although by definition IE can be done manually, in scenarios involving a large number of images and under different conditions and constraints, the task becomes complex and under specific scenarios unfeasible. In this work, we focus on creating an automatic IE approach. More specifically, we are concerned with creating an approach to automatic enhance images for the context of real estate marketing. Automatic IE brings significant challenges, especially when it comes to manipulating multiple aspects of the image simultaneously since the individual features are not independent of each other. There are multiple types of IE techniques with different purposes and characteristics. Some are more detailed static filters applied to the spatial domain, and others seek to adapt to the image context, avoiding heterogeneous results across multiple images. Based on this research, we explore IE methods focused on improving image aesthetics.

We propose a Genetic Programing (GP) approach that generates pipelines for image enhancement based on image processing filters, with decision components, which aim to alter the pipeline's output, depending on the input image state and features. We resorted to automatic fitness assignment schemes based on the response of an aesthetic evaluator, the Neural Image Assessment (NIMA) classifier [1]. We tested it using a provided dataset of various real-estate pictures of different quality. Furthermore, the outputted enhanced images are evaluated using image quality assessment tools to assess and validate the outputs' quality. We were able to evolve image enhancement pipelines that successfully enhance input images according to the aesthetic evaluate that assigned fitness and other image quality metrics that were only used for validation. However, some of the solutions that the GP approach was optimising were creating non-photorealistic renderings of the input images. The renderings resulted in aesthetically appealing images of arguably artistic merit. In the context of Computational Creativity, this can be viewed as a moment of serendipitous discovery [2], where a system was prepared with an objective in mind, and partially due to chance, an exciting and unexpected value output occurs. In this work, we also explore how the pipelines that create non-photorealistic renderings affect other types of images, further showing those solutions' value.

In this way, the contribution of our work can be summarised in the following main points: (i) the design of an approach that creates a sequence of image filters for image enhancement; (ii) analysis of results obtained with automatic fitness assignment schemes that quantify image aesthetics; (iii) comparison between evolved filters pipeline and a baseline subset of the state of the art image enhancement filters; (iv) analysis of the non-photorealistic effect detected during the experiments; and (v) exploration of the application of the non-photorealistic rendering pipelines on other types of images. The remainder of this paper is organised as follows. Section 2 presents related work. Section 3 present commonly used image filters in the area. Section 4 presents the approach. Section 5 lays the experimental setup and in Sect. 6 we present and discuss the experimental results. Finally, Sect. 7 draws final conclusions and points future work.

2 Related Work

There are multiple types of IE techniques with different purposes and characteristics. The research around this field revolves around machine learning models, computer vision pipelines by applying filters, or both. This section reviews the works related to image filters that are mostly related to our approach.

In terms of image processing and filter approaches, W. Wencheng et al. proposed an IE pipeline that aims to improve the overall brightness and contrast of low-illumination images [3]. C. Y. Wong et al. proposed another pipeline that tries to bridge the problem where approaches that are only based on intensity enhancement may produce artefacts in "over-enhance" regions, and lack enrichment on colour-based features [4]. Moreover, H. Talebi et al. proposed in [5] a novel way of improving an image detail and contrast by expanding on Laplacian operators of edge-aware filter kernels. Closing the classical techniques, we want to mention S. Zhuo et al. in [6], where a noise reduction pipeline is proposed.

Some works integrate evolutionary computation in IE. L. Rundo et al. proposes an evolutionary method based on genetic algorithms to improve medical imaging systems [7]. C. Munteanu also proposed an IE method that relies on evolutionary techniques to improve grey-scale images by evolving the shape of the contrast curve [8]. The work of Shan et al. [9] used an Immune Clone Algorithm (ICA) which make the enhancement method suppress noise and increase the visibility of the underlying signal at the same time on grayscale images.

Most of the approaches mentioned above are mostly non-modular pipelines with fixed parameterisation and applied to solve specific issues with the input image to the best of our knowledge. Based on the review, we moved to implement a set of image filters that will be included in our approach to providing flexibility to the pipelines that we aim to evolve. Thus, the selected filters are described in the next section.

3 Image Filters

A set of 7 previously reviewed image filters were implemented and used during this work. The implemented methods focus on five main aspects of IE approaches contrast adjustment, brightness adjustment, colour balance, noise removal and edge enhancement (also referred as *sharpening*). In this section, we present and explain each one individually.

The contrast in image processing is the range of intensity values available to an image. The contrast stretching is a point operation method that, as the name implies, tries to improve the image contrast by linearly increasing the difference between the maximum intensity value and the minimum intensity value in an image, therefore increasing the contrast level. The work of Bazeille et al.[10] discuss and report of such filter for IE.

Histogram Equalisation (HE) is another method that tries to improve an image's quality by manipulating the contrast. It does this by spreading the most common intensity values by the less common ones, increasing the global contrast of an image. This method is highly used, and there are multiple iterations and discussion about its results [11].

Contrast Limited Adaptive Histogram Equalization (CLAHE) is yet another contrast enhancement method and adaptable for different use cases [12]. It is an iteration of the Adaptive Histogram Equalization (AHE) technique that is an improved version of the regular histogram equalisation. CLAHE improves upon the AHE by clipping the maximum intensity values of each region and redistributing the clipped values uniformly throughout the histogram before applying the equalisation.

Gamma Correction (CB) tries to accommodate the fact that the Visual System (HVS) perceives brightness in a non-linear way. This is done by scaling each pixel brightness from *[0–255]* to *[0–1]* and applying an expression to map the original values.

Non-local Means Denoising (NLMD) [13], as the name implies, tries to reduce the existing noise in an image. It replaces the value of each pixel in each channel to the average of similar pixels.

Unsharp Masking (UM) is an IE technique that sharpens the edges of an image [14]. It does that by subtracting a blurred version of the original image from the original image to create a mask. This mask is then applied to the original image, enhancing edges and details.

Simplest Color Balance (SCB) was proposed by *N. Limare* et al. in [15]. The algorithm tries to remove incorrect colour cast by scaling each channel histogram to the complete *0–255* range via affine transform.

4 Evolving Image Enhancement Pipelines

The analysis of the related work and preliminary work with the implemented filters shown that individually the filters can perform well if under the right conditions but may struggle with versatility, i.e. input images under different conditions may require adjustments on parameters. Moreover, it is clear that applying different filters sequentially can produce unique results and that slight adjustments in the order of the filters may cause significant changes to the output. Furthermore, some images may require the application of one or two filters depending on different conditions. This led us to conclude, that for automatic IE we need to find a generic pipeline for application of image filters suitable to different input images. Contemplating these insights, we developed a way to automatically generate pipelines that compute image filters to be applied to the input images. We opted for GP using a tree representation since the problem inherently can be viewed as a program, a succession of steps and decisions of what filters should be applied on the input image and by what order. GP provides us with a representation suitable for exploring solutions in a structured and flexible way, with variation operators well-defined and adaptable to our problem.

Making use of GP requires the definition of the primitives and terminals that will be available to the population during the evolution. In our scenario, we want to evolve a sequence of filter functions that generally received at least an image and a numeric value as input, i.e. the filter's parameters. We defined a primitive set containing all the seven classical functions previously implemented, a terminal set containing the input image, and *ephemeral constants* ranging from -1 to 1. Each function then mapped the defined range in order to adapt it to the desired magnitude. Each function parameter range was empirically defined so that the function provided acceptable results. Additionally, we introduced an "*if-then-else*" function that, depending on the boolean value of a condition, returns the output of the "*then* tree" or the "*else* tree", allowing the same program solution to behave differently according to the input characteristics.

Since one of the primitives is the whole image, we required values that could be used for comparison to make conditionals. To make this possible, a set of "conditional functions" were introduced. Thus, we added to the primitive set, five functions that extracted relevant features from the image. The features are image-related features that capture characteristics of the perceived quality of the image: noise, contrast, saturation, brightness and sharpness. To extract the noise, we used the work proposed in [16] to quickly estimate the images Gaussian noise. For contrast, we calculated the *RMS* contrast [17], meaning standard deviation of pixel intensities. For saturation, we averaged the pixels' intensity in the *S* channel of the *HSV* colour system. For brightness, we used the *HSP* colour system [18], as it grants a brightness value closer to the real human perception when compared to the luminance (L) channel of the *HSL* or the value (V) channel from the *HSV*. We then averaged the perceived brightness (P) channel to obtain a final value. Finally, for sharpness, we applied a Laplacian filter, calculated the variance of the output and used that as a sharpness score, as proposed in [19].

All these functions were modified to expect an image as input and an ephemeral constant, which serves as a threshold for that condition. Figure 1 shows a graphical example of a possible individual. All the implementation was done using *DEAP* [20] for *Python* as the base evolutionary engine.

5 Experimental Setup

In this section, we present the experimental setup of our work. The GP approach creates image enhancement pipelines to create enhanced outputs of an input image. Without a lack of generality, we deploy the approach in online real-estate marketing scenarios, where the images should be aesthetically appealing to the audience. To validate the outputs we used a set of Image Quality Assessment (IQA) tools that are presented in Section . Afterwards, we present the test sample used in this work's experiments in Sect. 5.2. Furthermore, in the last section, we present the setup for our evolutionary approach in Sect. 5.3.

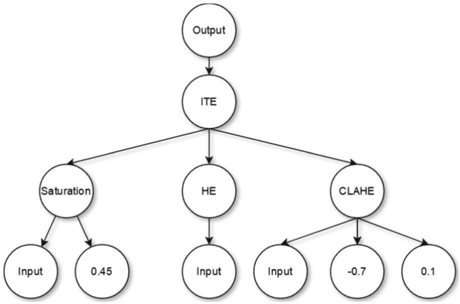

Fig. 1. Graphical example of a possible individual. The numbers represent the ephemeral constants, the *ITE* node represents the *if-then-else* primitive and the Saturation node represents the conditional function.

5.1 Image Quality Assessment

Although our work has its focus on IE, it is essential to understand how an image's quality can be measured. L. He et al. proposes the definition of image quality in three levels: fidelity, perception and aesthetics [21]. Fidelity is how well the image preserved the original information; Perception is how well the image is perceived according to every part of the HVS. Lastly, the image's aesthetic is the most subjective level because it varies from person to person. It is also the most difficult to measure objectively because *"aesthetics is too nonrepresentational to be characterised using mathematical models"* [21]. Since we are dealing with image quality measurements, which most often derives from a subjective appreciation, we must have a deterministic and automated way of qualifying an image quality. To tackle this problem, we made use of 3 distinct *no-reference* IQA tools, where *no-reference* means that the evaluation does not depend on a target or reference image to evaluate the quality of an input image.

PhotoILike (*PHIL* for short) is an IQA service provided by an external company *omitted for blind review*. This service is a closed source, third-party, black-box software that receives an image and returns a value from 1 to 10, where 1 means the worst quality and 10the best quality. Note that the calculated score is not solely based on the images aesthetic but also on multiple features considered relevant for real-estate marketing. For instance, the baseline score of a pool picture is much higher than the bathroom's baseline.

Blind Image Spatial Quality Evaluator [22] (*BRISQUE* for short), is a *no-reference* IQA tool, proposed by *A. Mittal* et al., used in image enhancement contexts [23,24]. As opposed to the previous methods, *BRISQUE* is based on a set of classical feature extraction procedures that computes a collection of 36 features per image. This tool originally outputs a value between 0 and 100, where 0 represents the best quality and 100 the worst. However, in order for the outputs to be in concordance with the previously presented methods, the output was mapped to a 1 to 10 range where 1 means is the worst score and 10 the best.

Neural Image Assessment [1] (NIMA for short), is a *no-reference* IQA tool based on a deep Convolutional Neural Network (CNN), proposed by H. Talebi and P. Milanfar. The paper highlights how the same architecture, trained with different datasets, leads to state-of-the-art performance predicting both technical and aesthetic scores. As the paper states, technical judgment considers noise, blur, and compression artefacts, among other image features. On the other hand, the aesthetic evaluation aims to quantify the semantic level characteristics associated with images' emotions and beauty. Both provided models predict the final score as an average of a distribution of scores between 1 and 10, where 1 means the worst score and 10 the best. Both models were used during the experiments, and require each input image to have a resolution of 224 by 224 pixels. The NIMA aesthetic response was used for evaluation of the individuals during the evolutionary process.

5.2 Test Dataset

To examine our results in an unbiased way, we separated in an early stage, 10% of the 12,090 images from the rest of the dataset. Thus, 1,209 images represent the test set. The rest served as a training set to our approaches, i.e. the dataset used by the evolutionary approach to be selected along the evaluation's evolutionary process. As a way of having a baseline for our measurements, we examined the original images of this test set, using all the IQA metrics presented in Sect. 5.1.

5.3 Evolutionary Setup

All the GP configuration used is presented in sum in Table 1, taking into account the standard GP operators and probabilities [25].

Table 1. Summary of the GP configuration used during the experiments.

Cross-over	One-point
Mutation	Sub Tree mutation - adds a tree with depth between 0 and 2
Selection	Tournament Size 3
Tree generation	Ramped half-and-half
Population size	80
Number of generations	150
Crossover probability	75%
Mutation probability	5%
Elite size	1
Max depth	10

As mentioned in Sect. 4, it is not the objective of these experiments to produce a solution that improves upon a specific image. Instead, we endeavour in searching for a solution as generalist as possible. To achieve this, it is necessary to do a fundamental change to the typical GP evaluation. Each solution will be evaluated on a set of 10 randomly selected images from the training subset. The average fitness of all images will be considered the fitness of the individual. In addition to that, to further prevent overfitting to a specific group of images, a new set of 10 images is selected in each generation. For this reason, it is expected a significant variation in fitness from one generation to another. In all performed experiments, the individual with the overall highest fitness was selected as the test subject for validation.

Defining a fitness function is one of the most crucial steps of building an evolutionary approach. In our case, we wanted a fitness function that evaluated each individual based on its output visual quality. We have selected the *NIMA* classifier tool as it produces results closer to *state-of-the-art* on visual quality evaluation. In the set of experiments presented in this work, we used values of the response from the *NIMA* classifier to assess the visual quality of the images produced by the individuals being evolved.

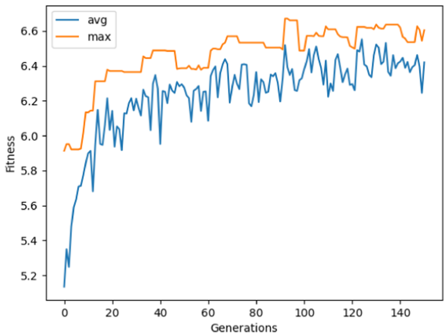

Fig. 2. Fitness evolution during 150 generations using the aesthetic model, with a depth of 10. The results are averages of 15 runs.

6 Experimental Results

In this subsection, we explore the experiments' results in a sub-dataset of real estate image, intending to enhance the image quality. Starting with the evolutionary approach, we can observe in Figure 2 that we can maximise the fitness function and the iterations. Note that since the set of images for evaluation is changing at each generation, the maximum value can oscillate even with elitism. Based on these results, we can see that the process is creating pipeline solutions that enhance the aesthetic score of images according to the NIMA aesthetic classifier output.

To further evaluate the pipeline solutions, we created a validation process where the best pipeline is applied to the test set, i.e. images not used in the evolutionary process. In this way, we can analyse the generalisation of the evolved pipelines. To establish a baseline approach, we conducted the same validation with a manually created pipeline, arranged based on expertise and state of the art references, with the default parameterisation. The pipeline by order of application is as follows: *Contrast Balance (CB), CLAHE, Unsharp Masking (UM, Non-local means denoising (NLMD) and Contrast Stretching (CS).* The best pipeline encountered uses the following pipeline (using to image filters acronyms from Sect. 3): $ITE(Sharp(x, -0.29), SCB(GC(x, 0.41), 0.69),$ $CB(NLMD(GC(GC(HE(CS(CLAHE(x, 0.36, 0.39))), 0.03), -0.97), 0.75),$ $-0.97))$.

The original scores according to the IQA tools are presented in Table 2. In Table 3, and Fig. 3 we can see the results after applying the manual and the best evolutionary pipeline. In benefit of readability, *NIMA* models are abbreviated to respective initials and *PhotoILike* is abbreviated to *PHIL*. All the results are presented in Table 3 and show the metrics improvement over the original images scores. A negative *improvement* score means the degradation of quality according to the respective metric.

Table 2. Average (μ) and standard deviation (σ) of the original images from the test dataset using 4 *no-reference* metrics. All the metrics values range from 1 to 10 where 1 means the lowest quality and 10 highest quality.

	NIMA A.	*NIMA T.*	*BRISQUE*	*PHIL*
Original - μ	4,91	5,36	7,77	5,56
Original - σ	0,45	0,43	1,15	1,41

Table 3. Average (μ) and standard deviation (σ) of the improvement a list of the manual pipeline and the evolutionary pipeline, on the test dataset. All the metrics values range from 1 to 10 where 1 means the lowest quality and 10 the highest quality, and the improvement is calculated by subtracting the original score from the resulting one.

	NIMA A.	*NIMA T.*	*BRISQUE*	*PHIL*
Manual - μ	0,43	−0,06	−0,08	0,25
Manual - σ	0,31	0,26	1,86	0,64
Evolutionary - μ	1,30	−0,55	0,40	0,47
Evolutionary - σ	0,45	0,43	0,10	0,96

The results in Table 3 and Fig. 3 indicate that the manual pipeline resulted in small alterations across all IQA scores, with the NIMA aesthetic response showing some improvements. In average the manually defined pipeline improves

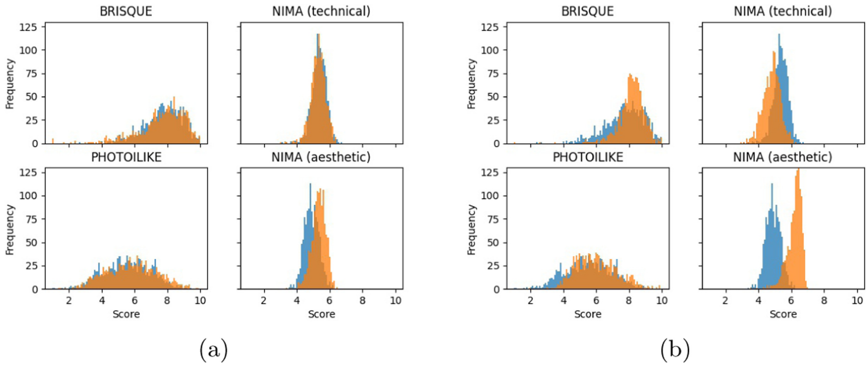

Fig. 3. (a) Graphical comparison between the original test dataset (blue) and the results from the list of classical functions manually selected (orange), computed by all four IQA tools used during the experiments. (b) Graphical comparison between the original test dataset (blue) and the results from using the best evolved solution using *NIMA* aesthetic as fitness, computed by all four *no-reference* IQA tools (Color figure online)

the test dataset in the NIMA and *PHIL*, with the other two metrics showing that it worsens the images. Evaluating the best-evolved pipeline, we can observe that all the metrics show greater increases when compared with the manual pipeline except for the NIMA technical, where for some images tend to be worse than the original and the manual pipeline on that metric, averaging a -0.55 difference for the original and -0.49 to the manual pipeline. Overall, we have significant improvements to the images' aesthetic component, considering the initial distribution of the scores, when using the aesthetic model as fitness. Based on the results, we can also say that the best-evolved pipeline has better results than the manual pipeline, the baseline approach.

Figure 4 show output examples for the experiments using the *NIMA* aesthetic model. All the images presented show an improvement in metrics and subjectively visual improvements compared with the original ones and the baseline. Using the *NIMA* aesthetic model can produce fascinating results, that almost resemble over stylised photos, almost like paintings in some details, instead of real pictures. The model's evaluation tends to associate higher scores with over-edited and saturated images. Another conclusion we can extrapolate from the results is the fact that improvements of the same magnitude as those of this test will also mean extreme adulteration of the original image even if they are considered good for the aesthetic model and *PhotoILike*, showing that this tool is also conductive to high aesthetic scores.

Besides the relevant results, we obtained pipelines where a non-photorealistic rendering effect occurs. The results shown in Fig. 5 were unexpected but a pleasant surprise, indicating a moment of computational serendipity [2]. The overall

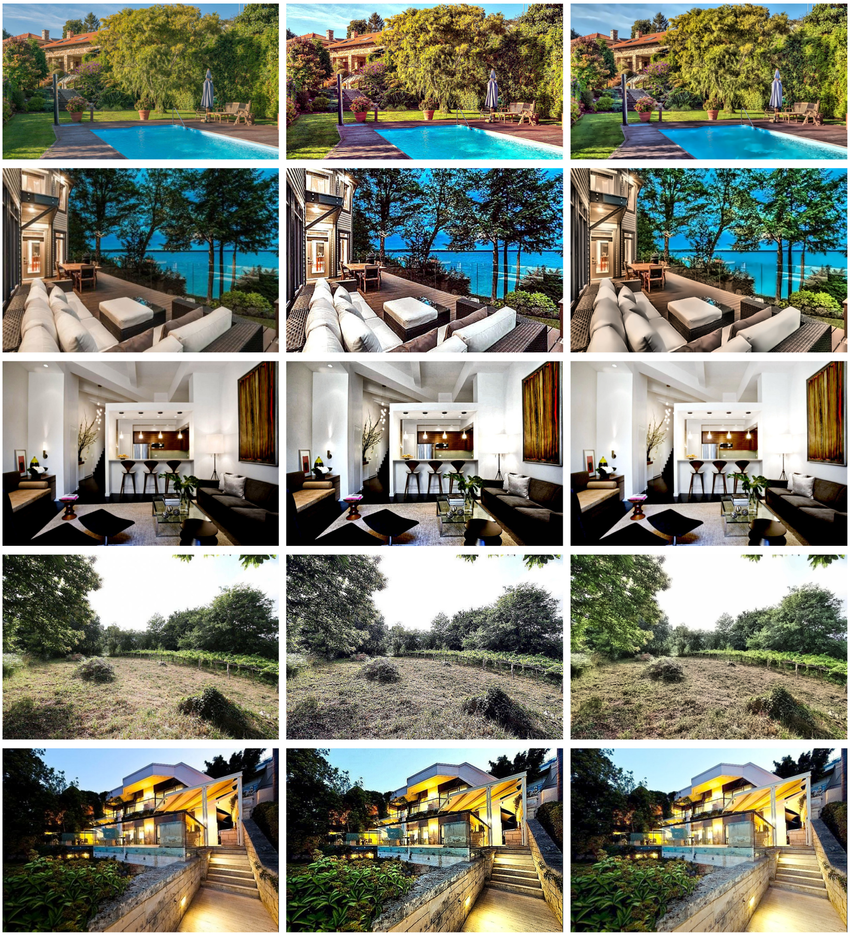

Fig. 4. Examples of images from the test set (left), processed from the manual pipeline (middle) and from the evolved pipeline (right).

idea and goal were to create a filter application pipeline that maximised aesthetic response by using NIMA aesthetic response to guide evolution. Indeed we can maximise the response, and the results indicate that the images are being enhanced across the IQA tools that take aesthetics into account. More importantly, we are improving the PHIL metric, which estimates aesthetics and value from a real-estate marketing perspective. So the system tends to evolve further these solutions which maximise the response of *PHIL* which is the metric more closely related to the problem that we are trying to solve. Moreover, a fascinating fact is that NIMA aesthetic was not trained with paintings; however, it tends

Fig. 5. Examples of images with non-photorealistic rendering effects.

to score higher for the more non-realistic images, which tend to be more like paintings.

Based on these results, we tested evolved pipelines in another type of imagery to evaluate the effect and assess if the results would render aesthetically pleasing images. Figure 6 show some of those outputs. As we can see, the pipelines are generic enough to alter the different type of images, abstract, minimalist and raw photographs. From a subjective standpoint, some of the outputs are more interesting and aesthetically pleasing than the inputs. The results further indicate our approach's potential to create non-photorealistic rendering pipelines for the creation of different artefacts.

For real estate marketing, based on these insights and results, we are convinced that we should incorporate the NIMA technical into the fitness function and some control mechanism to prevent the image from suffering many alterations compared with the original. This line of experimentation and research is already being pursued.

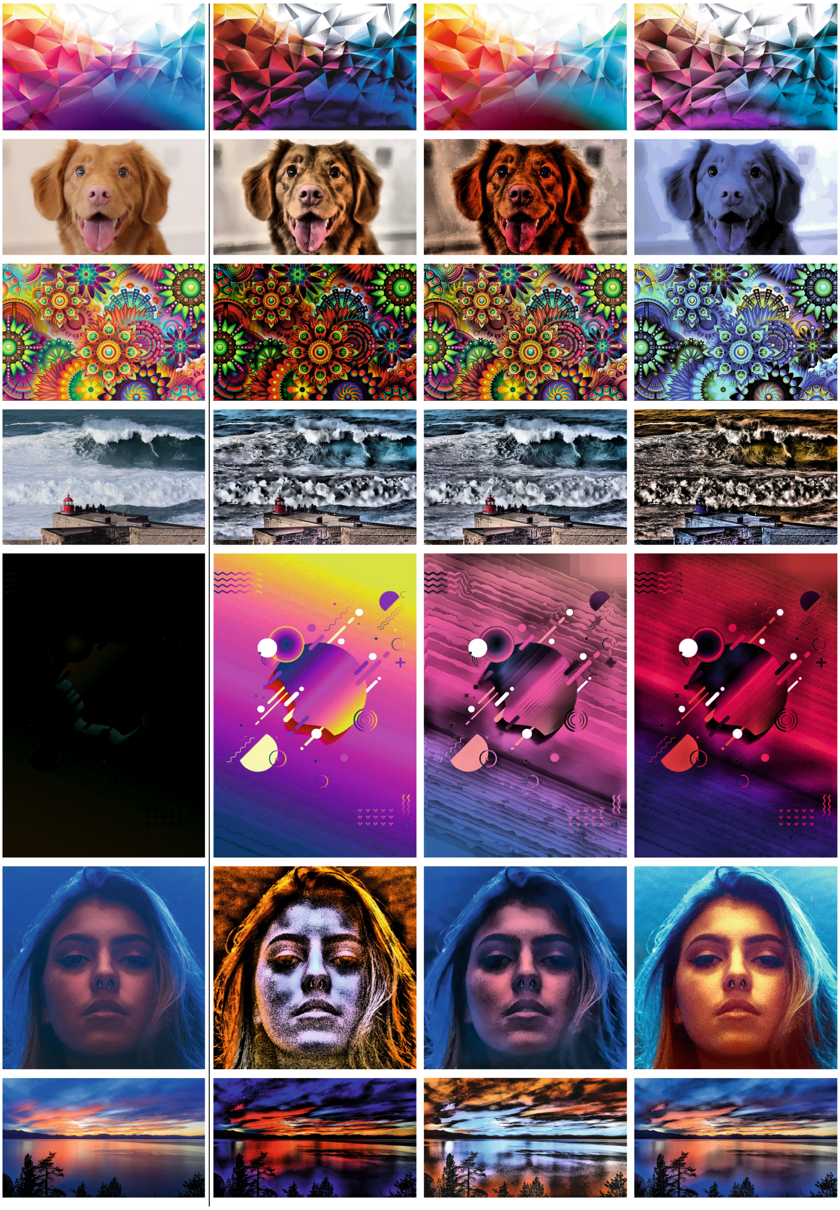

Fig. 6. Samples of images created with a selection of different evolved non-photorealistic pipelines. The original image is the first one on the left.

7 Conclusions

In this work, we presented an approach for automatic image enhancement using GP. The approach is instantiated in real estate marketing, to improve images from different types of real estate under diverse conditions, which requires a more modular approach. We propose an approach that relies on a set of 7 filters from the literature linked to image enhancement and image quality assessment. The context of the problem indicates that the image enhancement should aim towards aesthetically pleasing images. We explore that characteristic by resorting to the aesthetic classifier NIMA to evaluate the individuals. After evolving pipeline solutions, we tested the best solution in a subset of test images compared with the initial results and manual pipeline as the baseline. We show that the system can create filter application pipelines that improve the image quality in all the metrics chosen for image quality assessment. The technical IQA tool metric was the only metric that suffered a negative effect. We argue that the images tend to be altered much to the point that some parts can become more stylised than the original. We also argue that some mechanism to prevent the image from suffering many alterations could mitigate this effect.

During the experimentation, some of the results resulted on a non-photorealistic rendering effect that came to be unexpected and a computational serendipity phenomenon. The system maximised the objective function and the validation metrics response while creating effects that manipulate the images to the point that they became almost paintings. We moved to explore the application in another type of imagery and analysed the effect, showing this approach's potential to create non-photorealistic rendering and transforming effects.

As future work, we plan to expand and alter the set of functions available to the GP algorithm, fine-tuning evolution parameters, more experimentation with the fitness assignment alternatives and changing the evolution dataset, and ways of enlarging the scope of this work. We also plan to incorporate machine learning approaches to the pipeline to be used as filters for the pipeline's input images. Furthermore, in the context of the problem at hand, the results suggest that we should not exaggerate and alter the images' feature that much. We plan on doing a set of experiments using similarity metrics to control the pipelines and prevent them from altering the input images to much.

Acknowledgments. This work is funded by national funds through the FCT - Foundation for Science and Technology, I.P., in the scope of the project CISUC - UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020. This work is also funded by the INDITEX-UDC Program for predoctoral research stays through the Collaboration Agreement between the UDC and INDITEX for the internationalization of doctoral studies.

References

1. Esfandarani, H.T., Milanfar, P.: NIMA: neural image assessment. CoRR abs/1709.05424 (2017). http://arxiv.org/abs/1709.05424
2. Pease, A., Colton, S., Ramezani, R., Charnley, J., Reed, K.: A discussion on serendipity in creative systems. In: Maher, M., Veale, T., Saunders, R., Bown, O. (eds.) Proceedings of the 4th International Conference on Computational Creativity, ICCC 2013, 12 June 2013 Through 14 June 2013, pp. 64–71. University of Sydney, Faculty of Architecture, Design and Planning (2013). http://www.computationalcreativity.net/iccc2013/
3. Wang, W., Chen, Z., Yuan, X., Wu, X.: Adaptive image enhancement method for correcting low-illumination images. Inf. Sci. **496**, 25–41 (2019)
4. Wong, C.Y., et al.: Histogram equalization and optimal profile compression based approach for colour image enhancement. J. Vis. Commun. Image Represent. **38**, 802–813 (2016) http://dx.doi.org/10.1016/j.jvcir.2016.04.019
5. Talebi, H., Milanfar, P.: Fast multi-layer laplacian enhancement. IEEE Trans. Comput. Imaging (2016)
6. Zhuo, S., Zhang, X., Miao, X., Sim, T.: Enhancing low light images using near infrared flash images. In: Proceedings - International Conference on Image Processing, ICIP, pp. 2537–2540 (2010)
7. Rundo, L., et al.: MedGA: a novel evolutionary method for image enhancement in medical imaging systems. Expert Syst. Appl. **119**, 387–399 (2018)
8. Munteanu, C., Rosa, A.: Evolutionary image enhancement with user behaviour modeling. ACM SIGAPP Appl. Comput. Rev. **9**, 8–14 (2000)
9. Shan, T., Wang, S., Zhang, X., Jiao, L.: Automatic image enhancement driven by evolution based on ridgelet frame in the presence of noise. In: Rothlauf, F., et al. (eds.) EvoWorkshops 2005. LNCS, vol. 3449, pp. 304–313. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-32003-6_31
10. Bazeille, S., Quidu, I., Jaulin, L., Malkasse, J.P.: Automatic underwater image pre-processing. In: Proceedings of CMM 2006, October 2006
11. Xie, Y., Ning, L., Wang, M., Li, C.: Image enhancement based on histogram equalization. J. Phys.: Conf. Ser. **1314**, 012161 (2019)
12. Chang, Y., Jung, C., Ke, P., Song, H., Hwang, J.: Automatic contrast-limited adaptive histogram equalization with dual gamma correction. IEEE Access **6**, 11782–11792 (2018)
13. Buades, A., Coll, B., Morel, J.M.: Non-local means denoising. Image Process. Line **1**, 208–212 (2011)
14. Deng, Y., Loy, C.C., Tang, X.: Aesthetic-driven image enhancement by adversarial learning. In: MM 2018 - Proceedings of the 2018 ACM Multimedia Conference, pp. 870–878 (2018)
15. Limare, N., Lisani, J.L., Morel, J.M., Petro, A.B., Sbert, C.: Simplest color balance. Image Process. Line **1**, 297–315 (2011)
16. Immerkær, J.: Fast noise variance estimation. Comput. Vis. Image Underst. **64**(2), 300–302 (1996). https://doi.org/10.1006/cviu.1996.0060
17. Peli, E.: Contrast in complex images. J. Opt. Soc. Am. A **7**(10), 2032–2040 (1990)
18. Rex Finley, D.: HSP color model - alternative to HSV (HSB) and HSL (2006). http://alienryderflex.com/hsp.html
19. Pech-Pacheco, J.L., Cristobal, G., Chamorro-Martinez, J., Fernandez-Valdivia, J.: Diatom autofocusing in brightfield microscopy: a comparative study. In: Proceedings 15th International Conference on Pattern Recognition, ICPR-2000, vol. 3, pp. 314–317 (2000)

20. Fortin, F.A., De Rainville, F.M., Gardner, M.A., Parizeau, M., Gagné, C.: DEAP: evolutionary algorithms made easy. J. Mach. Learn. Res. **13**, 2171–2175 (2012)
21. He, L., Gao, F., Hou, W., Hao, L.: Objective image quality assessment: a survey. Int. J. Comput. Math. **91**(11), 2374–2388 (2014). https://doi.org/10.1080/00207160.2013.816415
22. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. IEEE Trans. Image Process. **21**(12), 4695–4708 (2012)
23. Lim, J., Heo, M., Lee, C., Kim, C.S.: Contrast enhancement of noisy low-light images based on structure-texture-noise decomposition. J. Vis. Commun. Image Represent. **45**, 107–121 (2017). http://www.sciencedirect.com/science/article/pii/S1047320317300603
24. Wang, G., Li, L., Li, Q., Gu, K., Lu, Z., Qian, J.: Perceptual evaluation of single-image super-resolution reconstruction. In: 2017 IEEE International Conference on Image Processing (ICIP), pp. 3145–3149 (2017)
25. Banzhaf, W., Francone, F.D., Keller, R.E., Nordin, P.: Genetic Programming: An Introduction: On the Automatic Evolution of Computer Programs and Its Applications. Morgan Kaufmann Publishers Inc., San Francisco (1998)