Evolutionary and interactive sketching tool for innovative car shape design

BERNARD YANNOU^{1,a}, FRANÇOIS CLUZEL¹ AND MARKUS DIHLMANN²

¹ Laboratoire Genie Industriel, École Centrale Paris Grande Voie des Vignes, 92290 Chatenay-Malabry, France

² Institute of Applied Analysis and Numerical Simulation, University of Stuttgart Pfaffenwaldring 57, 70569 Stuttgart, Germany

Received 5 October 2011, Accepted 20 March 2013

Abstract – So as to create innovative car silhouettes, we propose in this paper a model based on an Interactive Genetic Algorithm using an encoding of a design solution by a Fourier analysis approach. This model permits the designer to browse through generations of car profiles from an initial population of existing silhouettes. By qualitatively assessing each individual, the designer converges towards solutions complying with his/her requirements and preferences, possibly creating novelty and generating surprise. We describe here tests for assessing the efficiency of this innovative design platform. These tests are mainly based on a similarity index, a similarity measure being the perceived distance between two cars silhouettes. The results highlight a good convergence toward a satisfactory solution. In addition, this design process turns out to be very flexible because of the local and intuitive modifications allowed on a given individual solution at any moment of the design process.

Key words: Interactive genetic algorithm / evolutionary design / shape design / subjective evaluation / user tests / car profile

1 Introduction

1.1 Sketching in the first phases of a design process

Design is an engineering activity for creating new and innovative structures and shapes. Finding a new shape and style for an object can be seen as a profound human and sometimes artistic refinement process. Indeed, starting from an initial idea, the style designers continuously refine it through multiple sketches and drawings using their intuition and perception of their own production in a reflexive manner. Is it possible to help such style designers in their refinement process by stimulating creativity? Such an aiding tool should help him or her to explore more easily and systematically a large space of possible styles or shapes, and also to converge towards an ideal shape the designers could have more or less represented in their mind.

In the field of implementing this creative design process, evolutionary computation (EC) has become one of the primary approaches. A method in EC uses basically genetic algorithms (GA) [1,2], which were originally used to find solutions for complex optimization problems. For example, Poirson et al use GAs to optimize the design of

brass musical instruments considering mathematical and perceptual objectives [3]. Taking the evolution in nature as paradigm, the GAs work on the basis of a population of individuals, where each individual represents a possible solution for the initial problem. The structure and the qualities of each individual are encoded in their genomes. Through recombination of these genomes the individuals can reproduce themselves and produce new individuals (solutions), while by a sort of natural selection the individuals who are not adapted to the environment (whose properties do not satisfy the expectations) are not selected for procreation. In this way, the individuals display better and better qualities over the generations. Interactive Genetic Algorithms (IGA, see [4,5]) represent a special class of GAs where a human (here, the style designer) is a key player embedded within the task of selection of individuals of a generation. IGAs are then particularly adapted to situations where it is impossible to explicitly express a preference function (the fitting function) on individuals or even when it is hard to qualify expected properties. This is typically the case with style designers.

A major difficulty when using GAs in automatic design systems is the encoding of the genome (see [4]), which means the way of coding the phenotype (physical structure) of the individual into the genotype (genome). Most systems use a direct encoding where geometrical

^a Corresponding author: bernard.yannou@ecp.fr

Fig. 1. Cheutet models of a car using key lines [12].

dimensions and structures of the design object are directly represented in the genome (see for example [5]). When designing a bottle for example [6,7] or finding a design for cylinder shapes [8] the phenotype is represented in the genome by a sequence of geometrical parameters like radii, lengths and part locations. Consequently, the encoding is context dependent. Other works use tree structures [9] or shape grammars [10] to encode the genome. Kim and Cho [11] have used a set of predefined parts of clothes to find new designs in fashion by recombining these parts. But here the space of possible solutions is limited and we wonder, for all these methods, if an actual innovation results of these design processes.

In addition, all these systems are conceived for a given design domain. Implementing these methods in new fields of design is a difficult and time consuming process. However, a good design method should be applicable, as much as possible, on a large spectrum of situations. A direct change of the shape during the design process is not possible in other systems. If the designer would like to directly change an aspect of the shape, by for example stretching the form or adding a line, this is not possible with the foregoing systems.

1.2 Automatic car shape design techniques

In this paper, we first propose a method of encoding a 2D-closed-curve which is supposed to meet a desired style. This method can be applied to all possible objects represented by their 2D-silhouettes. For instance a car silhouette or profile is a primordial style feature of a car. Indeed, Cheutet [12] has shown that the character of a car profile is primarily expressed through a series of about ten lines (see Fig. 1). Five of them: hood line, windshield line, roof line, wheelbase line and wheel arch, may be merged into a silhouette closed line. These lines and especially the silhouette have been proved to have a strong determining influence on the car perception while embedding

perceptual attributes such as: sportiveness, aggressiveness or peacefulness, etc. In addition, it has been proved that the aesthetic aspects of a car amounts for 70% of purchase intents for customers [12].

Other approaches consider car shape design in a tool for assisting the designer. Petiot and Dagher [13] propose a tool to evaluate car front-end designs thanks to semantic attributes (see Fig. 2).

Osborn et al. [10] use shape grammars to assist the user in the design of new car profiles (see Fig. 3). Kelly et al. [5] describe a car silhouette with 12 points (8 fixed points and 4 varying points) and use an IGA to find new designs (see Fig. 4). These four approaches are synthesized in Table 1 according to their modeling technique, their level of interactivity, the need of an initial population and the needed knowledge level of the user to parameterize or start to use the design engine.

It comes up in Table 1 that no existing approach to design a car shape cumulates a high interactivity and no knowledge prerequisite to parameterize the design platform and further use the design system.

We propose in this paper to detail a principle for encoding the genes of a car silhouette after a Fourier decomposition. Next, an Interactive GA (IGA) has been developed in defining a crossing-over operation between genes. The interactivity consists in letting a style designer qualitatively assessing individuals at each generation. In this manner, new innovative designs are expected to emerge by a balanced collaboration between an automatic process of design space exploration and the interaction of a designer. Finally, we provide measures and user tests for proving that innovation and surprise may emerge from this process in an effective way. Indeed, we show that the initial population of individuals contains a sufficient richness of genes so as to be able to quickly converge towards a desired silhouette which is not an individual of this initial population. Moreover we propose a sketching functionality to permit the direct modifications of car silhouettes by the users.

Fig. 2. Perceptual evaluation of car front-end designs by Petiot and Dagher [13].

Fig. 3. Some stages of the building of a Ford Focus using the shape grammars defined by Osborn et al. [10].

Table 1. Comparison of some interactive car shape design approaches.

Approach	Modelling technique	Level of interactivity	Need of an initial population	Knowledge level of the user
Petiot and Dagher [13]	Design of experiment + Pairwise comparisons	Moderate: evaluation of the car profiles	Yes, based on existing car front-ends	Low: no prerequisites needed
Cheutet [12]	Shape grammars + NURBS curves	High: curves manipulation through shape grammars	No	Expert
Kelly et al. [5]	IGA + 12 points	Moderate: evaluation of the car profiles	No	Low: no prerequisites needed
Osborn et al. [10]	Shape grammars	High: curves manipulation through shape grammars	No	Expert
Proposed approach	IGA + Fourier decomposition	High: evaluation of the car profiles and possibilities of direct modifications	Yes, based on existing car profiles	Low: no prerequisites needed

Fig. 4. Car silhouette modelling from Kelly et al. [5]. Location of points 4, 5, 6 and 7 varies in vertical and horizontal plane. The other points are fixed.

In this way, innovative and new concepts are expected from the collaboration between an automatic process of design space exploration and the interaction of a designer. The main characteristics and advantages of this method appear in Table 1. This system is really open and does not need any knowledge or know-how prerequisite to parameterize the design quest nor formalized objectives – the designer may keep them for him or her –. But validations are necessary to prove that our system turns out to be relevant. That is why we have proposed to define a similarity index to quantify the resemblance between two individuals (car profiles) and to use it in a series of tests to answer to the following questions:

- Are the users satisfied by our model?
- Are the results better with our model than without?
- Does our model really create novelty?
- Are the results different from a user to another?

We first propose an automatic test to ensure the ability of our algorithm to converge. But user tests are also essential to prove that innovation and surprise may emerge from this process. Thus we perform a subjective evaluation workshop to collect user data. Then the results of two main tests and four post-processing tests using the similarity index are established. They show a satisfactory behavior of the model in terms of convergence, diversity, dependence to the initial population, but they also show interesting results about the user perception.

The paper presents in Section 2 the process of the Interactive Genetic Algorithm with the different operations

required for encoding a car profile, generating an initial population and combining the genomes. In Section 3, a similarity index is proposed to measure the perceived distance between two individuals. Section 4 deals with user and post-processing tests to ensure the validation of the model. Finally, Section 5 concludes on some forthcoming perspectives. More details are given previous conference papers [14, 15].

2 Model description

2.1 The genome

Concerning the encoding of a 2D-closed-curve, McGarva [16] has proposed its development into a Fourier series as a method for coding its phenotype. We have personally already used this theory in reference [17] for encoding 2D-closed-curves into the five first Fourier harmonics of this decomposition. In that way, we have been able to build an Artificial Neural Network for synthesizing four-bar linkage mechanisms following targeted trajectories. This approach is not as preassigned as the approach of parameterization for multiple reasons:

- This encoding is supposed to embrace a much vaster space of possible 2D-closed-curves – or 2D-silhouettes – than by a parameterization approach.
- All kind of shapes may be represented even with small details, which can be of the highest importance for provoking feelings and emotions.
- The encoding may be performed through a constant length of genotype, which simplifies a lot crucial GA stages such as the cross-over operation between parent individuals.
- Finally, the genes in our solution have proved to be closely associated to apparent characteristics, which is primordial to converge after several generations to the ideal 2D shapes.

The Mac Garva's theory of Fourier decomposition of a closed curve [16] considers that the position of each point belonging to this curve can be expressed by a complex function in the complex plane (see Eq. (1)).

$$z(t) = x(t) + iy(t) \quad (1)$$

As $z(t)$ is a closed curve, its function is periodic. The period is normalized with: $z(t+1) = z(t)$. This function $z(t)$ can be developed into a Fourier series given in Formula (2).

$$z(t) = \sum_{m=-\infty}^{\infty} a_m \exp(2\pi i m t) \quad (2)$$

where the complex Fourier coefficients can be calculated by Formula (3).

$$a_m = \int_0^1 z(t) \exp(-2\pi i m t) dt \quad (3)$$

Coefficient a_0 is called fundamental, a_1 and a_{-1} represent the first harmonic, a_2 and a_{-2} the second harmonic, etc.

As we will see later, function $z(t)$ is not known as an explicit function from the beginning. Instead, we assume that the curve has been initially defined by a set of successive points z_k ($k = 0, \dots, N$) which belong to the curve. So, in order to calculate a_m coefficients we need a numeric approximation. We obtain this approximation by dividing the curve into N segments connecting each point with its successor. We call t_k the length of the curve between the first point z_0 and point z_k . Under these conditions, the integral can be calculated by the trapezium Formula given in Formula (4).

$$a_m = \sum_{k=0}^N \left(\frac{t_{k+1} - t_k}{2} (z_{k+1} \exp(-2\pi i m t_{k+1}) + z_k \exp(-2\pi i m t_k)) \right) \quad (4)$$

while z is a periodic function, ($z_{N+1} = z_0$).

The value of t_k is the ratio of the length of the curve to point k , namely L_k , and the total length of the curve L . L_k and L are given by Formula (5), knowing that the periodicity imposes $x_{N+1} = x_0$ and $y_{N+1} = y_0$.

$$t_k = \frac{L_k}{L}, L = \sum_{i=0}^N \sqrt{(x_i - x_{i+1})^2 + (y_i - y_{i+1})^2} \text{ and} \\ L_k = \sum_{i=0}^{k-1} \sqrt{(x_i - x_{i+1})^2 + (y_i - y_{i+1})^2} \quad (5)$$

To construct the genome of an object, we develop its silhouette into a Fourier series and define the fundamental (the coefficient a_0) as gene number zero. The first harmonic (a_1, a_{-1}) is called the first gene, the second harmonic the second gene, etc.

On the basis of the genome, the original shape of the individual can be reconstructed. Every point P_k , with coordinates (x_k, y_k) on curve z^* which approximates the silhouette of the car, can be calculated by Formula (6).

$$z^*(t_k) = x_k + i y_k = \sum_{m=-p}^p a_m \exp(2\pi i m t_k) \quad (6)$$

Fig. 5. Decoding of a genome of a Smart car with different precisions.

where t_k ($0 \leq t_k \leq 1$) is the position on the curve and p fixes the number of harmonics used for the decoding. When p equals 1 for example, we use one harmonic to reconstruct the silhouette of the car. The more harmonics used for the decoding the more precise the approximation to the original curve is (as seen in Fig. 5). We call p the “precision” of decoding.

It can be easily proved that the fundamental (the sole complex coefficient a_0) represents the coordinates of the centre of gravity of the curve in a complex plane. The second gene (a_1 and a_{-1}) contains the information defining an ellipse. The influence of the other genes cannot be illustrated easily. But we can say that the first genes influence the very basic structure and shape of the silhouette while the higher genes bring in the details of the shape.

2.2 The process of the interactive genetic algorithm

The process of finding new design solutions can be divided into two phases (see Fig. 6). During phase 1 an initial population of individuals is created. Phase 2 consists of a loop where the user evaluates the current population and a genetic algorithm evolves the population respecting the evaluation of the user.

2.2.1 Phase 1: Creating an initial population

The genetic algorithm needs an initial population of individuals and their genetic code to start working. This

Fig. 6. Diagrammatic plan of the IGA process.

initial population consists of silhouettes of 30 already existing car bodies. In order to easily sketch these silhouettes we programmed an interface in Java which allows to draw curves on a plain and code them into a genome. To border a silhouette we display the image of an existing car in the background of the screen and draw a contour-chart around the car on the image by clicking on the screen. The result is a closed curve representing the silhouette of an existing car-body (see Fig. 7). During bordering, a sufficient amount of points must be used to represent as much as possible the peculiar aspects of the contour. In practice, we often result in 30 to 50 points per silhouette but these points are unevenly distributed.

However this amount M of 30 to 50 points unevenly distributed is not sufficient to calculate a genome which is precise enough to allow a highly detailed decoding into the phenotype because the series of points are generally densified in the regions of details and not added in regions of straight lines. Consequently we need to increase the number of points on the curve by smooth interpolations. The curve produced by interpolation must be very close to the original curve and must be continuously derivable at each point. If the curve is not continuously derivable, the decoding from the genotype into the phenotype produces high-frequency oscillations and is therefore useless. We solve this problem with bicubic splines linking three successive points (see Fig. 8), because this method provides a curve much close to the original one and without producing oscillations (as it is the case when using polynomial interpolations like Lagrange's interpolation Formula). With bicubic splines, the tangent of the spline at point i is parallel to the line passing by points $i - 1$ and

Fig. 7. After bordering we obtain a closed curve representing the car silhouette of an existing car.

$i + 1$. Within each spline, a given number of points is interpolated, leading to a total number of N points with $N > M$.

Taking care of the quality of the encoding amounts to find a satisfactory balance between the number N of points on the curve used for coding and the number p of harmonics used when decoding the genome into a curve. The number p of harmonics used for decoding has an

Fig. 8. Definition of bicubic splines upon the initial point set for further interpolation.

influence on the production of details. The more harmonics used for the decoding the more precise the approximation of the original curve. The number N of points on the curve used for coding the genome has also an influence on the precision of the Fourier coefficients. This is due to the fact that we use the trapezium Formula in (4) to approximate the integral during the calculation of the coefficients. We achieved numerous trials of (1) bordering a silhouette, (2) interpolating with N points, (3) encoding with p harmonics, (4) decoding, for finally comparing the initial and the resulting silhouettes (see Fig. 9). A full factorial design of experiments has been carried out (see Fig. 9 and Tab. 2) with N varying from 80 to 2000 and p varying from 5 to 200. This process has been repeated on a series of ten different car profiles. This is why Table 2 provides qualitative results of comparisons. Apart very smooth profiles where $p = 20$ would be sufficient, $p = 70$ is needed to have recover perfect shapes. We clearly noticed that if p is too small, the coding-decoding sequence visually fails to accurately represent the initial silhouette. In addition, for a given number p , there is a minimal number of points N beyond which the reconstructed curve displays strong oscillations (see Fig. 9 for some examples). In definitive, we found out that a satisfactory choice was achieved with a genome size of 71 ($p = 70$ due to the first gene numbered zero) and a number N of approximately 1500 points for the interpolation since both initial and resulting silhouettes were visually identical (see Tab. 2). This design of experiments in Table 2 probably illustrates the Shannon theorem linking the sampling frequency with the maximum frequency of the curve like in sounds sampling. But a theoretical proof is beyond the scope of the paper.

A last operation of normalization is necessary to the genomes so that the phenotypes – silhouettes – be independent of homothetical transformations, i.e. from a particular location, size or rotation so as to only be compared in terms of their shape. Coefficient a_0 is simply set to 0 to fix the centre of gravity of all individuals at the origin of the representation plane. The invariance by rotation is useless because car silhouettes of the initial population

Table 2. The design of experiments carried out for finding an ideal (p, N) couple. Initial and reconstructed silhouettes are visually compared to result in subjective assessments: I – inaccurate, O – oscillations, SO – strong oscillations, G – Good result, GG – very good result, GGG – optimal results.

$p \backslash N$	80	100	200	500	700	1000	1200	1500	2000
5	I	I	I	I	I	I	I	I	I
7	I	I	I	I	I	I	I	I	I
10	SO	I	I	I	I	I	I	I	I
15	SO	I	I	I	I	I	I	I	I
20	SO	O	G	G	G	G	G	G	G
30	SO	O	G	G	G	G	G	G	G
40	SO	SO	O	O	GG	GG	GG	GG	GG
50	SO	SO	O	O	GG	GG	GG	GG	GG
55	SO	SO	SO	O	O	GG	GG	GG	GG
60	SO	SO	SO	SO	O	O	GG	GG	GG
70	SO	SO	SO	SO	O	O	O	GGG	GG
80	SO	SO	SO	SO	O	O	O	GG	GG
90	SO	SO	SO	SO	SO	O	O	O	GG
100	SO	SO	SO	SO	SO	O	O	O	GG
120	SO	SO	SO	SO	SO	O	O	O	GG
140	SO	SO	SO	SO	SO	O	O	O	GG
170								O	GG
200								O	GG

are sketched horizontally and, consequently, the next generations of cars tend to stay horizontal. Mc Garva [16] proposes to normalize the size of the curve in setting to 1 the small axis of the ellipse defined by harmonics 2. It would amount in our case to fix to a constant height the car silhouettes which is not fair for short cars. We prefer to get a constant or normalized surface area instead. The calculus is then a bit more sophisticated but simply consists in dividing all coefficients a_m by a function of $|a_1|$ and $|a_{-1}|$ (Formula and proof not provided here).

2.2.2 Phase 2: Evolution of the population

We use an interactive genetic algorithm to evolve the population and create innovation. As Kelly et al. [5] say: “by using IGAs we hope to allow designers to enhance their creativity through design space exploration”. The individuals (car profiles) of a generation can reproduce among themselves and produce in this manner new solutions (see Fig. 10 for example).

In our case the genetic algorithm handles a population of individuals where each individual represents a possible design for a car body silhouette. A fitness value is assigned to each individual by the user. Consequently the fitness value f is a number between 0 and 6 according to the grade given by the user via an interface. The interface developed (see Fig. 11) displays six individuals at a time and the user can browse through all the individuals of a population. The user is supposed to evaluate all the individuals of a population on a scale from 0 to 6, where 0 is the worst and six the best evaluation. This fitness decides if an individual has a good chance to reproduce and generate children. Furthermore it influences

Fig. 9. Comparison of silhouettes after interpolation with N points, encoding with p harmonics and decoding.

the chance of an individual to survive and to live on in the next generation. This development is reached by applying the following genetic operators to the population:

- *Selection*: decides which individuals will reproduce and create children.
- *Crossover*: builds a child's genome from two parent genomes.
- *Mutation*: changes in a random way a genome after the crossover.
- *Killing*: decides which individuals from the parents' population survive in the new generation.

We have decided to adopt some conventional choices in term of selection and killing operators and to propose an original crossover operator. First, apart the initial population of 30 individuals, we have fixed the number

of individuals to 100 at each generation. We choose a turnover rate of 0.7, meaning that, for a coming generation, 30 individuals are kept from the previous one and 70 children are generated. In this way we do not lose potential good design solutions. The probability for an individual to be selected to be a parent is proportional to its fitness value (between 0 and 6). After choosing two individuals from the parents' population, their genomes are combined into the genome of a child by applying the crossover and the mutation operators. Afterwards the two individuals are re-put into the parents' population. Indeed, an individual can be selected more than once by the selection operator.

We have considered several possibilities to crossover the two genomes of parents into the one of the child. For

Fig. 10. Example of IGA interface applied on car silhouettes, taken from [5].

Fig. 11. The User Interface for the designer evaluation showing 6 individuals of a larger population. The designer can browse the individuals by clicking on the arrow buttons.

instance, we envisaged a “Two-Part-Crossover”-method which seemed promising at the beginning. It consists in randomly choosing a crossover point X , where X is a number between 2 and 69. The child’s genome was built by the first X genes from the genome of parent A and the last $(70-X)$ genes from the genome of parent B. This method has produced innovative designs for car silhouettes and few useless forms. However the method didn’t produce stable results over the generations, i.e. shapes that can be assigned to car contours in a realistic manner. For instance, after some generations the car silhouettes were useless because they had lost the tires or began to have oscillating contours.

The best choice is to operate a weighted average between the gene values of the two parents to build the genome of the child. A crossover weight W is chosen randomly between 0 and 100. A new gene g^* is obtained by calculating the weighted average of the genes $g_{m,1}$ and $g_{m,2}$ of the parents after Formula (7).

$$g_m^* = \frac{Wg_{m,1} + (100 - W)g_{m,2}}{100} \quad (7)$$

In function of the weight W we obtain different new design solutions which continuously interpolate a silhouette between the two parents’ silhouettes (see Figs. 12 and 13). The advantage of this method is the fact that a child

Fig. 12. Two parent individuals.

resembles a lot to its parents and that it produces almost no useless – or ill – car solutions (the tires keep their rounded shapes). The disadvantage is to stay within the convex hull of the gene space. As a consequence, the population of design solutions may have a tendency to quite rapidly converge. To enlarge the space of possible solutions and create novelty, we apply a mutation operator in a second stage in modifying a given gene – randomly chosen – within a defined range – but with a random value –.

The killing operator is applied to the original population and kills at first all the individuals who have a fitness of 0. These individuals are considered totally useless or totally non-satisfactory and shall no more contribute to the evolution of the population. All the other individuals have a chance to survive. The individuals to be killed are chosen by an inverse roulette wheel method. That means that the probability pk_i for an individual to be killed can be expressed by Formula (8).

$$pk_i = \frac{(7 - f_i)}{\sum_{j=1}^{N^*} (7 - f_j)} \quad (8)$$

where f_i is the fitness of individual i and N^* is the number of individuals in the population who have not been evaluated with a fitness of 0.

2.3 Enhanced interactivity and mutation role

2.3.1 Direct modifications of the car silhouettes by the users

The model previously presented in this paper is interactive in the sense that the user has the possibility to drive the algorithm process thanks to his or her perception and expectations of car profiles. But we believe that this level of interactivity, which is assessed in the next parts of the paper, is greatly improved when concurrently permitting direct graphical modifications by the users. The idea consists in offering to the designer the ability to select an appealing silhouette from the current population and to modify it graphically in an interactive manner. We propose the following scenario (see Fig. 14):

1. The user drives the IGA through several generations.

2. At generation n , he/she remarks an interesting profile but which would merit to be modified at different locations.
3. The given profile is then picked up from the given generation and put into a graphical editor.
4. Starting from the gene code of the chosen individual, a limited number of points (about 30) is generated along the profile so as to provide the user salient points to grasp and stretch. Then modifications are operated by the designer.
5. After the modifications, the new gene code is recomputed and re-normalized as well.
6. Then the new modified profile is reinjected in the current individual population in the place of the chosen one and the IGA process continues.

2.3.2 Some thoughts about the mutation operator

The mutation is the phenomenon that changes the genome of an individual randomly, without taking the genome of the parents into account. In this respect, it brings surprise to the evolution process. And like Kelly et al., we consider the mutation as “a way of introducing new variable values into the population, as well as exposing potential exciting design spaces” [5].

But it is crucial, however, that the mutation be carefully tuned, or “managed”. In order to reveal its full potential, the mutation must ideally be unconstrained, i.e. it might affect all the aspects of the genome representation.

But this freedom has to be moderated by the need for “reliability”: we must also avoid “degenerated” individuals who do not resemble to what an individual is supposed to be. Even if we would like to introduce variety and originality, the generation of new individuals still has to be constrained within some acceptable bounds, so that the generated shape will attract the eye of the designer.

The fact that our interactive genetic algorithm involves user interaction for the evaluation of every individuals of a generation necessarily limits the number of individuals that can be proposed for evaluation at each iteration, and also the number of generations that can be processed in a reasonable time. The designer being at the heart of the process, we must then limit the number of unacceptable shapes that will always be rejected, at the expense of freedom.

The mutation, as it is currently implemented, is based on the fact that “key” genes numbered 2 to 10 are the most influential code parts of the vehicle contour. The degree of mutation change, which is the probability that each individual mutates consists in equally choosing between one of the 9 key genes and to multiply its current value by an element of the set $[-2, -0.5] \cup [0.5, 2]$ (with a uniform probability distribution).

The current Fourier method for encoding the genome of a shape is such that changing one and only one gene at a time when an individual mutates may generate an odd geometric transformation. Whereas such alterations may be acceptable in the case of an automated genetic algorithm, this phenomenon makes the implementation of

Weight W	Child
100
90
80
70
60
50

40
30
20
10
0

Fig. 13. Results of a weighted average crossover between the genes of the two parents of Figure 12, using different weights W .

Fig. 14. Direct modification of a car silhouette by the user.

mutation in an interactive genetic algorithm prohibitive. Actually the designer will not accept any deformity, and straightforwardly eliminate the mutant individual from the current population, even though it exhibits other interesting features and create positive surprise.

It is therefore possible to make some propositions to improve the mutation operator. A solution that stems directly from the aforementioned remarks could be to identify the main features that the user usually recognizes in a class of shapes (e.g. for a car, the location of the wheels, or the flat shape of the floor), and then reshape these features automatically after the mutation. However, such a method may be complex to implement in practice because, given the choice of the genome representation, it would require some back and forth trips between the genome space and the geometrical space. From our work, other methods of mutation are possible and seem more interesting:

- A first method could be a combination of genes; since the mutation of a single gene causes most of the time a very noticeable geometric distortion, a combination of several genes in a single mutation step could reduce the number of unacceptable mutations.
- A second method could be to gather dynamic statistics about the designer's choices: the mutation algorithm could collect the history of the changes it has carried out in all the generations before the current one, and the evaluation of these changes given by the user. Hence, we would statistically favor mutations that please the user, to the detriment of mutations that would never have been accepted in the first place. Although it would seem attractive in the first place, this method would have two main drawbacks. Firstly, it would discard any chance of “invisible” mutations happening, i.e. mutant individuals included in a generation's population but not shown to the user, even though these individuals could be subsequently combined in perfectly viable shapes; and secondly, it requires a large number of samples (records of changes and corresponding evaluation) to be able to provide relevant statistics. However, an interactive genetic algorithm is commonly evaluated on only a few generations, while the user is making his or her choices. A solution could be that statistics be built on a common database, enriched by different users. But besides the technical difficulty of such an implementation, we may “dilute” the specific preference of each user which may vary depending on what he/she is looking for at a particular moment in time.

2.3.3 Synthesis

Combining the classical IGA scheme, a pertinent mutation operator and direct modifications functionality ensure at the same time the mastering of the process by the user and the possibility to create surprise and novelty. The convergence speed can be adjusted by increasing or decreasing the mutation rate and/or by letting the user make more or less direct modifications.

However and even if those two functionalities are available, we propose in the next parts of the paper to validate the model relevance using only the classical IGA scheme. We choose hereafter to consider a fixed mutation rate and not to propose to the user the direct modification tool. This permits to show that the basic model in itself is already better in terms of quality and quantity than a manual sketching method.

Section 3 describes the similarity index built to measure the perceived distance between two car profiles and to validate the model.

3 Similarity index

Is our system really capable to produce innovation and novelty? Is it possible for a user to design with the help of our system a new car body silhouette which is not part of the initial population?

To answer these questions, it is useful to have a way to automatically measure the perceived similarity between two car silhouettes. So we propose to automate a similarity index from two gene vectors coding car silhouettes. This similarity index must be based onto human appraisals of car contour discrepancies; it is not important in itself to contribute to the method but conversely this is an essential point to prove that our method is successful for creating surprise and converging efficiently towards a high quality design solution.

The authors have been working in the area of path generator linkage optimization (see [17]), i.e. optimizing planar mechanism on the basis of a given target closed curve to be described by a point linked to a linkage bar. Here, different solutions exist to compare the distance between two closed curves (the target one and the synthesized one), supposed to be resembling (see [18], pp. 88–92, and [19, 20]). These solutions are conventionally to compare the distance between two sets of NURBS curves. This solution can here also be used for calculating a similarity index and further a fitness rate, but it has not been investigated. In addition, there is an active community working in the more general area of shape similarity index.

We propose hereafter a description of the process to get this similarity index.

3.1 Mathematical expression

We first define $D(k, l)$ the distance between two genomes G_k and G_l . As only the modifications on the ten first genes are significant (modifications on other genes do not change anything on the car profile visual perception), the sum only considers those ten genes. And $D(k, l)$ is given by Formula (9).

$$D(k, l) = \sum_{m=1}^{10} \alpha(m) \|g_{k,m} - g_{l,m}\|^2 \quad (9)$$

The factor $\alpha(m)$ is a weighting factor which should give more importance to some genes according to their participation in the modification of the silhouettes. Here $g_{k,m}$ is

Fig. 15. Choice of the similarity index type.

gene number m from genome k and $g_{l,m}$ is gene number m from genome l . One gene consists of two harmonics, called a_m and a_{-m} , which are complex numbers. So they can be written as: $a_m = u_m + iv_m$.

Then we define the square distance between two genes of the same index by Formula (10).

$$\begin{aligned}
 \|g_{k,m} - g_{l,m}\|^2 = & \frac{(u_{k,m} - u_{l,m})^2}{(u_{\max,m} - u_{\min,m})^2} \\
 & + \frac{(u_{k,-m} - u_{l,-m})^2}{(u_{\max,-m} - u_{\min,-m})^2} \\
 & + \frac{(v_{k,m} - v_{l,m})^2}{(v_{\max,m} - v_{\min,m})^2} \\
 & + \frac{(v_{k,-m} - v_{l,-m})^2}{(v_{\max,-m} - v_{\min,-m})^2} \quad (10)
 \end{aligned}$$

where $u_{\max,k}$ and $u_{\min,k}$ (respectively $v_{\max,k}$ and $v_{\min,k}$) are the maximal and the minimal values of u_k and v_k on the whole initial population.

And we finally define the similarity index between two genomes k and l in Formula (11):

$$\begin{aligned}
 \text{SimInd}(k, l) &= \frac{100}{1 + D(k, l)} \% \\
 &= \frac{100}{1 + \sum_{m=1}^{10} \alpha(m) \|g_{k,m} - g_{l,m}\|^2} \% \quad (11)
 \end{aligned}$$

So with this definition, the similarity index is naturally comprised between 0 and 100%, where 100% means that the two individuals are identical.

We now have to define the factor series $\alpha(m)$. Two different types of functions have been investigated. The first is based on an exponential form. The second one is based on independent weights associated to each value of m . To find out the best version of the similarity index, we performed the tests described in Figure 15. 30 car silhouettes are used to build the two possible forms of similarity indexes. After that, in both procedures the same 20 car silhouettes are used to validate it and permit to choose the best one. We obtain 2 similarity matrices, and we ask the users to assess those 20 profiles in a third matrix. These three matrices permit to calculate the RMSE (Root Mean Square Error) associated to the two calculation approaches for the similarity index.

We obtain with the exponential form of the index an average RMSE of 7.04, and the value is 31.82 with the independent weights form. In addition we obtain with the exponential form a Mean Absolute Error (MAE) of 4.51. As the evaluation scale spans between 0% and 100%, the normalized NRMSE and NMAE are 7.04% and 4.51% for the exponential form, which shows a good precision of the model. We have conducted ten times this test with different sets of car silhouettes, the results obtained have been very similar. So we have chosen the exponential form for the similarity index. Its obtaining is detailed in the next paragraph.

3.2 Identification of the parameters

In this section we assume that $\alpha(m)$ can be written as an exponential expression which gives more importance

to the first genes than to higher order genes, because a modification of the first genes impacts more the car silhouette than a modification of the last ones. So $\alpha(m)$ is obtained after Formula (12).

$$\alpha(m) = a e^{bm} \quad (12)$$

where a and b are two constant terms. So $D(k, l)$ becomes in Formula (13):

$$\begin{aligned} D(k, l) &= \sum_{m=1}^{10} a e^{bm} \|g_{k,m} - g_{l,m}\|^2 \\ &= a \times \sum_{m=1}^{10} e^{bm} \|g_{k,m} - g_{l,m}\|^2 \end{aligned} \quad (13)$$

But now we need to find significant values of a and b .

We propose the following process to measure b :

- Choose a genome, which is copied 3 times: $G0$, $G1$, $G2$.
- Choose a gene i in $G1$ (better with a low weight to be significant): $g1_i$.
- Choose a gene j ($i \neq j$) in $G2$ (better with a low weight to be significant): $g2_j$.
- Modify gene $g1_i$ of $G1$ in an arbitrary way.
- Modify gene $g2_j$ of $G2$ in such a way that there is an *iso-similarity* (defined below) between $G0$ and $G1$ on the one hand, and $G0$ and $G2$ on the other hand.

Two couples of car silhouettes are *iso-similar* if the perceived level of similarity is the same for the two couples. For example, it would mean here that the level of similarity is the same between $G0$ and $G1$, and between $G0$ and $G2$. Practically, it means that the user must modify gene $g2_j$ until the level of perceived dissimilarity becomes the same between $G0$ and $G1$ as between $G0$ and $G2$. In that way, $G2$ and $G1$ are not identical, but their level of similarity according to $G0$ is the same.

From the mentioned principle, we can derive Formula (14) linking distances between genes.

$$\alpha(i) \times \|g0_i - g1_i\|^2 = \alpha(j) \times \|g0_j - g2_j\|^2 \quad (14)$$

with b given by Formula (15).

$$b = \frac{1}{(j-i)} \ln \frac{\|g0_i - g1_i\|^2}{\|g0_j - g2_j\|^2} \quad (15)$$

By making n times these tests with different car profiles and different users, we get n different b values. The final value of b adopted is the average.

The next step consists in measuring a . We propose the following process: for each of the previous comparison (between $G0$, $G1$ and $G2$), the user defines the qualitative level of similarity (“the similarity between $G0$ and $G1$ on the one hand and $G0$ and $G2$ on the other hand is 70%” for example). As it is very hard to express such a value, we propose to work with a 7 degrees qualitative scale; the user just chooses the level of similarity $x\%$ from Table 3 levels.

Table 3. Scale of similarity for user assessments.

Level of similarity	Value of similarity index
0	5%
1	30%
2	50%
3	65%
4	80%
5	90%
6	100%

Equation (16) provides then the expression of $x\%$.

$$x\% = \frac{100}{1 + a \sum_{m=1}^{70} e^{bm} \|g0_m - g1_m\|^2} \quad (16)$$

where b is the averaged value obtained by the previous tests. The second constant a is then obtained by Formula (17).

$$a = \frac{1}{\sum_{m=1}^{70} e^{bm} \|g0_m - g1_m\|^2} \times \left(\frac{100}{x\%} - 1 \right) \quad (17)$$

We also obtain $2n$ different a values. The calculation of the similarity index is now completed, using average values for a and b . In practice, $n = 8$ respondents likes in the user tests. a and b values vary significantly within $\pm 50\%$ around the average value. Finally adopting the average values is may be not the ideal solution but further results with this formula proved that the predictive precision of the similarity index formula is satisfactory. Practically, we use a Java interface (see Fig. 16) that permits to implement the processes described above by loading a population of different contours.

We can now perform user tests to validate our model thanks to this similarity index.

4 Validation tests

To validate our model we first performed simple user tests without the similarity index, based on Kim and Cho works [11].

4.1 First tests

4.1.1 Tests construction

Kim and Cho worked in 2000 on fashion design with an Interactive Genetic Algorithm [11]. They proposed a model to create innovative dresses from a catalogue of components. The main difference between our model and Kim and Cho’s one is the space of possible solutions. As Kim and Cho worked with a set of predefined parts, which is a finite and discrete set, this space is limited, whereas our possible solution space is infinite and continuous.

To validate their model, Kim and Cho performed two different tests:

- The first test is called *convergence test*. It permits to ensure that the average fitness value according to each generation of the Genetic Algorithm increases with the

Fig. 16. Java interface to carry out human appraisals leading to parameters a and b of the similarity index.

generations. They performed this test along 10 generations composed of 8 dresses. They used a panel of 10 users, working on two semantic attributes (*splendid* and *cool-looking*). It means that they asked the users to evaluate 10 generations of dresses according to the first attribute, and then to do the same process with the second attribute. The results showed a significant increase of the fitness.

- The second test is called *subjective test*. Its goal is to show that the results obtained with the model are better than without. Kim asked 3 users to find the 10 best dresses according to the two semantic attributes (*splendid* and *cool-looking*) among 500 individuals randomly created from the catalogue of permitted combinations (i.e. without the IGA engine). Then they asked the 10 previous users to find the best dress of the 10th generation of the previous *convergence test*, according to the two attributes. The last stage was to compare for each attribute and for each user their own best dress obtained with the model with the 10 best dresses obtained without. The comparisons were performed via pairwise comparisons, on

a 7 degrees scale (from -3 to 3). Finally, the results showed that the individuals obtained with the algorithm gained on average 2 degrees, compared to the individuals obtained without IGA algorithm. So the model has been proved to be truly satisfactory.

But two main criticisms can be formulated on these two tests:

- Are the evaluations hedonistic or not? A hedonistic evaluation includes the preference of the user (for example “Please rate these individuals according to your preferences in terms of sportiveness”; the user may be likes sportive individuals, but not too much), whereas a non-hedonistic evaluation does not (for example “Please rate these individuals according to their apparent degree of sportiveness”).
- What is the meaning given by Kim on the pairwise comparisons? He asked the users to compare their own best dress (from the algorithm) with the 10 best dresses found by 3 other people. So it seems logical that the first one (found by the user himself) is better than the others (found by other people).

		Best paper individuals

Best individuals from the IGA		>>	>>>	>
	>	>>>	=
	>	>>>	=

Fig. 17. Example of pairwise comparison matrix used for the tests.

We built our test taking into account these criticisms.

4.1.2 Workshop realization

We have chosen to work for the two tests with a non-hedonistic evaluation (“Please rate these individuals according to their apparent degree of sportiveness”). We have also chosen two semantic attributes: *friendly* and *sportive* (half of the users work with the first one, the others with the second one).

The first test that we have defined is the same that the *convergence test* defined by Kim. We called it test of “*designer satisfaction*”. We have worked with 10 generations composed of 20 individuals, and with a mutation probability of 0.05 and a selection rate of 0.7.

We have defined the second test (that we called test of “*satisfaction superiority of the IGA model*”) with regard to the problem of the pairwise comparisons (expressed in Sect. 4.1). The process is the following:

- 400 car silhouettes are created randomly from an initial population composed of contours of 20 existing cars. To clarify the notations in the next parts of the paper, these 400 individuals are called *paper individuals*. Each user has to find the 3 best individuals according to his/her assigned semantic attribute.
- Practically, these individuals are printed on a large format paper and numbered, and the users record their evaluation in an Excel sheet. Each user evaluates the 400 paper individuals on the same scale as in the IGA (from 0 to 6). The best car silhouettes in his/her mind have to receive a 6 grade.
- Only the individuals which receive a 6 are selected for the next evaluation. An Excel macro sorts out these individuals and a new evaluation table is presented to each user.
- The two first stages are repeated until each user finds the 3 best individuals. That means that the users must be more and more selective (to always have scores between 0 and 6 in order to progress).
- Then each user has to find the 3 best individuals of the 10th generation of the test of “*designer satisfaction*”.
- Finally a pairwise comparison matrix permits to each user to pairwise compare his/her own 3 best individuals of the algorithm with his/her own 3 best paper in-

Table 4. Pairwise comparison scale.

–3	<<<	highly inferior
–2	<<	inferior
–1	<	slightly inferior
0	=	equal
1	>	slightly superior
2	>>	superior
3	>>>	highly superior

dividuals obtained by a manual selection starting from 400 individuals. The scale used has 7 degrees (from –3 to 3). This stage can be compared to the process followed by Kelly et al to validate their IGA [7].

Pairwise comparisons permit to evaluate simply a set of individuals without any absolute scale. An example of such a matrix is given in Figure 17. The comparisons have a direction to follow: here the individuals in lines are compared to the individuals in columns. The evaluation scale for these comparisons is given in Table 4. The users work with the mathematic symbols (>, >>, = ...), which are then replaced by numbers to analyze the data.

The tests have been realized during a short workshop. The panel of users was composed of 7 students (6 men and one woman) and one professor. 4 users received the semantic attribute *friendly* and the 4 others the attribute *sportive*. About 2 h were necessary to complete the three stages.

4.1.3 Results

4.1.3.1 Test of “*designer satisfaction*”

The analysis of the car silhouettes obtained by the users during the first test shows interesting results.

For the two semantic attributes, the users reach in their own last generation the same type of car silhouettes. Moreover, those profiles are nearly the same for semantic attribute *sportive*, whereas they are really different from a user to another with attribute *friendly*. Figures 18 and 19 illustrate this phenomenon.

The displayed silhouettes are chosen from the 3 best ones of each user. It is also interesting to notice that the car silhouettes from the 6th generation are almost the

Fig. 18. Examples of best “sportive” car silhouettes (one from each user).

Fig. 19. Examples of best “friendly” car silhouettes (one from each user).

same for each user, which means that the diversity of the populations is weak from the 6th generation. In other words, the IGA seems to converge quite quickly. It could be interesting to know if this fast convergence is desirable or not. This observation is not tackled in this paper. However, a way to change this speed is the mutation operator. But for the tests, a quite low mutation rate has been chosen to already prove that our IGA approach is already valuable without imposing a gruelling work to the users since they not only have the IGA to use (this use represents a fourth of the test duration). With a high mutation rate, the convergence speed would probably be lower, because a lot of novelty would appear at each generation.

Finally, the functions of fitness evolution are obtained in Figure 20 for the two semantic attributes. The average fitness increases for the two semantic attributes showing a good global behavior of the model: the user satisfaction increases on average from 3.0 to 5.0 for attribute *friendly*, and from 2.3 to 3.9 for attribute *sportive*. So it is an increase of about 70%.

4.1.3.2 Test of “satisfaction superiority of the IGA model”

Table 5 shows the results of the test for each user and for each semantic attribute.

Some details have to be explained to understand the meaning of Table 5. The pairwise comparisons are made according to the following scheme: the best individuals from the IGA are compared to the best paper individuals. User 4 has for example an average evaluation of +1.00. According to Table 4, it means that user 4 found the best individuals from the IGA slightly superior to the best

Fig. 20. Fitness evolution for the convergence test.

Table 5. Average values of the results of the subjectivity test (for the evaluation scale, see Tab. 4).

Sportive	User 1	+1.78
	User 2	+1.67
	User 3	+1.56
	User 4	+1.00
	Average	+1.50
Friendly	User 5	-0.33
	User 6	+0.67
	User 7	+0.78
	User 8	+1.22
	Average	+0.58
Total average		+1.04

paper individuals. Thus positive numbers prove that the results obtained with our model are better than without.

The results show a good behavior for the semantic attribute *sportive*: the average score is +1.50, so the best IGA individuals are between “slightly superior” and “superior” to the best paper individuals.

For the attribute *friendly*, the results show a good behavior too, even if the difference is slightly less perceptible. All the users prefer the individuals from the IGA except user 5, who prefers the individuals without the model. Globally, the results for this attribute are not really homogenous, which can be explained by a more subjective comprehension (and so characteristic to each user) of the word *friendly* than the word *sportive* (this observation is highlighted in other tests presented hereafter).

4.1.4 Synthesis

Those first two tests show a satisfactory behavior of our model. The user satisfaction increases with the generations, and the results are globally better with the IGA than without.

However, some major differences can be noticed between the two semantic attributes. Indeed the user perception is really different from a user to another: *sportive* seems to be perceived in the same way by everyone, whereas each user has a different perception of *friendly*. So to confirm these observations we have proceeded to more sophisticated post-processing analysis, using the

Fig. 21. Comparison between the reference silhouette (a) and the final resulting silhouette (b).

similarity index described in Section 3. We propose here five user tests that use the similarity index to process the same data (except the first one) as those used for the previous tests.

4.2 Tests based on the similarity index

4.2.1 Test of “novelty emergence”

The goal of this test is to show that it is possible to reach a defined individual which is not part of the initial population.

To answer this question we can execute a simple test. A subject draws on a sheet of paper a car body silhouette which comes spontaneously to his or her mind and which is not part of the initial population. This car body silhouette is taken as a “target individual”. With the aid of our proposed IGA, the user must now try to obtain the “target silhouette” he had drawn on a sheet of paper by the end of the IGA process. The objective of this test is of course to prove that there is no limit to our system to converge toward the dreamed shape. Then, keeping in front of him or her this reference drawing, the user evaluates the car contour individuals, in providing higher grades to contours which look close to the target contour and lower grades to those that are more dissimilar. By considering both the number of generations needed to resemble to the target individual and the similarity value, we can estimate the quality of our design system.

Alternatively the target car silhouette may be an individual of the initial population that is removed from this initial population.

We have preferred to make abstraction of the designer subjectivity in automating the ability of the system to converge towards an ideal car silhouette, so as to measure the sole quality of the method. The role of the designer is played by an algorithm of similarity fitness calculation, which automatically evaluates the individuals of a generation in terms of their similarity to the target individual, thanks to the similarity index previously defined.

We repeated several times this test, with different target individuals extracted from the initial population for an average final similarity index of 90%. Let us provide a practical example with the target individual provided in Figure 21a. The parameters for the genetic algorithm were the following: population of 100 individuals, turnover rate of 0.7 and mutation probability of 0.3. The mutation can change a gene in a range of $\pm$ (50%–200%). After 10 generations our system reached the car body silhouette

Fig. 22. The average fitness of all individuals in the population over the generations.

in Figure 21b which has a similarity index of 92%, and which can be visually considered as a much satisfactory result.

The average fitness of the population converges over the generations to a high value (see Fig. 22), whereas the value of the best similarity index in the population (the fitness of the fittest individual) raises rapidly from relative low 44% to 92%.

4.2.2 Test of “diversity lowering”

The aim of this test, called test of “*diversity lowering*” is to study:

- The evolution of the similarity index for each user along the ten first generations.
- The similarities that could exist between the individuals of the last generation of each designer.

The idea simply consists in calculating the similarity index between all the individuals of all the users and for all the generations of the previous workshop, and to group together the data to observe potential correlations.

A high value of similarity for the last generation would mean that the IGA converges towards a single individuals family, and so that there is a diversity lowering.

These results concern the evolution of the similarity index of each user. The curves are shown in Figure 23.

The initial population has an average of 6.5% of similarity. In all the cases and for the two semantic attributes, the values increase quickly. At the 6th generation (Generation #5 on the graph), the averages of all the users

Fig. 23. Evolution of the similarity index along the generations for the two semantic attributes.

Table 6. Similarity values between users for semantic attribute friendly (in %).

	User 5	User 6	User 7	User 8
User 5	79.07	6.91	4.02	2.89
User 6		89.65	20.65	8.81
User 7			61.02	11.97
User 8				74.16
Total average			9.21	

Table 7. Similarity values between users for the semantic attribute sportive (in %).

	User 1	User 2	User 3	User 4
User 1	79.95	47.87	63.09	36.20
User 2		72.34	74.14	42.62
User 3			89.61	41.04
User 4				74.74
Total average			50.83	

are above 40% similarity, and at about 75% for the last generation. The averages of each user (in italic in Tables 6 and 7) are all included between 60 and 90% of similarity, which corresponds to a level of very strong similarities.

This test clearly shows that, for each user, the model converges towards a single family of car silhouettes, which are very close in terms of visual aspect. But it means too that there exists a significant diversity lowering.

4.2.3 Test of “inter-designer convergence”

The goal of this test, called test of “*inter-designer convergence*” is to study the similarities that can exist between the last generations of all the users. We clearly want a design system that preserves the freedom of thinking and creating of each designer, so as to be able to keep the personality of each of them within the last generation of individuals.

As in the previous test, the idea simply consists in calculating the similarity indexes for all pairwise comparisons between one individual of the last generation of one user and one individual of the last generation of the other

Table 8. Maxima of similarity values between best individuals of the IGA and paper individuals (in %).

Sportive	User 1	84.49
	User 2	21.26
	User 3	80.92
	User 4	97.93
	Average	71.15
Friendly	User 5	2.67
	User 6	98.09
	User 7	71.32
	User 8	26.94
	Average	49.75
Total Average		60.45

user. These inter-designer pairwise comparisons are averaged for all couples of designers to result in Tables 6 and 7 for the two considered semantic attributes. In the diagonals of the inter-designer matrices, one retrieves the results of the previous test.

The values are really consistent with the visual aspect of the individuals. For semantic attribute *friendly*, the values inter-designer (in bold) are not high, and the total average is only 9.21% similarity, whereas this total average is 50.83% for attribute *sportive*. It means that the similarity between the last population of the users for attribute *sportive* is high and that, on the contrary, there is few connection between individuals evoking *friendly*.

These results confirm rigorously that the perception of the word *sportive* is the same for all users (they all came to the same type of profiles), whereas each user has its own perception of the word *friendly* (they all have a different kind of car silhouettes at the end of the workshop).

4.2.4 Test of “superiority of the IGA model”

This test is called test of “*superiority of the IGA model*”. Its goal is to show that our IGA design system obtains better results than without, i.e. with a brute force solution enumeration, in particular in terms of time and number of individuals processed to obtain the results.

It consists in comparing for each user the 3 best individuals of the IGA with the 3 best individuals obtained without the IGA (printed on paper) with the similarity index.

The maxima of similarity are shown in Table 8. For 5 users out of 8, the maximum is above 70% of similarity. For 2 users, the maximum is included between 20 and 30%. For the last one, it is only 2.67% of similarity.

It means that, much of the time, the final solutions obtained are quite resembling, whatever the design method (IGA system or brute force enumeration). Indeed, without any judgment of value, 5 users have found with the IGA *at least* one individual that is very close to one of the paper individuals.

But, 400 individuals were proposed on paper and it takes about 45 min for one user to evaluate them. Comparatively, less than 200 individuals are used with the IGA process (10 generations of 20 individuals, minus

Table 9. Example of similarity values for the comparison between the best individuals of the IGA and the initial population (in %).

		User 2		
		Best Individual 1	Best Individual 2	Best Individual 3
Initial Population	Ind 0	2,04	1,93	1,84
	Ind 1	2,95	2,77	2,65
	Ind 2	10,02	9,02	8,80
	Ind 3	0,91	0,88	0,90
	Ind 4	3,45	3,24	3,35
	Ind 5	2,04	1,94	1,91
	Ind 6	1,42	1,36	1,36
	Ind 7	1,78	1,74	1,52
	Ind 8	13,90	13,73	10,75
	Ind 9	6,62	6,17	6,53
	Ind 10	2,76	2,60	2,47
	Ind 11	3,28	3,41	2,91
	Ind 12	2,08	1,97	1,97
	Ind 13	69,65	69,89	53,21
	Ind 14	3,50	3,30	3,46
	Ind 15	7,15	7,23	10,83
	Ind 16	2,07	1,97	1,93
	Ind 17	3,05	2,88	2,87
	Ind 18	1,65	1,58	1,47
	Ind 19	1,99	1,93	1,69

those who survive from one generation to the next one), and the evaluation only lasts about 20 min. So we can say that our system seems to be able to bring comparable results more quickly and with fewer individuals, this is clearly a proof of its utility

4.2.5 Test of “attraction in the surroundings of initial individuals”

The last test proposed in this paper is called test of “*attraction in the surroundings of initial individuals*”. It tries to answer the following question: is the user really able to design his/her own car silhouette, or are the final individuals influenced by the initial individuals and close to their genotypes?

We answer this issue by pairwise comparing the 3 best individuals produced by the IGA process with the 20 car silhouettes of the initial population.

The results of this test are very interesting, but too sizeable to be displayed here in their whole. That is why only the case of user 2 is presented in Table 9.

This table shows that the best individuals of user 2 are very close to Individual #13 of the initial population (more than 50% of similarity). Individual #13 represents the silhouette of a Porsche 911. Two other minor influences can be noticed (with Individuals #2 and #8), but the values are below 20%, which is not much significant. User 2 worked with the semantic attribute *sportive*, and it is important to notice that the same results are obtained with the 3 other users who worked with this attribute.

Porsche 911 acts as an absorbing genome in people’s minds, outperforming from the beginning all the other individuals and killing the competition. In other words, all the users perceive a *sportive* car as a Porsche 911. Porsche 911 has definitely a brand effect and a visual identity.

For semantic attribute *friendly*, the results are totally different. No car that influences all the users can be identified. In 2 cases out of 4, no significant influence of an initial individual is noticed. In the two other cases, such a car can be identified, with the scores of 40% of similarity, but this car is different for the two users (Porsche 911 for

the one, Chevrolet Corvette C4 for the other). It shows again that this attribute is differently perceived.

4.2.6 Synthesis

Those five tests based on the similarity index permit to prove rigorously some results. Indeed in all cases the system turns out to converge towards a uniformed population. For semantic attribute *sportive*, this profile is the same for any of the four users and is close to Porsche 911. For semantic attribute *friendly*, no similar influence from the initial population is identified, and the results really depend on the user. The perception of the attributes is really different from a user to another.

Moreover, our IGA model converges quickly and with fewer individuals than with the sole brute force enumeration method, which is really satisfactory and reveals its utility.

Finally, many issues would now merit to be investigated further:

- The model always converges towards one profile family. Is it not possible to converge towards at least two different families?
- What happens if Porsche 911 or the other influencing cars are not in the initial population? Do the users even find them?
- Why are there so many differences in the perception of the semantic attributes? Some concepts of Emotional Design and Kansei Engineering could be introduced to get a more sophisticated model, but may be at the detriment of simplicity.

5 General conclusions and perspectives

We have presented in this paper an innovative car silhouettes design model. Based on an encoding method of the genomes by Fourier decomposition, it offers to the car designers the possibility to create new car silhouettes from an initial population according to their preferences which are not needed to be made explicit before.

Several user tests have been carried out to study the convergence of the model in terms of user satisfaction, but also in terms of perceived distance, using a similarity index. The results show very satisfactory results. They permit to answer the four questions defined in the introduction:

- *Are the users satisfied by our model?* Yes, the satisfaction increases along the generations of the IGA, and the users comment that they really enjoy to work with our interactive GA system rather than picking up remarkable silhouettes in a catalogue.
- *Are the results better with our model than without?* Yes, in terms of user satisfaction, but also in terms of time and numbers of processed individuals.
- *Does our model really create novelty?* Potentially. It highly depends on the chosen semantic attribute or on the question or objective the user has in a more or less

explicit manner. Anyway, we proved that individuals which are considered as visually far from the initial population may be created.

- *Are the results different from a user to another?* Potentially. It also depends on the chosen semantic attribute too. But we at least proved that personal opinions can be expressed and preserved as long as people are not subject to the same image (like what a sportive car is).

In consequence, our system should allow style designers to converge towards intuitive ideas and to make emerging surprise in exploring large spaces of potential silhouettes.

The mutation operator and the ability to directly modify interesting car profiles during the IGA process have not been studied in detail in this paper. The setting of the mutation rate is a real issue and it will be studied in further work, it is the mean of regulating the convergence speed of the IGA. This speed is probably a bit fast for the moment, but the assessment effort asked to the users at each generation seems sufficient for us. Another way to slow down the convergence could be to keep in the next generation of the IGA a high number of the individuals of the previous generations.

Finally a last interesting perspective is to explore a new product field: the encoding method that we use is very flexible and applicable to many design objects. We think that excellent results could be obtained with simple shapes. The ideal form would be a simple 2D-closed-curve from a solid of revolution (which permits to describe the whole product with a single closed curve), like stemmed glasses or vases.

Acknowledgements. We gratefully thank François Bleibel, Nicolas Cordier, Gilles Foinet and Thomas Ricatte for their contributions to this work.

References

- [1] P.J. Bentley, D.W. Corne, Introduction to creative evolutionary systems, in Creative evolutionary systems, Morgan Kaufmann Publishers Inc., 2002, pp. 1–75
- [2] G. Renner, A. Ekart, Genetic algorithms in computer aided design, Computer-Aided Design 35 (2003) 709–726
- [3] E. Poirson, P. Dépincé, J.-F. Petiot, User-centered design by genetic algorithms: Application to brass musical instrument optimization, Engineering Applications of Artificial Intelligence 20 (2007) 511–518
- [4] F. Nicaise, E.K. Antonsson, Indirect encoding of structures for evolutionary design, in Proc. International Conference on Engineering Design, ICED'07, Cité des Sciences et de l'Industrie, Paris, France, 2007
- [5] J. Kelly, P.Y. Papalambros, C.M. Seifert, Interactive Genetic Algorithms for use as Creativity Enhancement Tools, in AAAI 2008 Spring Symposia, Palo Alto, California, USA, 2008
- [6] M.C. Ang, H.H. Chau, A. McKay, A. De Pennington, Combining Evolutionary Algorithms and Shape Grammars to Generate Branded Product Design, in Design Computing and Cognition '06, J.S. Gero, Ed. Dordrecht, Springer Netherlands, 2006, pp. 521–539

- [7] J.C. Kelly, G.H. Wakefield, P.Y. Papalambros, Evidence for using Interactive Genetic Algorithms in shape preference assessment, *International Journal of Product Development*, 13 (2011) 168–184
- [8] H. Yanagisawa, S. Fukuda, Interactive Reduct Evolutional Computation for Aesthetic Design, in *Proc. ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, IDETC/CIE2003*, Chicago, Illinois, USA, 2003, pp. 1063–1072
- [9] H. Liu, M. Tang, J.H. Frazer, Supporting evolution in a multi-agent cooperative design environment, *Advances in Engineering Software* 33 (2002) 319–328
- [10] S. Osborn, J. Cagan, R. Pawlicki, R.C. Smith, Creating cross-over vehicles: Defining and combining vehicle classes using shape grammars, *AIEDAM* 20 (2006)
- [11] H.-S. Kim, S.-B. Cho, Application of interactive genetic algorithm to fashion design, *Engineering Applications of Artificial Intelligence* 13 (2000) 635–644
- [12] F. Cheutet, 2D semantic sketcher for a car aesthetic design, in *Proc. CPI'2007: Conception et Production Intégrées*, Rabat, Maroc, 2007
- [13] J.-F. Petiot, A. Dagher, Preference-oriented form design: application to cars headlights, *International Journal on Interactive Design and Manufacturing (IJIDeM)* 5 (2011) 17–27
- [14] B. Yannou, M. Dihlmann, F. Cluzel, Indirect encoding of the genes of a closed curve for interactively create innovative car silhouettes, in *Proceedings of International Design Conference, DESIGN 2008*, Dubrovnik, Croatia, 2008
- [15] F. Cluzel, B. Yannou, Efficiency Assessment of an evolutive design system of car contours, in *Proceedings of International Conference on Engineering Design, ICED'09*, Stanford, CA, USA, 2009
- [16] J. McGarva, G. Mullineux, Harmonic representation of closed curves, *Appl. Math. Modell.* 17 (1993) 213–218
- [17] A. Vasiliu, B. Yannou, Dimensional synthesis of planar mechanisms using neural networks: application to path generator linkages, *Mech. Machine Theory* 36 (2001) 299–310
- [18] Sandor, A.G. Erdman, *Advanced Mechanism Design: Analysis and Synthesis*, Vol. II, 1^{re} édition Prentice Hall, 1984
- [19] J.R. McGarva, Rapid search and selection of path generating mechanisms from a library, *Mech. Machine Theory* 29 (1994) 223–235
- [20] C. Zhang, R.L. Norton, T. Hammonds, Optimization of parameters for specified path generation using an atlas of coupler curves of geared five-bar linkages, *Mech. Machine Theory* 19 (1984) 459–466