EAI Endorsed Transactions
on Creative Technologies
Research Article **EAI.EU**

Towards Artificial Intelligence Serving as an Inspiring Co-Creation Partner

Kevin German[1], Marco Limm[1], Matthias Wölfel[2,*], Silke Helmerdig[1]

[1]School of Design, Pforzheim University, Pforzheim, Germany

[2]Faculty of Computer Science and Business Information Systems, Karlsruhe University of Applied Sciences, Karlsruhe, Germany

Abstract

The promise of *artificial intelligence* (AI), in particular its latest developments in deep learning, has been influencing all kinds of disciplines such as engineering, business, agriculture, and humanities. More recently it also includes disciplines that were "reserved" to humans such as art and design. While there is a strong debate going on if creativity is profoundly human, we want to investigate if creativity can be enhanced by AI—not replaced. To be an inspiring co-creation partner by suggesting unexpected design variations and by learning the designer's taste. To do so we adopted AI algorithms, which can be trained by a small sample set of shapes of a given object, to propose novel shapes. The evaluation of our proposed methods revealed that it can be used by trained designers as well as non-designers to support the design process in different phases and that it could lead to novel designs not intended/foreseen. Besides the potentials of AI, we also point out and discuss moral threads caused by the latest developments in AI with respect to the creative sector.

Received on 18 February 2019; accepted on 30 March 2019; published on 26 April 2019

Keywords: inspirational AI, human-machine co-design, artificial neural network, genetic algorithm, design process

Copyright © 2019 Kevin German *et al.*, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution license (http://creativecommons.org/licenses/by/3.0/), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi:10.4108/eai.26-4-2019.162609

1. Introduction

The promises of AI assisted creation is "a world where creativity is highly accessible, through systems that empower us to create from new perspectives and raise the collective human potential" as Roelof Pieters and Samim Winiger pointed out. Recent developments in *artificial intelligence* (AI) have demonstrated that they are indeed capable to do things which in the past were restricted to humans. *Artificial neural networks* (ANN) and *genetic algorithms* (GA) are tools to make work easier for humans, for example through automatic speech translations (for instance simultaneous lecture translation has been demonstrated feasible already in 2008 by Kolss et al. [12]) or are even able to come up with solutions humans would never come up with effortlessly, see for instance the design of an "evolved antenna" using evolutionary algorithms published by Hornby et al. already in 2006 [8]. With further technological developments, of such processes there is a gradual transfer of competence from human beings to technical devices, namely, they serve as [24]:

1. tools: transfer of *mechanics* (material) from the human being to the device
2. machines: transfer of *energy* from the human being to the device
3. automatic machines[1]: transfer of *information* from the human being to the device
4. assistants: transfer of *decisions* from the human being to the device

We want to exemplify this concept with the field of mobility:

1. bicycle: feet are replaced by wheels
2. motor vehicle: propulsion is replaced by engine
3. self-driving rail vehicle: control is replaced by sensors and signal processing

[1]which is called *Automat* or *automate* in other languages such as German or French respectively

*Corresponding author. Email: Matthias@wolfel.de

4. autonomous vehicle: route planning or search for a parking space are replaced by artificial intelligence

Similarly, we can give an example from the field of art and design:

1. potter's kick-wheel: a tool used in the shaping of round ceramic ware driven by kicking a fly-wheel into motion
2. potter's electric-wheel: the kicking of the fly-wheel is replaced by a motor
3. construction & 3D printing: the object is constructed with a CAD-software according to given parameters and 3D printed
4. generated & 3D printing: the object is generated by an optimization process given particular constraints and 3D printed

In the coming years we are in the process of moving from Step 3. to Step 4. which raises—as it was the case from moving from Step 1. to Step 2. as well as from Step 2. to Step 3.—discussions, rejections, ethical issues (see Section 5), up to fears.

In the literature, some approaches to use AI in the creative sector have been presented. We review those approaches in the following section. Because the already introduced approaches are not available or were not fulfilling our requirements it was necessary to adopt given methods to intervene in the design process; either partially or in total. The investigated algorithms include genetic algorithms and different versions of ANNs namely convolutional neural networks, generative adversarial networks, and variational autoencoder. The developed algorithms can semi- or fully-automate the research, brainstorming and concept phase of the design process.

To evaluate and compare our different proposed approaches the entire development process was completed until the finished product for each approach. The approaches have been introduced within the School of Design at Pforzheim University, Germany and showcased to visitors of the Salone del Mobile in Milan, Italy, the Dutch Design Week in Eindhoven, Netherlands and the VDID Congress in Stuttgart, Germany. On these occasions, we were able to demonstrate that our proposed methods can be used by trained designers as well as non-designers to design semi-complex shapes with minimal user feedback.

2. Related Work

The idea of using algorithms to support the creation process is well established and frequently referred to as *generative design* or *procedural generation*. It is used to generate geometric patterns, textures, shapes, meshes, terrain or plants. The generation processes may include, but are not limited, to self-organization, swarm systems, ant colonies, evolutionary systems, fractal geometry, and generative grammars. McCormack et al. [15] review some generative design approaches and discuss how design as a discipline can benefit from those applications. While older approaches rely on generative algorithms which are usually realized by writing program code, ANN changes this process into data driven procedures. ANN can learn patterns from (labeled) examples or by reinforcement. Wölfel [25] points out that there are fundamental differences in the goals and reasons to use AI in art, design and cultural heritage: while in the former two AI should help to foster creativity and inspiration in the latter it should help to (re)discover or enhance given patterns; e.g. reconstruct part of an image which has been damaged over time. To create higher variations some artists randomly introduce glitches within the ANN. Due to the complex structure of the ANN these glitches (assuming that they are introduced at an early layer in the network) occur on a semantic level which causes the models to misinterpret the input data in interesting ways, some of which could be interpreted as glimpses of autonomous creativity; see for instance the artistic work 'Mistaken Identity' by Mario Klingemann.

AI or more precise ANNs has been introduced to support the creation process more recently. Leading software companies in engineering and design have already included AI-driven generative design paradigms which let humans input design goals. For instance, *Project Dreamcatcher* [2] is an engineering-based generative design program that takes into account how the forces will be directed best in the product and defines the best production method. Autodesk states the benefits of generative design to [1]:

- explore a wider range of design options
- make impossible designs possible
- optimize for materials and manufacturing methods

Most popular (at least in the mass media) are probably different variations of *image-to-image translation*. The most prominent example is *style transfer*—the capability to transfer the style of one image to draw the content of another. But mapping an input image to an output image is also possible for a variety of other applications such as *object transfiguration* (e.g. horse-to-zebra, apple-to-orange, *season transfer* (e.g. summer-to-winter) or *photo enhancement* [27]. While some of the just mentioned system seems to be toy applications, AI tools are taking over and gradually automate design processes which used to be time-consuming manual processes. Indeed, the most potential for AI in art and design is seen in its application to tedious, uncreative tasks such as coloring black-and-white images [26]. Cluzel et al. have proposed an interactive GA to progressively sketch the desired side-view of a car profile [3]. For this, the user has taken on the role of a fitness function[2] through interaction with the system. The *chAIr Project* [20] is a series of four chairs co-designed by AI and human designers. The project explores a collaborative creative process between humans and computers. It used a *generative adversarial network* (GAN) to propose new chairs which then have been 'interpreted' by trained designers to resemble a chair. *DeepWear* [9] is a method using deep convolutional GANs for clothes design. The GAN is trained on features of brand clothes and can generate images that are similar to actual clothes. A human interprets the generated images and tries to manually draw the corresponding pattern which is needed to make the finished product. Li et al. [14] introduced an ANN for encoding and synthesizing the structure of 3D shapes which—according to their findings—are effectively characterized by their hierarchical organization. Marco Kempf and Simon Zimmerman used AI in their work dubbed 'DeepWorld' to generate a compilation of 'artificial countries' using data of all existing countries (around 195) to generate new anthems, flags and other descriptors. Roman Lipski uses an *AI muse* (developed by Florian Dohmann et al.) to foster his/her inspiration. Because the AI muse is trained only on the artists previous drawings and fed with the current work in progress it suggests image variations in line with Romans taste.

Most of the related work is not ready yet to be used without a thorough understanding of the technology and is more an engineering approach using ANNs instead of common technology. Daniel Wikström [23] mentions that many designers do not yet know technology well enough and therefore perceive it as "magic". But he also explains how an intelligent assistant is perceived and would have to interact. What we are aiming for is different: The whole creation process—not its development—should be applicable to naive users without any profound understanding of design or engineering. The user has to only rely on his/her taste to cherry-pick examples he/she likes in an iterative process until he/she ends up with the final design.

3. Co-Creating Shapes with Artificial Intelligence

In this section we want to introduce different approaches how to use AI as a co-creation partner. Before we can start to do so, however, we first have to develop an approach which is able to suggest plausible shapes of bottles. We chose bottles because they have simple shapes and thus keep the necessary training effort low, but still are able to express unique features which can easily be recognized; just think of the iconic contour fluted lines of the Coca-Cola bottle.

3.1. Plausible Shape Representation

A "naive" approach to automatically generate bottle shapes would be to start by drawing randomly placed polygons or polylines and optimize it by targeted selection, mating, recombination and mutation by optimizing to a particular goal; e.g. has to look like a bottle. To determine if an image 'looks like a bottle' pre-trained classifiers could be used (see e.g. Krizhevsky et al. [13]). However, as has been pointed out by German et al. [4] this "naive" approach is not leading to satisfying results. The problem here is that not only shapes which are similar to bottles get a high score, but also shapes with random patterns are able to get similar scores. This is a problem, known as *adversarial examples* [6], and not uncommon in ANNs.

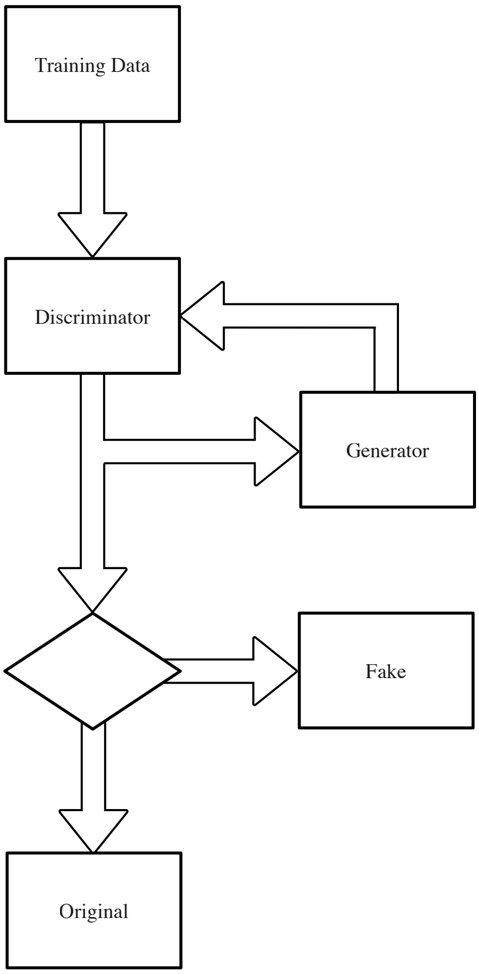

Figure 1. Flow chart of generative adversarial network and different instances according to the different steps.

[2]also referred to as objective function

Therefore, an approach is required which guarantees that the produced shapes are similar to the shape of bottles. In 2014 Goodfellow et al. proposed the special ANN architecture GAN which we have already mentioned before [5]. The main idea of their proposal is to use two ANNs that compete with each other. Fig. 1 demonstrates the basic principle and components: The *generator* tries to generate data from latent variables that are as similar as possible to the training data. The *discriminator* tries to classify the generated data according to the original training data. Both networks play a zero-sum game: As the system progresses, the generator as well as the discriminator are improving. This process continues until the discriminator can no longer distinguish between forgery and original. This is achieved when the discriminator is only correct in 50% of the cases.

Since the generator learns to generate data as similar as possible to the training data, it requires a training data set that corresponds as closely as possible to the desired output [5]. In our case we were interested in generating different variations of shapes resembling bottles. To train our system we converted 200 images of bottles into black and white silhouettes (see Fig. 2). As automatic segmentation did not lead to satisfactory results the conversation was done by hand. Since the data volume is small and GANs normally use data volumes in orders of magnitude of several 1,000 images, there is a risk of over-adaptation by the GAN [22]. To reduce over-adaptation, *data augmentation* is used by automatically generating variations of the available training data including shearing, enlarging, rotating and cropping. In order to further reduce the overfitting of all the methods presented here, a considerable amount of regularization [17] and dropout layers [21] were used.

Fig. 3 shows that the training loss in the first few generations quickly approaches zero. This is due to the fact that the network initially roughly maps the basic form of the input data. In higher epochs many bottles of an epoch have similar characteristics. This is a well-known problem in GAN architectures and is called *mode collapse*. The generator limits itself to generating only a few examples that the discriminator classifies as original. In the worst case, all images generated by the generator are almost identical [16]. Although in our example we see variations the problem is still visible. Different epochs can be considered to create more diverse bottles because the point of mode collapse shifts with each epoch. Although the training data set only consists of symmetrical bottles, the architecture is capable of generating asymmetric bottles. This is interesting because the net is able to generate something it did not know could be e.g. asymmetric. It is up to the designer to incorporate these unusual features such as asymmetrical elements into the product design or to rate them as a mistake and to correct them manually based on his/her taste.

Due to the required minimum complexity of the GAN architectures and the need for sharp high-resolution images in combination with the low amount of training data, overfitting inevitably occurs. However, subjective comparisons with the training data set did not rate the over-adaption as critical as the majority of the bottles are unique. Instead of treating the shape as one union it might be advantageous to separate the shape into different parts.

3.2. Semantic Shape Representation

The shape of an object can be decomposed into different features that can be assigned with particular “meanings” and semantically annotated[3]. In our particular application of a bottle the semantic shape representation can be separated and annotated into: lid, neck, wall, wall-to-neck transition and bottom[4]. The classification was done manually by cutting the existing 200 images into individual parts. In the future, however, this step can be automated using image segmentation.

One conceivable option for creating new shapes of bottles is the random permutation of the semantic parts and thus to overcome the limiting characteristics of the former approach where many generated bottles had similar characteristics. For this purpose, an ANN is to be conceptualized, which receives random features and assembles them to form a new object. The network learned, in the training phase, the relationship between the semantic features and the actual bottle. After this phase, the network is able to merge features seamlessly and to produce the shape of a consistent bottle. New permutations of features using the trained ANN are shown below in Fig. 4. The features were determined based on a discrete equal distribution. It can be observed that the features are transferred and combined successfully.

3.3. Introducing Personal Taste in Shape Representation

So far we have described the process of how to fully automatically generate plausible shapes by varying different features of the bottle. Now it’s time to bring back the user by having him/her intervene in the design process: The shape should advance iteratively towards the taste of the user. For optimization problems in which a solution approaches an optimum step by step,

[3]Semantic annotation is the process of attaching additional information to various concepts to be used by machines.

[4]In preliminary tests, this division turned out to be the most effective variant.

Figure 2. Black and white silhouettes of bottles.

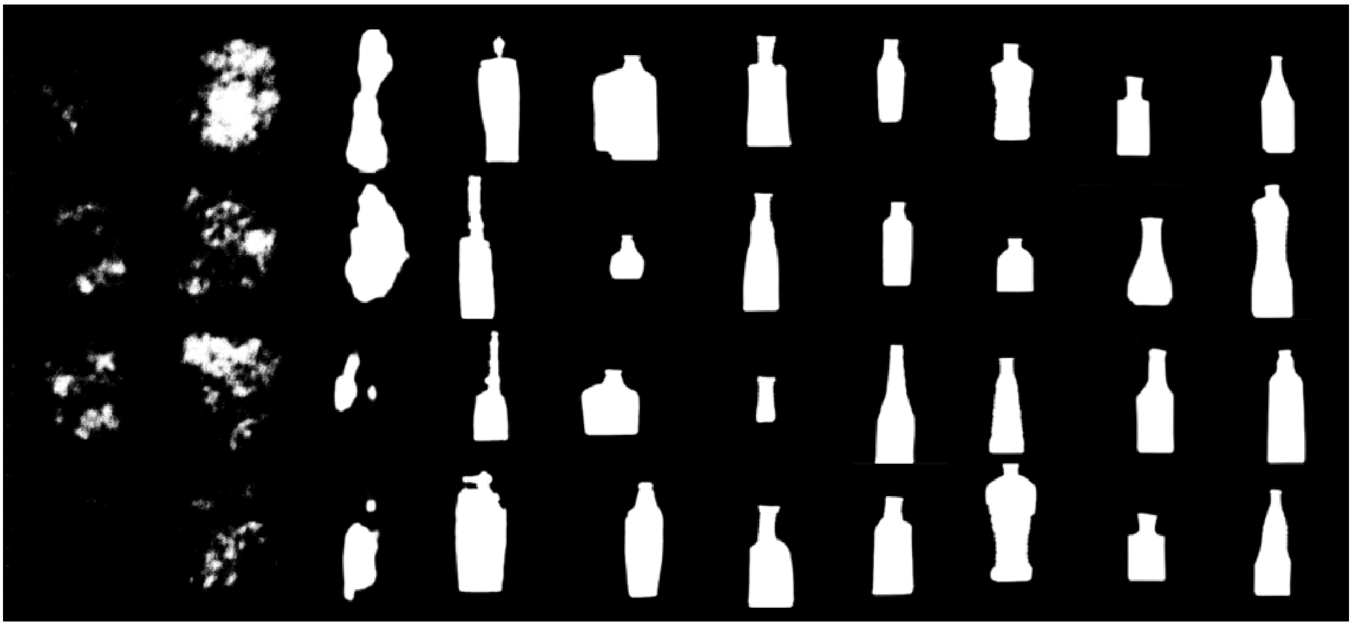

Figure 3. Different iterations of the learning process. From left to right, iteration 50, 100, steps of 100 until 1000. Four different examples are shown for each iteration.

GA has already proven to be an appropriate tool [8], which is also why a GA was used in this procedure. The basic idea is that you have a population of objects where each object is defined by its genes. Each gene represents a semantic feature, in this case, e.g., the bottleneck. To transform the genes into visible features, the ANN of the semantic shape representation is used. Another valid possibility that was not considered here would be to use the genes as an input vector for a generator or variational autoencoder to express the visible shape.

Similar to the biological model, the population gradually adapts to the environment through selection, mating, gene recombination and mutation [7]. To introduce the designer into the automatic algorithm the random permutations of features have to be evaluated by the designer instead of a genetic objective function. Therefore, the designer takes up the position of the fitness/objective function, similar to the ANN MobileNet, by sitting in front of the computer and by evaluating each instance individually; Fig. 5. The basic idea here is that the population gradually approaches the taste of the user until his/her ideal bottle is created. Therefore, each of the 20 individuals in the population is assigned a fitness value between zero and one by the user. The higher the fitness value the higher the probability of survival by an individual. Combined with the previously mentioned methods such as mutation, this results in a population which is more precisely adapted to the taste of the user. To cover a large search space, the population is initialized using a discrete equal distribution. Over a couple of iterations the final optimal shape is found.

3.4. Democratizing Shape Representation

To be able to democratize the design process we have to vary the proposed approaches so far to be able to do

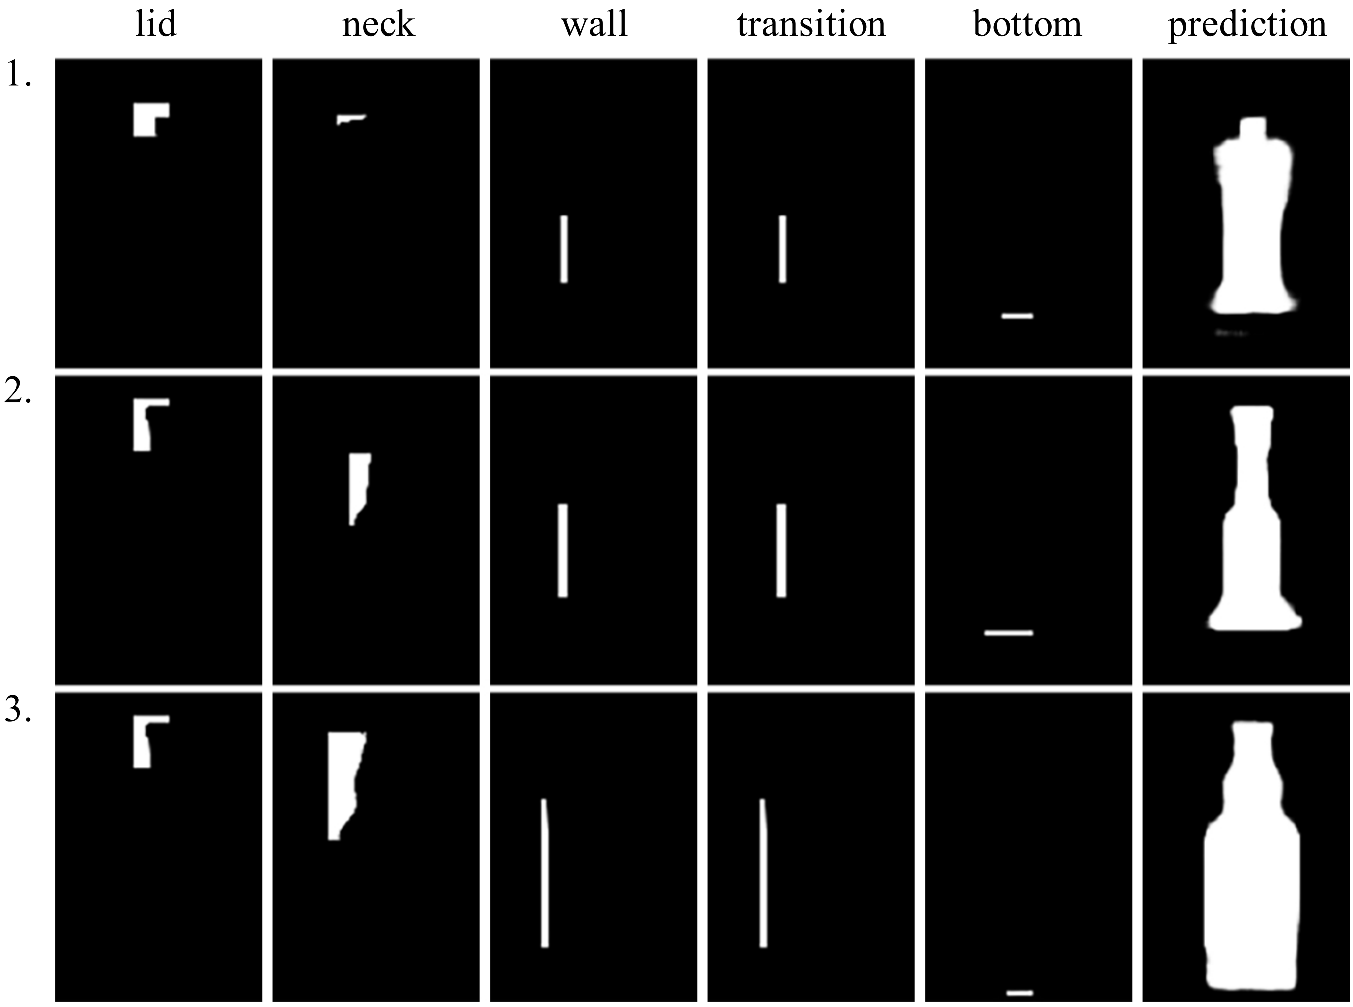

Figure 4. Three variations of bottle shapes as generated by merging the decomposed parts as given by the semantic shape representation approach.

some arithmetic's; e.g. to calculate the arithmetic mean of a set of bottles designed by different persons. To do so we use a *variational autoencoder* (VAE) [10]. It is an ANN that learns to produce the same output as input. A special feature here is that the network topology has a bottleneck between the input and the output layers. This bottleneck stores the compressed information as a vector of real numbers called *latent variables* (LV). As a result, the autoencoder must compress information of the input into the LV and then decompress it after the bottleneck. The VAE learns to extract the most relevant information from an input image as LV so that it can be used to regenerate the output image as correctly as possible [10].

The basic idea is this: After training, the LV can be accessed directly through sliders. The trained decoder would then convert the LV into a corresponding bottle. This would allow the non-designer with limited design skills to design an object in a playful way (Fig. 6). A number of eight LV have delivered satisfactory results in trials. A smaller number of LV leads to less detailed and more similar images. More LV, on the other hand, have not achieved any significant improvement in quality, but have worsened the user experience due to more necessary sliders.

It can be seen that by moving individual sliders, the bottle can be transferred into other forms. The transformation is done simultaneously with the slider movement, giving the user direct and intuitive feedback. A complete disentanglement of the LV could not be achieved. Consequently, a LV and thus the corresponding slider can be responsible for several semantic features of the object.

Because there are vectors behind the bottles, we can do bottle arithmetic with them [19]. This makes it possible to calculate the arithmetic mean of a set of bottles. This allows several individuals to democratically design a bottle by first creating a bottle for each individual using the sliders and then averaging all created bottles. There are two main pillars of

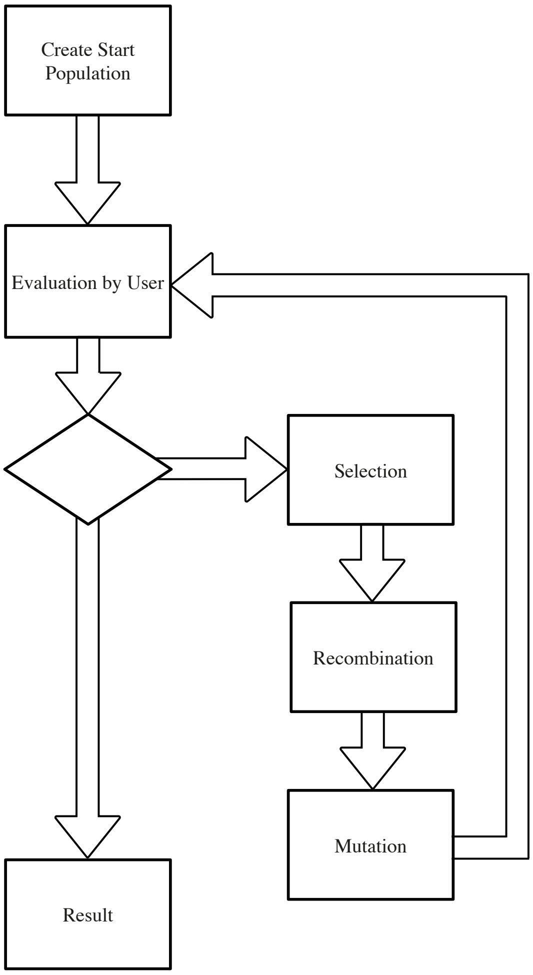

Figure 5. Flow chart of the genetic algorithm and different instances according to the different steps.

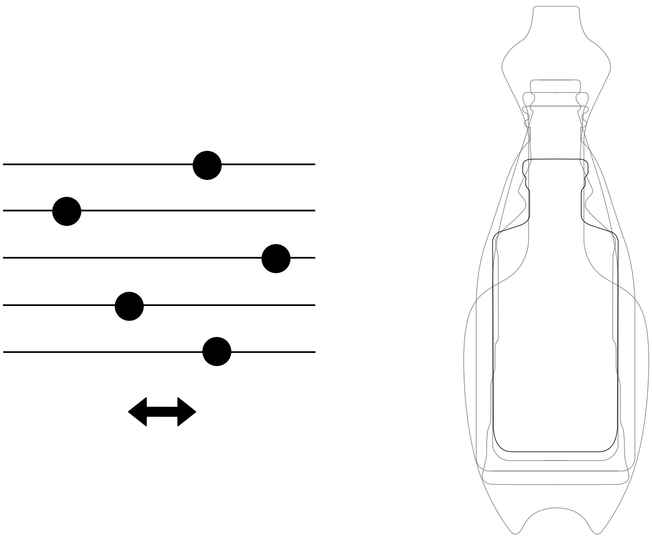

Figure 6. Transforming the bottles using eight parameters. Each slider corresponds to one latent variables.

democratic design. First, anyone can design objects now even without design skills and secondly, the taste of each individual can equally influence the final product.

4. Results, Evaluation & Limitations

Using the *plausible shape representation* (Section 3.1) method, it was shown that parts of the design process can be partially automated and thus speed up using ANNs. This architecture typically provides a good image quality. However, the algorithm does not allow direct access by the designer, so the output is heavily dependent on the training data. For instance, to specifically design a classic beer bottle, the designer would have to explicitly look for the shape in the output or to use only beer bottles as training data. Although novel bottle shapes are created, these usually do not deviate much from the training data set. For the design process, the user has received some suggestions from the algorithms and has decided on one of these in several iterations; see Fig. 7a and 8a.

Using the *semantic shape representation* (Section 3.2), new bottles could also be created automatically. In comparison to the previous method, these objects are more diverse and creative looking; see Fig. 7b and 8b. At present, there must be a database of the specific objects and their associated features available, which is not ideal. The image quality is slightly worse than in the plausible shape representation, but still at a very high level. In addition, as with the plausible shape representation, the problem is that the suggestions are not adapted to the user.

To tackle the latter problem, *personal taste* (Section 3.3) was introduced into the semantic shape representation. The bottles successfully adapted to the taste of the user through evolution; see Fig. 7c and 8c. A selection from a large amount of output data as in the last two algorithms is thereby eliminated (apart from the fitness score evaluation). In our opinion this is one of the most promising ways to liberate design processes in the future because designing personalized objects according to his/her taste becomes possible for everybody. The algorithm also adapts through the direct feedback dynamically to changes in the user's taste, for instance, during a lifetime. Since the architecture is based on the semantic shape representation, the image quality is at the same level and a database of associated features is also needed.

Through the *democratizing shape representation* (Section 3.4) method, see Fig. 7d and 8d, collectives can design objects together. With the introduction of variable parameters (sliders), every human being is able to design things, whether talented or not. This bypasses the designer and allows the end-user to take on the role of a designer directly. Secondly, the opinion of each individual can be incorporated into a final product.

Figure 7. 3D print of generated bottles using a. plausible shape representation, b. semantic shape representation, c. personal taste in shape representation, and d. democratized shape representation

There's no need for a central design instance anymore. The zeitgeist of the collective can (anonymously) create something together, on which the majority can agree on. Also, the manual sketches of the concept phase were eliminated. Within a few seconds, countless new variants could be created, for which otherwise individual manual sketches would be needed. However, the image quality and diversity are worse compared to the previous algorithms.

In Table 1 we compare the different approaches according to the parameters described next:

- *Affordance* (in data preparation) describes how much time has to be spent to prepare the data to train the ANN.

- *Automation* describes how much the process is automated and how much amount has to be done by the designer.

Figure 8. Rendering of generated bottles using a. plausible shape representation, b. semantic shape representation, c. personal taste in shape representation, and d. democratized shape representation

Table 1. Comparison of the different methods presented here.

	plausible shape	semantic shape	personal taste	democratic approach
Affordance	medium	high	medium	medium
Automation	full	full	semi	semi
Shape quality	very good	good	good	medium
Creativity	medium	very good	very good	medium
Personalization	low	low	high	medium

- *Shape quality* describes the subjective quality of the shape including detail density, image sharpness, resolution and number of image artifacts.

- *Creativity* describes to what extent the automatically generated results have a creative or inspiring effect on the designer.

- *Personalization* describes how much individuality is kept in the design process and how much of the personal taste is represented in the outcome.

As previously mentioned all variants shown here were trained with a well-defined data set consisting of 200 relatively simple 2D images. This procedure was sufficient to analyze the process. If the same procedures can be applied to more complex shapes and higher dimensionality is unclear because these variants might encounter additional problems. A possible solution in the future would be the use of voxels or a polygon mesh, which allows a 3D representation. However, experience shows that the necessary amount of training data increases with increasing complexity. A manually created data set is therefore no longer a valid option.

With automatically created 2D data sets, e.g. by web scraping, this leads to problems because these images often have a lack of quality for this application, for instance by having other objects in the image or through image artifacts (which is however desired for image classifications due to better generalization). For 3D objects, this is not to this extent the case, e.g. CAD files in most cases only depict the desired object. To get this data, there are already large databases that have high quality [11]. Because CAD is an industry-standard, companies can also use their existing data-sets. The disadvantage of the increased complexity due to the 3D representation can potentially be partly compensated by the high quality and quantity of the training data.

5. Moral Threads

For decades, AI has fostered (often false) future visions ranging from transhumanist utopia to "world run by machines" dystopia. Artists and designers explore solutions concerning the semiotic, the aesthetic and the dynamic realm, as well as confronting corporate, industrial, cultural and political aspects. The relationship between the artist and the artwork is directly connected through their intentions, although currently mediated by third-parties and media tools. Understanding ethical and social implications, in particular possible threads, of AI assisted creation is becoming a pressing need and include:

- *Wrong Expectations*: Only "working examples" are demonstrated in the media, therefore wrong expectations are raised. A lot of content claiming to be AI has indeed been produced by methods not containing AI (not only in the creative community). Wrong expectations are leading to worries about: design AI tools that replace us or design AI tools that shape us after we shape them (adapted from Marshall McLuhan)

- *Data Bias*: AI systems are sensitive to bias. As a consequence, the AI is not being a neutral tool, but has pre-decoded preferences. Bias relevant in creative AI systems are: algorithmic bias occurs when a computer system reflects the implicit values of the humans who created it; data bias occurs when your samples aren't representative of your population of interest; prejudice bias results out of cultural influences or stereotypes.

- *Authorship*: The authorship of AI generated content has not been clarified. For instance, is the authorship of a novel song composed by an AI trained exclusively on songs by Johann Sebastian Bach belonging to the AI, the developer, or Bach? See e.g. [18] for a more detailed discussion.

- *Replacement*: Do we still need a designer in times of AI and automation? Not only is this the first question that crosses the minds of non-designers, but it is an even more important question for the design world. Designers are not the only ones to feel the thread of AI. For instance, translators are concerned that they could be replaced through machine translation and truck drivers fear to lose their jobs because of autonomous driving.

6. Conclusion & Outlook

In this work, we set out to prove that most of the creation process could be automated or at least semi-automated and that a workflow from the first sketches to the final product can be supported by AI. This became possible by generating design proposals from different algorithms including ANN and GA. This drastically accelerated the creation process and saved tedious labor time. The potential of AI in creativity has just been started to be explored: AI is shifting the creativity process from crafting to generating and selecting; AI has a high potential in the creative sector it can lower the time between intention and realization; it can potentially lead to the democratization of creativity.

We chose a simple object—a bottle—to prove our concept. Any other object could, in principle, be designed the same way. It should be also possible to extend our proposed approach to include a third dimension. More complex shapes and higher dimensionality, however, raises complexity and therefore more data and other solutions might need to be introduced.

We live in an era of accelerating technological progress which is already influencing our daily lives. We cannot ignore technological developments and pretend these changes are not happening. Instead, we should embrace the development—but also reflect its impact—and see it as a new set of opportunities for us to explore and prosper. Widespread misuse (threads) can limit the social acceptance and requires an AI literacy—just like digital literacy—for everybody. We have to reflect on what makes us human and remember that we are still the ones who are conceiving something that we think of as beautiful and therefore value it. "Successful designs are not necessarily 'made': new functionality may 'evolve' through the use and interpretation of artifacts by an audience" [15]. There are many examples today where AI has influenced the creative process letting the designer cherry-pick and approve adjustments based on the proposed variations. Let us start exploring these possibilities today and see where they can take us.

References

[1] AutodeskResearch (2019) Generative design. URL https://www.autodesk.com/solutions/generative-design. Accessed: 2019-11-27.

[2] AutodeskResearch (2019) Project dreamcatcher. URL https://autodeskresearch.com/projects/dreamcatcher. Accessed: 2019-11-27.

[3] Cluzel, F., Yannou, B. and Dihlmann, M. (2012) Using evolutionary design to interactively sketch car silhouettes and stimulate designer's creativity. *Engineering Applications of Artificial Intelligence* **25**(7): 1413–1424.

[4] German, K., Limm, M., Wölfel, M. and Helmerdig, S. (2019) Co-designing object shapes with artificial intelligence. In *Proceedings of ArtsIT*.

[5] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. *et al.* (2014) Generative adversarial nets. In *Advances in neural information processing systems*: 2672–2680.

[6] Goodfellow, I.J., Shlens, J. and Szegedy, C. (2015) Explaining and harnessing adversarial examples.

[7] Gupta, D. and Ghafir, S. (2012) An overview of methods maintaining diversity in genetic algorithms. *International journal of emerging technology and advanced engineering* 2(5): 56–60.

[8] Hornby, G., Globus, A., Linden, D. and Lohn, J. (2006) Automated antenna design with evolutionary algorithms. In *American Institute of Aeronautics and Astronautics Conference on Space, San Jose, CA*: 19–21.

[9] Kato, N., Osone, H., Sato, D., Muramatsu, N. and Ochiai, Y. (2018) Deepwear: a case study of collaborative design between human and artificial intelligence. In *Twelfth International Conference on Tangible, Embedded, and Embodied Interaction* (ACM): 529–536.

[10] Kingma, D.P. and Welling, M. (2013) Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* .

[11] Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M. *et al.* (2019) ABC: A big CAD model dataset for geometric deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*: 9601–9611.

[12] Kolss, M., Wölfel, M., Kraft, F., Niehues, J., Paulik, M. and Waibel, A. (2008) Simultaneous german-english lecture translation. In *International Workshop on Spoken Language Translation*.

[13] Krizhevsky, A., Sutskever, I. and Hinton, G.E. (2012) Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*: 1097–1105.

[14] Li, J., Xu, K., Chaudhuri, S., Yumer, E., Zhang, H. and Guibas, L. (2017) Grass: Generative recursive autoencoders for shape structures. *Transactions on Graphics (TOG)* **36**(4): 52.

[15] McCormack, J., Dorin, A., Innocent, T. *et al.* (2004) Generative design: a paradigm for design research. *Proceedings of Futureground, Design Research Society, Melbourne* .

[16] Metz, L., Poole, B., Pfau, D. and Sohl-Dickstein, J. (2017) Unrolled generative adversarial networks. In *Proceedings of 5th International Conference on Learning Representations*.

[17] Ng, A.Y. (2004) Feature selection, l 1 vs. l 2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning* (ACM): 78.

[18] Rack, F. and Vettermann, O. (2019) KI-Kunst und Urheberrecht – die Maschine als Schöpferin? URL http://tlmd.in/a/3353. Accessed: 2019-11-27.

[19] Radford, A., Metz, L. and Chintala, S. (2015) Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* .

[20] Schmitt, P. and Weiss, S. (2018) The chair project—four classics. URL https://philippschmitt.com/work/chair. Accessed: 2019-11-27.

[21] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. and Salakhutdinov, R. (2014) Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1): 1929–1958.

[22] Wang, J. and Perez, L. (2017) The effectiveness of data augmentation in image classification using deep learning. *Convolutional Neural Networks Vis. Recognit* .

[23] Wikström, D. (2018) Me, myself, and AI: Case study: human-machine co-creation explored in design.

[24] Wölfel, M. (2018) Der smarte Assistent. *In: Smartphone-Ästhetik: zur Philosophie und Gestaltung mobiler Medien. Ruf, O. (editor). Transcript* : 269–288.

[25] Wölfel, M. (2020) Bringing cultural heritage to life through the lens of artificial intelligence. In *Proceedings of HCI International*.

[26] Zhang, G., Qu, M., Jin, Y. and Song, Q. (2019) Colorization for anime sketches with cycle-consistent adversarial network. *International Journal of Performability Engineering* **15**(3).

[27] Zhu, J.Y., Park, T., Isola, P. and Efros, A.A. (2017) Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*.