Marketing Science Institute Working Paper Series 2024

Report No. 24-103

Does That Car Want to Give Me a Ride? Bio-Inspired Automotive Aesthetic Design

Bowei Chen, Jingmin Huang, Mengxia Zhang, and Lan Luo

Does That Car Want to Give Me a Ride? Bio-Inspired Automotive Aesthetic Design

Bowei Chen, Jingmin Huang, Mengxia Zhang, and Lan Luo*

September, 2023

*Bowei Chen (e-mail: bowei.chen@glasgow.ac.uk) is an Associate Professor of Marketing Analytics and Data Science at the Adam Smith Business School of the University of Glasgow; Jingmin Huang (e-mail: jingmin.huang@glasgow.ac.uk) is a PhD candidate at the School of Computing Science of the University of Glasgow; Mengxia Zhang (e-mail: mezhang@ivey.ca) is an Assistant Professor of Marketing at the Ivey Business School of the Western University; Lan Luo (e-mail: lluo@marshall.usc.edu) is a Professor of Marketing at the Marshall School of Business of the University of Southern California. We thank participants at the INFORMS Marketing Science Conference, Carnegie Mellon University Conference on Digital Marketing and Machine Learning, China Marketing International Conference, Temple University Global Institute for Artificial Intelligence & Business Analytics Workshop Speakers Series, Triennial Invitational Choice Symposium, and University College Dublin Marketing Research Symposium. We also thank Nvidia Corporation for supporting this research through the Academic GPU Grant and the Accelerated Data Science Grant.

Does That Car Want to Give Me a Ride? Bio-Inspired Automotive Aesthetic Design

Abstract

Product aesthetic design is pivotal in shaping consumer evaluations, especially in automotive industry where captivating designs often secure a market advantage. Consistent with prior research on knowledge transfer and schema congruity, we discover that consumers attribute faster acceleration and power to a sedan when primed to believe that its design is inspired by cheetah, the world's fastest land animal. Building upon this insight, we propose a deep-learning-based computational framework to morph aesthetic features of the product domain (e.g., sports sedans) with features from a different domain (e.g., sleek body curves of running cheetahs) in a mid-generational automotive facelift. We discover that cheetah-inspired new automobile facelifts are consistently preferred over the original looks, suggesting that our framework allows automakers to elevate automobiles' visual appeal cost-effectively without overhauling their fundamental body structures. Furthermore, our research suggests that higher degrees of cheetah morphing are preferred for premium sedans, extra sporty models, and by high-income consumers, compared to their counterparts. Lastly, we learn that certain colors (e.g., red) or print patterns (e.g., cheetah prints) further boost the visual appeal of cheetah curve-inspired automotive facelifts. Our framework serves as a proof-of-concept of using AI to morph between two distinct subject domains for product aesthetic design.

Keywords: Automotive Aesthetics, Bio-Inspired Design, Facelift, Schema, Knowledge Transfer, Deep Learning, Generative AI

1 Introduction

> *Nature is incredibly inspiring, sharks, cats... it is the animals that have an appealing aesthetic that we often look at in our cars.*
>
> Mark Fetherston, Team Leader of Exterior Design, Mercedes-Benz

> *A car is not simply a mass of metal. Mazda believes it is more like a living creature.*
>
> Ikuo Maeda, Head of Global Design, Mazda

> *... because we were looking to establish McLaren as a very unique looking automobile, the secret I think would be to look back into nature for inspiration.*
>
> Frank Stephenson, automobile designer[a]

[a]Frank Stephenson has worked for BMW, Mini, Fiat, Ferrari and McLaren. Motor Trend magazine called him "one of the most influential automotive designers of our time".

It has been long recognized that automotive aesthetic design plays an integral role in an automobile's market success. According to a market report of JD Power (National Automobile Dealers Association 2016), aesthetic design has been considered as the top purchase motive of consumers. The German automaker Audi finds that exterior aesthetics can explain up to 60% of consumers' purchase decisions (Kreuzbauer and Malter 2005). Therefore, automobile designers have always been making tremendous efforts to develop appealing aesthetics for their consumers.

Within this context, not surprisingly, automobile designers have been turning to nature and animals for design inspirations (Bouchard et al. 2009). Today, many automobile models are named after animals, such as Jaguar, Ford Mustang, Ford Ranger Raptor, AMC Eagle, VW Beetle, Seat Leon, Dodge Viper, and so on.[1] Animal motifs are also often used on automobile logos like Ferrari, Lamborghini, Porsche, Alfa Romeo, and Peugeot.[2] Furthermore, some well-known automobile models are designed by directly incorporating certain characteristics of animals. Such practices are called bio-inspired designs (Aziz et al. 2016). For example, the McLaren P1 was created with inspiration from sailfish[3]; Kia adopted a tiger nose grille as its family signature for all of its

[1]`https://www.zeroto60times.com/2013/03/complete-list-cars-animal-names`.

[2]`https://www.carlogos.org/reviews/car-logos-with-animals.html`.

[3]The McLaren P1 is a limited-production mid-engine plug-in hybrid sports sedan produced by British automobile

car models[4]; and Mazda's KODO design uses Japanese aesthetics combined with animal-inspired forms[5].

Indeed, the psychology and consumer behavior literature on knowledge transfer (Gregan-Paxton and John 1997, Jia et al. 2020) and schema congruity (Fiske 1982, Aggarwal and McGill 2007) provides a solid theoretical foundation for these bio-inspired design practices. Both streams of literature suggest that consumers can transfer their knowledge from the base domains to novel instances/domains where the activated schema (e.g., baby schema) can determine the affective responses and evaluations of an instance perceived to fit the schema (e.g., car front with enlarged headlights/eyes). For example, prior marketing research on anthropomorphism finds that the human tendency to see the face schema in car fronts can affect consumer perception and evaluation towards the focal product (Landwehr et al. 2011b, Maeng and Aggarwal 2017). Kim et al. (2011) demonstrate that consumers' semantic and emotional responses to bio-inspired vehicle designs can be strongly correlated to the semantics (e.g., rapid speed, sportive, etc.) and emotions (e.g., positive, arousal, etc.) expressed by the inspirational source (e.g., an animal posture).

Although theoretically attractive, bio-inspired automotive design is costly to implement traditionally. Prior bio-inspired automotive design endeavors have been mostly driven by automobile designers manually drawing inspirations from nature, then implementing such inspirations via the aid of computer-aided design software. As such, the design generation process can be both labor-intensive and time-consuming. For example, General Motors designers often have to manually create several hundred sketches per new vehicle (Burnap et al. 2023).

In this paper, we propose a deep learning based framework to automatically generate candidate bio-inspired automotive aesthetic designs at scale. Compared to historical approaches commonly used in bio-inspired automobile design, our proposed framework can be readily automated, hence requiring fewer manual inputs from designers and making it easier to scale up. Using our framework, automobile designers can easily generate and screen through hundreds of possible design candidates.

In particular, our research focuses on the aesthetic design of the automotive *factlift* (also called

manufacturer McLaren Automotive. According to some tests, it is one of the fastest hypercars in the world, and it is among the most expensive street-legal road cars that money can buy (McLaren Automotive 2021).

[4] `https://www.kia.com/au/util/news/tiger-nose-grille-10-year-anniversary.html`.

[5] `https://www.mazda.com/en/innovation/design/`.

the *mid-generational refresh*), with the aim to effectively improve the aesthetic appeal of automobiles with bio-inspired features without dramatically altering their overall body lines and shapes. Such facelift often entails cosmetic changes (exterior or interior styling), new accessories, and/or safety options in between automobile generations.[6] The concept of automobile mid-generational facelift was first introduced by Alfred Sloan at General Motors in the 1920s and was soon adopted by a majority of automobile manufacturers (Schwartz 2000). There are two main reasons that contribute to the popularity of facelifts in this industry. First, a facelift can effectively boost consumer interest in an automobile model that has passed the first flush of youth but still has some years before a redesign. Second, a facelift is always considerably less expensive than a redesign (Blonigen et al. 2017). The latter is a full ground-up rework of the automobile, often accompanied by an entirely new look, many new features, a new interior, and a new powertrain.[7] Most automakers launch the next generation of a model every six to eight years but release the facelifted models annually or biennially.[8] Despite the practical importance and cost effectiveness of facelifts and the high frequency of such efforts, automotive facelift design has been highly under-researched. Our research is among the first empirical attempts in marketing with the aim to improve perceived aesthetics and aerodynamics from the consumer's perspective for mid-generational refreshes of automobiles.

More specifically, as a proof of concept, we demonstrate the proposed framework by facelifting the side-view design of sports sedans using features inspired by cheetahs. The general framework can apply to other domains (e.g., facelifting SUVs with rhino features). Sports sedans are a large and increasingly popular type of sedans that many marketers are interested in [9] and often feature a sporty, slim, and stylish look, which motivates us to turn to the world's fastest land animal cheetah for design inspiration.[10] In fact, Bill Thomas gained support from General Motors in 1963 to develop the concept car model Cheetah. It later became the Super Cheetah project at Chevrolet for Le Mans sports car race. Not surprisingly, cheetahs have always been a natural inspiration for

[6] https://en.wikipedia.org/wiki/Facelift_(automotive).

[7] https://bit.ly/3v1p79L.

[8] https://www.carwow.co.uk/guides/glossary/car-facelift-explained-0658

[9] https://www.supercars.net/blog/the-allure-of-super-sedans-and-their-popularity/

[10] https://en.wikipedia.org/wiki/Cheetah.

sedans' aesthetic design.[11] We hypothesize that using cheetahs as an inspiration for automotive aesthetic design might enhance consumers' evaluation of sports sedans. As discovered in a pretest (Section 4.1), when participants were primed to believe that a sedan's design was inspired by a cheetah, they perceived the identical sedan to be faster, more exotic, and more powerful.

We decide to focus on automobile side-view facelifts under the following two considerations. First, despite the importance of side views in automotive design (Buxton et al. 2000, Hucho and Sovran 1993) and brand identity (Karjalainen 2003, Catalano et al. 2007), enhancing the side view of an automobile has been under-researched in the existing marketing literature. To date, marketing researchers have emphasized either car front design (Aggarwal and McGill 2007, Landwehr et al. 2011a,b, Maeng and Aggarwal 2017, Miesler et al. 2011, Purucker et al. 2014) or complete redesign (Burnap et al. 2023). Second, if we were to capture the running cheetah's swift and consecutive movements and integrate these formed inspirations into the aesthetic design, a side-view design would be the most relevant initial attempt.

Figure 1: Side Views of the Recent Three Generations of BMW 3 Series

When it comes to side-view design, one particularly important feature is the exterior styling of shoulder and waist curves. Figure 1 provides an example of side-view designs of the three most recent generations of BMW 3 Series.[12] We observe the two main changes when comparing the side

[11] http://billthomascheetah.com/aluminum-super-cheetah.

[12] https://www.bmw-me.com/en/all-models/3-series.html.

views of these three models. First, the shoulder curve (or *bone line*[13]) – the line on the upper door panels that starts from the front fender, continues under the window area, and extends to the rear fender – becomes shorter, which makes the car look sportier and more dynamic. Second, the waist curve – the character line on or below the lower door panels that runs across the area between two wheels – becomes curvier and has a larger slope, also enhancing the sporty appeal of the model. We also collect additional examples of sporty sedan designs in Figure 2. The examples show that both mainstream (e.g., Hyundai; Toyota) and premium automotive brands (e.g., Prosche; Maserati) have recent models with sporty and curvy shoulder and waist curves. In response to such an emerging trend in the automobile industry, our study aims to use deep learning techniques to modify the shoulder and waist curves of existing sports sedans using the inspirations learned from running cheetahs. This endeavor is perfectly suitable for our goal of facelift design because such modifications allow automobile designers to alter the aesthetics of the car while retaining its basic exterior styling and platform chassis. In a nutshell, it is our belief that such aesthetic alterations will be much less costly to materialize compared to an entire redesign, yet still have the potential to considerably enhance consumers' perceived aesthetics and aerodynamics towards the focal automobile.

Our proposed design framework consists of three modules. The first module analyzes cheetah's movement and identifies curves that best capture its body dynamics. The second module processes sedan side-view images and extracts shoulder and waist curves for each sedan model in our study. The third module morphs the curves of the sedan with the learned curves of the running cheetah to generate new bio-inspired exterior designs for each target sedan model. The following three deep learning algorithms are used, tailored, and integrated into our framework: 1) the mask region-based convolutional neural network (Mask R-CNN) (He et al. 2017), 2) the holistically-nested edge detection (HED) algorithm (Xie and Tu 2017), and 3) the cycle consistency generative adversarial network (CycleGAN) (Zhu et al. 2017). The background and technical details about these algorithms will be reviewed in Sections 2-3.

We apply our proposed framework to 13 sports sedan models based on market popularity. We further recruit sedan intenders to evaluate the new designs. We demonstrate that our framework can

[13] `https://www.youtube.com/watch?v=haiMna8NBQs`.

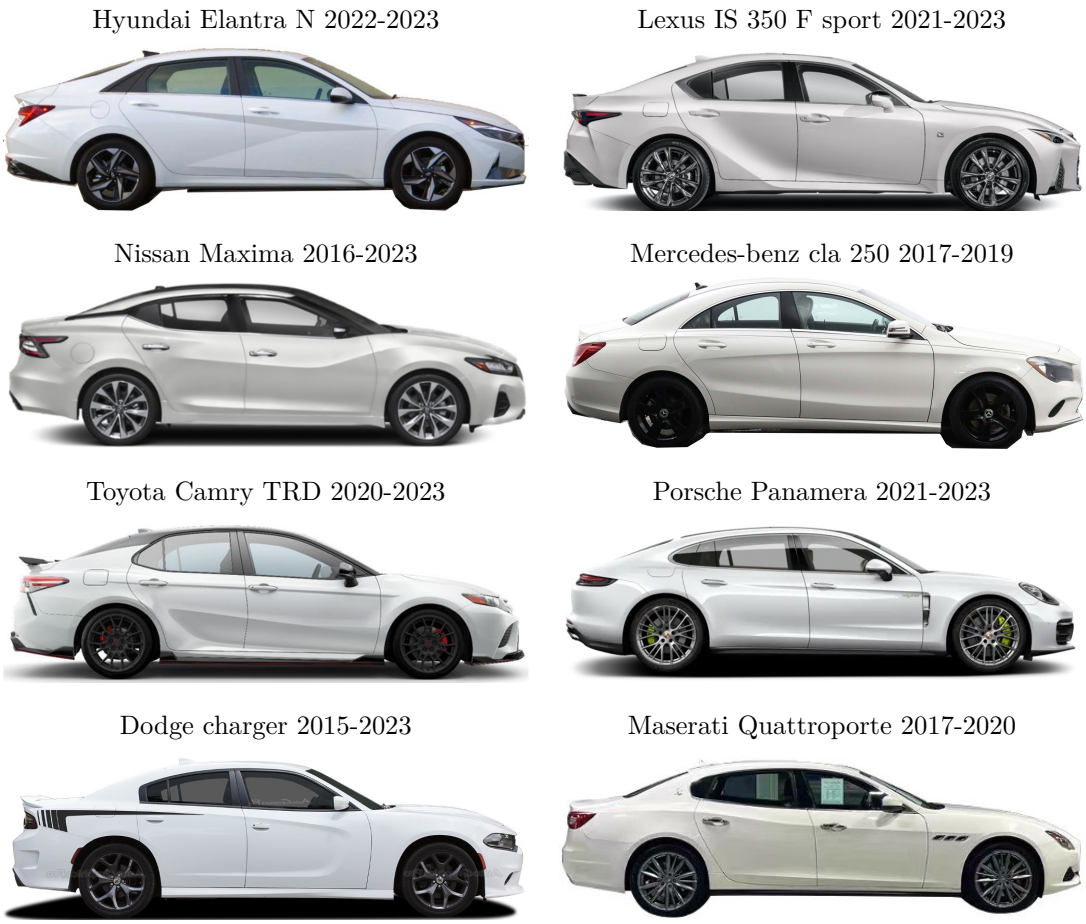

Figure 2: Examples of Sporty Automotive Side Views

augment the automotive aesthetic design process by granting automobile designers the flexibility to adjust the degree to which the sedan shoulder and waist curves should be morphed from those of running cheetahs, and by offering a range of new designs for users (e.g., designers, consumers) to choose from.

Our empirical investigations discover four main findings. First, the new designs preferred by consumers tend to have shorter and curvier shoulder and waist curves compared to original sedan models, consistent with the intra-brand industry trend observed in Figure 1 (see measurements in Appendix A). Second, we observe that consumers generally favor a higher degree of cheetah morphing for premium models compared to mainstream models, suggesting premium models may benefit more from such bio-inspired facelifts. Moreover, consumers tend to prefer a higher degree of morphing from cheetah curves for sportier automobile models compared to regular models. Third,

our analysis reveals that high-income consumers are attracted to a higher degree of morphing from cheetah curves compared to low-income consumers. Lastly, we find that certain automotive colors (e.g., red) and patterns (e.g., cheetah print) can serve as moderators to strengthen the attractiveness of cheetah-inspired designs.

To the best of our knowledge, this research is among the first to develop a computational framework to automatically generate a range of bio-inspired automotive aesthetic designs with a solid theoretical foundation. First, we bridge the psychology literature on knowledge transfer and schema activation and the literature on generative artificial intelligence. Prior psychology and consumer behavior research has used knowledge transfer (e.g., Jia et al. 2020, Gregan-Paxton and John 1997) and schema activation (e.g., Fiske and Linville 1980, Aggarwal and McGill 2007) theories to explain the phenomena of anthropomorphism and biomimicry in marketing. Meanwhile, computer science literature has developed numerous deep learning based methods to generate new images (e.g., Goodfellow et al. 2014, Zhu et al. 2017), which have the potential to be utilized for anthropomorphic and bionic applications. We contribute to both streams of literature by extending the literature on knowledge transfer and schema activation as the theoretical foundation for our proposed generative AI in bio-inspired designs.

Second, we contribute by demonstrating how to leverage deep learning techniques to morph aesthetic design features of the focal product domain with features from a different source domain within the context of bio-inspired automotive designs. To date, several marketing studies have developed deep learning based design tools to automatically generate new logo designs (Dew et al. 2022) and automotive designs (Burnap et al. 2016, 2023) based on existing design space. However, research that helps with sourcing inspirations from a source domain (e.g., sleek body curves of running cheetahs) into the focal domain (e.g., sports sedans) has been sparse. There are two unique challenges associated with such a computational bio-inspired product design: 1) how to quantify representative features of the inspirational source, and 2) how to morph the existing product design with features of the inspirational source. To address the first challenge, our research employs Mask R-CNN and HED to quantify the characteristic shoulder and waist curves of cheetahs. We further use heatmaps to identify the most representative regions of these curves. For the second challenge, we utilize interpolation between the identified cheetah shoulder and waist curves and those of the original sedans. We then apply CycleGAN to restore the shading and depth of curves, rendering

an aerodynamically appealing look. A similar framework can be used to morph the design of other products with other bio-inspired features (e.g., morphing SUVs with rhino features).

The proposed design framework provides practitioners with several practical benefits. First, our framework improves the efficiency of design generation by reducing the manual efforts involved in the traditional bio-inspired automotive design process, making the process easily scale up to different automotive models and inspirational sources. Second, our framework allows designers to source inspirations from nature via calibrating a specific region for a facelift and controlling the level of biological morphing. Third, technological advancements are making sporty features, such as sport-tuned suspensions and specialized tires, more accessible and less expensive to produce[14]. Coupled with the increasing availability of sports trims from automakers, our framework also presents customizable aesthetic design options that can be monetized by automakers. Last but not least, the proposed framework can also be extended to other design applications. For example, designing sports utility vehicles (SUVs), trains, and architecture by sourcing inspirations from other animals like rhinos, kingfishers, and bats, or other objects such as vases and trees (Ripley and Bhushan 2006).

The remainder of the paper is organized as follows. Section 2 provides a discussion of the relationship of our research to extant literature. Section 3 describes the technical details of the proposed design framework. Section 4 presents the results of our empirical investigation. Finally, we conclude the paper in Section 5.

2 Relationship to Extant Literature

Our research is built upon three streams of literature: (i) augmenting automotive aesthetic design with deep learning; (ii) theoretical foundations for our bio-inspired automotive aesthetic design; (iii) methodology foundations for the proposed framework. We discuss below how our paper relates to each research stream respectively.

[14] https://auto.howstuffworks.com/under-the-hood/trends-innovations/top-10-car-tech-from-racing.htm

2.1 Augmenting Automotive Aesthetic Design with Deep Learning

There are generally two types of automobile aesthetic design: 1) redesigns for new generations of automotive models; and 2) facelifts (mid-generational refreshes) in between automotive generations (Schwartz 2000). In both cases, the design team typically screens a range of potential designs to a smaller set of testable designs, followed by consumers evaluating the latter set in theme clinics. Finally, successful designs are advanced for further development. In the current highly competitive market, there is a need to continually reduce the development time and hefty costs of the automotive aesthetic design (Hirz et al. 2013).

In response to this need, several prior studies have utilized deep learning models to reduce the time and manual costs involved in automotive aesthetic design and make the process more scalable. For example, Pan et al. (2016) uses a deep learning based model to identify regions of visual attraction on automobiles. Pan et al. (2017) proposed a deep learning model to predict how consumers across different segments (e.g., wealthy males) perceive an automotive design on a given aesthetic attribute (e.g., sporty). More recently, Burnap et al. (2023) proposed a model that combines a variational autoencoder (VAE) and adversarial components from generative adversarial networks (GANs) to predict the average consumer evaluation of the aesthetic appeal of an automotive design and generate new automotive aesthetic designs accordingly. Our research differs from Burnap et al. (2023) both substantively and methodologically. Substantively, Burnap et al. (2023) generates new automotive aesthetic designs for the entire vehicle or a viewpoint of the vehicle that satisfies designer-specified attributes (e.g., red, Cadillac-like) and optimizes predicted consumer aesthetic rating based on the existing design space. In contrast, we source inspiration from animals whose characteristics may not exist in the current design space to facelift a specific region (i.e., shoulder and waist curves) of the vehicle. Methodologically, Burnap et al. (2023) utilize the VAE component to model the design space and use the GAN component to generate new designs given the design space. In comparison, we use Mask R-CNN and HED combined with heatmaps to extract the body curves of the cheetah and the target sedan, and then generate facelifts for the curves with interpolation and CycleGAN.

More closely related to our research, there have been some recent attempts in the engineering literature that use deep learning models to generate bio-inspired automotive design ideas. For

example, Zhu et al. (2023) utilized generative pre-trained language model GPT3 to create text descriptions of bio-inspired flying car designs (e.g., the flying car has a body that is similar in shape to pterosaurs), which then can be embodied with sketches drawn by human designers. Related, Deng et al. (2023) proposed a deep generative model based on StyleGAN that can fuse the contour shape of sharks with the shape of a sedan to generate bio-inspired automotive contours, which are then transformed into hand-drawn automotive sketches by human designers. Although we also source inspiration from animals, our research differs from these papers in the following important ways. First, both Zhu et al. (2023) and Deng et al. (2023) require human designers to generate the end designs, while our framework provides end-to-end solutions that morph the original automotive aesthetic design images into realistic new bio-inspired design images. Second, while Zhu et al. (2023) and Deng et al. (2023) generate futuristic and possibly wishful-thinking designs, they do not consider either implementation costs or engineering feasibility. In contrast, our research allows automakers to elevate automobiles' visual appeals realistically and cost-effectively without the need to majorly overhaul the fundamental body structure of the existing models.

2.2 Theoretical Foundations for Bio-inspired Automotive Aesthetic Design

In the past, automobile designers have sought inspiration from different sources when it comes to improving the aesthetic designs of their existing models. Some examples of inspirational sources are peer companies, human features, and animal features (Bouchard et al. 2009). Anthropomorphism, the human tendency to see faces in automobile fronts, has been discussed considerably in prior research and utilized to improve automotive aesthetic designs. For example, Purucker et al. (2014) used eye-tracking to demonstrate that humans tend to perceive automobile fronts similar to how they perceive human facial expressions. Anthropomorphism can, in turn, affect product perception and evaluation. Landwehr et al. (2011b) revealed that consumers infer emotional expressions (e.g., friendliness vs. aggressiveness) from the shape of an automobile grill and headlights because they look like a human's mouth and eyes. Based on the premise that automobile fronts are perceived similarly to human faces, Maeng and Aggarwal (2017) showed that automobiles with high face width-to-height ratios are perceived as more dominant, which in turn receive more positive evaluations. Prior literature also examined how anthropomorphism can be evoked. For example, Miesler et al. (2010) revealed that consumers' anthropomorphizing of products can be spontaneous.

In addition to obtaining inspiration from humans, automobile designers also source ideas from animals. This approach, called biomimicry or bionics, involves studying and imitating nature to solve human problems (Aziz et al. 2016). The basic premise is that biological systems in nature have evolved for millions of years to adapt and survive the environment, and thus humans can learn from nature to solve similar problems (Zhu et al. 2023, Kozlov et al. 2015). The wide application of biomimicry has shown benefits in various fields, including medicine, architecture, and transportation (Benyus 1997). In automotive design, biomimicry has been used to improve vehicle performance and fuel efficiency (Wijegunawardana and De Mel 2021). For example, the Mercedes-Benz Bionic car was inspired by the boxfish, a fish known for its boxy and aerodynamic shape (Kozlov et al. 2015), resulting in a car that was both aerodynamically efficient and visually striking.[15] Consumers' evaluation of bio-inspired vehicle designs can also relate closely to the semantics (e.g., rapid speed, sportive, aggressive, futuristic, retro, etc.) and emotions (e.g., positive, arousal, amused, etc.) demonstrated by the inspirational source (e.g., an animal posture) (Kim et al. 2011).

Both anthropomorphism and biomimicry have theoretical roots in schema activation and knowledge transfer. The schema concept refers to prior knowledge abstracted from experience with instances (Fiske and Linville 1980). The term schema also captures the process of matching new information to prior knowledge (Fiske 1982). The schema congruity model predicts that the activated schema (e.g., baby schema) can affect product evaluations in such a way that an instance perceived to fit the schema (e.g., automobile front with large "eyes") will receive the affect linked to the schema (Fiske 1982, Aggarwal and McGill 2007). For instance, Miesler et al. (2011) shows that morphing automobile fronts in accordance with the baby schema (e.g., enlarged headlights/eyes) leads to more positive affective responses compared to the original automobile fronts. Knowledge transfer literature also suggested that consumers can transfer their knowledge from familiar base domains to novel target domains (e.g., Gregan-Paxton and John 1997, Jia et al. 2020). For example, consumers can use their knowledge about human eyes to facilitate their understanding of digital cameras (Gregan-Paxton and John 1997). Jia et al. (2020) shows that consumers can apply norms (e.g., inverse size-movement-speed association) of animate agents (e.g., animals, humans) to their assessment of inanimate products (e.g., Swiss Army knife).

[15] https://newatlas.com/bercedes-benz-bionic-car-modern-art/8869/

Built upon this stream of research, we hypothesize that consumers may transfer their knowledge (e.g., fast, dynamic) about the inspiration source or the activated schema (e.g., cheetah) to perceive and evaluate automotive design models. As such, we adopt knowledge transfer (Gregan-Paxton and John 1997) and schema congruity (Aggarwal and McGill 2007) as the theoretical foundation for utilizing generative AI for bio-inspired design. Our pretest in Section 4.1 shows that displaying an animal image (i.e., cheetah or rhino) alongside a sedan side-view image without any alterations on the actual design indeed changes consumers' evaluations of how sporty the automobile looks. The results are consistent with research on schema congruity for product evaluation (Fiske 1982, Meyers-Levy and Tybout 1989, Aggarwal and McGill 2007).

2.3 Methodology Foundations for the Proposed Framework

In this section, we discuss relevant literature for the three different deep learning algorithms utilized in our proposed framework: Mask R-CNN (Region-Based Convolutional Neural Network), HED (Holistically-Nested Edge Detection), and CycleGAN (cycle consistency generative adversarial network). The Mask R-CNN (He et al. 2017) is a state-of-the-art algorithm for detecting and segmenting object instances in digital images. It has been widely used in autonomous driving (Huang et al. 2020, Xia and Sattar 2019), multi-person pose estimation (Dong et al. 2019), and neural stem cell differentiation (Zhu et al. 2021). We use Mask R-CNN to detect the running cheetah in video frames and separate sedans from the background in their original side-view images.

We then use HED (Xie and Tu 2017), a convolutional neural network based algorithm, to extract body curves of the segmented running cheetah. HED has been used as a middle edge-detection step in applications like semantic image segmentation (Chen et al. 2016, Yu et al. 2017) and salient object detection (Li and Yu 2016, Li et al. 2017). Compared to the Sobel edge detection (SED) that we use to extract body lines of sedan models (see Section 3.2), the HED is more robust in cheetah body curve detection because it can effectively eliminate cheetah print.

Finally, we use CycleGAN (Zhu et al. 2017) to generate realistic sedan side views from a new sedan side view sketch image, due to CycleGAN's superior performance in image-to-image translation tasks like image style transfer and photo enhancement.[16] The first GAN (or vanilla

[16]We do not use CycleGAN to translate a cheetah to a sedan design because the results would not be good for two reasons, according to Zhu et al. (2017). First, CycleGAN is for translation between two visually similar categories.

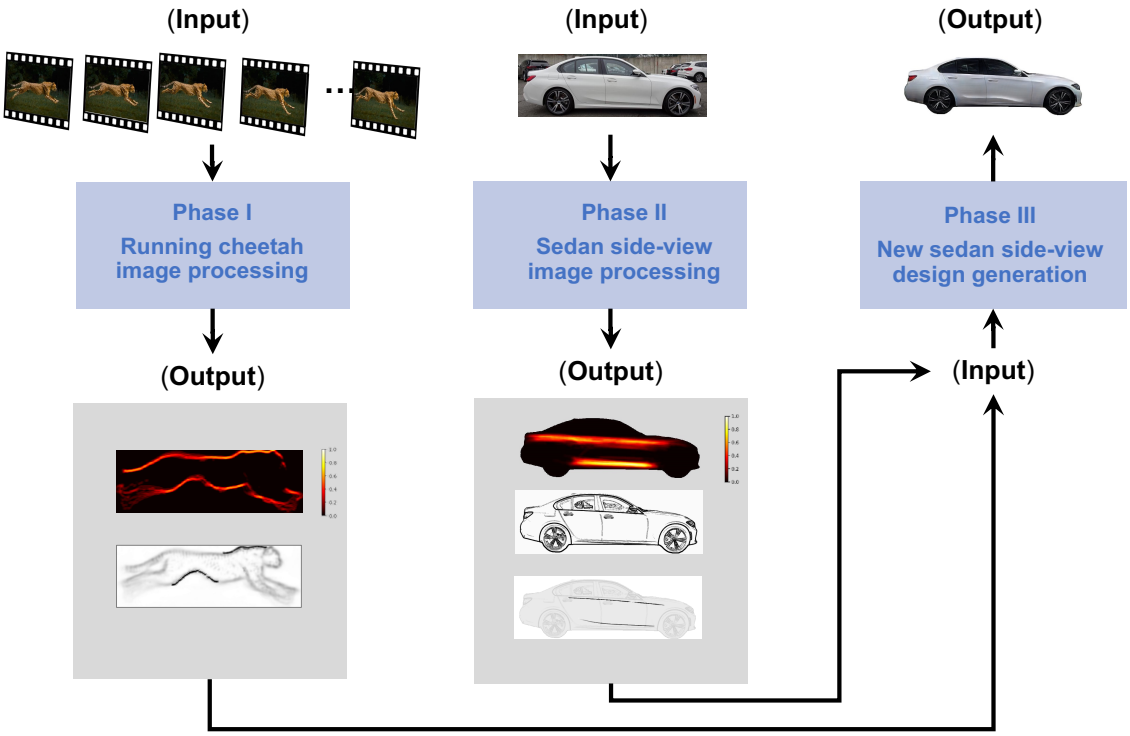

Figure 3: The Proposed Bio-Inspired Design Framework

GAN) was introduced by Goodfellow et al. (2014). Its principle is a min-max game between two neural networks: a generator that generates new data and a discriminator that detects whether the input data is real or generated. GANs have achieved tremendous success in face and landscape visual applications (Isola et al. 2017, Abdal et al. 2019, Karras et al. 2020) and have been recently used for product aesthetic design. For example, Quan et al. (2018) utilized GAN to incorporate given color and styling features into clothing design. Sbai et al. (2019) investigated whether GANs with an explicit loss term for creativity can be used to produce original and compelling fashion designs. In our context, we employ CycleGAN to restore shading and depth to morphed sedan shoulder and waist curves. As a result, our end-to-end design framework can achieve a naturally appealing and aerodynamic look for each of the target automobile models.

3 Deep Learning Framework for Bio-inspired Automotive Aesthetic Designs

Our proposed bio-inspired design framework contains three modules, as depicted in Figure 3. The representative shoulder and waist curves of the cheetah are extracted in the first module. The second model extracts the body curves of the sedan models. The two models can operate separately and in parallel. Then, in the third module, bio-inspired new designs are generated, which can accurately visualize various design outcomes and facilitate consumer evaluation.

3.1 Extracting Representative Body Curves of the Sprinting Cheetah

The first module aims to automatically learn the representative body curves from the world's fastest animal cheetah. Figure 4 presents the key steps in the module.

To form inspiration from the running cheetah, we need to capture its swift and consecutive movements. National Geographic magazine filmed the sprinting cheetah in June 2012 (Smith 2012) to record every nuance of the cheetah's movement as it reached top speeds. A 410-foot dolly track was constructed for a Phantom high-speed camera filming at 1200 frames per second, ensuring stable footage. The video "Cheetahs on the Edge - Director's Cut" was released by National Geographic magazine on YouTube. We downloaded this video with a duration of 7 minutes and 7 seconds.[17]

Because running cheetah does not appear in every frame of the video, we scan every single image frame in the video and use a pre-trained Mask-RCNN (He et al. 2017) to detect running cheetah. The algorithm takes the raw image as the input, and its output includes the detected objects with bounding boxes, the predicted classes for the detected objects, and the corresponding masks that each cover the shape of an object. The Mask R-CNN was pre-trained with the joint loss function

$$\mathscr{L}_{\text{MASK R-CNN}} = \mathscr{L}_{\text{CLS}} + \mathscr{L}_{\text{BOX}} + \mathscr{L}_{\text{MASK}}, \tag{1}$$

where $\mathscr{L}_{\text{CLS}}$ is the classification loss measuring how well the algorithm predicts the ground truth

Second, CycleGAN is good for color and texture transformation but not for geometric changes.

[17] `https://www.youtube.com/watch?v=THA_5cqAfCQ`

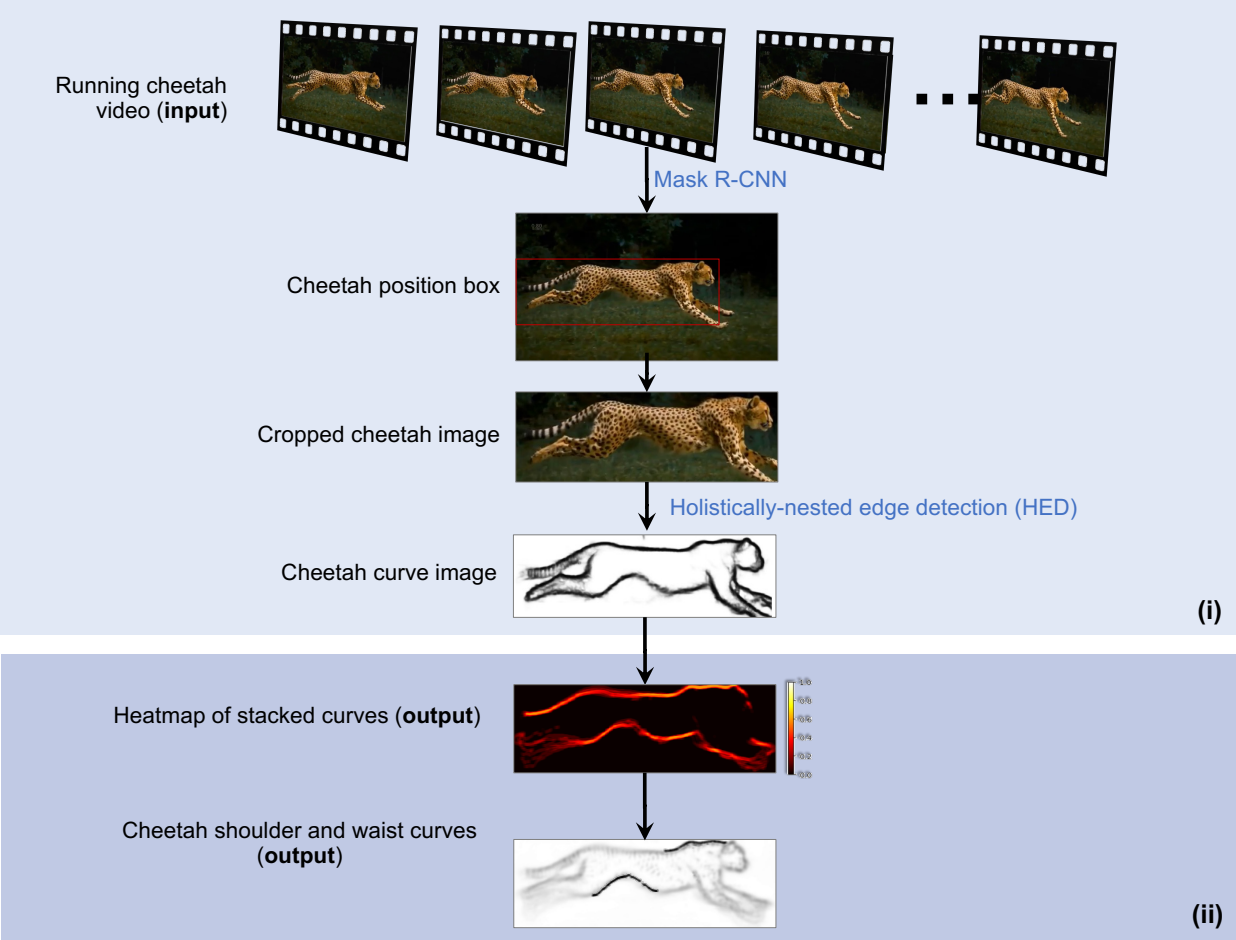

Figure 4: The Running Cheetah Image Processing Module

class of the object such as a cheetah, $\mathscr{L}_{\text{BOX}}$ is the bounding-box loss measuring the distance between the predicted bounding box (a rectangle drawn around the object) and the ground truth bounding box, and $\mathscr{L}_{\text{MASK}}$ is the mask loss that measures how well the algorithm identifies the correct location and shape of the object, such as a cheetah, within the image. Technical details of these loss functions are introduced in Appendix B.

In total, 8,222 frames contain the running cheetah over a duration of 4 minutes and 28 seconds. Due to the slow-motion nature of the cheetah video and the high degree of similarity between adjacent frames, we sample every fifth frame to reduce redundancy without sacrificing precision, resulting in 1,645 sampled images. Then, for each sampled image, we use the predicted bounding box to focus on the running cheetah. As highlighted in Figure 4, the rectangle is the bounding box enclosing the running cheetah.

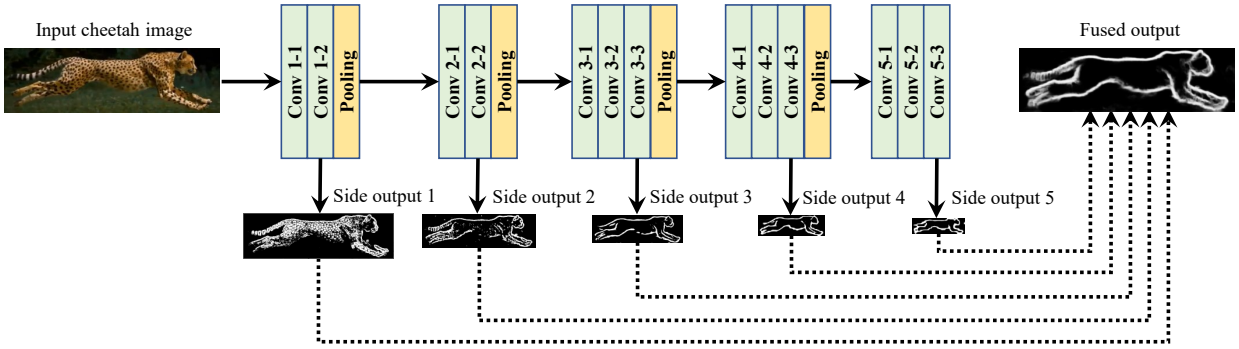

Figure 5: Illustration of HED Architecture for the Running Cheetah Edge Detection

Note: Conv = convolution.

Next, we extract the body curves of the running cheetah using the pre-trained HED (Xie and Tu 2017). This algorithm can effectively capture the contour of the running cheetah and remove the cheetah print. As shown in Figure 5, the HED is a deep convolutional neural network that automatically learns the progressively refined object edge maps produced as side outputs. The side outputs get smaller resolution from left to right due to standard pooling operations in neural networks, which helps to identify edges visible at different resolutions (e.g., subtle edges are visible under high resolution while main edges are clearer under low resolution). Using side outputs in neural networks has proven to be useful in achieving desirable results in image processing (Farabet et al. 2013, Simonyan and Zisserman 2015). The HED was pre-trained with a loss function that considers side outputs and the fused output, expressed as

$$\mathscr{L}_{\text{HED}} = \mathscr{L}_{\text{FUSE}} + \sum_{m} \mathscr{L}_{\text{SIDE}}^{(m)}, \tag{2}$$

where $m = 1, \ldots, M$, is the m^{th} side output, $\mathscr{L}_{\text{SIDE}}^{(m)}$ is the loss function for side output m, and $\mathscr{L}_{\text{FUSE}}$ is the loss function for the fused output. The loss functions measure how well the corresponding outputs predict the correct outline of the cheetah's body.

The side-output loss $\mathscr{L}_{\text{SIDE}}^{(m)}$ is defined as

$$\mathscr{L}_{\text{SIDE}}^{(m)} = -\left[\left[1 - \frac{1}{N}\sum_{n} z_n\right]\left[\sum_{n} z_n \cdot \log\{\hat{z}_n^{(m)}\}\right] + \frac{1}{N}\sum_{n} z_n \sum_{n} (1 - z_n)\log\{1 - \hat{z}_n^{(m)}\}\right], \tag{3}$$

where $n = 1, \ldots, N$, is the n^{th} pixel of the cheetah image and $z_n \in \{0, 1\}$ denotes the ground truth

whether the n^{th} pixel is on the outline of the cheetah body (i.e., an edge pixel). Thus, $\sum_n^N z_n$ is the number of edge pixels in the image, and the fractions $\frac{1}{N}\sum_n z_n$ and $(1-\frac{1}{N}\sum_n z_n)$ are proportions of edge and non-edge pixels, respectively, which are used as class-balancing weights. $\hat{z}_n^{(m)}$ is the m^{th} side output's predicted probability of pixel n being an edge pixel, computed using the sigmoid function $\hat{z}_n^{(m)} = \sigma(\gamma_n^{(m)})$ on the predicted activation value $\gamma_n^{(m)}$ at side output m pixel n.

Then the fused output that combined the predicted cheetah outlines from all side outputs is defined as $\sigma\Big(\sum_m \delta_m \gamma^{(m)}\Big)$, where δ_m is the fusion weight for the m^{th} side output. The fusion loss $\mathscr{L}_{\text{FUSE}}$ is specified as the cross-entropy

$$\mathscr{L}_{\text{FUSE}} = -\sum_n \left[z_n \log\left\{\sigma\Big(\sum_m \delta_m \gamma_n^{(m)}\Big)\right\} + (1-z_n)\log\left\{1-\sigma\Big(\sum_m \delta_m \gamma_n^{(m)}\Big)\right\}\right]. \tag{4}$$

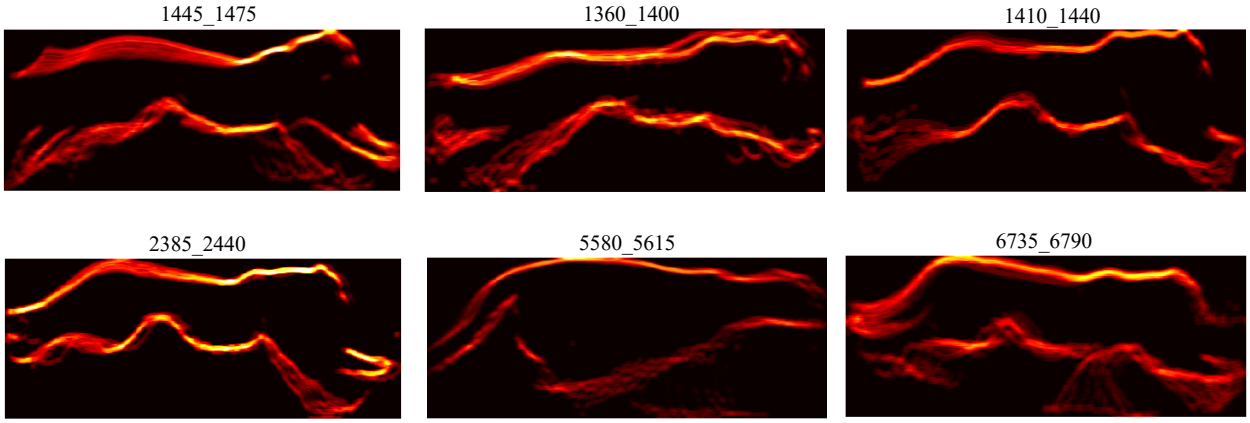

Figure 6: Examples for Stable Motion Moments of the Running Cheetah

With identified cheetah body curves in each video frame, our next goal is to identify representative cheetah curves. Given that the cheetah's body follows similar motion cycles when running, we aim to identify segments of similar consecutive video frames where the body maintains a similar posture in such cycles. We measure the similarity between two consecutive frames by counting the number of overlapped curve pixels. We use the lower 20% percentile of the measured similarity to identify transition points in between segments. These transition points typically differ from both preceding and following frames. Meanwhile, frames between two consecutive transition points display a similar posture. Given the scattered placement of transition points within the video, the 1,645 frames are divided into multiple segments, with many segments comprising only a few

consecutive frames. We focus on segments containing the greatest number of consecutive frames (i.e., roughly 10-20 sampled frames), as they represent the most stable form during the cheetah's running motion cycle. We identify 14 such segments, each representing a stable motion moment (SMM) of the running cheetah. Figure 6 illustrates the heatmaps of six SMM examples. The numbers above each example are the indices of the SMM's first and last video frames. For each SMM, the heat map is generated by stacking sampled frames in the same SMM, with the brighter pixels representing more overlap of identified curves. The heatmaps of SMMs are further stacked to identify the parts of the shoulder and waist curves with the highest degree of density. We deem such curves as the representative body curves of the running cheetah because they capture the most common body curve parts in the most typical posture of a running cheetah.

3.2 Identifying Shoulder and Waist Curves of Sedans

Our second module processes the side-view images of each target sedan model to identify its shoulder and waist curves. Figure 7 demonstrates the process using the BMW 3 Series 2019 as an example. We deliberately use online publicly available sedan images rather than limited numbers of "cleaned" blueprints or advertising images from automakers so more variations can be learned by CycleGAN (details in Section 4.2.1).

As shown in Figure 7, the original image contains a background irrelevant to aesthetic design. Hence, in the first step, we separate the sedan from its background by using the *Mask R-CNN* (He et al. 2017). Its mask output is a binary image (i.e., matrix) that indicates whether the sedan body presents in each pixel. Intuitively, we assign the pixels that lie outside the sedan mask to white.

Sedan's shape, color, and texture are much simpler than those of a cheetah, so we can then use the SED algorithm (Pratt 1978) to generate the corresponding sketch image.[18] At this stage, we convert sedan images into sketches. Sketches are a common tool in automotive design, helping designers gradually build up a design and understand the process (Bar-Eli 2013). In our context, sketch images are especially helpful for extracting and morphing the body curves of both a cheetah and a sedan. During this process, we remove any shading or depth from the curves in our first two models to reduce the influence of color themes and light reflection. In Section 3.3, we convert the

[18] In our study, compared to the HED, the SED can achieve comparable performance for sedan sketch generation while being simpler and computationally faster.

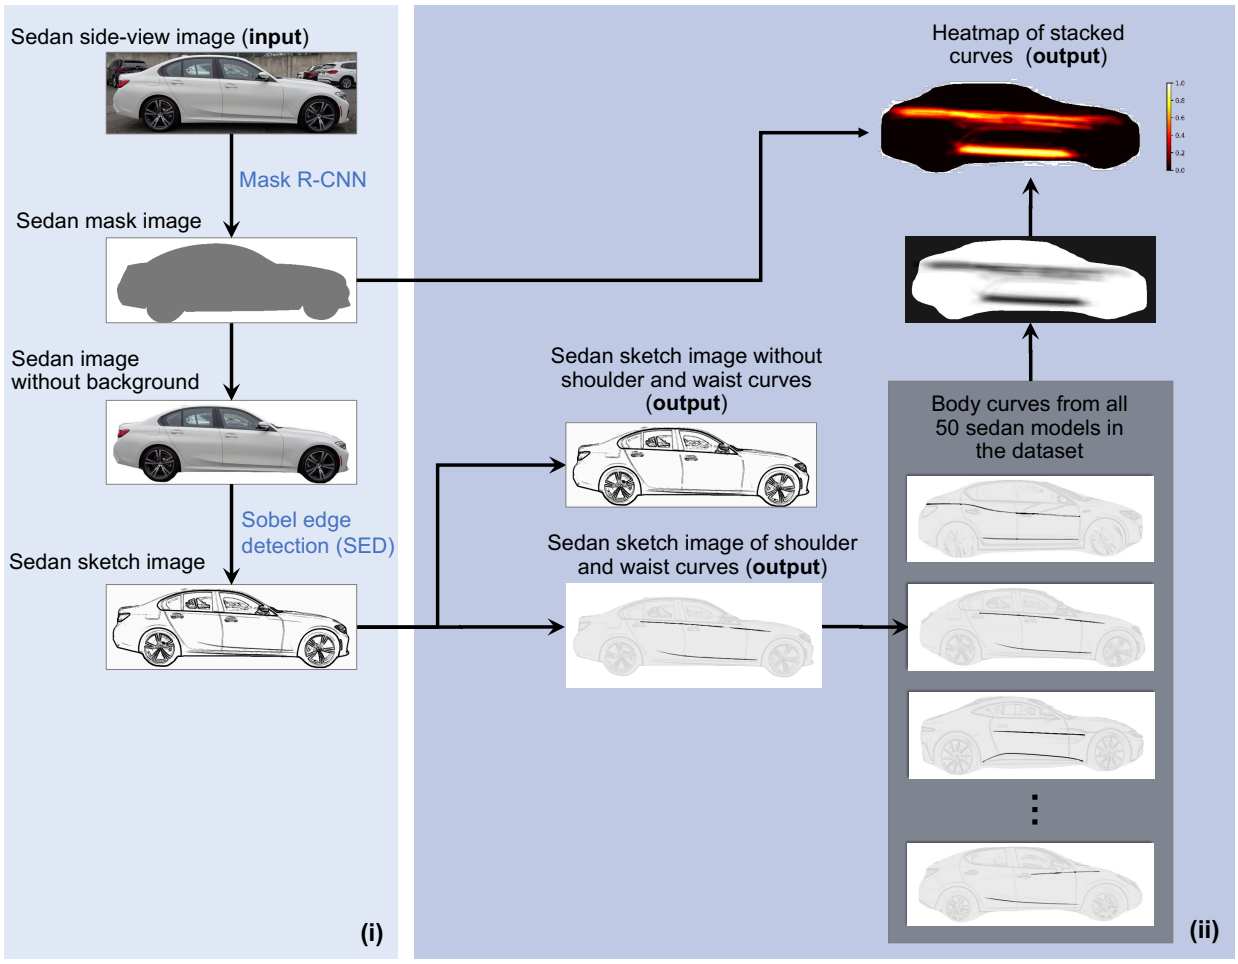

Figure 7: The Sedan Side-View Image Processing Module

sketch back to a realistic image of the sedan using CycleGAN (Zhu et al. 2017) to make it easier for designers and customers to evaluate the final design.

The input RGB image are first converted into grayscale by using the *ITU-R 601-2 luma transform*.[19] Let R, G, and B denote the red, green, and blue channel matrices of the input sedan side-view image without background, then the output grayscale image Θ can be obtained by

$$\Theta = \frac{1}{\lambda_R + \lambda_G + \lambda_B}\Big[\lambda_R R + \lambda_G G + \lambda_B B\Big], \tag{5}$$

where λ_R, λ_G, λ_B are constant parameters. The Sobel operators are two matrices that are convolved with the grayscale image. Let Φ_1 denote the Sobel operator responsible for horizontal edge

[19] `https://pillow.readthedocs.io/en/stable/reference/Image.html`

detection, and let Φ_2 denote the Sobel operator responsible for vertical edge detection. Then pixel (i, j) of the sketch image $\widetilde{\Theta}$ is obtained by

$$\widetilde{\Theta}_{i,j} = \sqrt{\Big[\sum_{i'=-\chi}^{\chi}\sum_{j'=-\chi}^{\chi} \Phi 1_{i',j'}\Theta_{i-i',j-j'}\Big]^2 + \Big[\sum_{i'=-\chi}^{\chi}\sum_{j'=-\chi}^{\chi} \Phi 2_{i',j'}\Theta_{i-i',j-j'}\Big]^2}, \tag{6}$$

where $\chi = \lfloor \frac{\nu-1}{2} \rfloor$, $\lfloor \cdot \rfloor$ is the floor function, ν is the size of Sobel operator, and a Sobel operator's column/row is indexed as $\{-\chi, \chi+1, ...\chi-1, \chi\}$. The technical details of SED are provided in Appendix C. Because we focus on morphing the shoulder and waist curves, each sketch image is decomposed into two images: a sedan sketch image without body curves and a sketch image of shoulder and waist curves.

Lastly, we pinpoint ideal regions to integrate cheetah curves. Traditionally, such decisions usually rest on the designer's aesthetic judgment. We attempt a more data-driven approach. We start with stacking sketch images of body curves from various sedan models in our dataset (details in Section 4.2.1) into a greyscale image where darker pixels indicate more typical curve regions in sedan designs. Images from different sedan models are normalized into the same size (i.e., sedan width and height) for stacking. Our approach is inspired by the "beauty in averageness" phenomenon, which posits that average or typical features are often perceived as more attractive—a principle also utilized in prior marketing studies for idea generation (Toubia and Netzer 2017). We then overlay the mask image of a specific target sedan model (e.g., BMW 3 Series 2019) on the greyscale image, producing a "heatmap" for that model. Because the generated heatmap is scaled back to the original size of the model-specific image, the heatmap generated with the model-specific mask image is unique for each model. In the context of our heatmap, higher values (i.e., higher intensity) indicate more typical curve regions in sedan designs, signifying the best regions for integrating the cheetah curves. Later in Section 3.3, we explain how we use this heatmap to morph the cheetah curves into these regions.

3.3 Generating Bio-Inspired Sedan Side-View Facelift Design

Our final model generates new sedan side-view aesthetic designs by morphing the identified body curves of the target sedan model with the identified shoulder and waist curves of the sprinting cheetah. Figure 8 presents a schematic view of the final module.

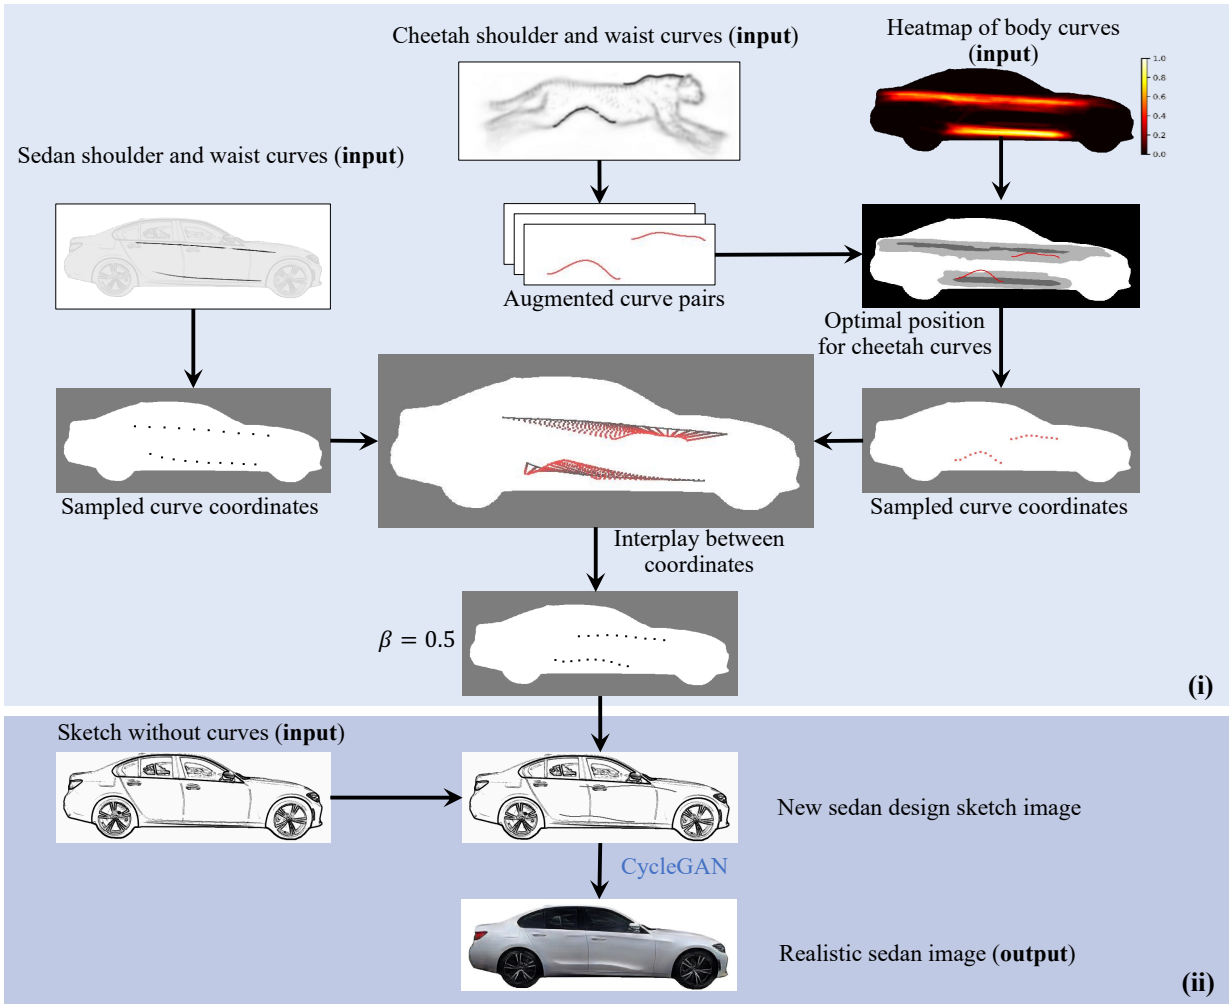

Figure 8: The New Sedan Side Facelift Design Generation Module

Note: The morphing parameter β ranges between 0 and 1. As an example, we let $\beta = 0.5$ in this figure.

To morph the target sedan's body curves with cheetah curves, we first find the best position on the sedan heatmap produced in Section 3.2 for incorporating the cheetah curves. This includes figuring out the right size, rotation, and location for the curves. In order to do so, the pair of cheetah shoulder and waist curves generated in Section 3.1 is re-scaled into ten different sizes and rotated with ten different angles.[20] This gives us 100 different combinations of size and angle. We next let each combination of cheetah curves move across the sedan heatmap. The overlap value between each combination at a certain location and the heatmap is computed by calculating the

[20] Extremely large or small curve sizes and dramatic rotation angles were avoided first as they would lead to noticeable distortions. To closely approximate the optimum, we then selected ten levels for both sizes and angles to ensure small but noticeable differences between adjacent levels.

Frobenius inner product between the binary matrix of the cheetah curve combination image and the greyscale (0-255) matrix of the heatmap. Higher values on the heatmap represent better spots for placing body curves, as discussed in Section 3.2. Then, the combination at a certain location that maximizes the overlap value is the best position for incorporating the cheetah curves for a given sedan. The best position is highlighted in red color, as illustrated in Figure 8. Our method for placing cheetah curves within a car side view design draws parallels to the video product placement method proposed by Yang et al. (2021). They calculate pixel-level engagement for each frame of an influencer video and recommend placing the product in the region with the highest engagement score. Similarly, we calculate the pixel-level typicality for curve integration via the sedan heatmap and place the cheetah curves in the region of the sedan heatmap with the highest typicality value.

We then fuse the corresponding body curves (e.g., the sedan's shoulder curve and the cheetah's shoulder curve). To do so, U points are sampled with a fixed horizontal step size from each body curve. In our empirical investigation described in Section 4.2, we set $U = 10$, which provides us with sufficient precision. Finally, we fuse the corresponding sedan and cheetah body curves with linear interpolation as described in Equation 7. Let $c_u = (c_u^x, c_u^y)$ denote the coordinates representing the u^{th} sampled point of the shoulder/waist curves of the sedan, and let $\widetilde{c}_u = (\widetilde{c}_u^x, \widetilde{c}_u^y)$ denote the coordinates representing u^{th} sampled point of the shoulder/waist curve of the running cheetah. The coordinates $r_u = (r_u^x, r_u^y)$ representing u^{th} point of the morphed sedan shoulder/waist curve can be obtained by:

$$r_u = (1-\beta)c_u + \beta\widetilde{c}_u, \tag{7}$$

where β is the coefficient that controls the degree of morphing for the shoulder and waist curves, subject to the constraint that $0 \leq \beta \leq 1$. We then obtain the smooth morphed sedan shoulder/waist curves by polynomial fitting (Bishop 2006) based on the vector of points representing the morphed curves.

Next, the new sedan design sketch image is generated by stacking the target sedan's side sketch image without the body curves and the newly morphed shoulder and waist curves. In this step, we use CycleGAN (Zhu et al. 2017) to turn the new sketches into realistic images of the sedans. This tool is great for converting one image into another, and by training it on real sedan images, it

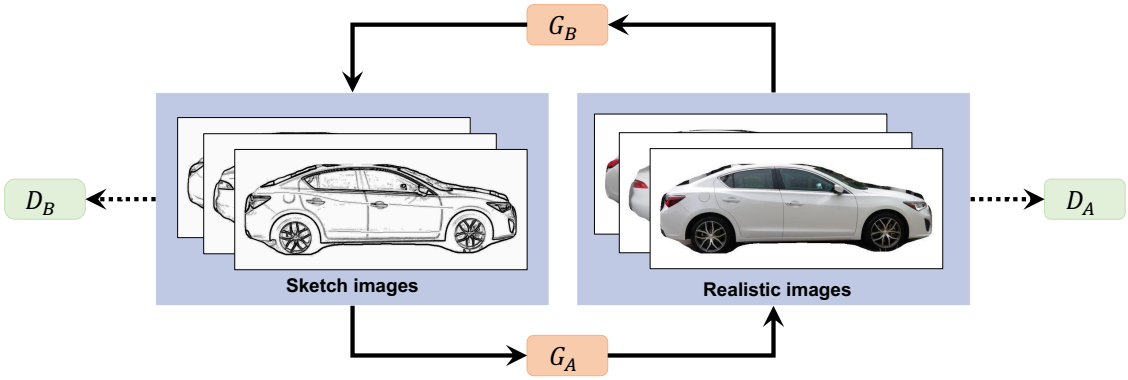

Figure 9: Illustration of Cyclegan Structure

can automatically learn how to add back the shading and depths to the curves, making them look natural and aerodynamically appealing.

The structure of CycleGAN is illustrated in Figure 9. In our framework, we aim to translate the new sedan design from the domain of sketch images into the domain of realistic images. We use $\mathscr{A}$ to denote the domain of the realistic images of sedan models and $\mathscr{B}$ to denote the domain of sketch images. CycleGAN contains two GANs (Goodfellow et al. 2014): one for the realistic images and one for the sketches. Each GAN has a generator and a discriminator. The former generates images, while the latter detects if an image is generated or not. The two generators translate images between the realistic image domain and the sketch image domain.

Mathematically, we denote G_A for the generator that maps from $\mathscr{B}$ to $\mathscr{A}$, G_B for the generator that maps from $\mathscr{A}$ to $\mathscr{B}$, D_A for the discriminator that detects if a realistic image is generated by G_A or not, and D_B for the discriminator that detects if a sketch image is generated by G_B or not. CycleGAN adopts the basic idea of GAN by constructing a minimax game between the discriminator and the generator. The CycleGAN is calibrated by solving the following optimization problem

$$\min_{G_A,G_B} \max_{D_A,D_B} \mathscr{L}_{\text{GAN-A}} + \mathscr{L}_{\text{GAN-B}} + \mathscr{L}_{\text{CYCLE}}, \tag{8}$$

where $\mathscr{L}_{\text{GAN-A}}$ is the loss function for the GAN that generates and distinguishes realistic images, $\mathscr{L}_{\text{GAN-B}}$ is the loss function for the GAN that generates and distinguishes sketches, and $\mathscr{L}_{\text{CYCLE}}$ is the loss function for ensuring the cycle consistency of image translation between two domains.

We adopt the specifications used by Zhu et al. (2017) for these loss functions:

$$\mathscr{L}_{\text{GAN-A}} = \mathbb{E}_{A\in\mathscr{A}}\Big[\log\{D_A(A)\}\Big] + \mathbb{E}_{B\in\mathscr{B}}\Big[\log\{1 - D_A(G_A(B))\}\Big], \tag{9}$$

$$\mathscr{L}_{\text{GAN-B}} = \mathbb{E}_{B\in\mathscr{B}}\Big[\log\{D_B(B)\}\Big] + \mathbb{E}_{A\in\mathscr{A}}\Big[\log\{1 - D_B(G_B(A))\}\Big], \tag{10}$$

$$\mathscr{L}_{\text{CYCLE}} = \mathbb{E}_{A\in\mathscr{A}}\Big[\|G_A(G_B(A)) - A\|_1\Big] + \mathbb{E}_{B\in\mathscr{B}}\Big[\|G_B(G_A(B)) - B\|_1\Big], \tag{11}$$

where $\|\cdot\|_1$ is the L^1 norm. We describe how we train the CycleGAN for our empirical investigation in Section 4.2.1.

4 Empirical Investigations

We start with a pretest to examine whether priming consumers to perceive the appearance of an automobile as inspired by an inspiration source (e.g., cheetah) impacts consumers' evaluations towards the focal sedan as we hypothesized in Section 2.2. We then conduct Study 1 to recruit sedan intenders to evaluate the new cheetah-inspired automobile designs. Lastly, we implement Study 2 to explore whether aesthetic features such as color and print pattern can serve as moderators to enhance the attractiveness of our bio-inspired design.

4.1 Pretest of Knowledge Transfer via Schema Activation

In this pretest, we explore whether consumers tap into their knowledge (e.g., fast, dynamic) about the inspiration source or activated schema (e.g., cheetah) when they are primed to believe that the focal automobile is modeled after the inspiration source. A total of 317 participants are recruited via Amazon Mechanical Turk (MTurk). To ensure the quality of the responses, we require that all participants of our study must reside in the United States and have more than 500 tasks completed with an approval rate no less than 99%.

4.1.1 Pretest Design

Our pretest adopts a between-subjects design with three conditions: control, cheetah, and rhino. We frame either a cheetah or a rhino as the inspiration source to test 1) whether the knowledge transfer happens and 2) if it is specific to the inspiration source. Participants are asked to evaluate

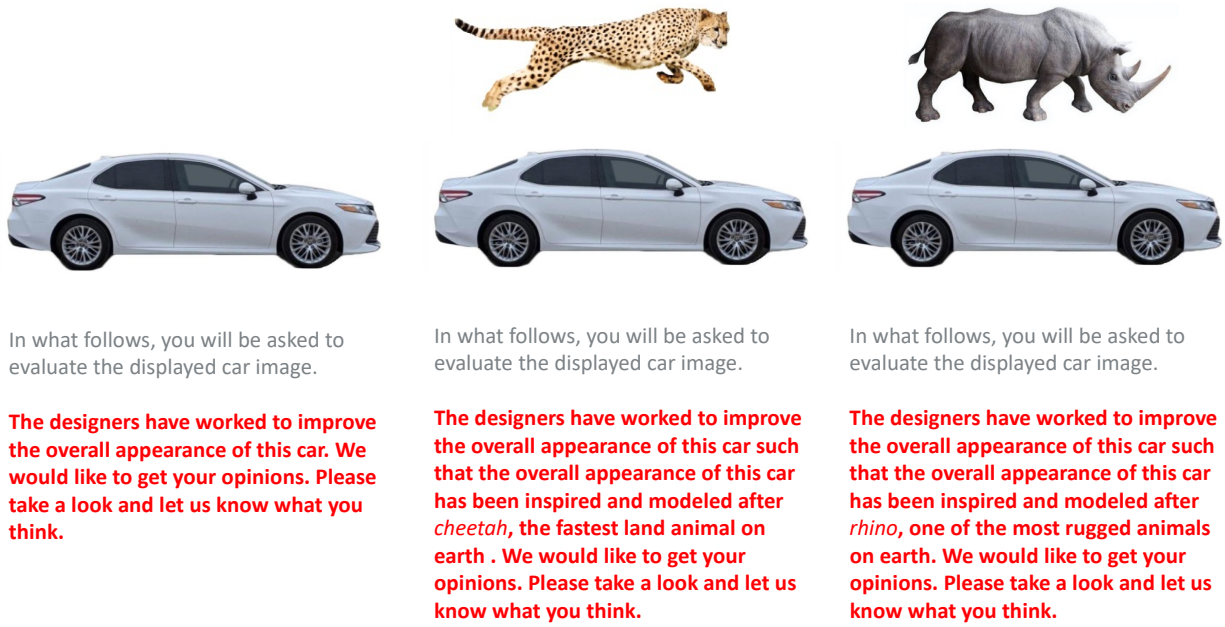

(a) Control condition **(b) Cheetah condition** **(c) Rhino condition**

Figure 10: Pretest Experimental Conditions

the design of a new sedan, and they are randomly assigned to the three conditions. In all three conditions, we display an identical sedan image to the participants and tell them that the designers have worked to improve the overall appearance of the sedan (see Figure 10). The only difference is that in the cheetah condition, participants were displayed with both sedan and cheetah images, and they were told that the overall appearance of the sedan had been inspired and modeled after the cheetah, the fastest land animal on Earth. In the rhino condition, participants were displayed with both sedan and rhino images, and they were told that the overall appearance of the sedan had been inspired and modeled after the rhino, one of the most rugged animals on Earth. We intentionally use both images and text to present the inspiration source to ensure the knowledge of the inspiration source is clearly communicated to the consumers.

If the perception of automotive designs inspired by different animals can be affected by the knowledge of the characteristics of specific animals, we expect participants in the three conditions to perceive the sedan differently in several dimensions: acceleration, exotic appeal, speed, and strength. Cheetahs are known for their incredible speed and agility and are often associated with exotic locations and rare beauty. Rhinos are known for their robust and powerful build. Unlike the cheetah, the rhino is not particularly known for its speed. In summary, cheetah-inspired

aesthetic designs may leave a stronger impression of fast acceleration, exotic appeal, fast speed, and strength than the control condition. Rhino-inspired aesthetic designs may have a strong expression of strength but may not be closely associated with fast acceleration, exotic appeal, or fast speed, compared to the control condition.

Participants in all three conditions evaluate the appearance of the sedan on four 9-point Likert scale questions (i.e., 1 = "Strongly Disagree" and 9 = "Strongly Agree"): "(i) this sedan can accelerate very fast; (ii) this sedan is exotic; (iii) this sedan is fast; and (iv) this sedan is powerful". On average, participants took 118 seconds to complete the survey. Following prior research (Burnap et al. 2023, Luo and Toubia 2015), responses from eight participants were dropped because the duration of their response time was either too short or too long.

4.1.2 Results

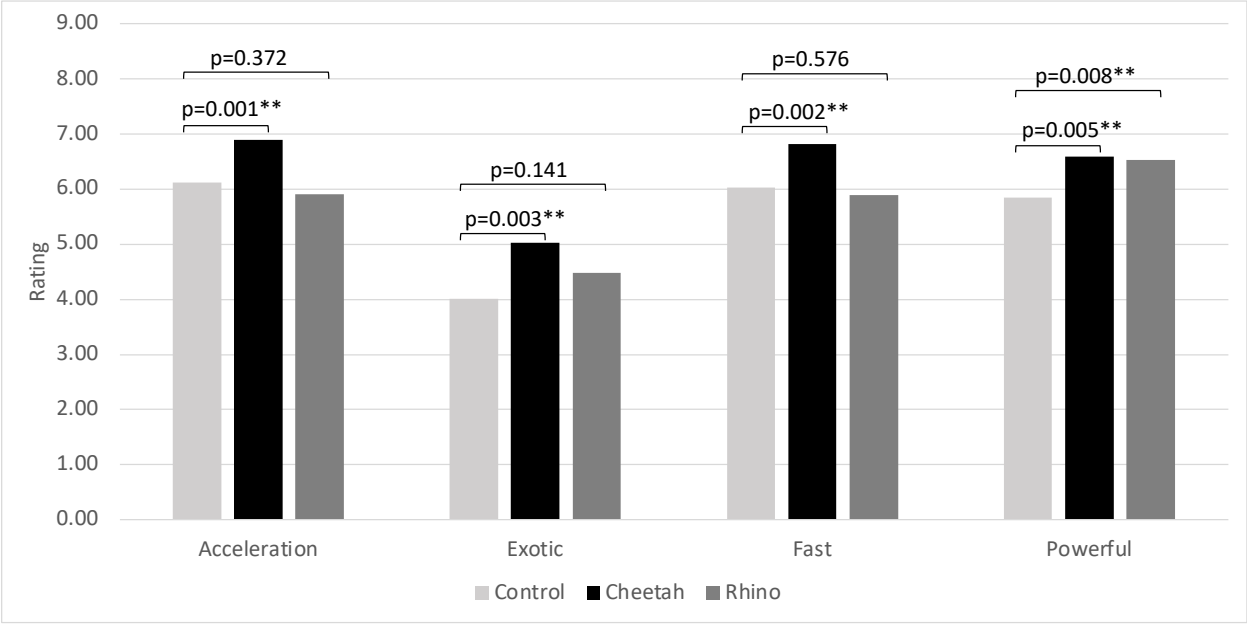

Figure 11: Users' Responses to Aesthetic Design

Significance levels: $* \ p < 0.05$, $** \ p < 0.01$, $*** \ p < 0.001$.

Figure 11 presents the results of our pretest. We discover that, although the same sedan image is shown in both the control and cheetah conditions, participants in the cheetah condition perceive the sedan to accelerate faster ($\text{Mean}^{\text{cheetah}} = 6.89$ vs. $\text{Mean}^{\text{control}} = 6.11$ and $p = 0.001$), more exotic ($\text{Mean}^{\text{cheetah}} = 5.02$ vs. $\text{Mean}^{\text{control}} = 4.01$ and $p = 0.003$), faster ($\text{Mean}^{\text{cheetah}} = 6.82$ vs.

$\text{Mean}^{\text{control}} = 6.03$ and $p = 0.002$) and more powerful ($\text{Mean}^{\text{cheetah}} = 6.59$ vs. $\text{Mean}^{\text{control}} = 5.85$ and $p = 0.005$). In comparison, participants in the rhino condition perceived the sedan as only more powerful than participants in the control condition ($\text{Mean}^{\text{cheetah}} = 6.53$ vs. $\text{Mean}^{\text{control}} = 5.85$ and $p = 0.008$). The results suggest that consumers indeed transfer specific knowledge about the inspiration source (cheetah or rhino) to evaluate the focal sedan model. Moreover, our pretest provides empirical support that framing cheetahs as an inspiration for automotive aesthetic design enhances consumers' perceived sportiness of such an automotive design, even when the appearance of the sedan is kept constant. The results from our pretest lead us to continue our investigation into whether consumers actually find the aesthetics of sedans more appealing when we employ a cheetah-inspired automobile design as proposed in Section 3.

4.2 Study 1 Applying the Proposed Framework to Generate Bio-Inspired Designs

We view Study 1 as a proof of concept to examine whether our proposed deep learning framework can be used to generate bio-inspired aesthetic designs preferred by actual sedan intenders. We first select ten popular sports sedans based on the 2019 popular automobile rankings from U.S. News and JD Power, including five premium models (Mercedes C Class, BMW 3 Series, Audi A4, Cadillac CT4, and Lexus IS) and five mainstream models (VW Passat, Ford Fusion, Toyota Camry, Honda Accord and Mazda 6).[21] To test the boundary condition of our cheetah-inspired sedan aesthetic designs, we further include three extra sporty sedan models with more extreme shoulder and waist curves (Nissan Maxima, Lexus IS 350 F Sport, and Maserati Quattroporte) as demonstrated in Figure 2. On one hand, our cheetah-inspired designs might not improve consumer perception of extra sporty models, since those sedans already have a fast and powerful appearance. On the other hand, the dynamic and sporty vibe of the cheetah could be a great context for extra sporty sedans. Therefore, we include a total of 13 sports sedan models (5 premium models, 5 mainstream models, and 3 extra sporty models) in Study 1. For each sedan model, we provide a spectrum of aesthetic designs, varying in the degree of cheetah morphing in shoulder and waist curves generated from the

[21] `https://cars.usnews.com/cars-trucks/rankings/used/2019-luxury-small-cars` (U.S.News); and `https://www.jdpower.com/cars/sedans/10-most-popular-midsize-cars` (J.D.Power). The ten models either market themselves as sports sedans or have a sports trim.

deep learning framework proposed in Section 3. We then ask consumers to select their preferred aesthetic design for each of the sedan models.

4.2.1 Training CycleGAN and Generating Candidate Automotive Designs

Our proposed design framework uses CycleGAN to generate realistic sedan side designs based on sketches. To ensure that our CycleGAN can generalize well beyond the 13 sedan models used in Study 1, we perform two operations in model calibration. First, we collect actual sedan side-view images of 50 different sedan models from 36 automakers from autotrader.com. The names of the 50 sedan models are listed in Appendix D. For each model, 20 side-view images from model years between 2010 and 2019 are sampled for each model, resulting in 1,000 total images.[22] Second, to incorporate additional variations, we augment the collected images by cropping each image with 280 possible positions, resulting in 280,000 realistic sedan images in total. Data augmentation is a commonly used technique in training machine learning models. It provides additional data points for model training and helps alleviate the problem of overfitting. Because training the cycleGAN requires both realistic images and sketches, we generate two sketches with different depths of shades using the SED algorithm for each realistic image. In total, we have 280,000 realistic images and 560,000 sketches for model calibration.

Next, we generated multiple bio-inspired facelift designs with varying degrees of morphing for each of the 13 sedan models in this study. The key processing steps of generating a bio-inspired facelift design for a morphing degree are demonstrated in Figure 12 using the BMW 3 Series as an example. Figure 13 provides three cheetah-inspired design examples for BMW 3 series (with β=0 being the original design with no morphing). We can see that when the value of the morphing parameter β increases (i.e., a higher degree of morphing from running cheetah), the shoulder and waist curves of the sedan become shorter and curvier, and the differences between the new design and the original design become more apparent. Figure 13 shows that our framework allows designers to finetune the degrees of biological morphing by offering a range of bio-inspired new designs for users (e.g., designers, consumers) to choose from.

[22] We prioritized more recent years, as using images from recent years ensures that CycelGAN generates styles that align closely with current market preferences. On the other hand, having a broad year range for selection mitigates image shortage for certain sedan models.

Figure 12: Key Steps of Generating a Bio-Inspired Facelift Design for a Given Morphing Degree

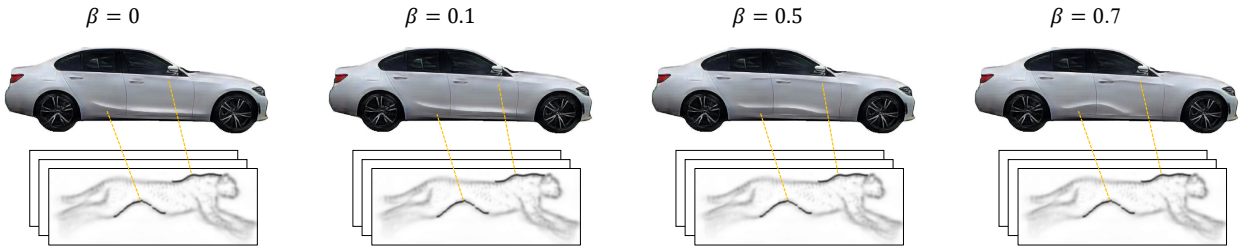

Figure 13: Illustration of New Designs of BMW 3 Series Generated by the Proposed Framework

4.2.2 Study Design

We then collect data to examine whether consumers prefer cheetah-inspired automobile designs more than their original look. Similar to prior research (Burnap et al. 2023), we recruit sedan intenders (consumers who intend to purchase sedans in the near future) from a professional panel called CloudResearch Prime Panel as participants in this study.[23] Given that we consider both premium and mainstream sedan models in this study, participants are uniformly and randomly recruited from five distinct income groups: less than 50k, 50k-75k, 75k-100k, 100k-150k, and more than 150k. Our survey includes video-based instruction and two attention-check questions to ensure that the participants understand and pay attention to our survey questions. The participants are disqualified from taking part in the study if they provide wrong answers to the attention check questions. A total of 275 participants are recruited.

After the introduction page, our study presents one car model on each screen. Figure 14 presents an example of the BMW 3 Series. On each screen, we present a slider with which the

[23] Prime Panels is a participant recruitment platform that was developed by CloudResearch, formerly TurkPrime, in 2017. Chandler et al. (2019) compared Prime Panels to MTurk samples and nationally representative samples. The research revealed that it is possible to collect high-quality data on Prime Panels. Examples of publications powered by Prime Panels include Davidai (2018), Pérez-Rojas et al. (2019), etc.

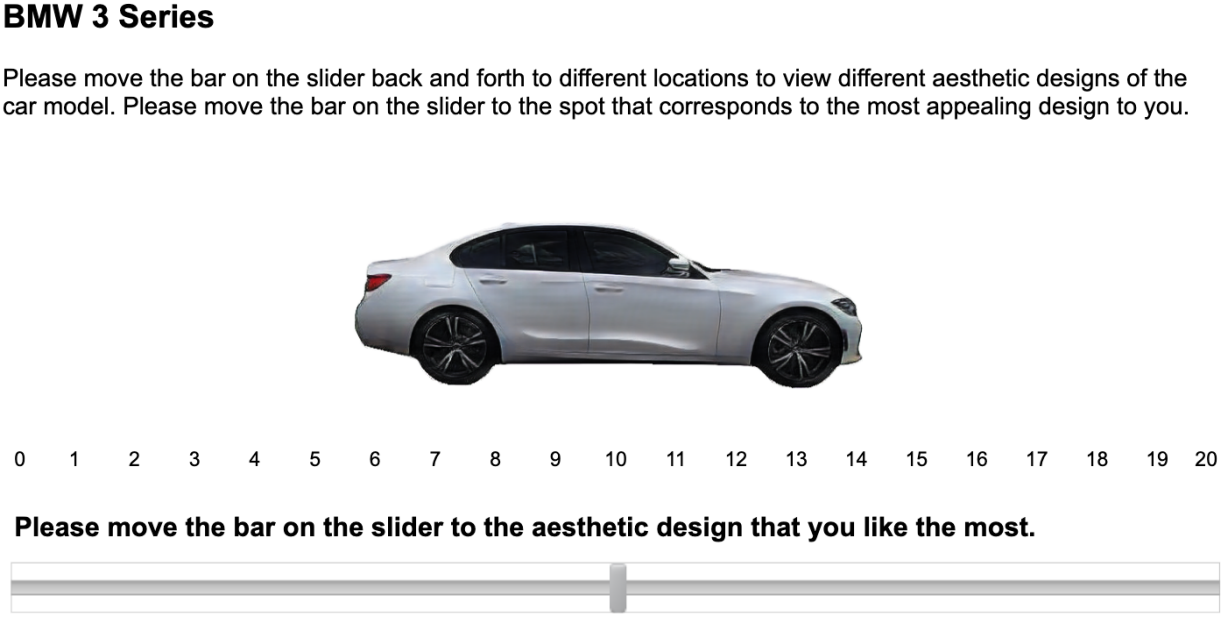

Figure 14: The Survey Page for BMW 3 Series

survey participant can move the bar back and forth to view different aesthetic designs of one target sedan model.[24] For each sedan model, each participant is presented with 21 different designs with different degrees of morphing, with 0 being the sedan's current look (i.e., $\beta = 0$) and 20 being the highest degree of morphing (i.e., $\beta = 0.7$) based on the approach outlined in Section 3. To avoid any potential confounder, we did not inform participants about which design is the current design or the most significant change.[25] We set the upper bound 0.7 for β in this study because we believe that the shoulder and waist curves with 100% morphing from cheetah curves seem overly dramatic from our visual inspection. Participants are instructed to move the bar back and forth until they find the most appealing look to them. The participants took 333.32 seconds on average to complete the survey. Following prior research (Luo and Toubia 2015, Burnap et al. 2023), responses from nine participants were dropped because they answered too quickly, too slowly, or in a "straight line" pattern.

[24] Following convention (e.g., Chan et al. 2018), the default position of the slider was set at the midpoint, and participants were required to move each slider at least once to indicate their response. To minimize noise, we did not use a random starting point.

[25] In an earlier version of this study, we informed participants about which design was the current one and which represented the most significant change; the results were qualitatively the same.

4.2.3 Results

For each sedan model included in Study 1, we identify the most preferred design by taking the average of the morphing parameter β corresponding to the participants' most preferred designs. In order to better illustrate the differences between the original and the new designs, we locate the shoulder and waist curves using rectangular boxes and report their width (W), height (H), and central position (C). The central position contains both horizontal and vertical locations of the box's centroid. Figure 15 presents an example of the BMW 3 Series. One notable change is that the waist and shoulder curves in the preferred new design become shorter. Another notable change is that the shoulder curve moves towards the front of the sedan while the waist curve moves a little backward.

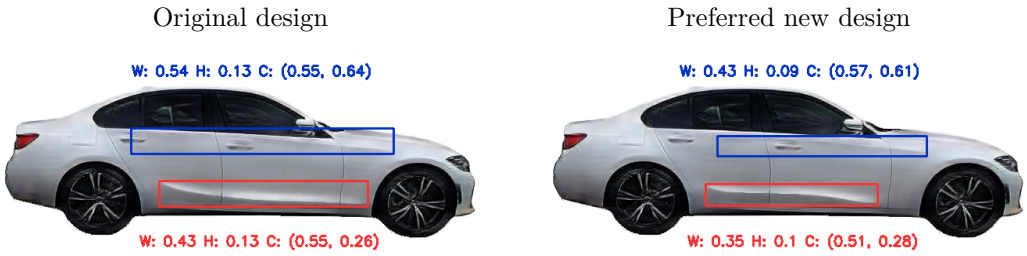

Figure 15: Comparison of Original and Preferred New Designs for BMW 3 Series

Note: The blue box locates the shoulder curve, and the red box locates the waist curve. W, H, and C represent the width, height, and central position of the corresponding box, relative to the sedan body.

Figures 16-18 compare the preferred new designs and the original designs of the 5 premium, 5 mainstream, and 3 extra sporty sedans with more extreme shoulder and waist curves, respectively. To better visualize differences between the original and preferred new designs, we omit the rectangle boxes in Figures 16-18 but still display the box measurements (i.e., width, height, and central position) for waist and shoulder curves.

Based on the width of the original shoulder curves, our target sedan models can be classified into three categories. The first category includes Mazda 6, Lexus IS 350 F Sport, and BMW 3 Series. They have short shoulder curves. Their new shoulder curves are even shorter. The shoulder curve moves a bit forward for Lexus IS 350 F Sport and BMW 3 Series while backward for Mazda 6. The second category includes Mercedes C Class, Toyota Camry, and Nissan Maxima, whose shoulder curves are in the medium range among all models included in our study. For this category, the

Figure 16: Original and Preferred New Designs of the Target Premium Sedans

lengths of shoulder curves are shortened significantly in the new designs, creating a sportier look. The sedans in the third category have very long shoulder curves (about 70% - 90% of the sedan body length), particularly Audi A4, VW Passat, and Ford Fusion, whose shoulder curves start from the sedan front and stop near the end of the sedan. The new shoulder curves of sedans in this category have a large length reduction at about 24%.

Based on the width of the original waist curves, our target sedan models can also be classified into three categories. Lexus IS and IS 350 F Sport form the first category, whose waist curves have

Figure 17: Original and Preferred New Designs of the Target Mainstream Sedans

the shortest lengths and relatively lower positions. In the new designs, both widths and heights of their waist curves reduce, and their positions move upward. Toyota Camry is in the second category, whose waist curve is slightly longer but has the highest position. In the new design, both the width and height of its waist curve reduce as well, but its position moves downward. The third category includes the rest of the sedan models. Their waist curve widths reduce, but their waist curve heights either increase or do not change (except for BMW 3 Series, whose waist curve height reduces).

Figure 18: Original and Preferred New Designs of Extra Sporty Sedans

By and large, the preferred designs favored by participants tend to feature shorter and curvier shoulder and waist curves, along with a more pronounced slope. These changes can be visually identified and are consistent with the trend over recent generations of BMW 3 Series demonstrated in Figure 1 (see measurements in Appendix A).

Moreover, compared with the original designs, we note that consumers prefer greater degrees of changes in their body curves for some sedan models in our study, while the others have less significant but still noticeable modifications. Notably, the average morphing parameter chosen by participants for premium sedans is slightly higher than that of mainstream sedans ($Mean_{premium} = 0.311$ vs. $Mean_{mainstream} = 0.297$, $p = 0.087$), suggesting that consumers prefer a higher degree of morphing from cheetah curves for premium vs. mainstream sedans. Furthermore, participants on average choose even higher morphing degrees for the three extra sporty sedan models with relatively more dramatic shoulder and waist curves as compared to the remaining 10 sedan models in our study ($Mean_{sporty} = 0.327$ vs. $Mean_{regular} = 0.304$, $p = 0.014$), suggesting our method is not limited to the boundary condition of this group of sedans.

Lastly, we also observe some interesting heterogeneous preferences among our participants in

terms of their preferences for the ideal degree of cheetah morphing. Given that we recruit participants uniformly at random from five income groups (<50k, 50k-75k, 75k-100k, 100k-150k, >150k), using a median split, we categorize those earning >= 75k as high income, and the rest as low income. We find that high-income consumers significantly prefer a higher degree of cheetah morphing compared to low-income consumers ($Mean_{high} = 0.325$ vs. $Mean_{low} = 0.284$, $p = 0.000$). This observation offers useful guidance for automakers in the premium segment of the sedan market. This finding is also managerially useful for automakers when it comes to consumer targeting, as many automakers provide several trim options for the same model. For example, the 2022 VW Passat has three trim options - SE, R-Line, and Limited Edition - that differ in aesthetics, functions, as well as prices.

4.3 Study 2 Moderating Role of Automotive Color and Print Pattern

Built upon the promising findings from our pretest and Study 1, we proceed to explore whether aesthetic features such as color and print pattern can make our bio-inspired design even more attractive. Answers to this question can be useful for automotive designers and dealerships that offer consumers with customization options. Specifically, we focus on two characteristics: color and pattern. For color, as an illustration, we test out the red color, because red color can be associated with power, arousal, and excitement (Labrecque and Milne 2012) and is often adopted for sports cars (e.g., Ferrari 458 and 488). For pattern, we test out the cheetah prints. Cheetah print is often used on clothes, shoes, accessories, and sometimes as automobile decals. In our case, combining cheetah prints and body curves for automotive aesthetic design may present a stronger linkage to the inspiration source - cheetah - which may in turn enhance consumers' perception of sportiness.

4.3.1 Study Design

To ensure that participants are representative of the target consumer segment, as in Study 1, we pre-screen for "sedan intenders" and recruit participants uniformly from the five income groups (<50k, 50k-75k, 75k-100k, 100k-150k, >150k) from Prime Panels of CloudResearch. This study includes similar instructions and two attention-check questions as in Study 1. The participants are disqualified from taking part in the study if they provide wrong answers to the attention check questions. A total of 284 participants are recruited.

(a) Control condition **(b) Red color condition** **(c) Cheetah print condition**

Figure 19: Study 3 Experimental Conditions

We employ a three-condition between-subjects (white sedans, red sedans, and sedans with cheetah print) randomized experiment to explore whether consumers prefer cheetah-inspired automobile designs more when it comes with red color or cheetah print. Similar to Study 2, after the introduction screen, each condition contains 13 screens of 9-point Likert scale questions, with each screen corresponding to one sedan model. And the slider and instructions on each page are the same as those in Study 1. Namely, the control condition with white sedans is the same as the survey in Study 1. The only difference is that, in the red color (cheetah print) condition, the target sedan models are presented in red color (with cheetah print). Figure 19 presents an example of the original look of the BMW 3 Series presented in the three conditions. The participants took 312.81 seconds on average to complete the survey. Following prior literature (Burnap et al. 2023, Luo and Toubia 2015), responses from four participants were dropped because they answered too quickly, too slowly, or in a "straight line" pattern.

4.3.2 Results

Figure 20 presents the results of the experiment. For simplicity, we only plot the average degrees of the morphing parameters across the three conditions. We discover that, compared to the control condition with white sedans, participants in the red color condition prefer sedans with higher cheetah ingredients ($\text{Mean}^{\text{control}} = 0.298$ vs. $\text{Mean}^{\text{red}} = 0.345$ and $p = 0.000$). Participants in the cheetah print condition also prefer sedans with body curves more like those of a running cheetah ($\text{Mean}^{\text{control}} = 0.298$ vs. $\text{Mean}^{\text{cheetah}} = 0.317$ and $p = 0.045$). Our results show that certain automotive colors or patterns can play as moderators for bio-inspired design, suggesting that automotive designers can combine bio-inspiration with certain colors or print patterns to further enhance the visual appeal of bio-inspired designs. These results provide actionable implications for automakers or automobile dealerships that provide customization options for consumers (e.g., offering several

color options or allowing consumers to create their own sedan decals).

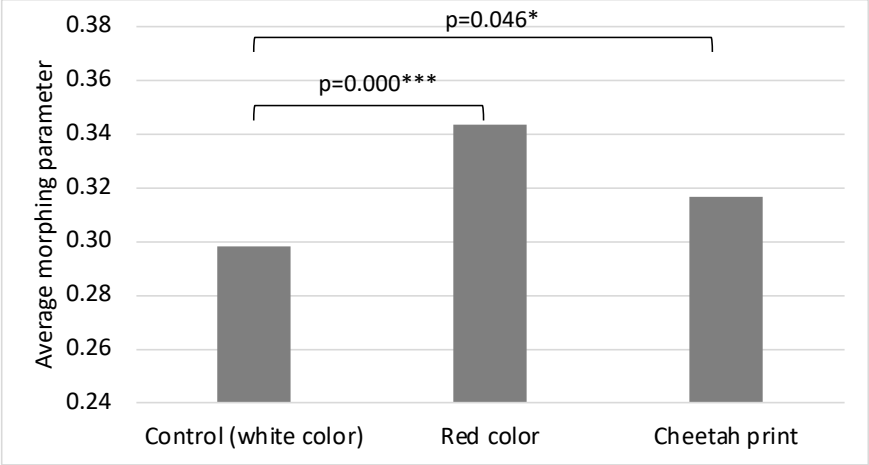

Figure 20: Users' Responses to Aesthetic Design

Significance levels: $* \ p < 0.05$, $** \ p < 0.01$, $*** \ p < 0.001$.

5 Conclusion

Our research proposes a deep learning based computational framework to automatically generate bio-inspired automotive aesthetic designs. As a proof of concept, we demonstrate this framework by facelifting the side views of sports sedans, drawing inspiration from the cheetah—the fastest land animal on Earth. Our pretest confirms that framing cheetahs as an inspiration for automotive aesthetic design enhances consumers' perceived sportiness of such a design, even when the appearance of the automobile is kept constant, which serves as the theoretical foundation for our bio-inspired design endeavor. Moreover, we demonstrate in Study 1 that the proposed framework is scalable, allows designers to control the degree of biological morphing, and offers a spectrum of new bio-inspired designs for users to choose from. As expected, we discover that consumers find the cheetah-inspired new automobile design to be more aesthetically appealing than all of the original looks in our study. In particular, higher degrees of cheetah morphing are preferred for premium sedan models (vs. mainstream models), for extra sporty models (vs. regular models), and by high-income consumers (vs. low-income consumers). Last but not least, we learn that certain colors (e.g., red) or print patterns (e.g., cheetah prints) can serve as moderators to enhance the attractiveness of cheetah-inspired automotive aesthetic designs when utilized together with cheetah-inspired

shoulder and waist curves.

Our research is among the pioneering efforts to craft a computational framework for bio-inspired automotive aesthetic design, melding theories from psychology about knowledge transfer and schema activation with generative artificial intelligence methodologies. We leverage deep learning techniques to blend design features from one domain with another, addressing the challenges of quantifying features of the inspiration source and seamlessly integrating them into product designs. This framework not only streamlines the bio-inspired design process but also offers automobile designers an efficient, scalable, and customizable tool to infuse nature-inspired aesthetic elements into automobile designs. It is our hope that our research offers a realistic, scalable, and cost-effective solution to improve the aesthetic appeal of automobiles using inspirations from running cheetahs.

Our research also provides fruitful directions for future research. First, the proposed framework can be extended to other automobile types or products for mid-generational design refreshes. For example, it can be adapted to SUVs. Because SUVs give consumers a higher driving position and off-road capability, we may incorporate the bold and brawny body curves of rhinos. Second, we can further extend the framework to automotive redesign, in which the automobile body shapes and functions will be redesigned holistically. Unlike facelifts, where the basic exterior styling and platform chassis are retained, a redesign entails extensive changes to both interior and exterior of the automobile. Future research may collaborate with automotive engineers to ensure critical engineering requirements (such as vehicle aerodynamics and safety standards) are met under such redesigns.

6 Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Abdal R, Qin Y, Wonka P (2019) Image2StyleGAN: How to embed images into the StyleGAN latent space? *Proceedings of the IEEE International Conference on Computer Vision* 2019-Octob:4431–4440.

Aggarwal P, McGill AL (2007) Is that car smiling at me? Schema congruity as a basis for evaluating anthropomorphized products. *Journal of Consumer Research* 34(4):468–479.

Aziz MS, et al. (2016) Biomimicry as an approach for bio-inspired structure with the aid of computation. *Alexandria Engineering Journal* 55(1):707–714.

Bar-Eli S (2013) Sketching profiles: Awareness to individual differences in sketching as a means of enhancing design solution development. *Design Studies* 34:472–493.

Benyus JM (1997) *Biomimicry: Innovation inspired by nature* (Morrow New York).

Bishop C (2006) *Pattern recognition and machine learning* (Springer).

Blonigen BA, Knittel CR, Soderbery A (2017) Keeping it fresh: Strategic product redesigns and welfare. *International Journal of Industrial Organization* 53:170–214.

Bouchard C, Mantelet F, Aoussat A, Solves C, Gonzalez J, Pearce K, Lottum C, Coleman S (2009) A european emotional investigation in the field of shoe design. *International Journal of Product Development* 7(1-2):3–27.

Burnap A, Hauser JR, Timoshenko A (2023) Product aesthetic design: A machine learning augmentation. *Marketing Science* .

Burnap A, Liu Y, Pan Y, Lee H, Gonzalez R, Papalambros PY (2016) Estimating and exploring the product form design space using deep generative models. *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 50107, V02AT03A013 (American Society of Mechanical Engineers).

Buxton W, Fitzmaurice G, Balakrishnan R, Kurtenbach G (2000) Large displays in automotive design. *IEEE Computer Graphics and applications* 20(4):68–75.

Catalano CE, Giannini f, Monti M, Ucelli G (2007) A framework for the automatic annotation of car aesthetics. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 1(21):73–90.

Chan HY, Boksem M, Smidts A (2018) Neural profiling of brands: Mapping brand image in consumers' brains with visual templates. *Journal of Marketing Research* 55(4):600–615.

Chandler J, Rosenzweig C, Moss AJ, Robinson J, Litman L (2019) Online panels in social science research: Expanding sampling methods beyond mechanical turk. *Behavior Research Methods* 51:2022–2038.

Chen L, Yang Y, Wang J, Xu W, Yuille AL (2016) Attention to scale: scale-aware semantic image segmentation. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 3640–3649.

Davidai S (2018) Why do americans believe in economic mobility? economic inequality, external attributions of wealth and poverty, and the belief in economic mobility. *Journal of Experimental Social Psychology* 79:138–148.

Deng Z, Lv J, Liu X, Hou Y (2023) Bionic design model for co-creative product innovation based on deep generative and bid. *International Journal of Computational Intelligence Systems* 16(1):8.

Dew R, Ansari A, Toubia O (2022) Letting logos speak: Leveraging multiview representation learning for data-driven branding and logo design. *Marketing Science* 41(2):401–425.

Dong J, Jiang W, Huang Q, Bao H, Zhou X (2019) Fast and robust multi-person 3D pose estimation from multiple views. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 7784–7793, ISSN 10636919.

Farabet C, Couprie C, Najman L, LeCun Y (2013) Learning hierarchical features for scene labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(8):1915–1929.

Fiske S (1982) Schema-triggered affect: Applications to social perception. Clark M, Fiske S, eds., *Affect and Cognition*, 55–78 (Psychology Press).

Fiske ST, Linville PW (1980) What does the schema concept buy us? *Personality and Social Psychology Bulletin* 6(4):543–557.

Girshick R (2015) Fast r-cnn. *Proceedings of the IEEE International Conference on Computer Vision*, 1440–1448.

Gonzalez RC, Woods RE, Eddins SL (2004) *Digital image processing using MATLAB* (Pearson Education India).

Goodfellow I, Bengio Y, Courville A (2016) *Deep learning* (MIT press).

Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial networks. *Advances in Neural Information Processing Systems*, 2672–2680.

Gregan-Paxton J, John DR (1997) Consumer learning by analogy: A model of internal knowledge transfer. *Journal of Consumer Research* 24(3):266–284.

He K, Gkioxari G, Dollar P, Girshick R (2017) Mask R-CNN. *Proceedings of the IEEE International Conference on Computer Vision* 2980–2988.

He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.

Hirz M, Dietrich W, Gfrerrer A, Lang J (2013) Integrated computer-aided design in automotive development. *Springer-Verl. Berl.-Heidelb. DOI* 10:978–3.

Huang X, Wang P, Cheng X, Zhou D, Geng Q, Yang R (2020) The apolloScape open dataset for autonomous driving and its application. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42(10):2702–2719.

Hucho W, Sovran G (1993) Aerodynamics of road vehicles. *Annual Review of Fluid Mechanics* 25(1):485–537.

Isola P, Zhu J, Zhou T, Efros A (2017) Image-to-image translation with conditional adversarial networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5967–5976.

Jia H, Kim BK, Ge L (2020) Speed up, size down: How animated movement speed in product videos influences size assessment and product evaluation. *Journal of Marketing* 84(5):100–116.

Karjalainen TM (2003) Strategic design language–transforming brand identity into product design elements. *Proceedings of the 10th International Product Development Management Conference*, 1–16.

Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T (2020) Analyzing and Improving the Image Quality of StyleGAN. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 8107–8116.

Kim J, Bouchard C, Bianchi-Berthouze N, Aoussat A (2011) Measuring semantic and emotional responses to bio-inspired design. *Design Creativity 2010*, 131–138 (Springer).

Kozlov A, Chowdhury H, Mustary I, Loganathan B, Alam F (2015) Bio-inspired design: aerodynamics of boxfish. *Procedia Engineering* 105:323–328.

Kreuzbauer R, Malter A (2005) Embodied cognition and new product design: changing product form to influence brand categorization. *Journal of Product Innovation Management* 22(2):165–176.

Labrecque LI, Milne GR (2012) Exciting red and competent blue: the importance of color in marketing. *Journal of the Academy of Marketing Science* 40(5):711–727.

Landwehr JR, Labroo A, Herrmann A (2011a) Gut liking for the ordinary: incorporating design fluency improves automobile sales forecasts. *Marketing Science* 30(3):416–429.

Landwehr JR, McGill A, Herrmann A (2011b) It's got the look: the effect of friendly and aggressive "facial" expressions on product liking and sales. *Journal of Marketing* 75(3):132–146.

Li G, Xie Y, Lin L, Yu Y (2017) Instance-level salient object segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 247–256.

Li G, Yu Y (2016) Deep contrast learning for salient object detection. *Proceedings of the IEEE Computer Conference on Computer Vision and Pattern Recognition* 478–487.

Lin TY, Dollár P, Girshick R, He K, Hariharan B, Belongie S (2017) Feature pyramid networks for object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 936–944.

Luo L, Toubia O (2015) Improving online idea generation platforms and customizing the task structure on the basis of consumers' domain-specific knowledge. *Journal of Marketing* 79(5):100–114.

Maeng A, Aggarwal P (2017) Facing dominance: anthropomorphism and the effect of product face ratio on consumer preference. *Journal of Consumer Research* 44(5):1104–1122.

McLaren Automotive (2021) McLaren P1 - specification. URL `https://cars.mclaren.com/gb-en/legacy/mclaren-p1/specification`.

Meyers-Levy J, Tybout A (1989) Schema congruity as a basis for product evaluation. *Journal of Consumer Research* 16(1):39–54.

Miesler L, Landwehr JR, Herrmann A, McGill A (2010) Consumer and product face-to-face: Antecedents and consequences of spontaneous face-schema activation. *ACR North American Advances* .

Miesler L, Leder H, Herrmann A (2011) Isn't it cute: an evolutionary perspective of baby-schema effects in visual product designs. *International Journal of Design* 5(3):17–30.

National Automobile Dealers Association (2016) Value discovery: How automotive brand affects used vehicle prices. Technical report.

Pan Y, Burnap A, Hartley J, Gonzalez R, Papalambros P (2017) Deep design: product aesthetics for heterogeneous markets. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1961–1970.

Pan Y, Burnap A, Liu Y, Lee H, Gonzalez R, Papalambros P (2016) A quantitative model for identifying regions of design visual attraction and application to automobile styling. *DESIGN: Proceedings of the 14th International Design Conference*, 2157–2174.

Pérez-Rojas AE, Bartholomew TT, Lockard AJ, González JM (2019) Development and initial validation of the therapist cultural comfort scale. *Journal of Counseling Psychology* 66(5):534.

Pratt W (1978) *Digital image processing* (New York: Wiley-Interscience).

Purucker C, Sprott DE, Herrmann A (2014) Consumer response to car fronts: eliciting biological preparedness with product design. *Review of Managerial Science* 8(4):523–540.

Quan H, Li S, Hu J (2018) Product innovation design based on deep learning and kansei engineering. *Applied Sciences* 8(12):2397.

Ripley R, Bhushan B (2006) Bioarchitecture: bioinspired art and architecture – a perspective. *Philosophical Transactions of the Royal Society A* 374(2073).

Sbai O, Elhoseiny M, Bordes A, LeCun Y, Couprie C (2019) DesIGN: Design Inspiration from Generative Networks. *Proceedings of the European Conference on Computer Vision* 11131 LNCS:37–44.

Schwartz M (2000) Markets, networks, and the rise of Chrysler in old Detroit, 1920-1940. *Enterprise and Society* 1(1):63–99.

Simonyan K, Zisserman A (2015) Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*.

Smith R (2012) Cheetahs on the edge. *National Geographic* URL: `https://www.nationalgeographic.com/magazine/article/cheetahs`.

Toubia O, Netzer O (2017) Idea generation, creativity, and prototypicality. *Marketing science* 36(1):1–20.

Wijegunawardana I, De Mel W (2021) Biomimetic designs for automobile engineering: a review. *International Journal of Automotive and Mechanical Engineering* 18(3):9029–9041.

Xia Y, Sattar J (2019) Visual diver recognition for underwater human-robot collaboration. *IEEE International Conference on Robotics and Automation* 6839–6845.

Xie S, Tu Z (2017) Holistically-Nested Edge Detection. *International Journal of Computer Vision* 125(1-3):3–18.

Yang J, Zhang J, Zhang Y (2021) First law of motion: Influencer video advertising on tiktok. *Available at SSRN 3815124* .

Yu Z, Feng C, Liu MY, Ramalingam S (2017) CASENet: Deep category-aware semantic edge detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 1761–1770.

Zhu J, Park T, Isola P, Efros A (2017) Unpaired image-to-image translation using cycle-consistent adversarial networks. *Proceedings of the IEEE International Conference on Computer Vision* 2223–2232.

Zhu Q, Zhang X, Luo J (2023) Biologically inspired design concept generation using generative pre-trained transformers. *Journal of Mechanical Design* 145(4):041409.

Zhu Y, Huang R, Wu Z, Song S, Cheng L, Zhu R (2021) Deep learning-based predictive identification of neural stem cell differentiation. *Nature Communications* 12(1):1–13.

Online Appendicies

A Changes Over the Recent Three Generations of BMW 3 Series

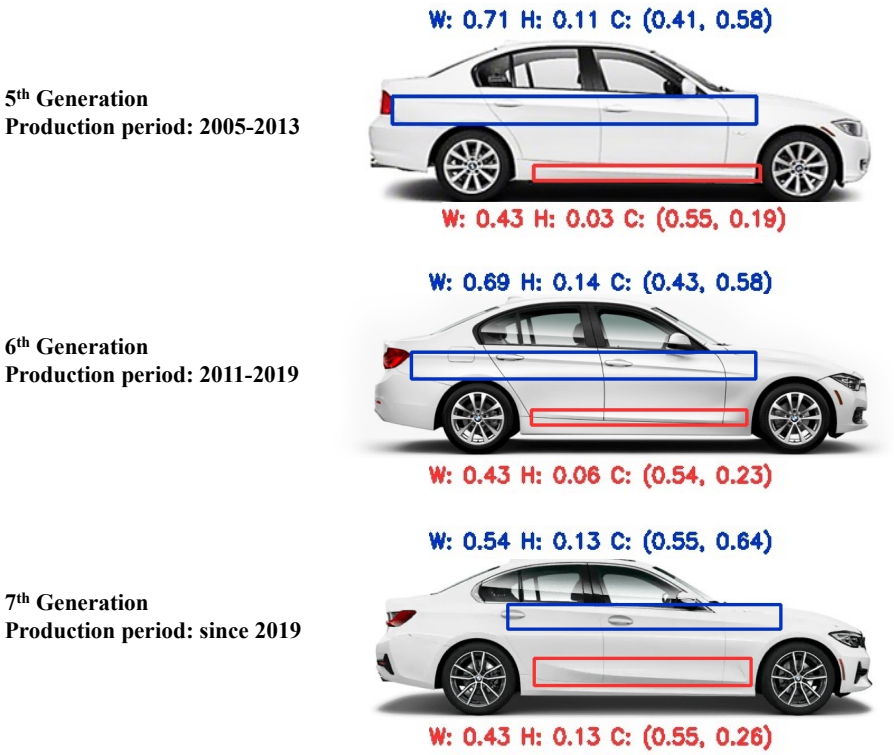

Figure A1: Changes Over the Recent Three Generations of BMW 3 Series

Note: The blue box locates the shoulder curve, and the red box locates the waist curve. W, H, and C represent the width, height, and central position of the corresponding box, relative to the car body.

B Mask R-CNN

This appendix provides technical details of the Mask R-CNN, which is a CNN-based image segmentation method. In this research, we use it in two ways: 1) detect the location of the cheetah in a video frame, and 2) identify the car body in a raw car image and remove the background.

The Mask R-CNN has two stages: (i) the region proposal network (RPN) that creates the candidate object bounding boxes; and (ii) the detection stage in which segmentation masks are predicted. In Figure A2, we use BMW 3 Series 2019 to illustrate how the Mask R-CNN works in our context.

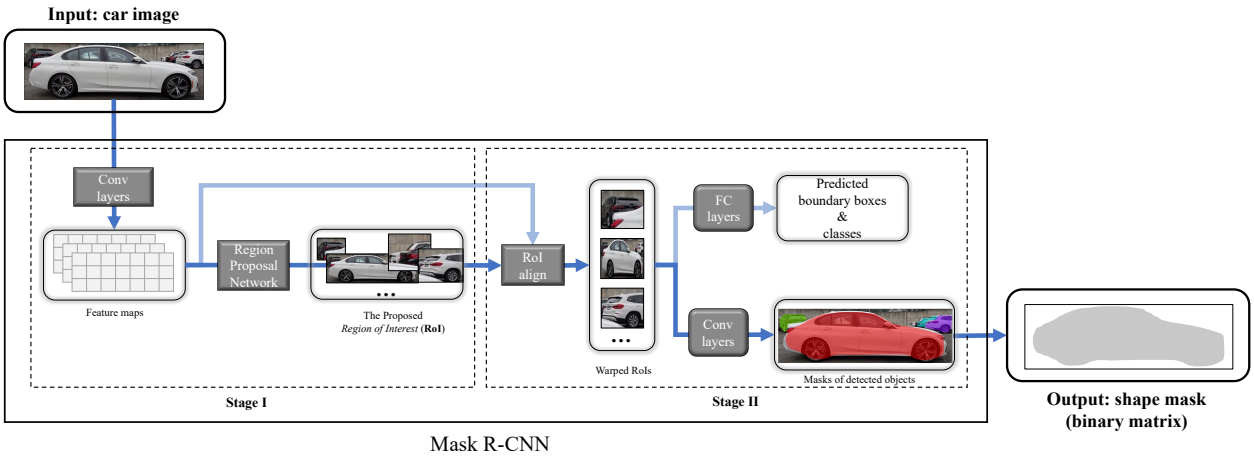

Figure A2: Schematic View of Using Mask R-CNN to Segment Car Body From Its Original Image

Note: FC = fully connected. Conv = convolution.

In the first stage, the main goal is to find bounding boxes that contain the car object. Patches, called anchor boxes, of the input image are randomly sampled. The input image will then go through a series of convolutional layers that are from well-trained existing CNNs such as ResNet (He et al. 2016) and feature pyramid network (Lin et al. 2017), resulting in matrices called *feature maps* (Goodfellow et al. 2016). Feature maps are then fed into the RPN, which proposes region candidates (i.e., updated anchor boxes) that contain car objects. Each candidate is called the *region of interest* (RoI). The RPN is often trained independently before fitting into the entire Mask R-CNN framework, and its loss function $\mathscr{L}_{\text{RPN}}$ is defined as

$$\mathscr{L}_{\text{RPN}} = \mathscr{L}_{\text{CLS}} + \mathscr{L}_{\text{BOX}}. \tag{A1}$$

The classification loss$\mathscr{L}_{\text{CLS}}$ (Girshick 2015) is specified as

$$\mathscr{L}_{\text{CLS}} = -\sum_i \left[\log(\hat{p}_i^{u_i})\right], \tag{A2}$$

where i is the index of RoIs. Each RoI is labeled with a ground-truth class u_i (e.g., a car). $\hat{p}_i^{u_i}$ is the predicted probability that anchor box i covers the ground truth class.

The bounding-box loss $\mathscr{L}_{\text{BOX}}$ (Girshick 2015) is defined as

$$\mathscr{L}_{\text{BOX}} = \sum_i \sum_j \left[\frac{1}{2} d_{(i,j)}^2\right]^{\mathbb{I}_{|d_{(i,j)}|<1}} \left[|d_{(i,j)}| - \frac{1}{2}\right]^{1-\mathbb{I}_{|d_{(i,j)}|<1}}, \tag{A3}$$

where $d_{(i,j)}$ is the difference between the predicted value and ground-truth value of RoI i's coordinate j.

In the second stage of the Mask R-CNN, our primary goal is to propose segmentation masks that indicate the location and shape of the detected car objects. The mask loss $\mathscr{L}_{\text{MASK}}$ (He et al. 2017) is specified as

$$\mathscr{L}_{\text{MASK}} = -\sum_i \sum_l \left[q_l^{u_i} \log\{\hat{q}_l^{u_i}\}\right], \tag{A4}$$

where l is the index of a pixel in an RoI. $q_l^{u_i} \in \{0,1\}$ is the ground truth label of pixel l for the ground-truth class u_i. $\hat{q}_l^{u_i}$ is the predicted probability. The Eq. (1) in the main paper specifies the linear combination of these three losses.

C SED Parameter Settings

As described in Section 3.2, the SED takes realistic car images as the input and generates sketch images. We follow Pratt (1978) and Gonzalez et al. (2004) and adopt the following parameter settings:

$$\lambda_R = 299, \quad \lambda_G = 587, \quad \lambda_B = 114,$$

$$\Phi_1 = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad \Phi_2 = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}.$$

The SED with the above settings works very well in our context. It captures the lines and shapes of the car in a given realistic image and converts them into sketches. Other more advanced edge detection algorithms (e.g., HED) can also be used in our framework. In our study, compared to HED, SED achieves comparable performance for car sketch generation while being simpler and computationally faster. Hence, we choose SED for car sketch generation.

D Sedan Models in the Training Dataset

Given our study's focus on mainstream and entry-level premium sports sedans, we selected these 50 models based on their market popularity and sales, representing the most common choices in the market. Table A1 lists all the 50 sports sedan models used for training the CycleGAN as described in Section 4.2.1.

Table A1: Sedan Models Used for Training the CycleGAN

Acura ILX	Fiat Tipo	Mazda 3
Alfa Romeo Giulia	Ford Fusion	Mazda 6
Alpine A110	Ford Mustang	Maserati Ghibli
Aston Martin Vantage	Genesis G70	Mercedes A Class
Audi A4	Honda Accord	Mercedes C Class
Audi A5	Honda Civic	Mitsubishi Mirage
Bentley Continental	Honda Insight	Nissan Sentra
BMW 3 Series	Hyundai Accent	Peugeot 508
BMW M2	Hyundai Elantra	Subaru Impreza
BuickLaCrosse	Hyundai Sonata	Tesla Model 3
Buick Verano	Infiniti Q50	Toyota Camry
Cadilac ATS	Infiniti Q60	Toyota Corolla
Cadillac CT4	Jaguar XE	Toyota Yaris
Chevrolet Camaro	Kia Forte	Volvo S60
Chevrolet Cruze	Lexus ES	Volkswagen Passat
Chevrolet Impala	Lexus IS	Volkswagen Jetta
Dodge Charger	Lincoln MKZ	