

# DVM-CAR: A Large-Scale Automotive Dataset for Visual Marketing Research and Applications

Jingmin Huang  
*School of Computing Science*  
*University of Glasgow*  
 Glasgow, UK  
 2421107h@student.gla.ac.uk

Bowei Chen  
*Adam Smith Business School*  
*University of Glasgow*  
 Glasgow, UK  
 bowei.chen@glasgow.ac.uk

Lan Luo  
*Marshall School of Business*  
*University of Southern California*  
 Los Angeles, USA  
 lluo@marshall.usc.edu

Shigang Yue  
*School of Computer Science*  
*University of Lincoln*  
 Lincoln, UK  
 syue@lincoln.ac.uk

Iadh Ounis  
*School of Computing Science*  
*University of Glasgow*  
 Glasgow, UK  
 iadh.ounis@glasgow.ac.uk

**Abstract**—There is a growing interest in product aesthetics analytics and design. However, the lack of available large-scale data that covers various variables and information is one of the biggest challenges faced by analysts and researchers. In this paper, we present our multidisciplinary initiative of developing a comprehensive automotive dataset from different online sources and formats. Specifically, the created dataset contains 1.4 million images from 899 car models and their corresponding model specifications and sales information over more than ten years in the UK market. Our work makes significant contributions to: (i) research and applications in the automotive industry; (ii) big data creation and sharing; (iii) database design; and (iv) data fusion. Apart from our motivation, technical details and data structure, we further present three simple examples to demonstrate how our data can be used in business research and applications.

**Index Terms**—Automotive analytics, product exterior design, visual marketing, big data, machine learning

## I. INTRODUCTION

The automotive industry is a major industrial and economic force worldwide. There is a rising interest in applying data mining, machine learning and artificial intelligence to automotive research and applications. Marketing researchers team up with automotive designers and computer scientists to solve the research questions and challenges in automotive exterior design, consumer analytics and sales forecasting. These are the key areas of marketing, and solving these research questions and challenges can have positive social consequences for automotive manufacturers and business researchers. On the other hand, data mining, machine learning and artificial intelligence are built upon data, in many cases, massive data. However, business details such as product sales are consistently reserved by companies and are rarely shared with researchers. Besides, data preparation can often be a lengthy, difficult and expensive process for many researchers, particularly, social science researchers' primary focuses are not on programming

and database management. Sharing reusable data with the academic community can be an effective and alternative way that makes this process more approachable.

We present our multidisciplinary initiative of creating a publicly available automotive dataset, which can be used to facilitate business research and market forecasting applications in the automotive market. Our initiative is motivated by the growing interest among the business school in big data and computer scientists are becoming interested in product aesthetics analytics and design but publicly available datasets that cover a wide range of product information, such as product appearance, model specifications and sales data, are extensively lacking. Therefore, we propose a data approach that integrates different online product information to obtain a comprehensive car dataset and hopes our work could boost the emergence of more product datasets for research purpose. In addition, three research samples are provided, illustrating the resulting dataset can comprehensively support the data needs in business analysis, product design and market forecasting studies. It is hoped the shared dataset can be used by researchers and practitioners from different disciplines to conduct economic or business related automotive research and applications.

This study makes four major contributions. First, in the context of **domain application**, our dataset meets the growing need for a comprehensive automotive dataset for economic and business research. For example, our dataset can be used for automotive exterior design and consumer analytics, which will benefit automotive manufacturers, help them better predict their targeted consumer segment preferences, and then use the obtained insights to direct the design of new exterior stylings. Our dataset can also be used for car sales forecasting which benefits all the participants in the automotive ecosystem including car dealers, consumers and marketers. Second, from a perspective of **big data**, to the best of our knowledge, our developed dataset is the very first large-scale automotive dataset. It contains 1.4 million images from 899 different car

models and corresponding specification and sales information from over ten years in the UK market. The four characteristics of big data (i.e., volume, variety, velocity, and veracity) are satisfied, which makes the dataset can be used for different types of analytics and forecasting tasks. The dataset can be used by multidisciplinary researchers to solve different tasks. Many interesting insights, data-driven models and forecasting applications can be derived. Third, in terms of **database design**, we conduct a survey study with the researchers working in either business or computer science field. The data challenges these researchers meet can be categorized into three major issues: coverage, accessibility and quality. We design the proposed dataset by addressing these issues with the hope that our data can be researcher-friendly and thus can have a large impact in research and practice. We also demonstrate a good practice of developing a dataset which can alleviate the common data issues faced by researchers. Last but not least, our work contributes to the multi-source **data fusion** as it includes different data formats from different sources. Car images, model specification and sales information are collected from different online sources while they are merged and stored in a flexible and hierarchical structure that allows it to be easily expanded with new data and used by researchers.

## II. DATA COLLECTION AND CLEANING

Nowadays researchers often collect online data and process them as new datasets for their research. Datasets like ImageNet [1] and OpenImages [2] are characterized by their enormous size and have achieved massive success in computer science research and applications. The two existing car image datasets, Stanford-Car<sup>1</sup> [3] and CompCars<sup>2</sup> [4], are also based on the web scraped contents. Motivated by these existing works, we develop the dataset by collecting and integrating data from different online sources. First, car images are collected from the popular automotive classified advertising platforms. These platforms are popular online marketplaces in the UK for buying and selling used cars, which host millions of car images and their selling prices for almost all car models from different automotive manufacturers. Second, car sales data is collected from the Driver and Vehicle Licensing Agency (DVLA), which is part of the Department for Transport, holding over 49 million driver records and over 40 million vehicle records in the UK. The DVLA publishes the statistics of newly registered vehicles in its seasonal reports. We extract and add the sales appropriately for various car models. Third, new car prices are collected from car review websites, which cover the selling prices of various car trims sold in past years.

In the preparation of data sharing, a number of data cleaning and fusion steps are performed. First, to comply with the General Data Protection Regulation (GDPR) [5], we use machine learning techniques to detect and remove privacy or unnecessary information. Mask R-CNN [6] (trained with

COCO [7]) is used to detect the appearing objects in car images. It helps to locate the positions of car bodies and detect whether other objects such as human body appear in raw images. Images with unexpected contents are dropped, and remaining images are cropped according to detected car positions. These image pre-processing steps aim to ensure resulting images are primarily occupied by car bodies without other unnecessary information. Moreover, we apply the algorithm from [8] to detect registration plates and then blur corresponding regions. Second, machine learning models are used to filter out non-exterior-viewing images. An ImageNet pre-trained convolutional neural network (CNN) [9] is fine-tuned with manually prepared samples where pictures are labelled according to their qualities and observation viewpoints. The trained CNN classifies images according to observation viewpoints. As illustrated in Fig. 1, only images taken from the eight targeting viewpoints are kept for later usage. This reduces the size of the raw image set from over six million to less than 1.5 million. Third, all the non-visual contents are processed into attributes and stored as data tables. The ambiguous values of observations are unified or corrected. All car models are assigned unique identifiers (IDs) used for information integration. Information such as annual sales is aggregated into more abstract levels. Images are listed in the table with the observation viewpoints and source advert IDs.

## III. TOWARDS A GOOD RESEARCH DATASET

During the data preparation, we deploy a survey using Qualtrics to explore the most common data issues that people face in their research and practice. We collect responses from 54 researchers, including 26 participants with a computer science background and 28 with a business studies background, including economics, marketing and management. The reported issues from our survey study can be broadly divided into three categories: coverage, accessibility, and quality.

**Coverage** refers to the issue of a dataset not containing the needed information. This happens in two situations. First, the dataset is not comprehensive enough. It does not contain the attributes or variables that a researcher is interested in for a specific study. Second, the dataset does not cover enough records or samples. Our dataset is designed to have excellent coverage of both variables and observed samples. For the former, as presented in Table I, we design the dataset to contain many important car specifications and sales variables used in the related marketing studies. In addition to these variables, as also demonstrated in Fig. 1, we provide corresponding images for various car models in various angles over a long-term period. The existing public car image datasets, such as the Stanford-car and CompCars datasets, do not include the mentioned car specification and sales variables as they are mainly designed for computer vision tasks. In terms of volume, the dataset covers millions of registered cars in the UK market in the past decades.

**Accessibility** refers to the difficulties regarding data usage caused by closed or proprietary datasets or the underlying complex data structure. Our participants in the survey

<sup>1</sup>[https://ai.stanford.edu/~jkrause/cars/car\\_dataset.html](https://ai.stanford.edu/~jkrause/cars/car_dataset.html)

<sup>2</sup>[http://mmlab.ie.cuhk.edu.hk/datasets/comp\\_cars/index.html](http://mmlab.ie.cuhk.edu.hk/datasets/comp_cars/index.html)

TABLE I: Summary of car specification and sales variables in related studies and our dataset.

Variables/features	Our dataset	Related work			
		[10]	[11]	[12]	[13]
Sales/market share	✓		✓	✓	✓
Price	✓	✓	✓	✓	✓
Brand	✓	✓	✓		✓
Exterior features	✓	✓	✓	✓	✓
Fuel economy	✓	✓		✓	
Horsepower	✓	✓		✓	
Engine	✓	✓		✓	
Transmission	✓	✓		✓	
Equipment features	✓	✓			✓
Life-cycle/model year	✓		✓	✓	✓
Advertising			✓	✓	
Reliability		✓	✓		
Safety		✓	✓		
Driving/handling		✓		✓	
Ergonomics/rooming		✓		✓	
Interior		✓			

highlighted that many datasets require license purchase, and in many cases, these datasets are sold at remarkably high prices. Researchers find that it is difficult to secure funding to purchase their needed datasets. On the other hand, the free shared datasets are often presented in a researcher-unfriendly way. The developers of those datasets may either put too much trivial information or leave data unprocessed. Even worse, many datasets are shared online without providing basic descriptions. A lot of efforts have been made to improve the accessibility of our dataset. First, we create an easily accessible data webpage on GitHub where researchers can download our dataset and find the needed description and usage instructions. Second, there are also no specific restrictions on our data usage. Moreover, we hope our data can be easily used by researchers from different backgrounds, so it is presented in a researcher-friendly manner. For example, all non-visual contents are processed into tabular attributes and organized into separate tables according to their categories. All the car models have been assigned a unique identifier (ID), which can be used for information integration and future data fusion. In addition, we remove the background of collected images, which can simplify the potential applications for researchers interested in automotive exterior design.

**Quality** refers to missing data and error values. If problems exist in the original data, they can hardly be resolved by researchers using the data. Therefore, a series of data cleaning procedures are adopted to ensure the resulting data quality. First, the ambiguous or inconsistent values are unified or corrected. For example, a car’s brand or name containing “Benz” is corrected into “Mercedes”. Second, data values at different granularities are aggregated according to certain groups. For instance, annual sales regarding various car trims, a trim level representing the equipment levels in a specific car model, are aggregated to the model levels. Third, uncertain contents are largely abandoned. The raw collected data has

more than six million images, but most of them are abandoned in later processes for the purpose of quality control.

Besides the aspects above, ethics issues are thoughtfully reviewed throughout the data preparation. We carefully preprocess and reproduce image contents to strictly comply with the GDPR (e.g., removing image background and covering plate numbers). It is worth noting that the UK government has special laws to encourage research studies related to creating and using web content based datasets<sup>3</sup>. Therefore, we do not have copyright concerns of our collected data if people use it for non-commercial purposes. As our dataset only contains car-related information (e.g., images, car model specification, sales), it seems unlikely it will generate negative societal impact.

#### IV. THE DVM-CAR DATASET

As our primary motivation to facilitate the visual marketing research and forecasting applications, we call the proposed dataset **Deep Visual Marketing Car** (in short **DVM-CAR**), which is publicly available under the **CC BY-NC** license at:

<https://deepvisualmarketing.github.io>.

The DVM-CAR dataset consists of two parts: image data and table data. The part of **image data** contains 1,451,784 car images (in JPEG format with resolution 300×300) which are compressed in a ZIP file (13.6 GB file size). These images have been selected from eight observation viewpoints and stored under the categorization “brand-model-year-color”. This structure allows researchers to locate target images easily. The backgrounds of all car images are removed. A table is prepared for image indexing, so researchers can select the images via this table. In total, our image data covers 899 car models that sold in the UK market over the last 20 years. As Fig. 2 (a) shows, the automotive classified advertising platforms have more data for newly-launched models than older car models. In spite of this, it contains 138 car models sold in the year 2000, each with an average of 28 images. The wide longitude of the dataset makes the observation of long-term trends easy, thus, is particularly useful for economic and marketing analytics and forecasting. For instance, through Fig. 2 (b), it is shown that SUVs are becoming more popular and taking higher market shares over the last two decades. The part of **table data** (i.e., non-visual part) comprises six tables in the CSV format (156 MB file size), which are named the basic, sales, price, trim, ad and image tables. Together, these tables cover various variables and form a relational database [14], as each two of them can be joined via the primary or secondary keys. More details of these tables can be found in Table II.

#### V. APPLICATION EXAMPLES

This section briefly illustrates three application examples for demonstrating how the DVM-CAR dataset could be applied to business research and applications.

<sup>3</sup><https://www.gov.uk/guidance/exceptions-to-copyright>




Fig. 1: Illustration of car images in our dataset: (a) angles of car image labeling; (b) example of background removal for Audi A2 from angle 270 degree; (c) example of background removal for Audi A2 from angle 45 degree.




Fig. 2: Summary of image data in the DVM-CAR dataset: (a) the number of car models across years and the average images for each model; (b) the annual percentage composition of images by body type.

### A. Understanding Automotive Exterior Aesthetics Design

Product aesthetics design is a determinant of consumer acceptance and product success [12], [15]–[17]. Marketing scholars have discussed aesthetics from various aspects, including the influence of aesthetics on product differentiation and new product development, and specific determinants of consumer responses to aesthetics. For example, morphing techniques were used to quantify and incorporate aesthetics design into empirical car sales models [13], [18]. This is an important step in modeling the effect of aesthetics design on sales. However, the used method is still limited in its ability to process image data. Car visual attributes are only extracted by pre-defined feature extractors, which are coarse-grained, and their respective implications for sales analysis and car appearance design are limited.

The recent advancement of machine learning technologies has provided marketing researchers with new tools to investigate product aesthetics. Several studies applied deep neural networks such as CNNs to interpret the perceived design features [19], [20]. Compared with the traditional quantitative methods used in marketing research, deep learning algorithms can automatically learn high-level representations of visual features from car image data. Thus, the deep models can be used end-to-end, which solely requires raw images and tagging

data as labels. Such advantages of deep models make them widely applied for visual based predictive studies [21].

The DVM-CAR dataset provides an excellent base for researchers to apply deep learning to extract visual attributes from car images. For instance, Fig. 3 (b) shows the sample using car images to infer the modernity score (i.e., labels computed from car registration years). By fine-tuning the ImageNet pre-trained VGG'16 [22], the results show deep learning models are capable of predicting the design fashion for family cars. Besides, we can further recognize the design language or patterns from the inside gradients. Fig. 3 (a) displays the predicted scores for Land Rovers. Although all these models are from the same automaker and were sold in the same year, the trained deep model rates them with entirely different modernity scores. The second row of Fig. 3 (a) presents, with the help of visualization methods such as guided back-propagation [23], researchers can identify specific designs (i.e., highlighted in blue color) that make the appearance outdated.

Based on the car models' modernity predictions, we can further investigate the associations between cars' modernity of appearance and their market performance by using the sales records in the DVM-CAR dataset. As Table III illustrates, we find the models' future withdrawn chances are correlated to their modernity scores (Note: the numbers in brackets of Table III are the sample sizes of the groups). The car group

TABLE II: Description of table data in the DVM-CAR dataset.

Name	Table description	Attribute name	Attribute description
Basic	It is mainly for indexing other tables.	Genmodel	Generic model name
		Genmodel ID	Generic model ID
		Automaker	Automaker name
		Automaker ID	Automaker ID
Ad	It contains more than 0.27 million used car advertisements information posted on the automotive classified advertising platforms, including variables like advertisement's creation time, used car registration year, cumulative mileage, selling price, etc.	Genmodel	Generic model name
		Genmodel ID	Generic model ID
		Maker	Automaker name
		Adv ID	Advertisement ID
		Adv year	Advertisement's creation year
		Adv month	Advertisement's creation month
		Color	This car's color
		Reg year	This car's first registration/selling year
		Bodytype	This car's body type
		Runned Miles	This car's runned mileage
		Engin size	This car's engin size
		Gearbox	This car's gearbox
		Fuel type	This car's fuel type
		Price	This car's selling price
		Seat num	This car's seats number
		Door num	This car's doors number
Image	It contains image data related information like predicted viewpoint and quality check result.	Genmodel ID	Generic model ID
		Image ID	Image ID
		Image name	Image name
		Predicted viewpoint	This image's predicted viewpoint
Price	It contains the entry-level new car prices. It is designed for people who only need the basic price of car models.	Quality check	Manually check result
		Genmodel	Generic model name
		Genmodel ID	Generic model ID
		Maker	Automaker name
Sales	It contains car sales data of the UK market (based on the released statics from the DVLA).	Year	Generic model's selling year
		Entry price	Generic model's entry-level price
		Genmodel	Generic model name
		Genmodel ID	Generic model ID
Trim	It includes 0.33 million trim level information such as selling price, fuel type and engine size. It is designed for people who are interested in the price of car model at a specific trim level.	Year 2001 to 2020	Generic model's annual sales
		Maker	Automaker name
		Trim	Trim name
		Year	Trim's selling year
		Price	Trim's price at selling year
		Gas emission	Trim's CO <sub>2</sub> emission
		Fuel type	Trim's fuel type
		Engine size	Trim's engine size

with low modernity scores have a higher chance of being withdrawn from the market, while the high modernity group tend to survive longer. As essential marketing variables such as new car prices and sales are all covered by the dataset, users can investigate the association between cars' exterior styling and market performance from diverse perspectives.

### B. AI-Powered Automotive Exterior Design


Anthropomorphism [24]–[26] refers to the attribution of human or animal characteristics to non-living objectives. As a common phenomenon in consumption environments, anthropomorphism has drawn wide attention from both social psychology and marketing researchers. Psychological studies [27] reveal that our brain is highly specified for face perception, which is so evolved that we often perceive faces

from non-living objects. A group of marketing studies extensively investigate how products with human-like or animal characteristics lead to face perception. Typically, in car related studies, existing investigations [24], [28], [29] show consumers have a strong tendency to anthropomorphize the car front. However, in these studies, human or animal characteristics are predefined and only used in user surveys. They cannot be further deployed for facial feature recognition for unseen images. Thus, providing limited insights on car front appearance design for car manufacturers.


As a group of generative models with neural network structures, deep generative models can learn to generate highly realistic representations from data, thus making them a hotspot in machine learning and statistics in past years [30]. These models have become extremely successful in the applications




(a)



(b)



(c)



(d)

Fig. 3: Illustration of application examples: (a) modernity scores of Land Rover models registered in 2015; (b) distribution of the predicted modernity scores of car models from 2000 to 2015; (c) new BMW front designs generated by CycleGAN according to given facial shapes; (d) automobile design sketches generated by CycleGAN.

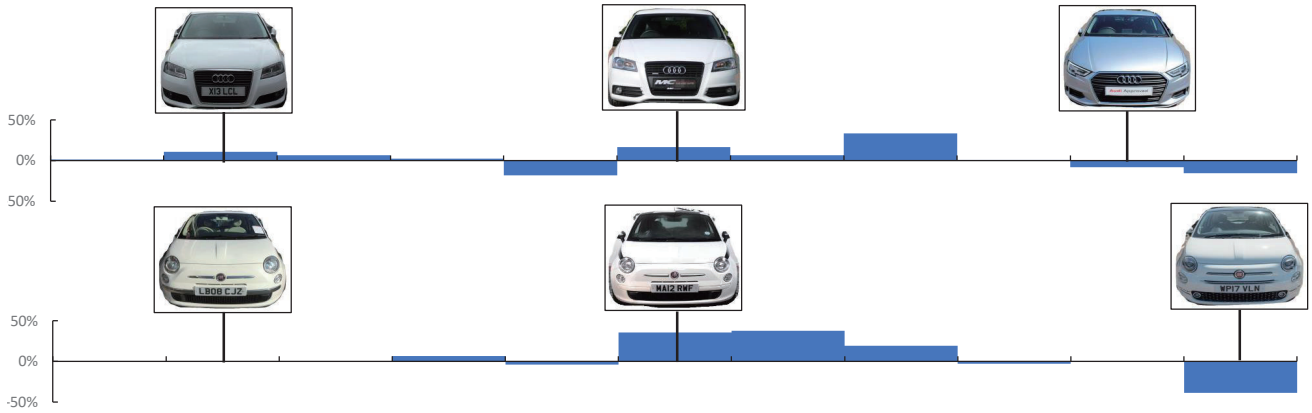


Fig. 4: Predicting the changes of market share changes according to design changes. The displayed two samples are the historical Audi A3 and Fiat 500 in the last decade. The position of displayed car images indicates the year for model redesign, the blue bar indicates the predicted market share change.

such as image translation and fake data generation. Particularly, marketing researchers apply these methods to generate novel product designs [31]–[33], including cars. Pan et al. [34] and Burnap et al. [35] used the deep generative models in the combination of design evaluators (i.e., deep convolution models trained for aesthetics evaluating) to generate appealing car appearance designs. The 1.4 million car images of the DVM-CAR dataset make it extremely useful for generative model training. Two examples are provided in Fig. 3 (c) and (d). The former shows new BMW 3 Series front designs generated by cycle-consistent adversarial networks (in short CycleGAN) [36] based on the fed semantic views. The latter presents sketches generated from real car images via the CycleGAN. These are from our working topics where we try to obtain bio-inspired designs by morphing the cars’ layouts after cheetah faces.

### C. Visual-based Car Sales Forecasting

Several existing studies [37]–[40] investigate the product exterior design as an optimization problem that searches for best-selling designs with considering aesthetic attributes as controllable factors in manufacturing. For example, Rajeev and Krishnamurti [38] simulated market sales using the share-of-choice setting that viewed each product as a profile with a collection of design attributes at different discrete levels. Due to their heterogeneous needs, the consumers are willing to pay more for the product with matched profile and thus more likely to buy. Differently, Kaul and Rao [39] considered consumers’ heterogeneous needs as coordinates distributed in a high-dimensional space of product attributes, where utilities of products are indicated by their distances from the coordinates. Their study tried to infer the best coordinates in the market that maximize the overall purchase likelihood. However, as designs are represented by the profiles in these existing studies, no actual designs are generated. The complexity of real scenarios is thus too simplified to be applicable. For instance, the existing studies didn’t clarify how to compute an aesthetic profile of a given design and whether particular aesthetic

TABLE III: Withdrawn chances within the next three years of car groups with different modernity scores.

Year	Modernity score				
	[0,1)	[1,2)	[2,3)	[3,4)	[4,5]
2000	12.0%(25)	2.8% (36)	0.0% (10)	-(2)	-(0)
2003	23.1%(13)	10.7%(56)	0.0%(28)	0.0%(10)	-(0)
2006	-(4)	13.9%(36)	8.3%(60)	7.4%(27)	-(1)
2009	-(1)	33.3%(9)	16.1%(62)	3.3%(61)	15.8%(19)
2012	-(0)	-(1)	17.9%(28)	9.1%(88)	5.4%(56)
2015	-(0)	-(0)	18.2%(11)	9.0%(67)	7.3%(109)

profiles are not achievable (e.g., a super sporty and cute design).

One of our working studies incorporates the deep learning-based design evaluation and generation procedure into product design optimization. Historical images and sales of the DVM-CAR are used to predict the changes in sales when adopting a new design. Being trained with massive car images, the GAN model is capable of generating various novel yet highly realistic car designs. Meanwhile, after collecting aesthetic ratings from surveys, a CNN model is trained with these aesthetic scores and car images to assess designs in terms of aesthetics. Then the trained CNN is used to evaluate both real designs from the dataset and the ones generated by the GAN model. For sales forecasting, an RNN model is trained with these obtained aesthetic scores and annual sales records. Thus, recognizing a car model’s prior sales and its aesthetic changes, the trained RNN can forecast the posterior sales (or market shares)., see Fig. 4 for a series of the predicted market share changes of Audi A3 and Fiat 500.

## VI. CONCLUSION

This paper demonstrates the design and development of a large-scale dataset for business research with online sources. As highlighted in Section I, our work has significant contributions for information and knowledge management, ranging from domain application to big data, to database design, and to data fusion. The targeted users of our dataset are business

researchers and computer scientists who work in visual-related research and applications, particularly on (but not limited to) the topics of automotive exterior design, consumer analytics and sales prediction. Through this paper, we would like to make an announcement that the DVM-CAR is now released publicly to research communities. We hope it can have a large impact. More importantly, we would like to maintain and keep updating the dataset by incorporating researcher's feedback over time.

#### ACKNOWLEDGMENT

The first author acknowledges the Adam Smith Business School and the School of Computing Science of University of Glasgow's funding support for this research. The second author would like to thank the funding support of the Region Bourgogne Franche Comté Mobility Grant, Nvidia Accelerated Data Science Grant and Google Cloud Academic Research Grant.

#### REFERENCES

- [1] J. Deng, W. Dong, R. Socher, L. Li, L. Kai, and L. FeiFei, "ImageNet: A large-scale hierarchical image database," *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- [2] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov, T. Duerig, and V. Ferrari, "The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale," *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1956–1981, 2020.
- [3] J. Krause, M. Stark, J. Deng, and L. Feifei, "3D object representations for fine-grained categorization," *IEEE International Conference on Computer Vision*, pp. 554–561, 2013.
- [4] L. Yang, P. Luo, C. C. Loy, and X. Tang, "A large-scale car dataset for fine-grained categorization and verification," *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3973–3981, 2015.
- [5] P. Voigt and A. Von dem Bussche, *The EU general data protection regulation (GDPR): A practical guide*, 1st ed. Springer International Publishing, 2017.
- [6] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask R-CNN," *IEEE International Conference on Computer Vision*, pp. 2980–2988, 2017.
- [7] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and L. Zitnick, "Microsoft COCO: Common objects in context," *European Conference on Computer Vision*, pp. 740–755, 2014.
- [8] S. M. Silva and C. R. Jung, "License plate detection and recognition in unconstrained scenarios," *European Conference on Computer Vision*, pp. 580–596, 2018.
- [9] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [10] M. Kukova, C. Diels, P. Jordan, M. Franco-Jorge, J. Anderson, and H. Kharouf, "Do we really know which vehicle attributes are important for customers?" *International Conference on Design Emotion*, 2016.
- [11] O. Korenok, G. Hoffer, and E. Millner, "Non-price determinants of automotive demand: Restyling matters most," *Journal of Business Research*, vol. 63, no. 12, pp. 1282–1289, 2010.
- [12] R. Jindal, K. Sarangee, R. Echambadi, and S. Lee, "Designed to succeed: Dimensions of product design and their impact on market share," *Journal of Marketing*, vol. 80, no. 4, pp. 72–89, 2016.
- [13] J. R. Landwehr, A. Labroo, and A. Herrmann, "Gut liking for the ordinary: Incorporating design fluency improves automobile sales forecasts," *Marketing Science*, vol. 30, no. 3, pp. 416–429, 2011.
- [14] E. F. Codd, "A relational model of data for large shared data banks," *Communications of the ACM*, vol. 26, no. 1, pp. 64–69, 1983.
- [15] G. E. Hoffer and R. J. Reilly, "Automobile styling as a shift variable: An investigation by firm and by industry," *Applied Economics*, vol. 16, no. 2, pp. 291–298, 1984.
- [16] P. H. Bloch, "Seeking the ideal form: Product design and consumer response," *Journal of Marketing*, vol. 59, no. 3, pp. 16–29, 1995.
- [17] J. P. Schoormans and H. S. Robben, "The effect of new package design on product attention, categorization and evaluation," *Journal of Economic Psychology*, vol. 18, no. 2-3, pp. 271–287, 1997.
- [18] I. Tseng, J. Cagan, K. Kotovsky, and M. Wood, "Form function fidelity," *Journal of Mechanical Design*, vol. 135, no. 1, p. 11006, 2013.
- [19] Y. Pan, A. Burnap, Y. Liu, H. Lee, R. Gonzalez, and P. Papalambros, "A quantitative model for identifying regions of design visual attraction and application to automobile styling," *International Design Conference*, 2016.
- [20] A. Burnap, Y. Pan, Y. Liu, Y. Ren, H. Lee, R. Gonzalez, and P. Y. Papalambros, "Improving design preference prediction accuracy using feature learning," *Journal of Mechanical Design*, vol. 138, no. 7, p. 071404, 2016.
- [21] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *International Conference on Learning Representations*, pp. 1–14, 2015.
- [23] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *ArXiv e-prints*, p. 1412.6806, 2014.
- [24] L. Miesler, H. Leder, and A. Herrmann, "Isn't it cute: An evolutionary perspective of baby-schema effects in visual product designs," *International Journal of Design*, vol. 5, no. 3, pp. 17–30, 2011.
- [25] A. Waytz, J. Heafner, and N. Epley, "The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle," *Journal of Experimental Social Psychology*, vol. 52, pp. 113–117, 2014.
- [26] S. Ku, "Is this car looking at you? How anthropomorphism predicts fusiform face area activation when seeing cars," *PLoS one*, vol. 9, no. 12, pp. 1–14, 2014.
- [27] C. Purucker, D. E. Sprott, and A. Herrmann, "Consumer response to car fronts: Eliciting biological preparedness with product design," *Review of Managerial Science*, vol. 8, no. 4, pp. 523–540, 2014.
- [28] P. Aggarwal and A. L. McGill, "Is that car smiling at me? Schema congruity as a basis for evaluating anthropomorphized products," *Journal of Consumer Research*, vol. 34, no. 4, pp. 468–479, 2007.
- [29] J. R. Landwehr, A. McGill, and A. Herrmann, "It's got the look: The effect of friendly and aggressive "facial" expressions on product liking and sales," *Journal of Marketing*, vol. 75, no. 3, pp. 132–146, 2011.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- [31] A. Burnap, Y. Liu, Y. Pan, H. Lee, R. Gonzalez, and P. Y. Papalambros, "Estimating and exploring the product form design space using deep generative models," *Design Automation Conference*, pp. 1–13, 2016.
- [32] H. Quan, S. Li, and J. Hu, "Product innovation design based on deep learning and kansei engineering," *Applied Sciences*, vol. 8, no. 12, p. 2397, 2018.
- [33] C. Yu, Y. Hu, Y. Chen, and B. Zeng, "Personalized fashion design," *IEEE International Conference on Computer Vision*, pp. 9045–9054, 2019.
- [34] Y. Pan, A. Burnap, J. Hartley, R. Gonzalez, and P. Papalambros, "Deep design: Product aesthetics for heterogeneous markets," *International Conference on Knowledge Discovery and Data Mining*, pp. 1961–1970, 2017.
- [35] A. Burnap, J. R. Hauser, and A. Timoshenko, "Design and evaluation of product aesthetics: A human-machine hybrid approach," *ArXiv e-prints*, p. 1907.07786, 2019.
- [36] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," *IEEE International Conference on Computer Vision*, pp. 2223–2232, 2017.
- [37] P. E. Green, J. D. Carroll, and S. M. Goldberg, "A general approach to product design optimization via conjoint analysis," *Journal of Marketing*, vol. 45, no. 3, p. 17, 1981.
- [38] K. Rajeev and R. Krishnamurti, "A heuristic approach to product design," *Management Science*, vol. 33, no. 12, pp. 1523–1533, 1987.
- [39] A. Kaul and V. R. Rao, "Research for product positioning and design decisions: An integrative review," *International Journal of Research in Marketing*, vol. 12, no. 4, pp. 293–320, 1995.
- [40] L. Shi, S. Ólafsson, and Q. Chen, "An optimization framework for product design," *Management Science*, vol. 47, no. 12, pp. 1681–1692, 2001.