Understanding Perceptual and Conceptual Fluency at a Large Scale

Shengli Hu[1(✉)] and Ali Borji[2]

[1] Cornell University, Ithaca, NY 14850, USA
sh2264@cornell.edu
[2] University of Central Florida, Orlando, FL 32816, USA
aborji@crcv.ucf.edu

Abstract. We create a dataset of 543,758 logo designs spanning 39 industrial categories and 216 countries. We experiment and compare how different deep convolutional neural network (hereafter, DCNN) architectures, pretraining protocols, and weight initializations perform in predicting design memorability and likability. We propose and provide estimation methods based on training DCNNs to extract and evaluate two independent constructs for designs: perceptual distinctiveness ("perceptual fluency" metrics) and ambiguity in meaning ("conceptual fluency" metrics) of each logo. We provide evidences of causal inference that both constructs significantly affect memory for a logo design, consistent with cognitive elaboration theory. The effect on liking, however, is interactive, consistent with processing fluency (e.g., Lee and Labroo (2004), and Landwehr et al. (2011)).

Keywords: Marketing application · Visual design
Cognitive information processing · Construal level theory

1 Introduction and Related Work

Which visual elements of a design increase memory for the design and which increase liking for it? Are these elements the same, or different—and what is the relationship between remembering and liking?

Ample research in cognitive psychology shows that an increased elaboration about the target is associated with a better memory for it, for instance, [28, 29, 31]. For a visual design, an increased elaboration could result from two sources—first, the perceptual fluency of the design and second, the conceptual fluency of the design. By perceptual fluency of a design, we refer to the purely perceptual elements of the design and the extent to which these might deviate from a visual prototype for the category that perceivers have in memory.

Existing research shows that visual representations tend to be stored in memory as deviations from category prototypes (e.g., [4, 10]). As people are exposed through their lifetimes to more exemplars of a category, they develop a prototype in memory. When new exemplars are encountered, the visual memory system

checks for and encodes differences from the prototype rather than remember every detail. The second aspect of a visual design that can affect memory for it is its meaningfulness. People remember information by associating its meaning to what they already know [5,26,27]. When the meaning of an attitude object[1] is very clear, it is associated easily into the knowledge a consumer already has in the area. But the attitude object draws little elaboration and its meaning is instantaneously incorporated into a gist of what the consumer already knows. The lack of elaboration might reduce the distinct memory for the object or the recognition probability.

How might these two aspects of visual design—its perceptual fluency and its conceptual fluency—impact liking for the logo? Research in cognitive psychology suggests that consumers like attitude objects that feel easy to process [1,7,20–22]. The idea is that information that feels easy to process because its information is perceptually clear and stands out or because its meaning is easy to elaborate on is liked more. These findings imply that logos that are visually distinct are likely to pop out and be liked more, and if they encourage meaningful elaboration, they will be liked even more. But if the meaning is very ambiguous they will be liked less.

To explore these ideas, we create a large dataset of 543,758 logo designs spanning 39 industrial categories and 216 countries. We annotate it by collecting and calculating likability and memorability scores from a series of online visual memory experiments adapted from procedures introduced by [11,12,17]. On a computing front, we first experiment with prediction tasks of design memorability and likability, using different DCNN architectures, pre-training, and transfer learning methods. On a cognitive front, we propose information measures to proxy for perceptual and conceptual fluency based on DCNN outputs. We then regress design memorability and likability on the proposed measures of fluency constructs, controlling various other aspects to identify the effects of perceptual and conceptual fluencies on memorability and likability. We ensure *identification* of both measures by running *controlled*[2] online experiments when eliciting annotations.

Major contributions of the present study include: it is the first study to (1) release a large-scale logo design dataset, valuable to possibly both the vision and business research communities; (2) investigate the effect of both fluencies on memorability and likability for designs at a large scale; and (3) compare how different deep neural network architectures, initializations, and transfer learning methods perform on the prediction tasks of design memorability and likability.

---

[1] Attitude objects are what you make a judgment about or have a positive or negative feeling toward.

[2] We follow the experimental economics protocols of randomization and clear treatment grouping based on perceptual and conceptual measures, on top of what was prescribed in [14–17].

2   Large-Scale Memorability and Likability Logo Design Dataset

We build a large dataset of 543,758 logo designs, with memorability and likability scores on a sub-sample of 123,928 logos.

2.1   Memorability and Likability Scores

We obtained memorability scores for a subset of the entire dataset using an online experimental procedure that is adapted from the efficient visual memory game developed in [14–17]. We elicit liking scores by asking subjects to rate how much do they like the logos being shown on a scale from 0 to 7. To guard against potential biases introduced by consumers' familiarity with the brand or the company associated with the logo, we first explicitly asked them not to base their decision on how much they like the brand for reasons other than esthetics. Secondly, likability scores were elicited as part of a post-experiment questionnaire, that included a question about subject's familiarity of the brands being shown. Surprisingly, the percentage of encounters of familiar logos or brands account for less than 3% ((3,711/123,928), we set the recognition threshold to 40% to reduce noise) in our annotated set, indicating that such a bias is secondary in our setting.

We recruited $38,542$ US-based subjects from Amazon Mechanical Turk and obtained 66 scores of both memorability and likability from each Turker, resulting in around 20 scores per logo. We calculated the estimated memorability and likability scores for each logo following the optimization procedure adopted in [17]. In the game, Turkers view a sequence of logos, each of which is displayed for 1 s, with a 1.4 s gap in between logo representations. Their task is to press the button whenever they see a repeat of a logo. Each task is designed to last about 4.5 min consisting a total of 186 logos divided into 66 targets, 30 filters, and 12 vigilance repeats. Targets are repeated after at least 35 logos, and at most 150 logos. Vigilance repeats are shown within 7 logos from the first showing. The vigilance repeats ensure that Amazon Mechanical Turkers are indeed paying attention. Turkers failing more than 25% of the vigilance repeats are blocked and all their answers are discarded. We refer readers to [15,17] for more procedural details.

Nonetheless, our experiments differ in at least two main ways. First, we use a $2 \times 2 \times 2 \times 2$ full factorial design and examine 16 different between-subject treatments varying the levels (above median or below median) of perceptual Fluency measures (Perceptual Complexity and Image Prototypicality, detailed in Sects. 1 and 3.2) and Conceptual Fluency measures (Entropy and Kullback-Leibler Divergence from deep neural network feature distributions) with a logo memory game, following experimental economics protocols, which provides greater control and facilitates causal inference. Second, rather than repeating identical natural images at variable time intervals, we repeat logos in different forms (logo pattern or plain text towards the end of the experimental

Fig. 1. Sample images arranged by their memorability and likability scores: top left logos are of high memorability scores, top right of low memorability scores, bottom left of high likability scores, bottom right of low likability scores

session) and elicit memory scores accordingly. Figure 1 shows sample designs from our dataset arranged by annotated scores.

2.2 Summary of the Dataset

We collected vector logos from various sources on the Internet courtesy of Brands of the World, Logo Types, World Vector Logo, Vector Me. Our dataset contains logo designs spanning 39 industry categories and 216 countries. Access is available at Logo Designs. We present the distributions of industrial categories and countries based on the annotated subset in Fig. 2. We detail the distribution of logos across the most concentrated countries and categories in the spider charts of Fig. 3. Most logos of the design category are from Brazil, US, Turkey, Mexico and Europe, whereas logos of food and drinks are from North and South America. Sports logos are more even distributed among European countries and South America with greater concentration on Brazil and US. Brazilian logos are mostly concentrated in sports, design, and business, whereas US logos are more evenly distributed across categories including food and drinks, technology, media, automobile, music and finance in our annotated sample from the larger dataset.
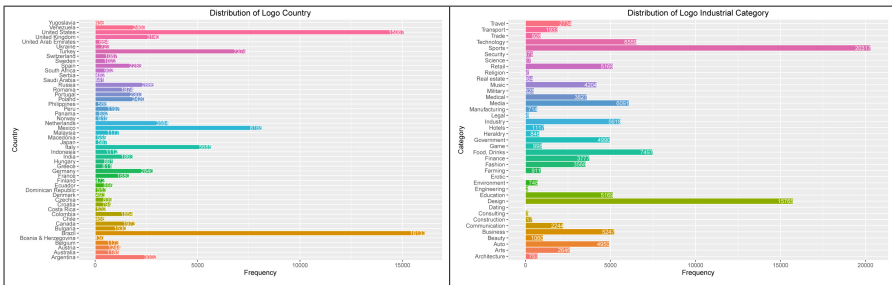

Fig. 2. Histograms of logo industrial categories and countries
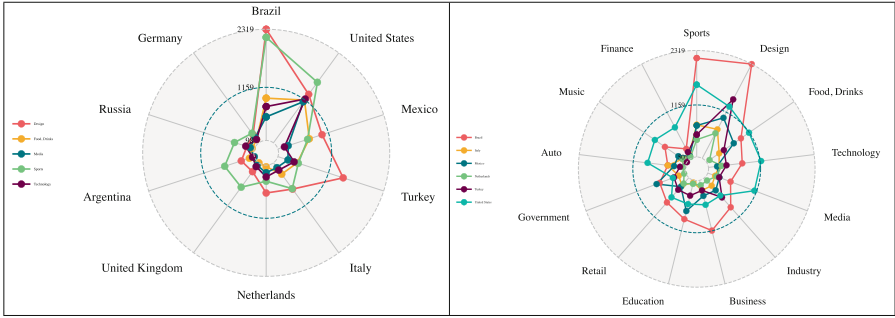
Fig. 3. Spider chart of logo distributions across countries (categories) by category (country)

2.3 Summary of Image Features

We explore a variety of image features of our dataset, including size (entropy), hue, saturation, value, the number of edges, the number of straight lines, the number of corners, the number of circles, number of polygons, and the presence of texts. We measure the size by taking the average of local entropy values of all pixels within an image, and plot the distribution by country and category. Figure 4 shows the comparative distribution of the most concentrated countries (top six) and categories (top ten). Interestingly, logos of the technology category appear to be much more concentrated on the lower end of the size spectrum, whereas media logos and food/drinks logos show the greatest variations in terms of size. Across countries, logos from European countries appear more concentrated towards the smaller sizes whereas the overall variation across countries are much less pronounced than that across industries, as are shown in Fig. 5.
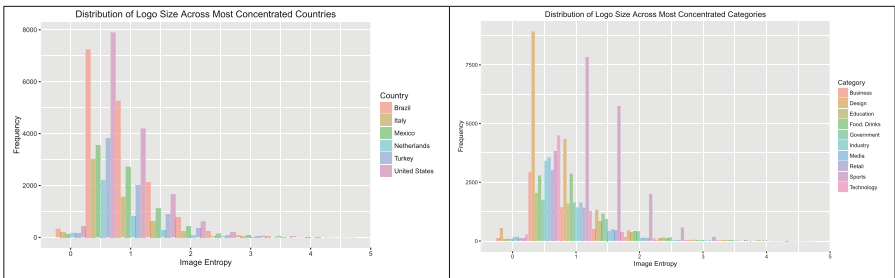

Fig. 4. Histograms of logo size distributions across most concentrated countries and categories

3D plots of representative countries and categories are shown in Fig. 6, where the hue, saturation, and value values are the averages taken across all pixels of
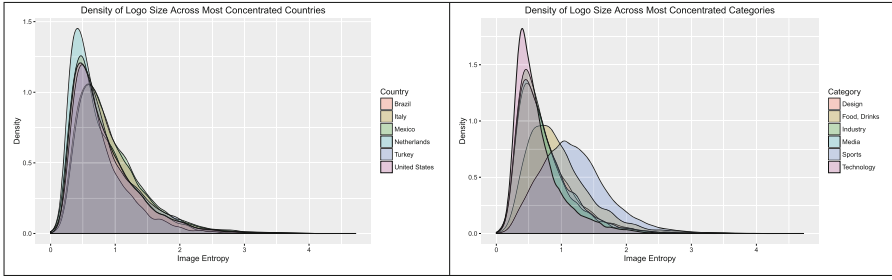
Fig. 5. Density plots of logo size distributions across most concentrated countries and categories

an image. Some descriptive patterns include: sports logos are distributionally more saturated than design logos, logos from North American countries are distributionally both more saturated and less concentrated in hues than logos from European countries.
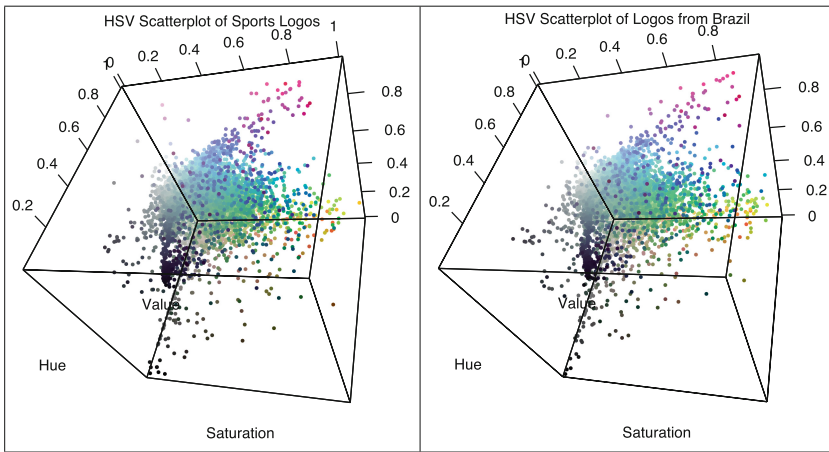
Fig. 6. 3D scatterplots of mean HSV values of some countries and categories

We also measure the number of edges with Canny detector, the number of corners with Harris detector, the number of straight lines, circles, and polygons with probabilistic Hough line transformation. We plot a selection of empirical densities of the resulting number of edges, circles, straight lines and polygons in Fig. 7. Some interesting patterns emerged: (1) Logos of the Design industry appear to adopt less straight lines and polygons, whereas Sports logos appear to exhibit polygons much more frequently; (2) Sports, Media, and Food/Drinks Logos appear to showcase more circles whereas the opposite goes for Retail and Technology logos; (3) American (both North and South) logos appear to adopt
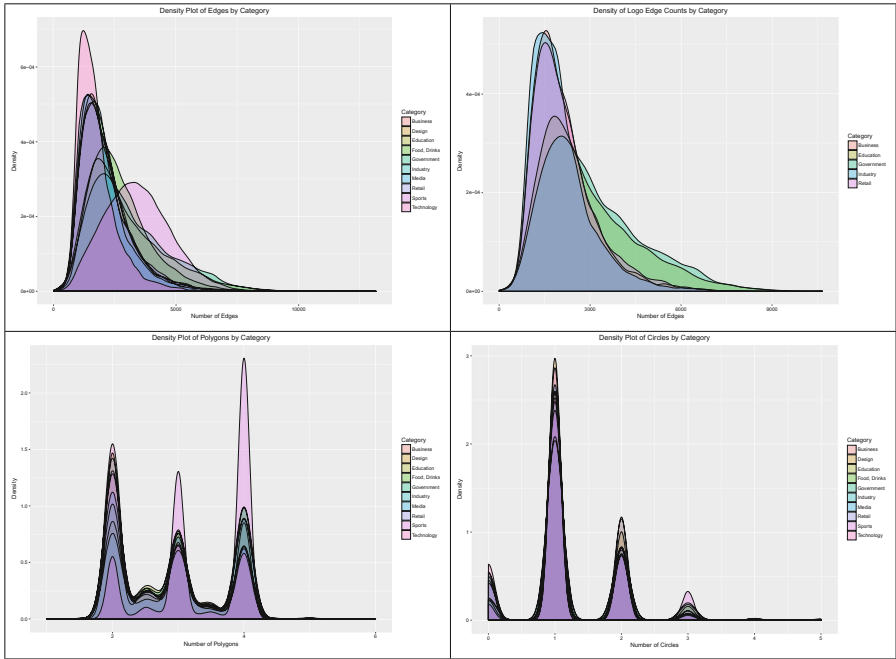
Fig. 7. Density plots of low-level features by category or country

more circles and polygons than European logos, with South American logos showcasing more edges.

2.4   Image Features and Memorability (Likability)

We did a preliminary analysis that correlats low level features with both memorability and likability. To simplify interpretations, we summarize identified relationships from our model-free results in Table 1, where − means negative, U means U-shaped curve, Inverse U means inverse U-shaped curve, and NS means Non-significant.

Table 1. Low-level features and memorability/likability

Relationship	#Edge	#Line	#Circle	#Polygon	Hue	Saturation	Value
Memorability	Linear +	U	Inverse U	Inverse U	−[a]	NS	NS
Likability	Linear +	U	Inverse U	U	NS	NS	NS

[a] On a representative subsample

3 Methods

To explore how different DCNNs perform on our design dataset in predicting
memorability and likability, we experiment with five network architectures, three
pre-training options and two initialization methods, which we detail in Sect. 3.1.
To better understand how *perceptual* and *conceptual fluency* affect design mem-
orability and likability, we propose DCNN-based measures for both fluency con-
structs and regress non-linearly both memorability and likability scores on flu-
ency measures, along with a myriad of control variables. We detail this more
economic design in Sect. 3.2. Compared to previous established measures of pro-
totypicality in perception literature, our proposed method is better suited given
the much bigger scale, and the abstract nature of our unique dataset distinct
from previous studies in perception.

3.1 Predicting Memorability and Likability

We distribute the annotated data into 5 random train and test splits. We run a
series of experiments varying three aspects:

1. The network architectures:
   AlexNet [18], VGG16 [24], VGG19 [24], Inception-V3 [25], and ResNet-50 [9];
2. Pre-training: ImageNet [2] pre-trained model with transfer learning (only the
   last layer was re-trained), ImageNet pre-trained model followed by fine-tuning
   (the last three layers were re-trained for 15 and 20 epochs for memorability
   and likability prediction, respectively), training from scratch without pre-
   training; a Euclidean loss layer is used since memorability or likability is a
   single real-valued output;
3. Weight initialization: random initialization drawn from Gaussian distribu-
   tions [19] or robust initialization proposed in [8];

## 3.2 Measuring and Identifying Perceptual and Conceptual Fluencies

Conceptual Fluency: Entropy and Kullback-Leibler divergence Recall
that we define conceptual fluency as if/how the underlying brand identity is
perceived, for which we propose two information-theory based measures using
DCNN prediction results: entropy and Kullback-Leibler divergence. Specifically,
we train a deep residual network [9] on our dataset in the same way as described
in Sect. 3.1, except that the output of the last fully-connected layer in fine-
tuning is fed to a 39-way softmax layer at the end, representing the 39 different
industrial identities in our setting. Cross-entropy losses replace the Euclidean
losses since the ground-truths are one-shot vectors. Given the output distribution
of our fine-tuned ResNet for logo category classification, denoted as Q, and the
true category label, denoted as P, we define our measure regarding conceptual
fluency as the Kullback-Leibler divergence from P to Q, written as $D_{KL}(P|Q)$.

The other dimension we extract for measuring *conceptual fluency* is the Shannon entropy of Q, $H(Q)$, the output distribution of the deep network as our measure:

$$D_{KL}(P|Q) = \sum_{c=1}^{C} P_c \log\left(\frac{P_c}{Q_c}\right), H(Q) = -\sum_{c=1}^{C} Q_c \log(Q_c) \tag{1}$$

where P represents the one-hot vector representation of the true category label and Q represents the inferred category label distribution given by the DCNN. We plot the empirical distributions of proposed measures of conceptual fluency in Fig. 8. Intuitively, KL on conceptual fluency should reflect the deviation (or surprise) of the perceived distributional identity from real identity, whereas entropy on conceptual fluency reflects only the perceived distributional identity. Thus KL measure is susceptible to industrial norms about logos, and its "power law" pattern indicates strong conformity within industry, whereas Entropy measure adheres to Law of Large Numbers, free of industrial constraints.
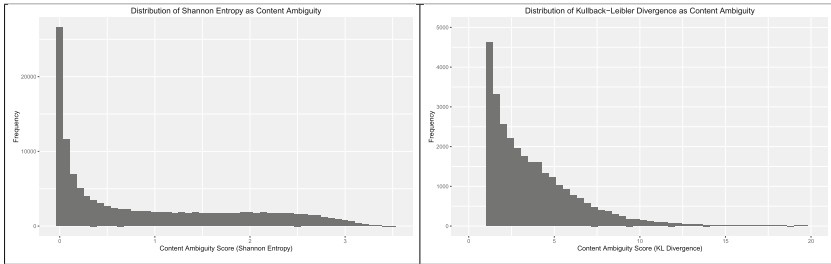
Fig. 8. Empirical distributions of two conceptual fluency measures (by intuition, they measure two dimensions of content ambiguity of logos)

Perceptual Fluency: Perceptual Complexity and Perceptual Prototypicality We calculate the image entropy of each design in our dataset as the measure for Perceptual Complexity, since perception research and algorithmic information theory posit that a compressed image file can accurately measure picture complexity. Such measures have been shown to correlate positively with subjective ratings of design complexity [3]. We measure the Perceptual Prototypicality in the same way as conceptual fluency except that the entropy is calculated based on the predicted distribution across 1000 concrete objects for each design (denoted as Q, the same Shannon Entropy of Q is detailed in Eq. 1, obtained from a fine-tuned ResNet-50 [9] based on ImageNet initializations.

We propose and measure these constructs based on previous findings that draw on an inherent and fundamental truth about the way the visual system operates: "the idea that people's preference for any design depends on the extent to which its visual processing is surprisingly fluent. This processing depends on two uncorrelated aspects of a visual design—processing expectation ex ante and processing efficiency ex post. Processing efficiency results whenever a design

is more prototypical, with fewer neural resources recruited, and is processed quickly. Such quick, efficient processing results in a gut-level positive affective response [29]. Whenever people have a processing expectation of difficulty, this gut-level affect increases perceived likability of a design. Processing expectation hinges on visual complexity of the design. When people expect difficulty in processing, they are unable to attribute the gut-positive affect evoked by efficient processing to specific design characteristics, and they therefore subconciously infer that the gut-positive reaction must imply that they like the design. On the other hand, when the processing expectation is low because a design is visually simple, people attribute the gut-level affective response arising from processing efficiency to design simplicity, and they correct for an increase of affect on their preference toward the design [30]" [20].

Table 2. Rank correlation results from test sets, averaged across five splits

Rank correlations	Transfer		Finetune	
Weight initialization	Random	ImageNet	Random	ImageNet
AlexNet	0.37	0.41	0.40	0.39
VGG16	0.40	0.38	0.44	0.40
VGG19	0.36	0.38	0.38	0.41
InceptionV3	0.50	0.51	0.54	0.55
ResNet50	0.50	0.50	0.58	0.58

Intuitively, this metric implies the extent to which the logo visual elements overlap with recognizable object categories as prototypes, and therefore, we refer to it as an Image Prototypicality measure, based on theories in perception research that we have detailed in Sect. 1. This metric might also appeal to the concreteness (abstractness) of the design pattern, which appears to be another open question yet to be resolved.

4   Results and Discussions

We document (1) the current results (rank correlations, following [17]) of the experiments of prediction tasks on the computing front in Table 2; and (2) the statistical significances of variables (our proposed measures of conceptual and perceptual fluencies) of interest regressed against annotated memorability and likability scores in Table 3 (where the significance matters rather than the prediction accuracy). The rank correlations were averaged across five splits twice—once for memorability prediction and once for likability prediction. The rank correlation peaking at 0.64 in [17] was benchmarked against human consistency at 0.68 in their setting of natural images. The disparity between ours and theirs is likely due to (1) the additional ambiguity and subjectivity of visual designs compared

to natural images or photos; (2) they added an extra support vector regression step on top of the last fully connected layers with Euclidean loss, which we could do shortly. We follow previous literatures on most choices of hyperparameters in the training processes. We chose the number of training epochs based on pilot experiments during which the value of the loss function had always plateaued after 15 epochs when predicting memorability scores and 25 epochs for likability scores. In Fig. 9, we plot the evolution of the values of loss functions over training epochs resulting from different configurations detailed in Sect. 3.1.

Table 3. Regression table on the effects of perceptual and conceptual fluency on memorability and likability

	Dependent variables:			
	mem	*mem*	*like*	*like*
Perceptual complexity (PC hereafter)	0.290***	0.278***	−0.027***	−0.012***
	(0.004)	(0.004)	(0.004)	(0.004)
Perceptual prototypicality (Inverse) (PPI)	−0.010***	−0.011***	0.003**	0.005**
	(0.002)	(0.002)	(0.001)	(0.001)
Category	−0.000***	−0.000***	−0.000***	
	(0.000)	(0.000)	(0.000)	
Country	0.000	0.000	0.000	
	(0.000)	(0.000)	(0.000)	
Conceptual entropy (More Ambiguous)	−0.018***	−0.088***	−0.037***	−0.035***
	(0.002)	(0.004)	(0.002)	(0.004)
Conceptual KL divergence	0.006***	0.013***	0.005***	0.004***
	(0.001)	(0.002)	(0.001)	(0.002)
Perceptual complexity^2	−0.074***	−0.071***	0.011***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)
Conceptual entropy^2	−0.002**	0.002**	0.009***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)
PP(I)^2	0.002***	0.002***	−0.001***	−0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Conceptual KL divergence^2	−0.000***	−0.000**	−0.000***	−0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
PC × Conceptual entropy		−0.001		0.011
		(0.003)		(0.003)
Perceptual complexity × KL divergence		0.000		−0.002
		(0.001)		(0.001)
PPI × Conceptual entropy		0.001		0.006***
		(0.001)		(0.001)
PPI × KL divergence		−0.000		−0.001***
		(0.000)		(0.000)
Constant	0.918***	1.010***	0.328***	0.299***
	(0.008)	(0.008)	(0.007)	(0.008)
Observations	123,928	123,928	123,928	123,928

Note: *p < 0.1; **p < 0.05; ***p < 0.01

We next proceed to interpret the results of the regression model from a perspective of economics and econometrics.

The consistently positively significant linear effect of perceptual complexity shown in Table 3 supports the visual processing expectation theory that when consumers experience cognitive difficulty processing certain design pattern, they mis-attribute greater perceived visual processing fluency due to adjusted higher expectation to greater liking for that design.

The consistent quadratic causal relationship that is significant between content ambiguity (conceptual entropy) and design likability in column 3 and column 4 provides evidence for our hypothesis that logos that are distinct in meaning are likely to pop out and be liked more but if the meaning is very ambiguous they will be liked less.

The other conceptual fluency measure—the KL Divergence between prediction and truth—yields qualitatively the same effect on design likability and memorability. Logos that are perceived with little bias in meaning are neither easily remembered or very much liked and neither do logos that are mostly mistakenly perceived in meaning.

The effects of perceptual fluency on design memorability and likability are flipped as posited: a moderate level of prototypicality increases memory but decreases affection, and the same with Perceptual Complexity. The more visually complex a design is, the greater visual processing expectation the perceiver will (mis-)attribute to the design, and the more positive propensities towards the design pattern. When the design pattern looks very generic or highly prototypical, it will draw less attention but increase affection since it is cognitively quick and easy to process. Conversely, low prototypicality in turn causes a visual overload and hamper visual encoding, resulting in less memorability yet greater positive affects since it encourages cognitive elaboration.

This likely suggests an interesting story about the interplay between the extent to which a logo design is ambiguous in terms of its content against its industrial content and its perceived memorability. When a logo conveys the industrial characteristics of the company in a most clear way—too *conceptually prototypical*, it is less likely to be remembered by consumers. Possible explanations include (1) visual processing expectation gets adjusted much lower, lowering the overall image processing fluency and therefore positive propensities towards the design pattern; (2) the pattern turns out generic within the category, making it difficult for perceivers to associate the design pattern with the identity of the focal company or brand. When a logo looks extremely ambiguous in terms of its industrial content, it too gets forgot about more often by consumers, possibly due to the extra cognitive processing difficulty involved that interferes processing fluency and exceeds the adjusted processing expectation gap.

The Kullback-Leibler Divergence measure of conceptual fluency with respect to the true identities of the focal companies shows a significantly positive effect on consumer memorability. It supports the reverse relative conceptual fluency effect that visual processing expectations are raised as the perceived conceptual fluency
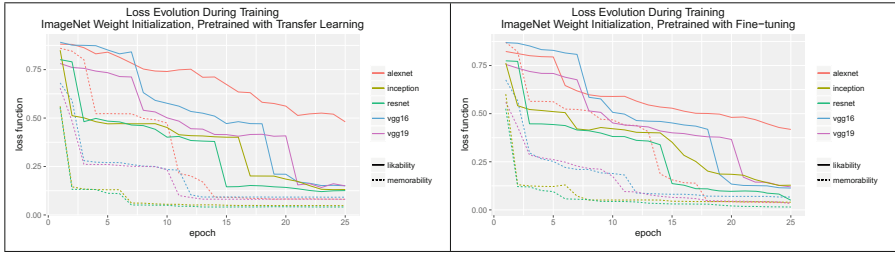
Fig. 9. Evolution of Euclidean loss during training

relative to the true company or brand identity is increased and therefore stronger preferences towards the design pattern thanks to consumers' mis-attribution of relative processing fluency given high expectations to liking.

The negatively significant interaction effect on design likability between image prototypicality and conceptual fluency measures suggests opposing moderator effects between the two pairs of constructs. When a logo design is highly ambiguous in meaning, or when the perceived industrial identity of a company represented by its logo is distinct from its true identity, the positive effect of a lower level of perceptual prototypicality on its likability within a certain range is moderated. Such moderation effects could be explained by the extra cognitive processing cost incurred by perceivers when the design pattern is very ambiguous in meaning. On the other hand, when a logo is complex in its visual style, the positive effect of content ambiguity relative to its true categorical content on likability is moderated whereas the first-order negative effect of a content-ambiguous design pattern on its likability is exacerbated, consistent with what we posit in Sect. 1 based on cognitive elaboration and processing fluency theories.

5   Robustness Analysis

For robustness checks, we considered alternative architectures, alternative datasets and alternative approaches.

– For extraction of perceptual fluency detailed in Sect. 3.2, we train VGG-16, GoogleNet, ResNet-101, ResNet-152, in addition to ResNet-50 [9], and calcualte the metrics in the same way as described. When used as regressors in our model specification, most of the resulting coefficients remain *qualitatively* unchanged except for third-order terms that become not significant. For extraction of conceptual fluency detailed in Sect. 3.2, we fine-tune ResNet-50 and GoogleNet to train the industrial category classifier as robustness checks. The resulting two measures of conceptual fluency based on Shannon entropy and Kullback-Leibler divergence yield coefficient estimates that remain qualitatively unchanged from previous model structures.
– We test the same idea on a random subsample of another design dataset we collected from Goodreads, which consists of 5,575,892 book cover images supplemented with rich meta data. Book genres were labeled as categories. Out

of a total of 1,005 genres, we randomly sampled 39 and within each genre, we randomly sampled 1,000 books, forming a subsample of 39,000 book cover design images. We collected both memorability and likability scores for a random subsample of size 11,041. The regression coefficients remain qualitatively the same when memorability is specified as the dependent variable. When the dependent variable is likability, the coefficients of style entropy squared and image entropy dispersion are not significant, whereas other first-order and second-order results remain qualitatively the same.

– We follow the approaches in [17] and train DCNNs on our datasets to predict annotated scores directly. We visualize the activated neurons at each layer for both prediction tasks of memorability and likability.

6   Conclusion

Logos are visual representations of companies and brands. The effective design of brand logos is a careful blend of art and science. While logos should be aesthetically appealing to target consumers, companies do spend a great deal of time and money crafting logo designs that reveal central messages about the brand. These efforts are not unwarranted since studies show that consumers attribute their inferences from logo designs to the associated companies and brands, or even the broader environment [6,13,23]. Therefore it remains relevant for both multi-national corporations and young entrepreneurs to better understand what makes a logo design memorable and likable. As one of the first steps towards this goal, we gather two large-scale visual design image datasets (one of business logos, the other of cover designs) with rich meta-data that will hopefully prove valuable to the business and vision research communities. We experiment and compare how different DCNN architectures perform in predicting design memorability and likability. We propose and provide estimation methods based on training DCNNs to extract and evaluate two independent constructs for designs: perceptual distinctiveness ("perceptual fluency" metrics) and ambiguity in meaning ("conceptual fluency" metrics) of each logo. We provide evidences of causal inference that both constructs significantly affect memory for a logo design, consistent with cognitive elaboration theory. The effect on liking, however, is interactive, consistent with processing fluency (e.g., [20,22]).

Potential intriguing extensions that leverage our datasets or results include:

– Developing generative adversarial networks for logo generation and slogan generation given the brand image and market position;
– Automatic storytelling or generating signature stories given the logos and business identities;
– Exploring what makes a logo-slogan combination memorable and/or likable, which appears to echo the emerging efforts that bridge language and vision.

References

1. Bornstein, R.F., D'Agostino, P.R.: The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. Soc. Cognit. **12**(2), 103–128 (1994)
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition CVPR 2009, pp. 248–255. IEEE (2009)
3. Donderi, D.C.: Visual complexity: a review. Psychol. Bull. **132**(1), 73–97 (2006)
4. Estes, W.K.: Array models for category learning. Cognit. Psychol. **18**(4), 500–549 (1986)
5. Flavell, J.H.: Metacognition and cognitive monitoring: a new area of cognitive-developmental inquiry. Am. Psychol. **34**(10), 906–911 (1979)
6. Hagtvedt, H.: The impact of incomplete typeface logos on perceptions of the firm. J. Market. **75**(4), 86–93 (2011)
7. Halberstadt, J., Rhodes, G.: The attractiveness of nonface averages: implications for an evolutionary explanation of the attractiveness of average faces. Psychol. Sci. **11**(4), 285–289 (2000)
8. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1026–1034 (2015)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
10. Hintzman, D.L.: Human learning and memory: connections and dissociations. Annu. Rev. Psychol. **41**(1), 109–139 (1990)
11. Isola, P., Xiao, J., Parikh, D., Torralba, A., Oliva, A.: What makes a photograph memorable? IEEE Trans. Pattern Anal. Mach. Intell. **36**(7), 1469–1482 (2014)
12. Isola, P., Xiao, J., Torralba, A., Oliva, A.: What makes an image memorable? In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 145–152 (2011)
13. Jiang, Y., Gorn, G.J., Galli, M., Chattopadhyay, A.: Does your company have the right logo? How and why circular-and angular-logo shapes influence brand attribute judgments. J. Consum. Res. **42**(5), 709–726 (2016)
14. Khosla, A., Bainbridge, W., Torralba, A.: Modifying the memorability of face photographs. In: Proceedings of the IEEE (2013)
15. Khosla, A., Xiao, J., Torralba, A.: Memorability of image regions. In: Advances in Neural (2012)
16. Khosla, A., Das Sarma, A., Hamid, R.: What makes an image popular? In: Proceedings of the 23rd International Conference on World Wide Web, pp. 867–876. ACM (2014)
17. Khosla, A., et al.: Understanding and predicting image memorability at a large scale. In: International Conference on Computer Vision (ICCV) (2015)
18. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
19. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)
20. Landwehr, J.R.: Gut liking for the ordinary: Incorporating design fluency improves automobile sales forecasts. Market. Sci. **30**(3), 416–429 (2011)

21. Langlois, J.H., Roggman, L.A.: Attractive faces are only average. Psychol. Sci. **1**(2), 115–121 (1990)
22. Lee, A.Y., Labroo, A.A.: The effect of conceptual and perceptual fluency on brand evaluation. J. Market. Res. **41**(2), 151–165 (2004)
23. Rahinel, R., Nelson, N.M.: When brand logos describe the environment: design instability and the utility of safety-oriented products. J. Consum. Res. **43**(3), 478–496 (2016)
24. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
25. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., et al.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–9 (2015)
26. Veenman, M.V., Van Hout-Wolters, B.H., Afflerbach, P.: Metacognition and learning: Conceptual and methodological considerations. Metacognition and learning **1**(1), 3–14 (2006)
27. Wang, M.C., Haertel, G.D., Walberg, H.J.: What influences learning? A content analysis of review literature. J. Educ. Res. **84**(1), 30–43 (1990)
28. Whittlesea, B.W.: Illusions of familiarity. J. Exp. Psychol. Learn. Mem. Cognit. **19**(6), 1235–1253 (1993)
29. Winkielman, P., Cacioppo, J.T.: Mind at ease puts a smile on the face: psychophysiological evidence that processing facilitation elicits positive affect. J. Personal. Soc. Psychol. **81**(6), 989–1000 (2001)
30. Winkielman, P., Halberstadt, J., Fazendeiro, T., Catty, S.: Prototypes are attractive because they are easy on the mind. Psychol. Sci. **17**(9), 799–806 (2006)
31. Winkielman, P., Schwarz, N., Fazendeiro, T., Reber, R., Musch, J., Klauer, K.C.: The hedonic marking of processing fluency: implications for evaluative judgment. In: The Psychology of Evaluation: Affective Processes in Cognition and Emotion, pp. 189–217 (2003)