

The Price is Right: Predicting Prices with Product Images

Steven Chen Edward Chou Richard Yang
Stanford University

Abstract

In this work, we build an array of machine learning models to predict the price of a product given its image, and visualize the features that result in higher or lower price predictions. We collect two novel datasets of product images and their MSRP prices for this purpose: a bicycle dataset and a car dataset. We set baselines for price regression using linear regression on histogram of oriented gradients (HOG) and convolutional neural network (CNN) features, and a baseline for price segment classification using a multiclass SVM. For our main models, we train several deep CNNs using both transfer learning and our own architectures, for both regression and classification. We achieve strong results on both datasets, with deep CNNs significantly outperforming other models in a variety of metrics. Finally, we use several recently-developed methods to visualize the image features that result in higher or lower prices.

1. Introduction

Online shopping is quickly becoming the norm, but the experience differs greatly from retail shopping, in which people have the opportunity to closely examine a product, weighing in the feel of a material or the scent of a cream before making a purchase decision. Online shoppers must rely entirely on the few product images to make a decision.

In this work, we build, optimize, and evaluate an array of machine learning models that can predict prices based on product images, for both regression and classification tasks. These models can be used by both buyers and sellers to suggest fair prices for products, or warn of inaccurate or unreasonable pricing. In this work, we also visualize which features tend to result in predicted higher or lower prices. Our proposed model can help sellers increase the perceived value of their products, helping guide product design and photo selection to improve a buyers impression.

Our project is split between CS229 and CS221. We build the linear regression models and SVM for CS221. For CS229, we build the array of deep CNNs for both regression and classification tasks, with our own novel architectures as

well as with transfer learning, and evaluate over multiple metrics using linear regression and SVM as baselines. For CS229, we also use multiple techniques to visualize the image features that result in higher or lower price predictions.

2. Related Work

Computer vision and supervised machine learning have been used in conjunction for a variety of pricing and regression tasks. Early work has used supervised learning to predict attractiveness given labeled faces [2]. Recent works have predicted age using face images [3, 5], and housing prices with satellite imagery [7, 13], tasks which are traditionally difficult for humans. In contrast to these works, we focus on the task of prediction using images of consumer products, with novel datasets curated specifically for this purpose.

ClickToPrice [8] proposes the most similar concept to our work. In ClickToPrice, the author explores the predictive power of product images for prices. Our project is similar in that we use machine learning to predict the prices. However, [8] uses very basic techniques to classify the general product category (e.g., towels, shoes), and uses that categorization alone to predict the average price for each item. We argue that such a model is functionally equivalent to image classification, and is not suited for price prediction. Our models are specifically designed for fine-grained price prediction for items of the same type and are significantly more sophisticated in technical implementation.

Recent research has delved into methods for visualizing what features and image parts CNNs use to make their predictions. Zeiler and Fergus [14] learn what visual features maximize hidden unit activations, and use obscuring sliding windows to determine which features influence prediction. Yosinski et al. [12] build live visualizations of activations, allowing for easier discovery of the inner workings of CNNs. Zhou et al. [15] use global average pooling to visualize what regions of images are most responsible for classification predictions. In contrast, Simonyan et al. [10] generate images that maximize the class score predicted by an object recognition network. We experiment with a subset of these methods to visualize the features that result in higher or lower prices for products.



Figure 1: Examples of images from bikes (left) and cars(right).




Figure 2: CDF of prices for bikes (left) and cars (right).

3. Approach

We first present the datasets we collect specifically for this work, then describe the algorithmic models used to predict prices with these datasets.

3.1. Datasets

In this work, we choose to use cars and bikes as target product datasets, due to the wide visual variances in bike and car models, close visual correlations to prices, and relevance of online shopping for cars and bikes.

Our first dataset, **bikes**, is curated from Bicycle Blue Book. We collect images and prices from the listings, and preprocess by filtering out low quality images and resizing to 224 by 224 pixels. Our final dataset contains solid white background, side view images. The dataset consists of 21,843 images, each with a model name and MSRP price.

Our second dataset, **cars**, is a dataset of vehicle images and their prices. We retrieve price data from Kaggle¹. We join these prices on images from Google Images, using search terms consisting of model and year, along with “Angular Front View”. We clean and resize the images, resulting in a final dataset of 1,400 examples.

The bike prices range between \$70 and \$17,000, and the car prices range between \$12,000 and \$2,000,000 (see Figure 2 for CDF). The prices closely follow an exponential distribution, in which there are significantly more models at the regular price segments than at the luxury segment.

3.2. Modeling

We approach price prediction through two different learning tasks: regression and classification. In the regression models, we attempt to directly predict the price given an image. In the price classification models, we split our data into various price segments and treat it as classification into price ranges.

¹www.kaggle.com/jshih7/car-price-prediction

3.2.1 Linear Regression Baselines

Our first baseline is multiple (multivariable) linear regression using histogram of oriented gradients (HOG) features, using PCA to reduce overfitting. Our second regression model is multiple linear regression using CNN features. For this model, we generate CNN features from the last convolutional layer of VGG-16 [11], a very deep CNN pretrained for object recognition, then use PCA-reduced-dimension features as input data. We report parameter values and evaluate performance for this baseline and the others in Section 4.

3.2.2 Multiclass SVM Baseline

Our baseline for classification is a multiclass SVM trained on price segments (see Section 4 for segments). Support vector machines can efficiently use powerful kernels such as the RBF (Gaussian) by optimizing the dual form:

$$\begin{aligned} \max_{\alpha} W(\alpha) &= \sum_{i=1}^m \alpha_i - \frac{1}{2} \sum_{i,j=1}^m y^{(i)} y^{(j)} \alpha_i \alpha_j \langle x^{(i)}, x^{(j)} \rangle \\ \text{s.t. } 0 &\leq \alpha_i \leq C, i = 1, \dots, m \text{ and } \sum_{i=1}^m \alpha_i y^{(i)} = 0 \end{aligned}$$

To support multiclass classification, we use the one vs. one approach, which trains one binary SVM between each pairwise combination of categories. Each binary classifier votes for a category, and the prediction of the model is the category that received the most votes.

3.2.3 Transfer Learning CNN

Our first deep CNN is trained using transfer learning. In particular, we use the pretrained ImageNet object recognition networks VGG-16 [11] and MobileNet [6]. VGG-16 is a very deep convolutional neural network consisting of many layers of small convolution and pooling filters, followed by two fully connected layers and a softmax output. MobileNet is an ImageNet CNN that uses separable depth-wise convolutions for higher efficiency. For both models, we use the Keras [1] framework, load the architecture and weights, and remove the networks’ softmax and dense layers. We set the remaining layers to be fixed, and add our own fully connected layer.

We use two different output layers, each designed for a specific task. For continuous regression, we add a single linear activation output unit after the fully connected layer. For segmented price classification, we add an output layer with an output unit for each class, and use a softmax activation. We optimize and tune each network and task pair separately (see Section 4 for more details).

Layer	Filter Size	Number Filters/Units
Conv2D	(5,5)	32
MaxPooling2D	(2,2)	-
Conv2D	(5, 5)	64
MaxPooling2D	(2, 2)	-
Conv2D	(5, 5)	64
MaxPooling2D	(2, 2)	-
Conv2D	(5, 5)	128
MaxPooling2D	(4, 4)	-
Dense (FC)	-	512
Dense (FC)	-	1

Table 1: Architecture for PriceNet-Reg. PriceNet-Class has an additional Conv and Dense layer and is not shown for brevity.

3.2.4 PriceNet

Lastly, we design our own novel deep learning architecture called PriceNet, and train it on the bikes dataset for price regression and classification. PriceNet is a deep CNN consisting of four blocks of convolutional and pooling filters, followed by fully connected layers. Each block has a dropout on its input to prevent overfitting, and all of our blocks are initialized randomly using Glorot uniform or normal [4].

We build two variations of PriceNet, PriceNet-Reg for price regression, and PriceNet-Class for price segment classification, which have slightly different layers and parameters. In particular, PriceNet-Class has an extra convolutional layer in the first block and two fully connected layers before the output layer, and uses softmax activation for the last layer compared to linear activation for PriceNet-Reg. Overall, we have 3.5M trainable parameters for PriceNet-Reg, and 1.2M trainable parameters for PriceNet-Class. We tune the two networks separately.

4. Experimental Results

We first describe our tuning process and parameter selections for our models. We then present our evaluations of our models in terms of several metrics. Finally, we present several different visualizations from our CNNs, along with our interpretations.

4.1. Model Tuning and Parameters

4.1.1 Linear Regression Baselines

For linear regression with HOG, we generate HOG features with 8 orientations per histogram and a window size of 32 by 32 pixels, selected as a reasonable balance between resolution and noise. We then run principal components analysis (PCA) and reducing the dimensionality to 200. For linear regression with CNN features, we generate features using the last convolutional layer of VGG16 [11], and use PCA to reduce the dimensionality to 256.

4.1.2 Multiclass SVM Baseline

We tune the SVM model with respect to two hyperparameters: C and γ , and run hyperparameter search on a log scale of the parameters. The top performance converges after using $C \geq 1$ with any γ , so we select $C = 1$ and $\gamma = 0.001$.

4.1.3 Transfer Learning (VGG16)

While training our transfer learning models, we use several techniques to tune our weights and parameters. For both regression and classification, we obtain the highest performance with the RMSprop optimizer, which divides learning rate by an exponentially decaying average of squared gradients. We also use dropout, dropping the effect of random hidden units during training, to help reduce overfitting. For parameter selection, we first tune parameters over a log scale and find the best candidates, and then fine tune over a smaller range around the candidates. We tune the learning rate, minibatch size, number of hidden units, and number of training epochs. Due to VGG-16’s stronger performance than MobileNet for the first task we evaluate, bike price regression, we decide to use VGG-16 for the rest of our transfer learning tasks.

4.1.4 PriceNet

We train PriceNet-Reg and PriceNet-Class on our bikes dataset. Due to of the small size of our car dataset (1,400 cars compared to more than 20,000 bikes), we did not have enough data to train a deep neural network from scratch for cars.

We train both PriceNet-Reg and PriceNet-Class using log-scale parameter selection to tune learning rate, minibatch size, number of hidden units, and number of training epochs. Initially, our PriceNet-Class utilized a similar architecture as PriceNet-Reg, which was training at 25 percent accuracy over 4 price segments meaning nothing was being learned. By adjusting the parameters and adding a few layers in experimentally, we were able to get a significant increase in our learning accuracy.

4.2. Evaluation

We split our datasets into training and testing splits, which are consistent across all models to ensure fair comparison. To create the split for both datasets, we first shuffle, then assign 90 percent of the points to train, and the remaining 10 percent to test.

4.2.1 Regression Models

We use three different metrics to evaluate and compare the performance of our models on price regression: root

Model	RMSE	MAE	R^2
Average Baseline	1810.19	1318.53	0.00
LinReg (HOG Features)	1274.99	833.02	0.50
LinReg (CNN Features)	1054.67	712.63	0.66
MobileNet Transfer CNN	764.21	429.43	0.82
VGG Transfer CNN	747.42	405.50	0.83
PriceNetReg CNN	702.29	380.21	0.85

Table 2: Regression results for the bikes dataset.

Model	RMSE	MAE	R^2
Average Baseline	76240.41	44410.57	0.00
LinReg (HOG Features)	41898.48	27588.70	0.70
LinReg (CNN Features)	37808.84	23929.67	0.75
VGG Transfer CNN	12363.65	7477.74	0.97

Table 3: Regression results for the cars dataset.

mean squared error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2). RMSE measures the root average squared error between the predicted and actual price $\sqrt{\sum_{i=1}^m (y^{(i)} - \hat{y}^{(i)})^2}$, while MAE is interpreted as the average absolute difference in price $\sum_{i=1}^m |y^{(i)} - \hat{y}^{(i)}|$. Lower values are better for both. Coefficient of determination measures the proportion of variance explained by the model, and lies between 0 and 1, where higher values are better.

We report results of our two linear regression models and the deep CNN in Table 2 (bikes) and 3 (cars), alongside a naive baseline that always predicts the average price. All models significantly outperform the naive baseline, with linear regression on CNN features showing a margin of improvement over HOG, likely due to the CNN features providing more discriminative visual cues. The deep CNNs in particular have very strong performance, significantly outperforming the other models in every metric. On the bikes dataset, our PriceNet architecture achieves the strongest results in each metric, with an MAE of \$380.21 on prices ranging from \$70 to \$1,700, while on the cars dataset, the transfer learning VGG network outperforms linear regression, achieving a very high R^2 of 0.97.

4.2.2 Classification Models

For classification, we assign class segments to each example using price cutoffs corresponding to percentiles of price. We assign labels of 25, 50, 75, 100 for the bikes dataset (4 classes), and 20, 40, 60, 80, 100 for the cars dataset (5 classes). While classification does not predict price directly like regression, we have two main reasons for using classification; price segmentation is useful in many business applications, and classification allows us to apply certain techniques such as class activation maps to visualize

Model	Precision	Recall	F1-Score
SVM (bikes)	0.80	0.45	0.43
VGG Transfer (bikes)	0.74	0.75	0.74
PriceNetClass (bikes)	0.81	0.80	0.80
SVM (cars)	0.83	0.82	0.82
VGG Transfer (cars)	0.82	0.82	0.82

Table 4: Classification results for the bikes and cars datasets.




Figure 3: Sliding window heatmap visualizing obscuring regions that result in higher (red) or lower (blue) prices. Obscuring the training wheels (bottom left), a visual cue of a low-price model, results in the predicted price increasing by 150 dollars.

features.

We evaluate our classification models based on three metrics: precision, recall, and the F-1 score. Precision measures the percent of positive predictions that were correct $P = \text{True Positives} / (\text{True Positives} + \text{False Positives})$. Recall measures the percent of positive predictions out of all the positive examples $P = \text{True Positives} / (\text{True Positives} + \text{False Negatives})$. The F-1 score is calculated as the harmonic mean of precision and recall $F1 = 2(P \times R) / (P + R)$.

4.3. Visualizations

In this section, we describe the three different methods we use to visualize how our deep CNN models see their input and which visual features affect the models' perception.

4.3.1 Sliding Window Heatmaps

We use obscuring sliding windows on a deep CNN with linear output (regression) to determine which features of input images are important to determining the predicted price, in the vein of [14]. We slide a 28 by 28 pixel window over an input image, obscuring that area of the image by replacing the pixel values with the average value for the network. We then run the obscured images through the network, and compare the predicted price of the obscured images to the original predicted price.

We visualize these changes using a heatmap, where each square of the heatmap corresponds to the region of the image that was obscured. An example is shown in Figure 3: obscuring the training wheels increases the predicted price by \$150.



Figure 4: A comparison of saliency maps for bike examples across four different classes. Note the highlights on the seat, handlebars, gearbox, and wheel spokes.



Figure 5: A comparison of saliency maps for car examples across four different classes. Note the highlights on the manufacturer logo, frame shape, and tires.



Figure 6: Saliency maps for two perspectives of the same bike.

4.3.2 Saliency Maps

For our classification networks, we create saliency maps to visualize how individual pixels contribute to the output. In [10], the authors showed that the class score S_C of an image I can be approximated by the first-order Taylor expansion $S_C(I) = w^T I + b$, where w is calculated by taking the gradient of the prediction with respect to an input image. w corresponds to the weight of each pixel in I , and shows much each pixel contributes to the class prediction.

In Figure 4, we show the original input images and the respective saliency map for bike examples across the four output classes. From these visualizations, we observe the most salient regions to our model are the seat shape, handlebars, gearbox, and brakes. Similarly, saliency maps for cars are shown in Figure 5. For cars, the most salient regions are the logo, body contour, and wheels.

Additionally, we use saliency maps to observe how well our CNNs generalize. We show that our model is invariant to angles by passing two photos of the same bike taken at different angles to our model, and it predicts the correct class for both images. (see Figure 6).




Figure 7: CAM visualizations overlaid on the original input image. Notable areas for bikes include seats, training wheels, and tire spokes. For cars, notable areas include convertible top, doors, and body contours.

4.3.3 Gradient-Weighted Class Activation Maps

Gradient-weighted Class Activation Maps (CAM) have been previously used for object locality detection [9]. When multiple output classes are present in an image, CAM will highlight parts of the image that contribute most to the selected output class. Since our classification segments by price, we use CAM to highlight areas of the bike or car that result in cheap or expensive price ranges.

The intuition of CAM is similar to that of saliency maps. However, instead of computing the gradient with respect to the output, we compute the gradient with respect to the feature maps generated by a specific convolutional layer. Then, we apply global average pooling on the gradients to create a weight vector w representative of the contributions of each unit on the class output. In Figure 7, we show CAM heatmaps overlaid on the input images. Our model focuses on visually diverse areas of bikes, such as the handlebars, seat, and tires, as well as important regions of cars, such as the convertible top or doors.

5. Conclusion and Future Work

In this work, we introduce two novel datasets and build multiple models for predicting prices of products using only image data. For regression, our custom network PriceNet outperforms transfer learning, as well as linear regression baselines. For classification, our transfer learning network outperforms a multiclass SVM. Additionally, we visualize what our CNNs see using three different methods, giving us insight on what visual features result in certain prices.

We have identified multiple future real-world applications. Using feature visualization, sellers can determine what features of objects are correlated to higher prices, and use this to help guide design. Our model can also be extended to assist valuations, such as for used car sales, where prices are difficult to determine. Finally, our model can be applied to auction sites such as eBay for recommendations of starting bids, and provide a tool to individuals for choosing better photo quality when selling products online.

6. Contributions

Steven Chen cleaned the bikes dataset, and built the CNN linear regression, PriceNet, and VGG transfer learning models. Edward Chou collected and cleaned the cars dataset, built the linear regression model on histogram of gradients features, the MobileNet transfer learning model, sliding windows visualization, and worked on the multi-class networks. Richard Yang collected the bikes dataset, generated CNN features, worked on VGG transfer learning, trained the SVMs, and generated the saliency map and class activation map visualizations. All three group members contributed to the writing and editing of this report.

References

- [1] F. Chollet et al. Keras. <https://github.com/fchollet/keras>, 2015.
- [2] Y. Eisenthal. *Facial Attractiveness: Beauty and the Machine*. PhD thesis, Tel-Aviv University, 2006.
- [3] Y. Fu and T. Huang. Human age estimation with regression on discriminative aging manifold. *IEEE Transactions on Multimedia*, 10(4), 2008.
- [4] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010.
- [5] H. Han, C. Otto, and A. Jain. Age estimation from face images: Human vs. machine performance. In *IAPR International Conference on Biometrics*, 2013.
- [6] A. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. In *arXiv:1704.04861*, 2017.
- [7] V. Limsombunchai. House price prediction: Hedonic price model vs. artificial neural network. In *NZARES*, 2004.
- [8] A. Maurya. Clicktoprice: Incorporating visual features of product images in price prediction. In *INFORMS*, 2016.
- [9] R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017.
- [10] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*, 2014.
- [11] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [12] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. In *IMCL*, 2015.
- [13] Q. You, R. Pang, L. Cao, and J. Luo. Image-based appraisal of real estate properties. *IEEE Transactions on Multimedia*, 19(12), 2017.
- [14] M. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014.
- [15] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016.