Quantifying Curb Appeal

Zachary Bessinger and Nathan Jacobs

http://cs.uky.edu/~zach/curbappeal

2016 IEEE International Conference on Image Processing

September 25-28, 2016 • Phoenix, Arizona, USA

Overview

Task: Estimate the price of a residential home.

Limitations: Existing methods rely on objective attributes and fail to account for differences in curb appeal.

Our Approach: Combine street-level photographs and objective attributes. We find that using images results in more accurate price estimates.

Dataset

We constructed a dataset of homes with photos, metadata, and prices [1].

Dataset details:

- 83,140 homes
- 15 objective attributes
- Each home has a front-facing image captured by an appraiser

[1] http://fayettepva.com/

[2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in NIPS, 2012.

[3] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, “Going deeper with convolutions,” in *CVPR*, 2015.

[4] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *preprint arXiv:1409.1556*, 2014.

[5] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *IJCV*, vol. 115, no. 3, 2015.

[6] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva, “Learning deep features for scene recognition using places database,” in *NIPS*, 2014.

[7] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Antonio Torralba, and Aude Oliva, “Places2: A large-scale database for scene understanding,” *Arxiv preprint: [pending]*, 2016.

Method

Learn models using linear, ridge, and random forest regression.

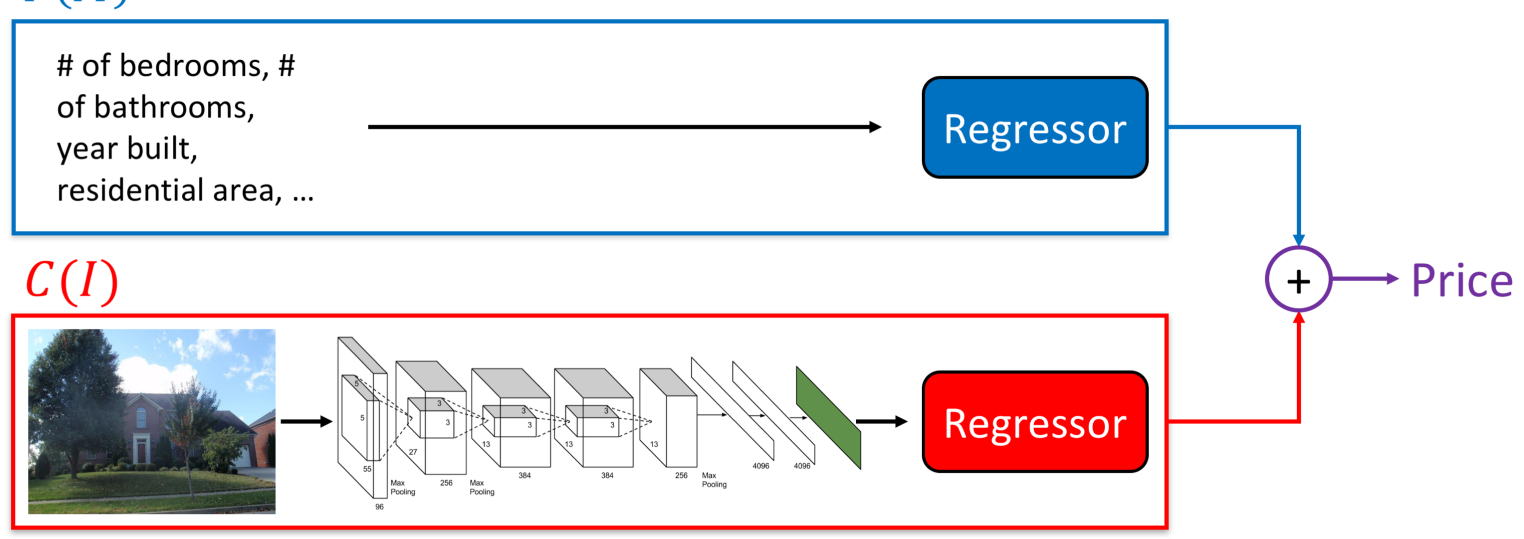

We compare our joint model, $P(M) + C(I)$, against two baseline models:

- $P(M)$, Predict price using only metadata
- $P(I)$, Predict price using only image features
- $P(M) + C(I)$, Predict price using metadata and a **curb appeal modifier.** $C(I)$ adjusts the price of a home, positively or negatively, based on its curb appeal.

Evaluation

Feature Selection

We apply linear regression to select the subset of attributes the highest R^2. These attributes are: the property area, # of bedrooms, # of bathrooms, year built, residential area, total fixtures, basement size, and garage size.

Image Features

We evaluated AlexNet[2], GoogLeNet[3], and VGG-16/VGG-19[4] trained on ImageNet[5], Places[6], and Places2[7] and found **VGG-16 Places2 FC8** features reduce the most error.

Evaluation (cont’d)

Does the joint model improve price estimates?

	Linear	Ridge	Random Forest
$P(M)$	\$37 058	\$37 127	\$29 365
$P(I)$	\$53 575	\$53 565	\$53 727
$P(M)+C(I)$	**\$34 538**	**\$34 606**	**\$28 281**

Evaluating on a 80/20 train/test split, our joint model reduces the RMSE by 4% over $P(M)$ and 90% over $P(I)$

How does semantic label affect home price?

Positive

Negative

Semantic Label	Mean C(I)
Yard	\$ 11,944.00
Lawn	\$ 5,485.00
Campus	\$ 3,565.00
Oast House	\$ 2,807.00
Residential Neighborhood	\$ 2,251.00
Cottage	\$ 1,158.00
Hunting Lodge (outdoor)	\$ 859.00
House	\$ 770.00
Inn (outdoor)	\$ 358.00
Library (outdoor)	\$ 77.00
Garage (outdoor)	\$ (27.00)
Church (outdoor)	\$ (2,201.00)
Manufactured Home	\$ (3,144.00)
Industrial Park	\$ (6,681.00)

Conclusions

Exterior appearance correlates with the price of a home and can be used to improve existing models. Applications of our work include automated exterior appraisal, architectural anomaly detection, and demographic prediction. Our work opens many avenues of future research towards not only estimating home prices, but also holistic understanding of urban regions.