

2018

WalkNet: A Deep Learning Approach to Improving Sidewalk Quality and Accessibility

Andrew Abbott

Southern Methodist University, anabbott@mail.smu.edu

Alex Deshowitz

Southern Methodist University, adeshowitz@mail.smu.edu


Dennis Murray

Southern Methodist University, dpmurray@smu.edu

Eric C. Larson

Southern Methodist University, eclarson@lyle.smu.edu

Follow this and additional works at: <https://scholar.smu.edu/datasciencereview>

 Part of the [Artificial Intelligence and Robotics Commons](#), and the [Numerical Analysis and Scientific Computing Commons](#)

Recommended Citation

Abbott, Andrew; Deshowitz, Alex; Murray, Dennis; and Larson, Eric C. (2018) "WalkNet: A Deep Learning Approach to Improving Sidewalk Quality and Accessibility," *SMU Data Science Review*: Vol. 1 : No. 1 , Article 7.

Available at: <https://scholar.smu.edu/datasciencereview/vol1/iss1/7>

This Article is brought to you for free and open access by SMU Scholar. It has been accepted for inclusion in SMU Data Science Review by an authorized administrator of SMU Scholar. For more information, please visit <http://digitalrepository.smu.edu>.

WalkNet: A Deep Learning Approach to Improving Sidewalk Quality and Accessibility

Andrew Abbott¹, Alex Deshowitz¹, Dennis Murray¹, Dr. Eric Larson¹

¹ Southern Methodist University, Dallas, Texas

Abstract. This paper proposes a framework for optimizing allocation of infrastructure spending on sidewalk improvement and allowing planners to focus their budgets on the areas in the most need. In this research, we identify curb ramps from Google Street View images using traditional machine learning and deep learning methods. Our convolutional neural network approach achieved an 83% accuracy and high level of precision when classifying curb cuts. We found that as the model received more data, the accuracy increased, which with the continued collection of crowdsourced labeling of curb cuts will increase the model's classification power. We further investigated a model using the TensorFlow Object Detection API, to construct a model that could accurately isolate curbs and identify whether a curb contained a proper cut or not. Our Results suggest that a municipality could pinpoint areas for infrastructure spend automatically. We leave implementation and data layering as an area of exploration.

1 Introduction

Municipalities across the United States continue to struggle to properly allocate infrastructure spending. Sidewalks represent a unique case because they are used by nearly all members of society yet the collection of data on inconsistencies and risks lies in the hands of users. Word of mouth remains the primary source of information although it is often biased. Many cities, such as the city of Dallas, hire contractors who rely on Excel spreadsheets to prioritize infrastructure spending. With open source algorithms and increased compute speed, there are new possibilities for solving these problems for cities across the U.S. in a cost-efficient manner that provides the optimal benefit to the users. There have been numerous sources of funding for sidewalks approved at both federal and local levels over the past 10 years. The American Recovery and Reinvestment Act of 2009 [1] allocated over \$1 billion to municipal infrastructure spending. North Texas received more than \$43 million in funding for quality of life enhancement projects [2]. However, this program has been dubbed by many as the sidewalk to nowhere due to the construction of sidewalks that literally lead into ditches and canals [3].

To rectify the lack of city-wide mapping of curb cuts, this paper demonstrates a deep learning based solution that leverages convolutional neural networks. Using a data set of images and coordinates of curb cuts within those images, several machine learning

and deep learning methods were evaluated for accuracy in detecting curb cuts in a sidewalk.

This initial data came from the University of Maryland's Project Sidewalk. The data consists of more than 40 gigabytes of Google Street View panoramic image data of Washington D.C. and coordinates of bounding box labels for curb cuts and accessibility challenges. The data was partitioned using a simple-split cross-validation methodology for model training and testing. Several machine learning algorithms were tuned and tested to gain a baseline for model accuracy. The best model achieved over 83% accuracy on the final test set of data

We further investigated means for fully automating the classification process using a TensorFlow object detection model. The TensorFlow Object Detection API was used to detect areas within an image where a curb cut should reside. Once the API detects curbs in a panoramic Street View image, the model is then able to accurately classify the curb as having a safe cut or not using our pre-trained model.

This type of model can be extended to any municipality that is mapped in Google's Street View infrastructure. Additionally, we suggest an approach to leveraging this model to make sound recommendations regarding areas of focus for the city planning commission.

2 A Primer on Sidewalk Accessibility Monitoring

The optimal allocation of infrastructure spending on sidewalks is a largely overlooked issue that has major effects on a share of the population. According to the U.S. Center for Disease Control, 13% of the U.S. adult population has some form of mobility impairment [4]. With the aging of the U.S. population, mobility issues continue to rise in importance. For individuals with mobility limitations, the availability of quality sidewalk ramps represents a necessity in navigating a city or urban environment. When navigating urban areas, wheelchair users often need to rely on their own experience and local knowledge to successfully and safely navigate in an efficient and safe manner. Without this knowledge of navigability, a wheelchair user may fail to reach their desired location, sustain an injury, or choose not to travel at all.

In addition to the mobility challenged, quality sidewalks and curb cuts aid in the safety and health of pedestrians and cyclists. In 2015 alone, 5,376 pedestrians and 818 bicyclists were killed in vehicle-related incidents [5]. With drivers becoming increasingly distracted, having well-designated curb cuts and quality sidewalks for non-vehicular modes of transportation will help reduce the impact of these incidents. Additionally, according to the U.S. Center for Disease Control, sidewalks can help limit obesity trends in urban areas. The CDC has declared the obesity problem in America to be an epidemic and has directly linked these risks to sedentary lifestyles [6]. With activity levels becoming increasingly sedentary, it is more important than ever to provide safe and quality ways for people to reach their destinations. A small investment in correctly-placed sidewalk cuts and infrastructure could help offset some of the rising medical costs from sedentary lifestyles.

Previous attempts at identifying curb cuts have been carried out through in-person (subjective) "Neighborhood Audits" or through information cataloged in Geographic

Information Systems [7]. The completeness, timeliness, and quality of the information gathered in these methods is often lower than what is needed for a comprehensive view of the navigability of curb ramps or sidewalks across a city. These manual approaches also tend to come with a high acquisition costs and still suffer from information bias.

The city of Dallas presents a prime example of the inefficiencies encountered in sidewalk infrastructure improvement. Currently, the city of Dallas responds to complaints about sidewalks through its street services program. As incidents are received, the city will place the inquiry on a list for assessment. There is no priority granted for severity of the situation, thus, all incidents are treated with the same level of priority. Incidents are also tracked at an individual level, without consideration of neighborhoods or zones that might represent systemic issues. Additionally, maintenance in suburban areas is the responsibility of the home or property owner and this is not necessarily considered in the ranking process. The current process takes 2-3 months in order to receive an assessment and cost estimate for each incident or property owner. Once the assessment has been completed, either the city or the property owner will plan and fund the project. One program in Dallas allows the city to reimburse homeowners up to \$500 or 50% of the repair cost, whichever is less. Under these circumstances, the city spends significant time and resources completing the assessments [8], without ever addressing an accessibility issue.

3 Related Work

The current reach of accessibility features in the urban landscape was researched by Bennett, Kirby, and MacDonald [9]. They surveyed 79 intersections in Halifax, Nova Scotia. Their scoring methodology asked 8 different questions that addressed both the presence and quality of curb ramps at these intersections. Each question required a binary response. Several of the questions would appear to be answerable from the research we propose – the presence of curb ramps, accessibility from the line of travel (that is, the chair user can access the ramp without exiting the crosswalk), that the ramp is “free from irregularities”, and free from drainage grates. Four additional questions address the question of slopes and dimensions of the curb ramp. Their findings in the limited scope of the survey was that 98.7% of intersections had curb ramps, but just more than half, 53.8, had a direct line of travel from the crosswalk. All of the ramps were free from drainage grates, and 85.9% were free from irregularity. The average intersection scored 5.6. The researchers proposed that wheelchair users must adapt to the lack of infrastructure by increasing their skill and dexterity in maneuvering the chair [9]. While improved skill among wheelchair users is desirable, it may also lie beyond the physical capabilities of the individual.

Bromley *et al.* [7] noted that legislation in the United Kingdom seeks to provide access to goods and services to all persons, but not necessarily the facilities containing goods and services. It is a fine distinction between the two, and within this context it could be judged that this is the granular difference that describes how accessibility isn't a special accommodation but provides equal access to all. Respondents in Bromley *et al.*'s survey-based study in Swansea, Wales found 60% thought that lack of curb ramps

were a “major” or “prohibitive” obstacle to access. Respondents used domain knowledge of the city to navigate around obstacles, and sometimes take much longer paths to access. Among the respondents, 60.8% agreed that “the way places are designed” is the major problem for wheelchair users. This attitude was more evident among younger users of wheelchairs than their older cohort. Wheelchair users recommended “more dropped kerbs” more often than any other improvement to the center city shopping experience.

Clarke *et al.*'s [10] audit of Street View images compared with an in-person audit of accessibility features. The study involved researchers in neighborhoods in Chicago walking each block from the inside to the outside, essentially walking the block twice, and assessing the quality of the sidewalks. The study then mirrored the street survey methodology using images sourced from Google Street View. The surveyors were asked to score attributes of the neighborhood on subjective scales. The results of the study concluded that questions that require determining the objective presence or absence of a feature are consistent across participants. Questions that required a subjective rating of a feature were less likely to be consistent. The level of detail required for answering specific features of a neighborhood also was a factor for the consistency between in person and Street View based studies. A question asking for the presence of litter or broken glass in the street had a consistency score of 0.347, while a question asking for the presence of check cashing scores had a consistency of 0.987. The temporal nature of some of the questions versus the age of the Street View images may also have been a factor in the outcomes.

Image recognition uses machines to recognize images. As early as 1963, the electrical engineering department at MIT began using computers to recognize 3D images. Lawrence Roberts's 1963 research [11] provided a method for a computer to process a photograph to a line drawing, reduce it to a set of coordinates, and then render the object from any direction. While these initial applications were somewhat simple compared to those that we currently use today, they paved the way for what has now become a commonplace practice across industries.

Bahlman, Zhu, and Pelkofer's work [12] provided meaningful advancements in image element detection and classification. The authors built upon their previous research involving shape and color recognition to help classify street signs and traffic signals. Their work utilizes a 2-step approach where if the model fails on the first classification step, the image is thrown out. The principle of the AdaBoost (adaptive boosting) algorithm is that it allows for the algorithm to select the most effective features for the classification task, and ignore the remaining features. The authors' model in this work found the color channel information, both in raw and normalized forms, to be highly important to the classification of road signage. This could be contrasted with methods like decision trees or logistic regression, which may require more tuning from the constructor to discover the most important and influential features.

Logistic regression and artificial neural networks have become benchmarks in classification tasks across problem types. Dreiseitl and Ohno-Machado [13] researched the methodology of machine learning methods across more than 70 papers. The authors state that the two methods, logistic regression and artificial neural networks, both have similar basis: statistical pattern recognition in large data sets. The authors reviewed 72 papers that compared outcomes of implementations of both logistic regression, and

artificial neural networks. Artificial neural networks outperformed logistic regression in 51% of the studies, but 42% of the studies provided no difference in outcome between the methods. They label logistic regression and several other methods as a “White box” method, where the parameters are clearly stated and the method that the model uses to assign importance and come to a conclusion are clearer. In contrast, artificial neural network and support vector machines are labeled as “black box” methods that do not provide interpretable markers of importance or provide methods to be verifiable.

Dean, Corrado and a group of Google researchers [14] created the predecessor for the modern open source library, TensorFlow, in a model labeled as “DistBelief.” DistBelief allows for parallel processing of training data both within a machine and across a network of machines. Likewise, the DistBelief process also allows “data parallelism” allowing for “multiple replicas of a model to optimize a single objective.” The result, the researchers conclude, is a method for training moderate sized models more quickly than before, and giving capability to training very large data set models.

DistBelief was the basis for the 2015 release of the open source TensorFlow machine learning system, documented in Abadi, Barham, et al. [15]. The purpose of TensorFlow is to provide a framework for large scale machine learning systems to be trained at higher speed through the use of many resources. These resources can stretch across multiple machines, or across the resources within a machine: CPU, GPU (Graphics Processing), or Application Specific Integrated Circuits (ASIC). The modeler can control TensorFlow from within the same environment that they use to define the model, whereas a parameter server would not have this capability. TensorFlow aligns all operations on a single dataflow graph, where the vertices are a computation and the edges are flows of data between computations. The paper reviewed the use of TensorFlow for image classification, and found a slight performance advantage for the ImageNet data set over another open source framework for deep learning, MXNet.

An aspect of previous work that is of high importance for this paper is the use and application of convolutional neural networks. As this is our model of choice, it was important for us to research the application of convolutional neural networks and their potential pitfalls. Goodfellow and a team from Google [16] showed an application of neural networks for image recognition. In this work, Goodfellow applied the DistBelief method for neural networks combined with Google Street View images to recognize multi-digit numbers, namely street addresses. In the model, the researchers first addressed training the model to identify house numbers. This was a very important step as many variables come into play with these image captures. For instance, lighting, obstructions, and changing conditions can provide potential issues when identifying numbers from images. Additionally, varying font sizes, colors, and styles can impact the ability of the algorithms to correctly identify an image. An important aspect of this type of recognition is that if a single digit is misidentified, the entire interpretation is irrelevant and meaningless. Once the model was trained on house numbers, a more complete Street View dataset was used. Using this approach, the researcher’s models were able to achieve 97.84% accuracy which was just short of the human benchmark of 98% that was the target of the project. This piece of research and the approach acted as an important catalyst for our approach to identifying curb ramps and grading accessibility.

Convolutional neural networks have also been used to improve the solutions submitted in the ImageNet Large-Scale Visual Recognition Challenge. In the work of

Simonyan, Karen, and Zisserman [17], the team used convolutional neural networks combined with several other approaches to achieve one of the highest accuracy levels seen in the competition. Their application of multiple models to solve the problem provides a solid reference point for the problem that we address in this paper.

These classification methods are able to very accurately classify or categorize an image into one of the trained classes but often it is first necessary to discover where in an image an object resides and then to classify the object or objects. This area of research is known as object detection. Uses of neural networks and deep learning for object detection fit into three major methods. The first, Faster R-CNNs [18] refers to a faster implementation of R-CNNs where the “R” refers to region proposal networks in a convolutional neural network. This method, while accurate, can be challenging to train and slower than other methods. A second major area of object detection is known as You Only Look Once (YOLO) [19]. The YOLO algorithm is much faster than R-CNN but also much less accurate. YOLO reframes object-detection as a regression problem to distinct bounding boxes and class probabilities from full images in one pass. The third type of object detection method is known as Single Shot Detection (SSD) [20]. SSDs which was originally a Google development, addresses the shortcomings of the other two methods. SSDs are much faster than Faster R-CNNs and more accurate than YOLO.

Overall, our research helped us layout the precedent for image classification and object detection, understand the application of the specific type of model that we are attempting to build, evaluate the effectiveness of proper infrastructure, and provide statistical affirmation of the health and societal benefits of proper infrastructure. With this knowledge established, we can make the case for our model to be used for the stated application problem.

4 Algorithm Design and Solution

The original images for this model were sourced from the Project Sidewalk team at the University of Maryland. Each image included a data set of bounding box coordinates enclosed in a separate text file. These bounding box coordinates were represented by x and y coordinates of the location of curb cuts existing within each image. The coordinates for each bounding box were applied to the images using Python’s image libraries, and the image of the curb was then extracted from the panoramic Street View image. Each image was then resized to a standardized 100 pixel by 100 pixel RGB color image. Each image was then flattened into a numpy array of 30,000 features (100x100x3).

Cross-validation was performed using randomized shuffle splits (with replacement). That is, images were allocated to training and test sets using an 80% training, 20% test split with a random selection mechanism.

As a baseline, several machine learning algorithms were tuned and trained using the Scikit-Learn API [21]. We trained and tuned a logistic regression classifier and linear support vector machine (SVM) classifier. Both models performed similarly and could correctly identify the lack of a curb cut in 68% accuracy (correctly assessing presence or absence of a curb cut) in the instances presented. We also attempted a number of

other models, including using the Naïve Bayes classifier and decision tree, however, none of these models provided a level of accuracy that was better than chance. Given that these simple models were performing so poorly, we next moved into ensemble learning.

Ensemble learning is an approach that allows for many models to be trained and a majority vote determines the classification. There are additional modifications that can be taken to this approach, including averaging and weighted averaging of the “votes” based on the training of the underlying models. For this approach, we first attempted a Random Forest classifier. After several rounds of hyper-parameter optimization, this model was able to perform at the same accuracy levels as the logistic regression, but we noticed that the Recall on the model was only 44%. This was cause for alarm because we generally look for models that are all within the same accuracy, precision, and recall ranges. After further inspection, it appeared that the model was missing many of the positive case classifications in the data. Therefore, it was determined that this normally robust model did not meet the criteria for our classification algorithm. The final machine learning model that we tried was Adaboost, using the default decision tree classifier as the base model. This model also had a Recall of about 46% and was less than satisfactory.

Given the poor performance in machine learning algorithms, it was decided that deep learning should be the next approach. Deep learning methods have revolutionized modern image processing, often performing significantly better than other methods on popular image recognition baselines. The classification approach selected implemented a convolutional neural network. Our network employs a simple 8-layer network with rectified linear units (ReLU) for the activation function for all layers. We also employ the popular and efficient Adam optimization algorithm, which is an extension of gradient descent that uses adaptive methods for back propagation gradient descents. Images were subjected to the same pre-processing algorithm to crop and resize as described previously. While there are many means for strong cross-validation procedures in existence, the chosen approach in this instance was a simple 80/20 split. This provided the model with plenty of data from which to learn and a modest amount of data for testing. For implementation we used the Keras sequential model API with Google’s TensorFlow backend. Our architecture used 32 filters for each convolution layer of an image, with a filter size of 3 by 3. This approach, while somewhat standard, proved to be fairly robust and prevented the model from over-learning. Max pooling was employed with a pool size of 2 by 2 to help increase resilience of the model to small perturbations in the images. We also performed a grid search of the number of filters, with the value of 32 found to be optimal. Additional filters did not tend to increase accuracy but greatly increased computation time. Therefore it was decided that the 32-filter approach was indeed a sound model parameter.


The resulting model identifies curb cuts correctly 83% of the time. This model clearly provided the highest level of accuracy and similar precision and recall. Therefore, the Convolutional Neural Network did indeed provide a level of precision that could not be attained in more basic machine learning approaches. It was also noticed that the time needed to train these models was considerably less than the amount of time needed for the tuning of the machine learning based approaches.

Table 1: Accuracy Measures by Method

Model type	Accuracy	Recall	Precision
Logistic Regression	68%	51%	53%
Support Vector Machine	68%	8%	65%
Random Forest classifier	74%	45%	67%
Decision Tree with Adaboost	69%	46%	55%
Convolutional Neural Network	83%	84%	83%

Having established that convolutional neural networks are able to accurately classify sidewalk curb ramps and curbs without ramps, we then turned our focus to the deployment of those CNNs. In order to obtain a fully automated system, the model needed to detect and classify curb ramps and missing curb ramps from the panoramic Google Street View images (not simple 100x100 images constructed from bounding box coordinates). For this task, we used Google's TensorFlow Object Detection API to segment areas of interest in the image. Specifically, we want the object detection API to identify areas in the images where curb cuts should be. Once segmented, we can use our existing CNN classifier to detect if a proper curb cut exists.

The API used the TFRecord file format, so we created a script to generate a TFRecord file for a training set and a test set from the UMD Project Sidewalk data set. For training, a pipeline is needed. The TensorFlow repository provides sample pipeline configuration files, and for our training we used the 'ssd_mobilenet_v1_pets.config' file with adjustments to fit our data set. The Mobilenet model is designed to be used when resources are limited, such as on mobile devices. Given our limited compute resources and the balance between speed and accuracy that the Single Shot Detector (SSD) method provides, we chose the SSD Mobilenet configuration. In actual deployment, given adequate resources, a different or custom network such as the classification network used above, and a larger training data set, may provide better performance. In order to speed up training, we also used a pre-trained model checkpoint as is recommended. The training job ran for just under 15 hours with 6440 steps. Figure 1 shows the Total Loss as training evolved. While we would have liked for the loss to be below 2.0, we have demonstrated the applicability of this method to our problem.

Figure 1. Total Loss versus Training Steps

After the training completed, we exported the trained model to be used for inference. This trained model can be used for many different applications, including streaming video and mobile uses. The example output in figures 2 and 3 shows how the model can now ingest panoramic images, identify curbs, and classify the curbs as good cuts versus poor cuts. Further evaluation, tuning and training on the full dataset are needed before production use.

Figure 2. Example model output showing correct detection of missing curb ramps.



Figure 3. Example of model output showing correct detection of curb ramps.



5 Results and Conclusions

In our research, we proved that deep learning can be used to correctly identify sidewalk features with curb cuts being a first example. We used data that had been pre-labeled; therefore, the step in the process of object identification was pre-solved in some sense thanks to the University of Maryland's Project Sidewalk data. However, this

model could not readily identify the portions of a panoramic image that represented sidewalk features. Therefore, the model was highly dependent on humans pre-defining bounding boxes to identify the portions of the image needed for classification scoring. To bridge this gap, we employed the Google Object Detection API. Once the model was trained to pick sidewalk features from panoramic Street View images, these cropped images were then able to be classified as having proper cuts or not having proper cuts.

While only a proof of concept, this exercise has proven that we can use these models to create applications for cities to focus infrastructure spend in an optimal manner, create applications for the handicapped, and an entire list of other options. Although the final SSD Mobilenet model took over 15 hours to train, employing faster hardware via cloud-based processing could drastically reduce this time. Additionally, once enough data is generated training only needs to happen in periodic intervals. However, as hardware speeds increase, this may be a moot point in several years.

6 Ethics

The work addressed in this paper is based primarily on Google Street View panoramic images. As of 2008, Google began to take steps to protect the privacy of people and their property captured on Street View by blurring faces and license plates found in the image [22]. While the expectation of privacy on a public street is beyond commonly accepted practice, it is a reasonable step that Google has taken to provide added privacy.

The labeled image data sourced from the University of Maryland Project Sidewalk team was created via crowd-sourced labor online. The project solicits unpaid volunteer labor to classify parts of the Street View image via a web interface without compensation.

Lower income neighborhoods may be under-represented in the Street View image data, or the frequency of image refresh may differ between different streets. If the virtual view provided by Street View is a conduit to improved infrastructure, steps must be taken to ensure there is equitable representation across neighborhoods in the image set.

No human subject testing was conducted as part of the research in this paper.

7 Future Areas of Research

As the model demonstrated in this paper currently stands, it could be deployed into an environment such as the city of Dallas to classify all streets in the city. The biggest stumbling block for productionizing of the model is the compute power that would be required to map an entire city. The output from training a model on an entire city could be incredibly powerful. City planners could deploy dashboards with geocoded mapping packages to enable visual scanning of cities for “hot spots” that need funding attention. Additionally, scoring areas that need attention via a search

algorithm would allow cities to deploy assets in an efficient manner rather than today's haphazard approach.

Now that the model demonstrated in this paper exists for curb cuts, the next logical extension is for the model to add features such as misplaced fire hydrants, benches, and other obstacles into the classification process. While proper curb cuts and cross-walks may be some of the most important features for proper sidewalk infrastructure, these other features can cause issues if not placed appropriately. Helping cities identify the entire universe of sidewalk hindrances remains one area that this model has yet to forge.

References

1. Congress, U. S. "U. S. Department of the Treasury." 17 February 2009. treasury.gov. 4 August 2017. <<https://www.treasury.gov/initiatives/recovery/Pages/recovery-act.aspx>>.
2. The North Central Texas Council of Governments. "Where the Money is Going: A report on the American Recovery and Reinvestment Act." 2010. The North Central Texas Council of Governments. 1 December 2017. <<http://www.nctcog.org/trans/outreach/arra.pdf>>.
3. Smith, Kyle. "The Obama Administration's Sidewalk to Nowhere." 23 March 2011. Forbes. 9 December 2017. <<https://www.forbes.com/sites/kylesmith/2011/03/23/the-obama-administrations-sidewalk-to-nowhere/#471d67132040>>.
4. The Morbidity and Mortality Weekly Report (MMWR). "Key Findings: Prevalence of Disability and Disability Type among Adults, United States – 2013." The Morbidity and Mortality Weekly Report (MMWR) 31 July 2015: 777-783. 1 December 2017. <<https://www.cdc.gov/ncbddd/disabilityandhealth/features/key-findings-community-prevalence.html>>.
5. National Center for Statistics and Analysis. "Traffic Safety Facts." n.d. Crashstats. National Highway Traffic Safety Administration. 1 December 2017. <<https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812375>>.
6. Barnes, Ann Smith and Stephanie A. Coulter. "Obesity and Sedentary Lifestyles: Risk for Cardiovascular Disease in Women." Texas Heart Institute Journal 39.2 (2012): 224-227.
7. Bromley, Rosemary DF, David L. Matthews and Colin J. Thomas. "City Centre Accessibility for Wheelchair Users: The Consumer Perspective and the Planning Implications." Cities 24.3 (2007): 229-241.
8. City of Dallas Department of Public Works. Sidewalk Replacement Program. n.d. City of Dallas. 1 December 2017. <<http://dallascityhall.com/departments/public-works/Pages/SidewalkReplacementProgram.aspx>>.
9. Bennett, Sean, Ronald Kirby and Blair Macdonald. "Wheelchair Accessibility: Descriptive Survey of Curb Ramps in an Urban Area." Disability and Rehabilitation. Assistive Technology 4 (2009): 17-23.
10. Clarke, Philippa and et al. "Using Google Earth to Conduct a Neighborhood Audit: Reliability of a Virtual Audit Instrument." Health & Place 16.6 (2010): 1224-1229.
11. Roberts, Lawrence G. "Machine Perception of Three-Dimensional Solids. Diss." (1963).
12. Bahlmann, Claus, et al. "A System for Traffic Sign Detection, Tracking and Recognition Using Color, Shape, and Motion." IEEE Xplore (2005).
13. Dreiseitl, Stephan and Lucila Ohno-Machado. "Logistic Regression and Artificial Neural Network Classification Models: A Methodology Review." Journal of Biomedical Informatics 35.5 (2002): 352-359.
14. Dean, Jeffrey, et al. "Large Scale Distributed Deep Networks." Advances in Neural Information Processing Systems. 2012. 1223-1231.

15. Abadi, Martin, et al. "A System for Large Scale Machine Learning." OSDI 16 (2016).
16. Goodfellow, Ian J. "Multi-digit Number Recognition from Street View Imagery using Deep Convolutional Neural Networks." arXiv Preprint arXiv:1312.6082 (2013).
17. Simonyan, Karen and Andrew Zisserman. "Very Deep Convolutional Networks for Large-scale Image Recognition." arXiv Preprint arXiv:1409.1556 (2014).
18. Ren, Shaoqing, et al. "Fast R-CNN: Towards Real-Time Object Detection with Region Proposal Networks." arXiv eprint arXiv:1506.01497 (2015).
19. Redmon, Joseph, et al. "You Only Look Once: Unified, Real-Time Object Detection." arXiv eprint arXiv:1506:02640 (2015).
20. Liu, Wie, et al. "SSD: Single Shot MultiBox Detector." arXiv eprint arXiv:1512.02325 (2015).
21. Pedregosa, F and et al. "Scikit-learn: Machine Learning in Python." JMLR 12 (2011): 2825-2830.
22. Shankland, Stephanie. "Google Begins Blurring faces in Street View." CNet 13 May 2013: 25.