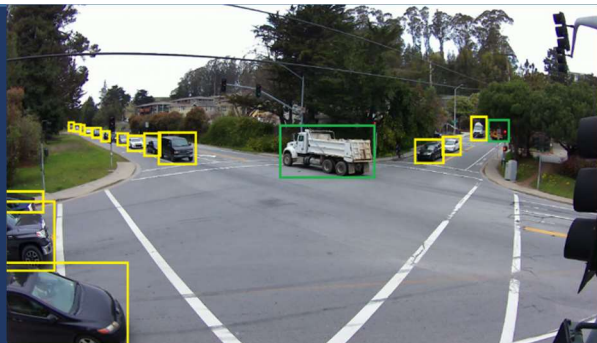


Project Story

Machine Learning Gives Santa Cruz Traffic the Green Light

www.cloudbrigade.com/smart-traffic-using-machine-learning



Summary

Santa Cruz, California, USA
 Transportation, Government and Smart Cities
 Machine Learning, IoT
 2020

Challenge

Cities all over America are experiencing road congestion at a growing rate each year, and the traffic light systems run on either timers or sensors. Neither of these systems are responsive to dynamic traffic conditions. Cloud Brigade wanted to use Machine Learning technology to maximize throughput while minimizing wait times of entire intersection traffic light systems, in real time.

Benefits

- Reduced traffic congestion
- Provides time for Traffic Engineers to be more proactive by eliminating the need for their constant attention to traffic signal patterns
- Reduction in payroll budget yields increased programs and equipment budget
- Future-proof and scalable
- Better reliability

Background

Ever wonder why it takes a half-hour or more to travel 2.5 miles on Mission Street on the West Side in Santa Cruz, California? The traffic lights only seem to laugh as you watch them cycle through from green to red without putting your car into gear. The Cloud Brigade team knew using a Machine Learning solution would deliver a better way to streamline the flow of traffic.

Business Challenges

- **Irresolvable Complexity** – Simply turning all the lights green at the same time on Mission Street ignores the traffic crossing or entering from the side streets
- **Inefficient Systems and Processes** – It should not take so long to get to the grocery store or home from work to see your family--this is a 4-lane road in a city of 65K
- **Skills and Staffing Gaps** – It really doesn't make monetary sense for the city to hire on a set of full-time Software and Traffic Engineers to solve this one problem
- **Antiquated Technology** – Using Traffic Signal Light (TSL) timers or metal/motion sensors at the intersections only exacerbates the existing congestion, especially during tourist seasons

Solution and Strategy

Cloud Brigade had already started looking at ways to use Computer Vision and a Reinforcement Learning model in tandem when it built its front-lawn defender, the [Poopinator](#). Cloud Brigade had proven it could think outside the box on complex projects and brought its "A team" to lead the development and design of an end-to-end IoT product that runs multiple Artificial Intelligence models to control the traffic light signals and respond to changes in real time, alleviating traffic wait times by reducing queueing and maximizing throughput.

The Locals Were Getting Restless

For years, Mission Street has become a bogged-down, carbon-monoxide-producing, loud-exhaust harboring nightmare for anyone trying to leave town north on route 1, or just get to the grocery store. The current technology that runs traffic regulation, Traffic Signal Lights (TSL's), is antiquated and to hire engineers and work crews to overhaul the whole system involves excessive operational costs at the local level. The Cloud Brigade team knew there was a better economical, scalable and reliable way.

This Was No Ordinary Project

Chris Miller, founder and CEO of Cloud Brigade, has been a Machine Learning guru for a few years now. Machine Learning is the process of building mathematical and software models that help a computer "learn" how to complete a task, and Computer Vision applies Machine Learning techniques to camera images. In the winter of 2019-2020, Chris found himself again waiting too long on Mission Street. Talking about his process, he explains "All these cars are just backed up from Bay Street, and it just doesn't make sense. There has to be a better way to streamline the flow of traffic."

Where it Began

He got the spark of interest in the traffic control problem from an RFP that the city of San Jose, California issued in 2019. They wanted to update all of their traffic lights from the current system.

"Updating the timing or sensor settings on a single intersection currently requires sending an entire team into the field to observe over many days, and then making changes and adjustments that might only be relevant for a specific time of year or day. The [current] systems are not dynamic in terms of responding to traffic in real time."

– Chris Miller, Founder and CEO

Although Chris didn't end up bidding on the San Jose project, Chris contacted Amazon in his exploration process and discovered the AWS DeepLens camera as a starting point in the Machine Learning solution.

Smarty-Pants Technology

The AWS DeepLens camera is essentially a computer within a video camera. As an IoT (internet of things) device with a Machine Learning edge, Amazon's DeepLens camera can deploy a computer vision model in real time. What does this mean? This means that the camera can be mounted in an intersection, can watch the traffic roll by, and without a human watching any footage, the camera can catalogue the number of vehicles entering and exiting the intersection, can discern how many of those vehicles are cars, trucks, buses, motorcycles, bicycles, human pedestrians, or tractor-trailers (or scooters, Segways, train-trolleys...I could keep going). Using a Machine Learning model that is deployed on the DeepLens camera's microprocessor, it can not only "see" what is happening, it can also detect the direction the vehicles are moving, what speed they are going, if they are turning or not, and even when there are traffic accidents. Pretty cool, huh?

Technology

- Software: Amazon SageMaker, Jupyter Notebooks, Simulation of Urban Mobility (SUMO) program
- Hardware: AWS DeepLens camera, traffic lights, sensors

Services

- Machine Learning & Artificial Intelligence: AWS DeepLens, SageMaker, EC2 & Lambda
- Internet of Things: Raspberry Pi
- Development: Python
- Cloud Infrastructure: AWS

Partners



Raspberry Pi

The data processing does take up a lot of personnel time and energy, and deploying multiple models to constantly run on the AWS DeepLens alone does get a bit space prohibitive, but our local Machine Learning engineer, Mark Davidson was happy to lend a hand. Using a Raspberry Pi, a small computer that can be adapted for many tasks including IoT projects, Mark was able to create a way to run the models in the field and upload the observations to the cloud.

How does this observation engine help us to make our traffic flow more smoothly? The answer is, this is only the first step. The majority of traffic lights in the United States are run on either a timer system, cycling through a certain number of seconds before switching the lights from green to red, or on a motion-sensor system, detecting the presence of traffic entering a waiting queue. This technology suits us fine in low-traffic situations, but leads to bottle-necking in heavy traffic times, as explained in queueing theory.

What is a Queueing Theory?

Queueing Theory, or the study of congestion and waiting in line, suggests that the complete absence of waiting in line is the symptom of a system that is too big for its demand, but that bottle-necking arises when the amount of vehicles coming into the system is too much for the system to handle. If this is a constant problem, then you need to increase the throughput to allow for the smoothest possible flow of traffic through the system. This can be done by building a larger system (adding a subway train, or maybe a vehicle road in the sky above the current road?), or by increasing the efficiency of the system regulation (traffic lights).

So, if all of our traffic signals are governed by old technology that either run timers or switch modes on a sensor, the bottle-necking that we see during rush hour—or during all daylight hours during the busy tourist seasons—will be too much for that technology to optimize. It is clear that revamping this system using that old technology will result in prohibitively excessive operational costs. This is where our second Machine Learning model comes in.

Building the Reinforcement Learning (RL) Model

You train a Reinforcement Learning (RL) Model to do its job by rewarding it with points when it does a good job, kind of like giving your dog a bone (positive reinforcement). Your dog can't understand (much) abstract thought, but it understands a piece of bacon. In this case, the bacon is a numeric point system, and we reward the model when the number of vehicles waiting at a red light is low. We also reward the model when the throughput, or total number of vehicles that run through the whole system is high. Since the rewards are multi-faceted, the model has to keep adjusting itself to find that sweet spot, and every time it takes an action and sees its reward or lack thereof, the model learns a little more and gets a little better.



This new Machine Learning technology will control for optimum traffic throughput while also minimizing the wait times of all travelers, may they be on Mission Street or on any of the cross-streets that intersect with Mission. For this, Cloud Brigade's internship program brought in UCSC Masters in Computer Science candidate Robert Sato during the summer of 2020. Robert describes himself as a person who is "interested in everything and anything related to Machine Learning and quantum computing," and so he was quite eager to employ his Python skills and to design our version of the model that I described above.

Technical Hurdles to Overcome

Robert mentioned "understanding Amazon's software" as his tallest hurdle. "While there are many example notebooks and decent documentation, the Amazon environment is so vast that it took a lot of time and effort to familiarize myself with building out the project in the Amazon environment." Still, he found using the AWS cloud infrastructure to be essential to connecting his RL Model to the virtual testing environment where his model would perform its training. "[Amazon] offered many well packaged services that were easy to use and bring together such as SageMaker, EC2 and Lambda...the many Amazon RL example notebooks, [had] a clear, consistent framework that I was able to apply to this project."

Vision Journey

When Chris had planned out the project back in February 2020, he'd envisioned creating a virtual world where he could feed in real world data using traffic information collected by the AWS Camera. This virtual world would be a place where the RL Model could learn the most efficient way to direct traffic. Robert found that he could use the AWS cloud infrastructure in concert with some existing traffic control software and a Python library.

Robert used Python's Gym library to interact with a simulation program called SUMO, or the Simulation of Urban Mobility program (built by German Aerospace Center and initially released in 2001). Through research, a lot of mathematics, software engineering and trial and error, Robert was able to use SUMO to create a simulation of traffic entering and exiting an intersection and to then use that simulation's data against his RL Model. Chris's plan was finally nearly complete.

Check out the video here: <https://youtu.be/Bm6Fo8XCahs>



A Successful Integration and a Technology Handoff

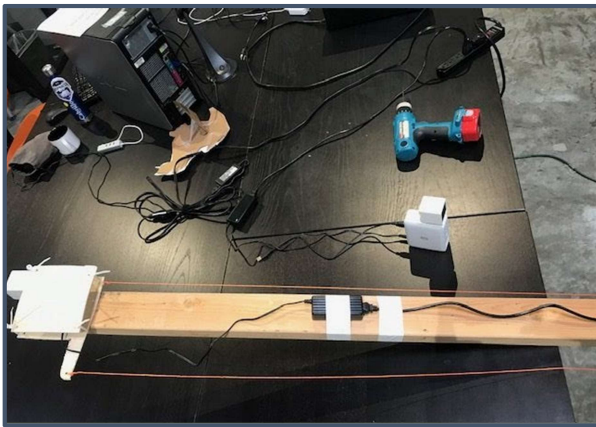
One thing we have learned through our work in Machine Learning is just how time consuming and intricate each step can be. A lot of times, the work of building up the infrastructure and lining up the data so that the different parts can all communicate with each other are actually more challenging than creating and using Artificial Intelligence. In this case, Robert was able to build out this model and its simulation environment in a way that we can deploy it to the AWS cloud infrastructure where the model can receive traffic data from the AWS DeepLens camera and in turn send its signal response out to the traffic lights in real time.

The final step will be to package the two Machine Learning models, the AWS DeepLens Camera, the AWS Cloud Computing instance, and the IoT traffic signal control product as a single end-to-end AI device that utilizes a Raspberry Pi and the AWS cloud infrastructure to respond to live changes in real time.



Opportunities

While this project is designed to fix our traffic problems where we live and work, it is not a niche product. From a one stop-light town to a complex traffic light system regulating large city surface streets, it can be scaled out easily. We look forward to the final product being useful in cities of every size all across the country. This is the perfect system for a small town that only sees an influx in traffic a few times a day, a large city that is in constant strain, or a medium sized city that is






constantly wondering where all the traffic is coming from in the first place. We're excited to take the next steps and can't wait to show more photos and videos of the process.



For more information:

cloudsales@cloudbrigade.com
+1 (831) 480-7199

101 Cooper St #218
Santa Cruz, CA 95060

 twitter.com/cloudbrigade
 facebook.com/cloudbrigade
 linkedin.com/company/cloud-brigade

1.28.2021