Basic video tracking for traffic analysis

Matt Bauman March 22, 2016

What it looks like right now



max_speed



Overwhelming response!

• I sent it to Bike Pittsburgh... and it took off!

• Definitely lots of interest

 Frequently asked: can I do this in front of my house / at stop signs / traffic lights?

Overwhelming response!



General image recognition is hard



Computer vision techniques

• Simplify the image as much as possible!

• Goal: reduce to core features necessary





Background subtraction: How I did it







Accentuating differences

- We don't care about colors
- Convert to grayscale by summing each color difference together





Threshold and blur

- We don't even care about levels of grey;
- Set a threshold to convert to black and white



Threshold and blur

- We don't even care about levels of grey;
- Set a threshold to convert to black and white



• We also don't care about windows and windshields, blur it to get uniform blobs:



Label components

• Library routine finds contiguous sections, and labels each with a unique identifier

										the for the herein 1
0	0	0	0	0	0	0	0	0	0	- Internet
6	0	0	0	0	0	0	0	0	0	
6	0	0	0	13	13	0	0	0	0	
6	0	13	13	13	13	13	0	0	0	
e	0	13	13	13	13	13	13	0	6	
0	0	13	13	13	13	13	13	13	0	
0	0	13	13	13	13	13	13	13	0	
6	0	13	13	13	13	13	13	13	6	
0	0	13	13	13	13	13	13	13	0	
0	0	0	13	13	13	13	13	13	0	
6	0	0	13	13	13	13	13	0	0	
0	0	0	0	13	13	13	13	0	0	

Label components

- Library routine finds contiguous sections, and labels each with a unique identifier
- Can think of it like the Paint Bucket tool:





Track across time

- The labelled components aren't stable across time
- New cars enter, identification numbers shift:



 Merge vehicle identifications by finding overlaps across video frames

Speed detection

• We have: speed of cars in pixels per second

• We want: speed of cars in terms of distance

• Need a ground truth



30 miles an hour?! But that's impossible! I've only been driving for 15 minutes!

Easy solution

• Time the cars across a measured distance



The wheels on the bus...

- I didn't have anypre-measured distances
- But PAT bus wheels are 279" apart:





Create a pixel density grid

• We have ten point estimates; use a linear model to estimate the conversion at all points



• Provides a consistency check: measurements agreed with each other to within 1mm/pixel.

• Cars can overlap and merge into each other:





• Cars can overlap and merge into each other:





• Solved: ignore areas between lanes





• Background changes over time



• Background changes over time



• Proposed solution: use inter-frame difference



- Storing HD video takes lots of space
 - 1GB per 8 minutes
 - 7.5GB per hour
- Processing videos offline takes lots of time
 - Approximately 6x slower than real-time on my old laptop in computation alone (if 100% automated)
 - Relatively un-optimized, could be faster
- Goal: Run in real-time (or close for rapid tuning)
- Goal: Only store resulting car positions?

How others have done it



Reynold's Video Speed Tracker

- Uses VASCAR-like fenceposts, tracking the timing of front bumper crossings
- Instead of background subtraction, uses interframe differences
- Even tries to track cars that cross in front of eachother
- C++ and OpenCV
- github.com/pfr/VideoSpeedTracker

Ideal outcome?

Handful of stationary traffic cameras in high value locations

• Pre-mapped distances

 Publicly available videos, maybe with semiautomated tracking.

Design Challenges

- What would an app look like? On a phone or computer? What knobs are needed for tuning?
- How would you mark known distances?
- Do you select the region of interest? Lanes? Directions?
- How would you display the results for validation?

Infrastructure challenges

• Is it possible to mount a permanent camera somewhere? Where? Physical access rights?

• Can we spec a camera? How expensive? How would you digitally access it?

• Where are high-value locations? Are there useful PENNDOT cameras? www.511pa.com

Let's do it!

github.com/pfr/VideoSpeedTracker github.com/mbauman/TrafficSpeed

- Coders:
 - Try using inter-frame difference
 - Try tracking new videos from last week
 - Put scale to videos of
 Fifth Ave & track them
 - Try compiling and installing Reynold's DVT app

www.511pa.com/CameraListing.aspx

- Others:
 - Brainstorm locations
 - Brainstorm outcomes (interventions and behavior changes)
 - Think about app design
 - Look up cameras,
 search for already
 available ones