S7170 - BICYCLE GREEN WAVES POWERED BY DEEP LEARNING

Edward Zimmermann (Nonmonotonic Networks) / joint R&D with (Michael Hartig, Mananging Director for Germany) GESIG. Gesellschaft fur Signalanlagen

GTC May 10 2017 San Jose, Ca. USA

Prologue

- Object recognition is solved. Object recognition is now more a less a commodity.
- A lot of AI too is now commodity.
- Low power embedded hardware is now available and relatively inexpensive.
- Thinking back just a few years. We've come a long way.

With these things now so widely available and relatively inexpensive a large range of new applications became possible– many unthinkable just 5 years ago.

Let me describe one to you:

“Smart Traffic Lights for bike paths”

Motivation

- The belief that the bicycles (both self-powered and e-bike/pedelecs) **are** the future of urban transport alongside (self-driving) electric (battery or hydrogen) cars, busses and rail services
  - There has been a great renaissance in bicycling.
  - In central London, for example, bicycles already account for nearly ¼ of rush hour traffic
  - In cities such as Copenhagen 41% of the population bicycle to work (or school)
  - Even cities such as Los Angeles and Tel Aviv where bicycles were once relatively rare and considered "at best" a children's toy we are seeing more and more people turning to them. (in Israel the trend is dominated by pedelecs).
  - To meet environmental goals and increase the attractiveness of cities many communities are increasingly looking to discourage use of motorcars in the city centers and improve their public transport systems.

“Green” Transportation Hierarchy

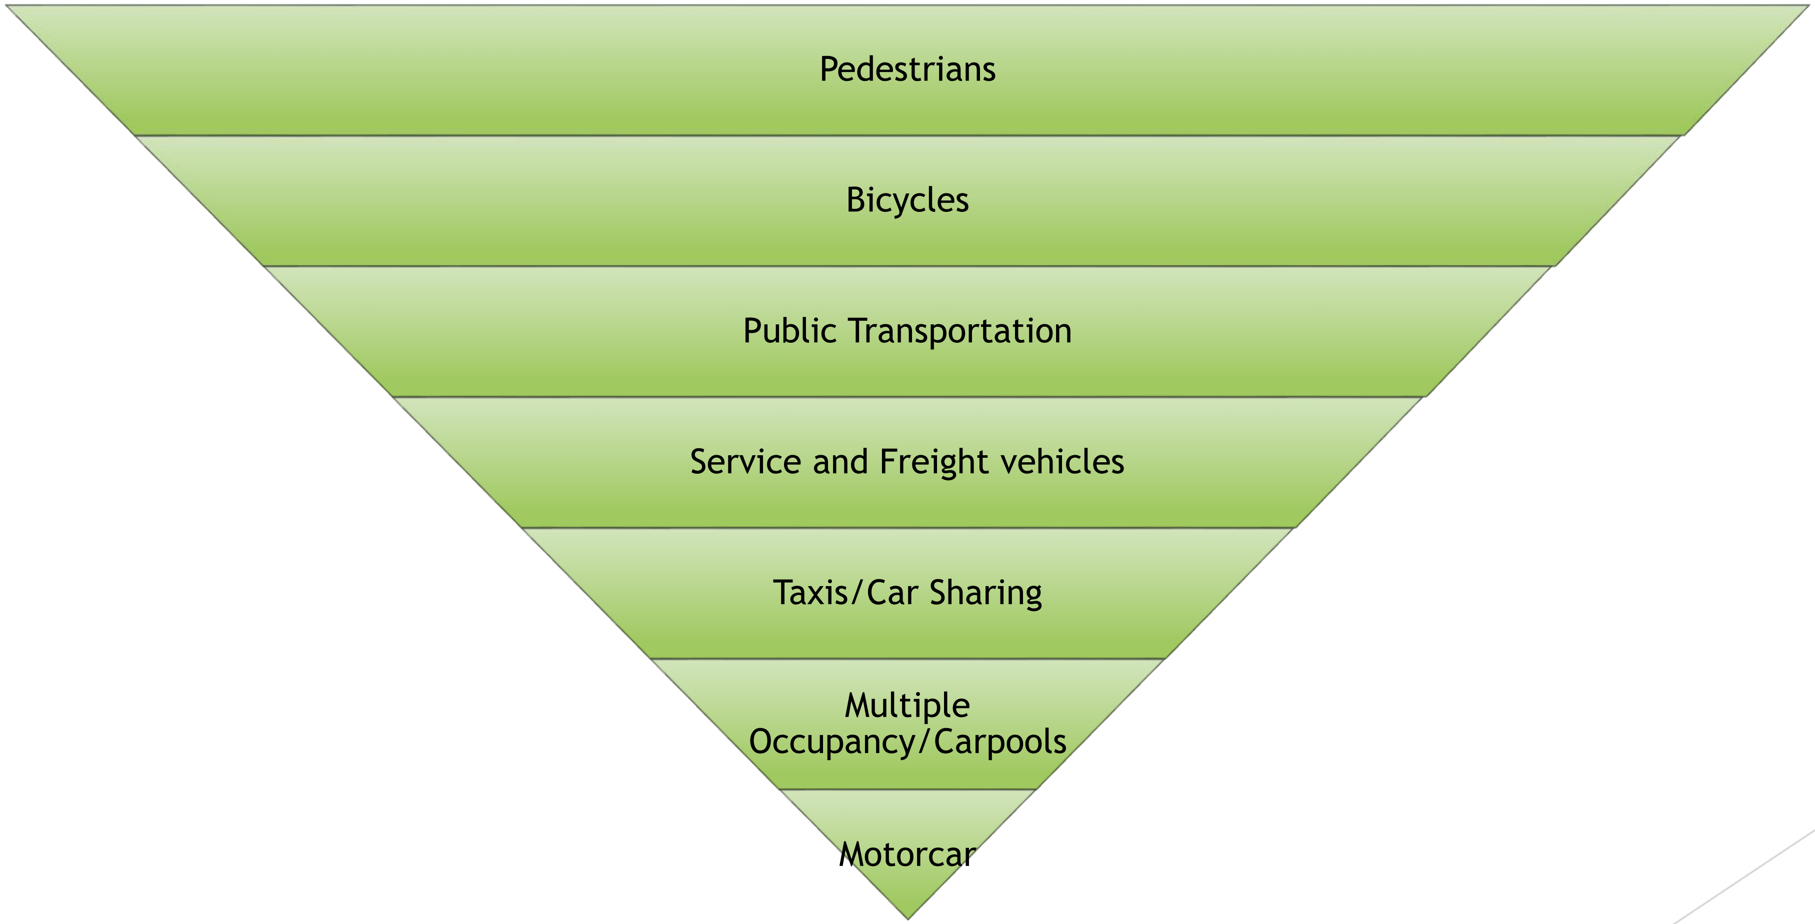

Bicycles are increasingly viewed as a component of public urban transportation

- Bicycling is increasingly viewed as part of the public transport concept– many communities have invested to bicycle sharing concepts. In Munich we even have two: one operated by the German Railroad (“Call a Bike”, launched in 1998, originally a start-up) and one operated by the Munich Transit Authority (“MVG Rad” since 2015) .
- Bike sharing is a global trend. Boston-Budapest-San Francisco Bay Area-Paris etc. Even car crazed China (where bicycles ruled just a few years ago) is experiencing a boom (more than 40 start-ups and growing).
- Bicycle ownership and use too are growing and bicycle lanes are being laid out at a rapid pace.

Communal Bicycle Sharing (Global Trend)

source: Wikipedia

v · t · e	Bicycle-sharing systems [hide]	
		List of bicycle-sharing systems
Asia		ADBC Bike share (Abu Dhabi) · CityBike (Kaohsiung) · Ddareungi (Seoul) · Hangzhou Public Bicycle · İsbike (Istanbul) · LinkBike (George Town, Penang) · Mobike (China) · Ofo (China) · Tel-O-Fun (Tel Aviv) · Tiruchirappalli Bicycle Share · YouBike (Taipei, New Taipei, Taoyuan, Hsinchu, Taichung and Changhua)
Europe		Aarhus City Bikes (Aarhus) · BicikeLJ (Ljubljana) · Bicing (Barcelona) · BikeMi (Milan) · BuBi (Budapest) · Bycyklen (Copenhagen) · **Call a Bike** (Germany) · Coca-Cola Zero Belfast Bikes · Dublinbikes · Helsinki City Bikes · Oslo Bysykkel · OYBike (Reading and Farnborough, UK) · Santander Cycles (London) · Sevici (Seville) · Stockholm City Bikes · Styr & Ställ (Gothenburg) · ToBike (Turin) · Valenbisi (Valencia) · Vélib' (Paris) · Vélivert (Saint-Étienne) · Vélo'v (Lyon) · Velobike (Moscow) · Vélopop' (Avignon) · Veturilo (Warsaw) · Villo! (Brussels)
North America	Canada	Bike Share Toronto · BIXI Montréal · VeloGo (Ottawa) · Mobi (Vancouver) · Sobi Hamilton
	Mexico	EcoBici (Mexico City) · Mi Bici (Guadalajara)
	United States	ArborBike (Ann Arbor) · Arcata Bike Library (Arcata) · Austin B-Cycle · Battle Creek B-Cycle · Bay Area Bike Share (San Francisco Bay Area) · Baltimore Bike Share · BICI Bike Share (Albuquerque) · Bike Chattanooga · Biketown (Portland) · Boise GreenBike · Boulder B-Cycle · Broward B-Cycle · Bublr Bikes (Milwaukee) · Capital Bikeshare (Washington, D.C.) · CAT Bike (Savannah) · Charlotte B-Cycle · Citi Bike (New York City; Jersey City) · Columbia County B-Cycle · CoGo Bike Share (Columbus) · Decobike (Miami) · Denver B-Cycle · Des Moines B-Cycle · Divvy (Chicago) · Fort Worth B-Cycle · Great Rides B-Cycle (Fargo) · GREENbike (Salt Lake City) · Greenville B-Cycle · Healthy Ride (Pittsburgh) · Heartland B-Cycle (Omaha) · Houston B-cycle · Hubway (Boston) · Indego (Philadelphia) · Kansas City B-Cycle · Indiana Pacers Bikeshare (Indianapolis) · Link Dayton Bike Share · Madison B-Cycle · Nashville B-Cycle · Nice Ride Minnesota (Minneapolis-Saint Paul) · Pronto Cycle Share (Seattle) · Rapid City B-Cycle · Red Bike (Cincinnati) · Richmond Bikeshare · Roseburg B-Cycle · San Antonio B-Cycle · Spartanburg B-Cycle · WE-cycle (Aspen) · Wolf Ride Bike Share (Stony Brook University, New York)
Oceania		Melbourne Bike Share · Nextbike (Auckland)
South America		Bike Rio (Rio de Janeiro) · EcoBici (Buenos Aires)
Defunct		Copenhagen City Bikes · SmartBike DC · Helsinki City Bikes 1st generation
Companies		8D Technologies · B-cycle · PBSC Urban Solutions · JCDecaux · Motivate · Nextbike · Smoove · Social Bicycles · Zagster · Serco

Transit Non-neutrality goals

- Increase system capacity without costly new infrastructure– use what one has in a more optimal manner.
- Increase the attractiveness of urban public transport and reduce energy demands.
- Improve the general quality of life and promote “greener” and “smarter” transport.

This increasingly also means prioritizing public transport and encouraging the use of bicycles.

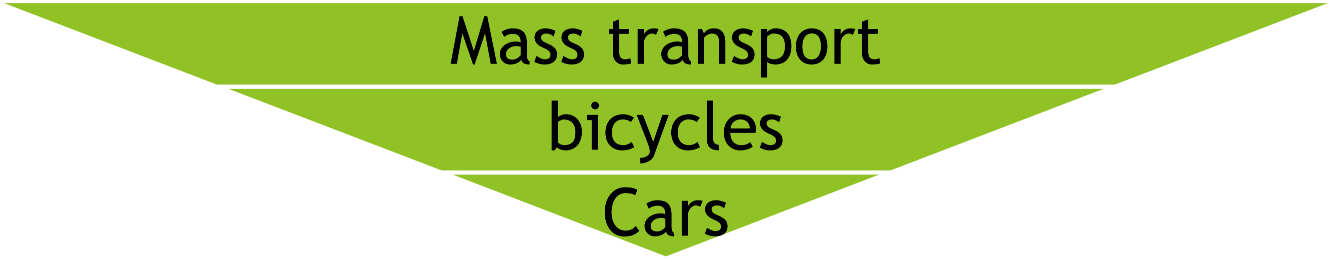

Green Waves

Green waves are when one crosses a series of traffic signals that remain green. The art of green waves are to time or coordinate the lights to allow for a more continuous flow through a series of signaled intersections. Green waves are highly desired.

- For automobiles they reduce emissions, fuel consumption, wear and tear.
- For bicycles green waves not only make cycling more efficient and attractive but also empirically reduce the likelihood that cyclists will endanger themselves by running red lights.
- For public transport they improve the attractiveness by measurably speeding up transit, improving reliability and reducing delays.

Communities are increasing focusing on prioritizing coordinated “green waves” for public transport through “Transit Signal Priority”.

- To increase the attractiveness and safety of bicycling there is an increased desire and demand for “green waves” for cycling.

In Germany TSP has the highest priority and this can conflict with bicycling.

Transit Signal Priority (among other techniques)

- Green Extension– extend the green interval to static (pre-set) max if a transit vehicle is approaching. (min. “just” missed green)
- Red Truncation/Early Green– shorten the phase when a bus approaches.
- Green Truncation/Early Red– if the bus is approaching during a green but too far away to reach it, the green the ended early. The phase is shortened a static (pre-set) amount. The idea is to min. the red cycle the bus would experience
- Actuated Transit Phase/Phase rotation/Phase Insertion

In many of these cases the timings are not altered only the phase is truncated. Red truncation is quite popular in Germany.

NOTE: TSP and visible phase countdown (as in San Jose) are not compatible.

“Dynamic Green Waves” for cars

Typically rule-based with human intervention (road in Munich going to the Allianz Arena where Bayern München play or around the main trade fair etc.)

- Times are shifted and readjusted in 15 min. periods.

Bicycle Green Waves (Prior Art)

Mainly passive: Fixed timings

Typically traffic lights are timed around a fixed transit speed (roughly 20 km/h).

Examples:

- Valencia Street in San Francisco: 2009 the lights were retimed to provide a 13 mph green wave (in both directions) from 16th Street to 25th Street (Mission Distict)– SFMTA has suggested it was the first ever bi-directional green wave. They added 14th Street in 2012, 11th Street in 2013. More streets have been since added (North Point from Stockton to Polk Streets, Folsom from 15th to 24th Streets, and Fulton from Laguna to Steiner Streets.
- Nørrebrogade in Copenhagen: 2.5km in length @ 20 km/h (~12 mph)

  (2.0 add green extension)
- Portland, Or: North Vancouver and North Williams corridor one way couplet.
- Amsterdam, Utrecht …

Demand Actuated Traffic Signals: Bicycle Detection

- Magnetometer: need ferromagnetic materials which are increasingly less common in bicycles (alu, plastics and composites, bamboo, wood etc.).
- Inductive Loop Sensor. Greatly improved over the years. Bicycles with metal rims are easily detected by properly designed and adjusted systems.
  - Properly designed they should not detect large vehicles in the adjacent traffic lane. (Unfortunately this is not always the case and sometimes they don’t detect things)
  - California enacted a law to require all new and upgraded traffic signal detector to handle bicycles (and motorcycles). Most are type “D” quadrupole loop.
- Doppler Radar (e.g. 25 Ghz). Less prone to problems.
- Image recognition: Video and other techniques (which are not “in road”).

Background

- Gesig’s intelligent traffic management system: VnetS.
  - VnetS was developed with the City of Munich. First launched in 1999 to cover the Munich Trade Fair its original intent was to increase the flow of traffic during peak periods and integrate with both the parking lot management and main highway control systems (“Autobahn Leitsystem”).
- Also extended to I2V (pilot with a German automotive company): information about traffic lights are sent to the vehicle (Infotainment integration).

Gesig was working on an idea of using (opt-in) cellphones to “track” bicycles- GPS data etc.– and detecting “bicycle swarms” (patent pending)– and better integrate real-time transit statistics.

Gesig had in Bonn a number of traffic lights outfitted with Raspberry Pi boards (B+) and cameras.

GTFS: General Transit Feed Specification

- Beyond traffic control the system integrated GTFS and GTFS-Realtime

GTFS was originally developed by Google (Chris Harrelson) in 2005 to help integrate transit data into Google maps– the "G" was originally "Google". It was first deployed in Portland, Or.

Since prior to GTFS there was no standard (not even de-facto) for public transit timetables it was quickly adopted and in 2009 the "G" became "General".

The format is a simple collected of CSV files (min. 6 and up to 13) collected into a ZIP archive.

- There is a European Standard (Transmodel) but even in Europe GTFS is increasingly popular.
- GTFS is already available in >500 cities. There are also a number of open interfaces. In Bonn, for example, via the VRS-interface (OpenService).

GTFS-Realtime

- Provides “real-time” updates about the fleet.
- Uses Protocol Buffers (a kind of binary encoding not much unlike BER in ASN.1)
- In addition to a number of positional information objects and things like occupancy, alerts (accident, weather, medical, strike etc.)
  - Traffic Congestion Level (CongestionalLevel):
    - Unknown
    - Running smoothly
    - Stop and Go
    - Congestion
    - Severe Congestion
- So we can not only know where the vehicle is (especially relative to a traffic light) but we also can have an indication of the current traffic conditions.

How I got involved

Video Detection

- Relatively “dumb” examining the “pixels” of the scene without bicycles to with bicycles.
- Can’t detect people (neither can inductive loop or magnetometer sensors).
- Sometimes can detect cars
- Tend not to be designed to distinguish objects but are used as an activation.
- Typically don’t see that well– and so it is quite common that they use the presence of bicycle headlight illumination to detect bicycles at night.

The system thinks it sees bicycles so activates the traffic light– “Green on demand”

Modern Radar Detection

- Can detect cars relatively well– especially those with metal bodies
- More reliable than inductive loop– detects bicycles that might not have been detected.
- Can detect bicycles well and in poor environmental conditions– rain, snow, fog.
- Does not depend upon lighting or vision.
- Can detect pedestrians (people contain, after all, water)

The State of the art (24 Ghz Doppler)

- Can detect stationary as well as moving objects. Can both count objects as well as detect speed.
- Often demands, like, inductive loop installation under the pavement ("in road")– in Germany the cost alone of the hole is roughly 700 EUROs.

Using CNNs to detect bicycles, people, cars.

- Current state of doing relatively fast “real time” object detection: R-CNN, SSD, YOLO

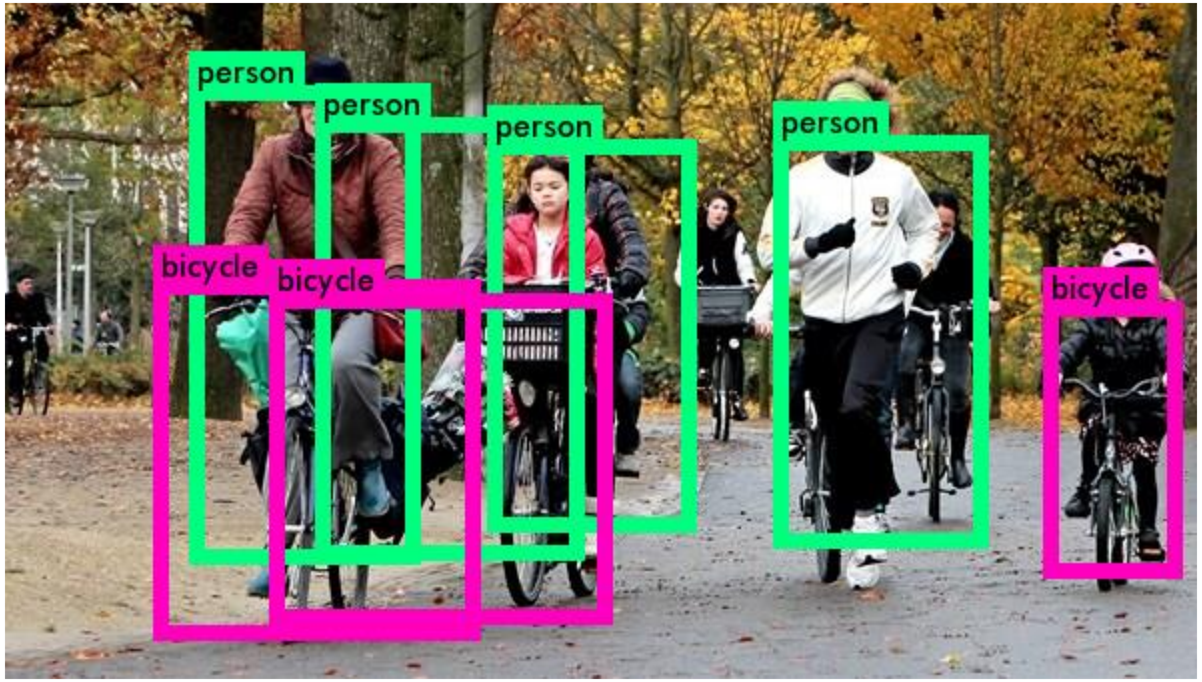

Advantages

- Cheaper
- Easily managed and updatable (flash a new system on demand)
- Really ideal for the application– we use our eyes after all to see traffic lights and traffic so it is just natural that a traffic light should it have eyes could spot us.
- Have computational resources that we can use to distribute other kinds of intelligence (decisions) to the individual lights.

Two years of so this would not have been possible. Before the Tegra X1 we really did not have a low power SoC with sufficient power.

Object Recognition

- Originally started off (initial proof of concept) with my own Yolo fork– part of a BI pipeline I was working on)
- The original YOLO was designed for the VOC competition using the Pascal training set with its 20 classes:
  - Persons, animals (bird, cat, cow, dog, horse, sheep), vehicle (airplane, bicycle, boat, bus, car, motorbike, train), indoor (bottle, chair, dining table, potted plant, soft, tv/monitor)
- Modified for the particular classes we wanted: bicycles, motorcycles, pedestrians, cars, trucks.
  - (I did also a version which included busses and trams– and a vastly enhanced training corpus to deal with some of the unique European designs) but it was more POC for another concept as don't expect to see trolley cars on bike paths)

Performance

- Took images captured under different conditions by the cameras in Bonn.
- Trained a number of bicycle images (to augment the corpus) and tested using a machine equipped with an NVIDIA Tesla K40 (an older machine which I had available).
- First crude test using a CPU Intel i7 for inferencing (and no special library such as Intel's MKL-DNN): 1 second.
- Originally tested (for fun and curiosity) on a Raspberry Pi B+ (the boards deployed in the Bonn traffic lights).

  Time to process a single image: 120 seconds (did not explore why).
- Initial tests on a NVIDIA Tegra X1 (Jetson Development board). Network was a bit large (4 GB DDR4 RAM). After some crude pruning: ½ - 1 second per image.
- After a bit or reorganization (and adjusting the power management of the X1): 3 (looking at 4) fps.
- Even with poor (blurred, odd angles, odd lighting etc.) was still was able to generally recognize bicycles– and if "missed" in one frame tended to get spotted in one of the adjacent frames.
- In packs of bicycles, of course, did not spot them all but generally always spotted some.

Performance interpretation

- Since bicycle detection is not the only form of activation available even under worse conditions– where the cyclist should anyway exert an extreme degree of caution due to cars turning– it is “tolerable” to not recognize a cyclist or pedestrian and perhaps require manual activation (“push the button”).

  Note: 81 percent of crashes between bicyclists and motorists occur at or within 50 feet of an intersection, according to data collected by the City of Minneapolis between 2000-2010 and published in a report titled, “Understanding Bicyclists-Motorist Crashes in Minneapolis, Minnesota.

- Looking now at speed, distance and processing speed

Speed		10 fps (100ms)	5 fps (200ms)	2 fps (500ms)
10 km/h	2.8 m/s	0.28m	0.56m (1’10”)	1.4m (4’7”)
20 km/h	5.6 m/s	0.56m	1.12m (3’8”)	2.8m (9’2”)
30 km/h	8.3 m/s	0.83m	1.66m	4.15m

- **Size_in_Image = Focal_Length * Object_Size / Object_distance**
- Ratio of the image size when traveling a distance x from d: **d** / (**d –x**)

Issues

- Tradeoffs between level of confidence and identification
  - Cost of miss recognition?
  - Better to be more certain?
- What do we expect to see on a bicycle path under different conditions?
- Important utility:
  - Cost of non-recognition is relatively low.
  - We have multiple chances to recognize (remember we are in the first instance interested in activation followed by data collection)

Can we quantify performance?

- Different lanes, cameras, conditions
- Non recognition of an object in a frame is not the same as non-recognition.

Strategies for active bicycle Green Wave

First off observations:

- “Bus priority” (Transit Signal Priority) conflicts with bicycle green waves.
- Many transit authorities demand TSP as the highest priority.
  - In Munich, for example, red truncation is order of the day. This tends to mean that cyclists (and pedestrians) need to often wait while a bus or tram approaches, while it boards and until it crosses (the middle of) the intersection.
- Transit authority priority is to keep to schedule: better wait at a boarding stop than at a traffic light. (remember that we also have GTFS-Realtime)
- Road traffic and bicycle lane traffic are de-coupled.
- Bicyclists would rather ride slower than wait at a stop light. Total speed (and time to destination) is often not prioritized. Bicyclists don’t like stopping at traffic lights.

Regret and Utility

- A control problem with hysteresis. Recall one can’t wildly change phase.
- First attempt: assume rule base green phase extension when TSP is not competing for phase. Need data.
- Use telemetry data to track losses.
- Use ML to generate alternative rules and policy goals.

- Initial simulation suggested that the thesis might be valid.
- Permission was granted in early Autumn 2016 to engage in a R&D field test.
- Business delay followed in early 2017 by…
- Hope to get back on-track as soon as possible.

Concurrent developments

(beyond finally going into the field)

Thank You