



Computer Vision Technologies and the Citizen

How visual sensing is changing “how we see ourselves”

João Paulo Costeira

Signal and Image Processing Group



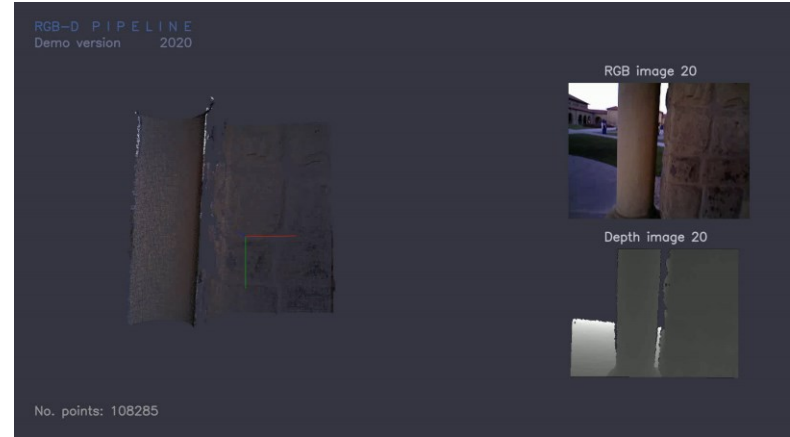
TÉCNICO
LISBOA



Institute for Systems
and Robotics
Lisboa



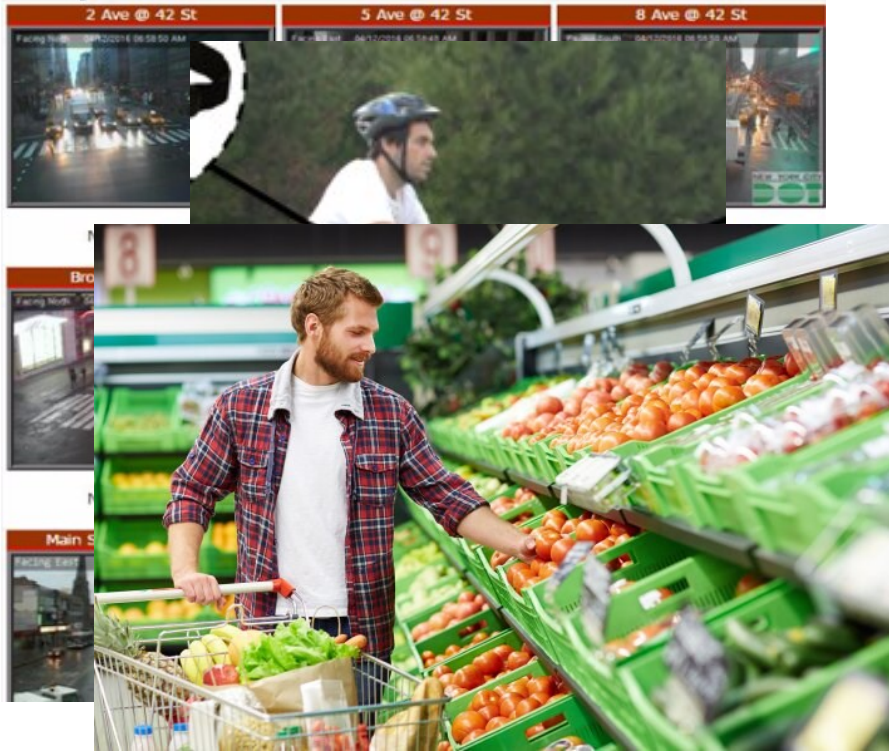
My “vision”



I attended ICCV #1 in London



How I got here ...



Transportation & Cities

The human factor gets in

... perplexity by human
decision-making

The main message



Category: Dress, Top, Skirt
Attributes: Sleeveless, Long-sleeved, Short-sleeved, Long-sleeved, Sleeveless, Long-sleeved, Short-sleeved, Long-sleeved

Docker container

TagMyOutfit

Tag-My-Outfit is a classifier that predicts the category and the attributes of a piece of clothing viewed in a given image. It is containerized and gRPC enabled so that it can be used standalone or in a pipeline using AI4EU experiments platform. [read more](#)



Docker container

Yolo V5 - Object Detection

YoloV5 model released by Ultralytics for object detection. The component takes one image as input and outputs the coordinates of the bounding boxes of all objects detected in the image. [read more](#)



Docker container

Openpose

This asset implements Openpose, a widely used open-source human pose estimation algorithm, as a containerized gRPC service. [read more](#)

Acknowledgements




Cláudia Soares
Univ. Nova



Manuel Marques
IST



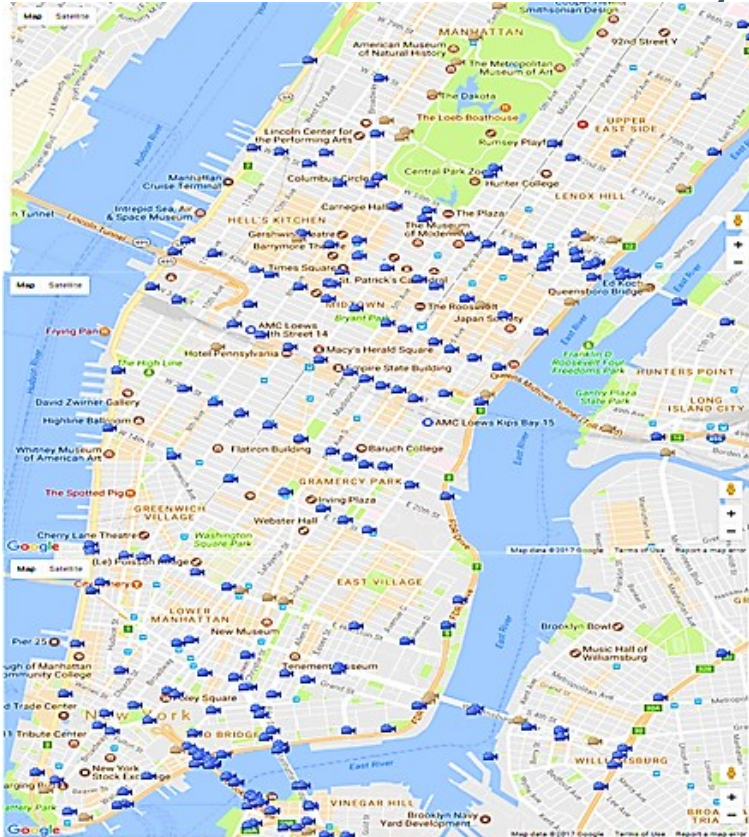
Qiwei Han
Nova SBE



Jose Moura CMU,
Shangan Zhang UC Berkeley

Eyes on the infrastructure

The City Scape Cameras



NYC DOT SEM



NYC DOT SEM



NYC DOT SEM



NYC DOT SEM



NYC DOT NYCWIN



NYS CAM



NYC SEM



NYC DOT SEM



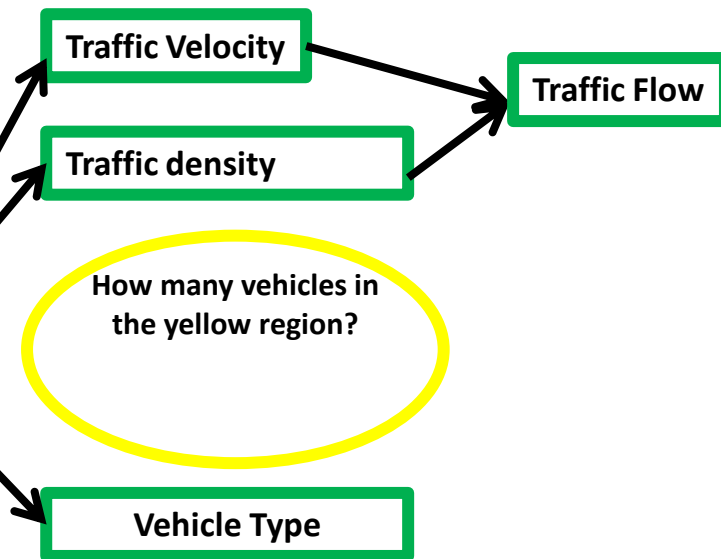
NYC DOT NYCWIN

From images to traffic understanding

Deep understanding of urban traffic


Extracting vehicle counts from streaming real-time videos

captured by **multiple low-quality city cameras**.




Count by Density !

Mapping pixels to vehicle density $\sum d(x,y)=1$



$$\sum d(x,y)=1$$

Images are “just numbers”



Training Data: Annotation vs “car density”

The screenshot displays a video annotation interface. The main window shows a street scene from a high-angle camera, facing south, dated 2016-04-29 at 18:00:21. The scene includes a pedestrian, a white truck, and a red car. Bounding boxes are drawn around these objects, labeled 1, 2, 3, and 4. A pink line connects the corners of the bounding boxes. A zoomed-in view of the truck is shown in the bottom right corner.

Metadata at the top: Face: South, Date: 2016/ 4/29, Time: 18:00:21, Weather: Cloudy, View: Normal, Flow: Horizont.

ID	Type	Dire...	Prev...
<input checked="" type="checkbox"/>	1	Mini ...	Down 0
<input checked="" type="checkbox"/>	2	Medi...	Down 0
<input checked="" type="checkbox"/>	3	Black	Down 0
<input checked="" type="checkbox"/>	4	Other	Down 0

ID
<input checked="" type="checkbox"/> 1

Bottom controls: Show annotation, Jump 1, First, Previous, Next, End, Play, Save, Inherit, Check.

On the right, a heatmap visualization shows the car density. The background is dark blue. A bright blue square is located in the upper right quadrant, and a smaller orange square is located in the lower right quadrant.

CityCam Dataset

- 60,000 annotated frames with around 900,000 labeled vehicles
- Rich annotations: vehicle count; type; bbox; re-id; orientation; time; weather
- 60,000,000 images for testing

The screenshot displays the CityCam dataset viewer interface. The main window shows a street scene with vehicles annotated with green bounding boxes and numbers. The interface includes a control panel at the top with fields for file path, camera direction, date, time, weather, and view. A table on the right lists vehicle details, and a bottom control bar contains playback and display options.

File path: C:\Users\Shaolong\Dropbox\City-code\数... 浏览... 打开

镜头方向: 朝南 日期: 2015/12/24 时间: 15:29:08 天气: 晴天 视野: 正常 车流方向: 垂直

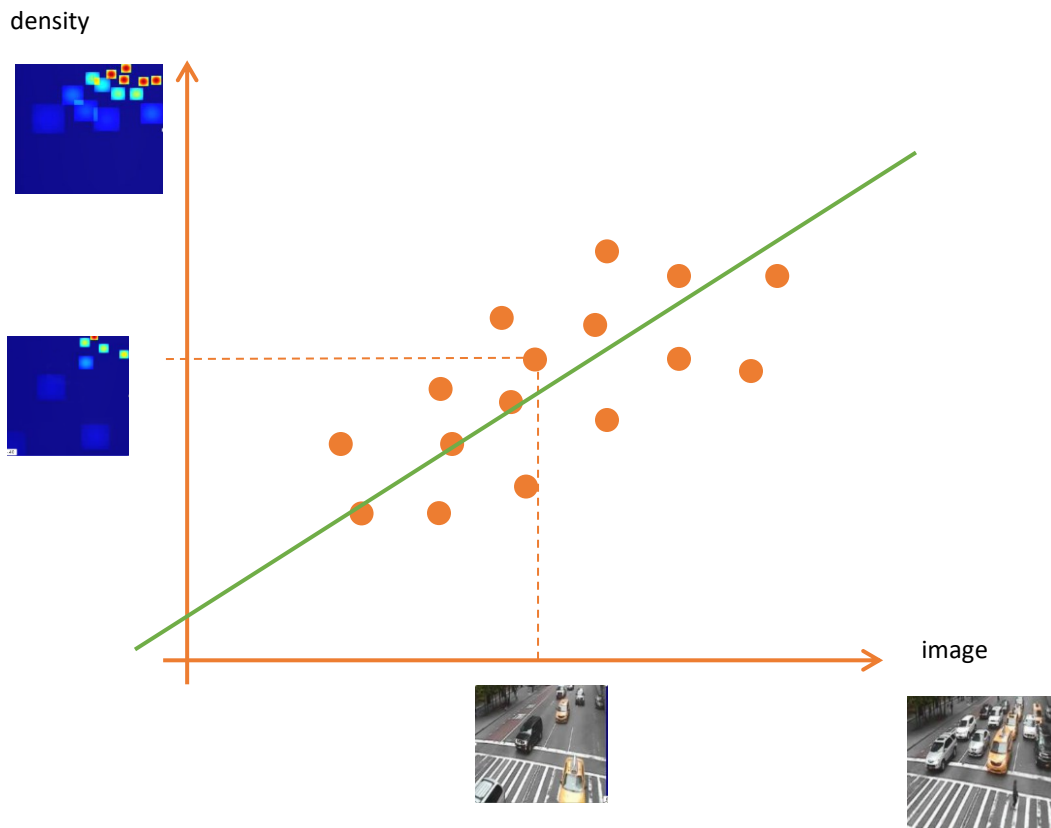
Facing South 12/24/2015 03:29:08 PM 上一帧

Facing South 12/24/2015 03:29:08 PM 当前帧

编号	类型	车头	上帧
<input checked="" type="checkbox"/>	1 小卡	下	2
<input checked="" type="checkbox"/>	2 面包	下	3
<input checked="" type="checkbox"/>	3 黑轿	下	0
<input checked="" type="checkbox"/>	4 面包	下	4
<input checked="" type="checkbox"/>	5 出租	下	5
<input checked="" type="checkbox"/>	6 黑轿	下	7
<input checked="" type="checkbox"/>	7 出租	下	8
<input checked="" type="checkbox"/>	8 出租	下	9
<input checked="" type="checkbox"/>	9 轿车	下	0
<input checked="" type="checkbox"/>	10 轿车	下	10
<input checked="" type="checkbox"/>	11 轿车	下	11
<input checked="" type="checkbox"/>	12 出租	下	12
<input checked="" type="checkbox"/>	13 出租	下	13
<input checked="" type="checkbox"/>	14 黑轿	下	0
<input checked="" type="checkbox"/>	15 出租	下	14
<input checked="" type="checkbox"/>	16 出租	下	15
<input checked="" type="checkbox"/>	17 轿车	下	16
<input checked="" type="checkbox"/>	18 出租	下	17
<input checked="" type="checkbox"/>	19 出租	下	18
<input checked="" type="checkbox"/>	20 出租	下	19
<input checked="" type="checkbox"/>	21 出租	下	0
<input checked="" type="checkbox"/>	22 出租	下	20
<input checked="" type="checkbox"/>	23 出租	下	21
<input checked="" type="checkbox"/>	24 出租	下	22
<input checked="" type="checkbox"/>	25 出租	下	23
<input checked="" type="checkbox"/>	26 出租	下	24

显示车框 显示编号 审核 显示标注 跳转 2 首帧 上帧 下帧 尾帧 播放 保存 继承 检查


Learning to count by (linear) regression



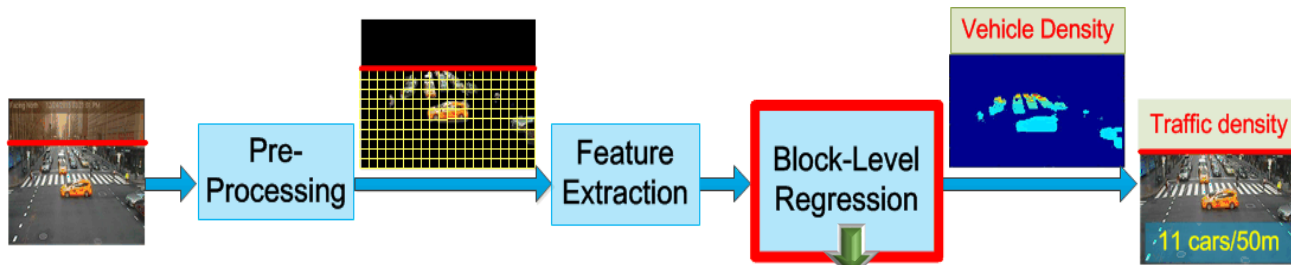
Algebraic solution – Large Scale **Low Rank** Regression

Block-level regression with rank constraint.

Embed road geometry in the weight matrix.




Optimization based vehicle density estimation




$$\min_{\mathbf{W}} \frac{1}{2N} \sum_{i=1}^N \sum_{j=1}^J (\mathbf{w}_j^T \mathbf{x}_j^{(i)} - D(B_j^{(i)}))^2 + \alpha \|\mathbf{W}\|_F^2 + \beta |\mathbf{W}|_1$$



s.t. $\text{rank}(\mathbf{W}) \leq r$



TrueCount: 16




EstCount: 16.8317



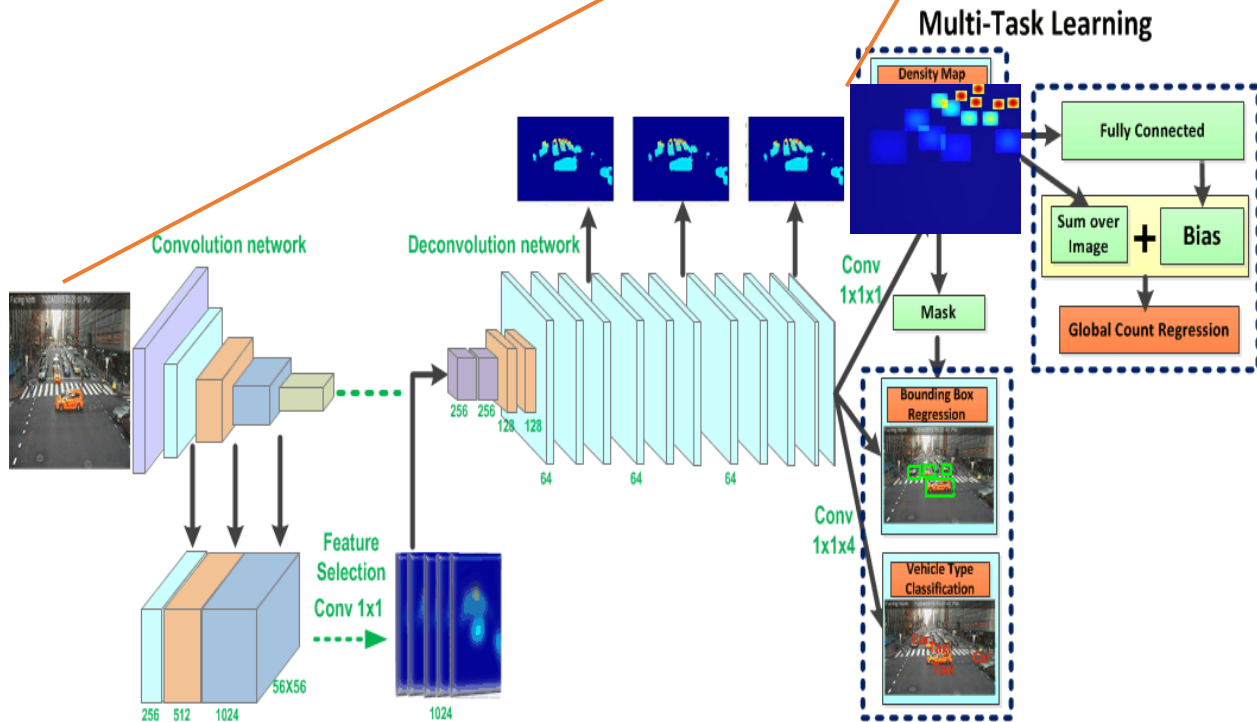
Deep Neural Nets, a new level in performance

Inspired by how the brain works




Regression with Deep Neural Networks


$$L_D(\Theta) = \frac{1}{2N} \sum_{i=1}^N \sum_{p=1}^P \|F(X_i(p); \Theta) - F_i(p)\|_2^2$$



■ Results of FCN-rLSTM for the parkway camera 691



Estimated vehicle count for different time and conditions **MAE = 1.67**



■ Results of FCN-rLSTM for the downtown cameras

Camera 410

Original Image Estimated Density Map




Camera 511


Original Image Estimated Density Map



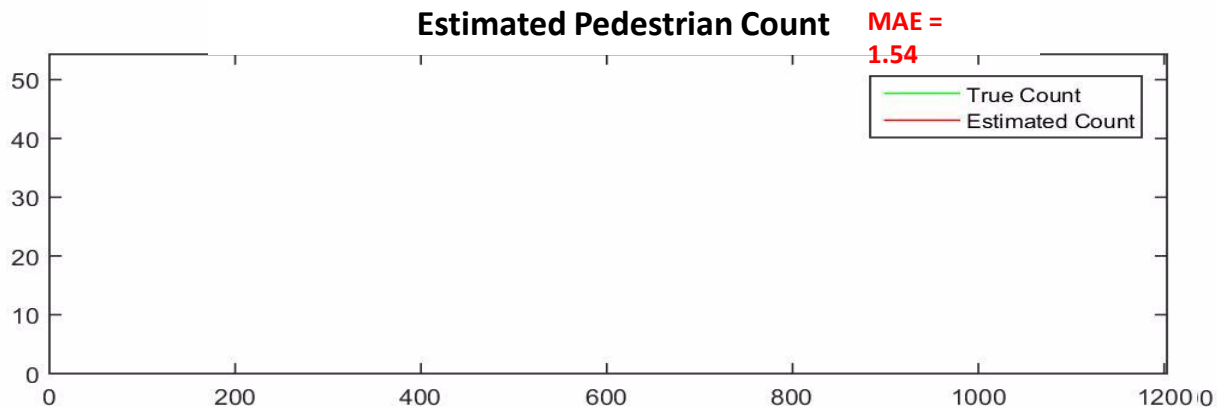

Estimated vehicle count for different time and conditions




Estimated vehicle count for different time and conditions



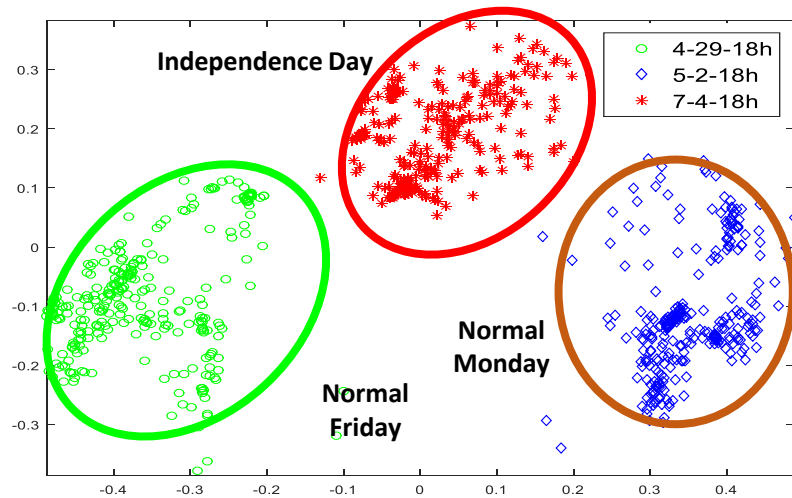
Can learn to count anything! It's the data ...



The pulse of the city

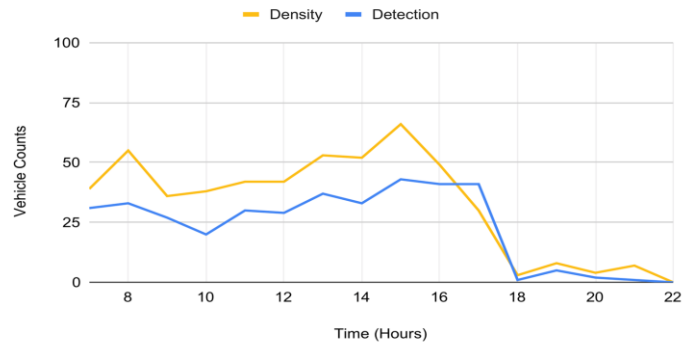

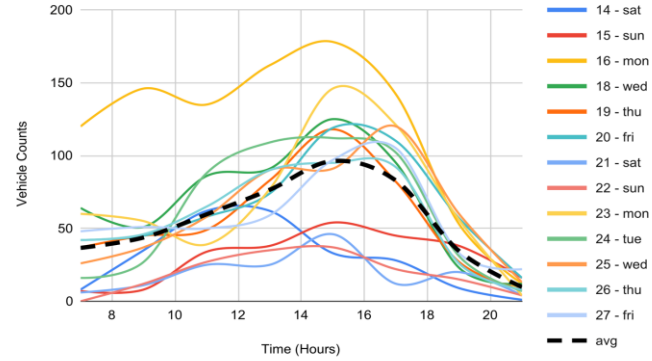


Estimated Count for different days



Isomap of vehicle density for different days

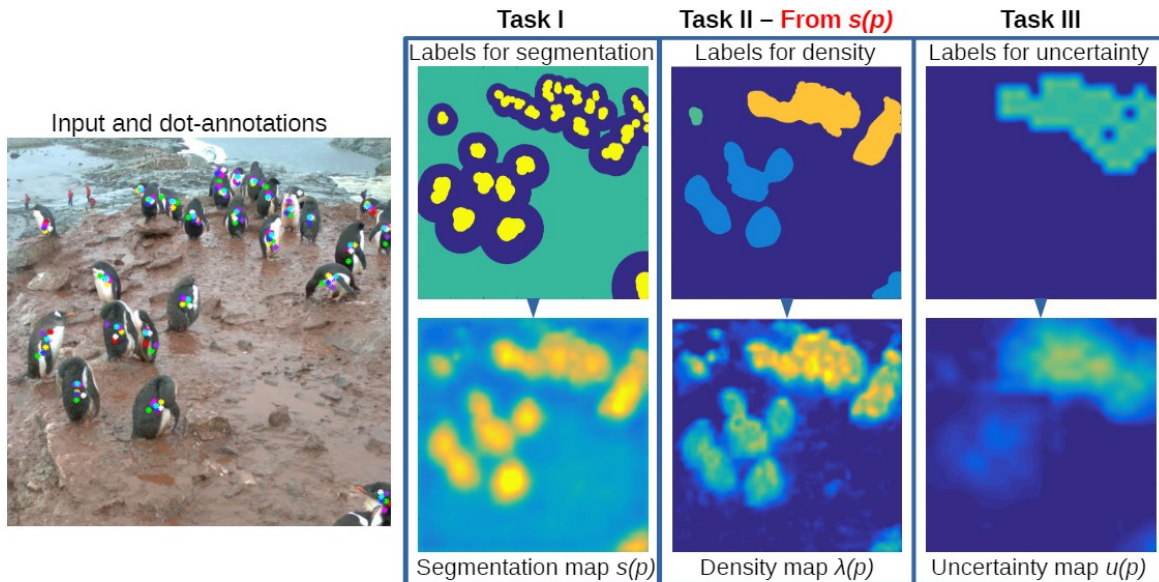
Main Avenue vs Highway (Tallin)



References

1. Shanghang Zhang, Guanhang Wu, João Paulo Costeira, José M. F. Moura, “*Understanding Traffic Density from Large-Scale Web Camera Data*,” **IEEE Conference on Computer Vision and Pattern Recognition (CVPR)**, Hawaii, USA, July, 2017
2. Shanghang Zhang, Guanhang Wu, João Paulo Costeira, José M. F. Moura, “FCN-rLSTM: Deep Spatio-Temporal Neural Networks for Vehicle Counting in City Cameras,” **IEEE International Conference on Computer Vision (ICCV)** , Venice, Italy, 2017.
3. Han Zhao · Shanghang Zhang · Guanhang Wu · José M. F. Moura · Joao P Costeira Geoffrey Gordon, “Adversarial Multiple Domain Adaptation,” **Conference on Neural Information Processing Systems (NIPS) 2018**, Montreal, 2018
4. **US Patent** ; Filed: April/5/2018 Title: Deep Learning Methods for Estimating Density and/or Flow of Objects, and Related Methods and Software Applicants: Carnegie Mellon University and Instituto Superior Técnico.

A. Zisserman (Univ. Oxford)



Density-based counting still an open problem

CNN-based Density Estimation and Crowd Counting: A Survey

Guangshuai Gao^{1,2}, Junyu Gao³, *Student Member, IEEE*, Qingjie Liu^{1,2*}, *Member, IEEE*, Qi Wang³, *Senior Member, IEEE*, and Yunhong Wang^{1,2}, *Fellow, IEEE*

Challenges (training)

Adapt to multiple domains

Pose


Geometry

Environment

Training with Simulator Data

IDDA: a large-scale multi-domain dataset for autonomous driving

Emanuele Alberti^{*1}, Antonio Tavera^{*1}, Carlo Masone² and Barbara Caputo¹



<https://www.ai4europe.eu/research/ai-catalog/gta-dataset>

Citycam Dataset

CityScape Cameras
TECNICO Universidade Nova de Lisboa

HOME PUBLICATIONS DATASET


A large annotated public dataset

Download our CityCam dataset

60.000 frames with 900.000 vehicles fully annotated

60 million images for testing

check on the download menu



LICENSE AND CITYCAM RELEASE AGREEMENT

The goal of the CityCam database is to provide a versatile platform for a wide range of computer vision research topics related to traffic videos. Therefore, the CityCam database is now made available for research purpose only. CMU serves as the technical agent and reserves the ultimate interpretation right for distribution of the database.

This dataset is released under the [Creative Commons Attribution 4.0 License](#) and researcher(s) agrees to the following restrictions on the CityCam database:

1. The CityCam database is available for non-commercial research purposes only.
2. All images of the CityCam database are obtained from the Internet which are not property of CMU. CMU is not responsible for the content nor the meaning of these images.
3. You agree not to reproduce, duplicate, copy, sell, trade, resell or exploit for any commercial purposes, any portion of the images and any portion of derived data.
4. You agree not to further copy, publish or distribute any portion of the CityCam database. Except, for internal use at a single site within the same organization it is allowed to make copies of the database.
5. CMU reserves the right to terminate your access to the database at any time.
6. All submitted papers or any publicly available text using the CityCam database must cite the following paper:
Shanghang Zhang, Guanhang Wu, João P. Costeira, and José MF Moura. "Understanding Traffic Density from Large-Scale Web Camera Data." In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4264-4273.
7. The ultimate explanation of this agreement refers to CMU.

REGISTRATION AND DOWNLOAD

To download our full dataset please fill the form below. A download link will be sent to the email. By submitting this form researchers agree with our license and user agreement. All published material with CityCam data must cite article

Shanghang Zhang, Guanhang Wu, João P. Costeira, and José MF Moura. "Understanding Traffic Density from Large-Scale Web Camera Data." In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4264-4273.

citycam

To download the citycam dataset please fill the form below. A link to the dataset will be sent to your email.

[Inicie sessão no Google](#) para guardar o seu progresso. [Saiba mais](#)

*Obrigatório

Name *

A sua resposta

<http://printart.isr.tecnico.ulisboa.pt/citycam>

What do cars tell about people !

a case of proxy-sensing

Visual Census: Using Cars to Study People and Society

Timnit Gebru¹ Jonathan Krause¹ Yilun Wang¹ Duyun Chen¹ Jia Deng² Li Fei-Fei¹
¹Stanford University ²University of Michigan

RESEARCH ARTICLE | COMPUTER SCIENCES | OPEN ACCESS



Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States

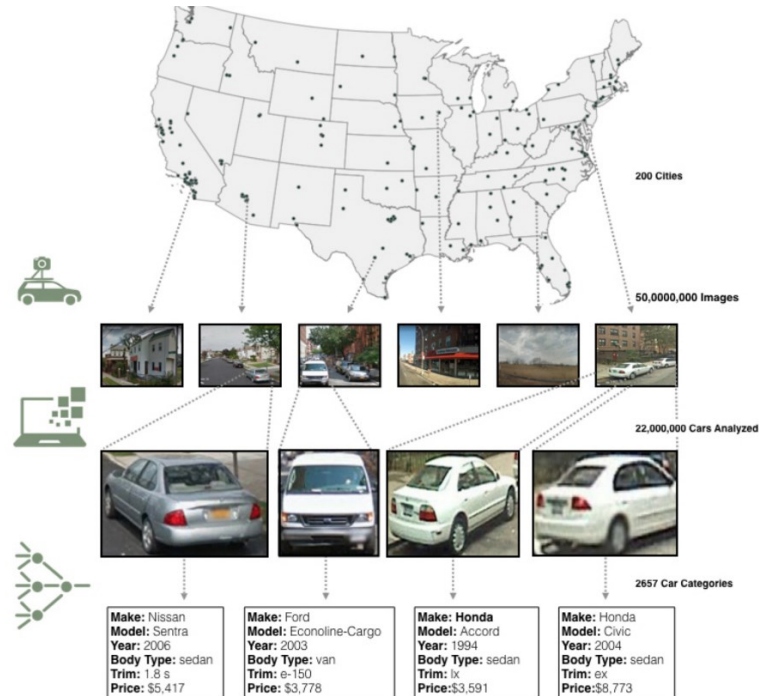
Timnit Gebru, Jonathan Krause, Yilun Wang, and Li Fei-Fei [Authors Info & Affiliations](#)

November 28, 2017 | 114 (50) | <https://doi.org/10.1073/pnas.1700035114>

200 cities

50 million images from Google Street View

US census



Source: Google Street View




What people drive ...




Google Street View


1997-2004 Dodge Dakota
extended-cab sport



2008-2010 BMW 1-Series
convertible 135i



Split in 2657 Classes



2657 Car Categories

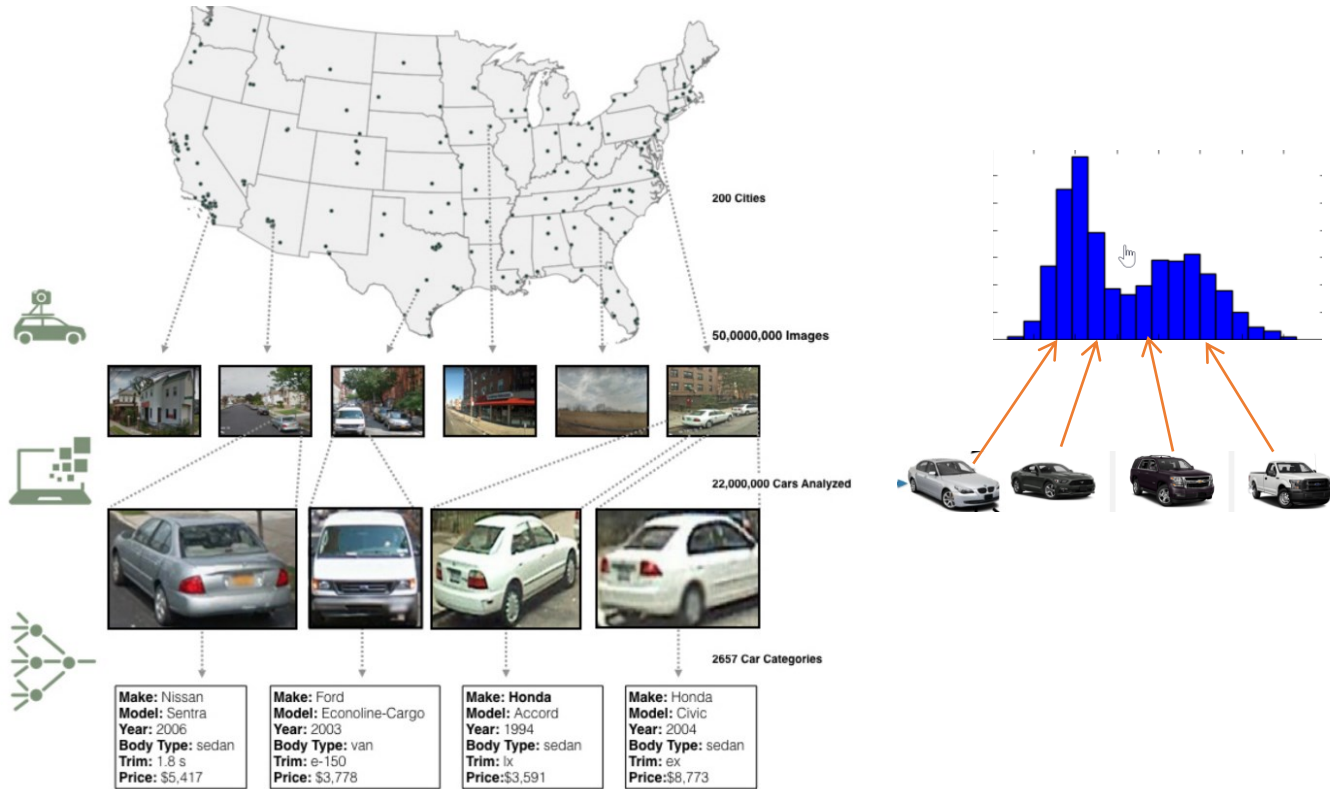
Make: Nissan
Model: Sentra
Year: 2006
Body Type: sedan
Trim: 1.8 s
Price: \$5,417

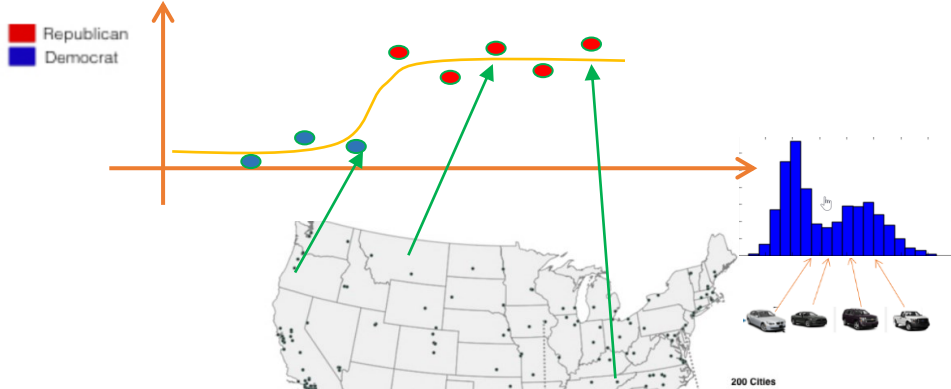
Make: Ford
Model: Econoline-Cargo
Year: 2003
Body Type: van
Trim: e-150
Price: \$3,778

Make: Honda
Model: Accord
Year: 1994
Body Type: sedan
Trim: lx
Price: \$3,591

Make: Honda
Model: Civic
Year: 2004
Body Type: sedan
Trim: ex
Price: \$8,773


Data representation


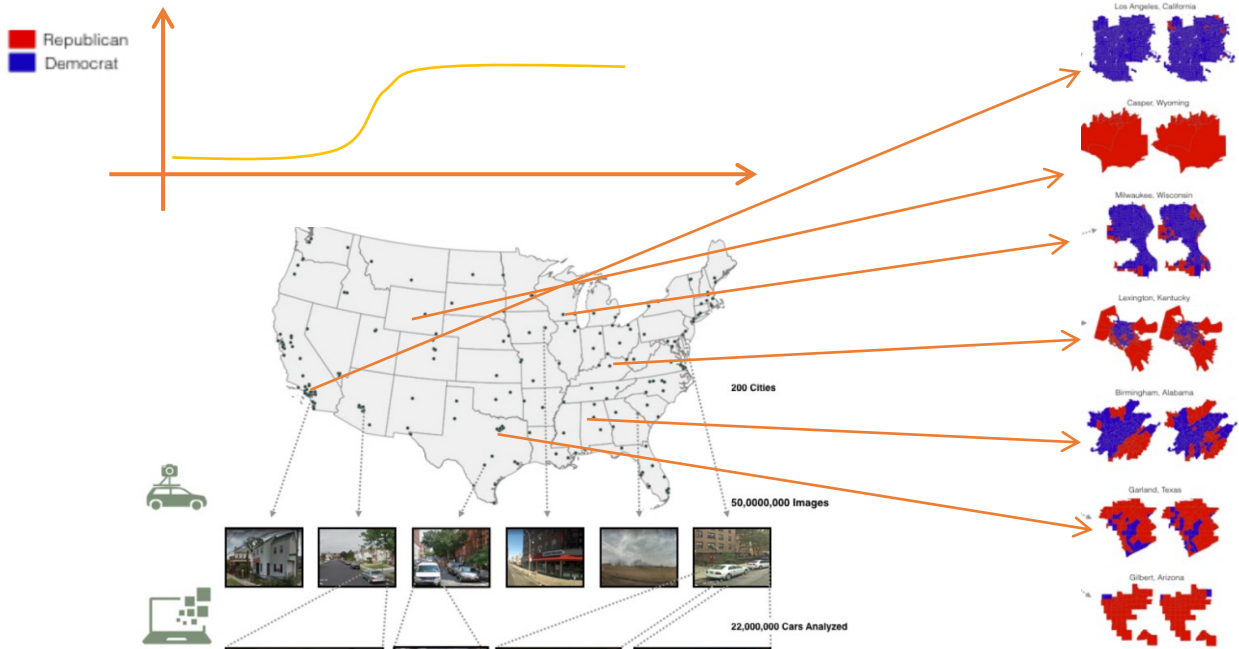




2657 Car Categories

Make: Nissan Model: Sentra Year: 2006 Body Type: sedan Trim: 1.8 s Price: \$5,417	Make: Ford Model: Econoline-Cargo Year: 2003 Body Type: van Trim: e-150 Price: \$3,778	Make: Honda Model: Accord Year: 1994 Body Type: sedan Trim: lx Price: \$3,591	Make: Honda Model: Civic Year: 2004 Body Type: sedan Trim: ex Price: \$8,773
--	---	--	---

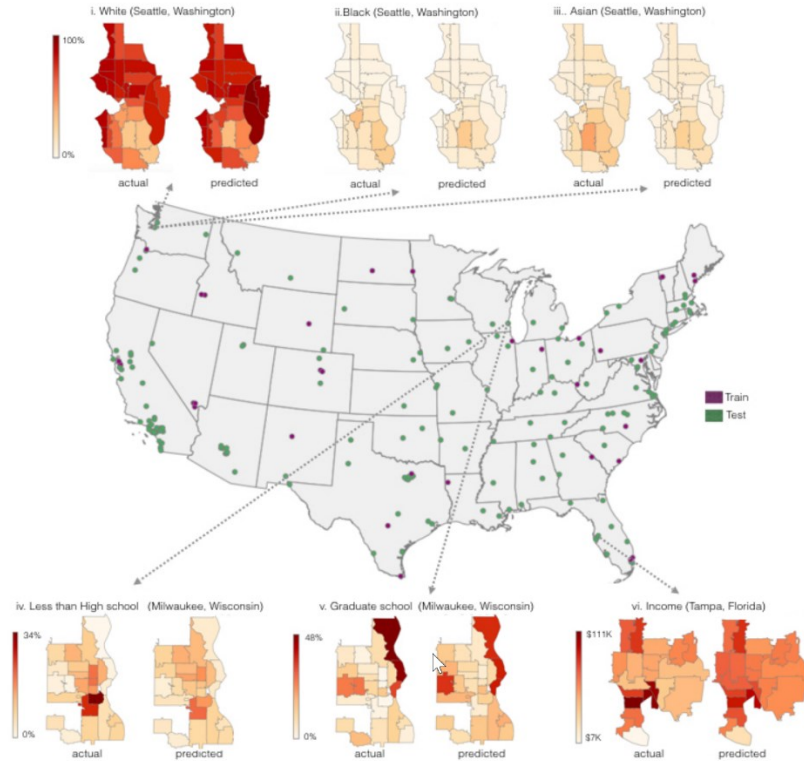




2657 Car Categories

Make: Nissan Model: Sentra Year: 2006 Body Type: sedan Trim: 1.8 s Price: \$5,417	Make: Ford Model: Econoline-Cargo Year: 2003 Body Type: van Trim: e-150 Price: \$3,778	Make: Honda Model: Accord Year: 1994 Body Type: sedan Trim: lx Price: \$3,591	Make: Honda Model: Civic Year: 2004 Body Type: sedan Trim: ex Price: \$8,773
--	---	--	---

Humans are “low rank” ?



Our model detects strong associations between vehicle distribution and disparate socioeconomic trends. For instance, several studies have shown that people of Asian descent are more likely to drive Asian cars ^[43], a result we observe here as well: the two brands that most strongly indicate an Asian neighborhood are Hondas and Toyotas. Cars manufactured by Chrysler, Buick and Oldsmobile are positively associated with African American neighborhoods, which is again consistent with existing research ^[44]. And vehicles like pickup trucks, Volkswagens and Aston Martins are indicative of mostly Caucasian neighborhoods. See Fig. [S2](#).

In some cases, the resulting associations can be easily applied in practice. For example, the vehicular feature that was most strongly associated with Democratic precincts was sedans, whereas Republican precincts were most strongly associated with extended-cab pickup trucks (a truck with rear-seat access). We found that by driving through a city for 15 minutes while counting sedans and pickup trucks, it is possible to reliably determine whether the city voted Democratic or Republican: if there are more sedans, it probably voted Democrat (88% chance) and if there are more pickup trucks, it probably voted Republican (82% chance). See Fig. [3\(a\)iii](#).

Let's put people in the equation



Photo by [Pavel Danilyuk](#) from [Pexels](#)




LARSyS
Laboratory of Robotics
and Engineering Systems



Institute for Systems
and Robotics | LISBOA

Bike Monitor

www.isr.ist.utl.pt/~manuel/smartbike



STRESS MAP

Find which streets you were
stressed

Our first encounter with “transportation”



João Paulo Cardoso




LABORATÓRIO NACIONAL
DE ENGENHARIA CIVIL



Carlos L. Azevedo



Filipe Moura



Why are there so few bikes ?



João Paulo Cardoso



LABORATÓRIO NACIONAL
DE ENGENHARIA CIVIL

Conflicts

- Cyclists vs Cars
- Cyclists vs Pedestrians



The textbook solution ?

Conflicts

- Cyclists vs Cars
- Cyclists vs Pedestrians

Surveys

- Imprecise
- Subjective
- People forget
- Hard to process



Where it all started

Smartphone App(2015)

Manuel Marques



Pedro Vieira



João P. Costeira



Susana Brandão



[SMARTcycling: Assessing cyclists' driving experience](#)

P Vieira, JP Costeira, S Brandão, M Marques

2016 IEEE Intelligent Vehicles Symposium (IV), 2016

Bike Monitor

- GPS Coordinates
- Acceleration on 3 axis (Accelerometer)
- Rotation (Gyroscope)
- Image (Camera)
- Date
- Heartbeat



Data I - Images



Spatial context from video

- Left/Frontal/Right obstacles
- Vehicles overtaking
- Pedestrians
- Road condition




Data II - accelerometers

- Ground irregularities (Pavement cracks, bumps etc..)





Who needs accelerometers if you have a camera ???

- Ground irregularities (Pavement cracks, bumps etc..)



Data III- Heart Rate Measurements

Heart monitoring chest strap:



Everything georeferenced






Image Descriptor




Image descriptor






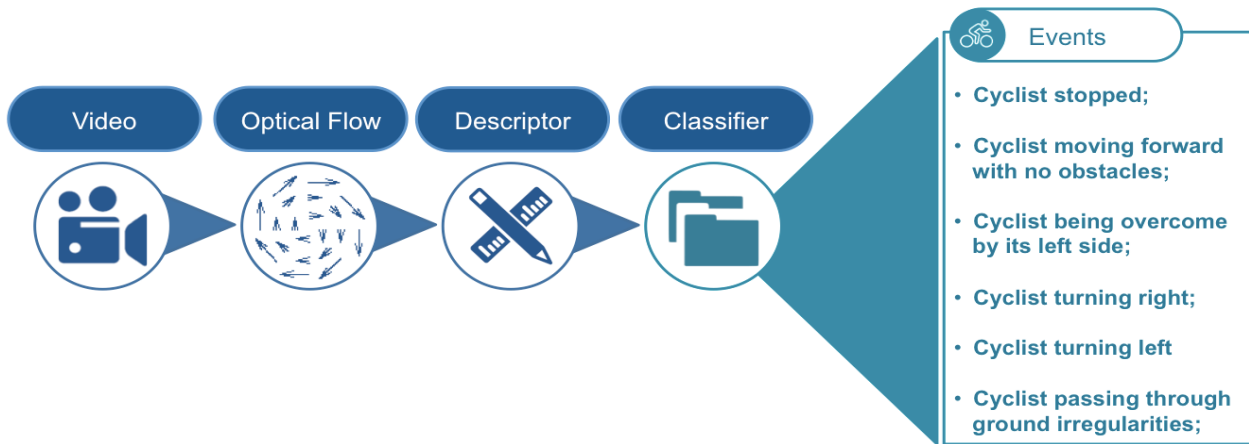


Image descriptor


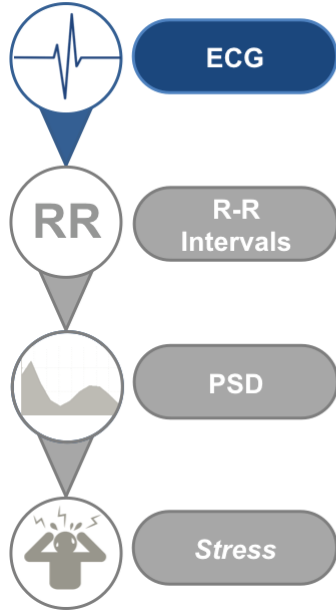


The whole visual pipeline


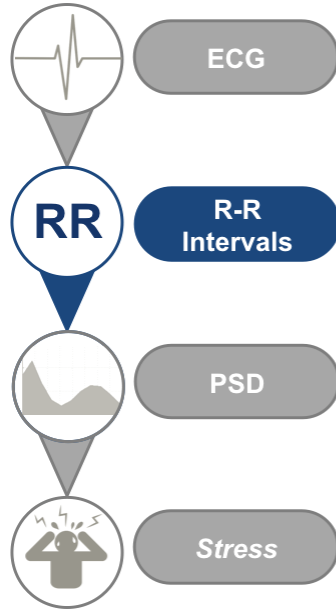
Visual Systems
© 2014 MIT



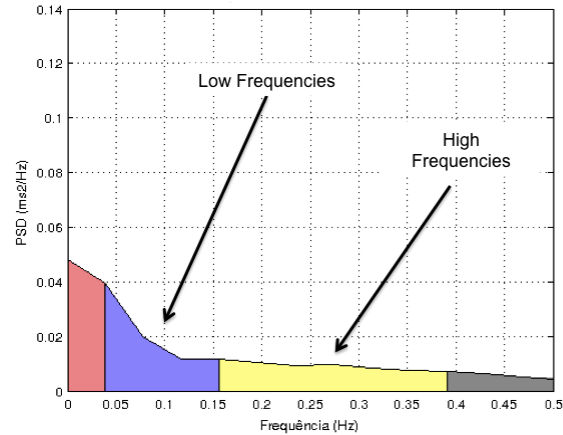
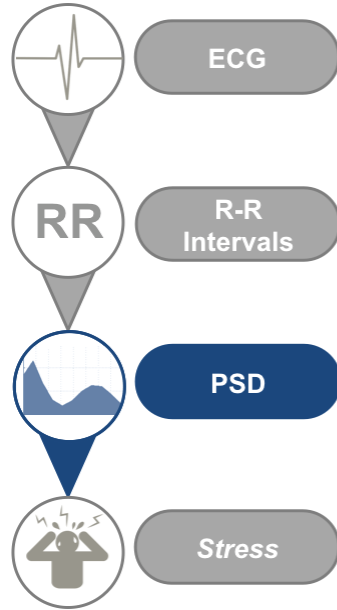
Evaluating “Stress”(HRV)



Evaluating “Stress”



Evaluating “Stress”




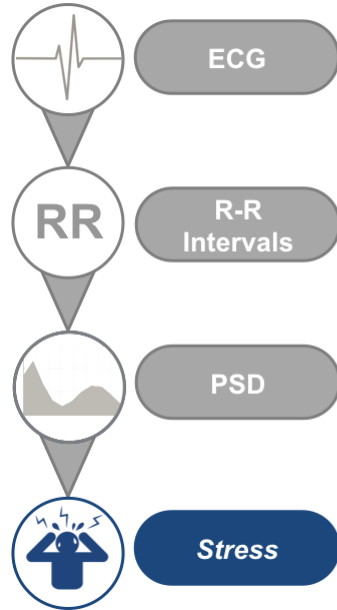
0.04 – 0.15 Hz

Sympathetic Stimulation (*Stress*)

0.15 – 0.4 Hz

Parasympathetic Stimulation

Measuring “Stress”




0.04 – 0.15 Hz

Sympathetic Stimulation (Stress)


0.15 – 0.4 Hz

Parasympathetic Stimulation







Not that stressful

Distribuição geográfica do Ratio LF/HF



Risky business





It's not your heart that tricks you...



- Assumptions without data are usually wrong (adoption)
- First person vision is crucial
- Doing good surveys is really difficult



Conscious choices !



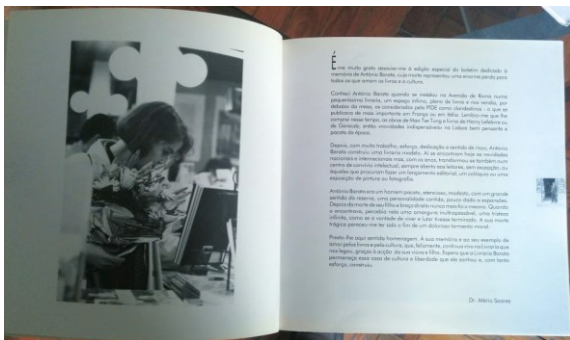
- One of Lisbon's historical bookshops
- located in the new expanding areas of the 40's
- Since the 50's became a central "upper-middle class?/affluent/ residential and shopping area.
- *Bloody* expensive these days 😊

A special Bookshop with a special history

Founded by António Barata in 1957 under the dictatorship



- originally a very small space selling magazines, newspapers and stationary
- “hidden in the back” was one of Lisbon’s largest supply of *clandestine books*



... bought “under the counter” the book of Mao-Tse-Tung, Henry Lefebvre and Garaudy...

Livraria Barata awarded “Medal for Municipal Merit”

Essays on Consumer Switching and Search Behavior

Qiwei Han

PhD in Engineering and Public Policy March, 2017

João Paulo Costeira, IST Pedro Ferreira, CMU EPP/Heinz (advisors)



Carnegie
Mellon
University

The Effect of Product Placement on Shopping Behavior at the Point of Purchase: Evidence from a Randomized Experiment Using Video Tracking in a Physical Bookstore,

Qiwei Han, Pedro Ferreira, João Paulo Costeira,

Conference on Digital Experimentation (CODE), MIT Sloan School of Management, Cambridge MA, October, 2016

The Effect of Product Placement on Shopping Behavior at the Point of Purchase: Evidence From Randomized Experiment Using Video Tracking in a Physical Bookstore (November 21, 2018). Available at Social Science Research Network: <https://ssrn.com/abstract=3288604>

The experimental site

Technical
Books(stable titles)



Main floor of the bookshop

- 29 tables with approx. 30 books each
- Tables display “recent books/editions”
- 100 thousand+ titles in shop
- 1000 customers/day
- Open 7 days/week (9AM-11PM)

“storage”

Long term storage in other floors

Table
Recent books

Other sections



Pocket editions (foreign)


Magazines/newspapers

Depth Camera 1


Counting incoming/outgoing costumers



Typical shopping day ...BC (Before Corona!)



Daily Traffic Volume




Camera 2: New releases !




RGB image



Depth image





Monitor the area close to the table of new books using 3D information







Detect and follow shopper in real-time using 3D representation of the space

3D View of the Scene




Count Pick and Take



Randomized Experiment on Book Placement Recommendations

- We **randomize book placement** on the table across 5 rows and 6 columns
- At cycle 0 we randomly shuffle the books already on the table. Then, each cycle we generate a random sequence of numbers between 1 and 30 and place the incoming books accordingly



Treated group: books in **edge positions** (first row and left/right most columns)

Control group: books in the **center positions**

Pick Rate by Positions (sample)

6	7.66	0.43	0.43	1.06	2.98	4.26
5	2.55	3.19	2.98	0.85	1.91	4.26
4	2.55	0.85	1.06	1.28	1.06	3.61
3	1.70	1.06	1.28	1.49	1.06	1.06
2	2.77	2.13	1.06	2.98	2.13	1.70
1	5.96	3.40	2.55	6.81	3.83	5.53
Left/Front	1	2	3	4	5	6

Pick Rate > 4

Pick Rate < 2

Empirical Results

- Books placed at the edge of the table are **picked 102%** more often and **taken 77%** more often per day than those at the center of the table
- However, conditionally on picking the book, shoppers are **equally likely** to take books placed at the edge and at the center of the table
- Book placement recommendations affects consumer choice **through its effect on the search process** but not through its effect on the consideration process
- The bookstore manager may maximize profit by placing **books with higher margins at the edge of the table**

Conclusions/policy

We show that **books placed at the edge of the table are more likely to be picked and taken** than those placed at the center of the table.

we also show **that conditional on being picked, shoppers are equally likely to take books** placed at the edge and at the center of the table.

Moreover, we demonstrate that the book placement on the table essentially resembles the display recommendations that online bookstores are heavily in use today.

So:


bookstore manager has incentives to **place books with higher margin at the edge of the table for longer time**, as she knows that this may incur higher sales.

To this end, **consumer welfare is likely to decrease** because consumers tend to purchase books that do not give them the highest utility, as a result of search cost obfuscation.

Unless you shop at Livraria Barata!



Coffee chat ...



About Human Decision-Making

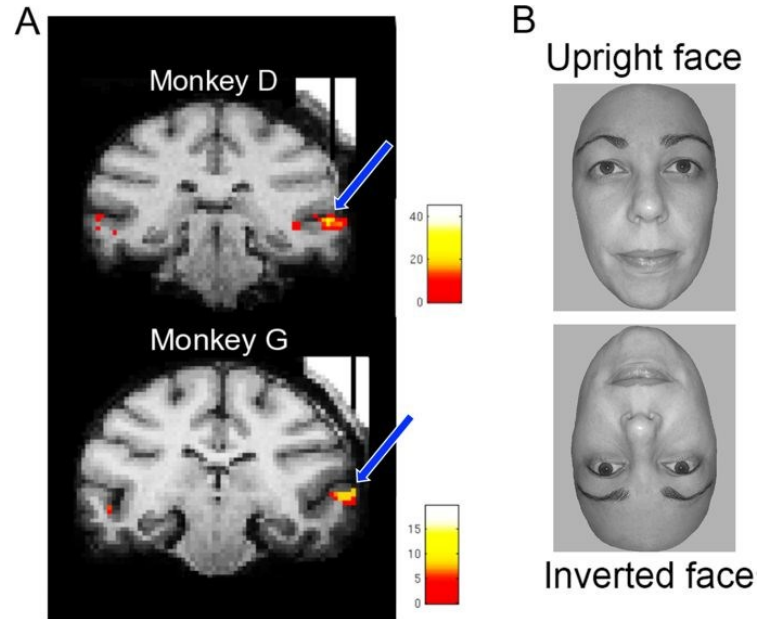
- Rationality is overrated : we decide then find a rational explanation ?
- Information is vital : made me a better consumer (at least in books)
- Cognitive overload (choice)
- Limitations to cognition (hardware!)

The Thatcher effect!



Different machinery for different tasks

accuracy depends on pose



The impact of stimulus size and orientation on individual face coding in monkey face-selective cortex
• [Jessica Taubert](#), [Goedele Van Belle](#), [Rufin Vogels](#) [Bruno Rossion](#), Nature, [Scientific Reports](#) volume 8,
Article number: 10339 (2018)

Can be fair and still increase revenue!




In summary



A.I./REAL LIFE

‘No-Code’ Brings the Power of A.I. to the Masses

A growing number of new products allow anyone to apply artificial intelligence without having to write a line of computer code. Proponents believe the “no-code” movement will change the world.

 Give this article



```
N_Hl7 VOucl>7]MxMTbi%d r9=}Rp,0{.fOB D |
P&/;a_2u /HA4=,^' P Cgz(_SpU?]4 C{Z
Sccc(. Ip>M[DT i 7Pe P[| "L !yQ[~~o c
nG.U'c}[mv<,D`m~ C`xb |1<E TT@%L }u60\v
we [L+x{D+<)Vsm#gN`zPvM5 oz / A,9 x'b,H{
d1Gt7G 9 +g;i6) P( \? jd:jjPoqg<d~@4l6
"hG Je P l1)duKd0 |M7 lYqheSyt"AR4>+
[Gf.vw3 zc R?8Y % +t ?5ad$>d/9i0{HO
K zqDOz'4! hbRt ,* )tz +S Ua#>a [$
oj w4 <no code> 5 pID6( Ou02&wG
fPkWIPDm] [$nQ .<F /p_QfK-qUj$S
```



BUILDING THE EUROPEAN AI ON-DEMAND PLATFORM

João Paulo Costeira

Signal & Image Processing Group

Instituto Sistemas e Robótica

jpcosteira@e-technico.ulisboa.pt

AI4EU National Contact Point (NCP)

co-leader of Task 7.5 Physical AI

www.ai4eu.eu




“ Desideratum: a European AI on-demand platform

- **THE Central access point: a shared resource for Europe**

- Integrate tools and resources
- Offer solutions and support to all users of AI to integrate AI into applications, products and services
- Mobilisation of the community
- Catch phrase: *“democratise AI”*

- **2018: Call ICT-26 (Apr 2018)**

- Large project: 20M€
- Cascade funds: 3M€
- 80 partners, 21 Countries



<http://ai4europe.eu>



[AI Community](#)

[Business & Industry](#)

[Research](#)

[Education](#)

[Ethics](#)

[Services](#)

[News & Events](#)



The European AI on Demand Platform

AI4EU is a one-stop-shop for anyone looking for AI knowledge, technology, tools, services and experts.

<http://aiexp.ai4europe.eu>

<https://github.com/ai4eu/tutorials>

AI4EU HOME MARKETPLACE ON-BOARDING MODEL DOCUMENTATION Q SIGN IN | SIGN UP NOW ?

AI4EU Experiments June Release

The June Release greatly streamlines the **onboarding** process and adds support for **shared folders** and the protobuf **stream** keyword. It also features the **training pipeline** example.

VIDEO ANALYTICS THREAT PREDICTION NETWORK SECURITY ANOMALY DETECTION

CLASSIFICATION REGRESSION CLUSTERING AGGREGATOR

Explore the AI4EU Experiments Marketplace

it's easy to **discover, download & deploy**

- AI4Industry Pilot Planning Component ASP+OWL
- TensorBoard1, NewsClassifier1, NewsTrainer1
- Hardware components (lights and disk)
- TensorFlow logo

AI Pipelines: <https://aiexp.ai4europe.eu> The Marketplace



The screenshot displays the AI4EU Marketplace interface. At the top, there is a navigation bar with the AI4EU logo and user profile information. Below this, a sidebar on the left contains navigation options: HOME, MARKETPLACE, MY MODELS, ON-BOARDING MODEL, DESIGN STUDIO BETA, Q AND A, and INTERFACES. The main content area is titled 'Marketplace' and shows a grid of AI models. The grid is filtered by 'All Catalogs' and 'My Favorite Catalogs'. The models listed include:

- pddl-planners-ffi** by Felix Ingrand (05/31/2021) | New. Category: Classification. Tags: Training, planning, deep learning.
- SentimentAnalysisModel** by Sajid (05/27/2021) | New. Category: Sentiment Analysis. Tags: industrial ai, knowledge representation, owl, Utility Node, integrative ai, manufacturing, GUI, answer set programming.
- LionForests** by Ioannis Mollas (05/27/2021) | New. Category: Binary Classification.
- Text2ImageSearch** by Nicola Messina (05/26/2021) | New. Category: cross modal.
- Sudoku Tutorial** by Peter Schueller (05/10/2021) | New. Category: hello world.
- Answer Set Solver: Clingo "One Shot" Interface** by Roland Kaminski (05/10/2021) | New. Category: answer set programming.
- AI4EU Sudoku Tutorial Component: Design Evaluator** by Peter Schueller, TU Wien.
- sudoku-tutorial-design-evalu...** by Peter Schueller (05/06/2021) | New. Category: hello world.
- sudoku-tutorial-gui** by Peter Schueller (05/06/2021) | New. Category: GUI.

The interface also features a search bar, a 'Filter By Category' section with checkboxes for Classification, Data Sources, Data Transformer, Prediction, and Regression, and a 'Tags' section with various categories like Training, planning, deep learning, industrial ai, knowledge representation, owl, Utility Node, integrative ai, manufacturing, GUI, answer set programming, Binary Classification, reasoning, pddl, Random Forests, rules, Explainable AI, retrieval, and similarity.

AI4EU AI Platform: The Design Studio

The screenshot displays the AI4EU AI Platform Design Studio interface. The top navigation bar includes the AI4EU logo, a search icon, a notification icon, a user profile icon labeled "PhysicalAI", and a help icon. The left sidebar contains navigation options: HOME, MARKETPLACE, MY MODELS, ON-BOARDING MODEL, DESIGN STUDIO BETA, Q AND A, and INTERFACES. The main workspace is titled "Acu-Compose BETA" and shows a breadcrumb path: Home / Design Studio / Acu-Compose. The workspace contains a workflow diagram with two nodes: "AudioSegmentation1" and "AudioSpeechToTextEnglish1", connected by a green arrow. The right sidebar shows the "Properties" panel for the selected node, with fields for "Link Name" (empty), "Source Node Name" (AudioSegmentation1), and "Target Node Name" (AudioSpeechToTextEnglish1). The bottom right corner shows a "My Solutions" section with a search bar. The top right of the workspace has a toolbar with buttons for "Probe", "Delete", "Clear", "Validate", "Save", and "Deploy".

Acu-Compose BETA
Home / Design Studio / Acu-Compose

Marketplace

Solutions Models

Classification

- AI-Panel (1.0.0)
- AudioPunctuationEnglish (1.0.0)
- AudioPunctuationGerman (1.0.0)
- AudioSegmentation (1.0.0)
- AudioSpeakerRecognition (1.0.0)
- AudioSpeechToTextEnglish (1.0.0)
- AudioSpeechToTextGerman (1.0.0)

Data Transform Tools

- ai4eu-icd-10 (1.0.0)
- ai4industry-gui (1.0.0)
- ai4industry-planner (1.0.0)
- ai4industry-skillmatcher-dummy (1.0.0)
- AIA-obfuscation (1.0.0)

Data Sources

- AudioFileBroker (1.0.0)
- databroker-houseprice (1.0.0)
- edm-env (1.0.1)
- edm-env (1.0.3)
- edm-env (1.0.0)

Untitled x New

Probe: Delete Clear Validate Save Deploy

AudioSegmentation1

AudioSpeechToTextEnglish1

Properties Matching Models

Link Name:
empty

Source Node Name:
AudioSegmentation1

Target Node Name:
AudioSpeechToTextEnglish1

My Solutions

AI4EU AI Platform: MarketPlace + Design Studio

The screenshot displays the AI4EU AI Platform interface, specifically the Design Studio. The top navigation bar includes the AI4EU logo, a search icon, a notification icon, a user profile icon labeled "PhysicalAI", and a help icon. The left sidebar contains a navigation menu with the following items: HOME, MARKETPLACE, MY MODELS, ON-BOARDING MODEL, DESIGN STUDIO BETA, Q AND A, and INTERFACES. The main workspace is titled "Acu-Compose BETA" and shows a breadcrumb path: Home / Design Studio / Acu-Compose. Below the breadcrumb, there are tabs for "Solutions" and "Models". The "Models" tab is active, displaying a list of models under three categories: "Classification", "Data Transform Tools", and "Data Sources". The "Classification" category includes: AI-Panel (1.0.0), AudioPunctuationEnglish (1.0.0), AudioPunctuationGerman (1.0.0), AudioSegmentation (1.0.0), AudioSpeakerRecognition (1.0.0), AudioSpeechToTextEnglish (1.0.0), and AudioSpeechToTextGerman (1.0.0). The "Data Transform Tools" category includes: ai4eu-icd-10 (1.0.0), ai4industry-gui (1.0.0), ai4industry-planner (1.0.0), ai4industry-skillmatcher-dummy (1.0.0), and AIA-obfuscation (1.0.0). The "Data Sources" category includes: AudioFileBroker (1.0.0), databroker-houseprice (1.0.0), edm-env (1.0.1), edm-env (1.0.3), and edm-env (1.0.0). The main workspace area shows a workflow diagram on a grid background. It features two nodes: "AudioSegmentation1" and "AudioSpeechToTextEnglish1". Both nodes are connected to a central data source icon (a blue box with a ship) via blue arrows pointing upwards. The "AudioSegmentation1" node is also connected to the "AudioSpeechToTextEnglish1" node via a green arrow pointing right. The top of the workspace has a toolbar with buttons for "Probe", "Delete", "Clear", "Validate", "Save", and "Deploy". The right sidebar contains two tabs: "Properties" and "Matching Models". The "Properties" tab is active, showing fields for "Link Name" (empty), "Source Node Name" (AudioSegmentation1), and "Target Node Name" (AudioSpeechToTextEnglish1). Below the tabs is a search bar for "My Solutions".

AI4EU AI Platform: MarketPlace + Design Studio

The screenshot displays the AI4EU AI Platform interface, specifically the Design Studio for the 'Acu-Compose BETA' project. The interface is divided into several sections:

- Left Sidebar:** Contains navigation options: HOME, MARKETPLACE, MY MODELS, ON-BOARDING MODEL, DESIGN STUDIO BETA, Q AND A, and INTERFACES.
- Marketplace Panel:** Lists various AI models and tools under categories like 'Classification', 'Data Transform Tools', and 'Data Sources'.
 - Classification:** AI-Panel (1.0.0), AudioPunctuationEnglish (1.0.0), AudioPunctuationGerman (1.0.0), AudioSegmentation (1.0.0), AudioSpeakerRecognition (1.0.0), AudioSpeechToTextEnglish (1.0.0), AudioSpeechToTextGerman (1.0.0).
 - Data Transform Tools:** ai4eu-icd-10 (1.0.0), ai4industry-gui (1.0.0), ai4industry-planner (1.0.0), ai4industry-skillmatcher-dummy (1.0.0), AIA-obfuscation (1.0.0).
 - Data Sources:** AudioFileBroker (1.0.0), databroker-houseprice (1.0.0), edm-env (1.0.1), edm-env (1.0.3), edm-env (1.0.0).
- Design Studio Canvas:** Shows a workflow diagram with two nodes: 'AudioSegmentation1' and 'AudioSpeechToTextEnglish1'. Below the nodes, there are illustrations of a person at a laptop, a server rack, and a person at a laptop. A mathematical optimization problem is displayed at the bottom right of the canvas:
$$\begin{aligned} & \underset{x}{\text{minimize}} && \sum_{n=1}^N (\epsilon - b_n^T x)_+ \\ & \text{subject to} && \mathbf{1}^T x = 0, \end{aligned}$$
- Right Panel:** Contains 'Properties' and 'Matching Models' tabs. The 'Properties' tab shows:
 - Link Name: empty
 - Source Node Name: AudioSegmentation1
 - Target Node Name: AudioSpeechToTextEnglish1
- Bottom Right Panel:** Labeled 'My Solutions', it contains a search bar.

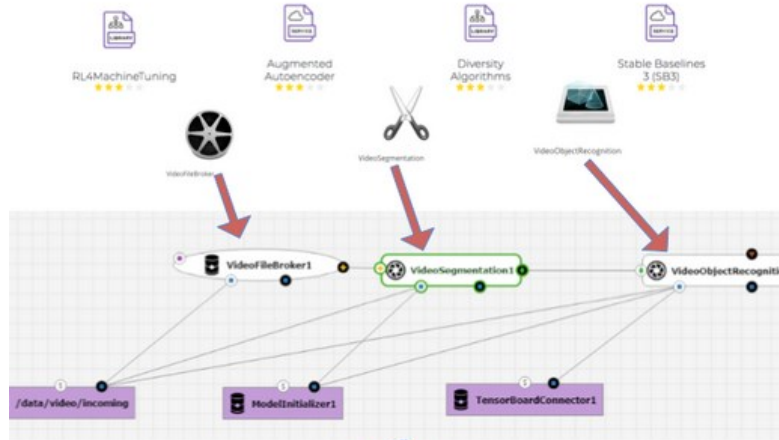
CATALOG

Highlights

edu2comAPI
edu2com is a team formation algorithm for allocating teams of students to internship programs
[Show more](#)

LORE - Local Rule-based Explanations
LORE (Local Rule-based Explanations) is a model-agnostic explainer for tabular data
[Show more](#)

IoT data analysis model
Analysis of IoT data in the logistic industry
[Show more](#)



Almost there !

Orchestrate & Deploy (distributed)

AI is a commodity ... and big business !

The Marketplace for AI models

MODELS

All models / C4I compatible

PointRend
Segment objects from the rest of the scene.
segmentation @ pytorch @ Chyngai @ Im

(Wild) Cotton Classifier
A model for differentiating cotton and velvet leaf.
classification @ sensorhub @ Inreal Clear >

Sign Language
Identification of American Sign Language Letters.
detection @ amannine @ American Sign L >

CenterNet
Identify objects and predict their bounding boxes in real time.
detection @ pytorch @ MICOCC @ Image

NAS-FPN
NAS-FPN is an object detection model.
detection @ amannine @ MICOCC @ Im

Iris
Detect and track 3D eye, eyebrows and its geometry from monocular video.
landmark.detection @ openvino @ Inreal D >

Face Mesh
Detect and provide detailed face mesh estimation from a single camera.
mesh.detection @ openvino @ Inreal Clear >

Blaze Face
Detect faces and track the ears, eyes, nose and mouth positions.
landmark.detection @ openvino @ Inreal D >

Text Spotting
Detect and recognize text without a dictionary.
see.detection @ openvino @ CTRW50 @ Im

Palm Detector
Detect and track 2D palm position in real time.
landmark.detection @ openvino @ HIRland >

Tiny YOLO v4
Detect many items in one image with high accuracy and speed.
detection @ openvino @ MICOCC @ Image

Lightweight OpenPose
Generate a 2D point body pose for multiple human subjects at once.
pose.estimation @ openvino @ MICOCC @ I >

DeepLabv3-
Segment people from the rest of the scene.
segmentation @ openvino @ Supervisey @ I >

DBFace
Detect faces in real time.
landmark.detection @ openvino @ WLVW @ I >

EAST
Detect text in arbitrary orientations.
see.detection @ openvino @ CTRW50 @ Im >

YOLO v3
Detect many objects in one-shot.
detection @ openvino @ MICOCC @ Image

aws

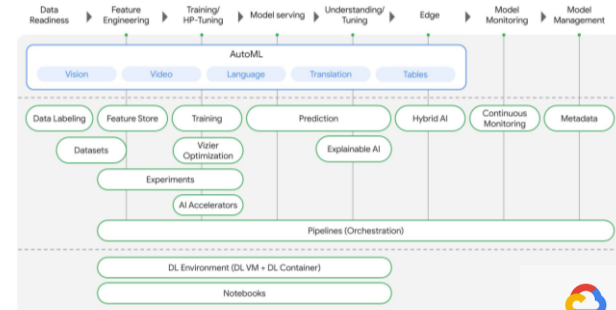
Machine Learning for Free on AWS

Free offers and services you need to build, deploy, and run machine learning applications in the cloud

[Learn more about AWS Free Tier](#)

Create a Free Account

What's included in Vertex AI?



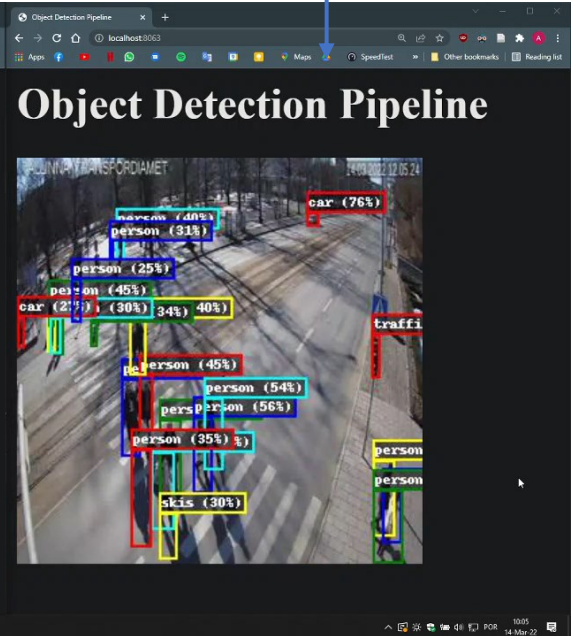
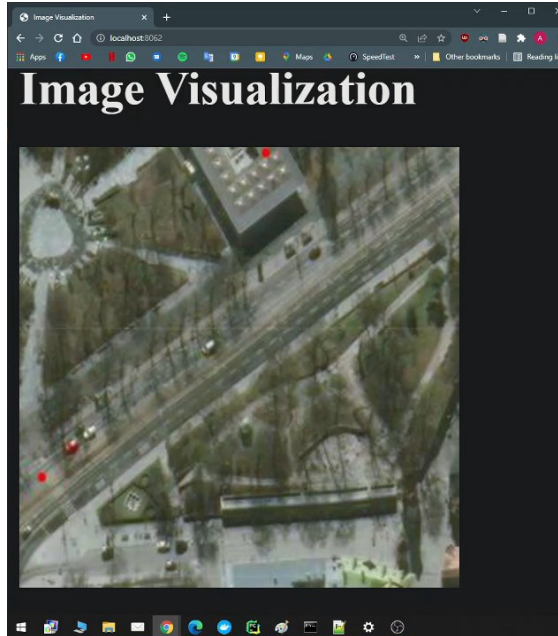
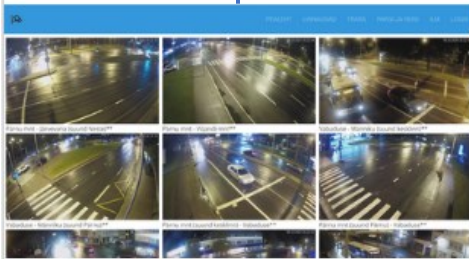
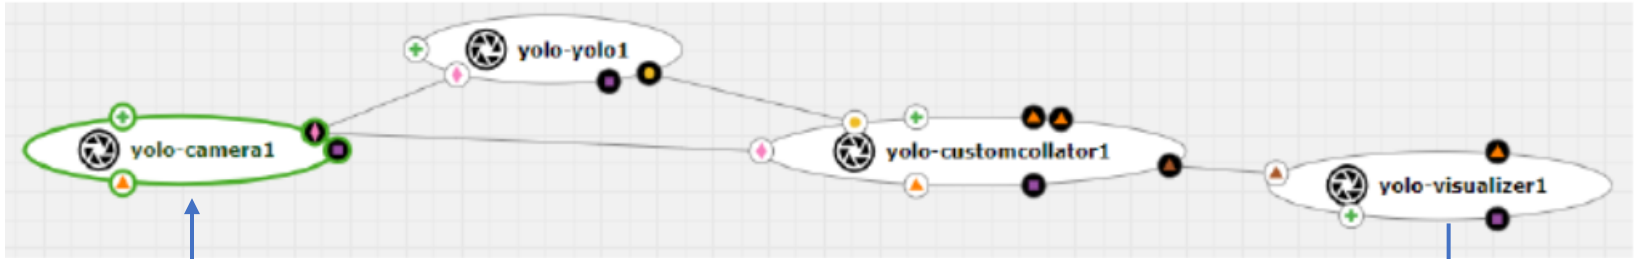
The more inefficient the market is ...the better

“two weeks old”!

The screenshot shows the arXiv website interface. At the top left is the Cornell University logo. The main navigation bar is red with the arXiv logo and the path 'cs > arXiv:2202.12566'. A search bar is on the right. Below the navigation bar, the page title 'Composing Complex and Hybrid AI Solutions' is displayed, along with the authors' names: Peter Schüller, João Paolo Costeira, James Crowley, Jasmin Grosinger, Félix Ingrand, Uwe Köckemann, Alessandro Saffiotti, and Martin Welsch. The submission date is noted as 'Submitted on 25 Feb 2022'. The abstract text begins with 'Progress in several areas of computer science has been enabled by comfortable and efficient means of experimentation, clear interfaces, and interchangeable components...'. On the right side, there is a 'Download:' section with links for 'PDF' and 'Other formats', a Creative Commons license icon (CC BY), and a 'Current browse context:' section showing 'cs.AI' and navigation links like '< prev | next >', 'new | recent | 2202', and 'Change to browse by: cs'.

6.3 Real time object detection in networked cameras

Developing a system for Urban Analytics in 10 minutes using the AI4EU Experiments Platform.



It will have impact

- The way we teach!
- With whom we collaborate !
- Less conflicts with NDA's and “borrowing/sharing data”
- Collaboration with industry
- Focus on “your thing” ...

“ How you can join, use and contribute to the platform

1. Publish your contents

- reviewed by an open editorial team



Alessandro Saffiotti

2. Contribute to further develop the portal

- reviewed by an inclusive technical board



Ioannis Mollas

3. Create and run “AI experiments”

- and put your assets in the Experiments Marketplace



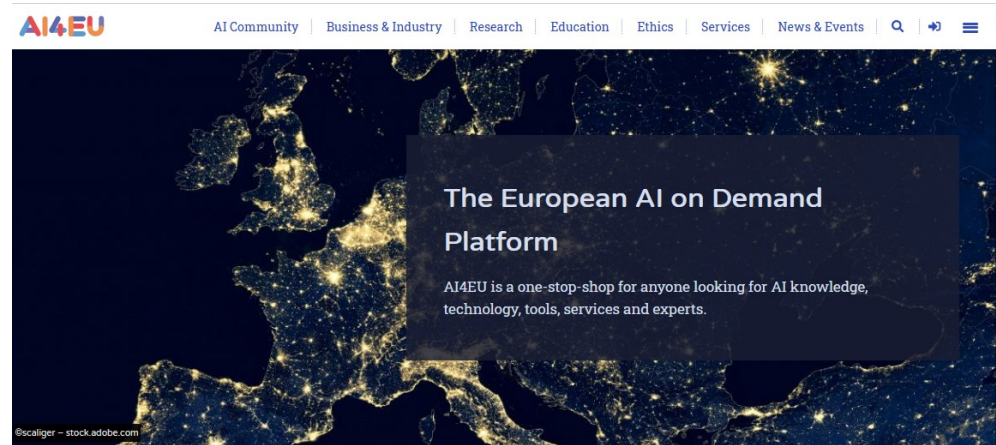
Martin Welss



Peter Schüller

Conclusions

- A paradigm shift in ICT's is underway
- It is critical that we appropriate the technology – DO IT and JOIN IT
- Unskilled “consumers” will be empowered
- The AI technology of highly skilled producers will “reign”



<http://ai4europe.eu>

Thank you !