



Visual Layer

From Mosquitoes to Trash Bins: Unbelievable Computer Vision Use Cases

Danny Bickson, Visual Layer

Joint Work



Dr. Danny Bickson
Co-founder & CEO

Co-founder and VP EMEA of Turi (acquired by Apple). CMU Researcher. Sr. Mgr at Apple.



Dr. Amir Alush
Co-founder & CTO

Co-founder and CTO of Brodmann17. Highly experienced in building CV/AI Groups and leading into production.

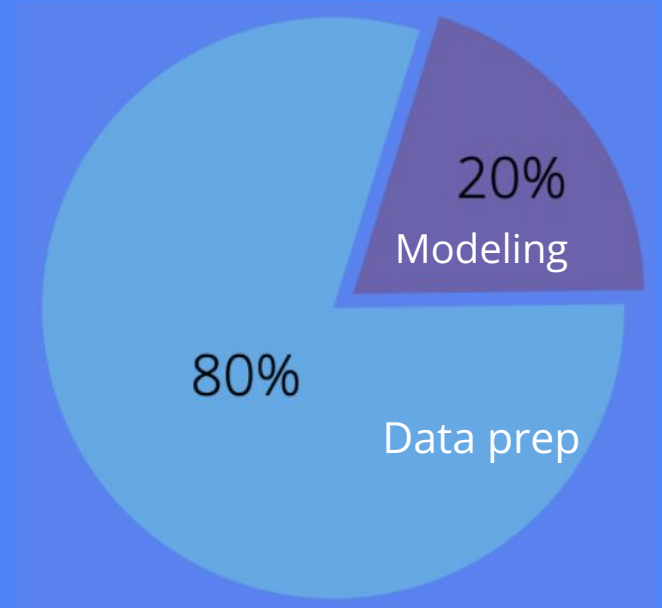



Prof. Carlos Guestrin
Co-founder & CSO

Stanford Prof, Co-founder and CEO of Turi (acquired by Apple). Sr. Dir. at Apple. Deep Learning Infra Team Pioneer.



My Personal Experience

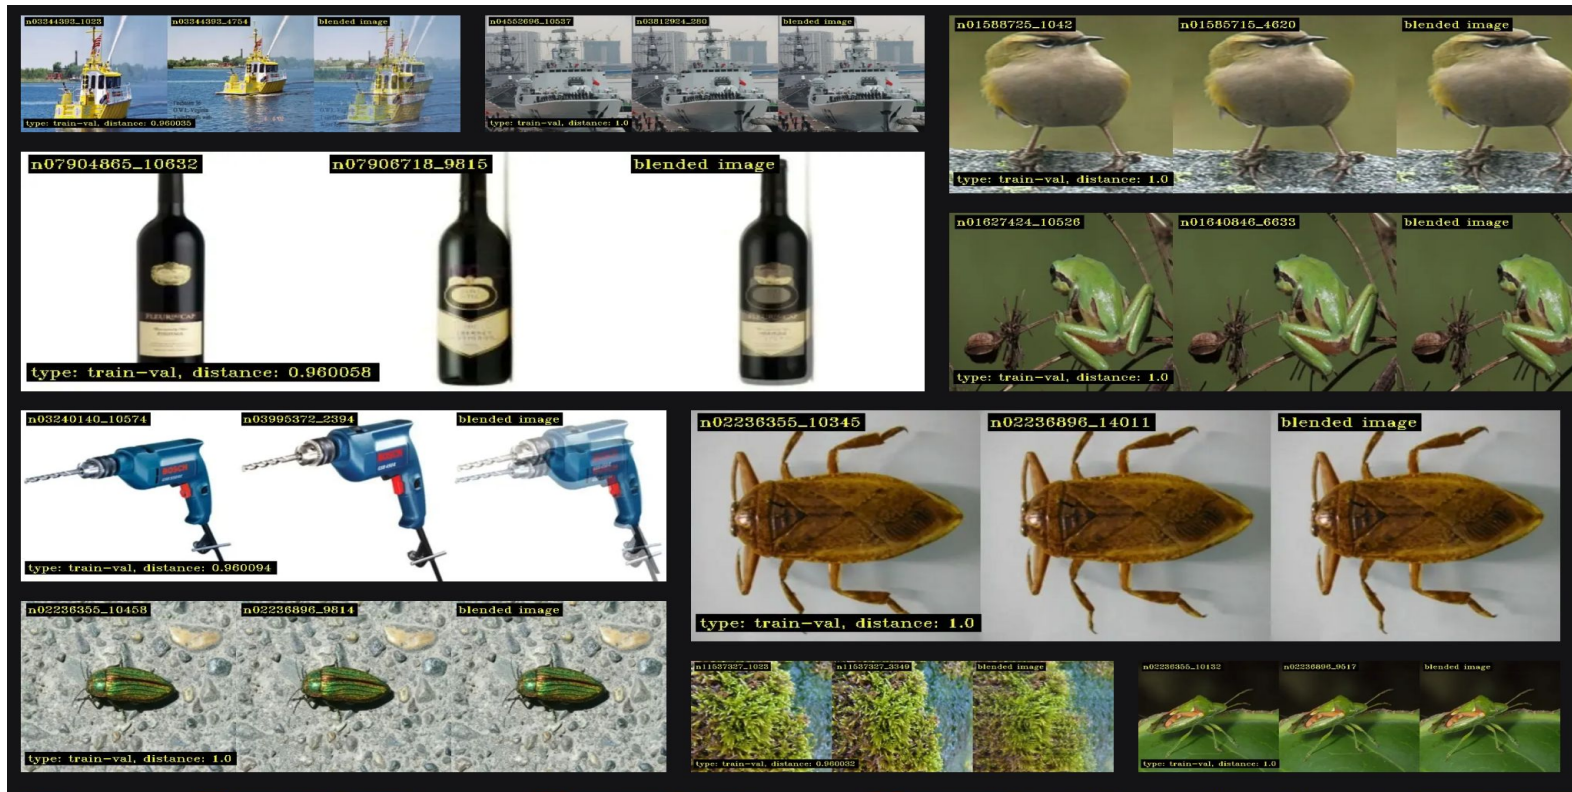


Case Study: ImageNet

32K+

Citations

1.7M Duplicates




Why should we care about duplicates?

- Recent study by Roboflow show 26% reduction in cloud training costs while keeping the same algorithm accuracy

<https://roboflow.com/integration/fastdup>

Train test/leak 100K



50K Wrong Labels



n07892512_15292
red_wine



n07895435_14
red_Bordeaux



n07894102_4847
Beaujolais.



n07895435_825
red_Bordeaux



n07895839_238
Merlot



n07896287_6971
dessert_wine



n07897116_5774
retsina



n07897600_1841
liebfraumlch



n07897750_2622
Rhone_wine



n07898333_10929
zinfandel



n07899003_6359
vin_ordinaire



n07899899_3812
Yquem.



n07900406_4012
fortified_wine



n07900225_4595
varietal,...



n07921615_17
hard_cider

Wrong Label Example



Scanner
C-clamp
Cinnamon
Fairlead
Fuel_gauge
Gas_thermomoter
Grapnel
Hypodermic syringe
Snatch block
Stabling
Wall unit
Winch
Windlass

BLIP Captions



“there is a man riding a surfboard on the beach”

https://github.com/visual-layer/fastdup/blob/main/examples/blip_laion_captions.ipynb

Guess the label ?



n04440963_5127.JPEG

Tire Chain



n04440963_5127.JPEG

Guess the label ?



Two-Spotted Ladybug



Guess the label ?



Challenges of Bad Visual Data

- Slow time to production
- High cost: storage, compute, annotations
- Inefficiency of CV teams
- Lower model quality

What can we
do about it?

Visual Layer Open Source: Trusted by 280,000+ Users

<https://github.com/visual-layer/fastdup>


fastdup



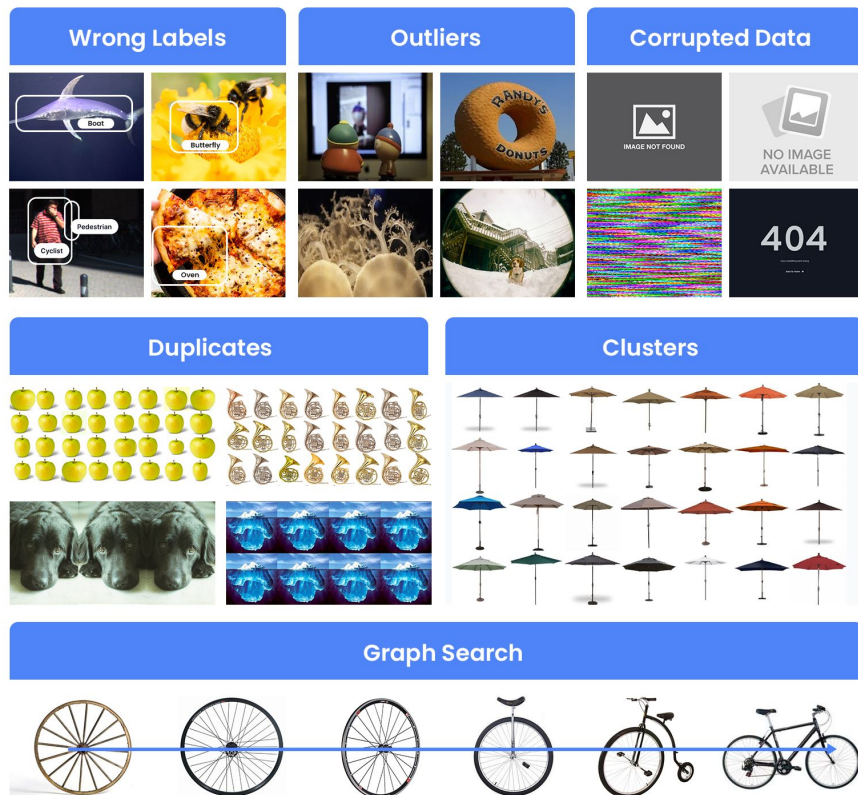
1.2K+
GitHub stars



280K+
Downloads



50B+
Images processed



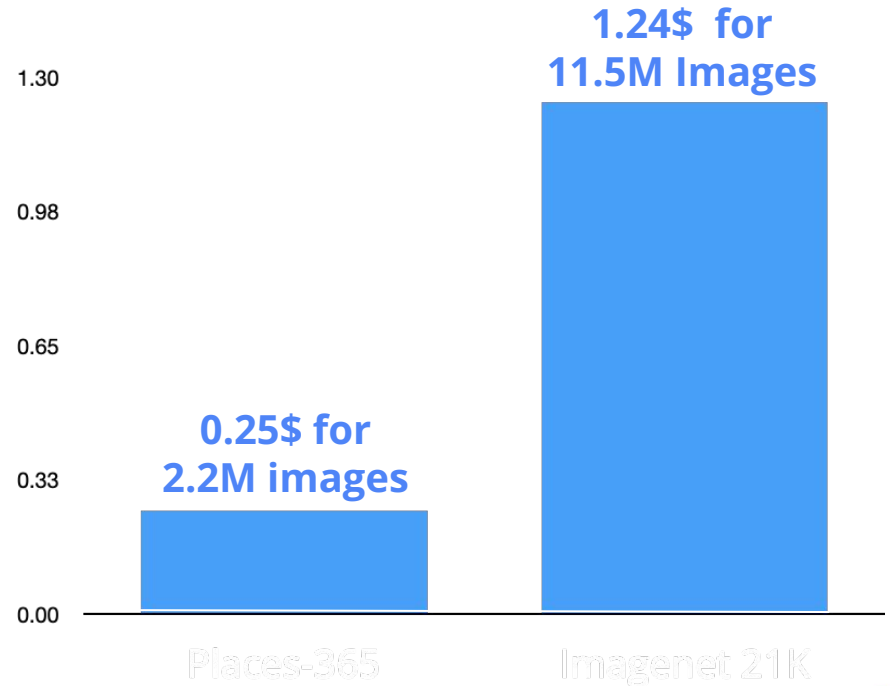
Trusted by users in the following companies:




fastdup Technology

Design principles

- Easy to use
- Efficient / Scalable
- No GPU
- Accurate
- Runs everywhere
- **FREE!!**



System Architecture



Behind the Scenes

2023 VLDB Test of Time Award

Distributed Graphlab: A framework for machine learning in the cloud. Yucheng Low, Joseph Gonzalez (UC Berkeley), Aapo Kyrola, Danny Bickson, Carlos Guestrin, Joseph M. Hellerstein (UC Berkeley). VLDB 2012

Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud

Yucheng Low
Carnegie Mellon University
ylow@cs.cmu.edu

Joseph Gonzalez
Carnegie Mellon University
jegonzal@cs.cmu.edu

Aapo Kyrola
Carnegie Mellon University
akyrola@cs.cmu.edu

Danny Bickson
Carnegie Mellon University
bickson@cs.cmu.edu

Carlos Guestrin
Carnegie Mellon University
guestrin@cs.cmu.edu

Joseph M. Hellerstein
UC Berkeley
hellerstein@cs.berkeley.edu

Case study: growing mosquitos

localhost:2080/dataset/0cf568ce-5964-11ee-bd1e-071b8deddb7

My Datasets → Kaggle Mosquitos

Kaggle Mosquitos

Source VL Images 4803 Size 19GB Created 4 Hours Ago

VL Quality Score **10** Total Clusters **26** out of **26** Sorted by **Severity**

Severity

- High 4827 (100.50%)
- Medium 0 (0.00%)
- Low 0 (0.00%)


Issue Type

- Duplicates 5K (100%)
- Outliers 24 (0.5%)

- > aegypti_0
- > aegypti_0_testing
- > aegypti_90
- > aegypti_90_testing
- > aegypti_180
- > aegypti_180_testing
- > aegypti_270
- > aegypti_270_testing

Cluster Name	Sub-Clusters	Images
Duplicate Images #1	9	4.2K
Duplicate Images #12	5	172
Duplicate Images #7	3	148
Duplicate Images #22	5	144
Duplicate Images #21	3	48
Outlier Images #2	0	20
Duplicate Images #8	2	12
Duplicate Images #4	2	8

VL Profiler Demo





`pip install fastdup`

<https://app.visual-layer.com>

`danny@visual-layer.com`