

How to navigate through the ML research information flood

Dmytro Mishkin. Faculty of Electrical Engineering, CTU in Prague



CFNet: Learning Correlation Functions for One-Stage Panoptic Segmentation

Yifeng Chen, Wenqing Chu, Fangfang Wang, Ying Tai, Ran Yi, Zhenye Gan, Liang Yao, Chengjie Wa...

Recently, there is growing attention on one-stage panoptic segmentation methods which aim to segment instances and stuff jointly within a fully convolutional pipeline efficiently. However, most of the existing works directly feed the backbone features to various segmentation heads ignoring the demands for semantic and instance segmentation are different: The former needs semantic-level discriminative features, while the latter requires features to be distinguishable across instances. To alleviate this, we propose to first predict semantic-level and instance-level correlations among different locations that are utilized to enhance the backbone features, and then feed the improved discriminative features into the corresponding segmentation heads,

Foreword



Dmytro Mishkin
@ducha_aiki

I am preparing for my lectures on How to nav CV/ML literature, how to read papers in an ef way, re-implement then, etc.

Which issues in this area do you personally h Questions? Whar is the most misleading advi get?

Please write in comments below

[Перекласти твіт](#)

5:13 пп · 15 січ. 2022 · Twitter for Android

||| Переглянути дії з твітом

7 Ретвітів 1 Цитувати твіт 135 Уподобань



Amy Tabb
@amy_tabb



У відповідь [@ducha_aiki](#)

Misleading advice: do *whatever* exactly the way someone else does it, whether it be read 10 papers a day, or none.

Your strategy should be allowed to change depending on the needs of the research problem, your own constraints & preferences, etc., and that's ok.

[Перекласти твіт](#)

5:27 пп · 15 січ. 2022 · TweetDeck

1 Ретвіт 9 Уподобань



Foreword: trust noone in ML

I mean, really

- You should develop your own style of reading ML papers, don't trust me
- and never trust authors of the paper
- ...even if it was published on CVPR
- ...especially if it was published on CVPR
- OK, let's go

Naive beliefs about publication

- Published papers are right, and good
- Rejected papers are wrong, or bad

<https://www.youtube.com/watch?v=9Y7NCdKdNyE>



What do I want to achieve?

Why do you need to read papers?

Use-cases

- (a) You are new to the field and want to get familiar with it
- (b) You are new to the field and need to implement something ASAP
- (c) You are familiar with the field, but not the specific area in particular
- (d) You know area well enough and want to keep yourself up-to-date
- (e) I don't want, but I have to - I am CVPR reviewer

Each of this cases requires a little bit different strategy

You are new to the field

and need to implement something ASAP (like, yesterday)

- That is really unfortunate situation.
- The only way to succeed is to be lucky and find ready implementation
- Check the following first:
 - Kaggle & DrivenData competitions
 - Papers With Code
 - fast.ai forums, pyimagesearch, etc.
 - Only then browse papers and github.

Kaggle

Competitions → keyword search → EDA kernels

≡

kaggle

+ Create

Home

Competitions

Datasets

<> Code

Discussions

Courses

More

Your Work

RECENTLY VIEWED

notebook_HPA_EDA

Human Protein Atlas - ...

1st place solution

Google Landmark Retr...

RECENTLY EDITED

notebookd63442bf65

View Active Events

Search

segmentation

Completed X

Results

Sartorius

Sartorius - Cell Instance

Detect single neuronal cells in m

Featured · Code Competition · 15

HuBMAP

HuBMAP - Hacking the

Identify glomeruli in human kidn

Research · Code Competition · 1:

Open Images Instance

Outline segmentation masks of c

Playground · 18 Teams · a year a

iMaterialist (Fashion) 2

Fine-grained segmentation task

Research · 56 Teams · 2 years a

Open Images 2019 - In

Outline segmentation masks of c

Research · 193 Teams · 2 years a

[TRAIN] Sartorius Segmentation EDA+EffDET [TF]

Notebook Data Logs Comments (35)

242 Copy & Edit 277

... RELEVANT DATAFRAME ROW - INDEX=42 ...

	id	annotation	width	height	cell_type	plate_time	sample_date	sample_id	elapsed_timedelta	im
42	13325f865bb0	[195920 2 196823 5 197327 6 198030 7 198733 7 ...	704	520	astro	09h00m00s	2020-09-14	astros[cereb]_G12-2_Vessel-361_Ph_4	0 days 09:00:00	/ki ce se

Cell Image

Instance Segmentation Mask

Cell image w/ Instance Segmentation Mask Overlay

Table of Contents

TABLE OF CONTENTS

0 IMPORTS ↑

1 BACKGROUND...

2 SETUP ↑

3 HELPER FUNCTION &...

4 DATASET CREATION AND...

5 MODELLING ↑

5 SUBMISSION ↑

Kaggle

If task & competition is relevant to you, check solutions

- Winning solutions
- Simple solutions

Research Code Competition

Google Landmark Retrieval 2021

Given an image, can you find all of the same landmarks in a dataset?

Google · 263 teams · 3 months ago

Overview Data Code Discussion Leaderboard Rules

1st place solution

Posted in [landmark-retrieval-2021](#) 3 months ago

Dieter
Topic Author
1st place

First, let me thank kaggle staff and google team for organ...

Overview Data Code Discussion Leaderboard Rules

New Topic ...

1st place solution

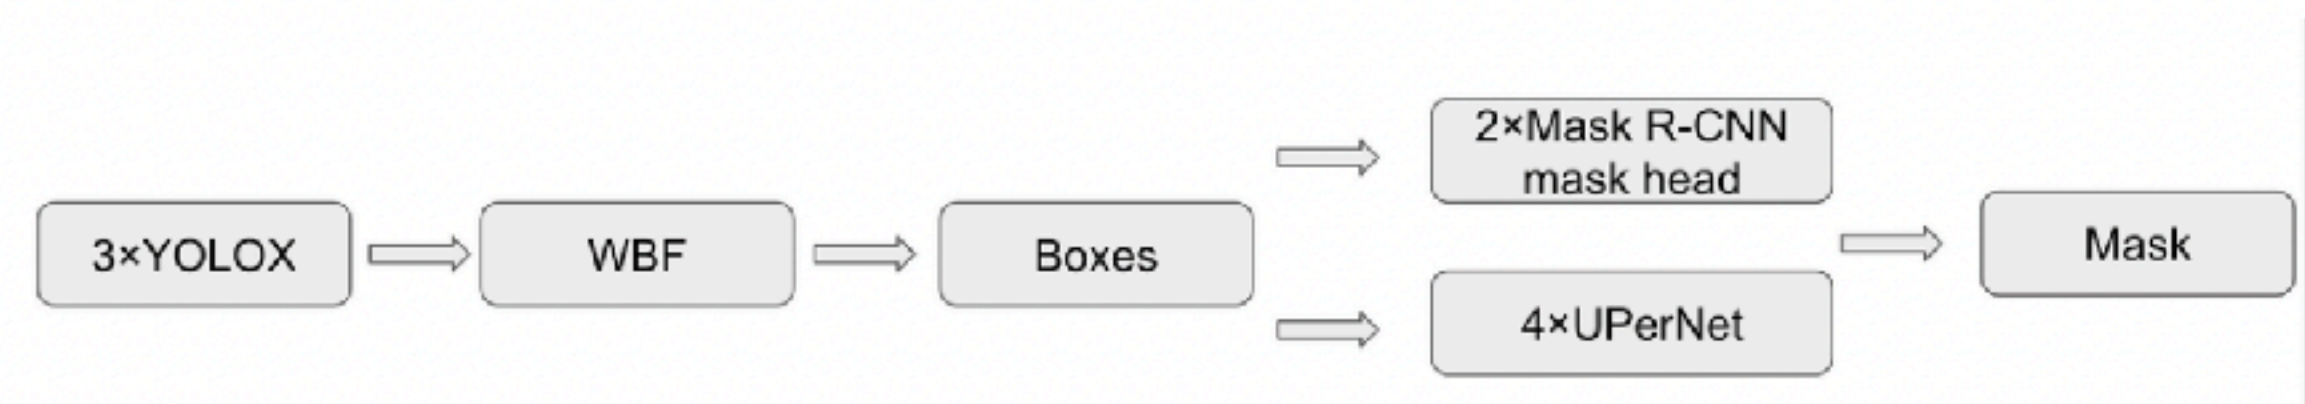
Posted in [sartorius-cell-instance-segmentation](#) 12 days ago

70

We would like to thank Kaggle and Sartorius for organizing such a great competition.

Overview

Our solution is as follows. In many respects, it is similar to the 2nd place solution.



```
graph LR; A[3xYOLOX] --> B[WBF]; B --> C[Boxes]; C --> D[2xMask R-CNN mask head]; C --> E[4xUPerNet]; D --> F[Mask]; E --> F;
```

At the very beginning of the competition, we decided to build a solution using box-based instance segmentation, and focus more on the bbox detection part. We think the mask prediction performance is mainly limited by annotation quality so we did not pay much attention to it.


During the competition, we used COCO mAP as our validation metric, we believe that high mAP and proper thresholding would give a high LB score. Following is the validation score we achieved at the end of the competition

DrivenData

- Much smaller community than Kaggle, with much less forum life
- But contains some unique challenges

PHASE 1 | Facebook AI Image Similarity Challenge: Matching Track

HOSTED BY FACEBOOK



Problem description

Summary

You will receive a reference set of 1 n images are derived from images in the

For this **Matching Track**, your task is one of the images in a large corpus of















Table of contents

[Challenge](#)
[Background](#)
[Phases](#)
[Data](#)
[Overview](#)
[Details](#)




Snowcast Showdown: Development Stage

HOSTED BY BUREAU OF RECLAMATION

[HOME](#) [PROBLEM DESCRIPTION](#) [DEVELOPMENT STAGE](#) [REPORT TEMPLATE](#) [ABOUT](#)

User or team			Best public ↓ RMSE ⓘ	R ² ⓘ	Timestamp ⓘ	Trend (last 10)	# Entries
	rasyidstat	1	3.4061	0.8305	2022-01-13 04:53:49		12
	FBykov	2	3.5919	0.8115	2021-12-29 17:13:44		75
	Team UA	3	3.6797	0.8021	2022-01-13 22:47:59		14
	separate	4	3.7156	0.7982	2022-01-13 23:54:41		36
	andrey1362010	5	3.8857	0.7793	2022-01-14 14:57:01		14
	NxGTR	6	4.0837	0.7563	2022-01-14 13:25:01		16
	Galeros93	7	4.1490	0.7484	2022-01-02 22:21:21		8

Papers With Code





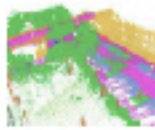
[Browse State-of-the-Art](#)[Datasets](#)[Methods](#)[More ▾](#)[We are hiring!](#)

[Browse SoTA](#) > Computer Vision

Computer Vision

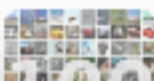
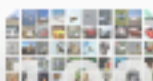



2594 benchmarks • 934 tasks • 1702 datasets • 22504 papers with code

Semantic Segmentation

 <div>Semantic Segmentation 📊 206 benchmarks 2411 papers with code</div>	 <div>Medical Image Segmentation 📊 86 benchmarks 254 papers with code</div>	 <div>Tumor Segmentation 📊 1 benchmark 113 papers with code</div>	 <div>Panoptic Segmentation 📊 10 benchmarks 77 papers with code</div>	 <div>3D Semantic Segmentation 📊 8 benchmarks 76 papers with code</div>
--	---	---	---	---

▶ [See all 19 tasks](#)

Image Classification

 <div>Image</div>	 <div>Knowledge</div>	 <div>Few-Shot Image</div>	 <div>Fine-Grained Image</div>	 <div>Semi-Supervised Image</div>
---	---	--	--	---

Papers With Code

Caveat: sparsely populated

Image Matching on IMC PhotoTourism

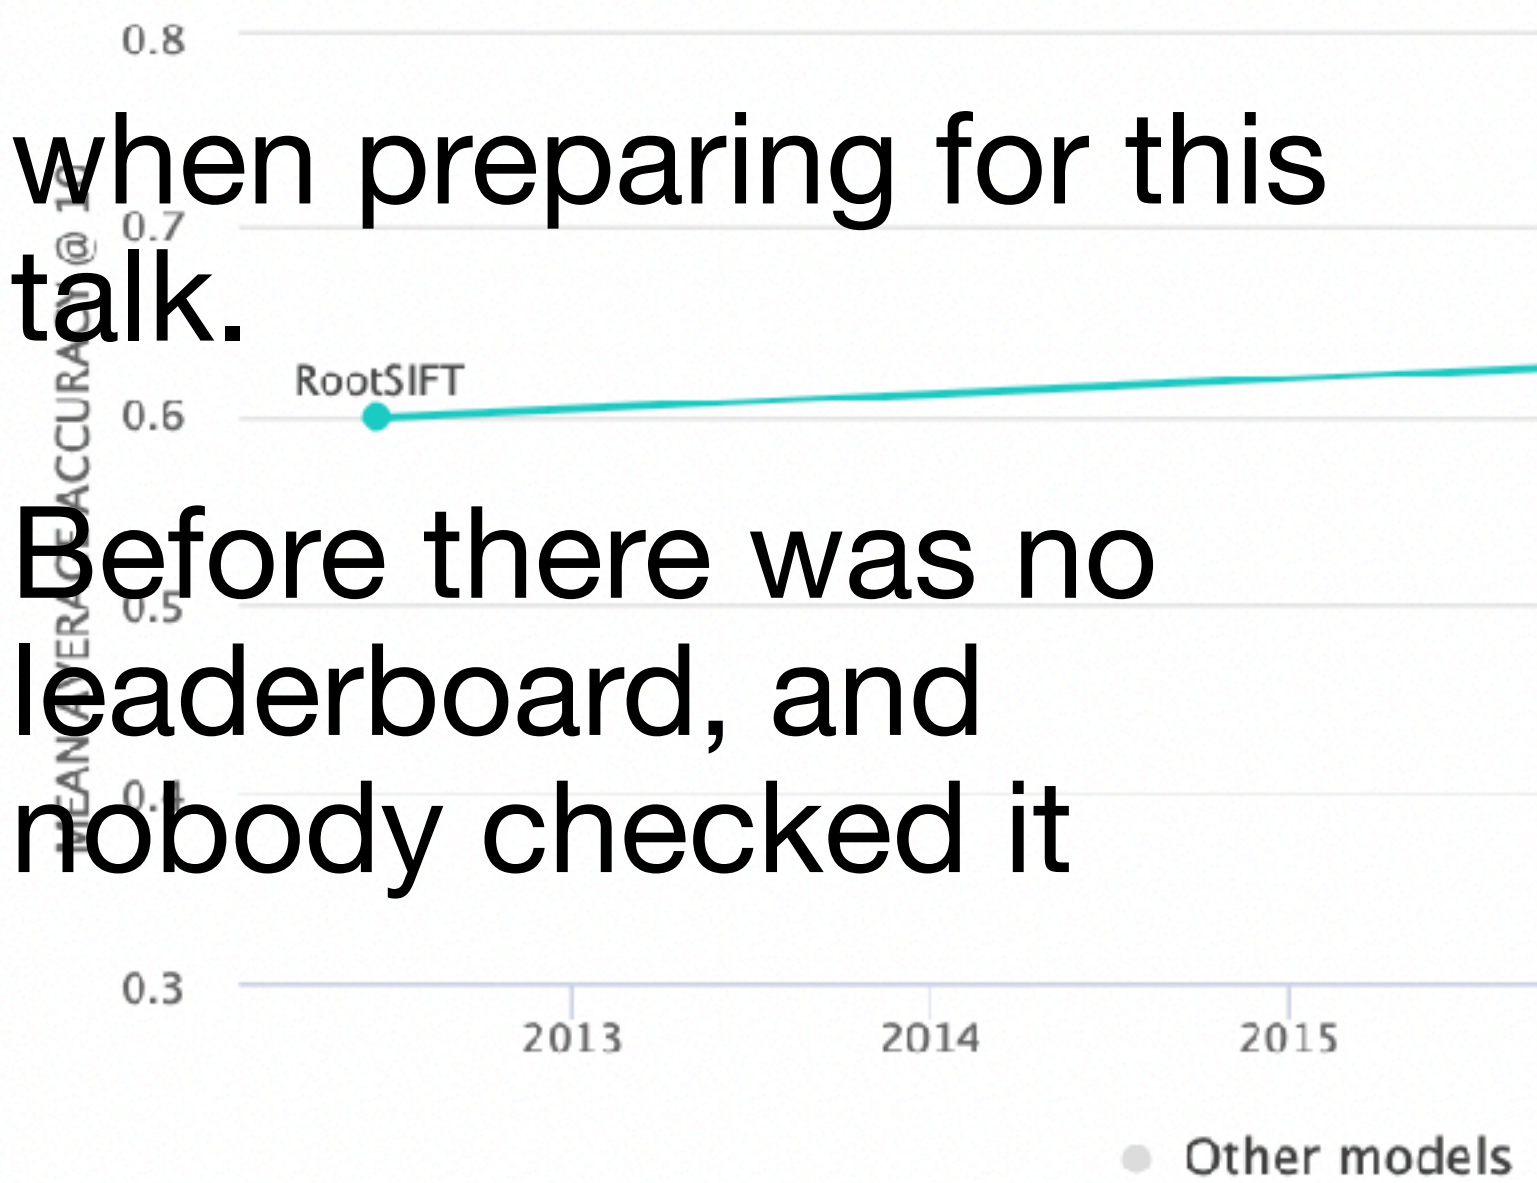
Leaderboard

Dataset

I have added this whole
leaderboard

when preparing for this
talk.

Before there was no
leaderboard, and
nobody checked it



Rank	Model	mean average accuracy @ 10	Extra Training Data	Paper	Code	Result	Year	Tags
1	LoFTR	0.68503	✓	LoFTR: Detector-Free Local Feature Matching with Transformers	Code	Result	2021	
2	HarrisZ improved, Blob DTM	0.65606	×	HarrisZ ⁺ : Harris Corner Selection for Next-Gen Image Matching Pipelines		Result	2021	
3	DISK	0.65435	✓	DISK: Learning local features with policy gradient	Code	Result	2020	
4	SuperGlue	0.65248	✓	SuperGlue: Learning Feature Matching with Graph Neural Networks	Code	Result	2019	
5	DoG-AffNet-HardNet8	0.64212	✓	Repeatability Is Not Enough: Learning Affine Regions via Discriminability	Code	Result	2017	
6	Key.Net-SOSNet	0.60285	×	Key.Net: Keypoint Detection by Handcrafted and Learned CNN Filters	Code	Result	2019	
7	PoSFeat	0.60072	×	Decoupling Makes Weakly Supervised Local Feature Better		Result	2022	
8	RootSIFT	0.59859	×	Three things everyone should know to improve object retrieval	Code	Result	2012	
9	R2D2	0.56345	✓	R2D2: Reliable and Repeatable Detector and Descriptor	Code	Result	2019	

Papers With Code

Caveat: sparsely populated

Methodology

Representation Learning


1643 papers with code 2 benchmarks • 2 datasets

Representation learning is concerned with training machine learning algorithms to learn useful representations, e.g. those that are interpretable, have latent features, or can be used for transfer learning.

Deep neural networks can be considered representation learning models that typically encode information which is projected into a different subspace. These representations are then usually passed on to a linear classifier to, for instance, train a classifier.

Representation learning can be divided into:

- **Supervised representation learning:** learning representations on task A using annotated data and used to solve task B



Content

Papers With Code

Caveat: sparsely populated

Methodology

Representation Learning

1643 papers with code

Representation learning is concerned with learning good representations, e.g. those that can be used for many different tasks.

Deep neural networks can learn representations which are projected onto a linear classifier.

Representation learning can be used to solve many tasks.

- **Supervised representation learning** and used to solve task B

Benchmarks

These leaderboards are used to track progress in Representation Learning.

Trend

Dataset

Best Model




Circle Data

🏆 Morphological Network



Sports10

🏆 Max Margin Contrastive




See all

If both Kaggle & PapersWithCode fail

Then go to github + arXiv search

- You might want to apply heuristics to decide which papers you should read first
- But that is later, now let's assume you have some time...



You are new to the field

and would like to get familiar with it

- The best idea is to start with some modern course (e.g. fast.ai)
- Starting with dataset to play with (see Kaggle, PapersWithCode) is also a good idea.
- Also a good idea is to start with a good benchmark paper
- Or a good survey paper
- But it is not always easy to find good survey paper, there many bad ones

How to spot good benchmark paper?

Citations

- It is used by community (unless it is just published 2 days ago), i.e. people cite it

<https://arxiv.org> › cs ▼ [Přeložit tuto stránku](#)

Revisiting Oxford and Paris: Large-Scale Image Retrieval ...

autor: F Radenović · 2018 · **Počet citací tohoto článku: 203** → In this paper we address issues with image **retrieval** benchmarking on standard and popular **Oxford 5k** and **Paris 6k** datasets...

<https://research.google> › pub49052 ▼ [Přeložit tuto stránku](#)

Google Landmarks Dataset v2 - A Large ... - Google Research

autor: T Weyand · 2020 · **Počet citací tohoto článku: 72** → We introduce the **Google Landmarks Dataset v2** (GLDv2), a new benchmark for large-scale, fine-grained instance recognition and...

<https://arxiv.org> › cs ▼ [Přeložit tuto stránku](#)

Image Matching across Wide Baselines: From Paper to Practice

autor: Y Jin · 2020 · **Počet citací tohoto článku: 78** → Abstract: We introduce a comprehensive benchmark for local features and robust estimation algorithms, focusing on the downstream ta...

Cite as: [arXiv:2003.01587](#)

<https://www.robots.ox.ac.uk> › vibes_ijcv2004 ▼ [PDF](#)

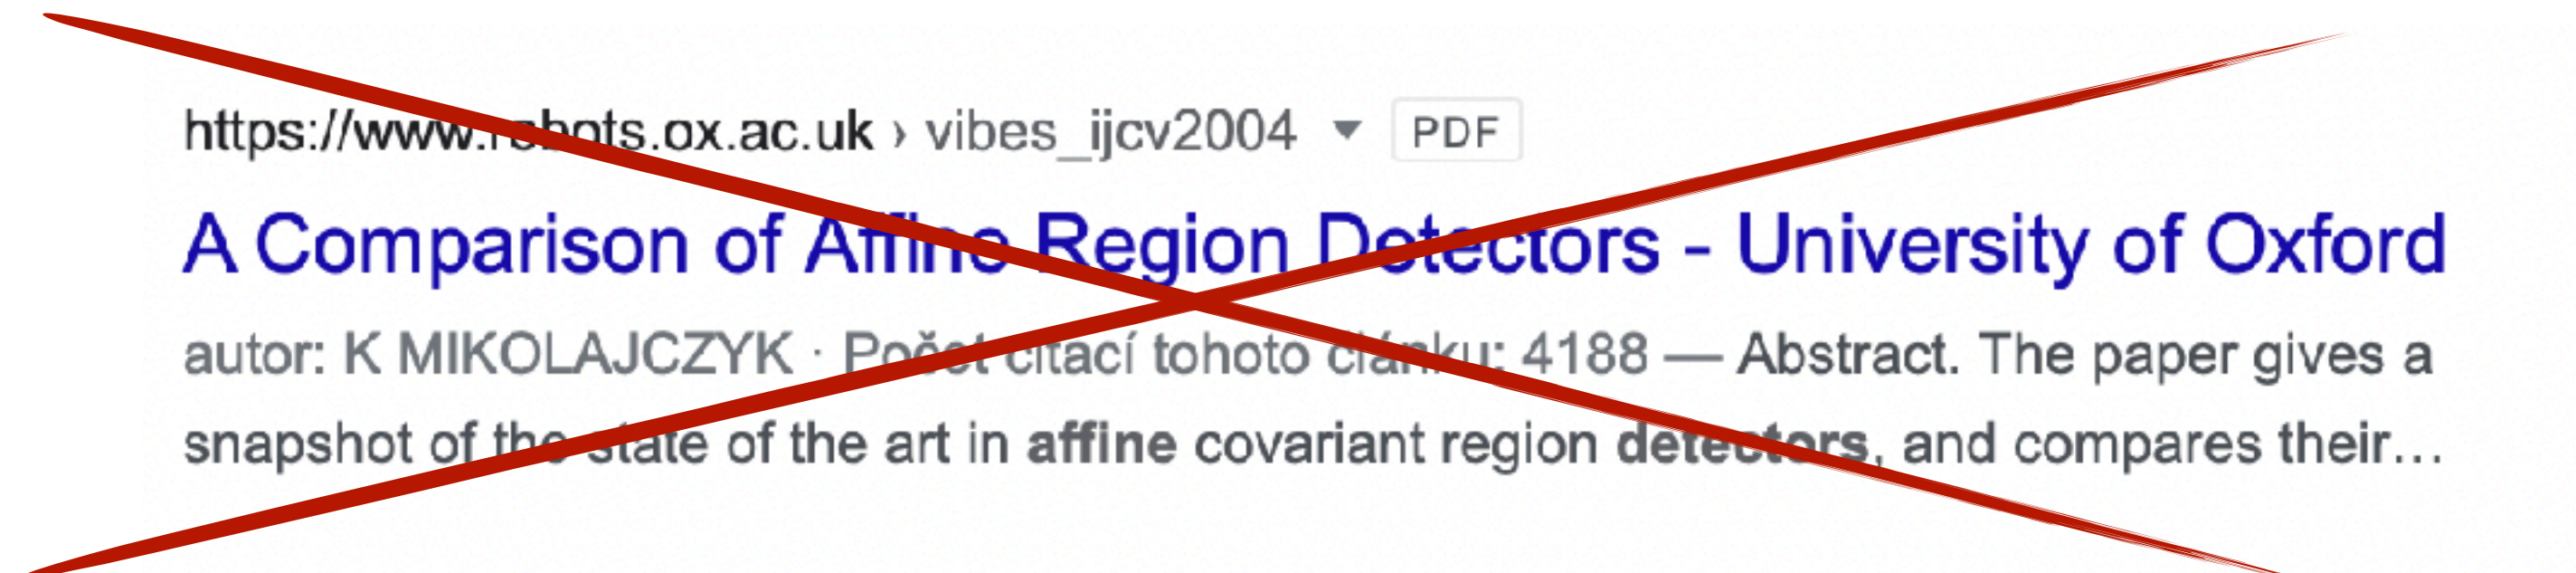
A Comparison of Affine Region Detectors - University of Oxford

autor: K MIKOLAJCZYK · **Počet citací tohoto článku: 4188** → Abstract. The paper gives a snapshot of the state of the art in **affine covariant region detectors**, and compares their...

How to spot good benchmark paper?

Year

- It should not be too old: benchmark from 2004 is a little bit obsolete
- Oxford-Affine is GREAT paper. But it would not help you find modern local features



How to spot good benchmark paper?

It proposes new dataset (or new annotation)

- There are TONS of papers, which just run methods from OpenCV on existing (and old) datasets



A Comparative Evaluation of Well-known Feature Detectors and Descriptors[#]

3.7. Effects of JPEG Compression

In this experiment, impacts of compression is examined in terms of comparison of results obtained from each method conducting on ubc dataset. For this purpose, the artifacts have introduced in Joint Photographic Experts Group (JPEG) compression by using a standard xv image browser with the image quality parameter changing from 40% to 2%.


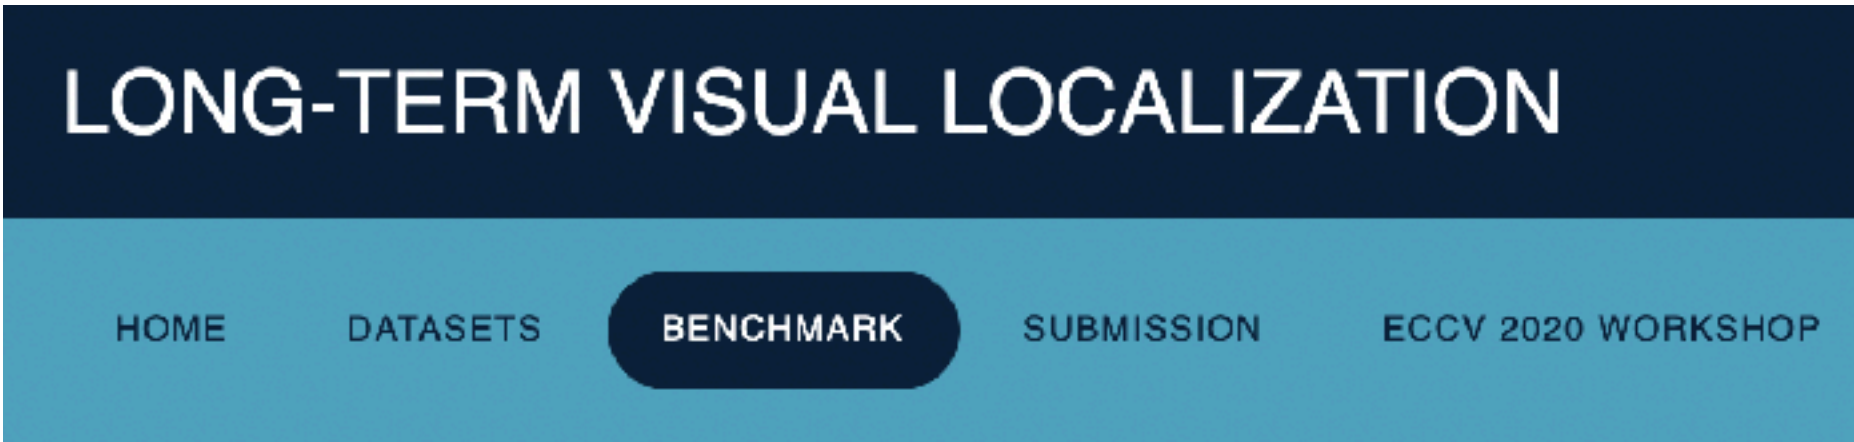


Figure 7: Precision and recall values for JPEG compression.

How to spot good benchmark paper?

It has a challenge linked to it, often — many years



Benchmark

Listed below are the public results on the three benchmark datasets. The localization results are reported as the percentage translation and rotation thresholds, for each condition.

Aachen Day-Night dataset

Localization thresholds:


All conditions: (0.25m, 2°) / (0.5m, 5°) / (5m, 10°)

Method	day
Combined tuned HLoc	89.7 / 95.8 / 98
Hierarchical Localization - SuperPoint + SuperGlue	89.6 / 95.4 / 98
SuperGlue + Patch2Pix (HLoc)	89.2 / 95.5 / 98
ONavi-V2.0	88.6 / 95.1 / 99
hloc	89.3 / 95.9 / 98
KAPTURE-R2D2-FUSION	89.4 / 96.4 / 99

<https://www.visuallocalization.net/benchmark/>


Benchmarks

Official results of past challenges are still available as a benchmark to researchers. We invite the visual tracking researchers to compare their methods with those results.




VOT2021 benchmark

The VOT2021 benchmark addresses short-term, long-term, real-time, RGB and RGBD trackers. Results were presented at the ICCV2021 VOT workshop.




VOT2020 benchmark

The VOT2020 benchmark addresses short-term, long-term, real-time, RGB, RGBT and RGBD trackers. Results were presented at the ECCV2020 VOT workshop.




VOT2019 benchmark

The VOT2019 benchmark addresses short-term, long-term, real-time, RGB, RGBT and RGBD trackers. Results were presented at the ICCV2019 VOT workshop.




VOT2018 benchmark

The VOT2018 benchmark introduced a long-term subchallenge VOT-LT2018. Results were presented at the VOT workshop at ECCV2018.



VOT2017 benchmark

The VOT2017 benchmark introduced a refreshed dataset and a real-time experiment. The winner will be determined on sequestered dataset. The results were presented at the VOT workshop at ICCV2017.



VOT2016 benchmark

The fourth challenge updated the dataset of 60 sequences with new annotations. The results were published in a joint paper presented at a workshop at ECCV2016.

<https://www.votchallenge.net>

How to spot good benchmark paper?


It re-evaluates existing method on the common ground

- Some papers do not propose new dataset, but they critically evaluate existing metrics
- And/or re-implement the methods from scratch on the common codebase


<https://arxiv.org> › cs ▼ Přeložit tuto stránku

[\[2003.08505\] A Metric Learning Reality Check - arXiv](#)

autor: K Musgrave · 2020 · Počet citací tohoto článku: 159 — Abstract: Deep metric learning papers from the past four years have consistently claimed great advances in accuracy, often...



(a) The trend according to papers



(b) The trend according to reality

Good Surveys: harder to distinguish

But there are couple of heuristics

- Good modern survey is getting updated
- available on arXiv/author page
- bad survey is “published” in obscure journals or behind the paywall

<http://www.ijeetc.com> › 2_NCETEC001_(p.9-17).... PDF

A SURVEY OF TRENDS IN LOCAL INVARIANT FEATURE ...

autor: JN Surekha · 2017 · Počet citací: 1 · článku: 2 — Many computer vision algorithms use

feature detection as the initial step, so as a result, a very large number of feature detectors...

9 stránek

arXiv:1704.05519 (cs)

[Submitted on 18 Apr 2017 (v1), last revised 17 Mar 2021 (this version, v3)]

Computer Vision for Autonomous Vehicles: Problems, Datasets and State of the Art

Joel Janai, Fatma Güney, Aseem Behl, Andreas Geiger

[Submitted on 26 Jul 2018 (v1), last revised 14 Jun 2019 (this version, v3)]

From handcrafted to deep local features

Gabriela Csurka, Christopher R. Dance, Martin Humenberger

[Submitted on 14 Sep 2020 (v1), last revised 16 Sep 2020 (this version, v2)]

Efficient Transformers: A Survey




Yi Tay, Mostafa Dehghani, Dara Bahri, Donald Metzler

Good Surveys: harder to evaluate at least without experience

- Cheat: you can use new edition (2021) of Szeliski book as a meta-survey.

Computer Vision: Algorithms and Applications, 2nd ed.

© 2021 [Richard Szeliski](#), The University of Washington



How to read benchmark/survey

How to read benchmark/survey paper if you are new to the field

- First, understand the task, by reading the intro and thinking. Write down:
- Who are the “users” of the algorithm?
 - Where is the result is going to be used?
 - What is important for the downstream task?
 - How one can measure it?
- Where the data comes from in production?
- Can we avoid solving this task at all?



Case study: image matching

Task understanding (1)

- What is the task?
 - Image matching.
- What are the inputs and output of the task?
 - **Missing in the intro!** Need to google



what is image matching

Vše

Obrázky

Nákupy

Videa

Zprávy

Více

Nástroje

Přibližný počet výsledků: 10 510 000 000 (0,58 s)

Image matching is **an important concept in computer vision and object recognition**. Images of the same item can be taken from any angle, with any lighting and scale. This as well as occlusion may cause problems for recognition. But ultimately, they still show the same item and should be categorized that way.

<https://ai.stanford.edu/~syyeong/cvweb/tutorial2>

[Tutorial 2: Image Matching - Stanford AI Lab](#)

Noname manuscript No.
(will be inserted by the editor)

Image Matching Across Wide Baselines: From Paper to Practice

Yuhu Jin · Dmytro Mishkin · Anastasiia Mishchuk · Jiri Matas · Pascal Fua ·
Kwang Moo Yi · Eduard Trulls

Received: date / Accepted: date

Abstract We introduce a comprehensive benchmark for local features and robust estimation algorithms, focusing on the downstream task – the accuracy of the reconstructed camera pose – as our primary metric. Our pipeline’s modular structure allows easy integration, configuration, and combination of different methods and heuristics. This is demonstrated by embedding dozens of popular algorithms and evaluating them, from seminal works to the cutting edge of machine learning research. We show that with proper settings, classical solutions may still outperform the *perceived state of the art*.

Besides establishing the *actual state of the art*, the conducted experiments reveal unexpected properties of Structure from Motion (SfM) pipelines that can help improve

This work was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant “Deep Visual Geometry Machines” (RGPIN-2018-03788), by systems supplied by Compute Canada, and by Google’s Visual Positioning Service. DM and JM were supported by OP VVV funded project CZ.02.1.01/0/0/0/0/16_019/0/000765 “Research Center for Informatics”. DM was also supported by CTU student grant SGS17/185/OHES/3/1/13 and by the Austrian Ministry for Transport, Innovation and Technology, the Federal Ministry for Digital and Economic Affairs, and the Province of Upper Austria in the frame of the COMET center SCCH. AM was supported by the Swiss National Science Foundation.

Y. Jin, K.M. Yi
Visual Computing Group, University of Victoria
E-mail: {yuhujin, kyii}@uvic.ca

A. Mishchuk, P. Fua
Computer Vision Lab, École Polytechnique Fédérale de Lausanne
E-mail: {anastasia.mishchuk, pascal.fua}@epfl.ch

D. Mishkin, J. Matas
Center for Machine Perception, Czech Technical University in Prague
E-mail: {dmishkin, matas}@cmp.felk.cvut.cz

E. Trulls
Google Research
E-mail: trulls@google.com




Fig. 1 Every paper claims to outperform the state of the art. Is this possible, or an artifact of insufficient validation? On the left, we show stereo matches obtained with D2-Net (2019) [37], a state-of-the-art local feature, using OpenCV RANSAC with its default settings. We color the inliers in green if they are correct and in red otherwise. On the right, we show SIFT (1999) [34] with a carefully tuned MAGSAC [13] – notice how the latter performs much better. This illustrates our take home message: to correctly evaluate a method’s performance, it needs to be embedded within the pipeline used to solve a given problem, and the different components in said pipeline need to be tuned carefully and jointly, which requires engineering and domain expertise. We fill this need with a new, modular benchmark for sparse image matching, incorporating dozens of built-in methods.

their performance, for both algorithmic and learned methods. Data and code are online¹, providing an easy-to-use and flexible framework for the benchmarking of local features and robust estimation methods, both *alongside* and *against* top-performing methods. This work provides a basis for the Image Matching Challenge².

¹ <https://github.com/veq-uvic/image-matching-benchmark>
² <https://vision.uvic.ca/image-matching-challenge>

arXiv:2003.01587v5 [cs.CV] 11 Feb 2021

Case study: image matching

Task understanding (1)

- Inputs: images of the “same item”.
- Output (after more googling): correspondences and relative camera pose
- Note: people can use different terms for the same thing:
 - "Image matching" == “Wide baseline stereo”

The wide multiple baseline stereo (WxBS) is a **process of establishing a sufficient number of pixel or region correspondences from two or more images depicting the same scene to estimate the geometric relationship between cameras, which produced these images.** Mar 27, 2020

<https://ducha-aiki.github.io> › 2020/03/27 › intro

The Role of Wide Baseline Stereo in the Deep Learning World



Case study: image matching

Task understanding (2)

- Who are the “users” of the algorithm? Where is the result is going to be used?
 - Image retrieval
 - 3D reconstruction
 - visual localization
 - SLAM
- Now we may, or may not need to check about this “users”

1 Introduction

Matching two or more views of a scene is at the core of fundamental computer vision problems, including image retrieval [54, 7, 78, 104, 70], 3D reconstruction [3, 47, 90, 122], re-localization [85, 86, 57], and SLAM [68, 33, 34]. Despite

Case study: image matching

Task understanding (2)

- What is important for the downstream task?
 - Image retrieval — if the result is relevant to query
 - 3D reconstruction — accuracy and completeness of the reconstruction
 - visual localization — camera pose accuracy, robustness to nuisance
 - SLAM — camera pose accuracy, realtime speed, blur-robustness

Is it relevant to me?

Case study: image matching

Task understanding (3)

- We have already googled and checked other papers to answer the questions, which are necessary to understand the paper we are reading now
- It is OK to deviate (a lot), when you are reading
- Especially when you are new to the area
- It is a good idea to “learn in public” and write answers to the question in a blog post.

Case study: image matching

How can we measure that?

- What is important for the downstream task?
 - Image retrieval — if the result is relevant to query
 - 3D reconstruction — accuracy and completeness of the reconstruction
 - visual localization — camera pose accuracy, robustness to nuisance
 - SLAM — camera pose accuracy, realtime speed, blur-robustness

Recall and precision

Lots of metrics

Pose error, accuracy

Pose error, fps

Case study: image matching

Task understanding (4)

- Where the data comes from in production?
 - Image retrieval — user upload, can be anything
 - 3D reconstruction — professional cameras, drones, smartphone, historical photos, etc
 - visual localization — robot camera, car, smartphone, internet photo (OSI)
 - SLAM — robot, UAV, smart glasses, car

Back to the paper

What does the paper evaluate?

Case study: Image matching

- What does paper measure?
- Where the data in the paper comes from?
- Where the ground truth labels come from?

What does the paper evaluate?

Again, we are speaking purely about the task now

- What does paper measure — mean average accuracy
- Where the data in the paper comes from?
 - YCC100M dataset, from flickr (user photos)
- Where the ground truth labels come from?
 - COLMAP (SIFT-based 3d reconstruction software)

3 The Image Matching Challenge PhotoTourism Dataset

While it is possible to obtain very accurate poses and depth maps under controlled scenarios with devices like LIDAR, this is costly and requires a specific set-up that does not scale well. For example, Strecha's dataset [95] follows that approach but contains only 19 images. We argue that a truly representative dataset must contain a wider range of transformations – including different imaging devices, time of day, weather, partial occlusions, etc. Phototourism images satisfy this condition and are readily available.

We thus build on 25 collections of popular landmarks originally selected in [47, 101], each with hundreds to thousands of images. Images are downsampled with bilinear interpolation to a maximum size of 1024 pixels along the long-side and their poses were obtained with COLMAP [90], which provides the (pseudo) ground truth. We do exhaustive image matching before Bundle Adjustment – unlike [92], which uses only 100 pairs for each image – and thus provide enough matching images for any conventional SfM to return near-perfect results in standard conditions.


How the paper data & metric matches user?

Case study: image matching, metric

Note: not clear, why paper says that image matching is used in image retrieval

Add to the ToDo list to figure out (low priority)

- User metrics:
 - Image retrieval — Recall and precision
 - 3D reconstruction — many metrics
 - visual localization — camera pose accuracy
 - SLAM — camera pose accuracy, realtime speed



- Paper metrics:
 - mean average accuracy on camera pose

Case study: image matching

Case study: image matching, data

- Where the data comes from in production?
 - Image retrieval — user upload, can be anything
 - 3D reconstruction — professional cameras, drones, smartphone, historical photos, etc
 - visual localization — robot camera, car, smartphone, internet photo (OSI)
 - SLAM — robot, UAV, smart glasses, car



- Paper:


- YCC100M dataset, from flickr (user photos)

Case study: image matching


Case study: image matching, data

- How the paper results are relevant for
 - Image retrieval
 - 3D reconstruction
 - visual localization
 - SLAM

Data



Metrics



Intermediate conclusion

- The paper findings are partially relevant for some use-cases in
 - visual localization
 - 3D reconstruction
- Loosely relevant for the SLAM
- Not relevant at all for image retrieval

Noname manuscript No.
(will be inserted by the editor)

Image Matching Across Wide Baselines: From Paper to Practice

Yihe Jin · Dmytro Mishkin · Anastasiia Mishchuk · Jiri Matas · Pascal Fua · Kwang Moo Yi · Eduard Trulls

Received: date / Accepted: date

Abstract We introduce a comprehensive benchmark for local features and robust estimation algorithms, focusing on the downstream task – the accuracy of the reconstructed camera pose – as our primary metric. Our pipeline's modular structure allows easy integration, configuration, and combination of different methods and heuristics. This is demonstrated by embedding dozens of popular algorithms and evaluating them, from seminal works to the cutting edge of machine learning research. We show that with proper settings, classical solutions may still outperform the *perceived state of the art*.

Besides establishing the *actual state of the art*, the conducted experiments reveal unexpected properties of Structure from Motion (SfM) pipelines that can help improve their performance, for both algorithmic and learned methods. Data and code are online¹, providing an easy-to-use and flexible framework for the benchmarking of local features and robust estimation methods, both *alongside* and *against* top-performing methods. This work provides a basis for the Image Matching Challenge².

¹ <https://github.com/eqy-uvio/image-matching-benchmark>
² <https://vision.uvic.ca/image-matching-challenge/>

arXiv:2003.01587v5 [cs.CV] 11 Feb 2021

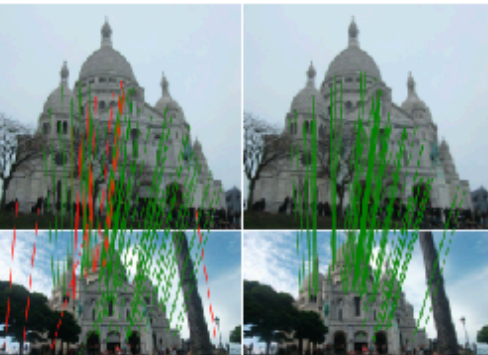


Fig. 1 Every paper claims to outperform the state of the art, is this possible, or an artifact of insufficient validation? On the left, we show stereo matches obtained with D2-Net (2019) [17], a state of the art local feature, using OpenCV RANSAC with its default settings. We color the inliers in green if they are correct and in red otherwise. On the right, we show SIFT (1999) [34] with a carefully tuned MAGSAC [13] – notice how the latter performs much better. This illustrates our take-home message: to correctly evaluate a method's performance it needs to be embedded within the pipeline used to solve a given problem, and the different components in said pipeline need to be tuned carefully and jointly, which requires engineering and domain expertise. We fill this need with a new, modular benchmark for sparse image matching, incorporating dozens of built-in methods.

If you need a truck,
do not watch racing car reviews



o



Research sometimes deviates from practice or often



Christoph Molnar
@ChristophMolnar

A lot of machine learning research has moved away from solving real problems, and creating "benchmark-islands".


How does this happen? And why are we not escaping this pattern?

A thread 

5. More researchers publish data and predictive models.

6. A community forms that cites and reviews each other.
7. Certain datasets become benchmarks and certain predictive models become state-of-the-art.
8. Predictive performance becomes the sole measure of progress, although improvements are becoming smaller.
9. Actual progress in solving the initial research questions has become irrelevant, even discouraged.
10. Decoupling from reality is complete.

OK, paper is relevant



Can we trust the paper?

*"The only correct attitude, when reading the paper is:
Why do these lying bastard lie to me?"*

David Forsyth

Do the metrics in the paper make sense?

If you don't have experience, you can skip this

- Even if the metric is measuring a proper thing, it can still be misleading
 - Is it robust to outliers? E.g. 1 failure can kill average pose error over dataset
 - Isn't it too robust to outliers? E.g. median error does not change if half of dataset is not solved at all
- What is the “null” performance according to the metric?




Figure 1. The difference between comparing algorithms in ROC vs PR space

How the ground truth labels were obtained?

Noname manuscript No.
(will be inserted by the editor)

- Check if the annotation process makes sense
- Is it biased?

ImageNet Large Scale Visual Recognition Challenge

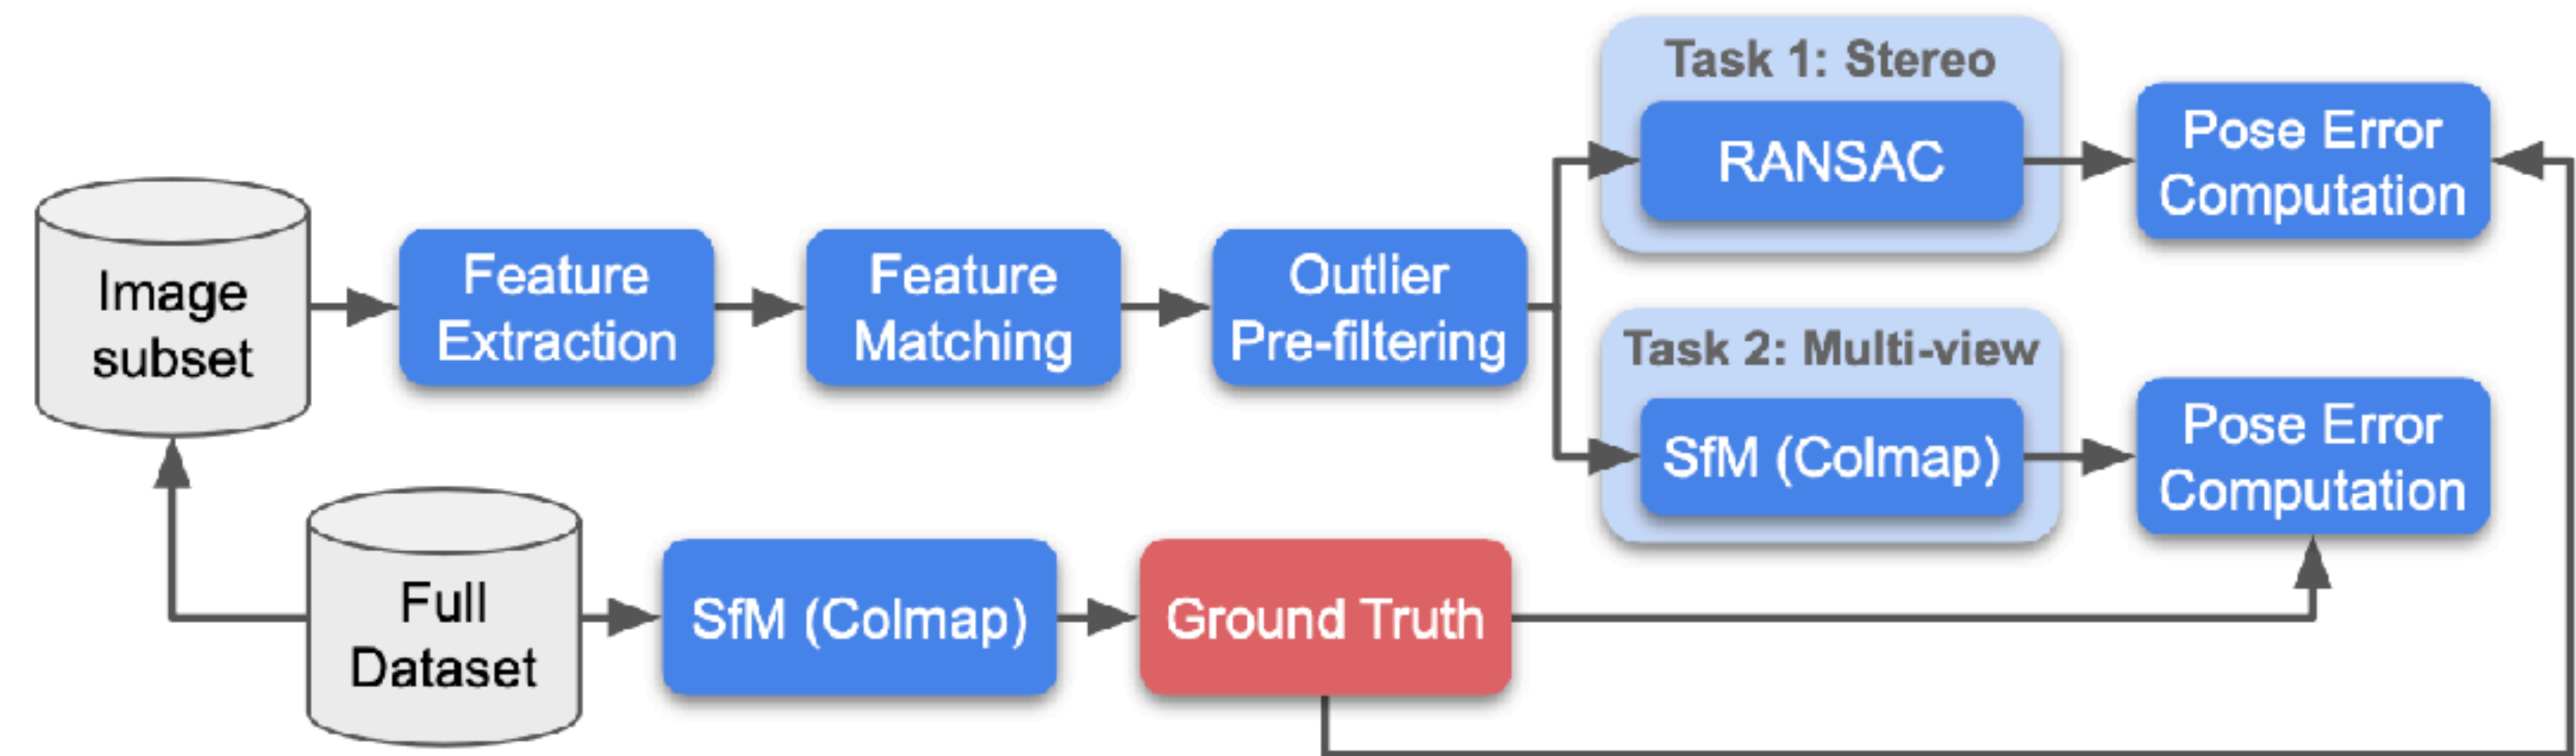
Olga Russakovsky* · Jia Deng* · Hao Su · Jonathan Krause ·
Sanjeev Satheesh · Sean Ma · Zhiheng Huang · Andrej Karpathy ·
Aditya Khosla · Michael Bernstein · Alexander C. Berg · Li Fei-Fei

1. **Drawing** A worker draws one bounding box around one instance of an object on the given image.
2. **Quality verification** A second worker checks if the bounding box is correctly drawn.
3. **Coverage verification** A third worker checks if all object instances have bounding boxes.

How the ground truth labels were obtained?

Case study: image matching

- Assumption: COLMAP reconstruction from the 1000 images outputs poses, which are good enough to evaluate stereo poses



How the ground truth labels were obtained?

Case study: image matching

- Assumption: COLMAP reconstruction from the 1000 images outputs poses, which are good enough to evaluate stereo poses
- Is it validated?

3.3 On the quality of the “ground-truth”

Our core assumption is that accurate poses can be obtained from large sets of images without human intervention. Such poses are used as the “ground truth” for evaluation of image matching performance on pairs or small subsets of images – a harder, proxy task. Should this assumption hold, the (relative) poses retrieved with a large enough number of images would not change as more images are added, and these poses would be the same regardless of which local feature is used. To validate this, we pick the scene “Sacre Coeur” and compute SfM reconstructions with a varying number of images: 100, 200, 400, 800, and 1179 images (the entire “Sacre Coeur” dataset), where each set contains the previous one; new images are being added and no images are removed. We run each reconstruction three times, and report the aver-

How the ground truth labels were obtained?

Case study: image matching

- Is assumption validated?




Fig. 7 Histograms of pose differences between reconstructions with different local feature methods. We consider five different local features – including rotation-sensitive and upright SIFT – resulting in 10 combinations. The plots show that about 80% percent of image pairs are within a 0.2° pose difference, with the exception of those involving SuperPoint.

How the ground truth labels were obtained?

Case study: image matching


- Is assumption validated?

Local featured type	Number of Images				
	100	200	400	800	all
SIFT [54]	0.06°	0.09°	0.06°	0.07°	0.09°
SIFT (Upright) [54]	0.07°	0.07°	0.04°	0.06°	0.09°
HardNet (Upright) [62]	0.06°	0.06°	0.06°	0.04°	0.05°
SuperPoint [34]	0.31°	0.25°	0.33°	0.19°	0.32°
R2D2 [80]	0.12°	0.08°	0.07°	0.08°	0.05°

Table 2 Standard deviation of the pose difference of three COLMAP runs with different number of images. Most of them are below 0.1°, except for SuperPoint.

It seems that we can trust the paper...



- Although we never can be sure



You should never be 100% sure

I am not joking

- On the Limits of Pseudo Ground Truth in Visual Camera Re-localisation. Brachmann et.al ICCV 2021



Paper is relevant and seems trustworthy



https://www.youtube.com/watch?v=_MeY1_ph2P4



Source

Checking results and methods

Start with baseline

- Find the method, which is the established “gold standard” and/or the baseline
- For image matching it is SIFT
- Check how it works
 - Note, that the “gold standard” might be complex and highly engineered
 - New methods might be simpler to implement
- Play selected (good and simple) methods
 - Maybe add their implementation to your ToDo




Fig. 17 Test – Stereo performance: 2k vs 8k features. We compare the results obtained with different methods using either 2k or 8k features – we use DEGENSAC, which performs better than other RANSAC variants under most circumstances. Dashed lines indicate SIFT’s performance. For LF-Net and SuperPoint we do not include results with 8k features, as we failed to obtain meaningful results. For R2D2, we use the best model for each setting.

How to “play with method”?

- First, try to make it work
 - Visualize, visualize, visualize
- Once it works, try to break it. Noise, transforms, data, etc.
 - that it is how understand its limits.
- Try to understand the ideas it is based on
- Only then try to implement it on your own

Playing with method

Case study


- I was working on some image-matching method with co-authors
- They sent me the training code, I had to plug it into the benchmark
- Instead I first wrote simple script to visualize keypoints

Playing with method


Case study

They were expected to look similar....

Model 1



Model 2



Playing with method

Case study

- If it is possible, gather some data yourself, not from benchmarks
- Make it easy and fun to run the method on your data
- You can even build web demo for it
- If it is not fun, it is not “playing”
- You can think of this as “micro-benchmark”

Playing with method



Dmytro Mishkin
@ducha_aiki



R2D2 fails here, others are fine
9/

[Перекласти твіт](#)



3:47 пп · 22 лист. 2020 · Twitter Web App

https://twitter.com/ducha_aiki/status/1330495426865344515




Dmytro Mishkin
@ducha_aiki



SIFT-HardNet and R2D2 fail here, SuperPoint rocks.
10/10

Overall, for random unknown upright pair of images with possible illumination change, I would go for SuperPoint, as least if we don't consider affine view synthesis.

[Перекласти твіт](#)



4:00 пп · 22 лист. 2020 · Twitter Web App

You know area well enough and want to keep yourself up-to-date

- You want to maximize useful information flow
- And minimize brain fuel & time spent

Where to look for papers on arXiv

arXiv itself?

- arxiv.org

Computer Science

- [Computing Research Repository \(CoRR\)](#) ([new](#), [recent](#), [search](#))
includes (see [detailed description](#)): [Artificial Intelligence](#); [Computation and Language](#); [Computational Complexity](#); [Computational Engineering, Finance, and Science](#); [Computational Geometry](#); [Computer Science and Game Theory](#); [Computer Vision and Pattern Recognition](#); [Computers and Society](#); [Cryptography and Security](#); [Data Structures and Algorithms](#); [Databases](#); [Digital Libraries](#); [Discrete Mathematics](#); [Distributed, Parallel, and Cluster Computing](#); [Emerging Technologies](#); [Formal Languages and Automata Theory](#); [General Literature](#); [Graphics](#); [Hardware Architecture](#); [Human-Computer Interaction](#); [Information Retrieval](#); [Information Theory](#); [Logic in Computer Science](#); [Machine Learning](#); [Mathematical Software](#); [Multiagent Systems](#); [Multimedia](#); [Networking and Internet Architecture](#); [Neural and Evolutionary Computing](#); [Numerical Analysis](#); [Operating Systems](#); [Other Computer Science](#); [Performance](#); [Programming Languages](#); [Robotics](#); [Social and Information Networks](#); [Software Engineering](#); [Sound](#); [Symbolic Computation](#); [Systems and Control](#)

Quantitative Biology

- [Quantitative Biology \(q-bio\)](#) ([new](#), [recent](#), [search](#))
includes (see [detailed description](#)): [Biomolecules](#); [Cell Behavior](#); [Genomics](#); [Molecular Networks](#); [Neurons and Cognition](#); [Other Quantitative Biology](#); [Populations and Evolution](#); [Quantitative Methods](#); [Subcellular Processes](#); [Tissues and Organs](#)

Quantitative Finance

- [Quantitative Finance \(q-fin\)](#) ([new](#), [recent](#), [search](#))
includes (see [detailed description](#)): [Computational Finance](#); [Economics](#); [General Finance](#); [Mathematical Finance](#); [Portfolio Management](#); [Pricing of Securities](#); [Risk Management](#); [Statistical Finance](#); [Trading and Market Microstructure](#)

Statistics

- [Statistics \(stat\)](#) ([new](#), [recent](#), [search](#))
includes (see [detailed description](#)): [Applications](#); [Computation](#); [Machine Learning](#); [Methodology](#); [Other Statistics](#); [Statistics Theory](#)

- CS.Computer Vision and Pattern Recognition
- CS.Artificial Intelligence
- CS.Neural and Evolutionary Computing
- CS.Machine Learning
- CS.Robotics
- Stat.Machine Learning

What you can do on arXiv.org

arXiv.org > cs > cs.AI

Search... All fields Search

Help | Advanced Search

Artificial Intelligence

Authors and titles for recent submissions

- [Tue, 25 Jan 2022](#)
- [Mon, 24 Jan 2022](#)
- [Fri, 21 Jan 2022](#)
- [Thu, 20 Jan 2022](#)
- [Wed, 19 Jan 2022](#)

[total of 331 entries: 1-25 | [26-50](#) | [51-75](#) | [76-100](#) | ... | [326-331](#)]
[showing 25 entries per page: [fewer](#) | [more](#) | [all](#)]

Tue, 25 Jan 2022 (showing first 25 of 79 entries)

[1] [arXiv:2201.09760](#) [[pdf](#), [other](#)]
Multi-Graph Fusion Networks for Urban Region Embedding
[Shangbin Wu](#), [Xu Yan](#), [Xiaoliang Fan](#), [Shirui Pan](#), [Shichao Zhu](#), [Chuanpan Zheng](#), [Ming Cheng](#), [Cheng Wang](#)
Subjects: Artificial Intelligence (cs.AI)

[2] [arXiv:2201.09708](#) [[pdf](#), [other](#)]
Towards Collaborative Question Answering: A Preliminary Study
[Xiangkun Hu](#), [Hang Yan](#), [Qipeng Guo](#), [Xipeng Qiu](#), [Weinan Zhang](#), [Jing Zhang](#)
Subjects: Artificial Intelligence (cs.AI); Computation and Language (cs.CL)

[3] [arXiv:2201.09694](#) [[pdf](#), [other](#)]
Scaling Up Knowledge Graph Creation to Large and Heterogeneous Data Sources
[Enrique Iglesias](#), [Samaneh Jozashoori](#), [Maria-Esther Vidal](#)
Subjects: Artificial Intelligence (cs.AI); Databases (cs.DB)

Dates, number of entries per page

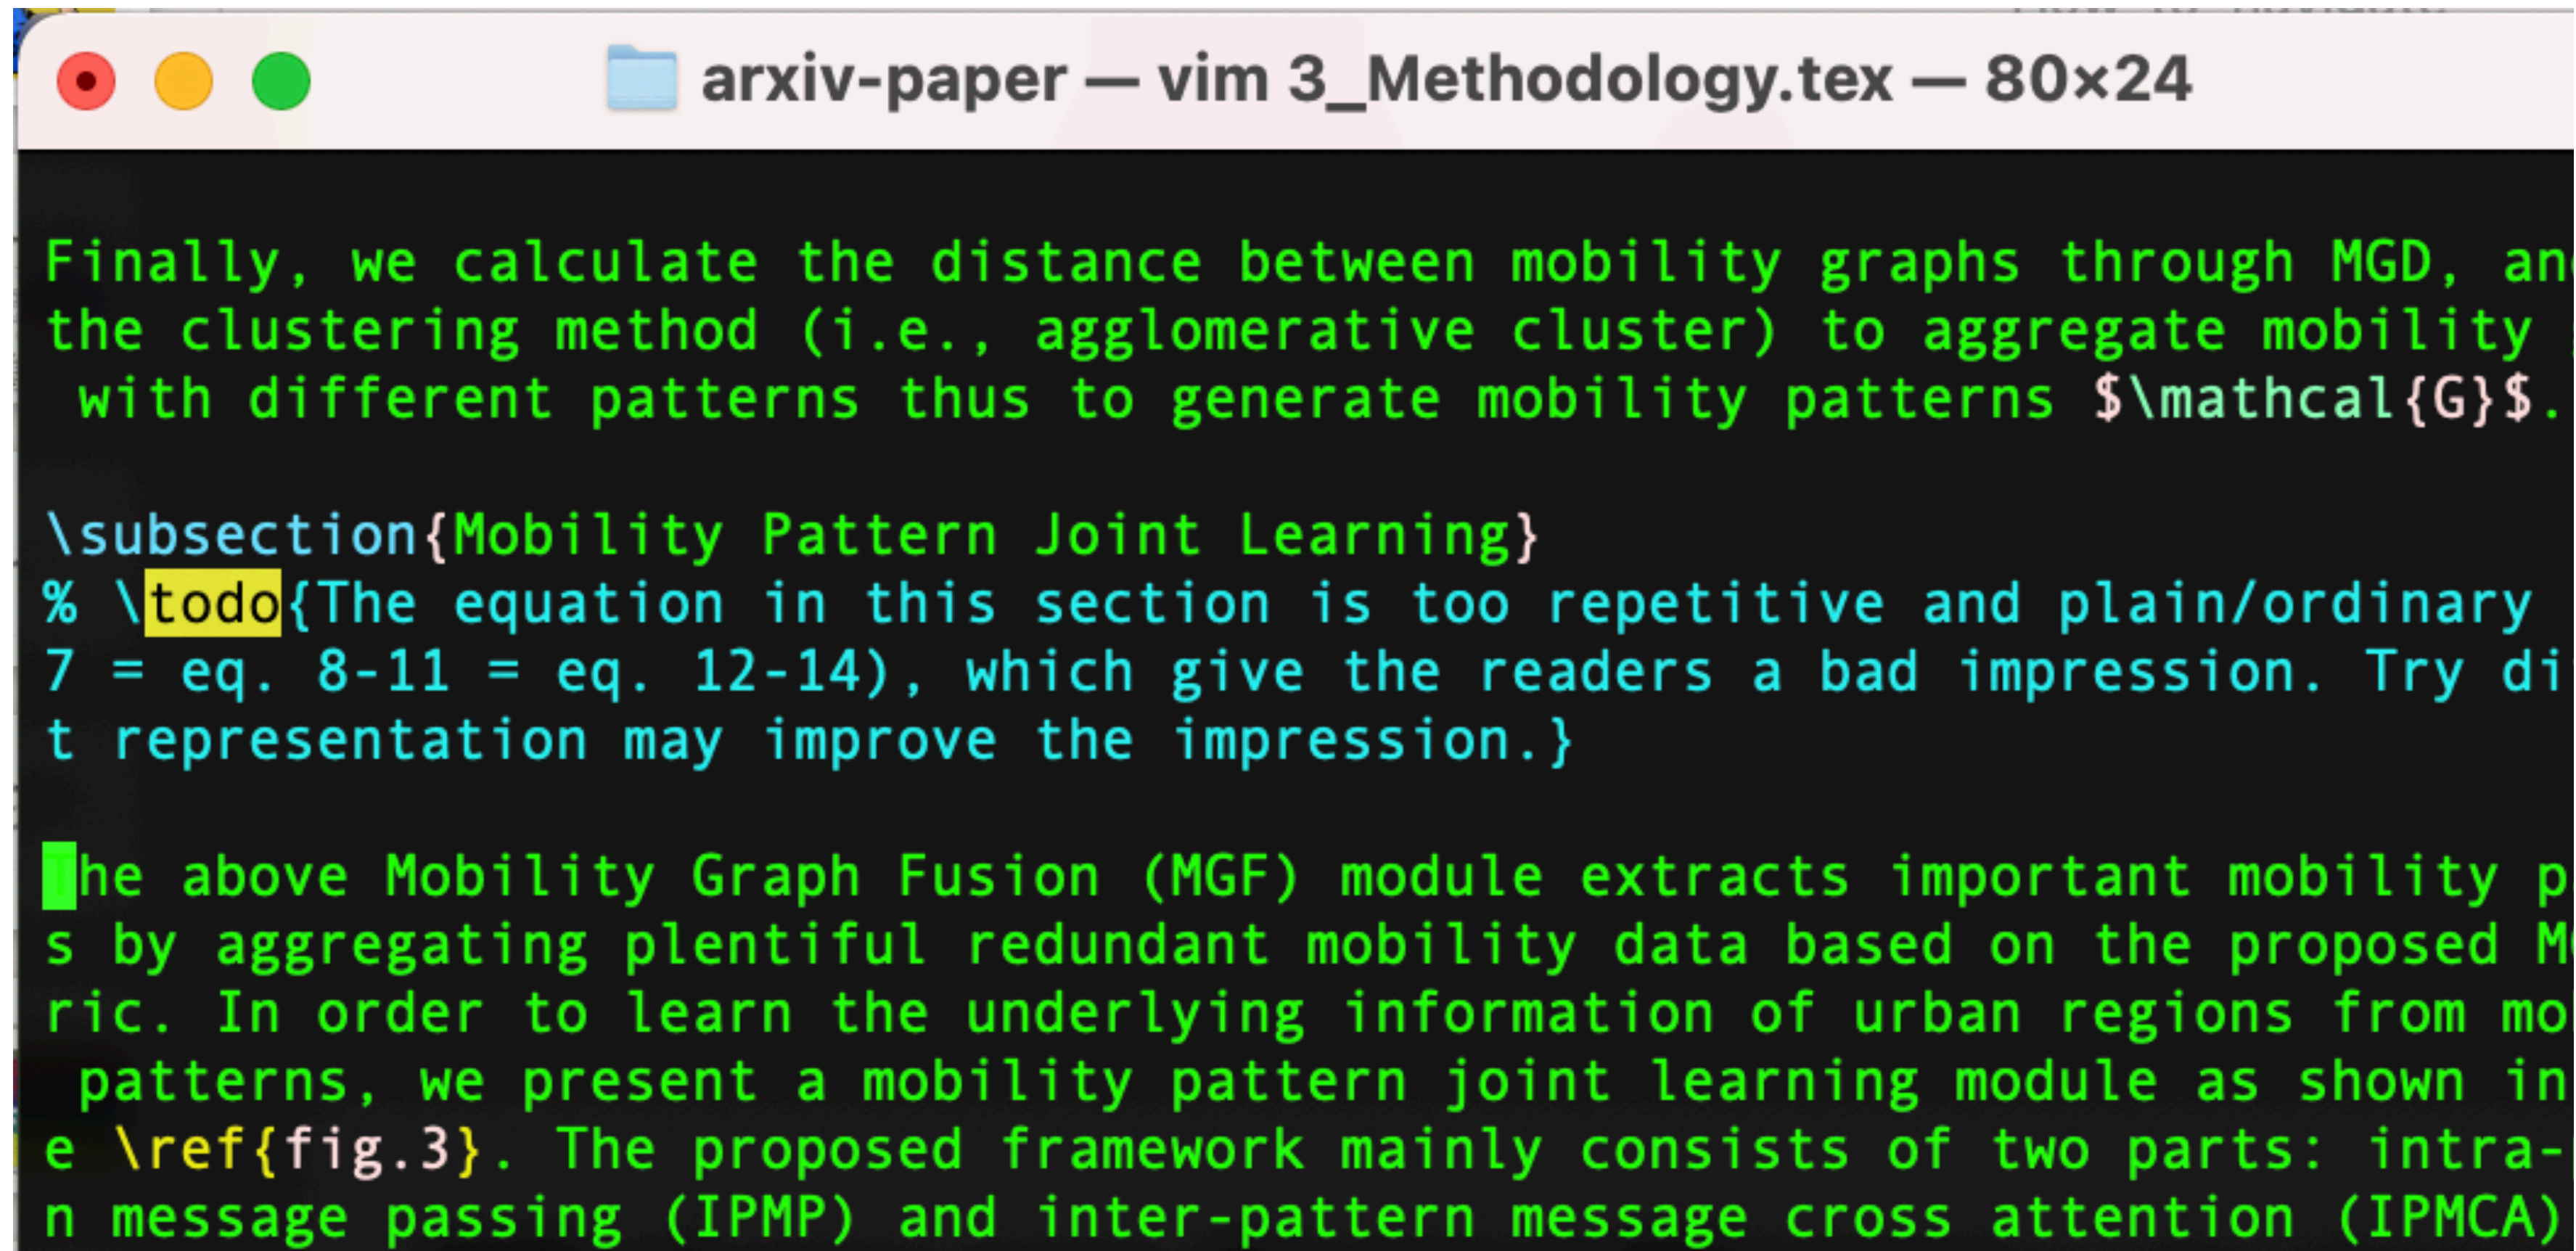
Some papers appear on several lists

You can see only title and authors

“other” means latex source, useful for copying tables, images, etc

What you can do on arXiv.org

“other” format

A screenshot of a vim editor window. The title bar shows 'arxiv-paper — vim 3_Methodology.tex — 80x24'. The editor content shows LaTeX code. The first paragraph describes calculating distance between mobility graphs using MGD and clustering. The second paragraph is a subsection titled 'Mobility Pattern Joint Learning' with a todo comment. The third paragraph describes the Mobility Graph Fusion (MGF) module.

```
Finally, we calculate the distance between mobility graphs through MGD, and the clustering method (i.e., agglomerative cluster) to aggregate mobility with different patterns thus to generate mobility patterns  $\mathcal{G}$ .
```

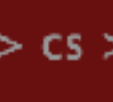
```
\subsection{Mobility Pattern Joint Learning}  
% \todo{The equation in this section is too repetitive and plain/ordinary  
7 = eq. 8-11 = eq. 12-14), which give the readers a bad impression. Try di  
t representation may improve the impression.}
```

```
The above Mobility Graph Fusion (MGF) module extracts important mobility p  
s by aggregating plentiful redundant mobility data based on the proposed M  
ric. In order to learn the underlying information of urban regions from mo  
patterns, we present a mobility pattern joint learning module as shown in  
e \ref{fig.3}. The proposed framework mainly consists of two parts: intra-  
n message passing (IPMP) and inter-pattern message cross attention (IPMCA)
```

- arXiv downloads “other” as a file without extension
- it is actually tar.gz file
- It will unpack to the current dir, so better to create a new dir
- You can find interesting points and todos there, if authors haven’t deleted them
- my todo: about arxiv-cleaner

What you can do on arXiv.org

Citations



Cornell University

We gratefully acknowledge support from
the Simons Foundation and member institutions.

arXiv.org > cs > arXiv:2201.09760

Search...

All fields

Search

Help | Advanced Search

Computer Science > Artificial Intelligence

[Submitted on 24 Jan 2022]

Multi-Graph Fusion Networks for Urban Region Embedding

Shangbin Wu and Xu Yan and Xiaoliang Fan and Shirui Pan and Shichao Zhu and Chuanpan Zheng and Ming Cheng and Cheng Wang,

Learning to fuse multi-modal information to reveal the spatio-temporal domain tasks such as multi-level spatio-temporal module. The multi-level module is trained from multi-modal messages. Experimental results demonstrate that the proposed MGRN outperforms the state-of-the-art methods by up to 12.35% improvement.

Download:

PDF

Other formats (license)

Current browse context:

cs.LG > Multi-modal Learning

2021 > January

Submitted by: Shangbin Wu

Citations

arXiv:2201.09760v1 [cs.AI]

Citation

Bibtex formatted citation

```
@misc{wu2022multigraph,
  title={Multi-Graph Fusion Networks for Urban Region Embedding},
  author={Shangbin Wu and Xu Yan and Xiaoliang Fan and Shirui Pan and Shichao Zhu and Chuanpan Zheng and Ming Cheng and Cheng Wang},
  year={2022},
  eprint={2201.09760},
  archivePrefix={arXiv},
  primaryClass={cs.AI}
}
```

Data provided by: [arXiv API](#)

Subjects: Artificial Intelligence (cs.AI)

Cite as: [arXiv:2201.09760](#) [cs.AI]

(or [arXiv:2201.09760v1](#) [cs.AI] for this version)

Submission history

From: Shanbin Wu [[view email](#)]

[v1] Mon, 24 Jan 2022 15:48:50 UTC (1,281 KB)

Where to look for papers on arXiv

arxiv-sanity.com

The screenshot shows the Arxiv Sanity Preserver website. At the top, a dark red header contains the site name, a description, a user login box with the email 'ducha.aiki@gmail.com' and a 'log out' button, and a 'Fork me on GitHub' link. Below the header is a search bar and a navigation menu with tabs: 'most recent', 'top recent', 'top hype', 'friends', 'discussions', 'recommended', and 'library'. A 'Only show v1' button is also present. The main content area displays 'Showing most recent Arxiv papers:' followed by a paper entry for 'CVAE-H: Conditionalizing Variational Autoencoders via Hypernetworks and Trajectory Forecasting for Autonomous Driving' by Geunseob Oh and Huei Peng, dated 1/24/2022. The paper's abstract is visible at the bottom of the screenshot, enclosed in a red oval. The abstract discusses the challenge of predicting stochastic behaviors of road agents in diverse environments for autonomous driving and introduces the CVAE-H model.

Arxiv Sanity Preserver
Built in spare time by @karpathy to accelerate research.
Serving last 185571 papers from cs.[CV|CL|LG|AI|NE]/stat.ML
(New: Nov 27, 2021) Have a look at the new arxiv-sanity-lite.com

[ducha.aiki@gmail.com](#) [log out](#) [Fork me on GitHub](#)

Search:

[most recent](#) [top recent](#) [top hype](#) [friends](#) [discussions](#) [recommended](#) [library](#)

[Only show v1](#)

Showing most recent Arxiv papers:

CVAE-H: Conditionalizing Variational Autoencoders via Hypernetworks and Trajectory Forecasting for Autonomous Driving
Geunseob Oh, Huei Peng
1/24/2022 [cs.LG](#) | [cs.AI](#) | [cs.RO](#) [2201.09874v1](#) [pdf](#) [show similar](#) | [discuss](#) [📄](#)

The task of predicting stochastic behaviors of road agents in diverse environments is a challenging problem for autonomous driving. To best understand scene contexts and produce diverse possible future states of the road agents adaptively in different environments, a prediction model should be probabilistic, multi-modal, context-driven, and general. We present Conditionalizing Variational AutoEncoders via Hypernetworks (CVAE-H); a conditional VAE that extensively leverages hypernetwork and performs generative tasks for high-dimensional problems like the prediction task. We first evaluate CVAE-H on simple generative experiments to show that CVAE-H is probabilistic, multi-modal, context-driven, and general. Then, we demonstrate that the proposed model effectively solves a self-driving prediction problem by producing accurate predictions of road agents in various environments.

- Developed by Andrew Karpathy
- Has uber-feature: page thumbnails
- Also shows paper abstract

Where to look for papers on arXiv

arxiv-sanity.com

The screenshot shows the Arxiv Sanity Preserver website. At the top, there's a red header with the text "Arxiv Sanity Preserver" and a description: "Built in spare time by @karpathy to accelerate research. Serving last 185571 papers from cs.[CV|CL|LG|AI|NE]/stat.ML (New: Nov 27, 2021) Have a look at the new arxiv-sanity-lite.com". There's also a user email "ducha.aiki@gmail.com" and a "log out" button. A diagonal banner on the right says "Fork me on GitHub". Below the header is a search bar. A row of tabs includes "most recent", "top recent", "top hype", "friends", "discussions", "recommended" (which is circled in red), and "library". Below these tabs are filters: "Only show v1", "Last day", "Last 3 days", "Last week", "Last month", "Last year", and "All time". A light blue box states: "Recommended papers: (based on SVM trained on tfidf of papers in your library, refreshed every day or so)". The main content area displays a paper titled "Avoiding Overfitting: A Survey on Regularization Methods for Convolutional Neural Networks" by Claudio Filipi Gonçalves dos Santos and João Paulo Papa, dated 1/10/2022, in the cs.CV and cs.LG categories. It has 27 pages. To the right of the title are links for "2201.03299v1 pdf", "show similar", and "discuss". Below the title is a preview of the paper's content, including a table of contents, a pie chart, and a bar chart. At the bottom, there's a green box with a snippet of the paper's abstract: "Several image processing tasks, such as image classification and object detection, have been significantly improved using Convolutional Neural Networks (CNN). Like ResNet and EfficientNet, many architectures have achieved outstanding results in at least one dataset by the time of their creation. A critical factor in training concerns the network's regularization, which prevents the structure from overfitting. This work analyzes several regularization methods developed in the last few years, showing significant improvements for different CNN models. The works are classified into three main areas: the first one is called 'data augmentation', where all the techniques focus on performing changes in the input data. The second, named 'internal changes', which aims to describe procedures to modify the feature maps generated by the neural network or the kernels. The last one, called 'label', concerns transforming the labels of a given input. This work presents two main differences comparing to other available surveys about regularization: (i) the first concerns the papers gathered in the manuscript, which are not older than five years, and (ii) the second distinction is about reproducibility, i.e. all works referred here have their code available in public repositories or they have been directly implemented in some framework."

- You can save paper to the library (only 500 last)
- Based on papers you save, arxiv-sanity recommends you new papers

The main problem with arxiv-sanity

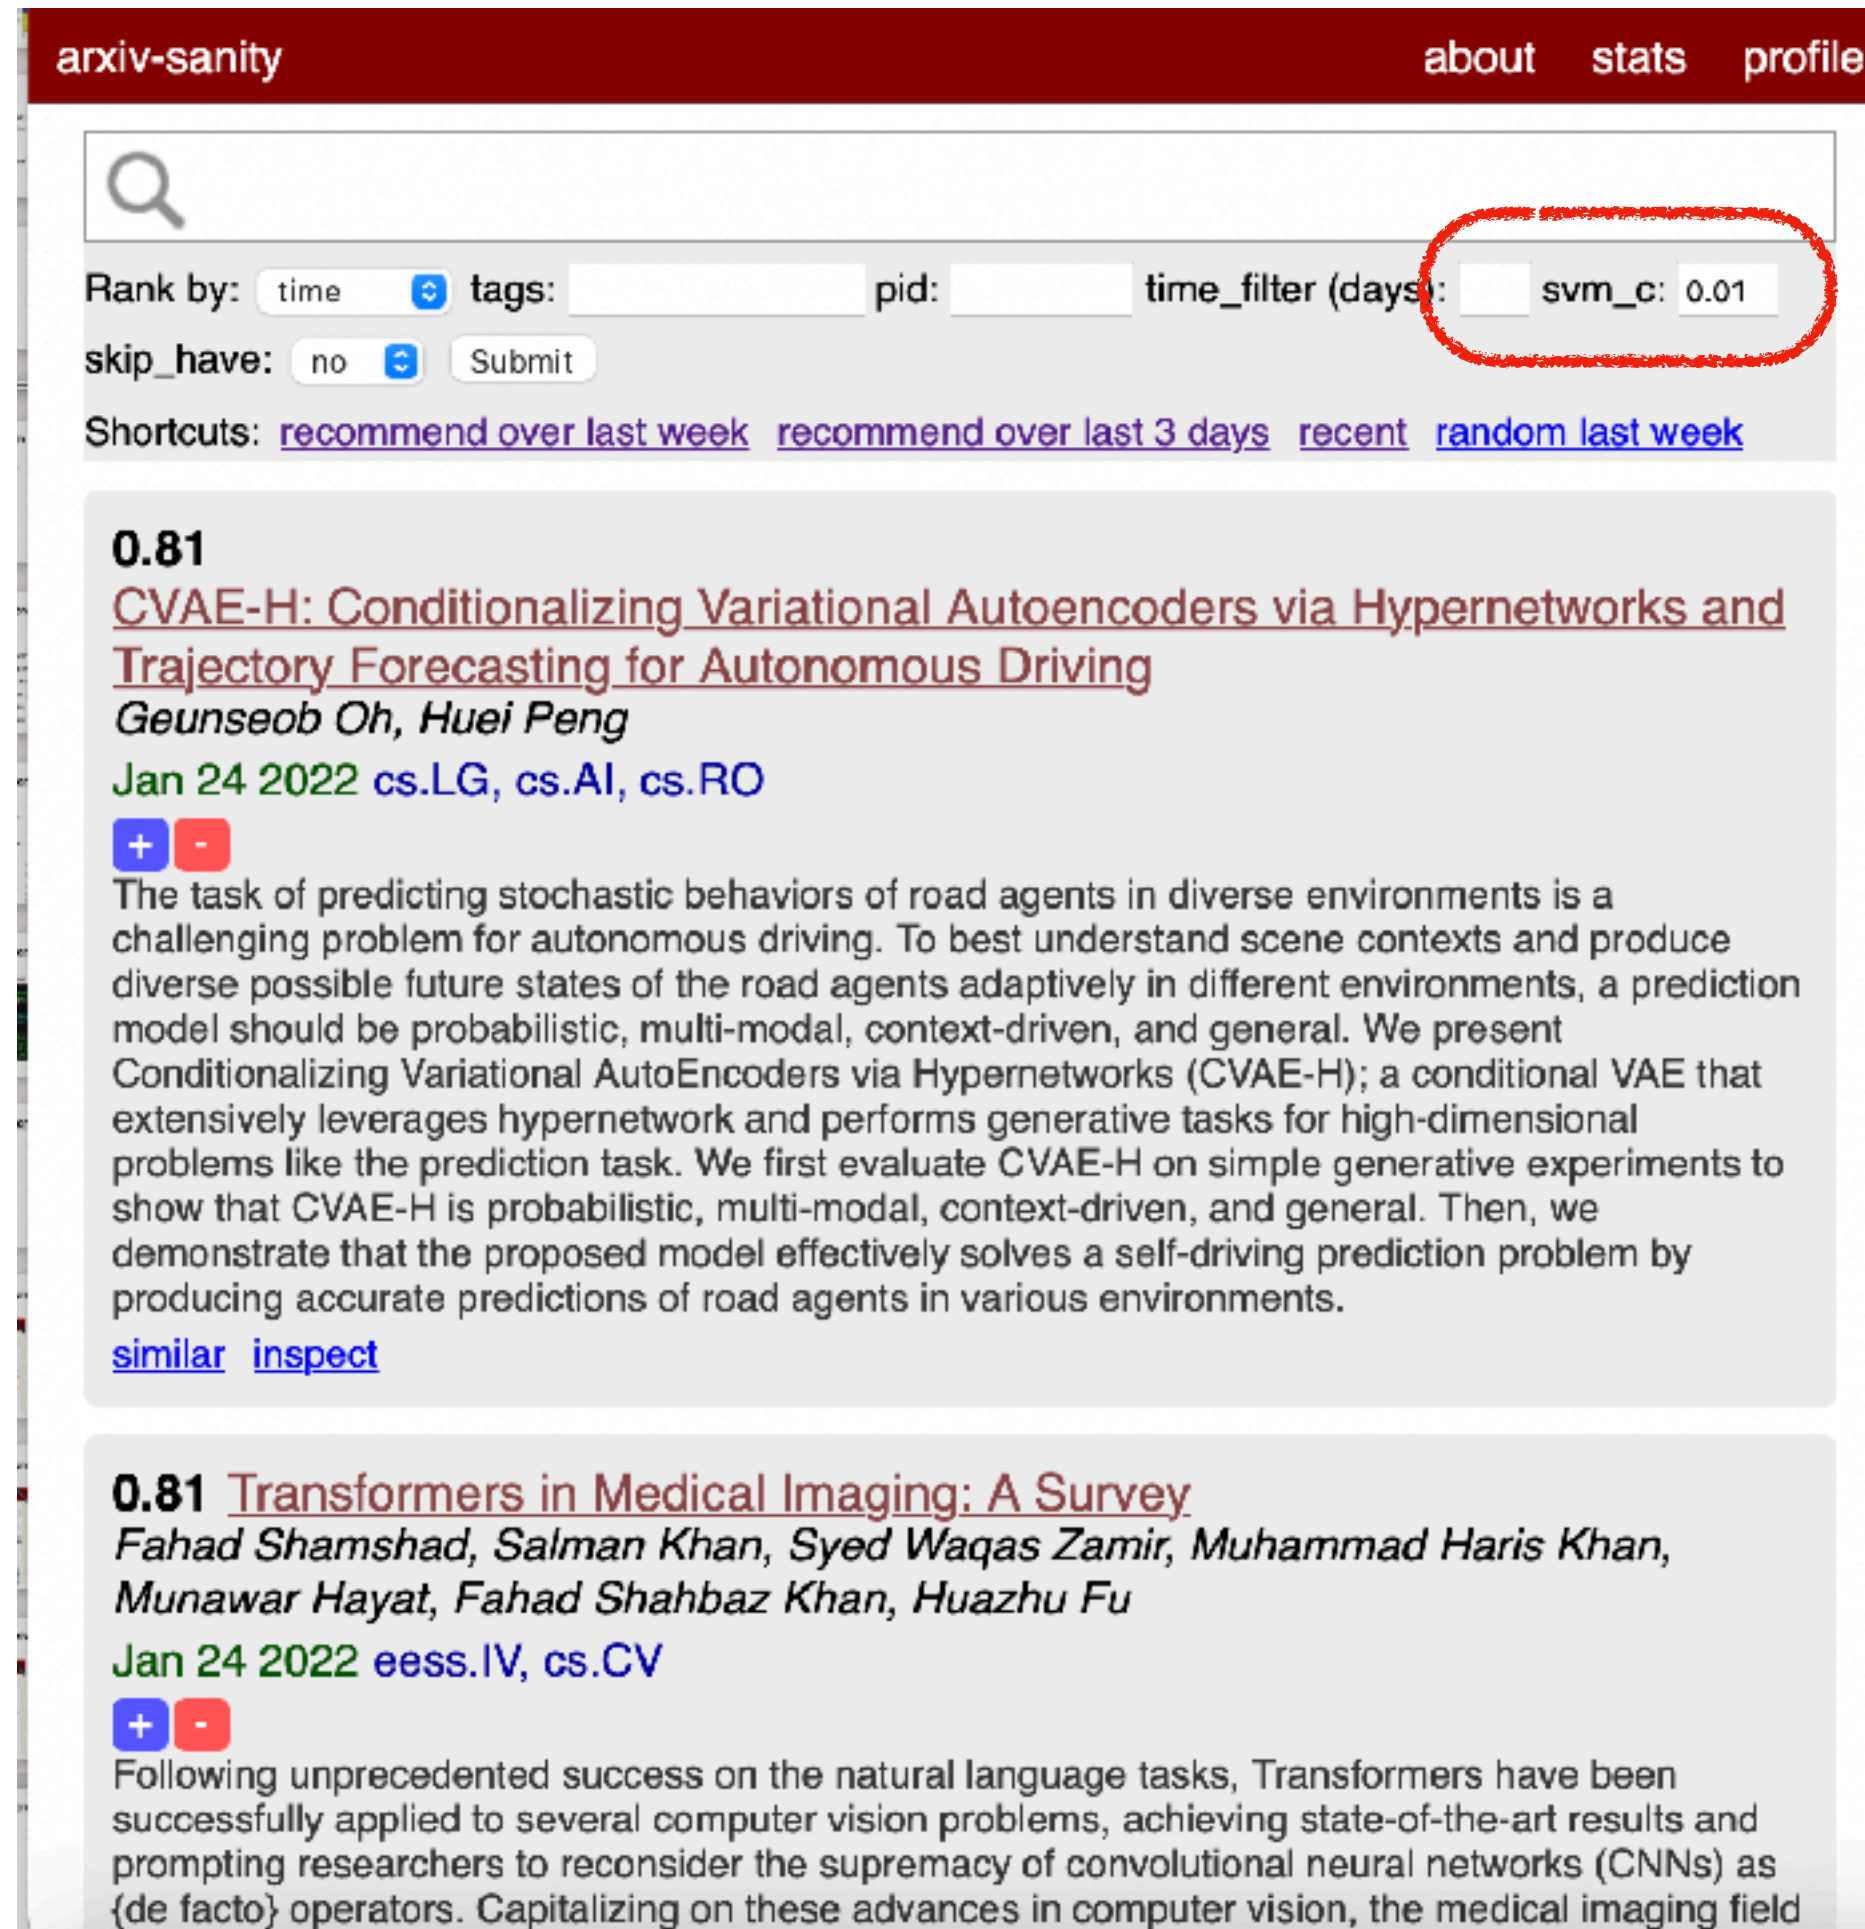
It is often offline



It is often offline :(

Where to look for papers on arXiv

arxiv-sanity-lite.com



- Re-vamp of arxiv-sanity
- No page thumbnails
- You can control recommendations from tf-idf SVM
- “Relevancy score”
- “Inspect” the terms used for relevancy score calculation

Where to look for papers on arXiv

arxiv-sanity-lite.com

arxiv-sanity

[about](#)

[stats](#)

[profile](#)

CVAE-H: Conditionalizing Variational Autoencoders via Hypernetworks and Trajectory Forecasting for Autonomous Driving

Geunseob Oh, Hui Peng

Jan 24 2022 [cs.LG](#), [cs.AI](#), [cs.RO](#)

The task of predicting stochastic behaviors of road agents in diverse environments is a challenging problem for autonomous driving. To best understand scene contexts and produce diverse possible future states of the road agents adaptively in different environments, a prediction model should be probabilistic, multi-modal, context-driven, and general. We present Conditionalizing Variational AutoEncoders via Hypernetworks (CVAE-H); a conditional VAE that extensively leverages hypernetwork and performs generative tasks for high-dimensional problems like the prediction task. We first evaluate CVAE-H on simple generative experiments to show that CVAE-H is probabilistic, multi-modal, context-driven, and general. Then, we demonstrate that the proposed model effectively solves a self-driving prediction problem by producing accurate predictions of road agents in various environments.

The following are the tokens and their (tfidf) weight in the paper vector. This is the actual summary that feeds into the SVM to power recommendations, so hopefully it is good and representative!

0.34 cvae 0.21 road 0.19 variational autoencoders 0.17 driving 0.17 autoencoders 0.16 agents
0.16 multi modal 0.16 autonomous driving 0.15 environments 0.15 trajectory forecasting
0.15 modal 0.15 hypernetwork 0.14 probabilistic 0.14 problems like 0.14 variational
0.14 dimensional problems 0.14 tasks high 0.14 task predicting 0.14 possible future
0.13 model effectively 0.13 prediction problem 0.13 different environments 0.13 autonomous
0.12 diverse 0.12 accurate predictions 0.12 driven 0.12 oh 0.12 generative 0.12 self driving
0.12 prediction task 0.11 prediction model 0.11 context 0.11 solves 0.11 vae 0.11 general
0.10 challenging problem 0.10 producing 0.10 contexts 0.10 adaptively 0.09 extensively
0.09 forecasting 0.09 behaviors 0.09 peng 0.09 demonstrate proposed 0.09 proposed model
0.09 trajectory 0.09 leverages 0.08 conditional 0.08 states 0.08 high dimensional
0.08 understand 0.08 predicting 0.08 scene 0.08 performs 0.08 produce 0.07 stochastic
0.07 predictions 0.07 effectively 0.07 dimensional 0.07 possible 0.07 future 0.07 like
0.07 accurate 0.06 self 0.06 best 0.06 simple 0.06 evaluate

Where to look for papers on arXiv

arxiv-sanity-lite.com

2.00 Learning and Crafting for the Wide Multiple Baseline Stereo

Dmytro Mishkin

Dec 22 2021 cs.CV



This thesis introduces the wide multiple baseline stereo (WxBS) problem. WxBS, a generalization of the standard wide baseline stereo problem, considers the matching of images that simultaneously differ in more than one image acquisition factor such as viewpoint, illumination, sensor type, or where object appearance changes significantly, e.g., over time. A new dataset with the ground truth, evaluation metric and baselines has been introduced. The thesis presents the following improvements of the WxBS pipeline. (i) A loss function, called HardNeg, for learning a local image descriptor that relies on hard negative mining within a mini-batch and on the maximization of the distance between the closest positive and the closest negative patches. (ii) The descriptor trained with the HardNeg loss, called HardNet, is compact and shows state-of-the-art performance in standard matching, patch verification and retrieval benchmarks. (iii) A method for learning the affine shape, orientation, and potentially other parameters related to geometric and appearance properties of local features. (iv) A tentative correspondences generation strategy which generalizes the standard first to second closest distance ratio is presented. The selection strategy, which shows performance superior to the standard method, is applicable to either hard-engineered descriptors like SIFT, LIOP, and MROGH or deeply learned like HardNet. (v) A feedback loop is introduced for the two-view matching problem, resulting in MODS -- matching with on-demand view synthesis -- algorithm. MODS is an algorithm that handles a viewing angle difference even larger than the previous state-of-the-art ASIFT algorithm, without a significant increase of computational cost over "standard" wide and narrow baseline approaches. Last, but not least, a comprehensive benchmark for local features and robust estimation algorithms is introduced.

[similar](#) [inspect](#)

- Similarity search by tf-idf

-69.66 Digging Into Self-Supervised Learning of Feature Descriptors

Iaroslav Melekhov, Zakaria Laskar, Xiaotian Li, Shuzhe Wang, Juho Kannala

Oct 10 2021 cs.CV



Fully-supervised CNN-based approaches for learning local image descriptors have shown remarkable results in a wide range of geometric tasks. However, most of them require per-pixel ground-truth keypoint correspondence data which is difficult to acquire at scale. To address this challenge, recent weakly- and self-supervised methods can learn feature descriptors from relative camera poses or using only synthetic rigid transformations such as homographies. In this work, we focus on understanding the limitations of existing self-supervised approaches and propose a set of improvements that combined lead to powerful feature descriptors. We show that increasing the search space from in-pair to in-batch for hard negative mining brings consistent improvement. To enhance the discriminativeness of feature descriptors, we propose a coarse-to-fine method for mining local hard negatives from a wider search space by using global visual image descriptors. We demonstrate that a combination of synthetic homography transformation, color augmentation, and photorealistic image stylization produces useful representations that are viewpoint and illumination invariant. The feature descriptors learned by the proposed approach perform competitively and surpass their fully- and weakly-supervised counterparts on various geometric benchmarks such as image-based localization, sparse feature matching, and image retrieval.

[similar](#) [inspect](#)

-72.01 SIFT Matching by Context Exposed

Fabio Bellavia

Aug 27 2021 cs.CV



This paper investigates how to step up local image descriptor matching by exploiting matching context information. Two main contexts are identified, originated respectively from the descriptor space and from the keypoint space. The former is generally used to design the actual matching strategy while the latter to filter matches according to the local spatial consistency. On this basis, a new matching strategy and a novel local spatial filter, named respectively blob matching and Delaunay Triangulation Matching (DTM) are devised. Blob matching provides a general matching framework by merging together several strategies, including rank-based pre-filtering as well as many-to-many and symmetric matching, enabling to achieve a global improvement upon each individual strategy. DTM alternates between Delaunay triangulation contractions and expansions to figure out and adjust keypoint neighborhood consistency. Experimental evaluation shows that DTM is comparable or better than the state-of-the-art in terms of matching accuracy and robustness. Evaluation is carried out according to a new benchmark devised for analyzing the matching pipeline in terms of correct correspondences on both planar and non-planar scenes, including several state-of-the-art methods as well as the common SIFT matching approach for reference. This evaluation can be of assistance for future research in this field.

[similar](#) [inspect](#)

How to read arXiv papers

Coarse-to-fine

- I skim through arxiv-sanity list (or [arXiv.org](https://arxiv.org) CS.CV) almost daily
- Check all titles
- Open in the new tab to check the abstract
- Based on abstract open from 1 to 10 pdfs

Should read this paper?

Questions I ask myself (1)

- My areas? Yes → **read**.
- Opens a new (sub-)area of research? Yes → **read**. E.g. GAN in 2014, NERF
- From the crowded area? Yes → **skip**. E.g GANs now
- Dataset or large-scale benchmark paper? Yes → **read**
- Simple baseline? Yes → **read**, regardless of the area.
- About understanding some aspect of machine learning? Yes → **read if have time**. E.g. padding, double descent, over-parametrization

Should read the Questions I ask myself

- Theory paper? Yes → **sk**
- Relevant for me as a use optimizer, etc.
- Isn't the title over-keyworded "Enhanced Generative Acoustic Classification". This kind

The image shows a mobile app interface. The top part is a 'Create New Submission' form with a blue header. It has a 'TITLE AND ABSTRACT' section with a 'Title' input field and a large 'Abstract' text area. Below this is an 'AUTHORS' section with a table. The table has columns for 'Primary Contact', 'Email', 'First Name', and 'Last Name'. There is a red heart icon above the table. At the bottom, there is a navigation bar with icons for home, search, a video player, a shopping bag, and a profile.

Primary Contact	Email	First Name	Last Name
old_ufo	ducha.alki@gmail.com	Dmytro	Mishchenko

me. E.g. new non-linearity,


Adaptive DropBlock
Hyperspectral Image
Verbalize.

Time for live example

Additional paper sources

Yannic Kilcher

Return of the Convolutions



Twitter AK feed — general purpose interesting

@ak92501

AK

@ak92501

paper tweets, dms are open

Перекласти відомості про себе

Дата приєднання: квітень 2014

1 208 Ви читаєте 45,9 тис. Читачі

Читають Hyun Soo Park, Ufuk Efe і ще 134 з-поміж тих, кого ви читаете

Твіти

Твіти й відповіді

Медіафайли

Вподобання

AK @ak92501 · 8 год

AlphaFold Accelerates Artificial Intelligence Powered Drug Discovery: Efficient Discovery of a Novel Cyclin-dependent Kinase 20 (CDK20) Small Molecule Inhibitor

abs: arxiv.org/abs/2201.09647

Figure 1. Our approaches to combine AlphaFold with Insilico medicine end-to-end, AI-powered drug discovery platform PandaOmics and Chemistry42 in the drug discovery for hepatocellular carcinoma from target selection, hit generation to hit identification.



AK @ak92501 · 9 год

Description-Driven Task-Oriented Dialog Modeling

abs: arxiv.org/abs/2201.08904

Model	Pretrain. Model (# Params.)	MW2.1	MW2.2	MW2.3	MW2.4
Transformer-DST (Zeng and Nie, 2021)	BERT Base (110M)	55.35	-	-	-
SOM-DST (Kim et al., 2020)	BERT Base (110M)	51.2	-	55.5	66.8
TripPy (Heck et al., 2020)	BERT Base (110M)	55.3	-	63.0	59.6
SAVN (Wang et al., 2020)	BERT Base (110M)	54.5	-	58.0	60.1
SimpleTOD★ (Hosseini-Asl et al., 2020)	DistilGPT-2 (82M)	50.3/55.7	-	51.3	-
Seq2seq (Zhao et al., 2021)	T5 Base (220M)	52.8	57.6	59.3	67.1
DaP (seq) (Lee et al., 2021a)	T5 Base (220M)	-	51.2	-	-
DaP (ind) (Lee et al., 2021a)	T5 Base (220M)	56.7	57.6	-	-
D3ST (Base)	T5 Base (220M)	54.2	56.1	59.1	72.1
D3ST (Large)	T5 Large (770M)	54.5	54.2	58.6	70.8
D3ST (XXL)	T5 XXL (11B)	57.8	58.7	60.8	75.9

(a) JGA on MultiWOZ 2.1-2.4.

Model	Pretrain. Model (# Params.)	JGA	Intent	Req slot
SGD baseline (Rastogi et al., 2020)	BERT Base (110M)	25.4	90.6	96.5
DaP (ind) (Lee et al., 2021a)	T5 Base (220M)	71.8	90.2	97.8
SGP-DST (Ruan et al., 2020)	T5 Base (220M)	72.2	91.8	99.0
paDST■ (Ma et al., 2020)	XLNet Large (340M)	86.5	94.8	98.5
D3ST (Base)	T5 Base (220M)	72.9	97.2	98.9
D3ST (Large)	T5 Large (770M)	80.0	97.1	99.1
D3ST (XXL)	T5 XXL (11B)	86.4	98.8	99.4

(b) JGA, active intent accuracy and requested slot F1 on SGD.

Table 1: Results on MultiWOZ and SGD datasets with full training data. “-” indicates no public number is available. Best results are marked in bold. ★: SimpleTOD results are retrieved from the 2.3 website <https://github.com/lexmen318/MultiWOZ-coref>, in which two numbers are reported for 2.1 (one produced by the 2.3 author, the other by the original SimpleTOD paper). ♦: No data pre-processing applied for MultiWOZ 2.1. ■: Data augmentation and special rules applied.



13



Twitter my feed: Image matching & hype

@ducha_aiki

 Редагувати профіль

Dmytro Mishkin
@ducha_aiki

Marrying classical CV and Deep Learning. I do things, which work, rather than being novel, but not working.

[Перекласти відомості про себе](#)

📍 Ukraine [dmytro.ai](#) 📅 Дата приєднання: липень 2017

441 Ви читаєте 5 852 Читачі

Твіти Твіти й відповіді Медіафайли Вподобання


Закріплений твіт

Dmytro Mishkin @ducha_aiki · 13 жовт. 2020

ArXiving before submission :


- helps professional identity building;
- protects against idea re-discovery/theft/gate-keeping
- facilitates open research distribution
- reduces inequality.

our refined statement, with [@amy_tabb](#) & Jiri Matas, now on arXiv



arxiv.org
ArXiving Before Submission Helps Everyone
We claim, and present evidence, that allowing arXiv publication before a conference or journal ...

9 44 160

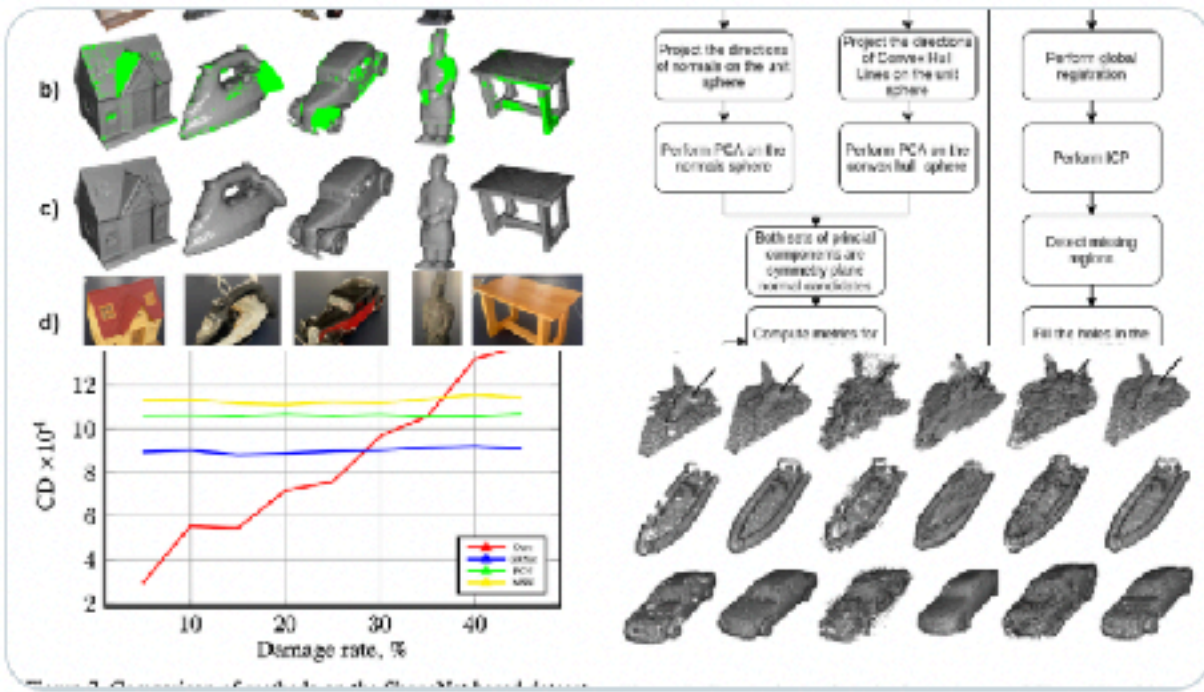
 **Dmytro Mishkin** @ducha_aiki · 7 січ.

Towards realistic symmetry-based completion of previously unseen point clouds

Taras Rumezhak, [@dobosevych](#), Rostyslav Hryniv, Vladyslav Selotkin, Volodymyr Karpiv, Mykola Maksymenko, [@ucu_apps](#)

tl;dr: good handcrafted method for point cloud completion

arxiv.org/abs/2201.01858



1 7 15

Dmytro Mishkin @ducha_aiki · 7 січ.

Bio-inspired Min-Nets Improve the Performance and Robustness of Deep Networks

Philipp Grüning, Erhardt Barth


tl;dr: $\text{out} = \min(\text{ReLU}(\text{convA}(\text{BN}(x))), \text{ReLU}(\text{convB}(\text{BN}(\langle x \rangle)))$
Idea: both filters should fire for reliable detection, kind of opposite of maxout idea.





[arxiv.org/abs/2201.02149...](https://arxiv.org/abs/2201.02149)

**Register on twitter and follow
relevant people**

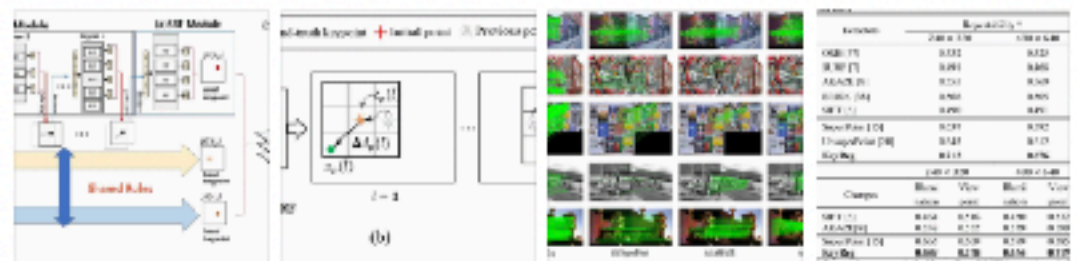
ResearchGate

way to share and request paywalled papers

HomeQuestionsJobs

Add new

Suggested research based on your interests




Self-Supervised Keypoint Detection Based on Multi-Layer Random Forest Regressor

Article · March 2021 · IEEE Access · 90 Reads · 2 Citations

Sangwon Kim · Mira Jeong · Byoungchul Ko

[Download](#) [Save](#) [Follow](#) [Recommend](#) [Share](#)

2 researchers recommend this research item



Fuchao Wu
was cited in a publication

8h ago

Learning Cross-Domain Descriptors for 2D-3D Matching with Hard Triplet Loss and Spatial Transformer Network

Chapter · September 2021 · 30 Reads

Baiqi Lai · Weiquan Liu · Cheng Wang · [...] · Ming Cheng

[Request full-text](#) [Save](#) [Follow](#) [Recommend](#) [Share](#)

We think we found your full-text online

Add your full-text to your profile to get more reads and citations.

[Preprint](#) Learning and Crafting for the Wide Multiple Baseline Stereo

[View full-text](#) [Ask me later](#)

Jobs you may be interested in

University Professor (Salary Scale W3 NBesO) "Computer Engineering" (m/f/d)

[New job](#)

Technische Universität Clausthal

Graduate teaching & research assistant, Graph neural networks

[New job](#)

University of Antwerp

Research Officer in Behavioural and Attitudinal Change in respect of Ireland's Water Resources

[New job](#)

Economic and Social Research Institute

Google Scholar recommendations



🔍

Рекомендовані статті



☆	Q-ViT: Fully Differentiable Quantization for Vision Transformer Z Li, T Yang, P Wang, J Cheng arXiv preprint arXiv:2201.07703 - 5 днів тому	▼	PDF
☆	Poseur: Direct Human Pose Regression with Transformers W Mao, Y Ge, C Shen, Z Tian, X Wang, Z Wang, A Hengel arXiv preprint arXiv:2201.07412 - 5 днів тому	▼	PDF
Інші статті, опубліковані 5 днів тому			
☆	A graph-matching approach for cross-view registration of over-view and street-view based point clouds X Ling, R Qin ISPRS Journal of Photogrammetry and Remote Sensi... - 6 днів тому	▼	HTML

☰

Сповіщення

Сповіщення для ducha.aiki@gmail.com

Andrej Karpathy – нові статті

"Learning discriminative affine regions via discriminability" – нові бібліографічні посилання

"Mods: Fast and robust method for two-view matching" – нові бібліографічні посилання

How to read a method paper?

The reading order of the method paper

Working hard to know your neighbor's margins: Local descriptor learning loss

- Abstract
- Method figure

Abstract

We introduce a loss for metric learning, which is inspired by the Lowe's matching criterion for SIFT. We show that the proposed loss, that maximizes the distance between the closest positive and closest negative example in the batch, is better than complex regularization methods; it works well for both shallow and deep convolution network architectures. Applying the novel loss to the L2Net CNN architecture results in a compact descriptor named HardNet. It has the same dimensionality as SIFT (128) and shows state-of-art performance in wide baseline stereo, patch verification and instance retrieval benchmarks.

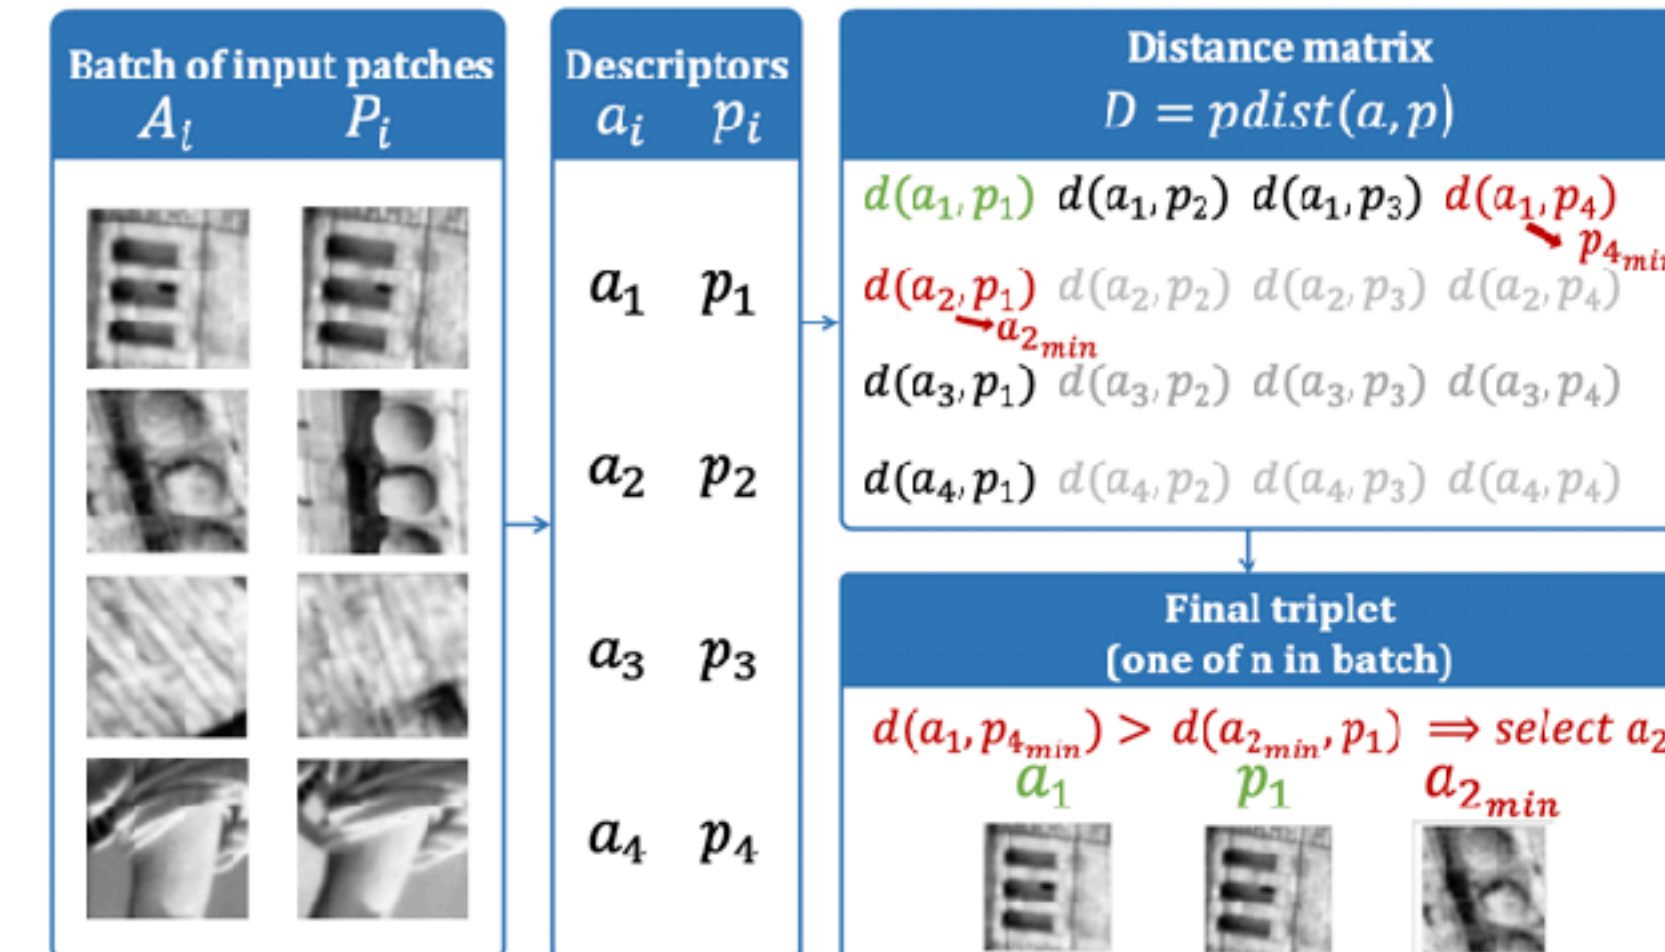


Figure 1: Proposed sampling procedure. First, patches are described by the current network, then a distance matrix is calculated. The closest non-matching descriptor – shown in red – is selected for each a_i and p_i patch from positive pair (green) respectively. Finally, among two negative candidates the hardest one is chosen. All operations are done in a single forward pass.

The reading order of the method paper

Working hard to know your neighbor's margins: Local descriptor learning loss

- Abstract

- Method figure

- Main plot of results

- Conclusion

- If it is not rephrased abstract

- Related work

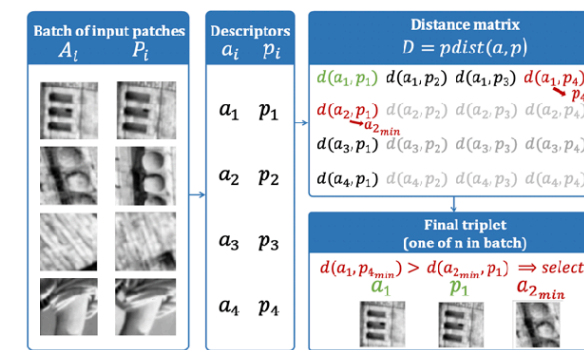


Figure 1: Proposed sampling procedure. First, patches are described by the current network, then a distance matrix is calculated. The closest non-matching descriptor – shown in red – is selected for each a_i and p_i patch from positive pair (green) respectively. Finally, among two negative candidates the hardest one is chosen. All operations are done in a single forward pass.



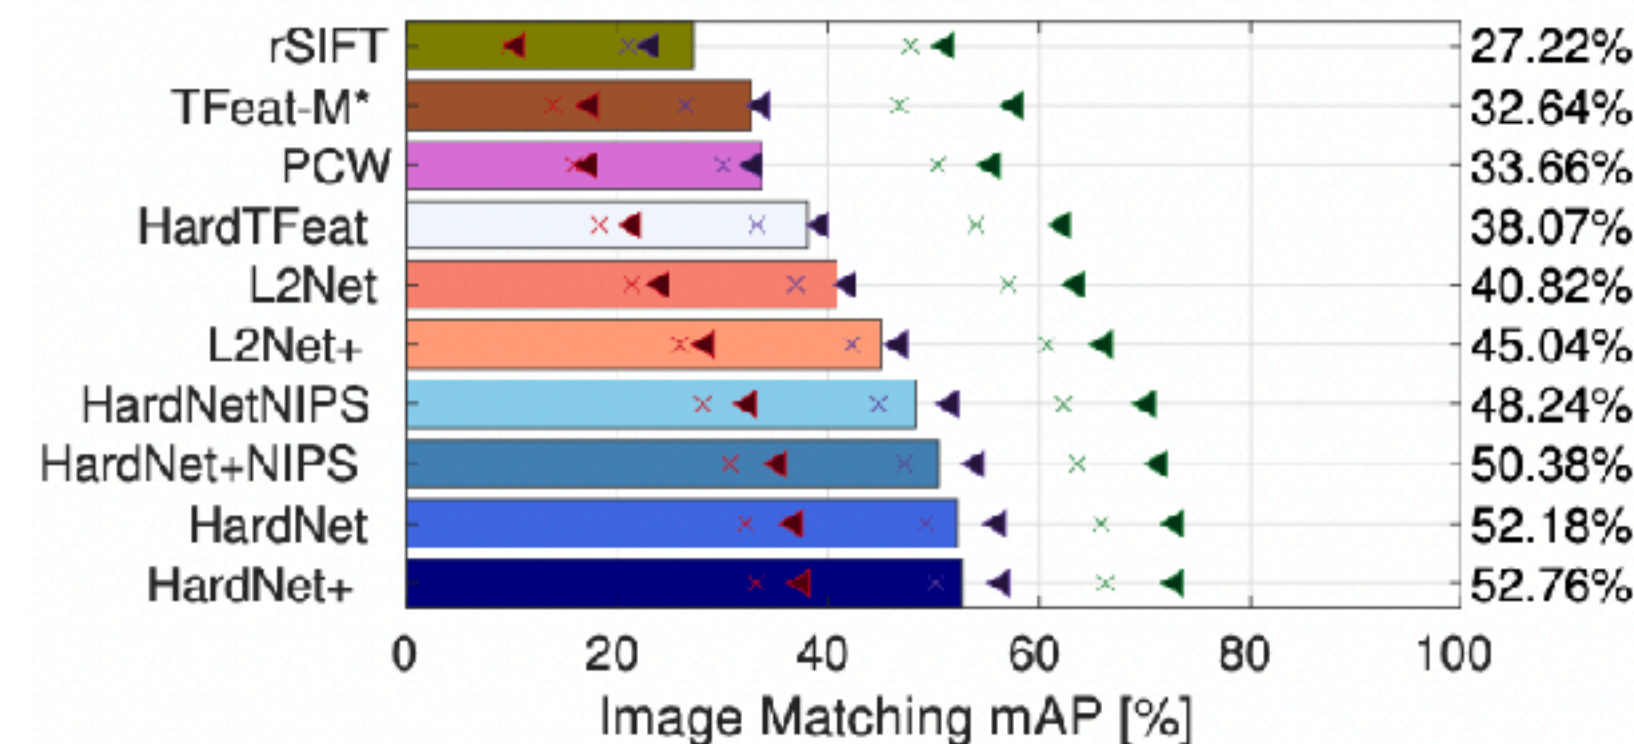


Figure 2: The architecture of our network, adopted from L2Net [24]. Each convolutional layer is followed by batch normalization and ReLU, except the last one. Dropout regularization is used before the last convolution layer.

■ EASY ■ HARD ■ TOUGH

◀ VIEWPT × ILLUM



Advice by Dushyant Mehta

From personal painful experience

1. Don't trust the story the paper tries to sell.
 - supposed intuition is often not true
 - Ask instead, what could be the other reasons for the claimed improvements.
2. Learn to ask yourself about simple and obvious baselines.
 - What should have been the simple and obvious baselines for the paper you are reading. How would you solve the task building from first principles?
 - What experiments would you design to validate the claims of the paper?

Advice by Dushyant Mehta

From personal painful experience

3. Don't take the results at their face value
 - Sure there is a big convenient table with all manner of methods conveniently compared.
 - Are they really comparable? Same data, same compute budget, fair training, fair evaluation?
4. Don't discount the simple methods even if they are slightly worse than overly complex (claimed) SoTA.
 - There is immense advantage to the simple. Less finicky training, less faff in deployment, faster experiment iteration,

Advice by Dushyant Mehta

From personal painful experience

5. Do NOT start building on top of their *method*
 - At least without having determined to your satisfaction that the method is sound
 - Reuse various parts from the codebase, without building on top of the method!
6. Data, data, data! For a lot of the problems, not a lot beats good data.
 - Not mindlessly collecting more data, but knowing the distribution of your data, and sampling smartly from it in a way that reflects your deployment setting.
 - Very simple/subtle sampling differences can cause a huge difference in your metrics, without anything about the method changing.

Advice by Dushyant Mehta

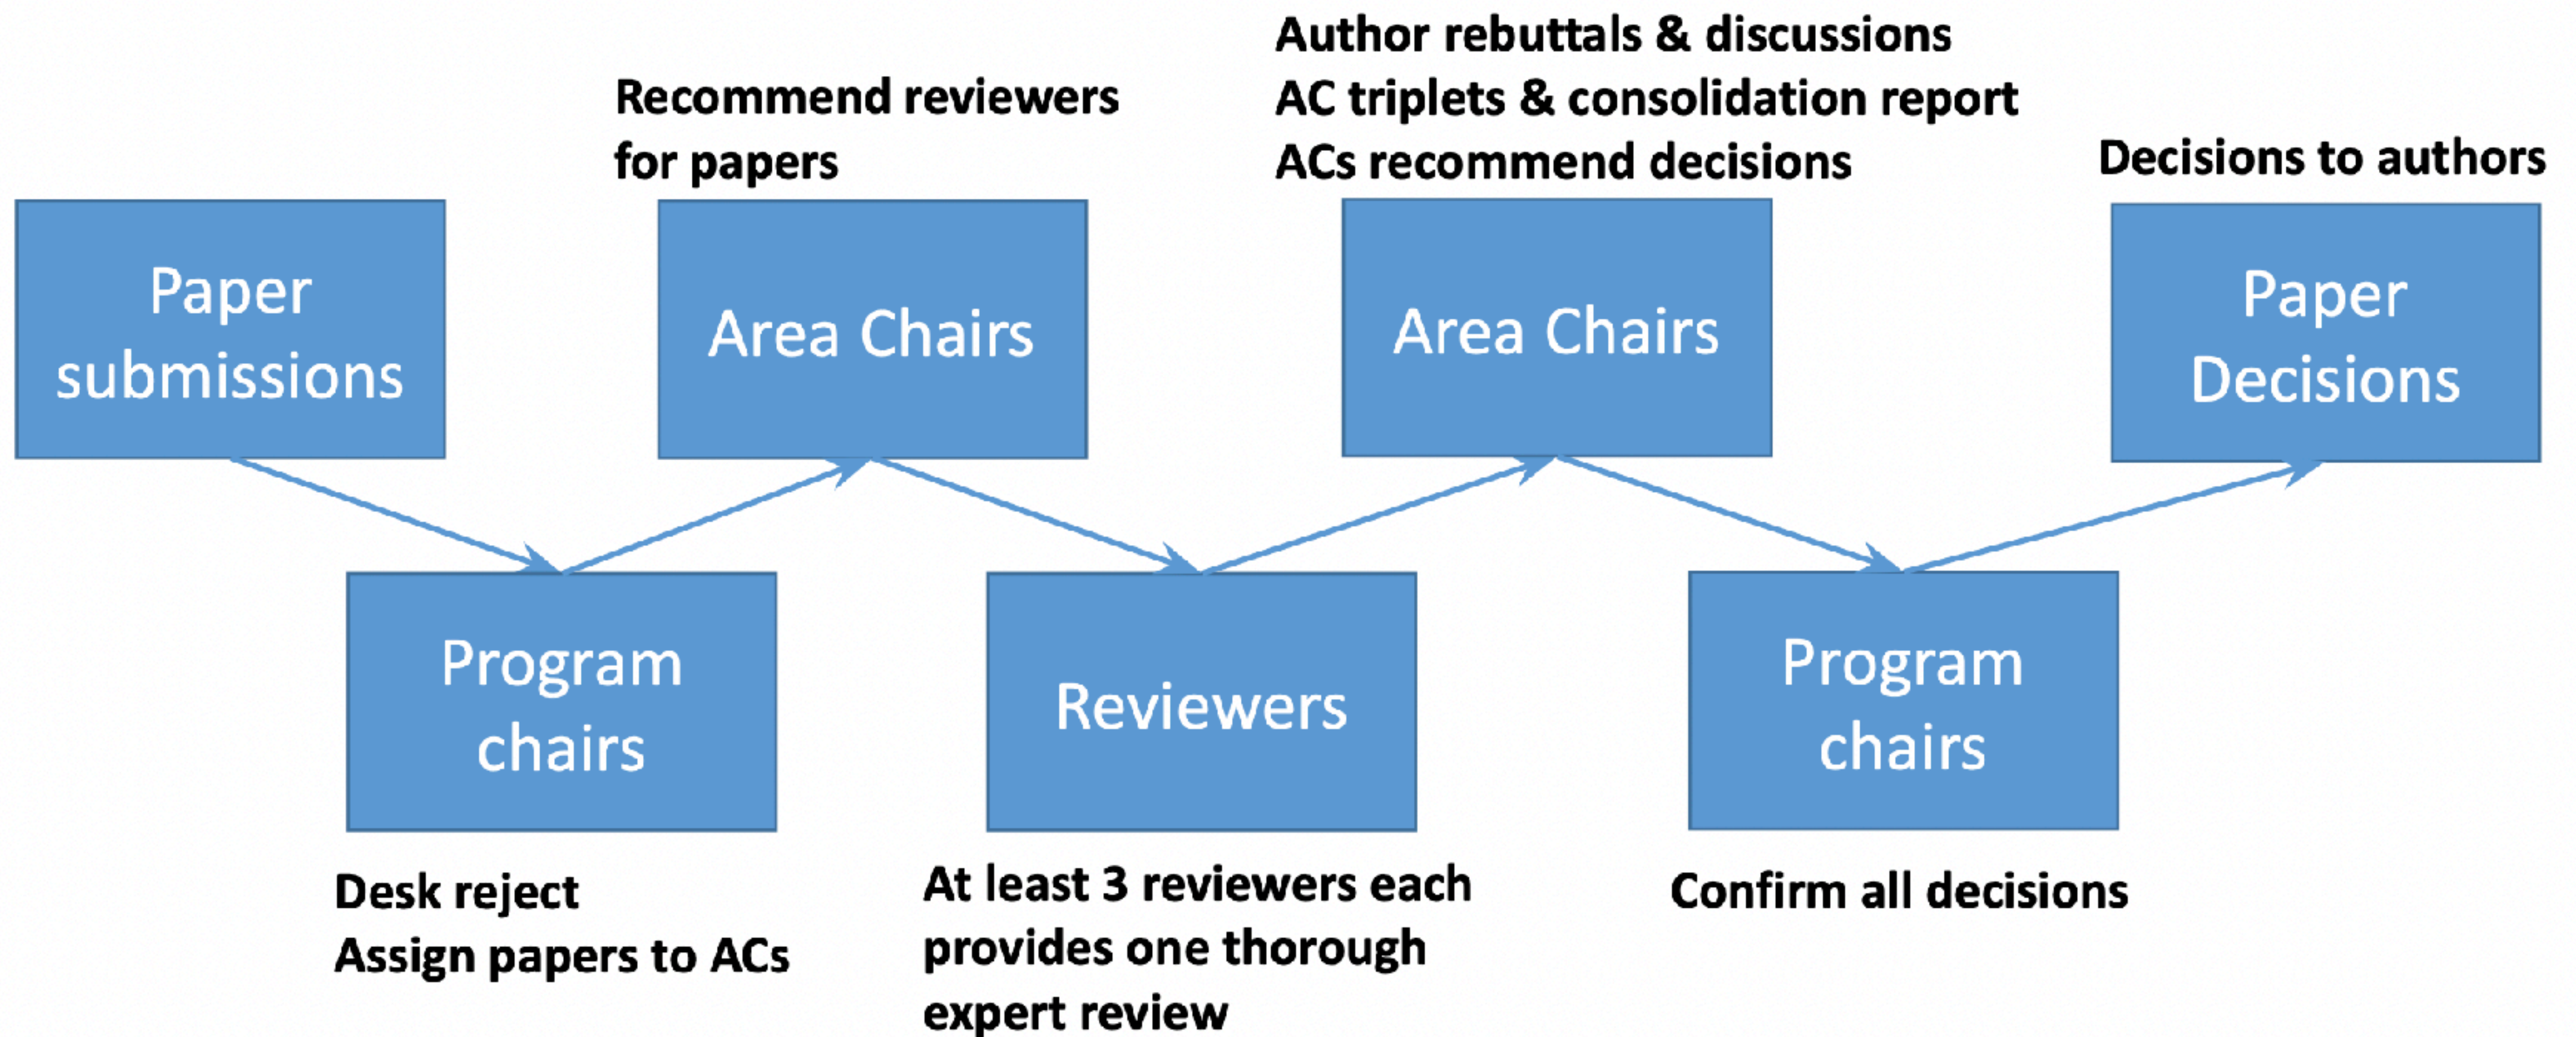
From personal painful experience

7. That brings us to the f**king metrics.

- If you mindlessly optimize for one number, look at leaderboards based on one number, you haven't understood the task.
- What does that one number tell you?
- More importantly, what does it hide from you?
- How relevant is it to the deployment scenarios for that task?
- What metrics would better reflect what matters the most for your deployment scenario?

What does peer review mean?

The process of accepting/rejecting a paper



The role of reviewers in the process

- Provide an independent, objective, critical, and comprehensive review
 - **Key:** What is the knowledge advancement in the paper?
- Discuss with AC and reviewer buddies to (hopefully) reach consensus
 - Explain clearly the basis of your review and recommendation
 - It is OK if the reviewers disagree with one another even after discussions
 - AC will form recommendations weighing in reviews, rebuttals, and discussions
- Make your final recommendations with solid justifications
 - Read the rebuttal and discussions. Do they change your position? Why?
 - This facilitates the ACs to make final recommendations for the paper

What paper should be accepted?

- Any paper that, in accordance with CVPR community standards,
 - presents **sufficient knowledge advancement** that is **well grounded**;
 - is of **sufficient interest** to some **CVPR audiences** who could **benefit** from it
- Note: CVPR is very inclusive
 - Historically rejection solely for out-of-scope is rather rare


What should be included in the review?

- A concise summary of the paper
 - What problem is addressed in the paper?
 - Is it a new problem? If so, why does it matter? If not, why does it still matter?
 - What is the key to the solution? What is the main contribution?
 - Do the experiments sufficiently support the claims?
- A clear statement of strengths and weaknesses
 - What are the key contributions and why do they matter?
 - What aspects of the paper most need improvement?
- A comprehensive check of potential fundamental flaws in the paper
 - Are the assumptions and theories (mathematically) sound?
 - Are the experiments scientifically sound and valid?
 - Is the problem addressed trivial?
 - Did the paper miss important prior work? Has it been done before? If yes, where?

What should be

- Common mistakes in a s
 - Arrogance, ignorance, ar
 - Be responsible!
 - Pure opinions
 - Be grounded!
 - Novelty fallacy
 - Be knowledge-driven!
 - Blank assertions
 - Be substantial!
 - Policy entrepreneurism
 - Be sensible!
 - Intellectual laziness
 - Be active!

e review?



Functions of the conferences

Why authors submit

- **Knowledge dissemination:** “I want people to know about the new knowledge I discovered.” The conference promises a certain minimum level of attention one’s work gets. In other words, if you don’t publish at top conferences, nobody would read your work.
- **Feedback:** I would like people to check my results through the review process and discussion.
- **Formal goodies:** checking boxes, required to defend Ph.D., for tenure package, performance review, to put into grant report, etc.
- **Certification.** Publishing at CVPR is hard, therefore valuable.
- **Reputation-building:** Listing certain conferences on C.V. as a way of building one’s name as a scientist.
- **Networking:** meeting with peers, potential employers, etc

Functions of the conferences

Why people attend

- **Prefiltering:** time is limited, so we outsource the selection of what we are reading to the reviewers.
- **Certification:** time is limited, so we outsource the quality control and result check to the reviewers. We create a basic classifier: “If the paper is published by a top-conference, it is true.”
- **Special case of certification:** for people outside the field without the basic qualifications to select work that meets basic quality guarantees.
- **Authors promise to answer our questions** (symmetrical to “attention for the author from audience,” and audience gets the guarantee that questions about the work will be answered at the talk or poster session).

Thank you for your attention



Dmytro Mishkin, PhD

ducha.aiki@gmail.com

<http://dmytro.ai/>

[ducha_aiki](https://twitter.com/ducha_aiki)

- Computer vision researcher at CTU in Prague.
- Member of the Expert Committee on Artificial Intelligence at Ministry of Digital Transformation Of Ukraine
- Computer vision and deep learning consultant
- Co-author of the first Ukrainian publications at CVPR and NeurIPS
- CVPR 2021 Outstanding reviewer
- Kaggle Master