Notes on CV Paper Writing

Xiang Gao, Lecturer

E-mail: xgao@ouc.edu.cn | Web: https://ouc-xgao.github.io/

College of Engineering, Ocean University of China

About Me

- Xiang Gao, Lecturer
  - Education
    - 2008.09-2012.06: College of Engineering, Ocean University of China, Bachelor
    - 2012.09-2015.06: College of Engineering, Ocean University of China, Master
    - 2015.09-2019.06: Institute of Automation, Chinese Academy of Sciences, Doctor

About Me

- Xiang Gao, Lecturer
  - Education
    - 2008.09-2012.06: College of Engineering, Ocean University of China, Bachelor
    - 2012.09-2015.06: College of Engineering, Ocean University of China, Master
    - 2015.09-2019.06: Institute of Automation, Chinese Academy of Sciences, Doctor
  - Research Interests
    - 3D Computer Vision
    - Large-Scale Structure from Motion
    - Multi-source Data Fusion-Based Large-Scale 3D Reconstruction

About Me

- Xiang Gao, Lecturer
  - Education
    - 2008.09-2012.06: College of Engineering, Ocean University of China, Bachelor
    - 2012.09-2015.06: College of Engineering, Ocean University of China, Master
    - 2015.09-2019.06: Institute of Automation, Chinese Academy of Sciences, Doctor
  - Research Interests
    - 3D Computer Vision
    - Large-Scale Structure from Motion
    - Multi-source Data Fusion-Based Large-Scale 3D Reconstruction
  - E-mail: xgao@ouc.edu.cn
  - Personal Homepage: https://ouc-xgao.github.io/

About Me

About Me

- Xiang Gao, Lecturer
  - Selected Publications
    - **Xiang Gao**, Lingjie Zhu, Zexiao Xie, Hongmin Liu*, and Shuhan Shen*. Incremental Rotation Averaging. *International Journal of Computer Vision (IJCV)*, 2021. (*CCF-A*, IF: 7.410, *h5*-index: 72)
    - **Xiang Gao**, Shuhan Shen*, Yang Zhou, Hainan Cui, Lingjie Zhu, and Zhanyi Hu. Ancient Chinese Architecture 3D Preservation by Merging Ground and Aerial Point Clouds. *ISPRS Journal of Photogrammetry and Remote Sensing (P&RS)*, 2018. (IF: 8.979, *h5*-index: 82)
    - **Xiang Gao**, Lihua Hu, Hainan Cui, Shuhan Shen*, and Zhanyi Hu. Accurate and Efficient Ground-to-Aerial Model Alignment. *Pattern Recognition (PR)*, 2018. (*CCF-B*, IF: 7.740, *h5*-index: 99)
    - **Xiang Gao**, Shuhan Shen*, Lingjie Zhu, Tianxin Shi, Zhiheng Wang, and Zhanyi Hu. Complete Scene Reconstruction by Merging Images and Laser Scans. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 2020. (*CCF-B*, IF: 4.685, *h5*-index: 77)
    - **Xiang Gao**, Lingjie Zhu, Hainan Cui, Zexiao Xie, and Shuhan Shen*. IRA++: Distributed Incremental Rotation Averaging. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 2021. (*CCF-B*, IF: 4.685, *h5*-index: 77)

Several Important Concepts

- Computer Vision (CV)
  - CV is a field of computer science that works on enabling computers to see, identify and process images in the same way that human vision does, and then provide appropriate output

Several Important Concepts

- Computer Vision (CV)
  - CV is a field of computer science that works on enabling computers to see, identify and process images in the same way that human vision does, and then provide appropriate output
  - Highly related areas:
    - Artificial Intelligence (AI)
    - Pattern Recognition (PR)
    - Computer Graphics (CG)
    - Image Processing (IP)
    - Robotics (RO)
    - Remote Sensing (RS)
    - Multimedia (MM)
    - …

Several Important Concepts

- Impact Factor (IF)
  - IF measures the average number of citations received in a particular year by papers in the journal during the two preceding years
    - Homepage: https://jcr.clarivate.com/

Several Important Concepts

- Impact Factor (IF)
  - IF measures the average number of citations received in a particular year by papers in the journal during the two preceding years
    - Homepage: https://jcr.clarivate.com/
  - Take the 2019 IF of IEEE TPAMI for example

Journal Impact Factor Calculation

$$\text{2019 Journal Impact Factor} = \frac{7{,}323}{410} = 17.861$$

How is Journal Impact Factor Calculated?

$$\text{JIF} = \frac{\text{Citations in 2019 to items published in 2017 (4,324) + 2018 (2,999)}}{\text{Number of citable items in 2017 (185) + 2018 (225)}} = \frac{7{,}323}{410}$$

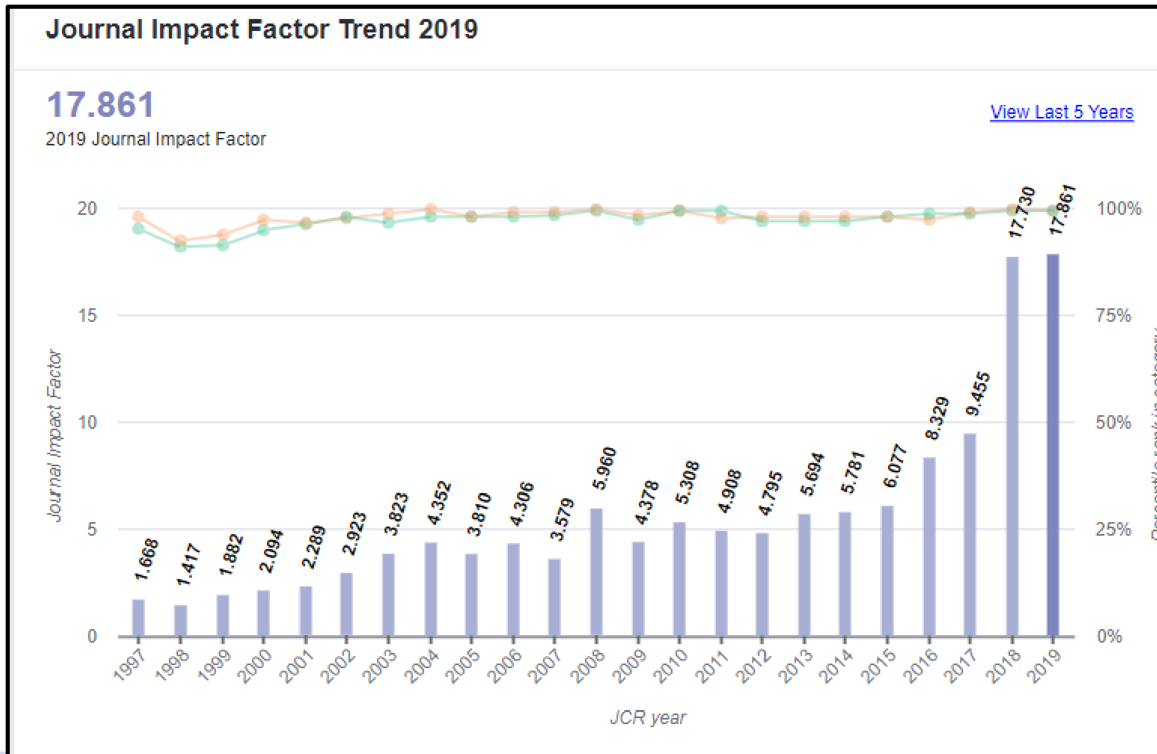

Several Important Concepts

- *h5*-index
  - *h5*-index is the *h*-index for articles published in the last *5* complete years. It is the largest number *h* such that *h* articles published in *2015-2019* have at least *h* citations each
    - Homepage: https://scholar.google.com/citations?view_op=top_venues&hl=en
    - Alternate Homepage: https://sc.panda321.com/citations?view_op=top_venues&hl=zh-CN

Several Important Concepts

	Publication	h5-index	h5-median
1.	Nature	376	552
2.	The New England Journal of Medicine	365	639
3.	Science	356	526
4.	The Lancet	301	493
5.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	299	509
6.	Advanced Materials	273	369
7.	Nature Communications	273	366
8.	Cell	269	417
9.	Chemical Reviews	267	438
10.	Chemical Society reviews	240	368
11.	Journal of the American Chemical Society	236	324
12.	Angewandte Chemie	229	316
13.	Proceedings of the National Academy of Sciences	228	299
14.	JAMA	220	337
15.	Nucleic Acids Research	219	475
16.	Physical Review Letters	209	288
17.	International Conference on Learning Representations	203	359
18.	Journal of Clinical Oncology	202	300
19.	Renewable and Sustainable Energy Reviews	201	263
20.	Energy & Environmental Science	199	289
21.	Neural Information Processing Systems	198	377
22.	ACS Nano	193	257
23.	Nature Materials	184	283
24.	The Lancet Oncology	183	300
25.	Nano Letters	183	241
26.	Advanced Energy Materials	181	250
27.	Nature Genetics	180	266
28.	Scientific Reports	178	226
29.	IEEE/CVF International Conference on Computer Vision	176	295
30.	PLoS ONE	175	237
31.	Nature Medicine	173	288
32.	Advanced Functional Materials	172	221
33.	International Conference on Machine Learning (ICML)	171	309
34.	The Astrophysical Journal	167	231
35.	Circulation	166	260
36.	Journal of the American College of Cardiology	164	232
37.	Journal of Materials Chemistry A	161	216
38.	Nature Nanotechnology	160	272
39.	ACS Applied Materials & Interfaces	160	200
40.	Journal of High Energy Physics	158	209
41.	Nature Biotechnology	154	269
42.	Journal of Cleaner Production	154	208
43.	Neuron	154	199
44.	European Heart Journal	153	245
45.	Applied Catalysis B: Environmental	153	189
46.	Nature Neuroscience	152	219
47.	Nature Methods	151	242
48.	BMJ	150	222
49.	Accounts of Chemical Research	149	220
50.	Gastroenterology	148	222
51.	Physical Review D	148	208
52.	Blood, The Journal of the American Society of Hematology	148	192
53.	Cochrane Database of Systematic Reviews	147	218
54.	Nano Energy	147	192
55.	American Economic Review	146	227
56.	ACS Catalysis	146	207
57.	Monthly Notices of the Royal Astronomical Society	146	193
58.	European Conference on Computer Vision	144	286
59.	Nature Photonics	144	245
60.	Computers in Human Behavior	144	198
61.	Applied Energy	143	185
62.	Science Advances	142	213
63.	Nature Physics	140	217
64.	Chemistry of Materials	140	189
65.	IEEE Communications Surveys & Tutorials	138	248
66.	Environmental Science & Technology	138	185
67.	Nature Reviews. Molecular Cell Biology	137	264
68.	Immunity	137	204
69.	Cell Metabolism	137	191
70.	Nature Climate Change	136	213
71.	Science Translational Medicine	136	202
72.	Meeting of the Association for Computational Linguistics (ACL)	135	220
73.	Chemical engineering journal	134	171
74.	Molecular Cell	133	181
75.	Clinical Cancer Research	133	177
76.	Chemical communications (Cambridge, England)	132	158
77.	IEEE Transactions on Pattern Analysis and Machine Intelligence	131	261
78.	Science of The Total Environment	131	176
79.	Nanoscale	131	169
80.	IEEE Communications Magazine	130	190
81.	Nature Immunology	130	189
82.	Journal of Hepatology	130	188
83.	European Urology	130	187
84.	The Journal of Clinical Investigation	130	179
85.	Nature Energy	129	235
86.	The Lancet Infectious Diseases	129	189
87.	IEEE Transactions on Industrial Electronics	129	174
88.	Cell Reports	128	165
89.	Physical Review B	128	156
90.	Nature Reviews Cancer	127	246
91.	Diabetes Care	127	209
92.	The Journal of Physical Chemistry Letters	127	193
93.	Circulation Research	127	187
94.	Annals of the Rheumatic Diseases	127	183
95.	eLife	127	159
96.	AAAI Conference on Artificial Intelligence	126	183
97.	Bioinformatics	125	207
98.	Annals of Oncology	125	199
99.	Nature Reviews Immunology	124	265
100.	Gut	124	193

Several Important Concepts

	Publication	h5-index	h5-median
1.	Nature	376	552
2.	The New England Journal of Medicine	365	639
3.	Science	356	526
4.	The Lancet	301	493
5.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	299	509
6.	Advanced Materials	273	369
7.	Nature Communications	273	366
8.	Cell	269	417
9.	Chemical Reviews	267	438
10.	Chemical Society reviews	240	368
11.	Journal of the American Chemical Society	236	324
12.	Angewandte Chemie	229	316
13.	Proceedings of the National Academy of Sciences	228	299
14.	JAMA	220	337
15.	Nucleic Acids Research	219	475
16.	Physical Review Letters	209	288
17.	International Conference on Learning Representations	203	359
18.	Journal of Clinical Oncology	202	300
19.	Renewable and Sustainable Energy Reviews	201	263
20.	Energy & Environmental Science	199	289
21.	Neural Information Processing Systems	198	377
22.	ACS Nano	193	257
23.	Nature Materials	184	283
24.	The Lancet Oncology	183	300
25.	Nano Letters	183	241
26.	Advanced Energy Materials	181	250
27.	Nature Genetics	180	266
28.	Scientific Reports	178	226
29.	IEEE/CVF International Conference on Computer Vision	176	295
30.	PLoS ONE	175	237
31.	Nature Medicine	173	288
32.	Advanced Functional Materials	172	221
33.	International Conference on Machine Learning (ICML)	171	309
34.	The Astrophysical Journal	167	231
35.	Circulation	166	260
36.	Journal of the American College of Cardiology	164	232
37.	Journal of Materials Chemistry A	161	216
38.	Nature Nanotechnology	160	272
39.	ACS Applied Materials & Interfaces	160	200
40.	Journal of High Energy Physics	158	209
41.	Nature Biotechnology	154	269
42.	Journal of Cleaner Production	154	208
43.	Neuron	154	199
44.	European Heart Journal	153	245
45.	Applied Catalysis B: Environmental	153	189
46.	Nature Neuroscience	152	219
47.	Nature Methods	151	242
48.	BMJ	150	222
49.	Accounts of Chemical Research	149	220
50.	Gastroenterology	148	222
51.	Physical Review D	148	208
52.	Blood, The Journal of the American Society of Hematology	148	192
53.	Cochrane Database of Systematic Reviews	147	218
54.	Nano Energy	147	192
55.	American Economic Review	146	227
56.	ACS Catalysis	146	207
57.	Monthly Notices of the Royal Astronomical Society	146	193
58.	European Conference on Computer Vision	144	286
59.	Nature Photonics	144	245
60.	Computers in Human Behavior	144	198
61.	Applied Energy	143	185
62.	Science Advances	142	213
63.	Nature Physics	140	217
64.	Chemistry of Materials	140	189
65.	IEEE Communications Surveys & Tutorials	138	248
66.	Environmental Science & Technology	138	185
67.	Nature Reviews. Molecular Cell Biology	137	264
68.	Immunity	137	204
69.	Cell Metabolism	137	191
70.	Nature Climate Change	136	213
71.	Science Translational Medicine	136	202
72.	Meeting of the Association for Computational Linguistics (ACL)	135	220
73.	Chemical engineering journal	134	171
74.	Molecular Cell	133	181
75.	Clinical Cancer Research	133	177
76.	Chemical communications (Cambridge, England)	132	158
77.	IEEE Transactions on Pattern Analysis and Machine Intelligence	131	261
78.	Science of The Total Environment	131	176
79.	Nanoscale	131	169
80.	IEEE Communications Magazine	130	190
81.	Nature Immunology	130	189
82.	Journal of Hepatology	130	188
83.	European Urology	130	187
84.	The Journal of Clinical Investigation	130	179
85.	Nature Energy	129	235
86.	The Lancet Infectious Diseases	129	189
87.	IEEE Transactions on Industrial Electronics	129	174
88.	Cell Reports	128	165
89.	Physical Review B	128	156
90.	Nature Reviews Cancer	127	246
91.	Diabetes Care	127	209
92.	The Journal of Physical Chemistry Letters	127	193
93.	Circulation Research	127	187
94.	Annals of the Rheumatic Diseases	127	183
95.	eLife	127	159
96.	AAAI Conference on Artificial Intelligence	126	183
97.	Bioinformatics	125	207
98.	Annals of Oncology	125	199
99.	Nature Reviews Immunology	124	265
100.	Gut	124	193

Categories > Engineering & Computer Science > Computer Vision & Pattern Recognition

Se

2020

	Publication	h5-index	h5-median
1.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	299	509
2.	IEEE/CVF International Conference on Computer Vision	176	295
3.	European Conference on Computer Vision	144	286
4.	IEEE Transactions on Pattern Analysis and Machine Intelligence	131	261
5.	IEEE Transactions on Image Processing	113	156
6.	Pattern Recognition	85	126
7.	IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops	73	110
8.	International Journal of Computer Vision	70	150
9.	Medical Image Analysis	67	115
10.	Pattern Recognition Letters	59	80
11.	British Machine Vision Conference (BMVC)	57	87
12.	Workshop on Applications of Computer Vision (WACV)	54	87
13.	IEEE International Conference on Image Processing (ICIP)	52	71
14.	IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)	51	75
15.	Computer Vision and Image Understanding	50	97
16.	Journal of Visual Communication and Image Representation	45	60
17.	IEEE International Conference on Automatic Face & Gesture Recognition	41	64
18.	International Conference on 3D Vision	37	65
19.	Image and Vision Computing	36	55
20.	International Conference on Pattern Recognition	35	55

类别 > Engineering & Computer Science > Computer Vision & Pattern Recognition

2021

	出版物	h5 指数	h5 中位数
1.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	356	583
2.	European Conference on Computer Vision	197	342
3.	IEEE/CVF International Conference on Computer Vision	184	311
4.	IEEE Transactions on Pattern Analysis and Machine Intelligence	149	275
5.	IEEE Transactions on Image Processing	123	187
6.	Pattern Recognition	99	141
7.	IEEE/CVF Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)	89	154
8.	Medical Image Analysis	76	149
9.	International Journal of Computer Vision	72	173
10.	British Machine Vision Conference (BMVC)	66	102
11.	Pattern Recognition Letters	66	93
12.	IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)	62	121
13.	IEEE International Conference on Image Processing (ICIP)	60	89
14.	IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)	57	83
15.	Computer Vision and Image Understanding	52	91
16.	Journal of Visual Communication and Image Representation	47	64
17.	International Conference on 3D Vision (3DV)	44	89
18.	International Conference on Pattern Recognition	43	78
19.	Asian Conference on Computer Vision (ACCV)	43	69
20.	IEEE International Conference on Automatic Face & Gesture Recognition	42	66

Several Important Concepts

- IF *vs.* *h5*-index
  - Publications with more papers per year tend to have higher *h5*-index
    - IF: **11.079**, *h5*-index: **111**, papers in 2017 and 2018: **649** (IEEE TCYB)
    - IF: **11.148**, *h5*-index: **67**, papers in 2017 and 2018: **263** (Elsevier MIA)

Several Important Concepts

- IF *vs. h5*-index
  - Publications with more papers per year tend to have higher *h5*-index
    - IF: **11.079**, *h5*-index: **111**, papers in 2017 and 2018: **649** (IEEE TCYB)
    - IF: **11.148**, *h5*-index: **67**, papers in 2017 and 2018: **263** (Elsevier MIA)
  - Several **extraordinary** papers would result in high IF
    - However, for high *h5*-index, much more **good** papers are need

Several Important Concepts

- IF *vs. h5*-index
  - Publications with more papers per year tend to have higher *h5*-index
    - IF: **11.079**, *h5*-index: **111**, papers in 2017 and 2018: **649** (IEEE TCYB)
    - IF: **11.148**, *h5*-index: **67**, papers in 2017 and 2018: **263** (Elsevier MIA)
  - Several **extraordinary** papers would result in high IF
    - However, for high *h5*-index, much more **good** papers are need
  - Only the citations of **official publications** are used for IF computation
    - However, for *h5*-index, citations in **arXiv** papers, and even in **zhihu**, are used

Several Important Concepts

- China Computer Federation (CCF) List
  - List of International academic conferences and periodicals recommended
    - Homepage: https://www.ccf.org.cn/en/Bulletin/2019-05-13/663884.shtml
    - Computer Architecture/Parallel and Distributed Computer/Storage
    - Computer Networks
    - Network and Information Security
    - Software Engineering/System Software/Programming Language
    - Database/Data Mining/Content Retrieval
    - Computer Science Theory
    - CAD & Graphics and Multimedia
    - Artificial Intelligence
    - Human Computer Interaction and Pervasive Computing
    - Cross-disciplinary/Comprehensive/Emerging

Several Important Concepts

- China Computer Federation (CCF) List
  - List of International academic conferences and periodicals recommended
    - Homepage: https://www.ccf.org.cn/en/Bulletin/2019-05-13/663884.shtml
    - Computer Architecture/Parallel and Distributed Computer/Storage
    - Computer Networks
    - Network and Information Security
    - Software Engineering/System Software/Programming Language
    - Database/Data Mining/Content Retrieval
    - Computer Science Theory
    - **CAD & Graphics and Multimedia**
    - **Artificial Intelligence**
    - Human Computer Interaction and Pervasive Computing
    - Cross-disciplinary/Comprehensive/Emerging

中国计算机学会推荐国际学术期刊

The List of International Academic Periodicals Recommended by CCF

（计算机图形学与多媒体）

CAD and Graphics & Multimedia

1. Class A

No.	Abbr. of Journal	Full Name of Journals	Publishing House	Website
1	TOG	ACM Transactions on Graphics	ACM	http://dblp.uni-trier.de/db/journals/tog/
2	TIP	IEEE Transactions on Image Processing	IEEE	http://dblp.uni-trier.de/db/journals/tip/
3	TVCG	IEEE Transactions on Visualization and Computer Graphics	IEEE	http://dblp.uni-trier.de/db/journals/tvcg/

中国计算机学会推荐国际学术会议

The List of International Academic Conferences Recommended by CCF

（计算机图形学与多媒体）

CAD and Graphics & Multimedia

1. Class A

No.	Abbr. of Conf.	Conferences	Organizer	Website
1	ACM MM	ACM International Conference on Multimedia	ACM	http://dblp.uni-trier.de/db/conf/mm/
2	SIGGRAPH	ACM SIGGRAPH Annual Conference	ACM	http://dblp.uni-trier.de/db/conf/siggraph/index.html
3	VR	IEEE Virtual Reality	IEEE	http://dblp.uni-trier.de/db/conf/vr/
4	IEEE VIS	IEEE Visualization Conference	IEEE	http://dblp.uni-trier.de/db/conf/visualization/index.html

中国计算机学会推荐国际学术期刊

The List of International Academic Periodicals Recommended by CCF

（人工智能）

Artificial Intelligence

1. Class A

No.	Abbr. of Journal	Full Name of Journals	Publishing House	Website
1	AI	Artificial Intelligence	Elsevier	http://dblp.uni-trier.de/db/journals/ai/
2	TPAMI	IEEE Trans on Pattern Analysis and Machine Intelligence	IEEE	http://dblp.uni-trier.de/db/journals/pami/
3	IJCV	International Journal of Computer Vision	Springer	http://dblp.uni-trier.de/db/journals/ijcv/
4	JMLR	Journal of Machine Learning Research	MIT Press	http://dblp.uni-trier.de/db/journals/jmlr/

中国计算机学会推荐国际学术会议

The List of International Academic Conferences Recommended by CCF

（人工智能）

Artificial Intelligence

1. Class A

No.	Abbr. of Conf.	Conferences	Organizer	Website
1	AAAI	AAAI Conference on Artificial Intelligence	AAAI	http://dblp.uni-trier.de/db/conf/aaai/
2	NeurIPS	Annual Conference on Neural Information Processing Systems	MIT Press	http://dblp.uni-trier.de/db/conf/nips/
3	ACL	Annual Meeting of the Association for Computational Linguistics	ACL	http://dblp.uni-trier.de/db/conf/acl/
4	CVPR	IEEE Conference on Computer Vision and Pattern Recognition	IEEE	http://dblp.uni-trier.de/db/conf/cvpr/
5	ICCV	International Conference on Computer Vision	IEEE	http://dblp.uni-trier.de/db/conf/iccv/
6	ICML	International Conference on Machine Learning	ACM	http://dblp.uni-trier.de/db/conf/icml/
7	IJCAI	International Joint Conference on Artificial Intelligence	Morgan Kaufmann	http://dblp.uni-trier.de/db/conf/ijcai/

Several Important Concepts

- China Computer Federation (CCF) List
  - List of **Chinese** academic periodicals recommended

CCF 推荐中文科技期刊目录

A 类

序号	期刊名称	主办单位	网址
1	软件学报	中国科学院软件研究所 中国计算机学会	http://www.jos.org.cn
2	计算机学报	中国计算机学会 中国科学院计算技术研究所	http://cjc.ict.ac.cn
3	中国科学：信息科学	中国科学院 国家自然科学基金委员会	http://infocn.scichina.com
4	计算机研究与发展	中国科学院计算技术研究所 中国计算机学会	http://crad.ict.ac.cn
5	计算机辅助设计与图形学学报	中国计算机学会 中国科学院计算技术研究所	http://www.jcad.cn
6	电子学报	中国电子学会	http://www.ejournal.org.cn
7	自动化学报	中国自动化学会 中国科学院自动化研究所	http://www.aas.net.cn

Recommended Journals & Conferences

- Journals
  - ***IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)***
    - ***IEEE, IF: 17.861, h5-index: 131, CCF-A***
  - ***International Journal of Computer Vision (IJCV)***
    - ***Springer, IF: 5.698, h5-index: 70, CCF-A***

Recommended Journals & Conferences

- Journals
  - *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*
    - *IEEE, IF: 17.861, h5-index: 131, CCF-A*
  - *International Journal of Computer Vision (IJCV)*
    - *Springer, IF: 5.698, h5-index: 70, CCF-A*
  - **IEEE Transactions on Image Processing (TIP)**
    - **IEEE, IF: 9.340, h5-index: 113, CCF-A**
  - **Pattern Recognition (PR)**
    - **Elsevier, IF: 7.196, h5-index: 85, CCF-B**

Recommended Journals & Conferences

- Journals
  - *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*
    - *IEEE, IF: 17.861, h5-index: 131, CCF-A*
  - *International Journal of Computer Vision (IJCV)*
    - *Springer, IF: 5.698, h5-index: 70, CCF-A*
  - **IEEE Transactions on Image Processing (TIP)**
    - **IEEE, IF: 9.340, h5-index: 113, CCF-A**
  - **Pattern Recognition (PR)**
    - **Elsevier, IF: 7.196, h5-index: 85, CCF-B**
  - Computer Vision and Image Understanding (CVIU)
    - Elsevier, IF: 3.121, h5-index: 50, CCF-B
  - Image and Vision Computing (IVC)
    - Elsevier, IF: 3.103, h5-index: 36, CCF-C
  - Pattern Recognition Letters (PRL)
    - Elsevier, IF: 3.255, h5-index: 59, CCF-C

Recommended Journals & Conferences

- Related Journals
  - Artificial Intelligence: ***IEEE TNNLS***, **IEEE TCYB**
  - Computer Graphics: ***ACM TOG***, **IEEE TVCG**, Wiley CGF
  - Robotics: ***SAGE IJRR***, **IEEE TRO**, IEEE RAL, Elsevier RAS
  - Remote Sensing: **ISPRS P&RS**, **IEEE TGRS**, IEEE GRSL
  - Multimedia: **IEEE TMM**, **IEEE TSCVT**

Recommended Journals & Conferences

- Conferences
  - ***IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)***
    - ***h5-index: 299, CCF-A***
  - ***IEEE/CVF International Conference on Computer Vision (ICCV)***
    - ***h5-index: 176, CCF-A***
  - ***European Conference on Computer Vision (ECCV)***
    - ***h5-index: 144, CCF-B***

Recommended Journals & Conferences

- Conferences
  - *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*
    - *h5-index: 299, CCF-A*
  - *IEEE/CVF International Conference on Computer Vision (ICCV)*
    - *h5-index: 176, CCF-A*
  - *European Conference on Computer Vision (ECCV)*
    - *h5-index: 144, CCF-B*
  - **British Machine Vision Conference (BMVC)**
    - **h5-index: 57, CCF-C**
  - **International Conference on 3D Vision (3DV)**
    - **h5-index: 37, CCF-C**
  - **Asian Conference on Computer Vision (ACCV)**
    - **h5-index: 33, CCF-C**

Recommended Journals & Conferences

- Conferences
  - *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*
    - *h5-index: 299, CCF-A*
  - *IEEE/CVF International Conference on Computer Vision (ICCV)*
    - *h5-index: 176, CCF-A*
  - *European Conference on Computer Vision (ECCV)*
    - *h5-index: 144, CCF-B*
  - British Machine Vision Conference (BMVC)
    - h5-index: 57, CCF-C
  - International Conference on 3D Vision (3DV)
    - h5-index: 37, CCF-C
  - Asian Conference on Computer Vision (ACCV)
    - h5-index: 33, CCF-C
  - IEEE International Conference on Image Processing (ICIP)
    - h5-index: 52, CCF-C
  - International Conference on Pattern Recognition (ICPR)
    - h5-index: 35, CCF-C

Recommended Journals & Conferences

- Related Conference
  - Artificial Intelligence: ***ICLR***, ***NeurIPS***, ***ICML***, **AAAI**, **IJCAI**
  - Computer Graphics: ***SIGGRAPH***, ***SIGGRAPH Asia***
  - Robotics: ***RSS***, **ICRA**, **IROS**
  - Remote Sensing: ISPRS Archives, ISPRS Annals
  - Multimedia: **ACM MM**, ICME

Reviewing Process

- Journal Papers
  - Manuscript submitted
  - Associate Editor-in-Chief (AEiC) assigns papers to Associate Editor (AE)
  - AE assigns papers to reviewers
  - First round review: several months (or years)
    - Accept as is (rare cases)
    - Accept with minor revision (rare cases)
    - Major revision
    - Resubmit as new
    - Reject
  - Second round review: several months
    - Accept as is
    - Accept with minor revision
    - Major revision (rare cases)
    - Resubmit as new
    - Reject
  - Editor-in-chief (EiC) makes final decision

IEEE Transactions on PATTERN ANALYSIS AND MACHINE INTELLIGENCE

Editorial Board

Reviewing Process

- Conference Organization
  - General Chairs (GCs): administrating conference ≈ Journal EiC
  - Program Chairs (PCs): handling papers ≈ Journal AEiCs
  - Workshop chairs
  - Tutorial chairs
  - Website chairs
  - Publication chairs
  - …
  - Area Chairs (ACs) ≈ Journal AEs
    - Assign reviewers
    - Read reviews and rebuttals
    - Recommendation

Reviewing Process

- Conference Papers
  - Manuscript submitted
  - PCs assign papers to ACs
  - ACs assign papers to reviewers
  - First round reviewing: About two months
  - Rating and Rebuttal: About one week
  - Second round reviewing: About one month
  - Author notification

Reviewing Process

- AC Meetings
  - Each paper is reviewed by 2/3 ACs
  - ACs make recommendations
  - PCs make final decisions
  - ACs know the reviewers and the reviews are weighted
  - Based on reviews and rebuttal
    - Accept: decide oral later
    - Reject: don’t waste time
    - Go either way: lots of papers
  - Usually agree with reviewers but anything can happen as long as there are good justifications

Reviewing Process

- More About Conference Papers
  - Double or Single blind review
    - Double blind review for relatively good conferences
      - CVPR, ICCV, ECCV, BMVC, ACCV, 3DV, *etc.*
    - Single blind review for others
      - ICIP, ICPR, *etc.*

Reviewing Process

- More About Conference Papers
  - Double or Single blind review
    - Double blind review for relatively good conferences
      - CVPR, ICCV, ECCV, BMVC, ACCV, 3DV, *etc.*
    - Single blind review for others
      - ICIP, ICPR, *etc.*
  - Accept rate
    - 20%~30% for top conferences (CVPR, ICCV, ECCV, *etc.*)
    - 40%~50% for good conferences
    - 99%~99.99% for 'purchasable' conferences

By X. Gao

Reviewing Process

- More About Conference Papers
  - Oral or Poster

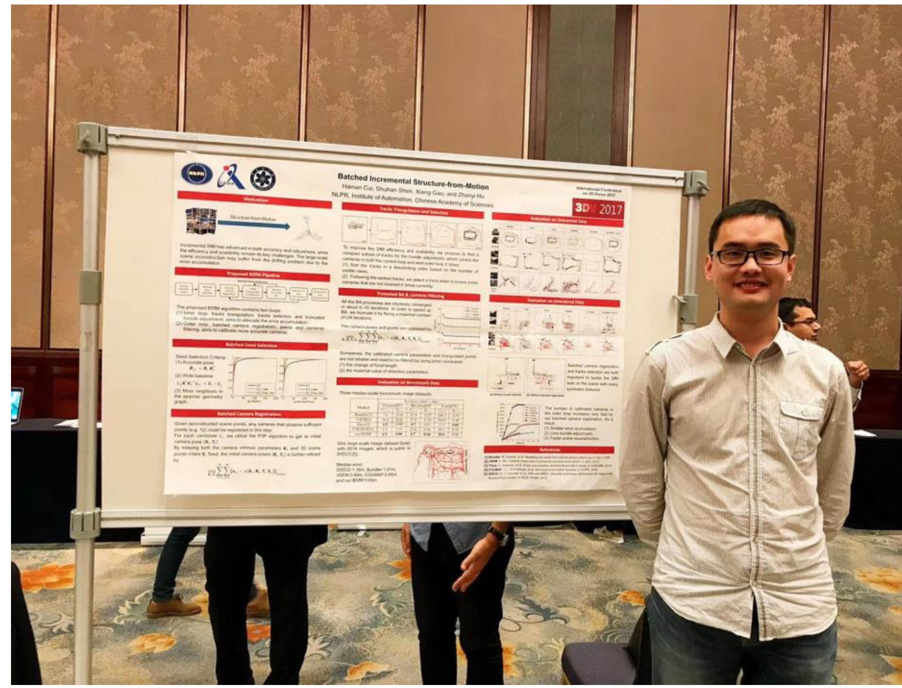

Reviewing Process

- More About Conference Papers

Posters/Orals

As before, papers were accepted as orals and posters purely based on the quality. There were no caps set in the paper decision process.

- **6,424 registered** (vs. 5,165 in 2019)
- **5,865 valid submissions** (vs. 4,538 in 2019)
- **1,467 accepted (25.0%)**
- **335 orals (5.7%)**

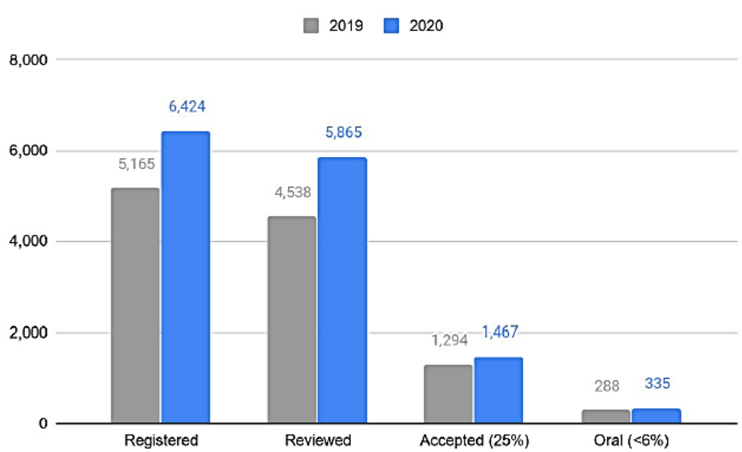

Registered vs accepted last 10 years

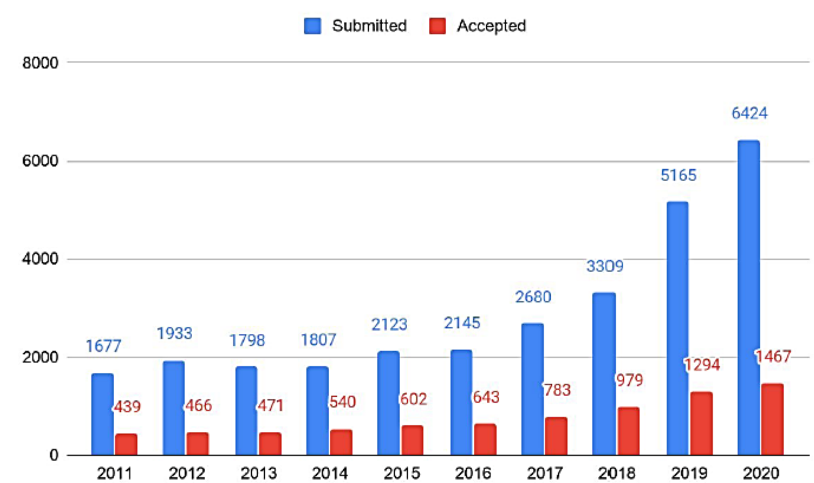

Reviewing Process

- More About CVPR 2020

Author Distribution

Authors by country/region

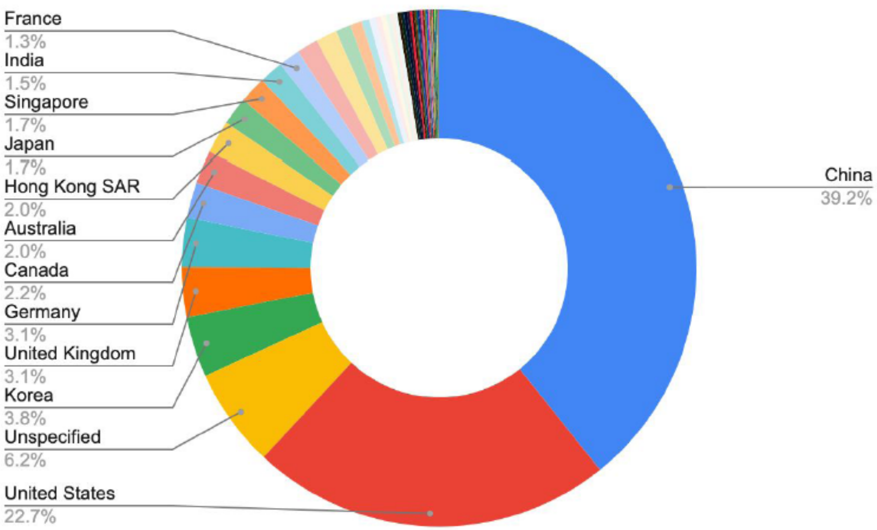

Authors by organization (top 10)

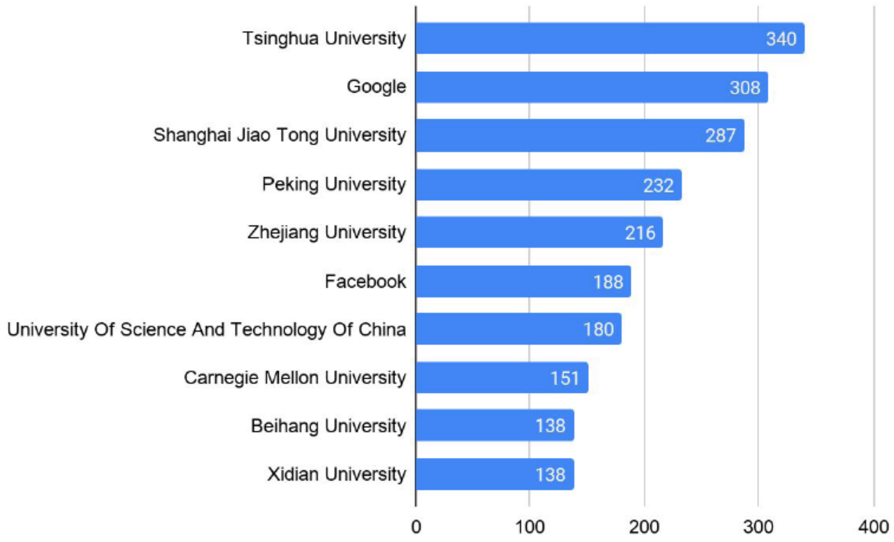

Reviewing Process

- More About CVPR 2020

Reviewer Distribution

Reviewers by country/region

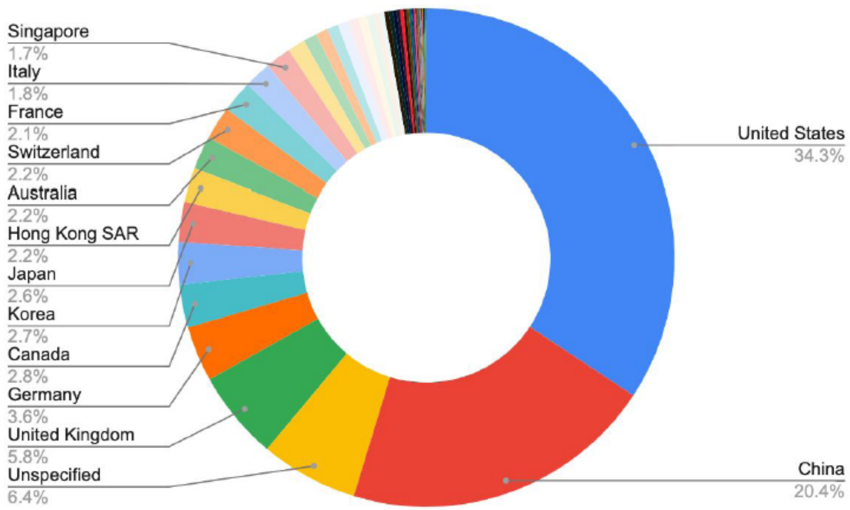

Reviewers by organization (top 10)

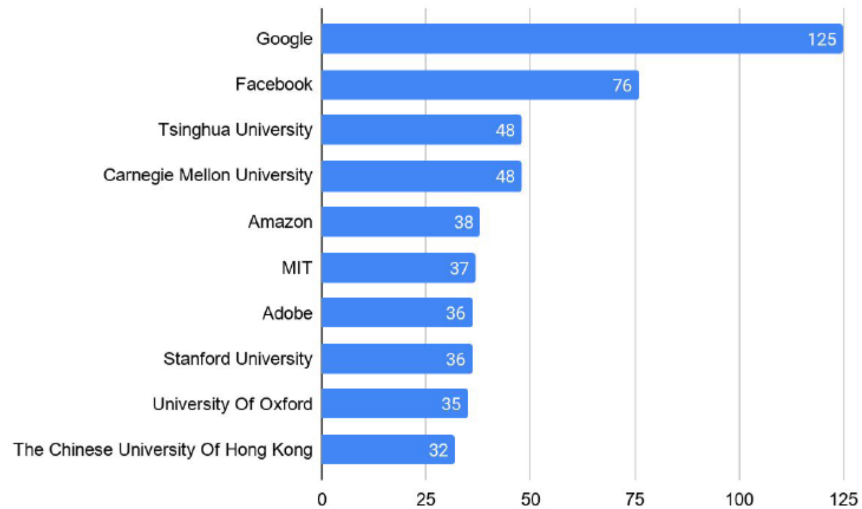

Reviewing Process

- More About CVPR 2020

AC Distribution

ACs by country/region

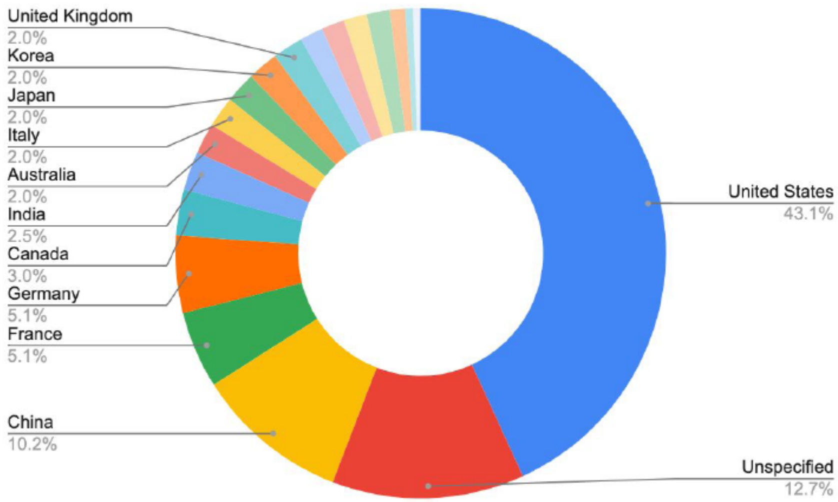

ACs by organization (top 10)

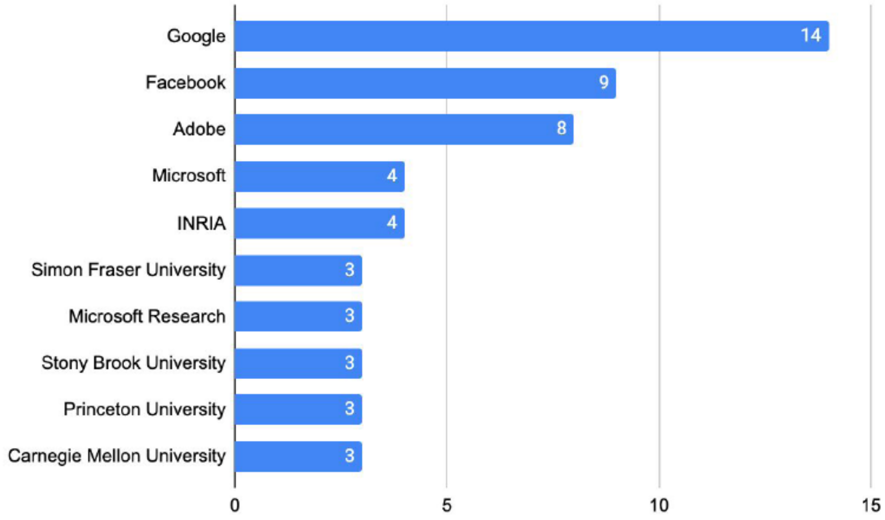

- 35 women
- Increasingly more ACs in Asia

Reviewing Process

- More About CVPR 2020

General Chairs

Terry Boult
UCCS

Gerard Medioni
Amazon & USC

Ramin Zabih
Cornell & Google

Program Chairs

Ce Liu
Google

Greg Mori
SFU & Borealis AI

Kate Saenko
Boston University

Silvio Savarese
Stanford University

Reviewing Process

- More About CVPR 2020

Area Chairs

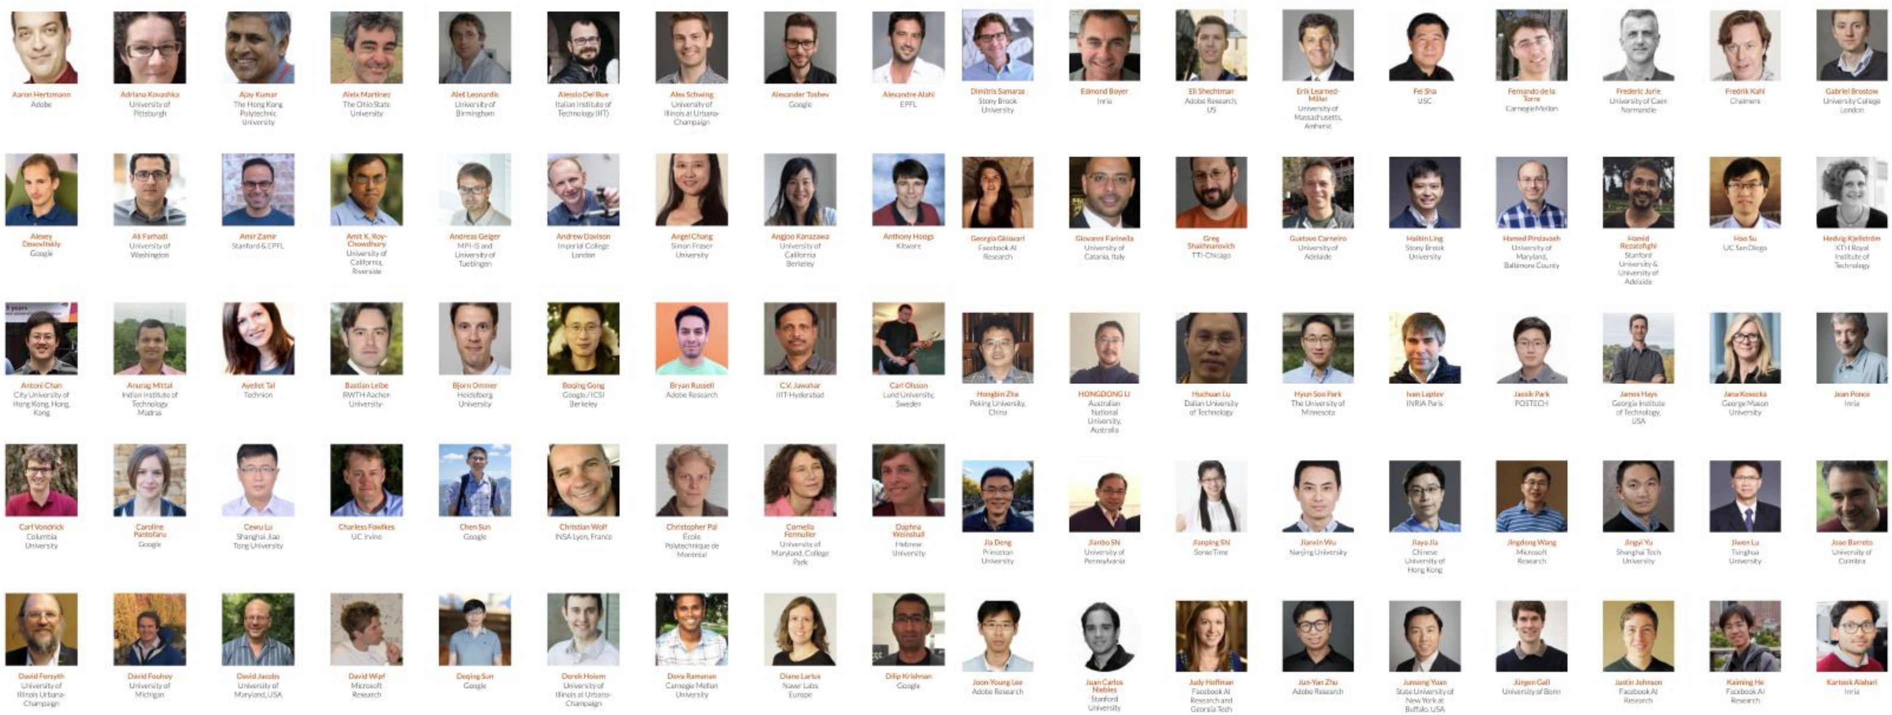

Reviewing Process

- More About CVPR 2020

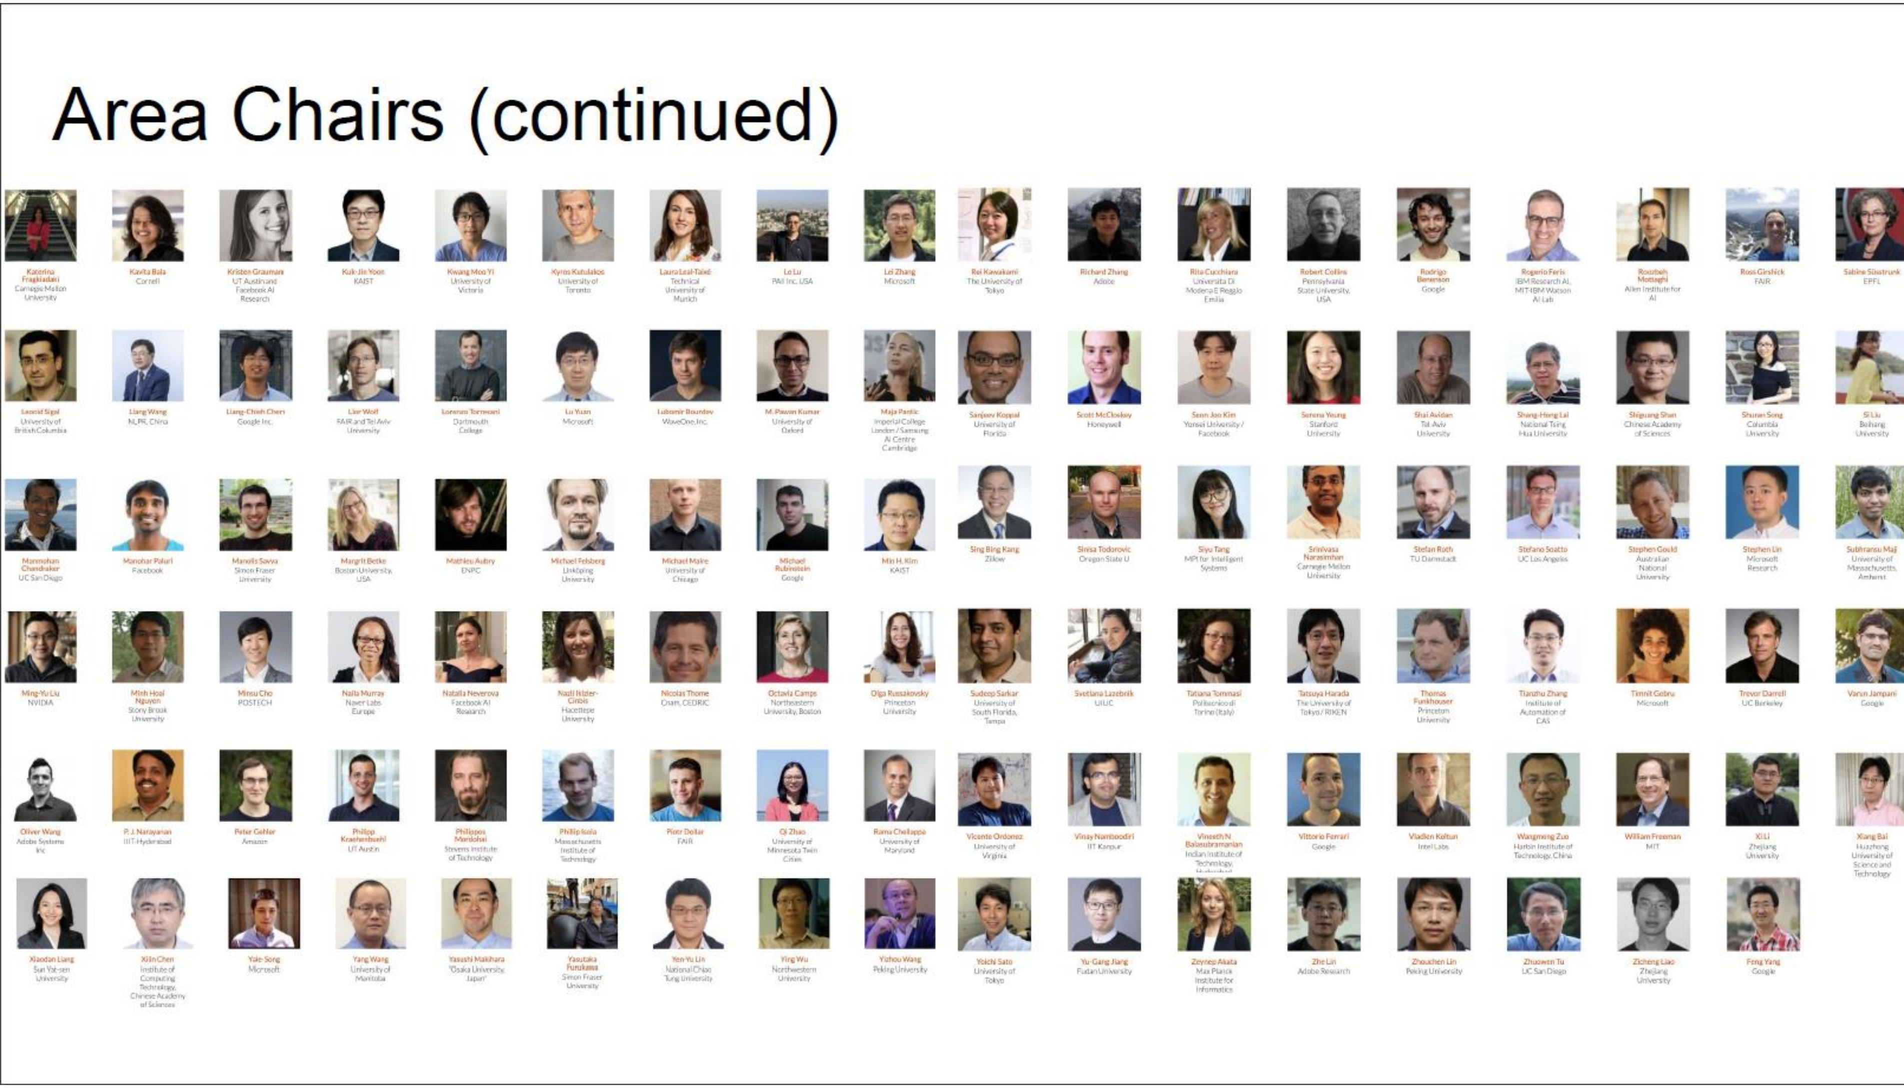

Reviewing Process

- More About CVPR 2020

AC meeting at UCSD

- More About CVPR 2020

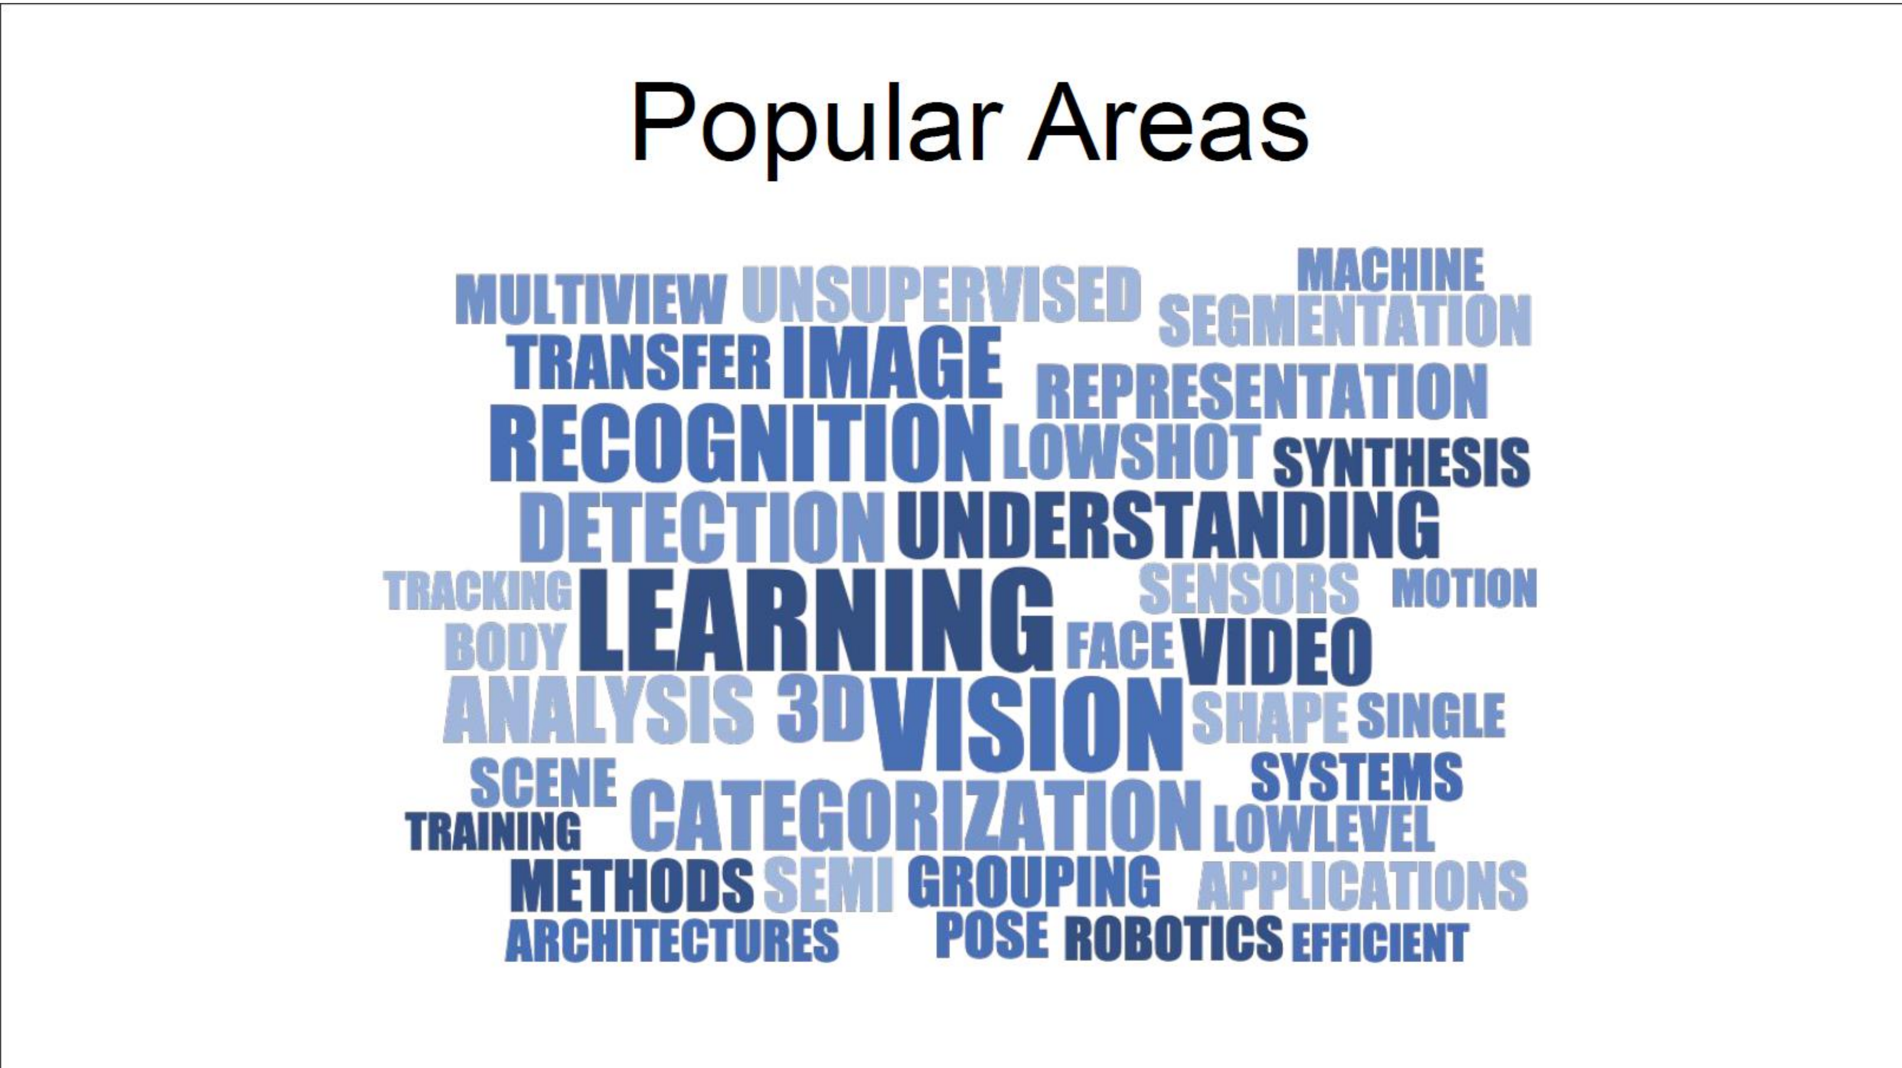

Reviewing Process

- More About CVPR 2020

Distribution of subject Areas

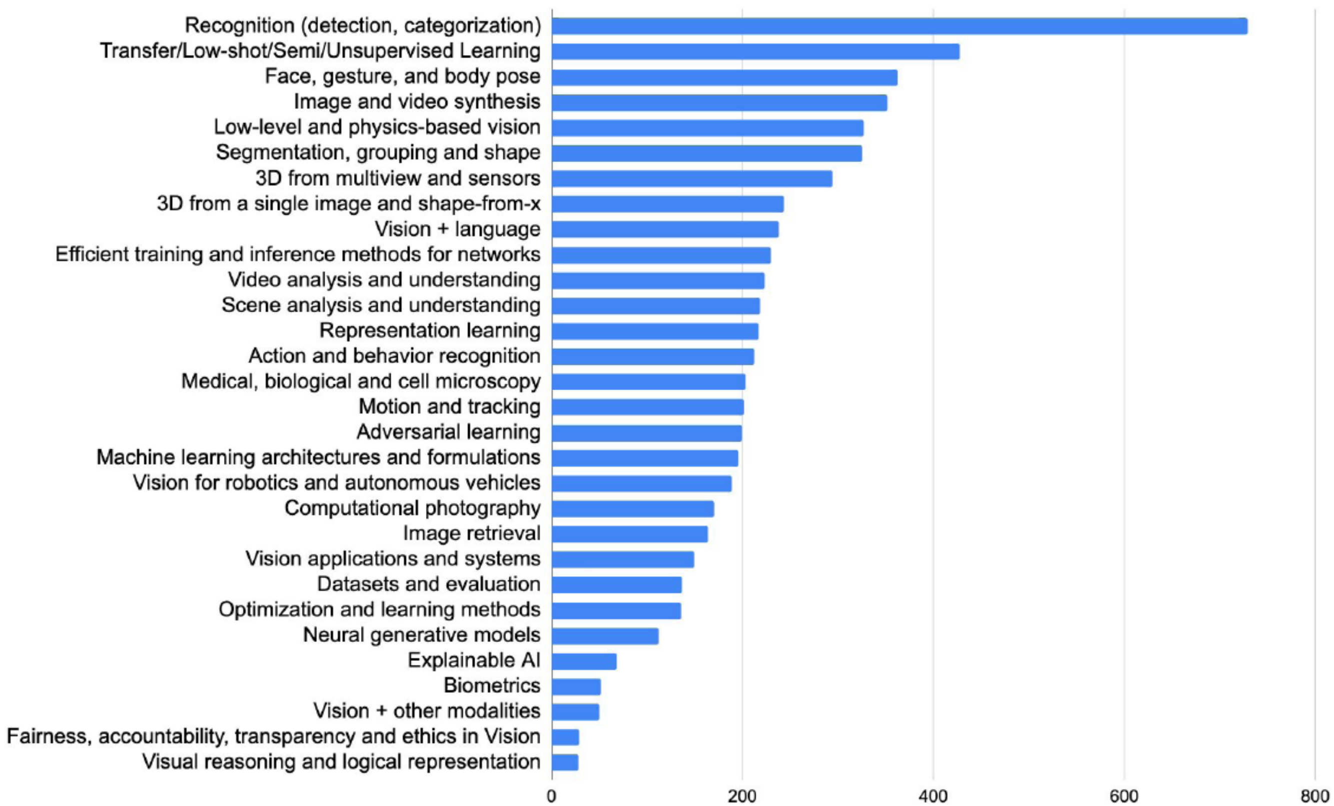

Reviewing Process

- ## More About CVPR 2020

  - 31 tutorials

Title (Program)	Organizers	Tutorial Website	Date	Schedule
Deep Learning and Multiple Drone Vision	Ioannis Pitas	http://icarus.csd.auth.gr/cvpr2020-tutorial-deep-learning-and-multiple-drone-vision/	19-Jun	Half-day (morning)
RANSAC in 2020	Jiri Matas, Ondrej Chum, Tat-Jun Chin, René Ranftl, Dmytro Mishkin, Dániel Baráth	http://cmp.felk.cvut.cz/cvpr2020-ransac-tutorial/	14-Jun	Full-day
Vision Models for Emerging Media Technologies and Their Impact on Computer Vision	Marcelo Bertalmío	https://www.upf.edu/web/marcelo-bertalmio/cvpr-2020-tutorial	19-Jun	Half-day (afternoon)

  - 67 workshops

Workshop Name	Organizers Names	Workshop URL	Full Schedule
15th IEEE Computer Society Biometrics Workshop	Bir Bhanu, Ajay Kumar	https://vislab.ucr.edu/Biometrics2020/index.php	June 19th, Full Day
16th IEEE CVPR Workshop on Perception Beyond the Visible Spectrum	Riad I. Hammoud, Michael Teutsch, Angel D. Sappa, Yi Ding	http://vcipl-okstate.org/pbvs/20/	June 14th, Full Day
2nd CVPR Workshop on 3D Scene Understanding for Vision, Graphics, and Robotics	Siyuan Huang, Chuhang Zou, Hao Su, Alexander Schwing, Shuran Song, Jiajun Wu, Siyuan Qi, Yixin Zhu, David Forsyth, Derek Hoiem, Leonidas Guibas, and Song-Chun Zhu	https://scene-understanding.com	June 15th, Full Day

Reviewing Process

- More About CVPR 2020

Code submission!

- Opportunity for authors to voluntarily submit their code
- Out of all of submitted papers, 730 were coupled with code uploads.

Reviewing Process

- More About CVPR 2020

Code submission!

- Opportunity for authors to voluntarily submit their code
- Out of all of submitted papers, 730 were coupled with code uploads.
- **Talk is cheap, show me the code!**

Reviewing Process

- More About CVPR 2020

Code submission!

- Opportunity for authors to voluntarily submit their code
- Out of all of submitted papers, 730 were coupled with code uploads.
- **Talk is cheap, show me the code!**
- **P话少说，放‘码’过来！**

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Novel idea and convincing evaluation

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Longer reviewing and revising period

- Novel idea and convincing evaluation
- Shorter reviewing (and rebuttal) period

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Longer reviewing and revising period
- Unscheduled reviewing process

- Novel idea and convincing evaluation
- Shorter reviewing (and rebuttal) period
- Scheduled reviewing process

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Longer reviewing and revising period
- Unscheduled reviewing process
- Less uncertainty in quality

- Novel idea and convincing evaluation
- Shorter reviewing (and rebuttal) period
- Scheduled reviewing process
- More uncertainty in quality

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Longer reviewing and revising period
- Unscheduled reviewing process
- Less uncertainty in quality
- Time for next submission preparation

- Novel idea and convincing evaluation
- Shorter reviewing (and rebuttal) period
- Scheduled reviewing process
- More uncertainty in quality
- Meeting and presentation opportunity

Journal Papers *vs.* Conference Papers

- Comprehensive study and evaluation
- Longer reviewing and revising period
- Unscheduled reviewing process
- Less uncertainty in quality
- Time for next submission preparation
- Included in SCI and with IF

- Novel idea and convincing evaluation
- Shorter reviewing (and rebuttal) period
- Scheduled reviewing process
- More uncertainty in quality
- Meeting and presentation opportunity
- Higher h5-index for top conferences

How to Structure a Paper?

- State which problem you are addressing, keeping the audience in mind
  - They must care about the problem, which means that sometimes you must tell them why they should care about it

How to Structure a Paper?

- State which problem you are addressing, keeping the audience in mind
    - They must care about the problem, which means that sometimes you must tell them why they should care about it
- State briefly what the other solutions are, and why they aren't satisfactory
    - If they were satisfactory, you wouldn't need to do the work

How to Structure a Paper?

- State which problem you are addressing, keeping the audience in mind
  - They must care about the problem, which means that sometimes you must tell them why they should care about it
- State briefly what the other solutions are, and why they aren't satisfactory
  - If they were satisfactory, you wouldn't need to do the work
- Explain your solution, compare it with other solutions, and say why it's better
  - Only better performance is not enough, its reason is more important

How to Structure a Paper?

- How to write a IEEE/ACM T-level CS paper?

恭喜你已经有了一个A类别的idea和相应的实现以及数据，那么接下来就是写文章的问题了。计算机论文写作还是有一定的规律可以遵循的。

首先，你需要判断自己的文章是投往哪个A类期刊或者会议，是IEEE的还是ACM的。通常IEEE会议的格式和ACM会议的文章格式要求有所不同，我建议你在投论文之前先把该会议的Call For Papers好好研读一番，弄清楚文章长度，需要使用的Latex或者Word模板以及匿名方式等一系列非技术性问题，然后再开始写作。

在写作之前，先问问自己如果这篇文章写好以后给整个领域（community）的贡献是什么，这实际上是整个文章的灵魂，也就是你解决某个问题（problem）的方案（idea）。想清楚以后把它（们）按照重要性顺序写下来，这些就是你在Introduction里面告诉读者包括审稿人的contributions。贡献可能是新算法，新架构，新实现或者是前人没有的insights。你在写contributions的时候面向的读者很有可能是自己，所以可能忽略了problem背景和定义，这些就可以慢慢在Introduction里面填充。

这样Introduction就写好了，比如说：某某问题是实际中存在的一个问题，这个问题重要性是blah，blah，blah。之前发表的论文针对这个问题提出了三个有代表性的解决方案（此处引用可能至少三篇论文）。第一个解决方案甲大概做了一二三，但是没做四；第二个解决方案乙做了一四，但是没做二三；第三个解决方案丙做了一二三四，但是性能比较差。在这篇文章中，我们提出一个性能比较好并且同时做一二三四的解决方案。接下来写我们这个解决方案是如何实现同时支持一二三四的情况下提升性能的。比如说用了新算法，新架构或者新的实现，都可以。讲完基本技术创新点以后就是contributions，把之前想好的贴上去就可以了。最后在Intro里面加上后续内容组织，比如说第二章是相关工作，第三章是综述，。。。

一般来说Intro写完以后会写一章相关工作（Related Work）。从最Related的论文开始写起，比如说以上提到的三篇。这里需要着重讲的是，Related Work不是记流水账（e.g.，甲用了idea A，乙用了idea B，丙用了idea C），而是要比较这些论文，阐述她们各自的优缺点。

How to Get Your CVPR Paper Rejected?

- Do not

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
    - What does *the* reviewer knows so far?
    - What does *the* reviewer expect next and **why**?

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process

By X. Gao

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process
  - Deliver what you promise

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process
  - Deliver what you promise
  - Completely provide important references

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process
  - Deliver what you promise
  - Completely provide important references
  - Carry out sufficient amount of experiments

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process
  - Deliver what you promise
  - Completely provide important references
  - Carry out sufficient amount of experiments
  - Compare with state-of-the-art algorithms

How to Get Your CVPR Paper Rejected?

- Do not
  - Put yourself as a reviewer
  - Pay attention to review process
  - Deliver what you promise
  - Completely provide important references
  - Carry out sufficient amount of experiments
  - Compare with state-of-the-art algorithms
  - Pay attention to writing

By X. Gao

Review Form

- Summary
- Overall Rating
  - Definite accept, weakly accept, borderline, weakly reject, definite reject
- Novelty
  - Very original, original, minor originality, has been done before
- Importance/relevance
  - Of broad interest, interesting to a subarea, interesting only to a small number of attendees, out of scope
- Clarity of presentation
  - Reads very well, is clear enough, difficult to read, unreadable
- Technical correctness
  - Definite correct, probably correct, contains rectifiable errors, has major problems
- Experimental validation
  - Excellent validation, limited but convincing, lacking in some aspects, insufficient validation
- Additional comments

Review Form

- Summary
- Overall Rating
  - Definite accept, weakly accept, borderline, weakly reject, definite reject
- **Novelty**
  - Very original, original, minor originality, has been done before
- Importance/relevance
  - Of broad interest, interesting to a subarea, interesting only to a small number of attendees, out of scope
- Clarity of presentation
  - Reads very well, is clear enough, difficult to read, unreadable
- Technical correctness
  - Definite correct, probably correct, contains rectifiable errors, has major problems
- **Experimental validation**
  - Excellent validation, limited but convincing, lacking in some aspects, insufficient validation
- Additional comments

Novelty

- How to carry out a novel work?
  - Someone told you it is not so hard:
    - Read the papers of your research area extensively
    - Find a research direction you interested in and has never been done
    - That is what you need!

Novelty

- How to carry out a novel work?
  - Someone told you it is not so hard:
    - Read the papers of your research area extensively
    - Find a research direction you interested in and has never been done
    - That is what you need!
  - However, …

Novelty

- How to carry out a novel work?
  - Then, you continue reading and reading, …
  - When you really have read large numbers of papers, you will see:
    - The ideas of high-impact papers are usually different from each other
    - The **routines** of low-impact papers are usually similar

 By X. Gao

Novelty

- What is new in this work?
  - New pipeline, new method, new data, new metric, *etc.*

Novelty

- What is new in this work?
  - New pipeline, new method, new data, new metric, *etc.*
- What are the contributions (over prior art)?
  - Higher accuracy, significant speed-up, scaleup, stronger robustness, ease to implement, less sensitive to parameter, generalization, wide application domain, connection among seemingly unrelated topics, *etc.*

Novelty

- What is new in this work?
  - New pipeline, new method, new data, new metric, *etc.*
- What are the contributions (over prior art)?
  - Higher accuracy, significant speed-up, scaleup, stronger robustness, ease to implement, less sensitive to parameter, generalization, wide application domain, connection among seemingly unrelated topics, *etc.*
- Make a compelling story with strong supporting evidence
  - Then, how to tell a compelling story?

Novelty

- What is new in this work?
  - New pipeline, new method, new data, new metric, *etc.*
- What are the contributions (over prior art)?
  - Higher accuracy, significant speed-up, scaleup, stronger robustness, ease to implement, less sensitive to parameter, generalization, wide application domain, connection among seemingly unrelated topics, *etc.*
- Make a compelling story with strong supporting evidence
  - Then, how to tell a compelling story?
  - Learn from **the voice of China**!
    - Compelling Story
    - Sufficient Evidence
    - Amazing Demonstration

Novelty

- The importance of telling a compelling story
  - **ICML with one line of code!** $\tilde{J}(\boldsymbol{\theta}) = |J(\boldsymbol{\theta}) - b| + b$

Novelty

- The importance of telling a compelling story
  - **ICML with one line of code!** $\tilde{J}(\boldsymbol{\theta}) = |J(\boldsymbol{\theta}) - b| + b$
  - Compelling story

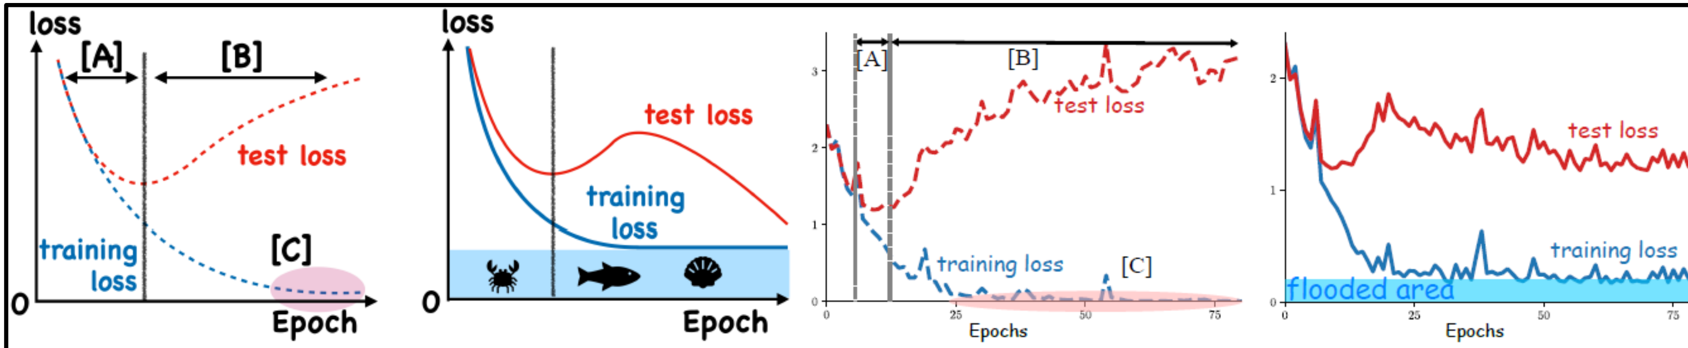

(a) w/o Flooding **(b)** w/ Flooding **(c)** CIFAR-10 w/o Flooding **(d)** CIFAR-10 w/ Flooding

Since it is a simple solution, this modification can be incorporated into existing machine learning code easily: Add one line of code for Eq. (1), after evaluating the original objective function $J(\boldsymbol{\theta})$. A minimal working example with a mini-batch in PyTorch [Paszke et al., 2019] is demonstrated below to show the additional one line of code:

```
outputs = model(inputs)
loss = criterion(outputs, labels)
flood = (loss-b).abs()+b  # This is it!
optimizer.zerograd()
flood.backward()
optimizer.step()
```

Novelty

- The importance of telling a compelling story
  - **ICML with one line of code!** $\tilde{J}(\boldsymbol{\theta}) = |J(\boldsymbol{\theta}) - b| + b$
  - Sufficient evidence

function. ℓ can be the *zero-one loss*,

$$\ell_{01}(\boldsymbol{v}, z') := \begin{cases} 0 & \text{if } \arg\max_{z\in\{1,\ldots,K\}} v_z = z', \\ 1 & \text{otherwise,} \end{cases} \quad (2)$$

where $\boldsymbol{v} := (v_1, \ldots, v_K)^\top \in \mathbb{R}^K$, or a surrogate loss such as the softmax cross-entropy loss,

$$\ell_{\mathrm{CE}}(\boldsymbol{v}, z') := -\log \frac{\exp(v_{z'})}{\sum_{z\in\{1,\ldots,K\}} \exp(v_z)}. \quad (3)$$

For a surrogate loss ℓ, we denote the *classification risk* by

$$R(\boldsymbol{g}) := \mathbb{E}_{p(\boldsymbol{x},y)}[\ell(\boldsymbol{g}(\boldsymbol{x}), y)] \quad (4)$$

where $\mathbb{E}_{p(\boldsymbol{x},y)}[\cdot]$ is the expectation over $(\boldsymbol{x}, y) \sim p(\boldsymbol{x}, y)$. We use $R_{01}(\boldsymbol{g})$ to denote Eq. (4) when $\ell = \ell_{01}$ and call it the *classification error*.

The goal of multi-class classification is to learn $\boldsymbol{g}$ that minimizes the classification error $R_{01}(\boldsymbol{g})$. In optimization, we consider the minimization of the risk with a almost surely differentiable surrogate loss $R(\boldsymbol{g})$ instead to make the problem more tractable. Furthermore, since $p(\boldsymbol{x}, y)$ is usually unknown and there is no way to exactly evaluate $R(\boldsymbol{g})$, we minimize its empirical version calculated from the training data instead:

$$\widehat{R}(\boldsymbol{g}) := \frac{1}{n}\sum_{i=1}^{n} \ell(\boldsymbol{g}(\boldsymbol{x}_i), y_i), \quad (5)$$

where $\{(\boldsymbol{x}_i, y_i)\}_{i=1}^n$ are i.i.d. sampled from $p(\boldsymbol{x}, y)$. We call $\widehat{R}$ the *empirical risk*.

We would like to clarify some of the undefined terms used in the title and the introduction. The "train/test loss" is the empirical risk with respect to the surrogate loss function ℓ over the training/test data, respectively. We refer to the "training/test error" as the empirical risk with respect to ℓ_{01} over the training/test data, respectively (which is equal to one minus accuracy) [Zhang, 2004].

Finally, we formally define the Bayes risk as

$$R^* := \inf_{\boldsymbol{h}} R(\boldsymbol{h}), \quad (6)$$

where the infimum is taken over all vector-valued functions $\boldsymbol{h}: \mathbb{R}^d \to \mathbb{R}^K$. The Bayes risk is often referred to as the *Bayes error* if the zero-one loss is used:

$$\inf_{\boldsymbol{h}} R_{01}(\boldsymbol{h}). \quad (7)$$

3.2 Algorithm

With flexible models, $\widehat{R}(\boldsymbol{g})$ w.r.t. a surrogate loss can easily become small if not zero, as we mentioned in Section 1; see [C] in Fig. 1(a). We propose a method that "floods the bottom area and sinks the original empirical risk" as in Fig. 1(b) so that the empirical risk cannot go below the flooding level. More technically, if we denote the flooding level as b, our proposed training objective with flooding is a simple fix:

7

Definition 1. *The flooded empirical risk is defined as*[4]

$$\widetilde{R}(\boldsymbol{g}) = |\widehat{R}(\boldsymbol{g}) - b| + b. \quad (8)$$

Note that when $b = 0$, then $\widetilde{R}(\boldsymbol{g}) = \widehat{R}(\boldsymbol{g})$. The gradient of $\widetilde{R}(\boldsymbol{g})$ w.r.t. model parameters will point to the same direction as that of $\widehat{R}(\boldsymbol{g})$ when $\widehat{R}(\boldsymbol{g}) > b$ but in the opposite direction when $\widehat{R}(\boldsymbol{g}) < b$. This means that when the learning objective is above the flooding level, we perform gradient *descent* as usual (gravity zone), but when the learning objective is below the flooding level, we perform gradient *ascent* instead (buoyancy zone).

The issue is that in general, we seldom know the optimal flooding level in advance. This issue can be mitigated by searching for the optimal flooding level b^* with a hyper-parameter optimization technique. In practice, we can search for the optimal flooding level by performing the exhaustive search in parallel.

3.3 Implementation

For large scale problems, we can employ mini-batched stochastic optimization for efficient computation. Suppose that we have M disjoint mini-batch splits. We denote the empirical risk (5) with respect to the m-th mini-batch by $\widehat{R}_m(\boldsymbol{g})$ for $m \in \{1, \ldots, M\}$. Then, our mini-batched optimization performs gradient descent updates in the direction of the gradient of $\widetilde{R}_m(\boldsymbol{g})$. By the convexity of the absolute value function and Jensen's inequality, we have

$$\widetilde{R}(\boldsymbol{g}) \le \frac{1}{M}\sum_{m=1}^{M}\left(|\widehat{R}_m(\boldsymbol{g}) - b| + b\right). \quad (9)$$

This indicates that mini-batched optimization will simply minimize an upper bound of the full-batch case with $\widetilde{R}(\boldsymbol{g})$.

3.4 Theoretical Analysis

In the following theorem, we will show that the mean squared error (MSE) of the proposed risk estimator with flooding is smaller than that of the original risk estimator without flooding.

Theorem 1. *Fix any measurable vector-valued function $\boldsymbol{g}$. If the flooding level b satisfies $\widehat{R}(\boldsymbol{g}) < b < R(\boldsymbol{g})$, we have*

$$\mathrm{MSE}(\widehat{R}(\boldsymbol{g})) > \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})). \quad (10)$$

If $b \le \widehat{R}(\boldsymbol{g})$, we have

$$\mathrm{MSE}(\widehat{R}(\boldsymbol{g})) = \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})). \quad (11)$$

A proof is given in Appendix A. If we regard $\widehat{R}(\boldsymbol{g})$ as the training loss and $R(\boldsymbol{g})$ as the test loss, we would want b to be between those two for the MSE to improve.

[4] Strictly speaking, Eq. (1) is different from Eq. (8), since Eq. (1) can ignore constant terms of the original empirical risk. We will refer to Eq. (8) for the flooding operator for the rest of the paper.

8

A Proof of Theorem

Proof. If the flooding level is b, then the proposed flooding estimator is

$$\widetilde{R}(\boldsymbol{g}) = |\widehat{R}(\boldsymbol{g}) - b| + b. \quad (12)$$

Since the absolute operator can be expressed with a max operator with $\max(a, b) = \frac{a+b+|a-b|}{2}$, the proposed estimator can be re-expressed as

$$\widetilde{R}(\boldsymbol{g}) = 2\max(\widehat{R}(\boldsymbol{g}), b) - \widehat{R}(\boldsymbol{g}) = A - \widehat{R}(\boldsymbol{g}). \quad (13)$$

For convenience, we used $A = 2\max(\widehat{R}(\boldsymbol{g}), b)$. From the definition of MSE,

$$\begin{aligned} \mathrm{MSE}(\widehat{R}(\boldsymbol{g})) &= \mathbb{E}[(\widehat{R}(\boldsymbol{g}) - R(\boldsymbol{g}))^2] & (14)\\ \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})) &= \mathbb{E}[(\widetilde{R}(\boldsymbol{g}) - R(\boldsymbol{g}))^2] & (15)\\ &= \mathbb{E}[(A - \widehat{R}(\boldsymbol{g}) - R(\boldsymbol{g}))^2] & (16)\\ &= \mathbb{E}[A^2] - 2\mathbb{E}[A(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g}))] + \mathbb{E}[(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g}))^2]. & (17) \end{aligned}$$

We are interested in the sign of

$$\mathrm{MSE}(\widehat{R}(\boldsymbol{g})) - \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})) = \mathbb{E}[-4\widehat{R}(\boldsymbol{g})R(\boldsymbol{g}) - A^2 + 2A(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g}))]. \quad (18)$$

Define the inside of the expectation as $B = -4\widehat{R}(\boldsymbol{g})R(\boldsymbol{g}) - A^2 + 2A(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g}))$. B can be divided into two cases, depending on the outcome of the max operator:

$$\begin{aligned} B &= \begin{cases} -4\widehat{R}(\boldsymbol{g})R(\boldsymbol{g}) - 4\widehat{R}(\boldsymbol{g})^2 + 4\widehat{R}(\boldsymbol{g})(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g})) & \text{if } \widehat{R}(\boldsymbol{g}) \ge b \\ -4\widehat{R}(\boldsymbol{g})R(\boldsymbol{g}) - 4b^2 + 4b(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g})) & \text{if } \widehat{R}(\boldsymbol{g}) < b \end{cases} & (19)\\ &= \begin{cases} 0 & \text{if } \widehat{R}(\boldsymbol{g}) \ge b \\ -4\widehat{R}(\boldsymbol{g})R(\boldsymbol{g}) - 4b^2 + 4b(\widehat{R}(\boldsymbol{g}) + R(\boldsymbol{g})) & \text{if } \widehat{R}(\boldsymbol{g}) < b \end{cases} & (20)\\ &= \begin{cases} 0 & \text{if } \widehat{R}(\boldsymbol{g}) \ge b \\ -4(b - \widehat{R}(\boldsymbol{g}))(b - R(\boldsymbol{g})) & \text{if } \widehat{R}(\boldsymbol{g}) < b \end{cases}. & (21) \end{aligned}$$

The latter case becomes positive when $\widehat{R}(\boldsymbol{g}) < b < R(\boldsymbol{g})$. Therefore, when $\widehat{R}(\boldsymbol{g}) < b < R(\boldsymbol{g})$,

$$\begin{aligned} \mathrm{MSE}(\widehat{R}(\boldsymbol{g})) - \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})) &> 0 & (22)\\ \mathrm{MSE}(\widehat{R}(\boldsymbol{g})) &> \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})). & (23) \end{aligned}$$

When $b \le \widehat{R}(\boldsymbol{g})$,

$$\begin{aligned} \mathrm{MSE}(\widehat{R}(\boldsymbol{g})) - \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})) &= 0 & (24)\\ \mathrm{MSE}(\widehat{R}(\boldsymbol{g})) &= \mathrm{MSE}(\widetilde{R}(\boldsymbol{g})). & (25) \end{aligned}$$

□

19

B Bayes Risk for Gaussian Distributions

In this section, we explain in detail how we derived the Bayes risk with respect to the surrogate loss in the experiments with Gaussian data in Section 4.1. Since we are using the logistic loss in the synthetic experiments, the loss of the margin is

$$\ell(yg(\boldsymbol{x})) = \log(1 + \exp(-yg(\boldsymbol{x}))), \quad (26)$$

where $g(\boldsymbol{x}) : \mathbb{R}^d \to \mathbb{R}$ is a scalar instead of the vector definition that was used previously, because the synthetic experiments only consider binary classification. Take the derivative to derive,

$$\begin{aligned} \frac{\partial \mathbb{E}[\ell(yg(\boldsymbol{x}))]}{\partial g(\boldsymbol{x})} = \frac{\mathbb{E}[\log(1 + \exp(-yg(\boldsymbol{x})))]}{\partial g(\boldsymbol{x})} &= \mathbb{E}\left[\frac{-y\exp(-yg(\boldsymbol{x}))}{1 + \exp(-yg(\boldsymbol{x}))}\middle|\boldsymbol{x}\right] p(\boldsymbol{x}) & (27)\\ &= \mathbb{E}\left[\frac{-y}{1 + \exp(yg(\boldsymbol{x}))}\middle|\boldsymbol{x}\right] p(\boldsymbol{x}) & (28)\\ &= \mathbb{E}\left[\frac{y+1}{2}\frac{1}{\exp(-g(\boldsymbol{x})) + 1} + \frac{y-1}{2}\frac{-1}{\exp(g(\boldsymbol{x})) + 1}\middle|\boldsymbol{x}\right] p(\boldsymbol{x}) & (29)\\ &= \mathbb{E}\left[\frac{y+1}{2}\middle|\boldsymbol{x}\right]\frac{1}{\exp(-g(\boldsymbol{x})) + 1}p(\boldsymbol{x}) + \mathbb{E}\left[\frac{y-1}{2}\middle|\boldsymbol{x}\right]\frac{-1}{\exp(g(\boldsymbol{x})) + 1}p(\boldsymbol{x}) & (30)\\ &= p(y = +1|\boldsymbol{x})\frac{1}{\exp(-g(\boldsymbol{x})) + 1}p(\boldsymbol{x}) + p(y = -1|\boldsymbol{x})\frac{-1}{\exp(g(\boldsymbol{x})) + 1}p(\boldsymbol{x}) & (31) \end{aligned}$$

Set this to zero, divide by $p(\boldsymbol{x}) > 0$ to obtain,

$$\begin{aligned} p(y = -1|\boldsymbol{x})\frac{1}{\exp(-g(\boldsymbol{x})) + 1} &= p(y = +1|\boldsymbol{x})\frac{1}{\exp(g(\boldsymbol{x})) + 1} & (32)\\ \exp(g(\boldsymbol{x})) &= \frac{p(y = +1|\boldsymbol{x})}{p(y = -1|\boldsymbol{x})} & (33)\\ g(\boldsymbol{x}) &= \log\frac{p(y = +1|\boldsymbol{x})}{p(y = -1|\boldsymbol{x})} & (34) \end{aligned}$$

Since we are interested in the surrogate loss under this classifier, we plug this into the logistic loss, to obtain the Bayes risk,

$$\mathbb{E}[\ell(yg(\boldsymbol{x}))] = \mathbb{E}\left[\log(1 + \frac{p(-y|\boldsymbol{x})}{p(y|\boldsymbol{x})})\right] = \mathbb{E}\left[\log(\frac{1}{p(y|\boldsymbol{x})})\right] = \mathbb{E}[-\log p(y|\boldsymbol{x})]. \quad (35)$$

In the experiments in Section 4.1, we report the empirical version of this with the test dataset as the Bayes risk.

20

Novelty

- The importance of telling a compelling story
  - **ICML with one line of code!** $\tilde{J}(\boldsymbol{\theta}) = |J(\boldsymbol{\theta}) - b| + b$
  - Amazing demonstration

Table 3: Results with benchmark datasets. We report classification accuracy for all combinations of weight decay (✓ and ×), early stopping (✓ and ×) and flooding (✓ and ×). The second column shows the training/validation split used for the experiment. W stands for weight decay, E stands for early stopping, and F stands for flooding. "–" means that flooding level of zero was optimal. "N/A" means that we skipped the experiments because zero weight decay was optimal in the case without flooding. The best and equivalent are shown in **bold** by comparing "with flooding" and "without flooding" for two columns with the same setting for W and E, e.g., the first and fifth columns out of the 8 columns. The best performing combination is highlighted.

Dataset	tr/va split	W: / E: / F:	× / × / ×	× / ✓ / ×	✓ / × / ×	✓ / ✓ / ×	× / × / ✓	× / ✓ / ✓	✓ / × / ✓	✓ / ✓ / ✓
MNIST	0.8		98.32%	98.30%	**98.51%**	98.42%	**98.46%**	**98.53%**	98.50%	**98.48%**
	0.4		97.71%	97.70%	97.82%	**97.91%**	**97.74%**	**97.85%**	–	97.83%
Fashion-MNIST	0.8		89.34%	89.36%	N/A	N/A	–	–	N/A	N/A
	0.4		88.48%	88.63%	88.60%	88.62%	–	–	–	–
Kuzushiji-MNIST	0.8		91.63%	91.62%	91.63%	91.71%	**92.40%**	**92.12%**	**92.11%**	**91.97%**
	0.4		89.18%	89.18%	89.58%	89.73%	**90.41%**	**90.15%**	**89.71%**	**89.88%**
CIFAR-10	0.8		**73.59%**	73.36%	73.65%	73.57%	73.06%	**73.44%**	–	**74.41%**
	0.4		66.39%	66.63%	69.31%	69.28%	**67.20%**	**67.58%**	–	–
CIFAR-100	0.8		42.16%	42.33%	42.67%	42.45%	**42.50%**	**42.36%**	–	–
	0.4		34.27%	34.34%	37.97%	38.82%	**34.99%**	**35.14%**	–	–
SVHN	0.8		92.38%	92.41%	93.20%	92.99%	**92.78%**	**92.79%**	–	**93.42%**
	0.4		90.32%	90.35%	90.43%	90.49%	**90.57%**	**90.61%**	**91.16%**	**91.21%**

helps the late-stage training improve test accuracy.

We also conducted experiments with early stopping, meaning that we chose the model that recorded the best validation accuracy during training. The results are reported in sub-table (B) of Table 2. Compared with sub-table (A), we see that early stopping improves the baseline method without flooding well in many cases. This indicates that training longer without flooding was harmful in our experiments. On the other hand, the accuracy for flooding combined with early stopping is often close to that with early stopping, meaning that training until the end with flooding tends to be already as good as doing so with early stopping. The table shows that flooding often improves or retains the test accuracy of the baseline method without flooding even after deploying early stopping. Flooding does not hurt performance but can be beneficial for methods used with early stopping.

11

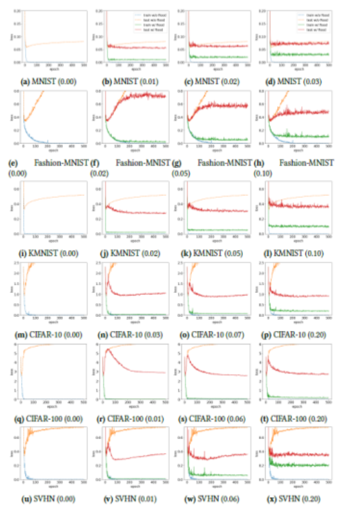

Figure 4: Learning curves of training and test loss. The first figure in each row is the learning curves without flooding. The 2nd, 3rd, and 4th columns show the results with different flooding levels. The flooding level increases towards the right-hand side.

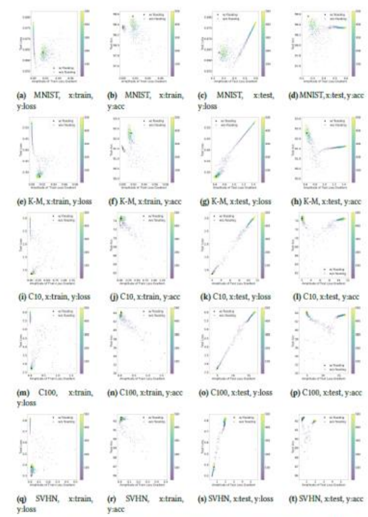

Figure 5: Relationship between test performance (loss or accuracy) and amplitude of gradient (with training or test loss). Each point ("o" or "+") in the figures corresponds to a single model at a certain epoch. We remove the first 5 epochs and plot the rest. "o" is used for the method with flooding and "+" is used for the method without flooding. The large black "o" and "+" show the epochs with early stopping. The color becomes lighter (purple → yellow) as the training proceeds. K-M, C10, and C100 stand for Kuzushiji-MNIST, CIFAR-10, and CIFAR-100.

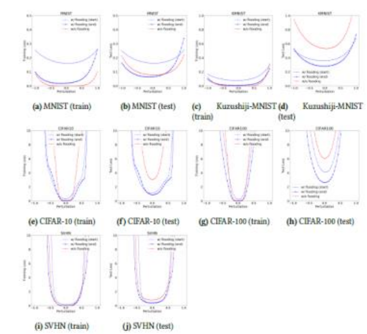

Figure 6: One-dimensional visualization of flatness. We visualize the training/test loss with respect to perturbation. We depict the results for 3 models: the model when the empirical risk with respect to training data is below the flooding level for the first time during training (dotted blue), the model at the end of training with flooding (solid blue), and the model at the end of training without flooding (solid red).

27

Notes on CV Paper Writing

Xiang Gao, Lecturer

E-mail: xgao@ouc.edu.cn | Web: https://ouc-xgao.github.io/

College of Engineering, Ocean University of China

Previous Review

- Several Important Concepts
  - CV, IF, h5-index, CCF recommended list
- Recommended Journals & Conferences
  - Journals and conferences in CV and other related areas
- Reviewing Process
  - Journal/conference reviewing process
  - Something more about conference papers and CVPR2020
- Journal Papers *vs.* Conference Papers
  - Evaluation, reviewing period, quality, *etc.*
- How to Structure a Paper?
  - **Introduction**, related work, method and evaluation
- *How to Get Your CVPR Paper Rejected?*
- Review Form of a Conference Paper
  - **Novelty**: Make a compelling story with strong supporting evidence

Experimental Validation

- Common dataset

Experimental Validation

- Common dataset
- Baseline experiment

Experimental Validation

- Common dataset
- Baseline experiment
- Killer dataset

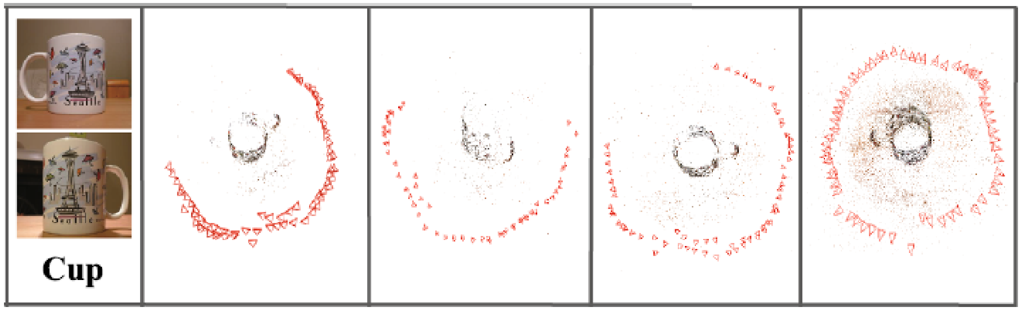

Experimental Validation

- Common dataset
- Baseline experiment
- Killer dataset
- Large-scale experiment

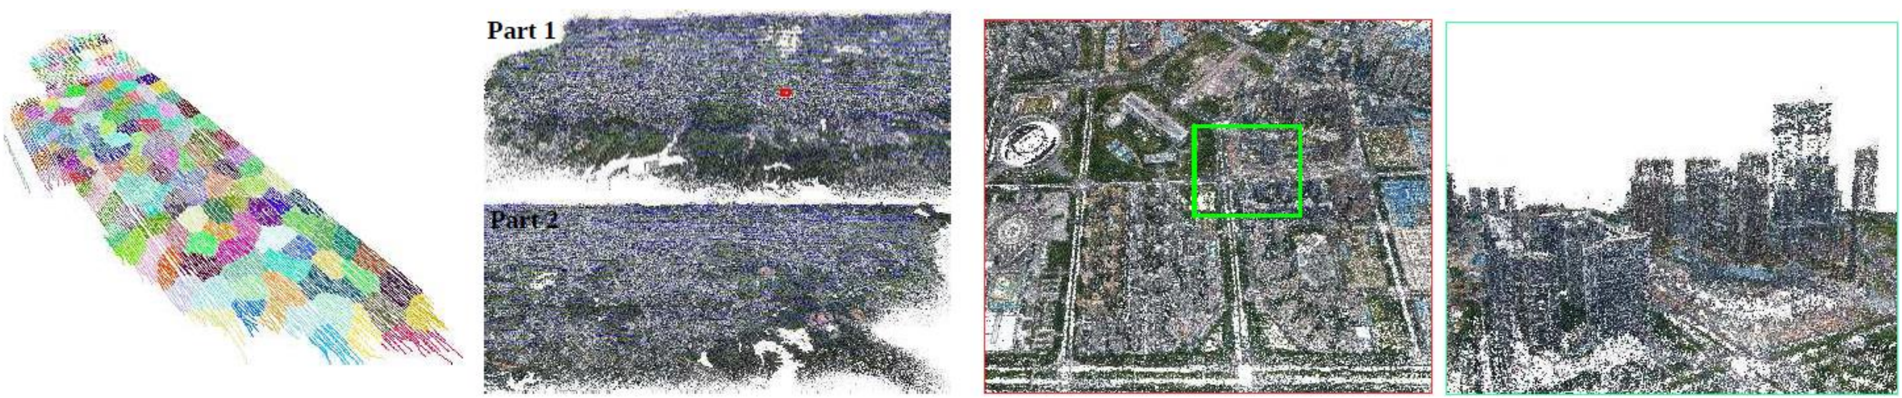

Experimental Validation

- Common dataset
- Baseline experiment
- Killer dataset
- Large-scale experiment
- Evaluation metric

Experimental Validation

- Common dataset
- Baseline experiment
- Killer dataset
- Large-scale experiment
- Evaluation metric
- Friendly fire

Learn from Reviews

- You will never know what would happen in your reviews

Learn from Reviews

- You will never know what would happen in your reviews
  - Me: Here is a faster horse

Learn from Reviews

- You will never know what would happen in your reviews
  - Me: Here is a faster horse
  - R1: You should have used my donkey

Learn from Reviews

- You will never know what would happen in your reviews
  - Me: Here is a faster horse
  - R1: You should have used my donkey
  - R2: This is not a horse, it’s a mule

Learn from Reviews

- You will never know what would happen in your reviews
  - Me: Here is a faster horse
  - R1: You should have used my donkey
  - R2: This is not a horse, it's a mule
  - R3: I want a unicorn!

Learn from Reviews

- You will never know what would happen in your reviews
  - Me: Here is a faster horse.
  - R1: You should have used my donkey.
  - R2: This is not a horse, it's a mule.
  - R3: I want a unicorn!
- Be able to find valuable comments/suggestions (even few) from a mass of (useless) ones
  - Submit your manuscript to good journals/conferences
  - Focus on the process rather than the outcome
  - Accept the valuable comments/suggestions to improve the manuscript and unhesitatingly discard others

Learn from Reviews

- My experience
  - ICIP 2017 (4 pages): Rejected (2017.01 → 2017.04)

Learn from Reviews

- My experience
  - ICIP 2017 (4 pages): Rejected (2017.01 → 2017.04)
  - 3DV 2017 (8 pages): Rejected (2017.04 → 2017.07)

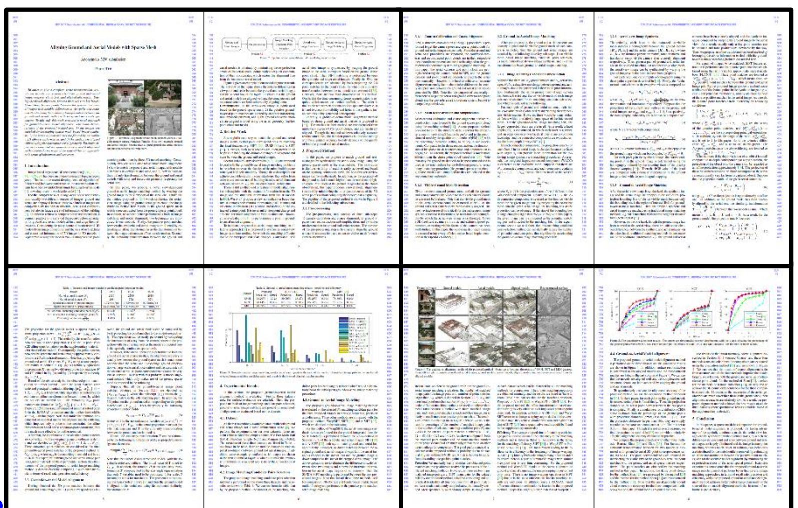

Learn from Reviews

- My experience
  - ICIP 2017 (4 pages): Rejected (2017.01 → 2017.04)
  - 3DV 2017 (8 pages): Rejected (2017.04 → 2017.07)
  - ISPRS P&RS (12 pages):
    - Submitted (2017.10) → Major Revision (2018.01) → Minor Revision (2018.04) → Accepted (2018.05)

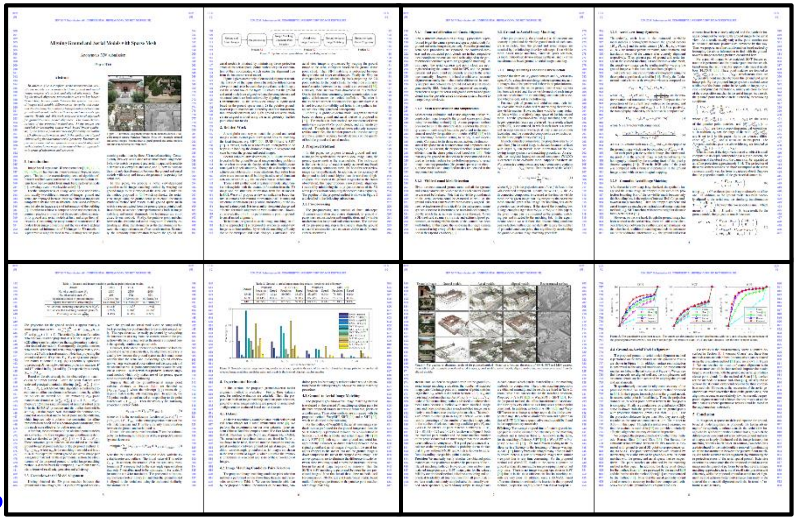

Learn from Reviews

- My experience
  - Submission (2021.08.06, **7 pages**) → Minor Revision (2021.09.08)

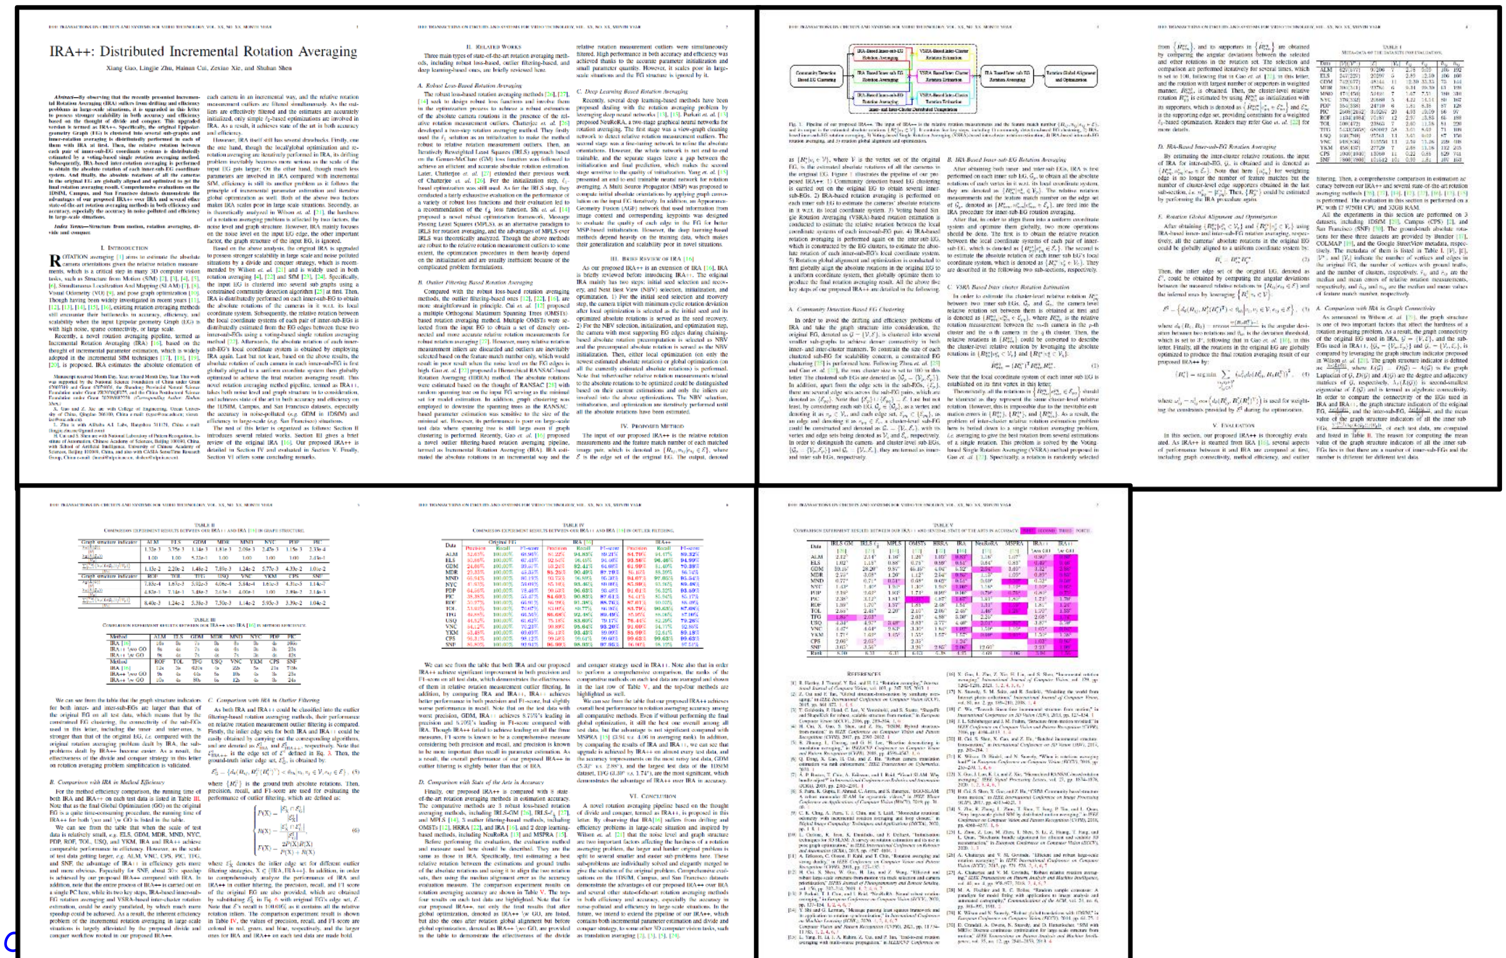
IRA++: Distributed Incremental Rotation Averaging

Learn from Reviews

- My experience
  - Submission (2021.08.06, **7 pages**) → Minor Revision (2021.09.08)
  - Revision Submitted (2021.09.17, **8 pages**) → Accept (2021.10.06)

Learn from Reviews

- My experience
  - Submission (2021.08.06, **7 pages**) → Minor Revision (2021.09.08)
  - Revision Submitted (2021.09.17, **8 pages**) → Accept (2021.10.06)

Dear Editors and Reviewers,

Many thanks for the valuable comments, the manuscript has been revised accordingly. The following is our reply to the specific comments by the editors and reviewers. In the revised text, the changed parts are marked in RED.

For the Reviewer #1:

R1Q1: It seems that the proposed IRA++ is a simple combination of HRRA and IRA. In particular, IRA++ uses the same framework of HRRA (including finding the cluster graph and estimation of inter-cluster rotations). The only difference is that IRA++ replaces the RANSAC-based rotation averaging in some places of HRRA with the IRA solver. Thus, it looks to me that the novelty of the proposed method is limited as the proposed method is a heuristic combination of existing methods. I hope the authors can clarify the contribution and novelty of this work.

A: Thanks for the comment. In the revised text, the contributions of this work are made more clarified. They lie in the following three folds: *1) A upgraded version of IRA, termed as IRA++, is presented based on the thought of divide and conquer, to deal with the drifting and efficiency issues of IRA in large-scale situations. 2) The original IRA method is tightly integrated into this novel pipeline to solve the low-level inner-sub-EG rotation averaging problems locally and the high-level inter-sub-EG rotation averaging problem globally. 3) Our proposed IRA++ is thoroughly evaluated and achieves overall best performance in both efficiency and accuracy compared with several other state-of-the-art rotation averaging methods.* In addition, note that though IRA++ and HRRA share similar divide and conquer strategy and framework, their motivations and technical insights are different. For HRRA, the core rotation averaging solver, RANSAC-based rotation averaging, could be regarded as a global solver, as it performs outlier filtering and rotation optimization globally. In that way, the clustering operation in HRRA is used to constrain the size of the random spanning tree, by which the all-inlier minimal set is more likely to be selected. However, For IRA++, the core rotation averaging solver is IRA, which is an incremental solver and suffers from drifting and efficiency issues in large-scale situations. As a result, the divide and conquer strategy in IRA++ is mainly employed to deal with these exclusive issues of incremental parameter estimation pipeline.

R1Q2: The estimation accuracy of IRA++ without global optimization is still comparable to that of IRA. Thus, even though it achieves the top accuracy, the improvement seems incremental compared to IRA. This implies that the final global optimization plays an important role on improving the accuracy. Indeed, the difference in accuracy between IRA++ with/without GO is more significant than the difference between IRA++ without GO and IRA. However, this observed importance of the final global optimization is a bit contradictory to the motivation of this work. In particular, from table 2 the authors claim that partitioning the graph nodes improves the graph connectivity statistics of Wilson et.al. and thus the divide-and-conquer strategy makes rotation averaging easier. From this point of view, the final global optimization would make the connectivity statistics worse again, but it turns out that this step has more significant improvement on the overall accuracy and I find this a bit confusing. It is possible that the inter-cluster rotations are not estimated accurately due to the sparse connections between the clusters, or it is possible that the graph structure indicator does not fully explain the hardness of the rotation averaging problem in the new distributed setting (for example statistics based on graph connection Laplacian could be better suited, as it simultaneously considers the graph connectivity and rotation consistency). In any case, I believe that a more careful investigation and explanation is needed in this regard.

A: Thanks for the comment. Please note that global optimization is usually a necessary step in divide and conquer strategy-based pipeline. That is because without considering all the constraints globally, the divide and conquer strategy-based pipeline' solution is easy to fall into local minima, which would result in the loss of estimation accuracy. In addition, for the IRA++ method considered here, though the accuracy of IRA++ \wo GO is not extremely high, it is good enough to **serve as initialization** and **filter out outlier** for the global optimization to achieve obvious accuracy improvement. Please note also that global optimization is iteratively performed on the original IRA, including the final global optimization with all absolute rotation estimated, and thus it is not quite fair to directly compare the accuracy of IRA and IRA++ \wo GO. As for the graph connectivity evaluation, the indicator we used stems from the graph Laplacian, and it is only used to compare the graph connectivity of the EGs used in IRA and IRA++. In addition, we think that the accuracy improvement after global optimization of IRA++ is not contradictory to the problem hardness discussed in this work. That is because when discussing the hardness of a rotation averaging problem, the estimation initial guess and measurement outlier label is supposed to be unknown. However, the problem of IRA++ with only global optimization left is not the case. With the quite accurate estimation initial guess and measurement outlier label, its hardness is undoubtedly much lower than the original problem of IRA and the sub-problems of IRA++.

R1Q3: In table 5, none of other distributed methods (except HRRA) are compared. It seems that other global methods can also be applied to inner/inter-cluster rotation estimation, following the same framework of this work and HRRA. Have the authors tried other methods (e.g. IRLS-$\ell_{\frac{1}{2}}$) using the same divide and conquer strategy? I am curious about the speed and accuracy of such implementation and how that compares with IRA++. This could verify the critical component (IRA or distributed implementation) of the proposed method.

A: Thanks for the comment. We locally implement a distributed version of the IRLS-$\ell_{\frac{1}{2}}$, which is achieved by alternating the IRA-based inner- and inter-sub-EG rotation averaging steps in IRA++ to IRLS-$\ell_{\frac{1}{2}}$-based ones. In the revised text the distributed IRLS-$\ell_{\frac{1}{2}}$ is compared with IRA, IRA++, and IRLS-$\ell_{\frac{1}{2}}$ in both efficiency (TABLE IV) and accuracy (TABLE VI). From the comparison experiments we can see that though distributed IRLS-$\ell_{\frac{1}{2}}$ achieves better performance in both efficiency and accuracy than IRLS-$\ell_{\frac{1}{2}}$, our proposed IRA++ still achieves top performance, which demonstrate the effectiveness of both the divide and conquer pipeline and the IRA-based key steps.

R1Q4: The comparison between IRLS-GM and IRLS-$\ell_{\frac12}$ is contradictory to that of "robust rotation averaging, 2018". In that paper, IRLS-$\ell_{\frac12}$ works significantly better than IRLS-GM on 1DSfM dataset. However, table 5 shows a different story using the same dataset. Is there any difference between the implementation of IRLS in the two papers?

A: Thanks for the comment. The results of IRLS-GM of the originally submitted manuscript are obtained by our locally running the source code provided by the authors. In the revised text, they have been modified to the ones listed in the paper of "Robust relative rotation averaging, 2018".

R1Q5: It would be better if the authors can report the size of the clusters (or minimal, maximal, and median size) and the number of clusters for each dataset in the revised manuscript. The size and number of clusters will certainly affect the speed (and possibly accuracy) of the

Learn from Reviews

- My experience
  - Submission (2021.08.06, **7 pages**) → Minor Revision (2021.09.08)
  - Revision Submitted (2021.09.17, **8 pages**) → Accept (2021.10.06)

Dear Editors and Reviewers,

Many thanks for the valuable comments, the manuscript has been revised accordingly. The following is our reply to the specific comments by the editors and reviewers. In the revised text, the changed parts are marked in RED.

For the Reviewer #1:

R1Q1: It seems that the proposed IRA++ is a simple combination of HRRA and IRA. In particular, IRA++ uses the same framework of HRRA (including finding the cluster graph and estimation of inter-cluster rotations). The only difference is that IRA++ replaces the RANSAC-based rotation averaging in some places of HRRA with the IRA solver. Thus, it looks to me that the novelty of the proposed method is limited as the proposed method is a heuristic combination of existing methods. I hope the authors can clarify the contribution and novelty of this work.

A: Thanks for the comment. In the revised text, the contributions of this work are made more clarified. They lie in the following three folds: *1) A upgraded version of IRA, termed as IRA++, is presented based on the thought of divide and conquer, to deal with the drifting and efficiency issues of IRA in large-scale situations. 2) The original IRA method is tightly integrated into this novel pipeline to solve the low-level inner-sub-EG rotation averaging problems locally and the high-level inter-sub-EG rotation averaging problem globally. 3) Our proposed IRA++ is thoroughly evaluated and achieves overall best performance in both efficiency and accuracy compared with several other state-of-the-art rotation averaging methods.* In addition, note that though IRA++ and HRRA share similar divide and conquer strategy and framework, their motivations and technical insights are different. For HRRA, the core rotation averaging solver, RANSAC-based rotation averaging, could be regarded as a global solver, as it performs outlier filtering and rotation optimization globally. In that way, the clustering operation in HRRA is used to constrain the size of the random spanning tree, by which the all-inlier minimal set is more likely to be selected. However, For IRA++, the core rotation averaging solver is IRA, which is an incremental solver and suffers from drifting and efficiency issues in large-scale situations. As a result, the divide and conquer strategy in IRA++ is mainly employed to deal with these exclusive issues of incremental parameter estimation pipeline.

R1Q2: The estimation accuracy of IRA++ without global optimization is still comparable to that of IRA. Thus, even though it achieves the top accuracy, the improvement seems incremental compared to IRA. This implies that the final global optimization plays an important role on improving the accuracy. Indeed, the difference in accuracy between IRA++ with/without GO is more significant than the difference between IRA++ without GO and IRA. However, this observed importance of the final global optimization is a bit contradictory to the motivation of this work. In particular, from table 2 the authors claim that partitioning the graph nodes improves the graph connectivity statistics of Wilson et.al. and thus the divide-and-conquer strategy makes rotation averaging easier. From this point of view, the final global optimization would make the connectivity statistics worse again, but it turns out that this step has more significant improvement on the overall accuracy and I find this a bit confusing. It is possible that the inter-cluster rotations are not estimated accurately due to the sparse connections between the clusters, or it is possible that the graph structure indicator does not fully explain the hardness of the rotation averaging problem in the new distributed setting (for example statistics based on graph connection Laplacian could be better suited, as it simultaneously considers the graph connectivity and rotation consistency). In any case, I believe that a more careful investigation and explanation is needed in this regard.

A: Thanks for the comment. Please note that global optimization is usually a necessary step in divide and conquer strategy-based pipeline. That is because without considering all the constraints globally, the divide and conquer strategy-based pipeline' solution is easy to fall into local minima, which would result in the loss of estimation accuracy. In addition, for the IRA++ method considered here, though the accuracy of IRA++ \wo GO is not extremely high, it is good enough to **serve as initialization** and **filter out outlier** for the global optimization to achieve obvious accuracy improvement. Please note also that global optimization is iteratively performed on the original IRA, including the final global optimization with all absolute rotation estimated, and thus it is not quite fair to directly compare the accuracy of IRA and IRA++ \wo GO. As for the graph connectivity evaluation, the indicator we used stems from the graph Laplacian, and it is only used to compare the graph connectivity of the EGs used in IRA and IRA++. In addition, we think that the accuracy improvement after global optimization of IRA++ is not contradictory to the problem hardness discussed in this work. That is because when discussing the hardness of a rotation averaging problem, the estimation initial guess and measurement outlier label is supposed to be unknown. However, the problem of IRA++ with only global optimization left is not the case. With the quite accurate estimation initial guess and measurement outlier label, its hardness is undoubtedly much lower than the original problem of IRA and the sub-problems of IRA++.

R1Q3: In table 5, none of other distributed methods (except HRRA) are compared. It seems that other global methods can also be applied to inner/inter-cluster rotation estimation, following the same framework of this work and HRRA. Have the authors tried other methods (e.g. IRLS-$\ell_{\frac{1}{2}}$) using the same divide and conquer strategy? I am curious about the speed and accuracy of such implementation and how that compares with IRA++. This could verify the critical component (IRA or distributed implementation) of the proposed method.

A: Thanks for the comment. We locally implement a distributed version of the IRLS-$\ell_{\frac{1}{2}}$, which is achieved by alternating the IRA-based inner- and inter-sub-EG rotation averaging steps in IRA++ to IRLS-$\ell_{\frac{1}{2}}$-based ones. In the revised text the distributed IRLS-$\ell_{\frac{1}{2}}$ is compared with IRA, IRA++, and IRLS-$\ell_{\frac{1}{2}}$ in both efficiency (TABLE IV) and accuracy (TABLE VI). From the comparison experiments we can see that though distributed IRLS-$\ell_{\frac{1}{2}}$ achieves better performance in both efficiency and accuracy than IRLS-$\ell_{\frac{1}{2}}$, our proposed IRA++ still achieves top performance, which demonstrate the effectiveness of both the divide and conquer pipeline and the IRA-based key steps.

R1Q4: The comparison between IRLS-GM and IRLS-$\ell_{\frac12}$ is contradictory to that of "robust rotation averaging, 2018". In that paper, IRLS-$\ell_{\frac12}$ works significantly better than IRLS-GM on 1DSfM dataset. However, table 5 shows a different story using the same dataset. Is there any difference between the implementation of IRLS in the two papers?

A: Thanks for the comment. The results of IRLS-GM of the originally submitted manuscript are obtained by our locally running the source code provided by the authors. In the revised text, they have been modified to the ones listed in the paper of "Robust relative rotation averaging, 2018".

R1Q5: It would be better if the authors can report the size of the clusters (or minimal, maximal, and median size) and the number of clusters for each dataset in the revised manuscript. The size and number of clusters will certainly affect the speed (and possibly accuracy) of the

R2Q2: From part 3, considering that the whole pipeline is constructed with the existing methods or frameworks, the technique novelty may be limited. As the highlight of this method: a new divide-and-conquer strategy, deserving more discussion. (e.g. the technique intuition, the

A: Thanks for the comment. Table II has been modified in the revised text for better clarification according to the reviewer's suggestion.

Learn from Reviews

- [A copy of a valuable review]

1. The noun " method" is used many many times through abstract. I would try to reformulate some sentences and use the nouns " scheme" or " discretization" as alternatives. 这条意见是讲英文论文中词汇要尽可能丰富多彩一些，不要总是用一个词。所以后来我平时读英语文章的时候会注意搜集一些近义、同义词，比如: 采用，adopt use utilize employ;表示、描述\说明， represent show indicate depict illustrate give elucidate;提供、提出， contribute provide propose represent; （谁）发展、提供、提出、展示（了什么方法...): report show give offer develop image devise; 要求：require ask demand call. 这些词汇记录小本子上，等论文基本完稿之后拿出来对照看一看，将论文中重复的相关词汇替换下来，很有帮助。
2. Second line: " ...to represent THE solution IN each cell..." ; i.e. I would add " the" and replace " at" with " in" . 原文中缺少the。一般特指的、第一次出现的名词重复出现的时候都要用the。其它的用法般语法书中都有，不详述。然后是介词in,at,of,on,for之类的要尽量准确。
3. Third line: please add a space between freedoms and the first parenthesis. 原文The degree of freedoms(DOFs)，(DOFs)前面要留下空格。（那位审稿人细心程度可见一斑）。
4. 如果水平不足以把握长句，尽量用短句表明清楚意思。好几条意见都是关于长句的理解问题。
5. The adverb " So" is used to start and connect two sentences. I would replace it with formal adverbs like " Therefore, In fact, etc." .
6. I believe that " differentiate operators" is not correct. I think the Author should use " differential operators" or " differentiation operators" . Please, apply this correction through all the manuscript.一些词汇不准确。我后来对于自己感觉比较模糊的词汇，都会上网搜索一下，看看一般大家都怎样用的。
7. End of the fourth line/beginning of the fifth line: the comma is preceded by a wrong additional space. 标点不能出现在行首。
8. The acronyms ENO and WENO are immediately used without explicitly state what they mean. Moreover, the full name for discontinuous Galerkin is used without introducing the acronym which will be useful later on. I think the Authors should try to be

Writing: General Ideas

- Write early
  - Write the paper *as early as possible*, sometimes even *before* starting the research work
  - Good writing is *re-writing*, and it often helps to put the paper down and return to it later

Writing: General Ideas

- Write early
  - Write the paper *as early as possible*, sometimes even *before* starting the research work
  - Good writing is *re-writing*, and it often helps to put the paper down and return to it later
- Writing modes
  - Conventional mode: idea → research → writing

Writing: General Ideas

- Write early
  - Write the paper *as early as possible*, sometimes even *before* starting the research work
  - Good writing is *re-writing*, and it often helps to put the paper down and return to it later
- Writing modes
  - Conventional mode: idea → research → writing
  - Simon Peyton Jones: idea → writing → research

Writing: General Ideas

- Write early
  - Write the paper *as early as possible*, sometimes even *before* starting the research work
  - Good writing is *re-writing*, and it often helps to put the paper down and return to it later
- Writing modes
  - Conventional mode: idea → research → writing
  - Simon Peyton Jones: idea → writing → research
  - Ming-Hsuan Yang: idea → writing → research → revising

Writing: Some Details

- Title
  - Cool titles grab people's attention
    - How to get your CVPR paper accepted/**rejected**?
    - Shiftable multiscale transforms/**What is wrong with wavelets?**

Writing: Some Details

- Title
  - Cool titles grab people's attention
    - How to get your CVPR paper accepted/**rejected**?
    - Shiftable multiscale transforms/**What is wrong with wavelets?**
- Equations
  - The formulation of equations should conform to the specification
  - The number of equations should be in a proper range
  - Many good papers have no or few equations: CVPR 2013 best paper
  - Paper is not product manual, there is no need to provide every detail

Writing: Some Details

- CVPR 2013 best paper

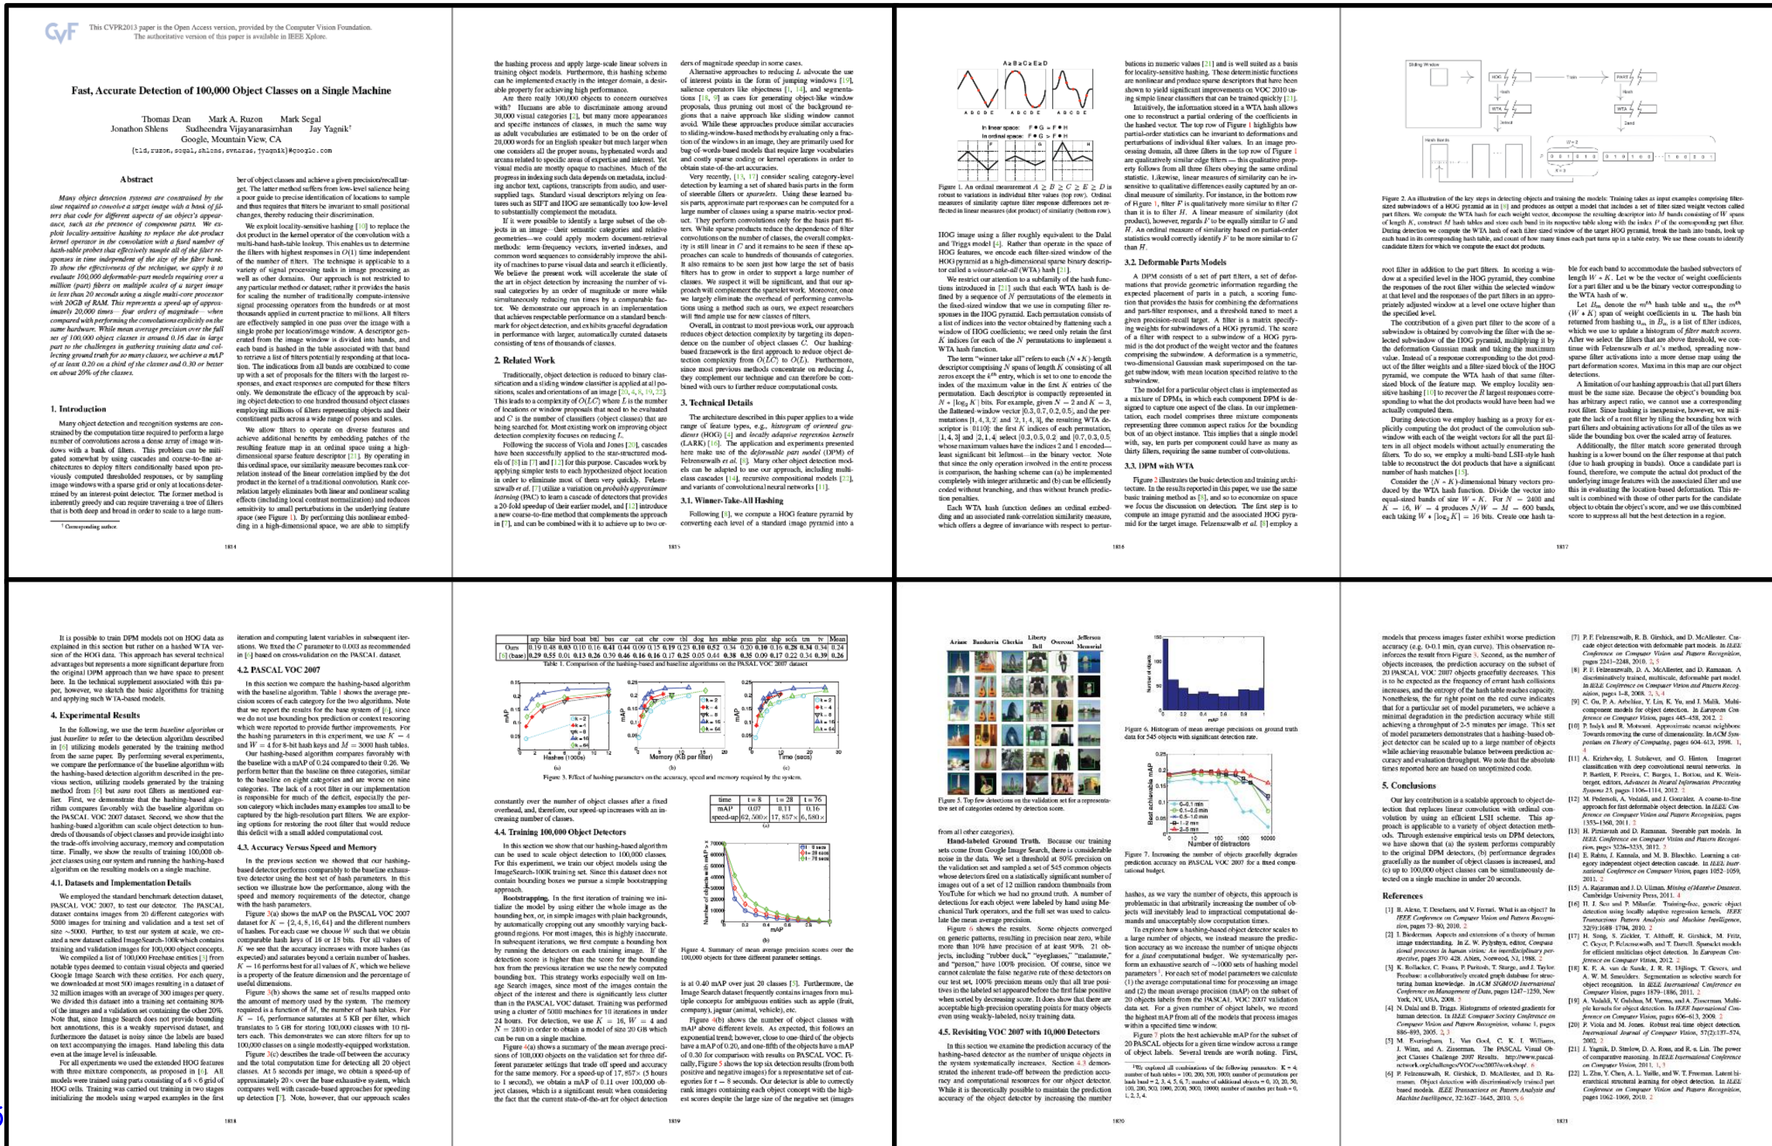
Fast, Accurate Detection of 100,000 Object Classes on a Single Machine

Thomas Dean, Mark A. Ruzon, Mark Segal, Jonathon Shlens, Sudheendra Vijayanarasimhan, Jay Yagnik

Google, Mountain View, CA

Writing: Some Details

- Figures and captions
  - **Clear** and sufficient

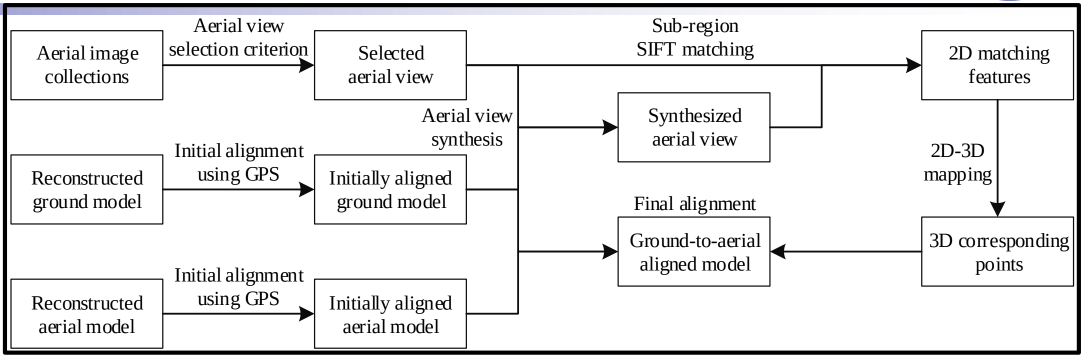

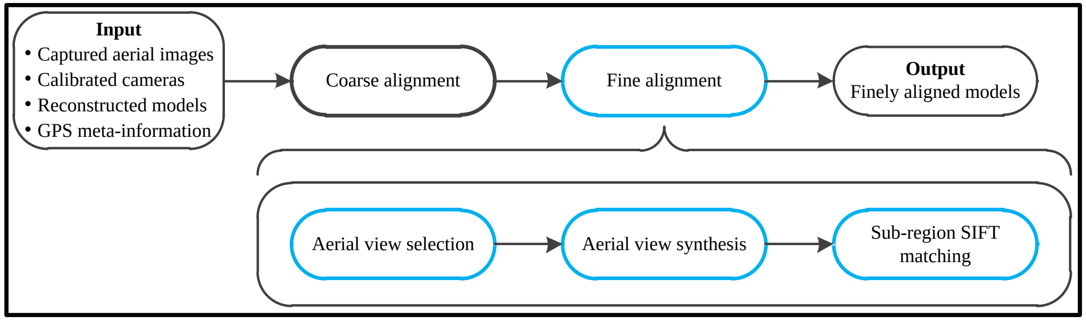

Writing: Some Details

- Figures and captions
  - **Clear** and sufficient
  - Many CV papers prefer putting a figure in the **top-right** corner of the first page for overall illustration


Height and Uprightness Invariance for 3D Prediction from a Single View

Manel Baradad Antonio Torralba
Massachusetts Institute of Technology

Where am I looking at? Joint Location and Orientation Estimation by Cross-View Matching

Information-Driven Direct RGB-D Odometry

Dynamic Fluid Surface Reconstruction Using Deep Neural Network

Simron Thapa Nianyi Li Jinwei Ye
Louisiana State University, Baton Rouge, LA 70803, USA

Writing: Some Details

- Figures and Captions
  - Paper should be easy to read in a big hurry and still learn its main points
  - Probably most of your readers will **skim** the paper
  - The figures and captions can help tell the story

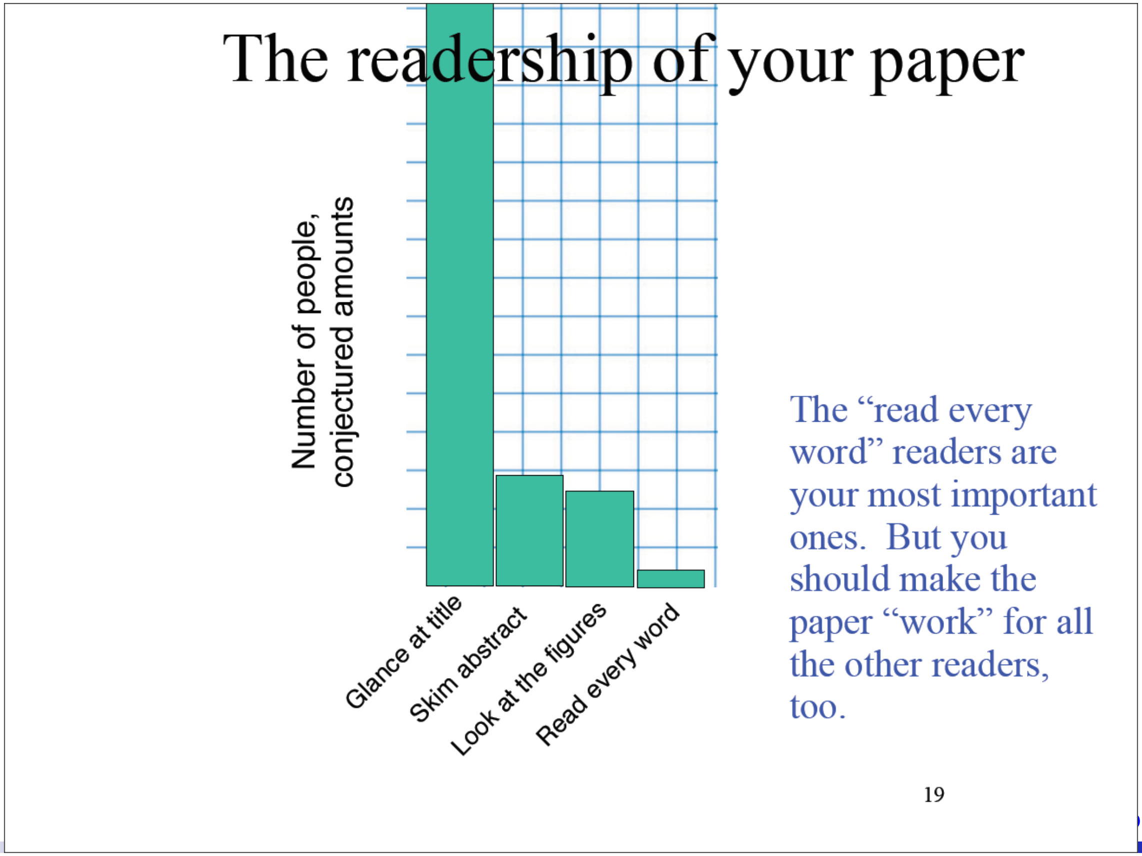

Writing: Some Details

- Figures and Captions
  - So the figure captions should be **self-contained** and the caption should tell the reader what to notice about the figure

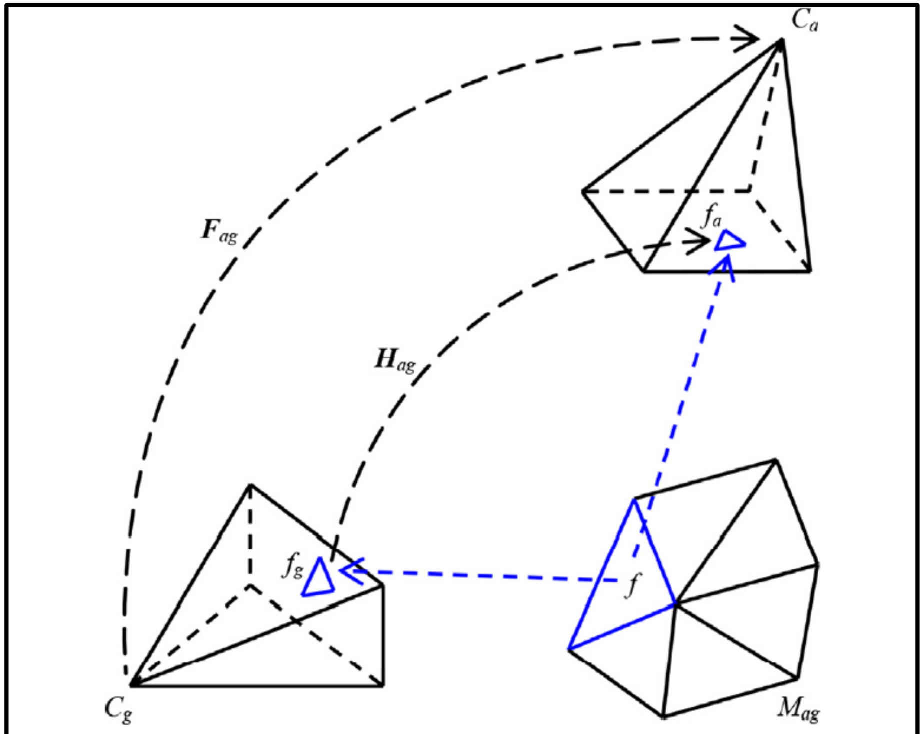

Fig. 3. Schematic diagram of the proposed aerial-view synthesis method. C_g and C_a are a pair of ground and aerial cameras, and $\boldsymbol{F}_{ag}$ is the fundamental matrix between them. M_{ag} is the co-visible mesh of C_g and C_a. f is a facet in M_{ag}, and f_g and f_a are the projections of f in C_g and C_a, respectively. $\boldsymbol{H}_{ag}$ is the homography between f_g and f_a induced by the facet f. Note that each facet in M_{ag} induces a unique homography.

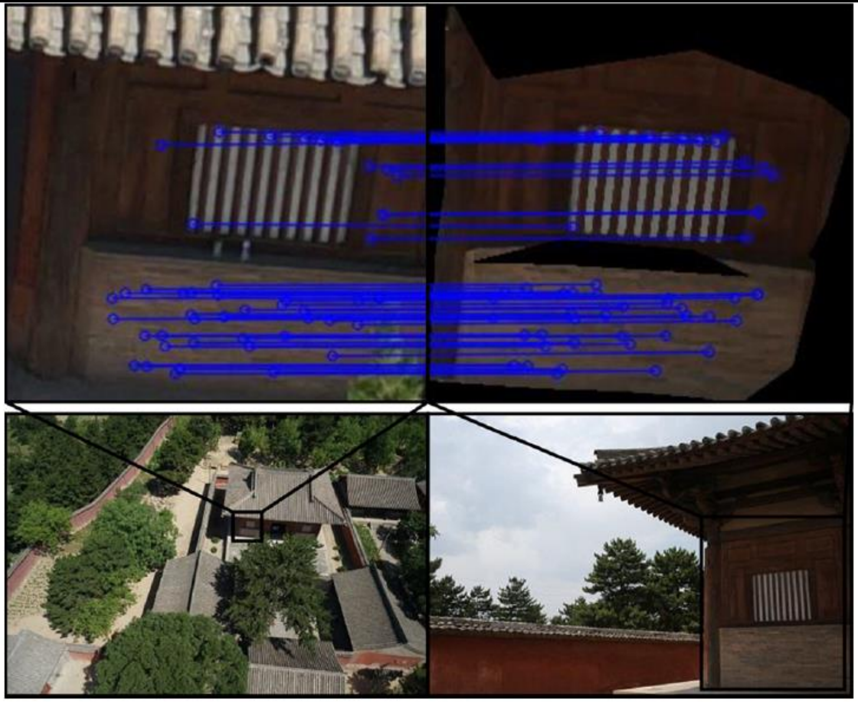

Fig. 4. An example of ground-to-aerial image matching result. The first row is the matching result between the RoIs (defined in Section 3.2.3) of the aerial and synthetic images, where the blue segments denote the point matches. The second row is the original aerial and ground image matching pair, where the black rectangles denote the RoIs for image matching.

By X. Gao

Writing: Some Details

- Common mistakes
  - Typos
  - Unsupported claims
  - Unnecessary adjectives (superior!)
  - “the” or not: whether **definite**
  - Inanimate objects with verbs
  - Inconsistent usage of words
  - Bad references
  - Laundry list of related work
  - Repeated boring statements
  - Needless words

Writing: Some Details

- Common mistakes
  - Typos
  - Unsupported claims
  - Unnecessary adjectives (superior!)
  - "the" or not: whether **definite**
  - Inanimate objects with verbs
  - Inconsistent usage of words
  - Bad references
  - Laundry list of related work
  - Repeated boring statements
  - Needless words

there is no doubt that *vs.* **no doubt (doubtless)**

used for … purposes *vs.* **used for …**

he is a man who *vs.* **he**

in a … manner **vs. …ly**

Paper Gestalt

- Most CVers are the members of appearance club
  - Motivation

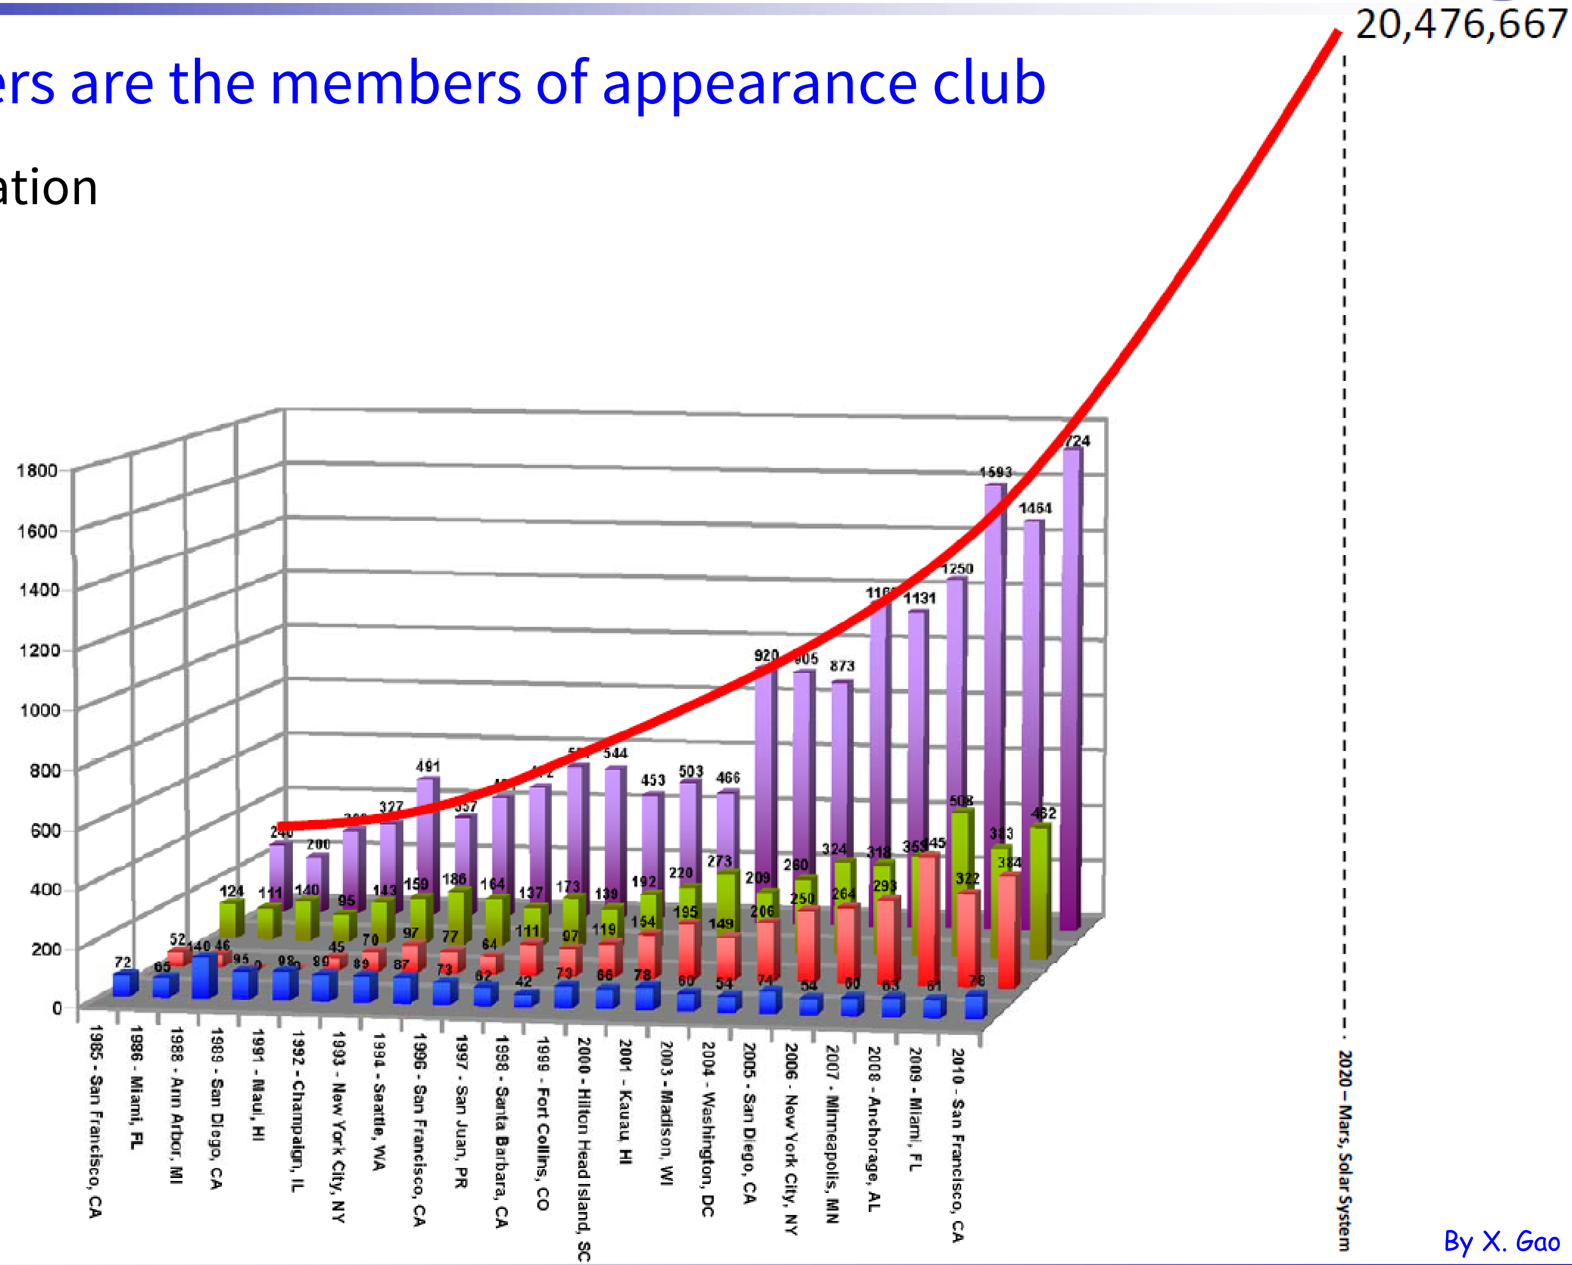

Paper Gestalt

- Most CVers are the members of appearance club
  - Motivation
  - Method

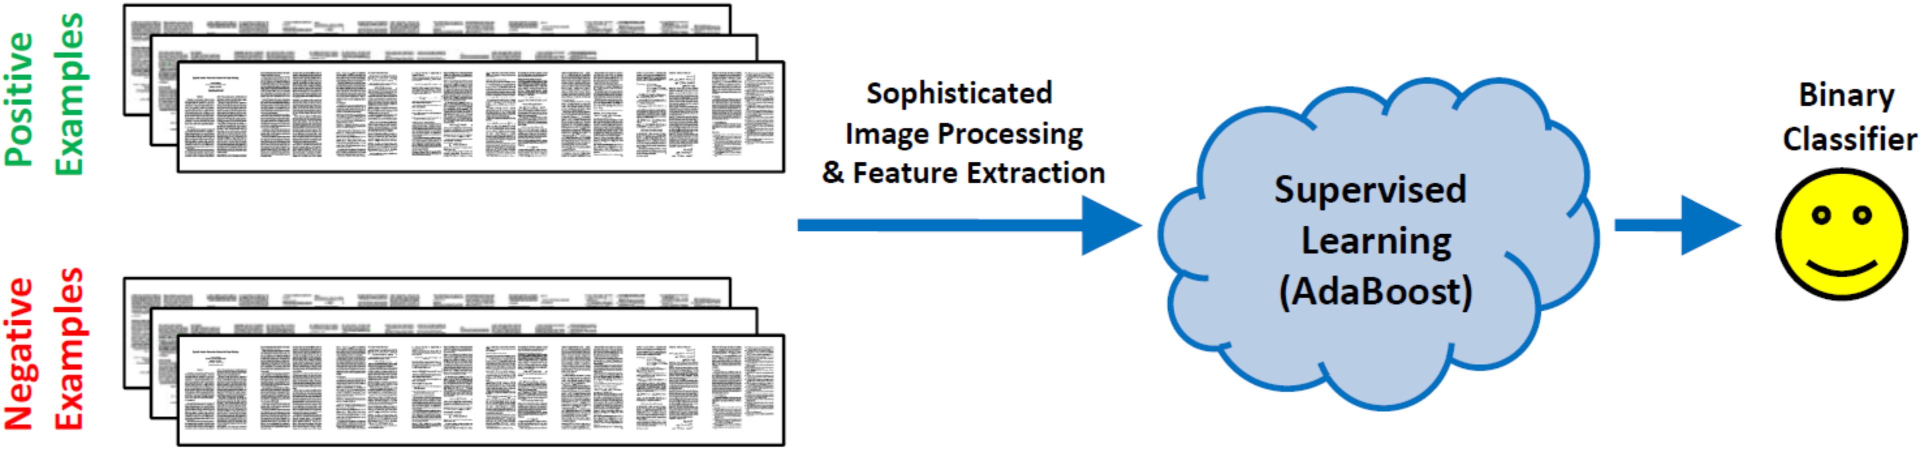

Paper Gestalt

- Most CVers are the members of appearance club
  - Main point
    - Get your paper looking pretty with right mix of equations, tables and figures

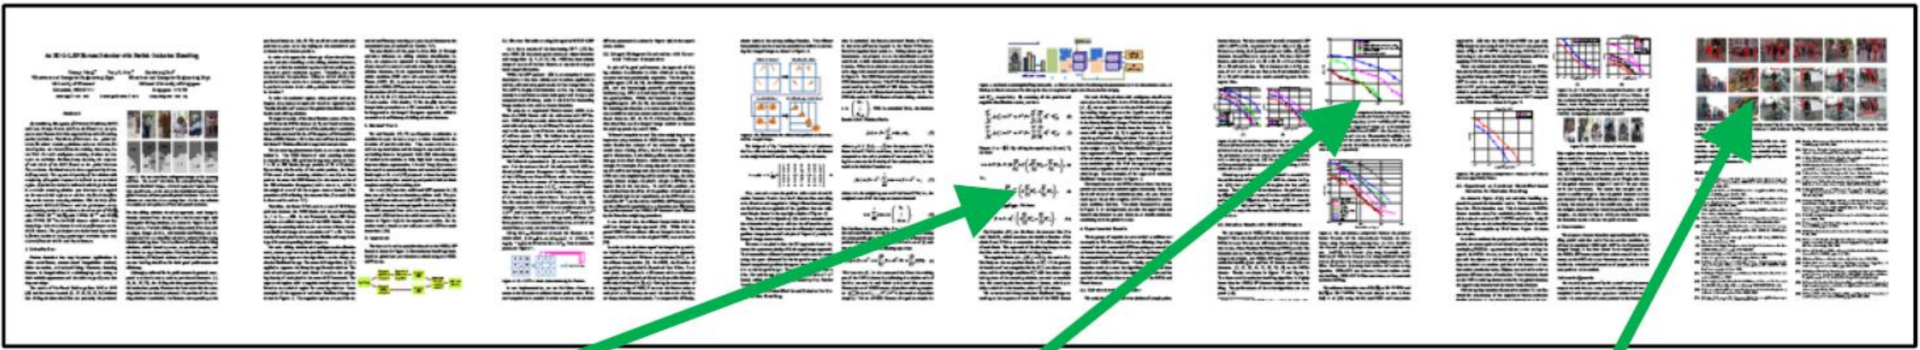

Math: Sophisticated mathematical expressions make a paper look technical and make the authors appear knowledgeable and "smart".

Plots: ROC, PR, and other performance plots convey a sense of thoroughness. Standard deviation bars are particularly pleasing to a scientific eye.

Figures/Screenshots: Illustrative figures that express complex algorithms in terms of 3rd grade visuals are always a must. Screenshots of anecdotal results are also very effective.

Paper Gestalt

- Most CVers are the members of appearance club
  - Main point
    - Get your paper looking pretty with right mix of equations, tables and figures

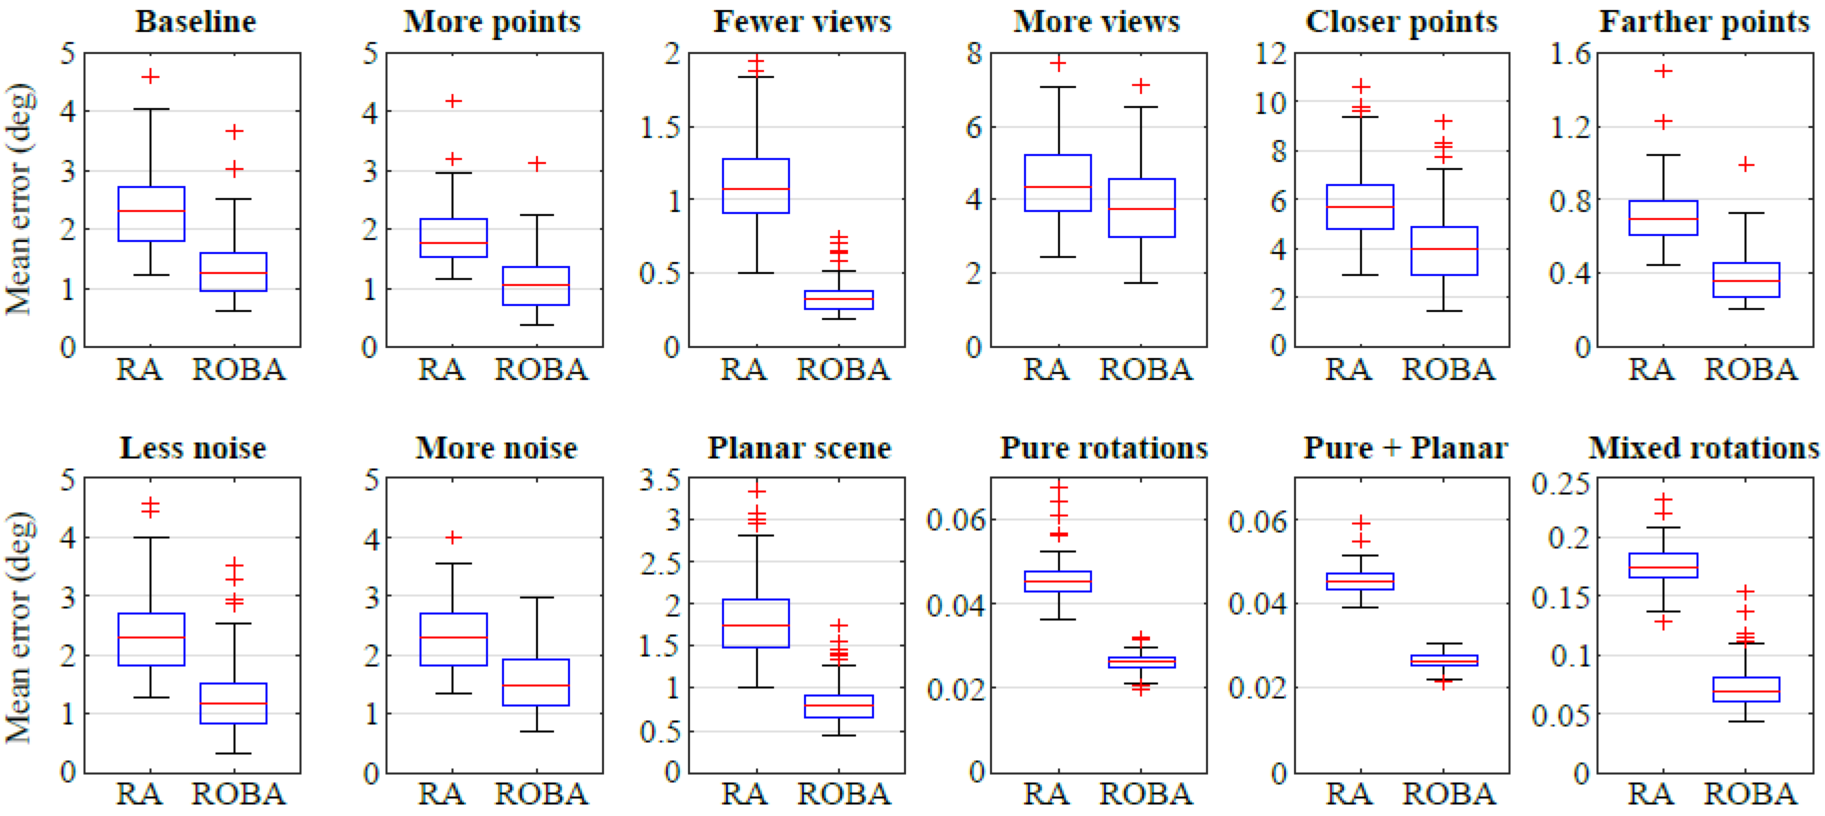

Paper Gestalt

- Most CVers are the members of appearance club
  - Main point
    - Get your paper looking pretty with right mix of equations, tables and figures

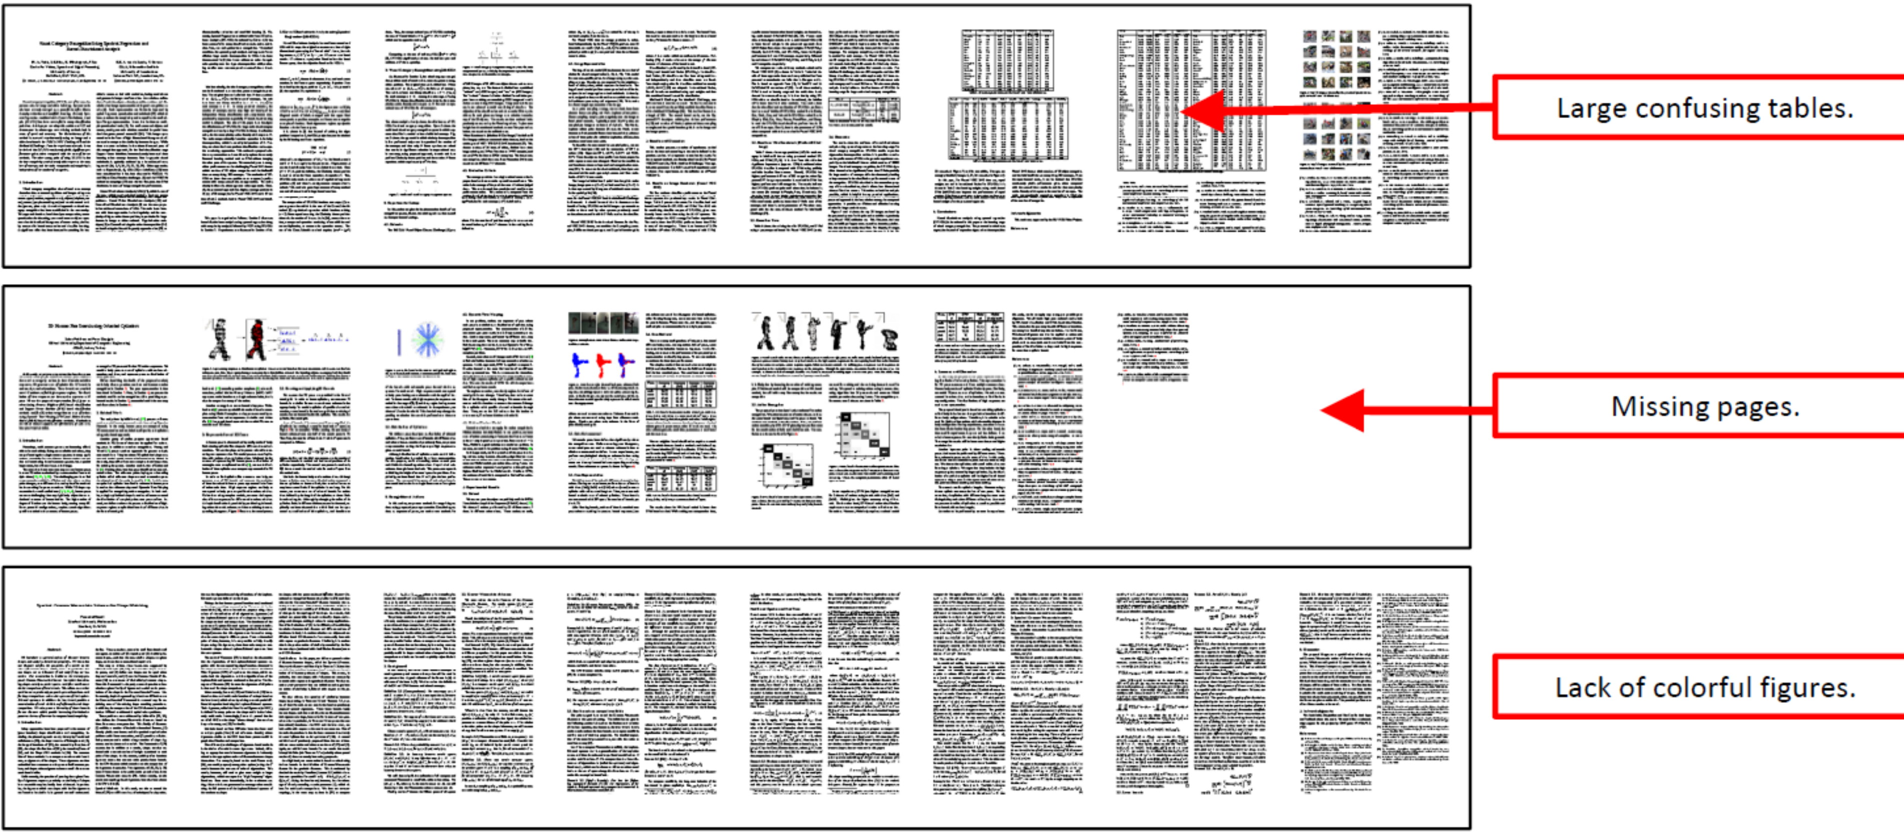

Additional Rules and Lessons

- Rules
  - Use **LaTeX**

Additional Rules and Lessons

- Rules
  - Use **LaTeX**
  - Read **author guidelines**

Submission Guidelines

All submissions will be handled electronically via the conference's CMT Website. By submitting a paper, the authors agree to the policies stipulated in this website. The paper submission deadline is November 15, 2019. Supplementary material can be submitted until November 22, 2019.

Papers are limited to eight pages, including figures and tables, in the CVPR style. Additional pages containing only cited references are allowed. Please refer to the following files for detailed formatting instructions:

- Example submission paper with detailed instructions
- LaTeX/Word Templates (tar)
- LaTeX/Word Templates (zip)

Papers that are not properly anonymized, or do not use the template, or have more than eight pages (excluding references) will be rejected without review.

1) Paper submission and review site:

Submission Site (bookmark or save this URL!)

Please make sure that your browser has cookies and Javascript enabled.

Please add "email@msr-cmt.org" to your list of safe senders (whitelist) to prevent important email announcements from being blocked by spam filters.

Log into CMT3 at https://cmt3.research.microsoft.com. If you do not see "2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)" in the conference list already, click on the "All Conferences" tab and find it there.

2) Setting up your profile: You can update your User Profile, Email, and Password by clicking on your name in the upper-right inside the Author Console and choosing the appropriate option under "General".

3) Domain Conflicts: When you log in for the first time, you will be asked to enter your conflict domain information. You will not be able to submit any paper without entering this information. We need to ensure conflict-free reviewing of all papers. At any time, you can update this information by clicking on your name in the upper-right and entering "Domain Conflicts" under CVPR 2020.

It is the primary author's responsibility to ensure that all authors on their paper have registered their institutional conflicts into CMT3. Each author should list domains of all institutions they have worked for, or have had very close collaboration with, within the last 3 years (example: mit.edu; ox.ac.uk; microsoft.com). DO NOT enter the domain of email providers such as gmail.com. This institutional conflict information will be used in conjunction with prior authorship conflict information to resolve assignments to both reviewers and area chairs. If a paper is found to have an undeclared or incorrect institutional conflict, the paper may be summarily rejected.

4) Creating a paper submission: This step must be completed by the paper registration deadline. After this deadline, you will not be able to register new papers, but you will be able to edit the information for existing papers.

Additional Rules and Lessons

- Rules
  - Use **LaTeX**
  - Read **author guidelines**
  - Read **reviewer guidelines**

The CVPR 2020 Reviewer Guidelines

Thank you for volunteering your time to review for CVPR 2020! To maintain a high-quality technical program, we rely very much on the time and expertise of our reviewers. This document explains what is expected of all members of the Reviewing Committee for CVPR 2020.

Benefits for Reviewers: 100 of our top reviewers will receive a CVPR Top Reviewer certificate and a gift certificate of $100 USD. In addition, all reviewers who did a good job (on time in submitting reviews, no reviews with very few words) will be guaranteed a registration ticket for a period of time after registration opens.

In addition to the guidelines below, you should read this CVPR 2020 Reviewer Tutorial for a summary of the decision process, annotated examples of good/bad reviews, and tips. You may also be interested in the CVPR 2020 Area Chair Tutorial to give you an overview of the process from the Area Chairs' point of view.

The CVPR 2020 Reviewing Timeline

Paper Submission Deadline	November 15, 2019
Papers Assigned to Reviewers	December 7, 2019
Reviews Due	January 17, 2020
Start of Post-Rebuttal Discussion Period	February 9, 2020
Final Recommendations Due	February 17, 2020
Decisions Released to Authors	February 23, 2020

Blind Reviews

Our Author Guidelines have instructed authors to make reasonable efforts to hide their identities, including omitting their names, affiliations, and acknowledgments. This information will of course be included in the published version. Likewise, reviewers should make all efforts to keep their identity invisible to the authors.

With the increase in popularity of arXiv preprints, sometimes the authors of a paper may be known to the reviewer. Posting to arXiv is NOT considered a violation of anonymity on the part of the authors, and in most cases, reviewers who happen to know (or suspect) the authors' identity can still review the paper as long as they feel that they can do an impartial job. An important general principle is to make every effort to treat papers fairly whether or not you know (or suspect) who wrote them. If you do not know the identity of the authors at the start of the process, DO NOT attempt to discover them by searching the Web for preprints.

Please read the FAQ at the end of this document for further guidelines on how arXiv prior work should be handled.

Check your papers

As soon as you get your reviewing assignment, please go through all the papers to make sure that (a) there is no obvious conflict with you (e.g., a paper authored by your recent collaborator from a different institution) and (b) you feel comfortable to review the paper assigned. If either of these issues arise, please let us know right away by emailing the Program Chairs (program-chairs-cvpr2020@googlegroups.com).

Additional Rules and Lessons

- Rules
  - Use **LaTeX**
  - Read **author guidelines**
  - Read **reviewer guidelines**
- Lessons
  - Several influential papers have been rejected once or twice: SIFT
  - Some best papers make little impact
  - Never give up in the process

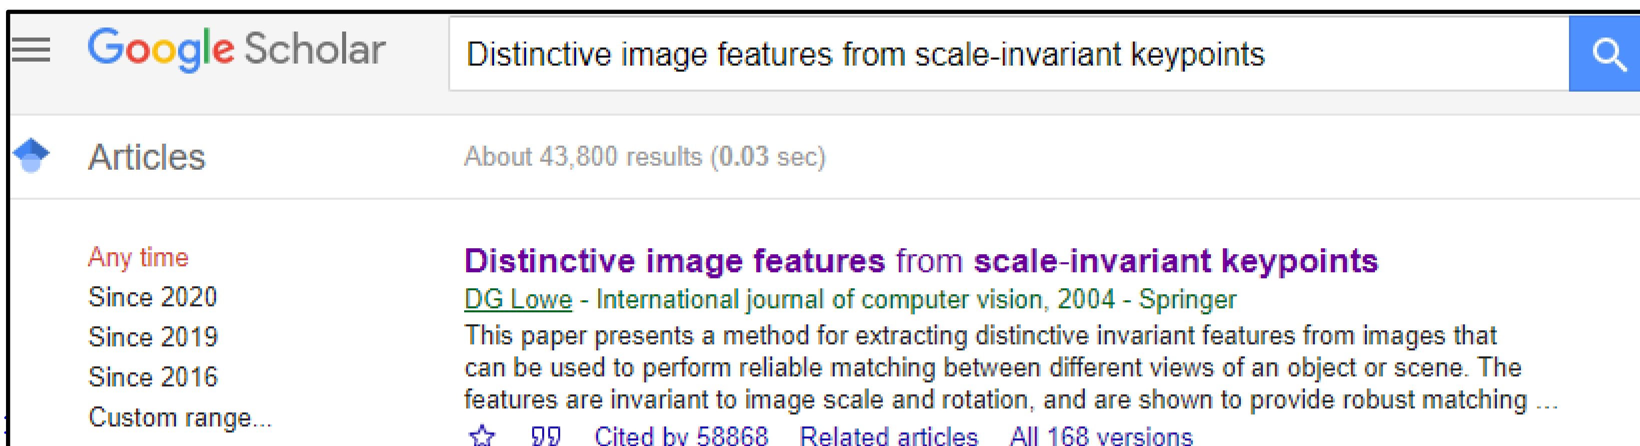

Your Advisor and You

Two More Things

- One Possibly Effective Way to Improving English Writing
  - Find several papers in the top-tier publications related to your research field, *e.g.*, IEEE T-PAMI, IJCV, Automatica, IEEE T-AC
  - Read the abstracts carefully and translate them into Chinese

Two More Things

- One Possibly Effective Way to Improving English Writing
  - Find several papers in the top-tier publications related to your research field, *e.g.*, IEEE T-PAMI, IJCV, Automatica, IEEE T-AC
  - Read the abstracts carefully and translate them into Chinese
  - Forget about this and do other jobs on your schedule
  - After several days, find out your Chinese translation version and translation them back to English

Two More Things

- One Possibly Effective Way to Improving English Writing
  - Find several papers in the top-tier publications related to your research field, *e.g.*, IEEE T-PAMI, IJCV, Automatica, IEEE T-AC
  - Read the abstracts carefully and translate them into Chinese
  - Forget about this and do other jobs on your schedule
  - After several days, find out your Chinese translation version and translation them back to English
  - Perform a careful cross check between the original abstracts and your circle translation ones and find out what are the differences and **why**

Two More Things

- 'Cold Humor' of CVers

USAC: A Universal Framework for Random Sample Consensus

Rahul Raguram, Ondřej Chum, *Member, IEEE*, Marc Pollefeys, *Member, IEEE*, Jiří Matas, *Member, IEEE*, and Jan-Michael Frahm, *Member, IEEE*

Abstract—A computational problem that arises frequently in computer vision is that of estimating the parameters of a model from data that have been contaminated by noise and outliers. More generally, any practical system that seeks to estimate quantities from noisy data measurements must have at its core some means of dealing with data contamination. The random sample consensus (RANSAC) algorithm is one of the most popular tools for robust estimation. Recent years have seen an explosion of activity in this area, leading to the development of a number of techniques that improve upon the efficiency and robustness of the basic RANSAC algorithm. In this paper, we present a comprehensive overview of recent research in RANSAC-based robust estimation by analyzing and comparing various approaches that have been explored over the years. We provide a common context for this analysis by introducing a new framework for robust estimation, which we call Universal RANSAC (USAC). USAC extends the simple hypothesize-and-verify structure of standard RANSAC to incorporate a number of important practical and computational considerations. In addition, we provide a general-purpose C++ software library that implements the USAC framework by leveraging state-of-the-art algorithms for the various modules. This implementation thus addresses many of the limitations of standard RANSAC within a single unified package. We benchmark the performance of the algorithm on a large collection of estimation problems. The implementation we provide can be used by researchers either as a stand-alone tool for robust estimation or as a benchmark for evaluating new techniques.

Index Terms—RANSAC, robust estimation

1 Introduction

A computational task that arises in a number of application scenarios is the estimation of model parameters from data that may be contaminated with measurement noise and, more significantly, may contain points that do not conform to the model being estimated. These points, called outliers, have a dramatic effect on the estimation process—a nonrobust technique, such as least squares regression, can produce arbitrarily bad model estimates in the presence of a single outlier. Consequently, the field of *robust estimation* has been well studied over the years, both in the statistics community [1], [2], [3], [4], as well as in computer vision [5], [6], [7]. A wide variety of algorithms have been proposed over the past four decades, varying in the degree of robustness that they provide to outliers, the assumptions they make about the data, and their computational complexity, among other aspects. Of these many algorithms, perhaps the one that is used most widely, particularly in computer vision, is random sample consensus, or RANSAC [7].

The RANSAC algorithm is a remarkably simple, yet powerful, technique. One compelling reason for its widespread adoption, in addition to its simplicity, is the ability of the algorithm to tolerate a tremendous level of contamination, providing reliable parameter estimates even when well over half the data consists of outliers. However, while robust, the basic RANSAC algorithm has its drawbacks, impacting its accuracy, efficiency, and stability. Recent years have seen exciting advances in dealing with each of these problems. Indeed, these improvements in computational efficiency and robustness have helped drive forward the state of the art, particularly as the computer vision and robotics communities push toward more challenging problems on massive real-world datasets [8], [9], [10], [11], [12] and seek real-time performance [13], [14], [15], [16]. However, while a number of recent efforts have focused on addressing issues with RANSAC, relatively less attention has been paid to a unified review of these developments. Some recent efforts in this direction are those of [17], [18], which analyze and compare the performance of some recent RANSAC variants on a selection of geometric estimation problems. We seek to extend this idea further. Our goals in this work are twofold:

- To present a comprehensive overview of recent research in RANSAC-based robust estimation, and to provide a common context within which to study these disparate techniques. To do so, we propose a generalization of the standard *hypothesize-and-verify* structure of standard RANSAC, extending it to incorporate a number of important practical and computational considerations. We term this *Universal RANSAC*, to emphasize the fact that most of the important RANSAC variants

- *R. Raguram is with Apple, Inc., 1 Infinite Loop, Cupertino, CA 95014. E-mail: rraguram@apple.com.*
- *O. Chum and J. Matas are with the Czech Technical University, Karlovo náměstí 13, Czech Republic. E-mail: {chum, matas}@cmp.felk.cvut.cz.*
- *M. Pollefeys is with the Department of Computer Science, ETH Zurich, CNB G105, Universitatstrasse 6, CH-8092 Zurich, Switzerland. E-mail: marc.pollefeys@inf.ethz.ch.*
- *J.-M. Frahm is with the Department of Computer Science, University of North Carolina at Chapel Hill, Office: FB 246, Chapel Hill, NC 27599. E-mail: jmf@cs.unc.edu.*

Manuscript received 15 Aug. 2011; revised 30 Jan. 2012; accepted 14 Oct. 2012; published online 29 Nov. 2012.
Recommended for acceptance by T. Tuytelaars.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2011-08-0550.
Digital Object Identifier no. 10.1109/TPAMI.2012.257.

VSAC: Efficient and Accurate Estimator for H and F

Maksym Ivashechkin[1], Daniel Barath[2], and Jiri Matas[1]
[1] Centre for Machine Perception, Czech Technical University in Prague, Czech Republic
[2] Computer Vision and Geometry Group, Department of Computer Science, ETH Zürich
{ivashmak, matas}@cmp.felk.cvut.cz dbarath@ethz.ch

Abstract

We present VSAC, a RANSAC-type robust estimator with a number of novelties. It benefits from the introduction of the concept of independent inliers that improves significantly the efficacy of the dominant plane handling and, also, allows near error-free rejection of incorrect models, without false positives. The local optimization process and its application is improved so that it is run on average only once. Further technical improvements include adaptive sequential hypothesis verification and efficient model estimation via Gaussian elimination. Experiments on four standard datasets show that VSAC is significantly faster than all its predecessors and runs on average in 1-2 ms, on a CPU. It is two orders of magnitude faster and yet as precise as MAGSAC++, the currently most accurate estimator of two-view geometry. In the repeated runs on EVD, HPatches, PhotoTourism, and Kusvod2 datasets, it never failed.

1. Introduction

The Random Sample Consensus (RANSAC) algorithm introduced by Fischler and Bolles [14] is one of the most popular robust estimators in computer science. The method is widely used in computer vision, its applications include stereo matching [33, 35], image mosaicing [15], motion segmentation [33], 3D reconstruction, detection of geometric primitives, and structure and motion estimation [28].

The textbook version of RANSAC proceeds as follows: random samples of minimal size sufficient to estimate the model parameters are drawn repeatedly. Model consistency with input data is evaluated, *e.g.*, by counting the points closer than a manually set inlier-outlier threshold. If the current model is better then the *so-far-the-best*, it gets stored. The procedure terminates when the probability of finding a better model falls below a user-defined level. Finally, the estimate is polished by least-squares fitting of inliers.

Many modifications of the original algorithm have been proposed. Regarding sampling, PROSAC [8] exploits an a priori predicted inlier probability rank. NAPSAC [27] samples in the neighborhood of the first, randomly selected, point. Progressive NAPSAC [2] combines both and adds gradual convergence to uniform spatial sampling.

In textbook RANSAC, the model quality is measured by its support, *i.e.*, the number of inliers, points consistent with the model. MLESAC [34] introduced a quality measure that makes it the maximum likelihood procedure. To avoid the need for a user-defined noise level, MINPRAN [32] and A-contrario RANSAC [13] select the inlier-outlier threshold so that the inliers are the least likely to occur at random. Reflecting the inherent uncertainty of the threshold estimate, MAGSAC [5] marginalizes the quality function over a range of noise levels. MAGSAC++ [4] proposes an iterative re-weighted least-squares optimization of the *so-far-the-best* model with weights calculated from the inlier probability of points. The Locally Optimized RANSAC [9] refines the *so-far-the-best* model using a non-minimal number of points, *e.g.*, by iterated least-squares fitting. Graph-Cut RANSAC [3], in its local optimization, exploits the fact that real-world data tend to form spatial structures. The model evaluation is usually the most time-consuming part as it depends both on the number of models generated and the number of input data points. A quasi-optimal speed-up was achieved by the Sequential Probability Ratio Test (SPRT) [25] that randomizes the verification process itself.

In many cases, points in degenerate configuration affect the estimation severely. For example, correspondences lying on a single plane is a degenerate case for **F** estimation. DEGENSAC [11] detects such cases and applies the plane-and-parallax algorithm. USAC [29] was the first framework integrating many of the mentioned techniques, including PROSAC, SPRT, DEGENSAC, and LO-RANSAC.

In this paper, we present VSAC[1], a RANSAC-type estimator that exploits a number of novelties. It is significantly faster than all its predecessors, and yet as precise as MAGSAC++, the currently most accurate method both in our experiments and according to a recent survey [23]. The accuracy reaches, or is very near, the geometric error of the

[1] VSAC has multiple novelties and we found no natural abbreviation reflecting them. We chose "V" as the letter following "U", as in USAC.

Two More Things

- ‘Cold Humor’ of CVers

USAC: A Universal Framework for Random Sample Consensus

Rahul Raguram, Ondřej Chum, *Member, IEEE*, Marc Pollefeys, *Member, IEEE*, Jiří Matas, *Member, IEEE*, and Jan-Michael Frahm, *Member, IEEE*

Abstract—A computational problem that arises frequently in computer vision is that of estimating the parameters of a model from data that have been contaminated by noise and outliers. More generally, any practical system that seeks to estimate quantities from noisy data measurements must have at its core some means of dealing with data contamination. The random sample consensus (RANSAC) algorithm is one of the most popular tools for robust estimation. Recent years have seen an explosion of activity in this area, leading to the development of a number of techniques that improve upon the efficiency and robustness of the basic RANSAC algorithm. In this paper, we present a comprehensive overview of recent research in RANSAC-based robust estimation by analyzing and comparing various approaches that have been explored over the years. We provide a common context for this analysis by introducing a new framework for robust estimation, which we call Universal RANSAC (USAC). USAC extends the simple hypothesize-and-verify structure of standard RANSAC to incorp... provide a general-purpose C++ software l... various modules. This implementation th... We benchmark the performance of the al... used by researchers either as a stand-al...

Index Terms—RANSAC, robust estimati...

1 Introduction

A computational task that arises in a... tion scenarios is the estimation of... from data that may be contaminated... noise and, more significantly, may contai... conform to the model being estimated. ... outliers, have a dramatic effect on the es... nonrobust technique, such as least squa... produce arbitrarily bad model estimates in the presence of a single outlier. Consequently, the field of *robust estimation* has been well studied over the years, both in the statistics community [1], [2], [3], [4], as well as in computer vision [5], [6], [7]. A wide variety of algorithms have been proposed over the past four decades, varying in the degree of robustness that they provide to outliers, the assumptions they make about the data, and their computational complexity, among other aspects. Of these many algorithms, perhaps the one that is used most widely, particularly in computer vision, is random sample consensus, or RANSAC [7].

impacting its accuracy, efficiency, and stability. Recent years have seen exciting advances in dealing with each of these problems. Indeed, these improvements in computational efficiency and robustness have helped drive forward the state of the art, particularly as the computer vision and robotics communities push toward more challenging problems on massive real-world datasets [8], [9], [10], [11], [12] and seek real-time performance [13], [14], [15], [16]. However, while a number of recent efforts have focused on addressing issues with RANSAC, relatively less attention has been paid to a unified review of these developments. Some recent efforts in this direction are those of [17], [18], which analyze and compare the performance of some recent RANSAC variants on a selection of geometric estimation problems. We seek to extend this idea further. Our goals in this work are twofold:

- To present a comprehensive overview of recent research in RANSAC-based robust estimation, and to provide a common context within which to study these disparate techniques. To do so, we propose a generalization of the standard *hypothesize-and-verify* structure of standard RANSAC, extending it to incorporate a number of important practical and computational considerations. We term this *Universal RANSAC*, to emphasize the fact that most of the important RANSAC variants

- R. Raguram is with Apple, Inc., 1 Infinite Loop, Cupertino, CA 95014. E-mail: rraguram@apple.com.
- O. Chum and J. Matas are with the Czech Technical University, Karlovo náměstí 13, Czech Republic. E-mail: {chum, matas}@cmp.felk.cvut.cz.
- M. Pollefeys is with the Department of Computer Science, ETH Zurich, CNB G105, Universitatstrasse 6, CH-8092 Zurich, Switzerland. E-mail: marc.pollefeys@inf.ethz.ch.
- J.-M. Frahm is with the Department of Computer Science, University of North Carolina at Chapel Hill, Office: FB 246, Chapel Hill, NC 27599. E-mail: jmf@cs.unc.edu.

Manuscript received 15 Aug. 2011; revised 30 Jan. 2012; accepted 14 Oct. 2012; published online 29 Nov. 2012.
Recommended for acceptance by T. Tuytelaars.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2011-08-0550.
Digital Object Identifier no. 10.1109/TPAMI.2012.257.

0162-8828/13/$31.00 © 2013 IEEE Published by the IEEE Computer Society

VSAC: Efficient and Accurate Estimator for H and F

Maksym Ivashechkin[1], Daniel Barath[2], and Jiri Matas[1]
[1] Centre for Machine Perception, Czech Technical University in Prague, Czech Republic
[2] Computer Vision and Geometry Group, Department of Computer Science, ETH Zürich
{ivashmak, matas}@cmp.felk.cvut.cz dbarath@ethz.ch

Abstract

ples in the neighborhood of the first, randomly selected, point. Progressive NAPSAC [2] combines both and adds ...ce to uniform spatial sampling. ...NSAC, the model quality is measured by ... number of inliers, points consistent with ...AC [34] introduced a quality measure that ...mum likelihood procedure. To avoid the ...efined noise level, MINPRAN [32] and ...SAC [13] select the inlier-outlier thresh... ...liers are the least likely to occur at ran... the inherent uncertainty of the threshold ...C [5] marginalizes the quality function ...ise levels. MAGSAC++ [4] proposes an ...ed least-squares optimization of the *so-*... ...l with weights calculated from the inlier ...ts. The Locally Optimized RANSAC [9] ...*he-best* model using a non-minimal num... ... by iterated least-squares fitting. Graph-... ... in its local optimization, exploits the fact that real-world data tend to form spatial structures. The model evaluation is usually the most time-consuming part as it depends both on the number of models generated and the number of input data points. A quasi-optimal speed-up was achieved by the Sequential Probability Ratio Test (SPRT) [25] that randomizes the verification process itself.

In many cases, points in degenerate configuration affect the estimation severely. For example, correspondences lying on a single plane is a degenerate case for F estimation. DEGENSAC [11] detects such cases and applies the plane-and-parallax algorithm. USAC [29] was the first framework integrating many of the mentioned techniques, including PROSAC, SPRT, DEGENSAC, and LO-RANSAC.

In this paper, we present VSAC[1], a RANSAC-type estimator that exploits a number of novelties. It is significantly faster than all its predecessors, and yet as precise as MAGSAC++, the currently most accurate method both in our experiments and according to a recent survey [23]. The accuracy reaches, or is very near, the geometric error of the

1. Introduction

The Random Sample Consensus (RANSAC) algorithm introduced by Fischler and Bolles [14] is one of the most popular robust estimators in computer science. The method is widely used in computer vision, its applications include stereo matching [33, 35], image mosaicing [15], motion segmentation [33], 3D reconstruction, detection of geometric primitives, and structure and motion estimation [28].

The textbook version of RANSAC proceeds as follows: random samples of minimal size sufficient to estimate the model parameters are drawn repeatedly. Model consistency with input data is evaluated, *e.g.*, by counting the points closer than a manually set inlier-outlier threshold. If the current model is better then the *so-far-the-best*, it gets stored. The procedure terminates when the probability of finding a better model falls below a user-defined level. Finally, the estimate is polished by least-squares fitting of inliers.

Many modifications of the original algorithm have been proposed. Regarding sampling, PROSAC [8] exploits an a priori predicted inlier probability rank. NAPSAC [27] sam-

[1] VSAC has multiple novelties and we found no natural abbreviation reflecting them. We chose "V" as the letter following "U", as in USAC.

VSAC has multiple novelties and we found no natural abbreviation reflecting them. We chose ‘V’ as the letter following ‘U’, as in USAC

Several Minor Notes on Figures, Tables, Mathematical Notations, and Numbering

Xiang Gao, Lecturer

E-mail: xgao@ouc.edu.cn | Web: https://ouc-xgao.github.io/

College of Engineering, Ocean University of China

Figures

- Definite Purpose
  - Each figure should have a definite purpose
    - This might be to help **clarify** the text, or **demonstrate** a particular experimental result
    - Figures included **just** to look more appealing are not appropriate in **scientific writing**
    - Figures should be used for information which is **hard** to explain in **words**, and the reader will find **easier** to grasp that by means of **figures**

Figures

- Definite Purpose
  - For figures with curves or broken lines:
    - Make sure that the axes are **labelled** to state what they represent;
    - Make sure that the range of values are shown with **units**

Figures

- Definite Purpose
  - For figures with curves or broken lines:
    - Make sure that the axes are **labelled** to state what they represent;
    - Make sure that the range of values are shown with **units**
  - **Bad** example

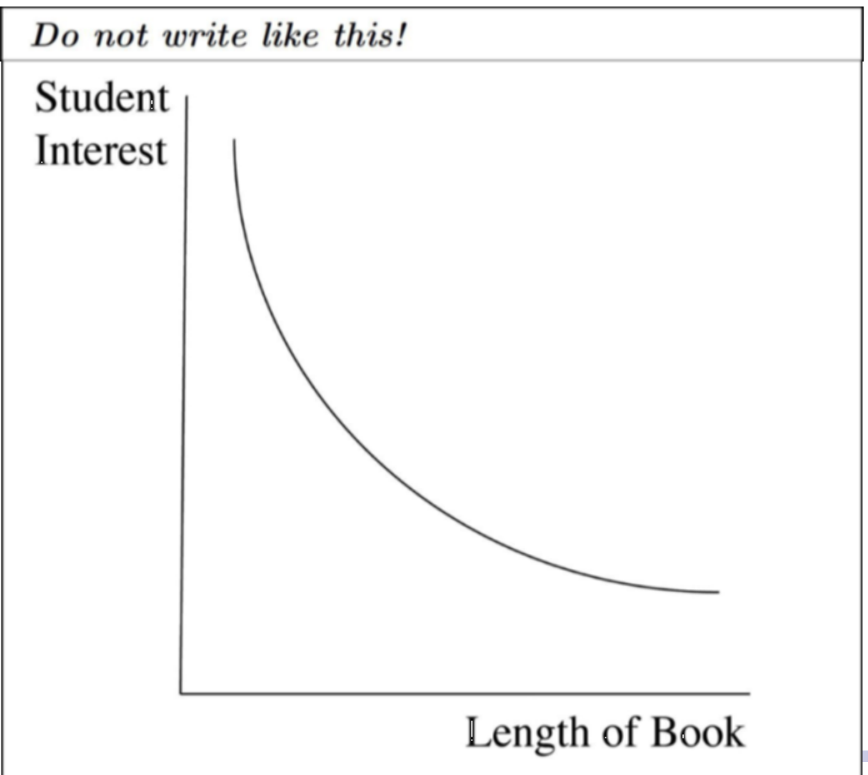

❌ Vague and suggestive

❌ Fail to give units

❌ No mathematical basis

Figures

- Definite Purpose
  - For figures with curves or broken lines:
    - Make sure that the axes are **labelled** to state what they represent;
    - Make sure that the range of values are shown with **units**
  - **Good** example

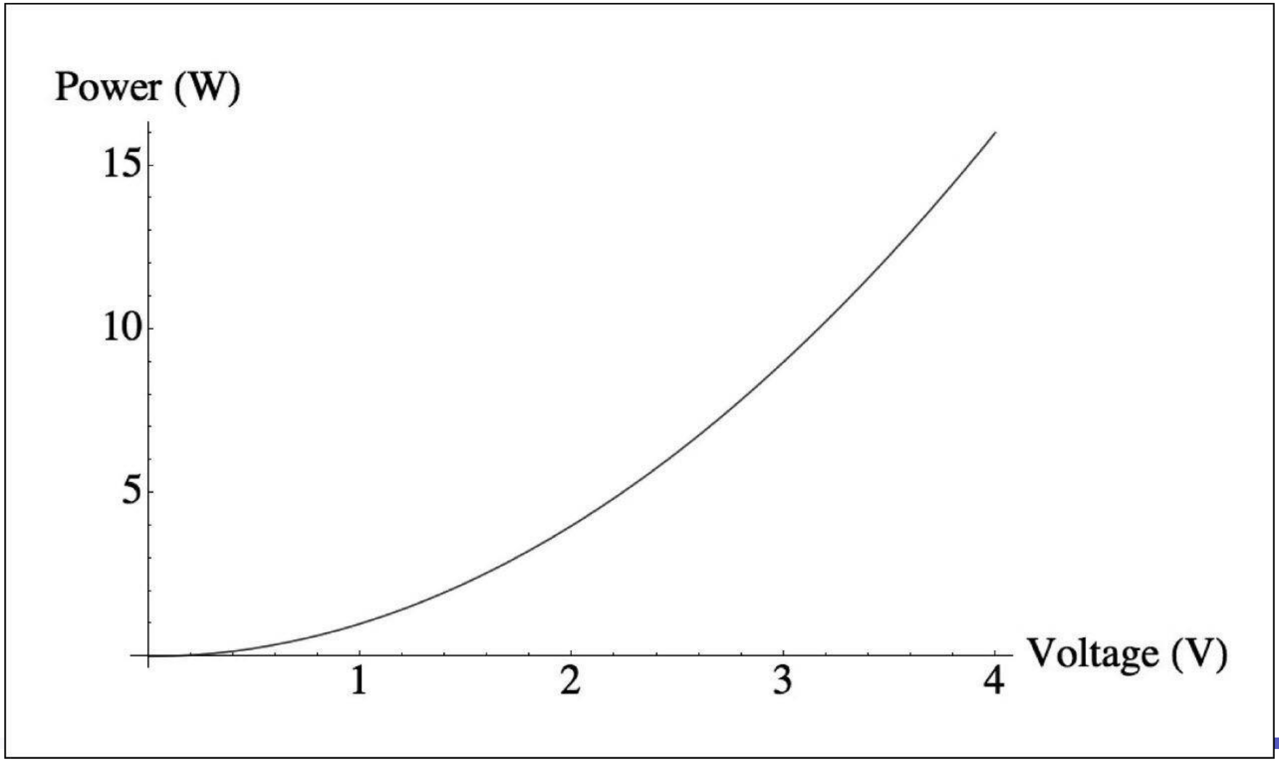

Figures

- Reference and Explanation
  - Every figure should be **referred** to in the main text explicitly. Do not include figures without saying what they show
  - **Explain** how it adds to the text, and what the reader is **supposed** to understand

Figures

- Reference and Explanation
  - Every figure should be **referred** to in the main text explicitly. Do not include figures without saying what they show
  - **Explain** how it adds to the text, and what the reader is **supposed** to understand
  - **For example**: **Fig. 1** shows **how** power delivered to the battery varies with voltage in our supercharger circuit. As the voltage increases, the power delivered also increases. **Thus**, for rapid charging, the supercharger should be operated at as high a voltage as possible

Figures

- Size
  - Do **not** make figures too **small**
  - Try to **avoid shrinking** figures to accommodate more text

Figures

- Size
  - Do **not** make figures too **small**
  - Try to **avoid shrinking** figures to accommodate more text
  - If you **must** use **small** figures, at least show a **sub-figure** which **zooms in** on the important part to show the difference in detail
  - Make sure that the smallest text in any figure is **no smaller** than the main **font size** used in the paper

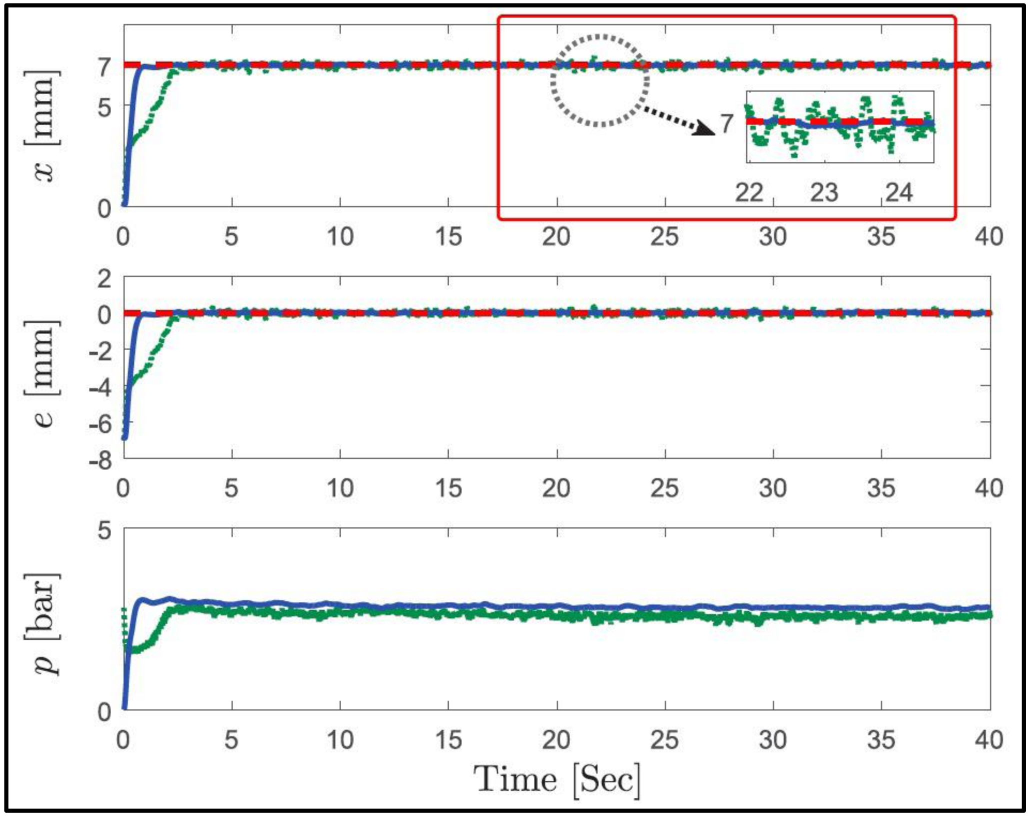

Figures

- Consistency
  - Make sure that text within figures, and the caption, is **consistent** with the main text:
    - Any terminology used should **match** that in the main text
    - Symbols should look the **same**, ideally in the same **font**

F

•

ne

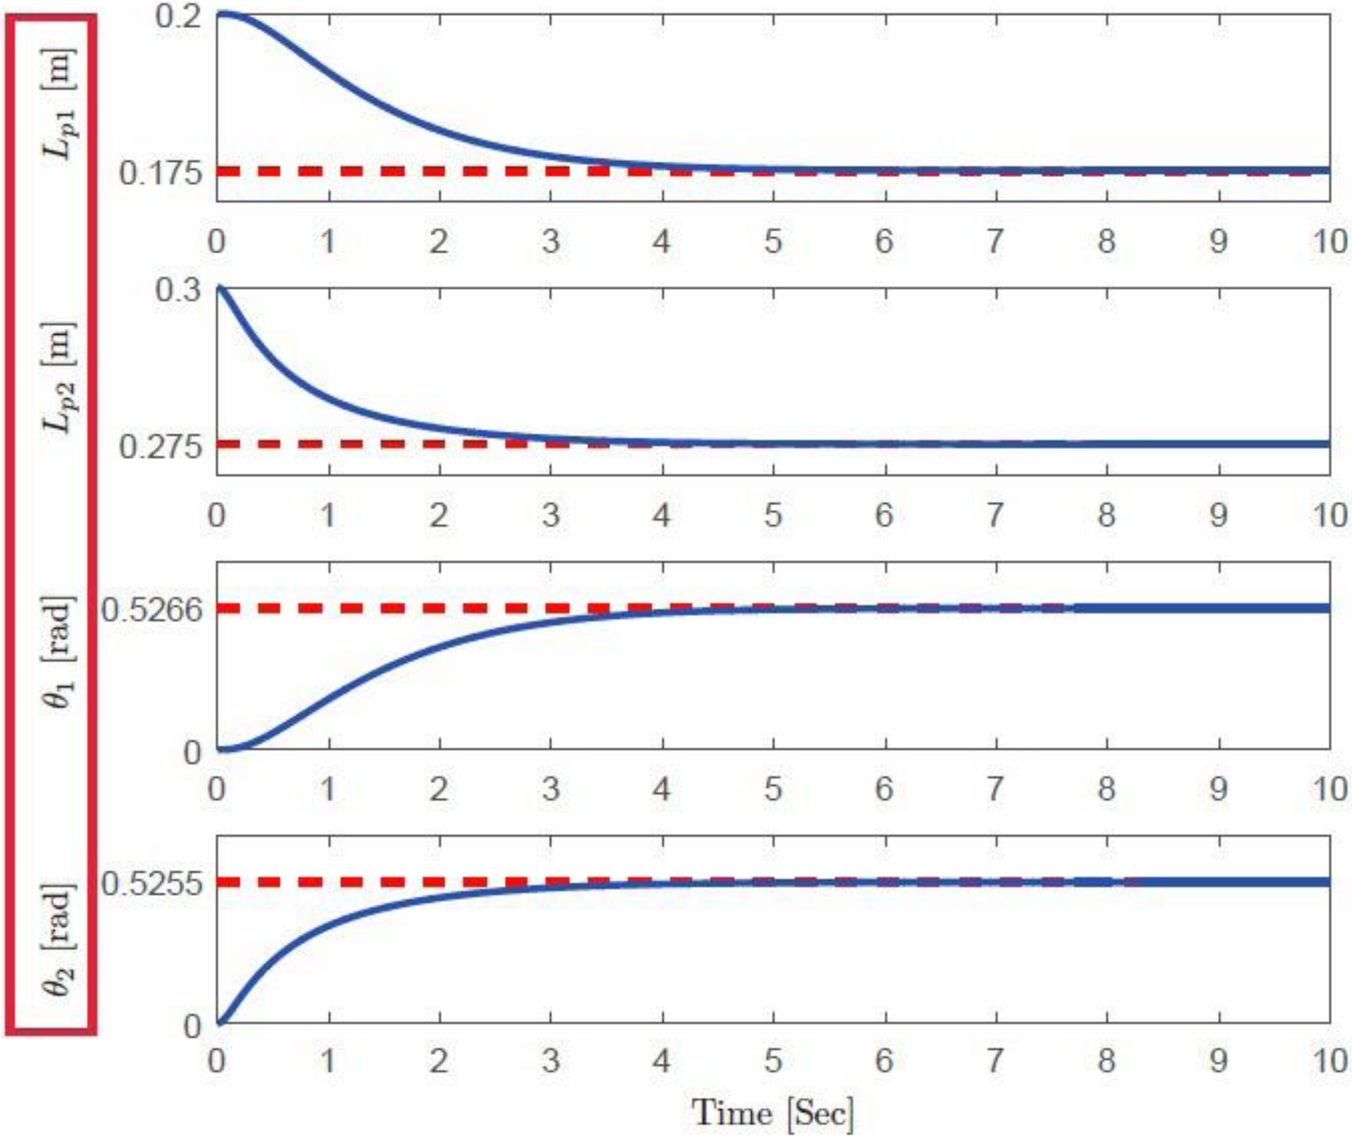

Fig. 2. Simulation results of $\theta_1(t)$, $\theta_2(t)$, $L_{p1}(t)$, $L_{p2}(t)$ of the dual-PAM system (reference values–red dashed line; simulation results–blue solid line).

Figures

- Placement
  - **Place** figures as **near** as possible to where they are **first** mentioned in the text, **ideally** on the **same** page, or at least the next page
    - It is **distracting** to readers to have to **skip** forwards and backwards between the text and figures
    - The linear flow of ideas should **not** be **disrupted**
    - Ensure that figures are **numbered** in the **same** order that they **appear** in the paper

Figures

- ## Placement
  - For ease of reading, figures (and tables) should **normally** be **placed** at the **top** of the page (or column), rather than in the **middle** of it, **except** for small figures which fit into the flow of the text

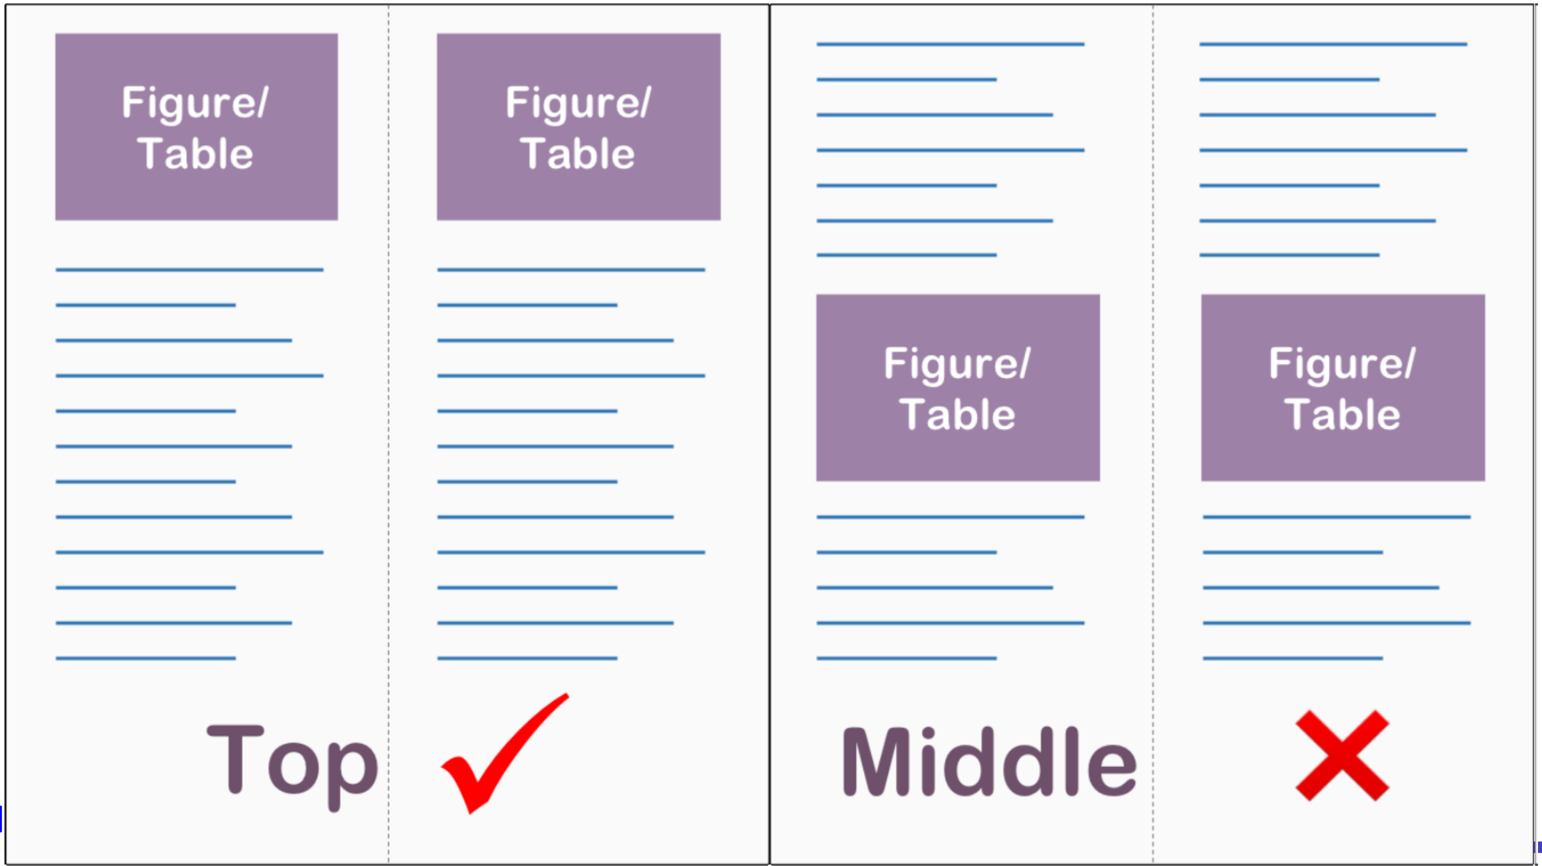

Tables

- Special Figures
  - Most of the above comments about figures **equally apply** to tables
  - They are really a **particular kind** of figure containing **textual** information.
  - Tables and figures conventionally **numbered** separately
  - Algorithm pseudocode listings could also be presented as another **special** kind of figure (or table), again with their **own** numbering sequence

Tables

- Presentation
  - The typical use of tables is to present **numerical results** or some other numerical information
  - Make sure **each** row and column has appropriate **headers** to explain what that row or column contains
  - If numbers are **physical** quantities, the table should **state** the units with each number

Tables

- Decimal Places
  - Do **not** give numbers to **more** significant digits than are necessary to make your point
  - **For example**: if comparing the success rate of alternative approaches, and these numbers vary from **40% to 95%**, you do not need to give any decimal places at all. On the other hand, if they vary between **98% to 99%** you may need **one or even two** decimal places

Mathematical Notations

- Follow the Standard
  - If your research field conventionally uses **standard** notations for various mathematical values, make sure you follow it
  - **For example**: It is **standard** practice to call the principal curvatures in differential geometry k_1 and k_2, and it would be **unhelpful and confusing** to refer to them as c_a and c_b
  - If previous papers have all used the **same** symbols for some quantities you **also** need, use the same symbols
  - Readers often **consider** and **compare** several papers on the **same topic**

Mathematical Notations

- ## Definition and Placement
  - Make sure that **all** mathematical notations used are **defined**, apart from commonly understood ones like π and e
  - The definition should come as **close** as possible to the place where the symbol is **first** used in your paper

Mathematical Notations

- Definition and Placement
  - An example

Considering an EG, denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, is formed by $|\mathcal{V}|$ cameras and $|\mathcal{E}|$ relative rotation measurements. A vertex $v_i \in \mathcal{V}$ corresponds to a camera with absolute rotation $\boldsymbol{R}_i$ and an edge $e_{ij} \in \mathcal{E}$ links an image pair with relative rotation $\boldsymbol{R}_{ij}$. Then, the rotation averaging problem is defined as:

$$\{\boldsymbol{R}_i^*\} = \arg\min \sum_{e_{ij} \in \mathcal{E}} \rho\left(d(\boldsymbol{R}_{ij}, \boldsymbol{R}_j \boldsymbol{R}_i^T)\right), \qquad (1)$$

where $\{\boldsymbol{R}_i^*\}$ is the estimated absolute rotations, $\rho(.)$ is the loss function for robust optimization, and $d(.,.)$ is the distance measure between the measured and re-computed relative rotations. For loss function $\rho()$, thanks to our effective outlier filtering strategy, the simple ℓ_2 loss is used in this letter. For distance measure $d(.,.)$, we choose the angular distance $d_\theta(.,.)$, which is used in most of the related works [10], [11], [17],

Mathematical Notations

- Single Meaning
  - Make sure that each notation is used with only a **single meaning** in a given paper
  - When you are defining your own symbols, use **easily remembered** names as much as possible
    - **For example**, Use P for a point and L for a line, rather than, say A for the point and B for the line
  - If you have **several** related items, give them related **names**
    - **For example**, If you have several related points, use **subscripts**, and call them P_1, P_2, and P_3, or failing that, call them P, Q, and R

Mathematical Notations

- Advice
  - Computer scientists are usually **advised** that variables and functions in programs should be given **long names**
  - However, in mathematics, the convention is to (usually) use **single letter** names for such **quantities**, and for **subscripts**
  - Do **not** express ideas **entirely** through **mathematical notation**.
  - Trying to put the ideas into **words** in the main text
  - Explain ideas informally in English **first**, further giving more **precise** details in mathematical notation

Numbering

- Sections, Figures, Tables, *etc.*
  - Sections and subsections should be **hierarchically** numbered throughout the paper. The first section of the paper, **Section 1**, should have subsections numbered **1.1**, **1.2**, and so on
  - **All** figures should be sequentially numbered using a **single sequence** throughout the paper, rather than a hierarchical approach
  - Tables should have **their own** separate sequential numbering sequence, as should any other type of special items such as **algorithms**, **theorems**, *etc*

Numbering

- Examples of Figures or Tables
  - When **cross-referring** to sections, figures, tables, and equations, refer to them **precisely** by number, rather than more **vaguely**
    - **For example**: ‘Figure 7 shows …’ rather than ‘The **above** figure shows …’
    - ‘The above tables’ could refer to **any** number of previous tables
    - However, refer to ‘The **next** section’ in cases where the meaning is **unambiguous**

Numbering

- Examples of Figures or Tables
  - When **cross-referring** to sections, figures, tables, and equations, refer to them **precisely** by number, rather than more **vaguely**
    - **For example**: ‘Figure 7 shows …’ rather than ‘The **above** figure shows …’
    - ‘The above tables’ could refer to **any** number of previous tables
    - However, refer to ‘The **next** section’ in cases where the meaning is **unambiguous**
  - The **exact** format used to refer to figures, tables and equations is **determined** by the **publisher’s** house style
    - **For example**, it might be as in ‘**Figure 2**’, ‘**Fig. 2**’, or ‘**Figure (2)**’

Numbering

- Examples of Equations
  - When you are referring back to equations, you should summarize what they mean, rather than simply referring to them by number
  - Do not write like this!
    - We substitute **Eqn. (2)** into **Eqn. (6)** to obtain the following equation
  - Instead, write
    - We substitute the **locality constraint** in **Eqn. (2)** into the **similarity function** in **Eqn. (6)** to give the neighborhood similarity, as follows

References

- Ming-Hsuan Yang
  - How to Get Your CVPR Paper Rejected?
- Bill Freeman
  - How to Write a Good CVPR Submission?
- Ming-Ming Cheng
  - Optimizing the Presentation of Paper
- CVPR 2020
  - Opening Slides
- Zhihu
  - https://www.zhihu.com/

If you are not disappointed enough,

you have not tried hard enough!

Xiang Gao, Lecturer

E-mail: xgao@ouc.edu.cn | Web: https://ouc-xgao.github.io/

College of Engineering, Ocean University of China