Kaggle Kuzushiji Recognition Challenge

5th place solution

Background

- Kaggle Top 100
- Gold medals in computer vision and speech recognition competitions

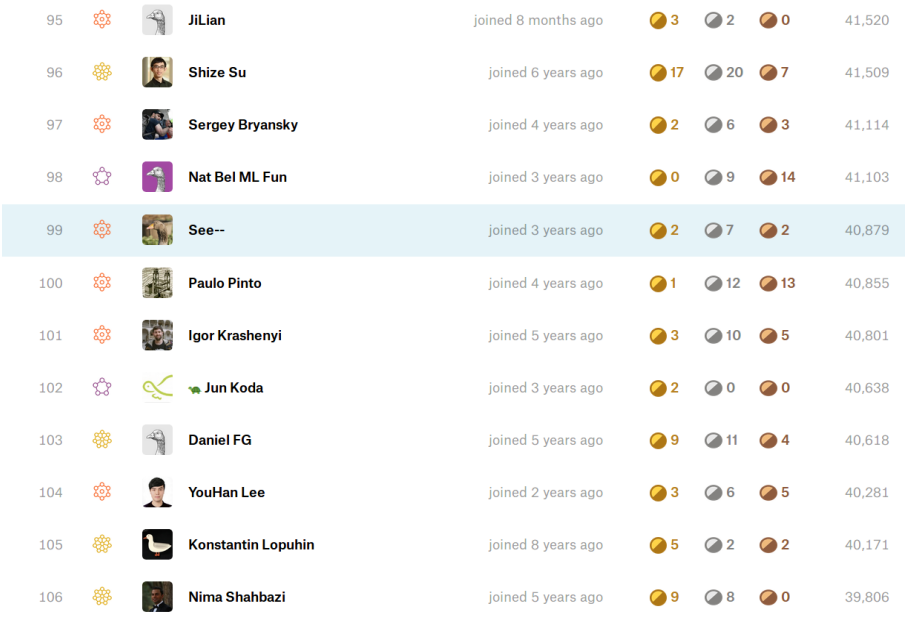

95	JiLian	joined 8 months ago	3	2	0	41,520
96	Shize Su	joined 6 years ago	17	20	7	41,509
97	Sergey Bryansky	joined 4 years ago	2	6	3	41,114
98	Nat Bel ML Fun	joined 3 years ago	0	9	14	41,103
99	See--	joined 3 years ago	2	7	2	40,879
100	Paulo Pinto	joined 4 years ago	1	12	13	40,855
101	Igor Krashenyi	joined 5 years ago	3	10	5	40,801
102	Jun Koda	joined 3 years ago	2	0	0	40,638
103	Daniel FG	joined 5 years ago	9	11	4	40,618
104	YouHan Lee	joined 2 years ago	3	6	5	40,281
105	Konstantin Lopuhin	joined 8 years ago	5	2	2	40,171
106	Nima Shahbazi	joined 5 years ago	9	8	0	39,806

Motivation

- Created an inference Keras port of CenterNet[1]
- Implement the training part
- https://github.com/see--/keras-centernet

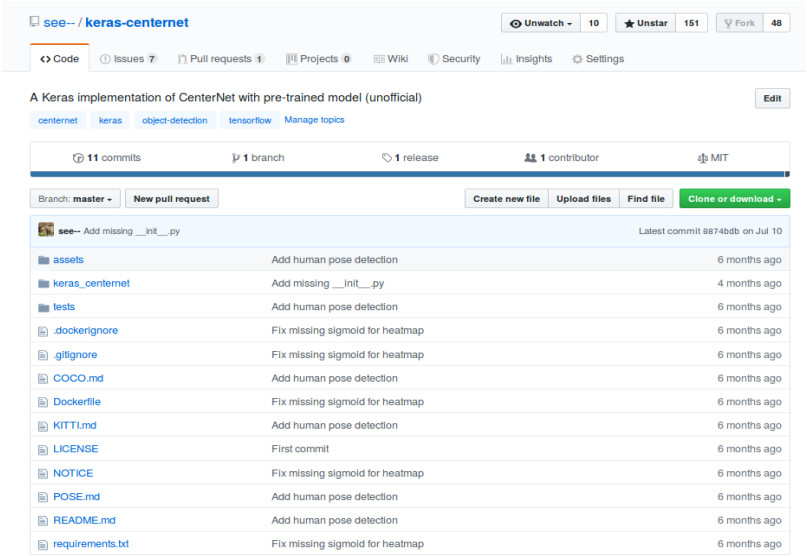

[1] https://arxiv.org/abs/1904.07850

CenterNet

- Strong and fast one-stage detector: 45.1% AP on COCO
- Meta-algorithm for various object detection tasks

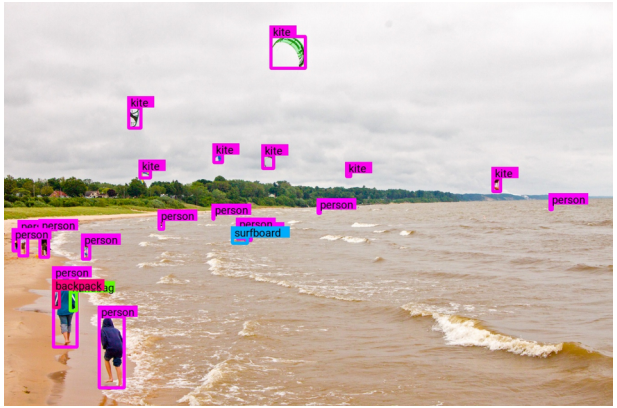

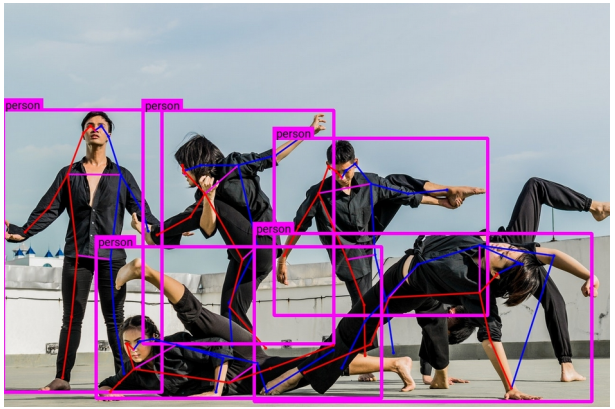

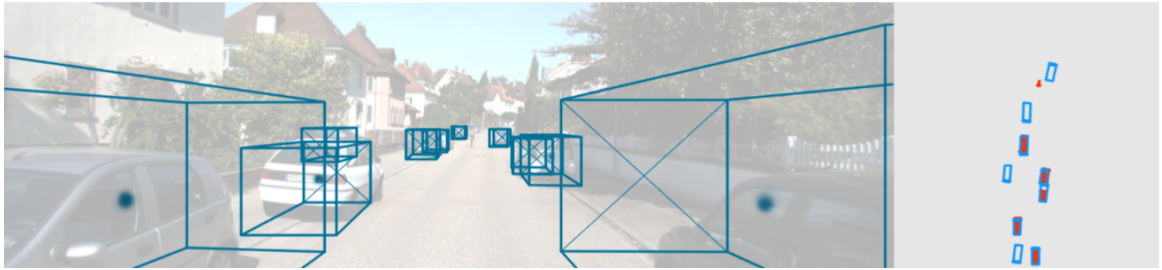

Summary

- Kuzushiji Recognition is 2d-detection with many classes: 4212
- 5th place with 0.94 F1-score
- 8 hours using one V100 GPU and mixed precision training[1]

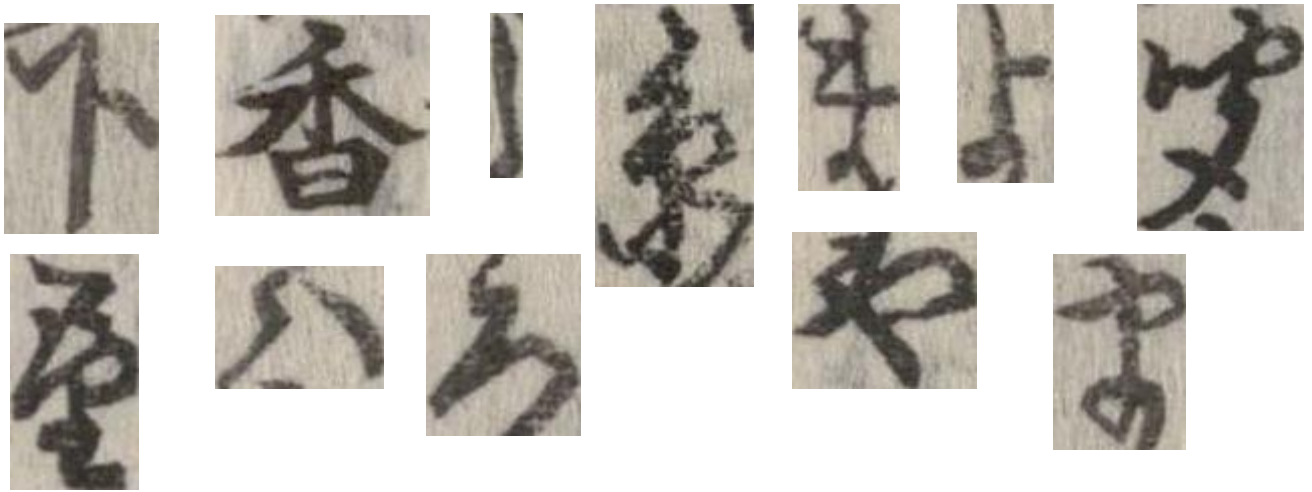

[1] https://github.com/NVIDIA/apex

- No pre- or post-processing
- Train on smaller crops and inference on full images
- Lots of augmentations

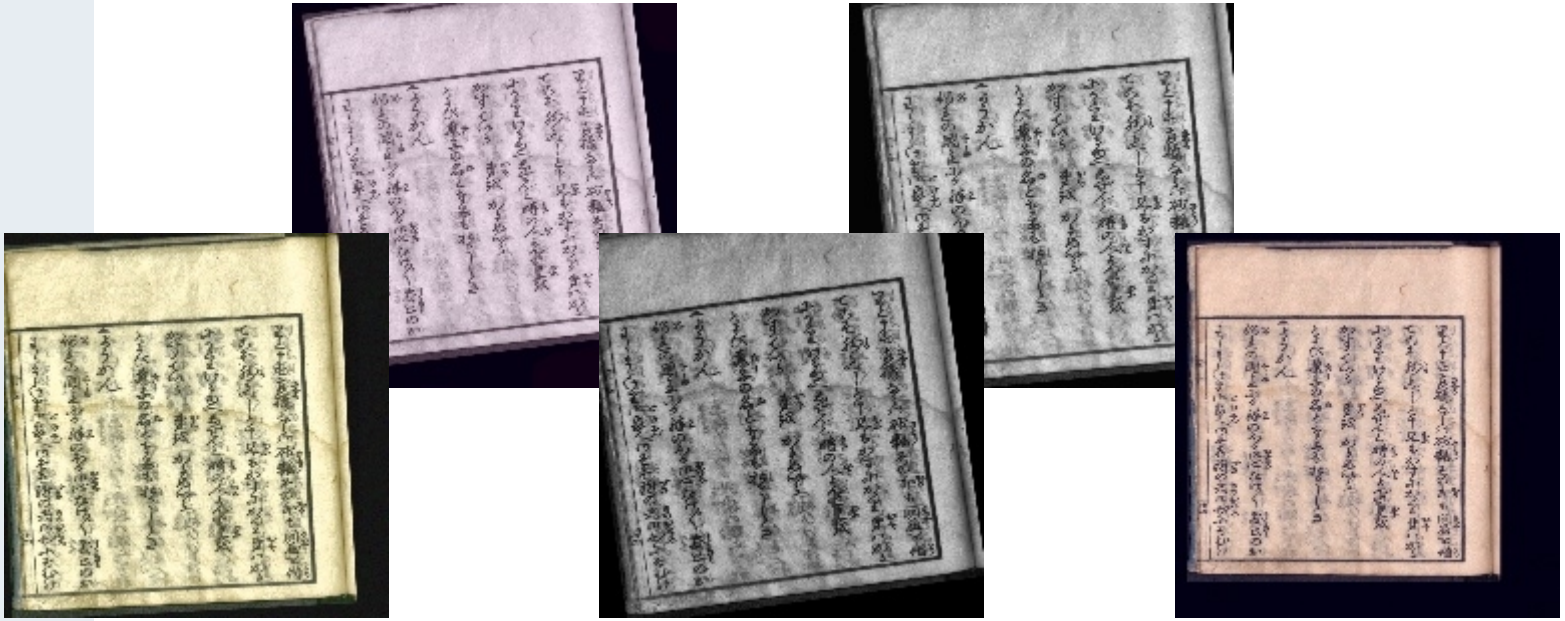

Lessons learned

- Many classes are no problem (4212 vs 80 on COCO)
- Context is very important for Kuzushiji Recognition
- Crop based classifier only reached 0.85 F1-score
- CNNs work well with very imbalanced datasets

Shakeup

- Little shakeup
- Public and private scores are close
  → Great dataset!

#	△pub	Team Name	Notebook	Team Members	Score	Entries	Last
1	–	tascj			0.950	13	24d
2	–	Konstantin Lopuhin			0.950	60	23d
3	–	Kenji			0.944	161	23d
4	▲1	YoudaoOCR			0.942	49	23d
5	▼1	See--			0.940	42	25d
6	–	abc			0.939	15	23d
7	–	K_mat			0.934	20	23d
8	–	t-hanya			0.920	21	23d
9	–	Ollie, Nanashi, and Tom			0.910	35	23d
10	–	Zenkei_R&D			0.903	144	23d
11	–	masayai			0.903	12	23d
12	▲5	Kirill Brodt (shad nsk)			0.901	4	1mo
13	▲1	James Day			0.901	33	23d

Solution sharing

- Code and trained models are available
- https://github.com/see--/kuzushiji-recognition

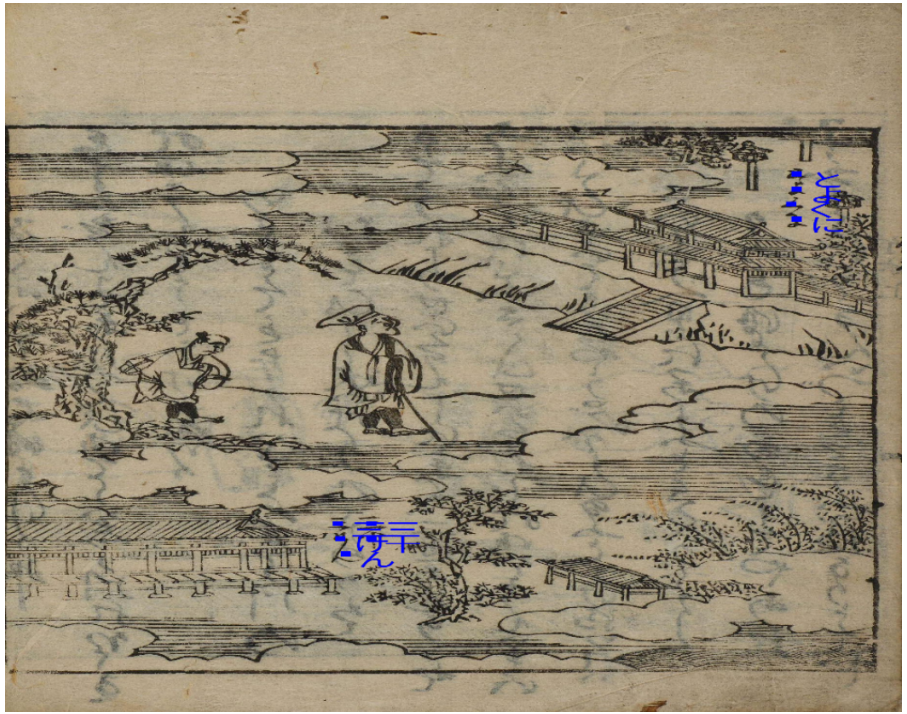

Thanks