



Picture Rotating and Labelling

Final Project Presentation

December 14th 2020

The Team

- Lennart Kitt (project owner, SEB)
- Arnel Pällo
- Karl-Kristjan Kõverik
- Sander Perens
- Oliver Ossip



The Problem

SEB gets different documents and real estate pictures that they need to

- Categorize
- Rotate upright
- Archive



Facade 90%



Document 60%

Acquired over 8000 photos from real estate portals



Preparing Training Data

- Removed borders and margins
- Labelled each one manually (gave up after 1000)
- Rotated all of them 90, 180 and 270 degrees (with a script)



Tested Different Models and Training Data

- Fast-AI with ResNet-34
- Build neural network ourselves with Keras
- Image-Net dataset
- One model with 12 classes vs several models with fewer classes
- Different image sizes
- Which augmentations improve performance?
- How big is our patience (can we tolerate training a model for 4 hours)?

Choosing the Best Approach

Confusion matrix

document0	25	23	0	0	0	0	0	0	0	0	0	0	0	0	0	
document180	25	26	0	0	0	0	0	0	0	0	0	0	0	0	0	
document270	1	1	29	10	0	0	0	0	0	0	0	0	0	0	0	
document90	0	1	18	25	0	0	0	0	0	0	0	0	0	0	0	
facade0	1	0	0	0	69	1	1	1	59	1	0	0	0	0	0	
facade180	0	0	0	0	11	63	1	2	1	34	1	0	0	0	0	
facade270	0	0	0	0	2	5	28	48	1	1	22	15	0	0	0	
facade90	0	0	0	1	0	1	29	28	2	0	18	29	0	0	0	
interior0	0	0	0	0	21	3	2	1	63	1	1	2	0	0	0	
interior180	0	0	0	0	0	46	2	2	0	49	1	0	0	0	0	
interior270	0	0	0	0	0	5	17	20	2	3	18	31	0	0	0	
interior90	0	0	0	0	0	2	8	20	6	1	39	24	0	0	0	
	document0	document180	document270	document90	facade0	facade180	facade270	facade90	interior0	interior180	interior270	interior90				

Actual

Predicted

Confusion matrix

document0	102	37	0	0	0	0	0	0	0	0	0	0	0	0	0	0
document180	34	80	0	0	0	0	0	0	0	0	0	0	0	0	0	0
document270	0	0	128	14	0	0	0	0	0	0	0	0	0	0	0	0
document90	0	0	19	139	0	0	0	0	0	0	0	0	0	0	0	0
facade0	0	0	0	0	195	0	0	0	0	0	0	0	0	0	0	0
facade180	0	0	0	0	2	191	0	0	0	0	0	0	0	0	0	0
facade270	0	0	0	0	0	1	119	78	0	0	0	0	0	0	0	0
facade90	0	0	0	0	1	0	87	91	0	0	0	0	0	0	0	0
interior0	0	0	0	0	0	0	0	0	86	2	0	0	0	0	0	0
interior180	0	0	0	0	0	0	0	0	1	196	0	0	0	0	0	0
interior270	0	0	0	0	0	0	0	0	0	2	24	88	0	0	0	0
interior90	0	0	0	0	0	0	0	0	2	2	83	142	0	0	0	0
	document0	document180	document270	document90	facade0	facade180	facade270	facade90	interior0	interior180	interior270	interior90				

Actual

Predicted

Training data



"It's a driver's license"

Model



KITCHEN

Results

- Best performance with Fast-AI using ResNet-34
- Used one model with 12 classes
- Model classified between documents, facades and interiors successfully
- Also classified 0 and 180 degree rotations successfully
- Confused at 90 vs 270 degrees
- Solution: if classified as rotated 90 or 270 degrees, rotate 90 degrees, then classify again. Finally rotate 180 degrees if needed

Architecture and workflow

- ipynb notebook for training and exporting the model (pickle)
- Command line script imports the model, then runs the classification and rotation
- Automated process in SEB systems.
 - Embed confidence in filenames of exported pictures
 - If model $\geq 90\%$ confident, proceed
 - Else delegate for manual review

Next Steps

- Code committed to GitHub
 - `https://github.com/perens/picture-labelling`
- SEB engineers will take ownership

Lessons Learned

- Labelling is laborious. Maybe look for a career in unsupervised learning?
- Smaller images classify just as successfully, but train much faster...
- ... except when it's not the case (images of documents)
- Augmentation does not always help. It did not make our model more accurate
- FastAI worked wonderfully for us. Invest time into learning it!
- Some things don't work and we have no idea why
- Some things do work but we have no idea why

Fast-AI



[classified as easy-peasy]

Thank you!

- SEB and Lennart Kitt for an interesting and very well defined problem
- Dima and TAs for a VERY well-organized and easy to follow course