View Reviews

Paper ID
40

Paper Title
CIL Road Segmentation Project Report

Reviewer #1

Questions

2. Quality of Paper (30%)
6.0
The paper is well structured. The figures are diverse and help visualise the models' behaviour effectively. The explanation of the CRF extension is particularly detailed and insightful.

3. Creativity of Idea (20%)
6.0
The authors base their solution on a publicly available neural network, augment it with multi-task learning as well as a custom postprocessing CRF module and train it using augmented images.

4. Quality of Implementation (20%) (Note: this is not just code, but an assessment of how the team went about implementing their idea, including, for example, which experiments they have chosen to run)
6.0
The comparison to naive baselines as well as base models is well made. The experiments performed show how each addition affects the model's predictions, and the ablation study shows the improvements achieved through the modifications. The code documentation is particularly complete.

Reviewer #2

Questions

2. Quality of Paper (30%)
5.5
- Well-written text overall, but some issues in language and several typos in some sections.

- Baselines are clearly stated and the contributions of the authors are highlighted well.

- The rather mathematical nature of the idea is accompanied well with suitable mathematical notation and is, as a result, very easy to follow.

- Very illustrative figures are used to explain the core idea on an intuitive, geometrical level as well.

- Very interesting read!

3. Creativity of Idea (20%)
6.0
- The authors combine lots of different previous works. They use PSPnet which won some competition four years ago. They introduce auxiliary losses based on center line and edge prediction and they leverage the CRF literature

to enforce more smoothness in the output geometry.

- They introduce a new kernel that is adapted to the geometry of the roads considered in this dataset.

- The idea of measuring closeness of pixel lines on roads makes a lot of sense! Using this idea you came up with a well-defined similarity measure and performed nice reductions of the problems (for instance only consider horizontal, vertical and diagonal lines etc). Moreover, you also derived a convincing way to measure these pixel lines and implemented everything into your framework. Great job!!

4. Quality of Implementation (20%) (Note: this is not just code, but an assessment of how the team went about implementing their idea, including, for example, which experiments they have chosen to run)
5.5
- Code is well-documented with instructions in the README file on how to reproduce the results.

- Lots of ablation studies are performed to investigate the effect of post processing with CRF, data augmentation and adding the new kernel. Great that you can show improvements for the considered factors!

- I'm missing a comparison to a different, maybe weaker baseline (from another model class). As far as I can tell, you only have the PSPNet without modifications as a baseline. A simple convolutional network approach for comparison would have been informative.