Scene Representation II
Lane boundaries

D.A. Forsyth

Huh?!?!

- Not even in "Computer Vision for Autonomous Vehicles"
  - (recent review by Janai et al - very good)
- Lane boundaries are very important
  - lots of money in good lane boundary detection
  - easy cases are firmly solved; hard cases remain hard
- Interplay of detection, geometry
 - variance and bias
- Firmly scene understanding

FIGURE 4A

Strategy: detect markers (reflective paint), join up
exercise in robust fitting of curves

US 9081385

Issues

- You have to do it fast
- You have to do it right
- Paint detection problems
- Geometric model problems

Missing paint, dirt

Curvy road

No paint

Labelled data methods

- Generally, rack up a labelled dataset and regress
  - Datasets - see
- But there are issues here, too
  - interaction between lane and boundary?
  - rectify?
  - explicit sequence modeling for boundaries?
  - image vs video?

Road and boundary interact

Fig. 1. Reciprocal constraints with geometric relation. **Left:** With an image input, traditional methods generate a binary segmentation mask for lane areas (green) or lane boundaries (red), which are severely affected by outlier situations. **Right:** Our approach introduces a geometric constraint into a multi-task network, which is capable to restore the missing lane area and lane boundaries (blue) mutually.

Zhang et al

Fig. 2. The proposed multi-task framework. Input images are fed into a shared encoder, which extracts the critical features for lane segmentation and lane boundary detection. Two inter-link encoders connected to each decoder provide complementary information for tasks. The overall performance is enhanced by introducing a structure loss, assuming that the lane boundary is predicted as the outer contour of the lane area while the lane area is predicted as the area integration result within the lane boundaries.

Zhang et al

Use IoU loss on predicted road
area - this strongly penalizes errors
at the boundaries

From predicted boundaries, compute
predicted road area, then compare
to true road area

Fig. 5. Boundary-aware loss and area-aware loss. **Left:** An illustration of our boundary-aware loss. The blue area indicates boundary inconsistency. **Right:** An illustration of our area-aware loss. Different intensities in prediction areas indicate different prediction confidence. The difference between restored area and ground truth indicates the area aware loss.

(a) Ours

(b) Up-Conv-Net

(c) U-Net

(d) Our Lane Boundary Detection

Fig. 7. Lane area segmentation and lane boundary detection results on KITTI dataset. Green corresponds to true positives, blue to false positives and red to false negatives.

Zhang 18: Notice

- No explicit geometric model of
  - boundary
  - drivable region
- Implicit models from
  - labelled data
  - interaction
- No motion cues or filtering
  - weird
- Q:
  - can we use biased geometric models to
    - improve performance
    - avoid problems with dataset frequencies
    - deal with missing, dirty, etc. paint

Khan et al 20

- Strategy
  - detect keypoints in image
  - rectify
    - using estimated horizon from vanishing points
 - impose structural model on keypoints in rectified image

Rectification

- Imagine a camera at a fixed height
  - moving rigidly over a textured ground plane
  - bottom half of image is distorted ground plane texture
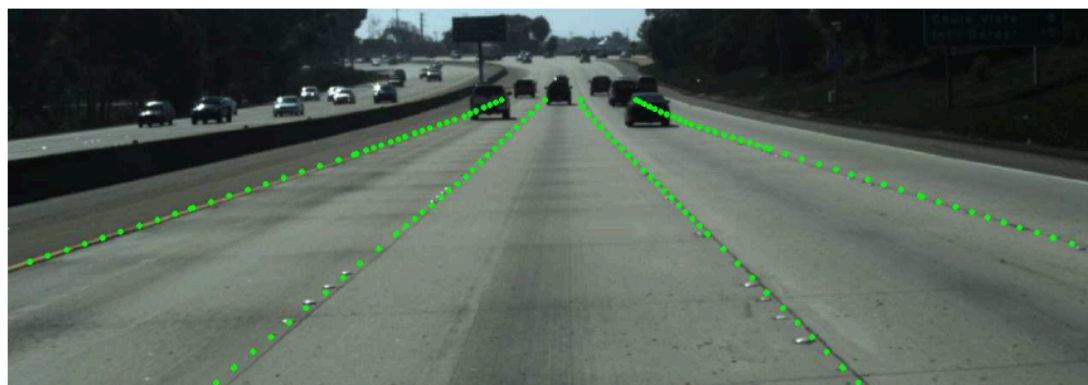 - If we know the camera, we can map image plane texture to ground plane

z

y

x

horizon (v=0)

(u, v)=(-s/t, -h/t)

(s, t, -h)

Plane   z=-h

Estimating the camera

- ## Height
  - from car (calibrated and known)
- ## Roll and pitch
  - from horizon
    - roll is why horizon isn't parallel to image plane
    - pitch is why it isn't centerline

z

y

x

horizon (v=0)

(s, t, -h)

(u, v)=(-s/t, -h/t)

Plane   z=-h

Encoder Network   Decoder Network

3xHxW   2048xH/6xW/6   x4   (K*L)xH/2xW/2   Lane Marker Map

Convolution layer   BatchNorm layer   Transpose Convolution layer   Relu layer   Pooling layer
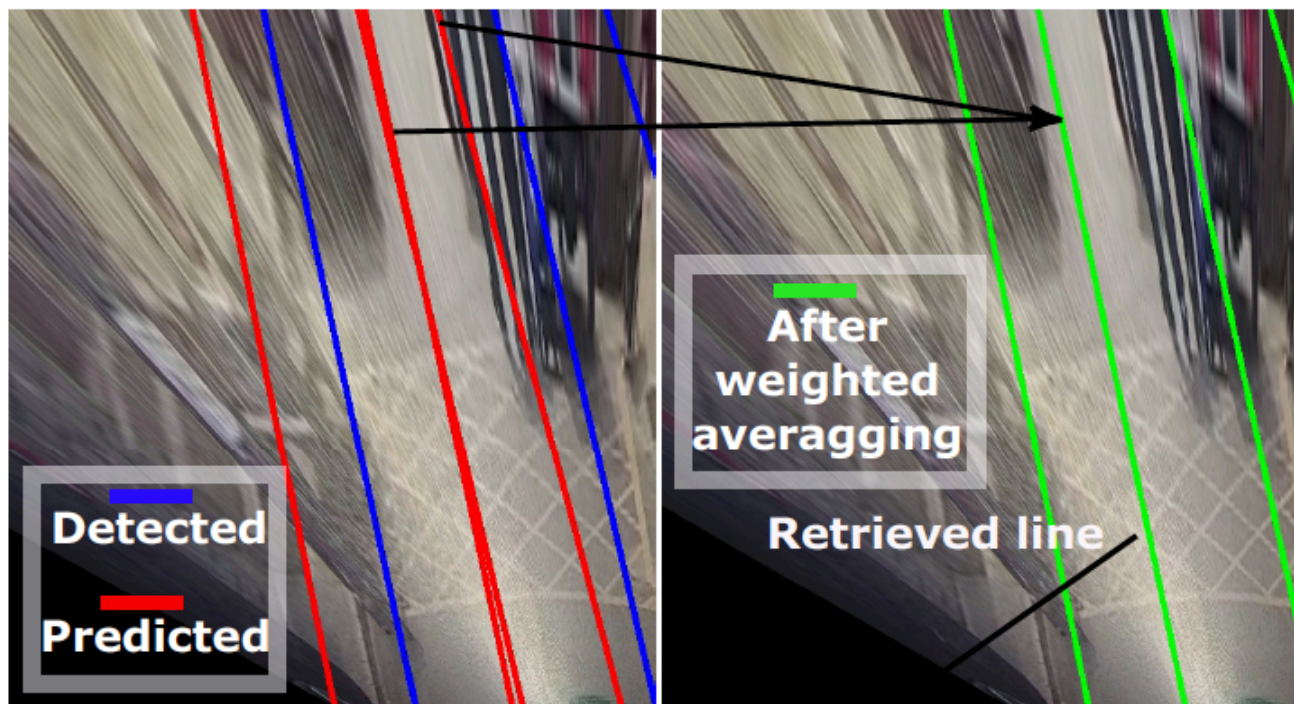
(a)

(b)

Figure 2. (a) The figure shows the architecture of the lane boundary marker network. (b) The sampled keypoints from the ground truth lane line are shown here.

Key point: it's easier to impose a geometric model on keypoints in rectified frame.

(a)

(b)

(c)

Figure 5. Detecting missing lane boundries. (a) Rectified view. Lane boundaries are predicted in red using c_o from section 3.3.2. (b) Filtered lane boundaries after weighted averaging (c) Recovered perspective view with all the four lane boundaries.

Geometric model

- Lane markers lie on
  - a quadratic curve OR a straight line (in rectified frame)
    - fitted using a version of RANSAC
  - lane boundaries are parallel
    - lines - easy
    - curves - look at tangents
    - Q: why not use a (latent) center curve?
- Search:
  - There are four boundaries (three lanes)
   - or three
    - or two

Figure 1. Sample results of our algorithm on examples from four different classes of CULane dataset [33] are shown here. Cyan lines are the detected lane boundaries, green region represents the ego lane and magenta line displays the estimated horizon. In the *No Line* class, there is actually no line markings on the road but the ground truth carries the lines shown.

Khan 20: Notice

- Current SOA on many datasets
  - for list of datasets, see Khan 20 - v. good on this
- Q:
  - what about less structured drivable regions?
  - can this be learned with less data? or none?
    - need data to learn keypoint finder
  - can rectification estimate be improved
    - better horizon finders out there - see Jacobs papers on website
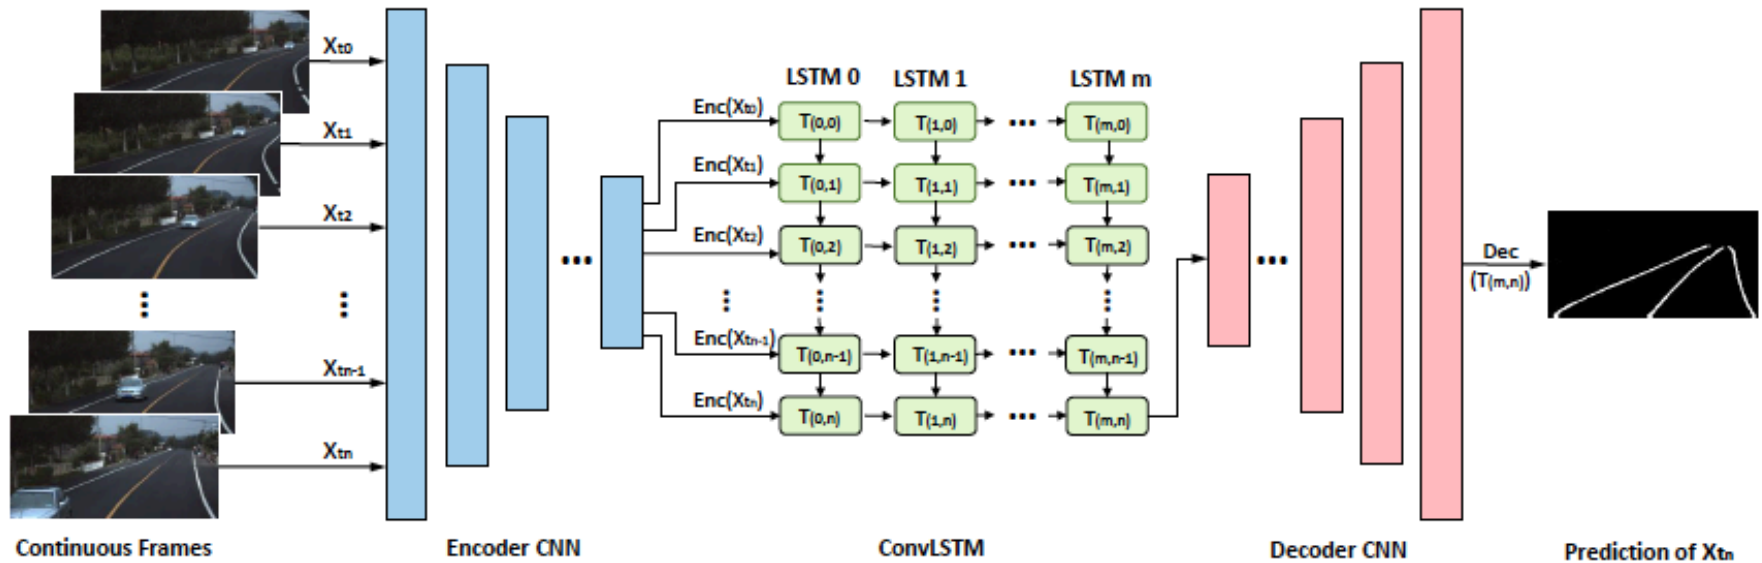
Zou 19: sequence models

Fig. 2: Architecture of the proposed network.

Training - piles of marked up data

(a)   (b)

Fig. 7: Results obtained by UNet-ConvLSTM on challenging scenes in (a) Testset #1 and (b) Testset #2 without post-processing. Lanes suffered from complications of occlusion, shadow, dirty, blur, curve, tunnel and poor illuminance are captured in rural, urban and highway scenes by automobile data recorder at different heights, inside or outside the front windshield.

TABLE IV: TuSimple lane marking challenge leader board on test set as of March 14, 2018 [52].

Rank	Method	Name on board	Using extra data	Accuracy(%)	FP	FN
1	Unpublished	leonardoli	–	96.9	0.0442	0.0197
2	Pan et al. [50]	XingangPan	True	96.5	0.0617	**0.0180**
3	Unpublished	astarry	–	96.5	0.0851	0.0269
4	Ghafoorian et al. [52]	TomTom EL-GAN	False	96.4	**0.0412**	0.0336
5	Neven et al. [43]	DavyNeven	Flase	96.2	0.2358	0.0362
6	Unpublished	li	–	96.1	0.2033	0.0387
	Pan et al. [50]	N/A	False	96.6	0.0609	**0.0176**
	Pan et al. [50]	N/A	True	96.6	0.0597	0.0178
	SegNet_ConvLSTM	N/A	False	97.1	0.0437	0.0181
	SegNet_ConvLSTM	N/A	True	97.2	0.0424	0.0184
	UNet_ConvLSTM	N/A	False	97.2	0.0428	0.0185
	UNet_ConvLSTM	N/A	True	**97.3**	0.0416	0.0186

Seo 14: Simple detection

- Strategy:
  - rectify, threshold to get rough road (?!?), estimate boundary points
  - fit curve to boundary points using unscented kalman filter
   - and so repeatedly update while moving

Fig. 2: (a) A part of an input image (i.e., the yellow rectangle area) is transformed into a bird-eye view image (or an inverse perspective image) (b) to remove perspective effects and to facilitate image processing. (c) Our algorithm analyzes the image region at the front of our vehicle and detects drivable-region by applying intensity-thresholding.

Fig. 4: The road-shape model and the measurement model.
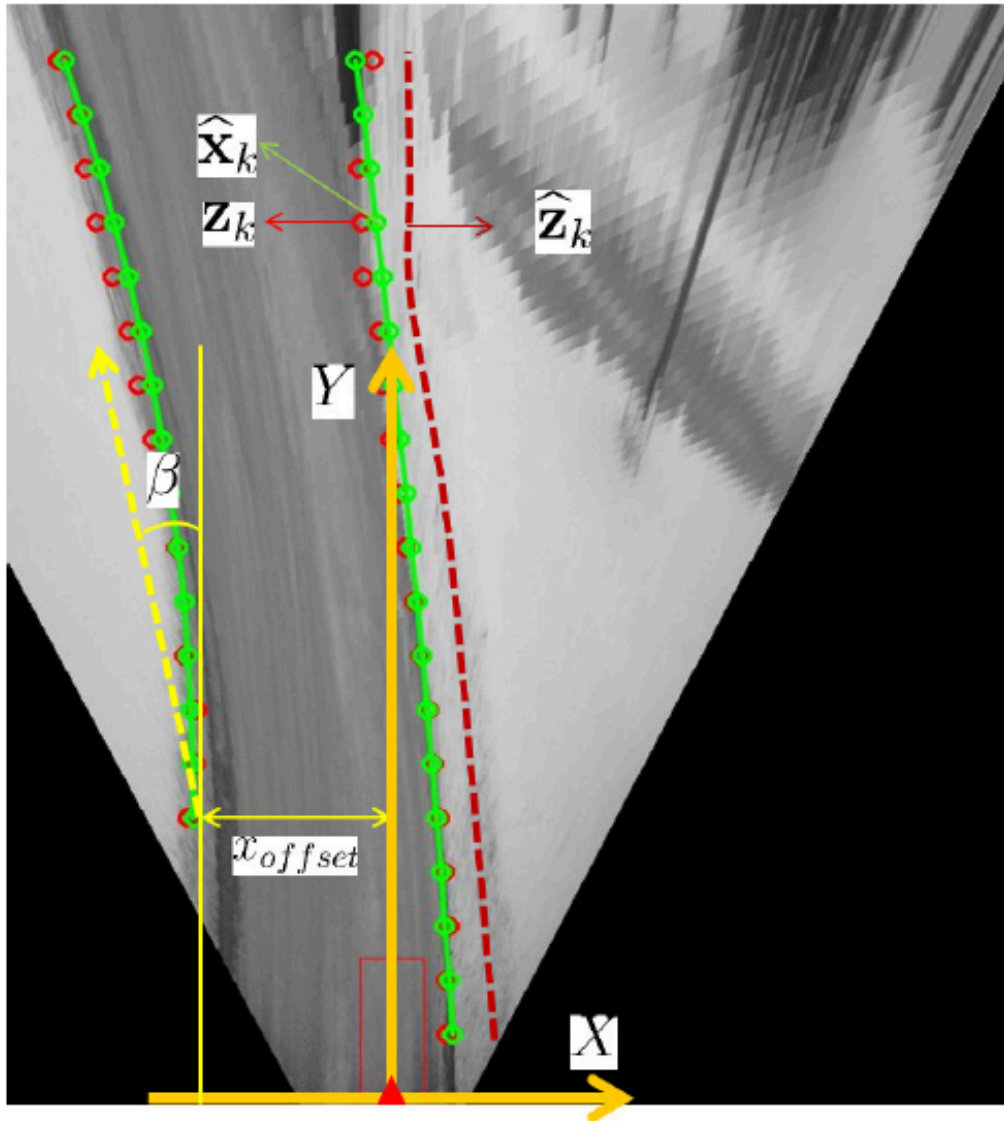
Clothoid

$$x = \frac{1}{6} c_1 y^3 + \frac{1}{3} c_0 y^2 + b y + x_0$$

curvature rate   curvature

heading angle (of road?)

offset of boundary from centerline of car

Filter state:

$$\mathbf{x} = [x_{offset}, \beta, c_0, c_1]^T$$

Fig. 5: An example sequence of drivable-regions' boundary detection and tracking.

Fig. 6: Example outputs of boundary detection and tracking.

Seo 14: Notice

- The detection process is minimal
  - essentially, a form of clustering
  - Q: is this a route to unsupervised lane detection?
    - boundary has rigid form
    - interior has coherent appearance
    - interior different from exterior