What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?

Alex Kendall • Yarin Gal

University of Cambridge • University of Oxford • The Alan Turing Institute
yarin@cs.ox.ac.uk

Bayesian deep learning

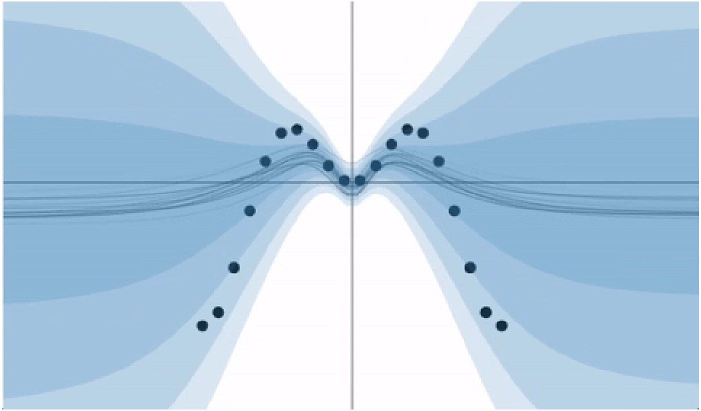

For example:

```
y = []
for _ in xrange(10):
    y.append(model.output(x, dropout=True))
y_mean = numpy.mean(y)
y_var = numpy.var(y)
```

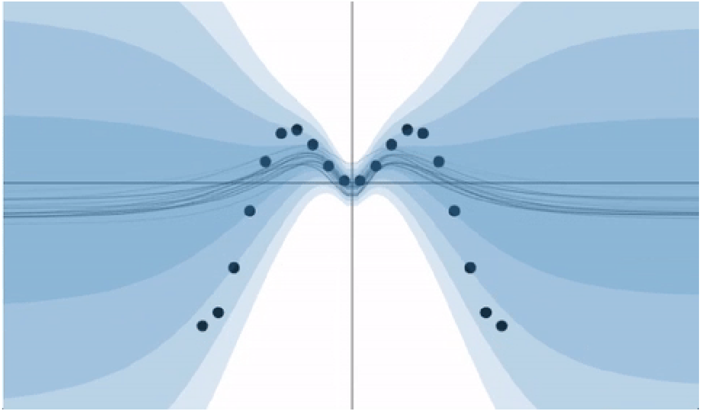

- What if we could capture uncertainty in modern computer vision?
  - Detect anomalies with image data
  - Identify adversarial examples
  - Learn with small amounts of labelled image data

Uncertainty in Computer Vision

- Not a new idea...
  - Particle filtering [Blake, Curwen, and Zisserman, 1993],
  - Conditional random fields [He, Zemel, and Carreira-Perpinan, 2004]
- Using BDL we can estimate uncertainty for modern computer vision models.
  E.g., Segnet: [Badrinarayanan, Kendall, and Cipolla, 2015]

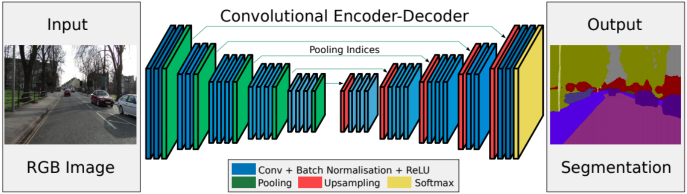

- But what uncertainty do we even want?
  There are many different types of uncertainty, including:
  - *Aleatoric uncertainty*, capturing inherent noise in the data
  - *Epistemic uncertainty*, capturing model's lack of knowledge

Uncertainty in Computer Vision

- Not a new idea...
  - Particle filtering [Blake, Curwen, and Zisserman, 1993],
  - Conditional random fields [He, Zemel, and Carreira-Perpinan, 2004]
- Using BDL we can estimate uncertainty for modern computer vision models.
  E.g., Segnet: [Badrinarayanan, Kendall, and Cipolla, 2015]

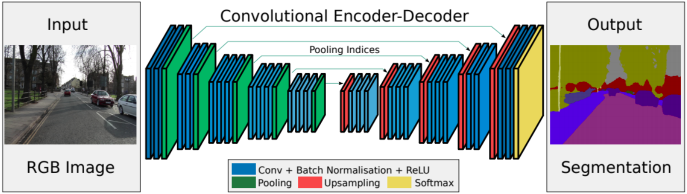

- But what uncertainty do we even want?
  There are many different types of uncertainty, including:
  - *Aleatoric uncertainty*, capturing inherent noise in the data
  - *Epistemic uncertainty*, capturing model's lack of knowledge

Uncertainty in Computer Vision

- Not a new idea...
  - Particle filtering [Blake, Curwen, and Zisserman, 1993],
  - Conditional random fields [He, Zemel, and Carreira-Perpinan, 2004]
- Using BDL we can estimate uncertainty for modern computer vision models.
  E.g., Segnet: [Badrinarayanan, Kendall, and Cipolla, 2015]

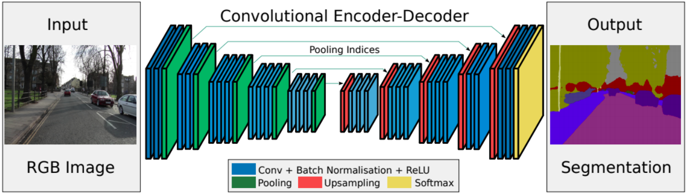

- But what uncertainty do we even want?
  There are many different types of uncertainty, including:
  - *Aleatoric uncertainty*, capturing inherent noise in the data
  - *Epistemic uncertainty*, capturing model's lack of knowledge

Uncertainty in Computer Vision

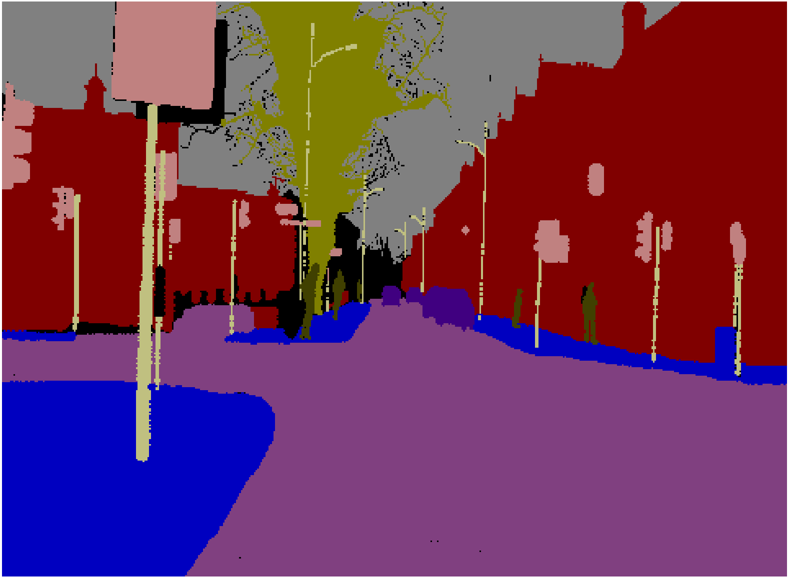

Uncertainty in Computer Vision

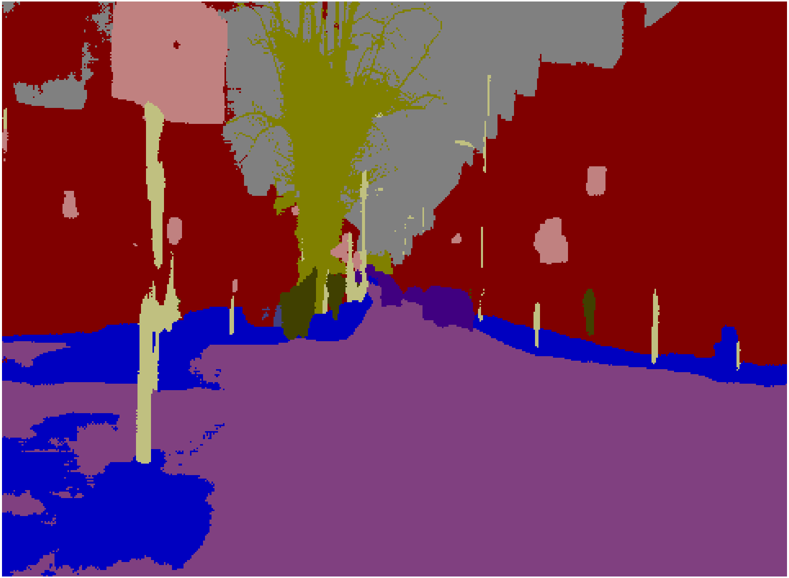

Uncertainty in Computer Vision

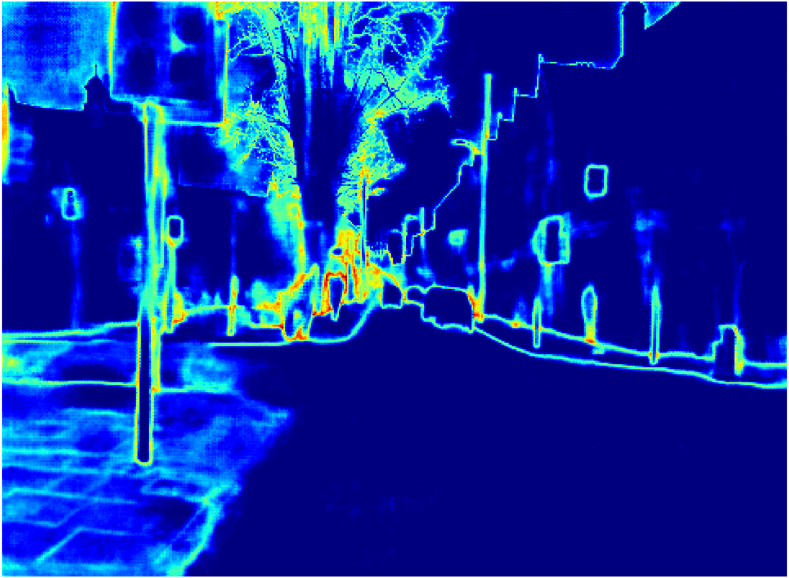

Uncertainty in Computer Vision

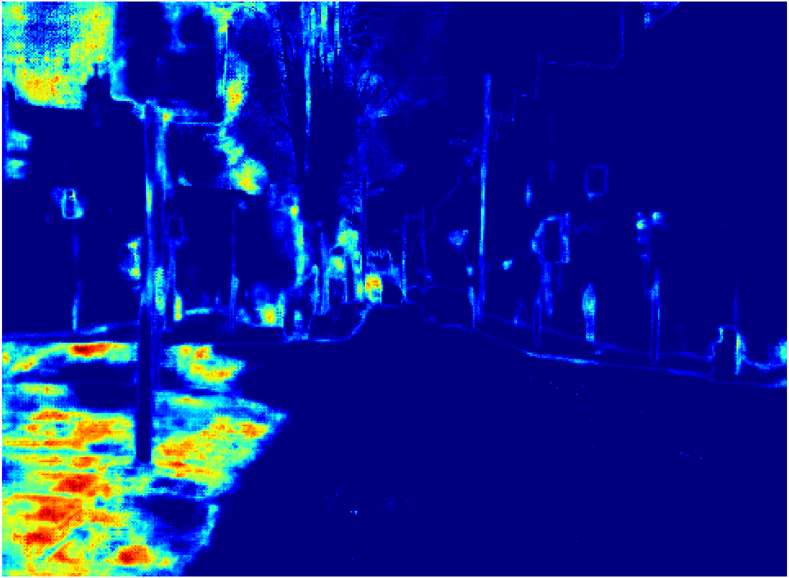

- *Aleatoric uncertainty*, capturing inherent noise in the data
- *Epistemic uncertainty*, capturing model's lack of knowledge

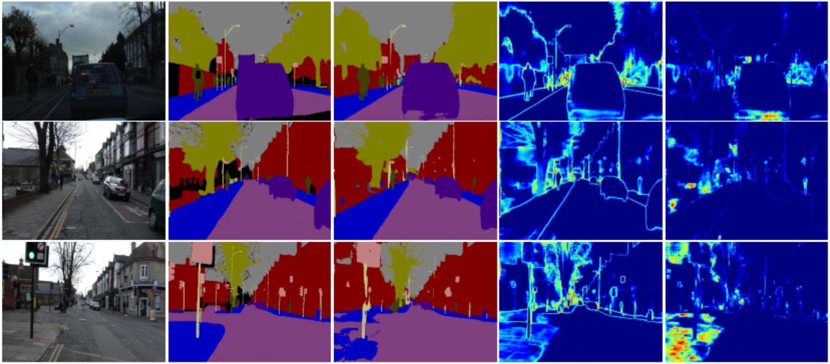

(a) Input Image (b) Ground Truth (c) Semantic Segmentation (d) Aleatoric Uncertainty (e) Epistemic Uncertainty

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *[Kendall & Gal, NIPS, 2017]*

Uncertainty in Computer Vision

- *Aleatoric uncertainty*, capturing inherent noise in the data
- *Epistemic uncertainty*, capturing model's lack of knowledge

Can we detect anomalies with Segnet?

Train dataset	Test dataset	RMS	Aleatoric variance	Epistemic variance
Make3D / 4	Make3D	5.76	0.506	7.73
Make3D / 2	Make3D	4.62	0.521	4.38
Make3D	Make3D	3.87	0.485	2.78
Make3D / 4	NYUv2	-	0.388	15.0
Make3D	NYUv2	-	0.461	4.87

(a) Regression

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *[Kendall & Gal, NIPS, 2017]*

- *Aleatoric uncertainty*, capturing inherent noise in the data
- *Epistemic uncertainty*, capturing model's lack of knowledge

Can we detect anomalies with Segnet?

Train dataset	Test dataset	RMS	Aleatoric variance	Epistemic variance
Make3D / 4	Make3D	5.76	0.506	7.73
Make3D / 2	Make3D	4.62	0.521	4.38
Make3D	Make3D	3.87	0.485	2.78
Make3D / 4	NYUv2	-	0.388	15.0
Make3D	NYUv2	-	0.461	4.87

(a) Regression

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *[Kendall & Gal, NIPS, 2017]*

- *Aleatoric uncertainty*, capturing inherent noise in the data
- *Epistemic uncertainty*, capturing model's lack of knowledge

Can we detect anomalies with Segnet?

Train dataset	Test dataset	RMS	Aleatoric variance	Epistemic variance
Make3D / 4	Make3D	5.76	0.506	7.73 ↓
Make3D / 2	Make3D	4.62	0.521	4.38
Make3D	Make3D	3.87	0.485	2.78
Make3D / 4	NYUv2	-	0.388	15.0 ↑
Make3D	NYUv2	-	0.461	4.87

(a) Regression

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *[Kendall & Gal, NIPS, 2017]*

Come to our Poster!

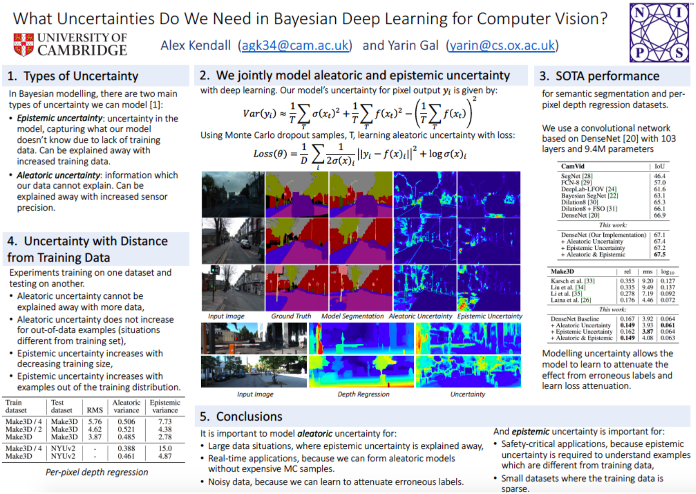

- Tue Dec 5th 06:30 – 10:30 PM Pacific Ballroom #95

What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *[Kendall & Gal, NIPS, 2017]*