



#### Summary

#### Coding up uncertainty models at CIFAR+ scale leads to bad times.

How do you build a Bayesian ResNet? Uncertainty for model-based RL? Support GPU/TPU and flexible training?

Core contributions:

- No reinventing the wheel. No new language. No new abstractions.
- Extend neural network library semantics to distributions over functions.
- Everything composes!

# What goes into a neural network library?

units,
activation=None,
use_bias=True,
<pre>kernel_initializer='glorot_uniform',</pre>
<pre>bias_initializer='zeros',</pre>
kernel_regularizer=None,
<pre>bias_regularizer=None,</pre>
activity_regularizer=None,
<pre>kernel_constraint=None,</pre>
<pre>bias_constraint=None,</pre>
**kwargs

Optimizers, losses, metrics, and GPU/TPU strategies are designed around them.

Initializers. **Regularizers. Compositionality.** 

## Example: Deep Gaussian Process

model = tf.keras.Sequential([tf.keras.layers.Flatten(), ed.layers.SparseGaussianProcess(units=256, num\_inducing=512), ed.layers.SparseGaussianProcess(units=256, num\_inducing=512), ed.layers.SparseGaussianProcess(units=10, num\_inducing=512), **def** loss\_fn(features, labels): predictions = model(features) $nll = tf.reduce_mean((labels-predictions.mean())**2)$ kl = sum(model.losses) **return** nll + kl/dataset\_size model.compile('adam', loss\_fn) model.fit(features, labels, batch\_size=64, epochs=100)

In all libraries, layers are the core building block.

### Bayesian Layers **A Module for Neural Network Uncertainty** Dustin Tran<sup>1</sup> Michael W. Dusenberry<sup>1</sup> Mark van der Wilk<sup>2</sup> Danijar Hafner<sup>1</sup> <sup>1</sup>Google Brain <sup>2</sup> PROWLER.io

# Example: Bayesian LSTM

lstm = ed.layers.LSTMCellReparameterization(512) $output_layer = tf.keras.layers.Dense(10)$ 

def loss\_fn(features, labels, dataset\_size): state = lstm.get\_initial\_state(features) nll = 0.for t in range(features.shape[1]):

- net, state = lstm(features[:, t], state) logits = output\_layer(net) nll += tf.reduce\_mean( tf.nn.softmax\_cross\_entropy\_with\_logits(
  - labels[:, t], logits))

kl = sum(lstm.losses) / dataset\_size **return** nll + kl

# Bayesian Layers

There are several design considerations for uncertainty models.

Computing the integral. We need to compute often-intractable integrals. For example in Bayesian neural networks:

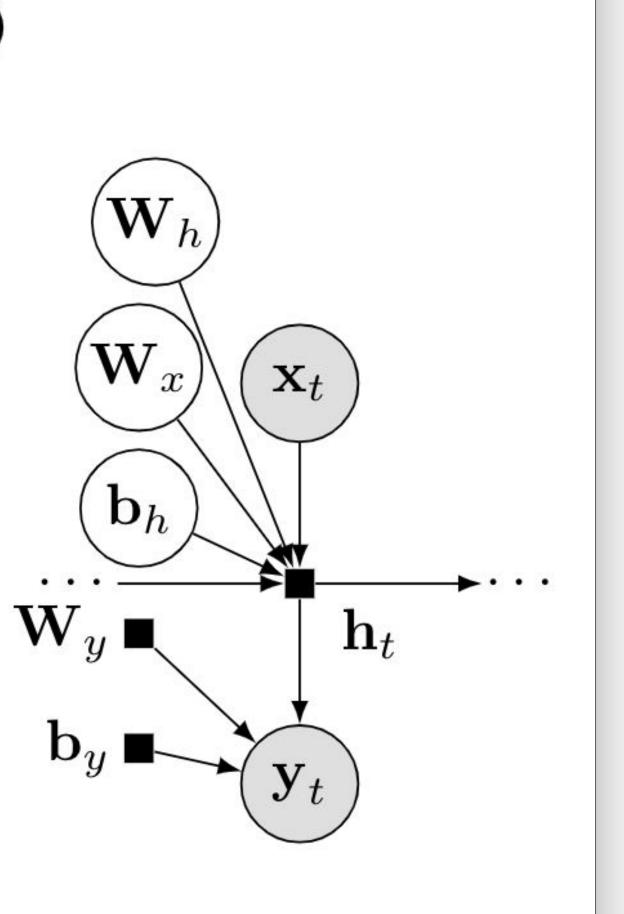
$$\text{ELBO}(\theta) = \int q(\theta) \log p(\mathbf{y} \mid q(\mathbf{y} \mid \mathbf{x})) = \int q(\theta) p(\mathbf{y} \mid f_{\theta}(\mathbf{y} \mid \mathbf{x})) d\theta$$

Make each estimator its own Layer.

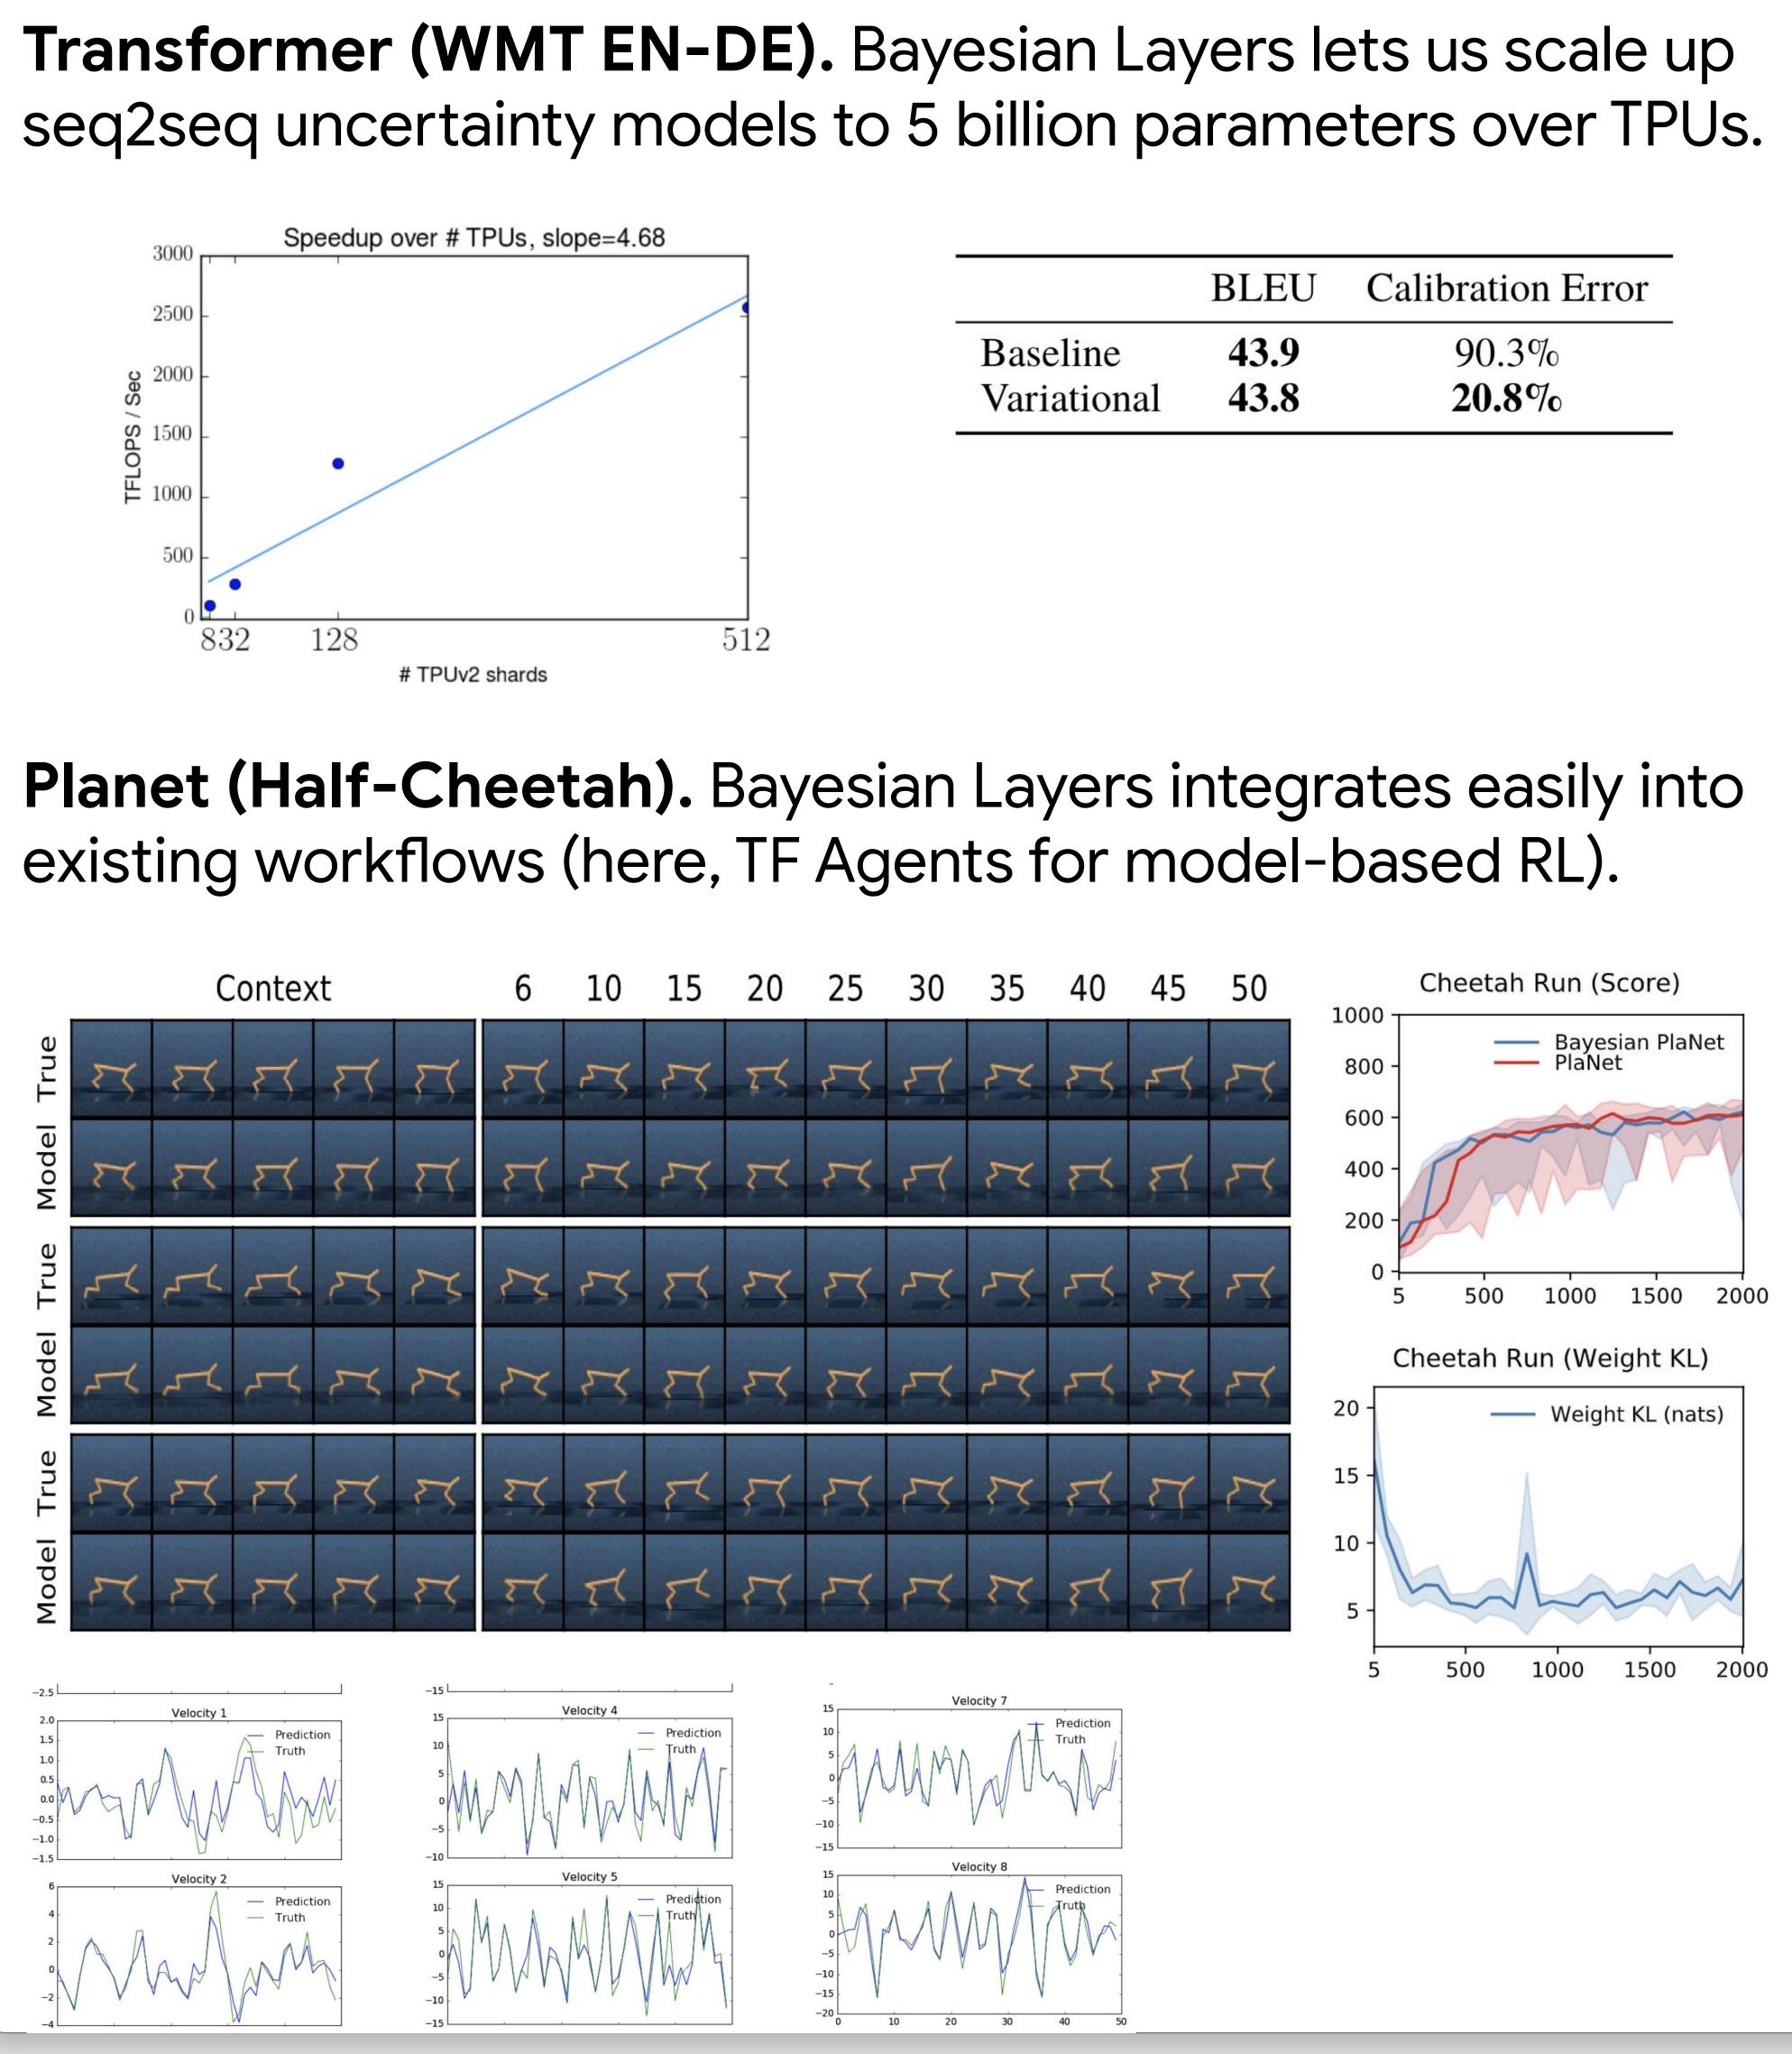
**Type Signature**. Bayesian layers are **drop-in replacements**.

- if FLAGS.be\_bayesian: Conv2D = layers.Conv2DReparameterizationelse: Conv2D = tf.keras.layers.Conv2D
- model = tf.keras.Sequential([ Conv2D(32, kernel\_size=5, strides=1, padding="SAME"), tf.keras.layers.BatchNormalization(), tf.keras.layers.Activation("relu"), Conv2D(32, kernel\_size=5, strides=2, padding="SAME"), tf.keras.layers.BatchNormalization(), . . .

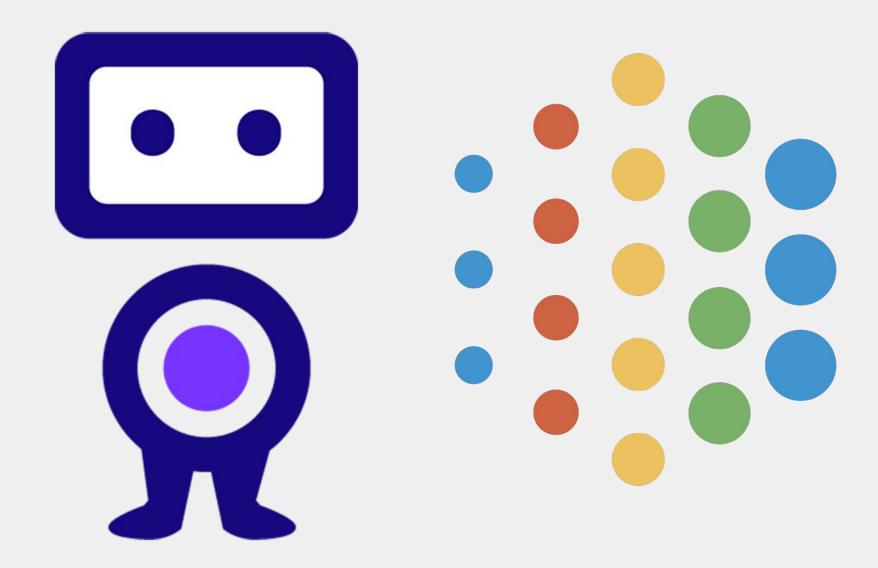
**Initializers**. Extend to enable learnable distributions. **Regularizers**. Extend to enable distributional penalties.



- $f_{\theta}(\mathbf{x}) d\theta \mathrm{KL} \left[ q(\theta) \| p(\theta) \right],$
- $(\mathbf{x})$  d $\theta$ .



**Uncertainty baselines.** High-quality, SOTA implementations of uncertainty models. (note: We already have CIFAR & ImageNet SOTA!)



### Experiments

	BLEU	Calibration Error
Baseline	43.9	90.3%
Variational	43.8	20.8%

# Current Work

 How does model uncertainty correlate with generalization? • How do we combine SGD randomization with explicit prior/posterior randomization? • How can uncertainty models achieve not just better robustness and reliability but better test accuracy?

#### Code: <u>github.com/google/edward2</u>