



Keras

Keras Community Meeting

August 5, 2022

Agenda

- **KerasCV** update [Luke, Zhenyu]
- **KerasNLP** update [Chen]
- **TensorFlow 2.10** release (RC0) [Haifeng]
- **Keras** optimizers [Chen]
- New **keras.io** examples [Haifeng]
- Community feedback & questions

KerasCV updates

- 0.2.10 released! Check the release notes [here](#).
 - First pre-trained weights now available for some image classification models
 - Ready for community training script contributions in next 2 weeks
- API sneak peek of the KerasCV 1-stage object detection API featuring
 - A fit() friendly RetinaNet, PascalVOC
 - In-graph train time COCO metric evaluation
 - KerasCV data augmentation
- [API Sneak Peek: Train an Object Detection Model on Pascal VOC 2007 using KerasCV](#)
- [RFC of 2-stage object detection API](#)

KerasNLP updates

- 0.3 released! Check the release notes [here](#).
- [Metrics](#) added: BLEU, ROUGE, Edit distance
- [Text generation utilities](#): greedy, beam search, top-N, top-P.

TensorFlow 2.10 release (RC0)

- Update to Attention and AdditiveAttention layer
 - `use_causal_mask` is in `call()` not in the constructor.
 - This is for consistency with other layers.
- `tf.keras.optimizers.experimental.Optimizer` will graduate in Release 2.11,
 - `tf.keras.optimizers.experimental.Optimizer` -> `tf.keras.optimizers.Optimizer`.
 - `tf.keras.optimizers.Optimizer` -> `tf.keras.optimizers.legacy.Optimizer`
- Behavior change for `tf.keras.initializers`
 - Keras initializers will now use stateless random ops to generate random numbers.
 - Both seeded and unseeded initializers will always generate the same values every time they are called (for a given variable shape).

TensorFlow 2.10 release (RC0)

- `tf.keras.layers.experimental.EinsumDense` -> `tf.keras.layers.EinsumDense`.
- `tf.keras.utils.audio_dataset_from_directory`
- `tf.keras.utils.***_dataset_from_directory(subset="both")`
- `tf.keras.utils.split_dataset`
- BackupAndRestore now support step recovery
- `tf.keras.dtensor.experimental.optimizers.AdamW`
- Improved masking support for `tf.keras.layers.MultiHeadAttention`.
- `ignore_class` in `SparseCategoricalCrossentropy`, `IoU` and `MeanIoU`
- `tf.keras.models.experimental.SharpnessAwareMinimization`.

Keras optimizer

- [Experimental optimizer](#), e.g., ``tf.keras.optimizers.experimental.Adam`` will graduate in 2.11, the old optimizer will keep being supported in ``legacy`` namespace.
- New optimizer only supports TF2.
- Easier to fork and customize.
- Clean code
 - No more slot variables.
 - Distributed training code isolated from other part.
 - No more fused ops, performance is guaranteed by XLA.

New **keras.io** examples

- keras.io has ~150 code examples now
 - Short and illustrative
 - Covering wide range of topics
 - Showing Keras best practices
- A big shout out to the contributors!
- New examples:
 - [Denoising Diffusion Implicit Models](#)
 - [Text Classification using FNet](#)
 - [English-to-Spanish translation with KerasNLP](#)

New **keras.io** examples

- [Denoising Diffusion Implicit Models](#)



Community feedback & questions