

**Desmond Elliott:** *Language Modelling from Pixels*

Language models are usually defined over a finite set of inputs, which creates a bottleneck if we attempt to scale the number of languages supported by a model. Tackling this bottleneck often results in a trade-off between what can be represented in the model and computational issues in the output layer. I will present the Pixel-based Encoder of Language, which suffers from neither of these issues by rendering text as images, making it possible to transfer representations across languages based on the co-activation of pixels. I will discuss the results of various models, pretrained on only English text, ranging from just 5M parameters up to 86M parameters on a variety of downstream syntactic and semantic tasks in 32 typologically diverse languages across 14 scripts.