Learning to Design CNNs with Reinforcement Learning

Aashima Arora[1], Jason Krone[1]

[1]Columbia University Department of Computer Science

Why is this a good idea?

CNNs are getting more complex and harder to design

VGG

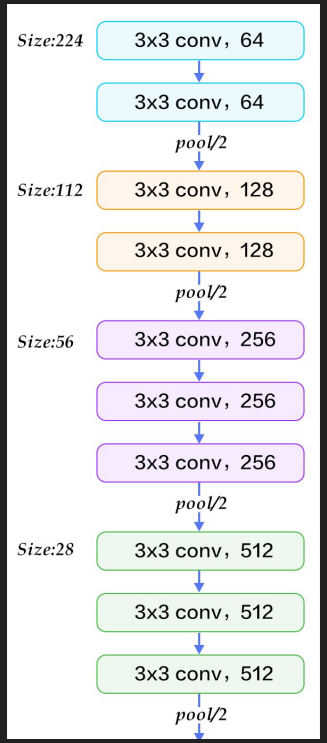

Inception

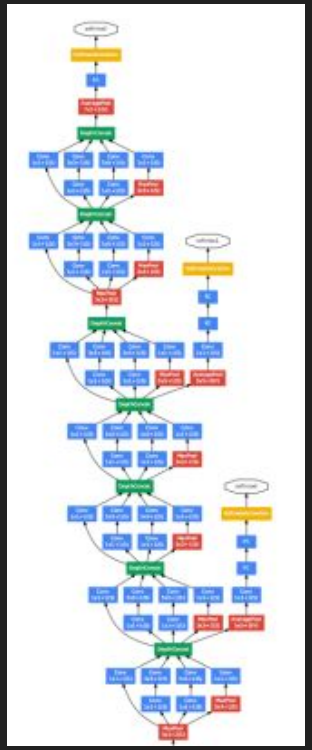

ResNet

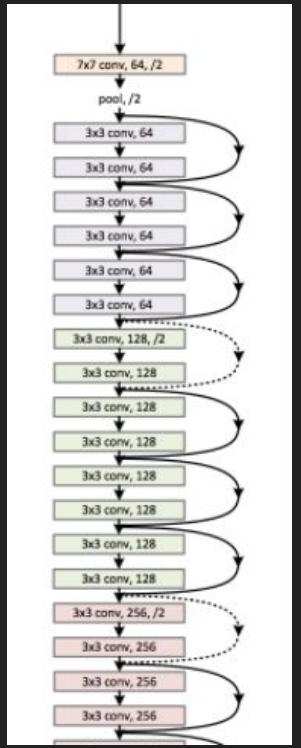

The best models are learnt end-to-end

- R-CNN
- LSTM+HMM
- RNN+SMT

Mask R-CNN (Object Detection)

LSTM + CTC (Speech Recognition)

Only Attention (Machine Translation)

Can we also learn cell architectures end-to-end?

Definition: CNN “cell” architectures

- “Cell”: pattern of repeated operations

VGG

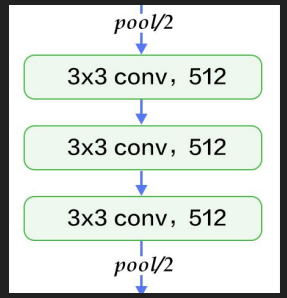

Inception

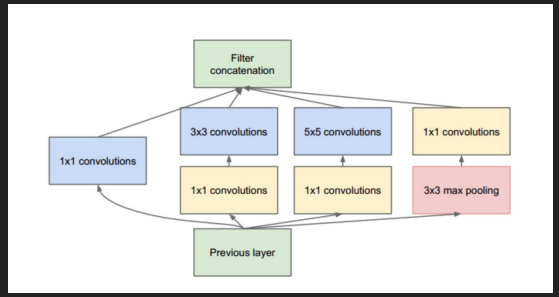

ResNet

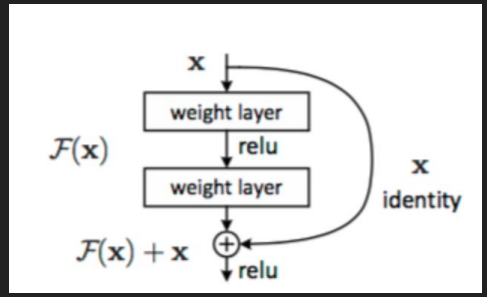

Idea: optimize cell for validation accuracy using RL

RL Framework

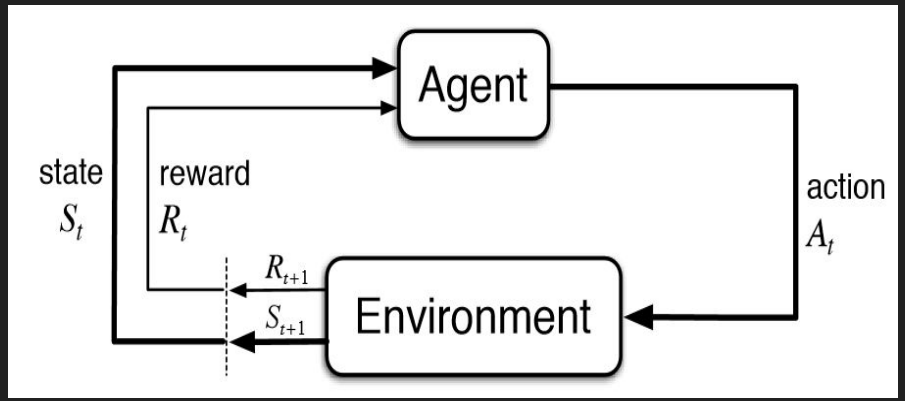

RL applied to Architecture Search

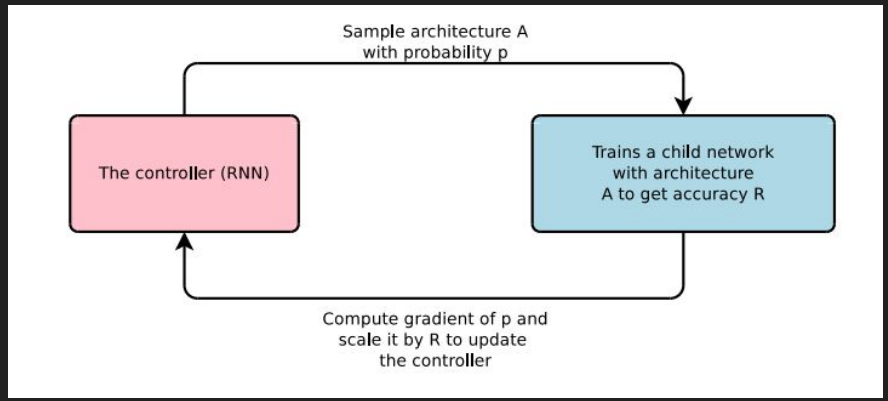

Training loop

1. Sample **actions** (child architectures) from **policy** (controller LSTM)
2. Train child networks on MNIST and report reward **reward** (val accuracies)
3. Form dataset D = (**actions, rewards)**
4. Update **policy**

Update policy:

Take K steps of SGD using proximal policy optimization (PPO) loss where mini-batches are drawn from D

Controller architecture dictates “search space”

We explore:

- A small search space
  - 64 possible child architectures
  - 2 operations applied consecutively define a cell
- A large search space
  - Thousands of possible child architectures
  - Cells allow skip-connections and branching (multiple operations applied to the same input)

Operations

- 1x3 then 3x1 conv
- 3x3 conv
- 3x3 depthwise-separable conv
- 5x5 conv
- 7x7 conv
- 2x2 max pool
- 2x2 average pool
- identitty

Controller network (small search space)

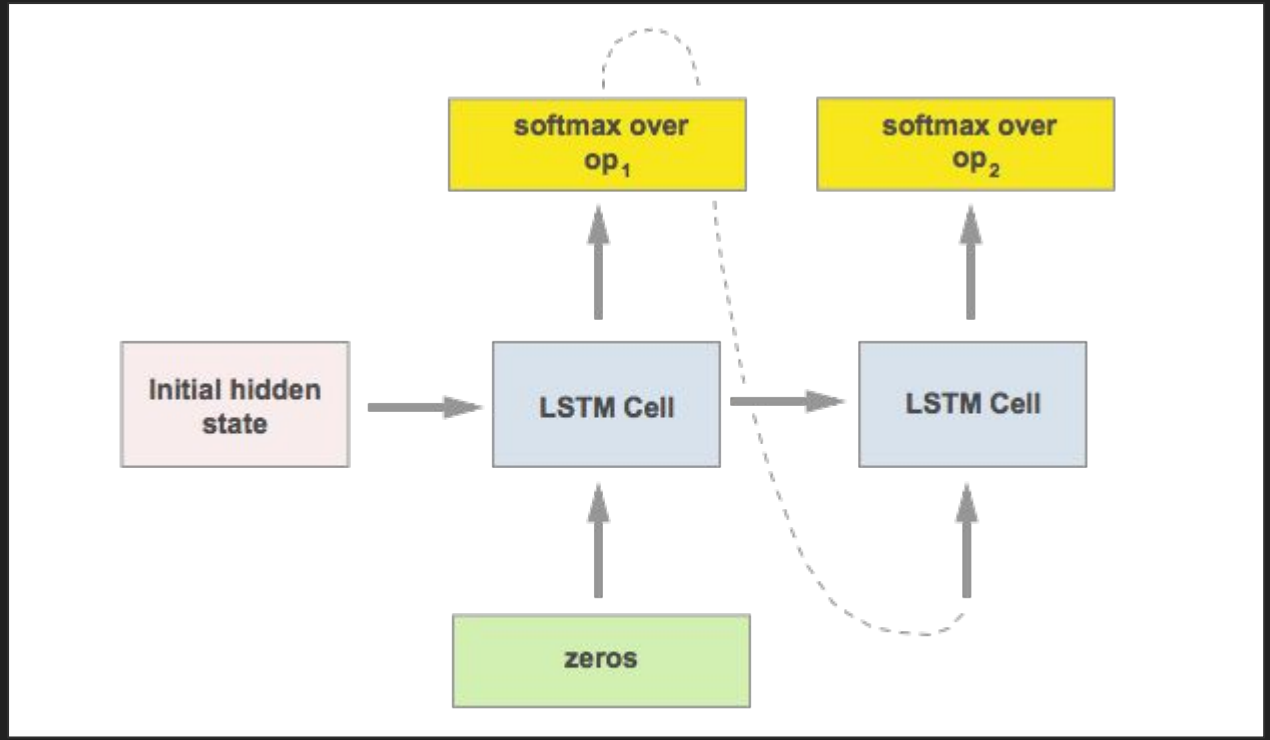

Controller network (large search space)

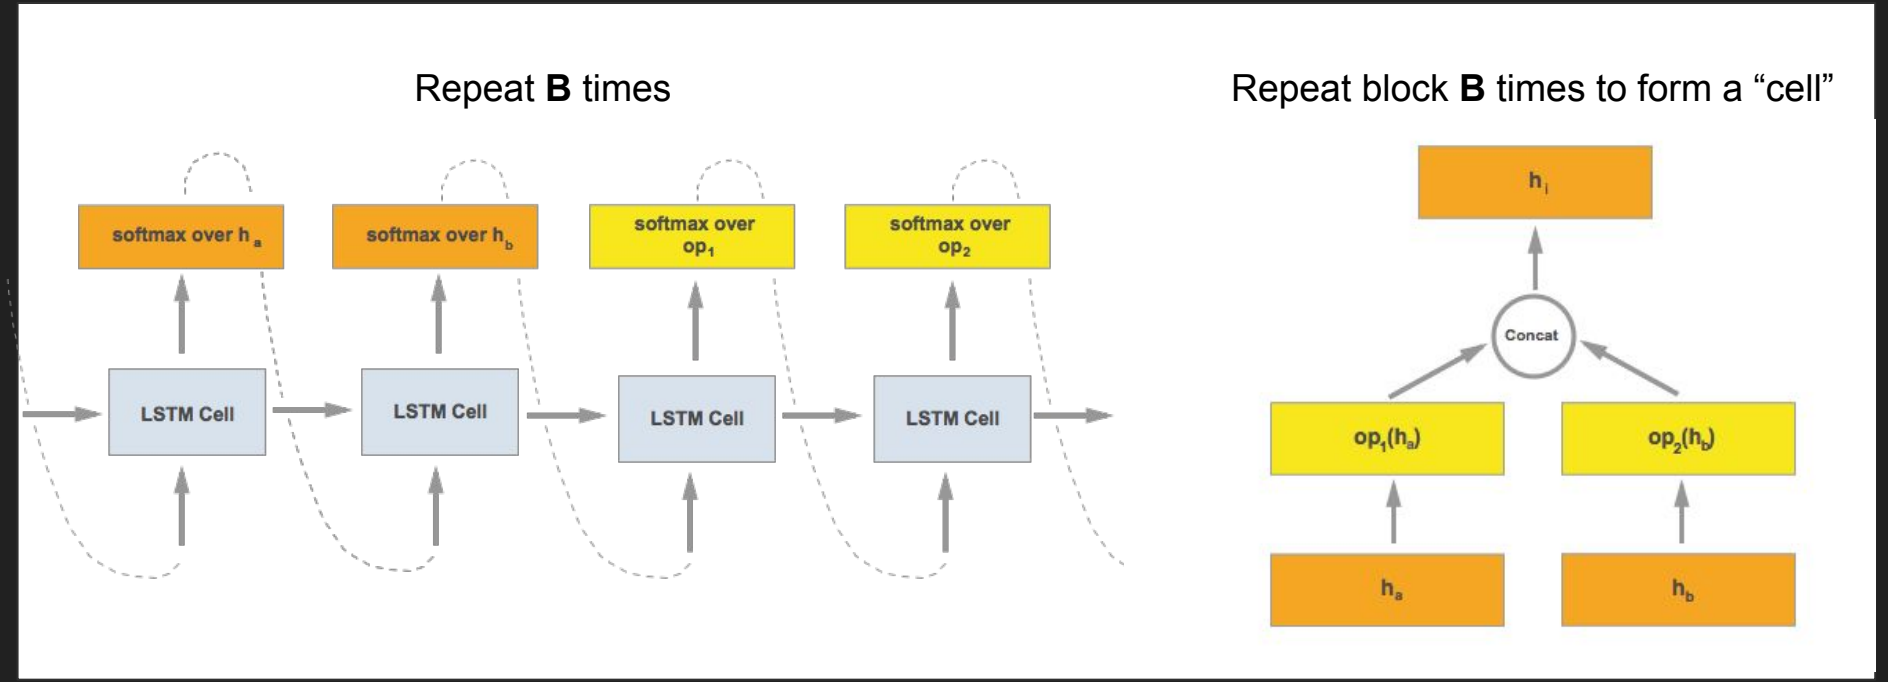

Child network

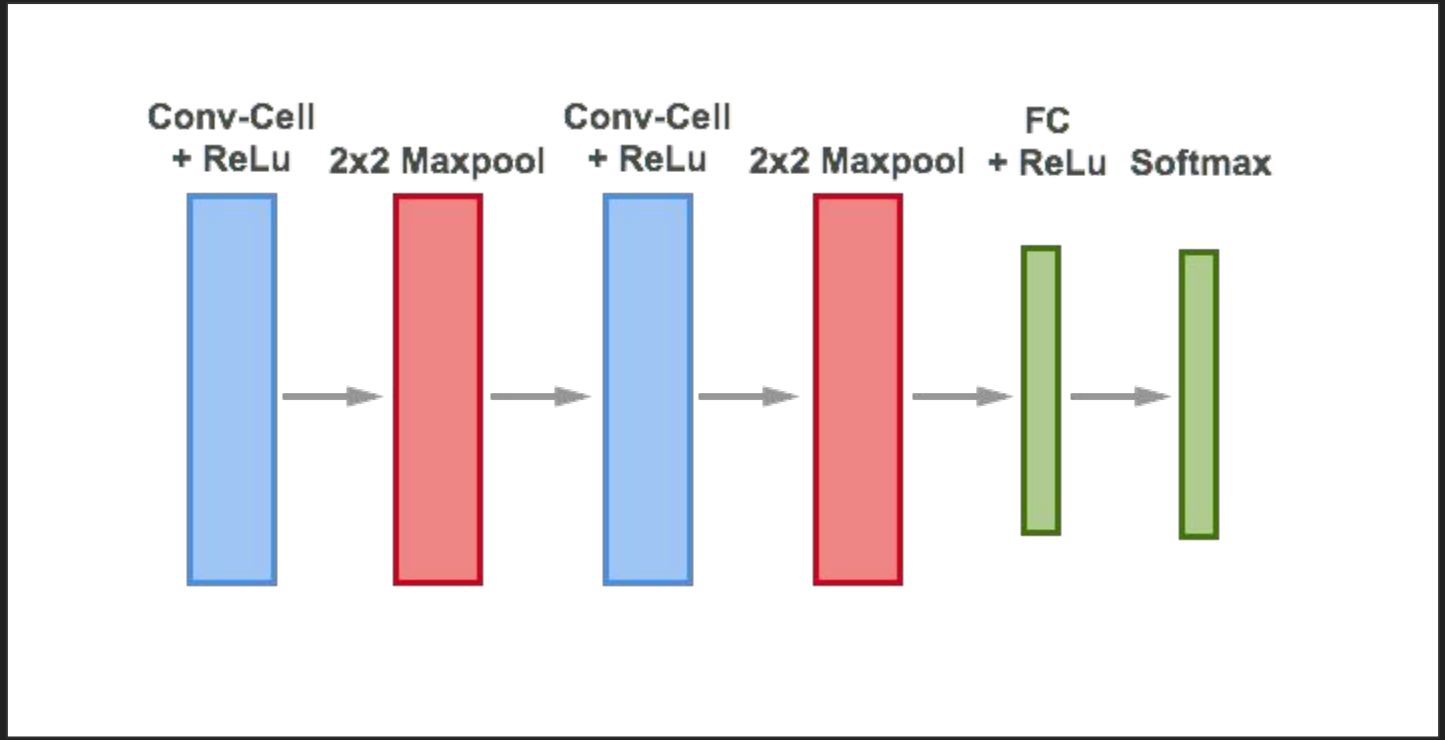

Pre-computed reward

- Only 8x8 = 64 possible child networks in small search space
- We compute reward ahead of time to speed up training

Small search space results (pre-computed reward)

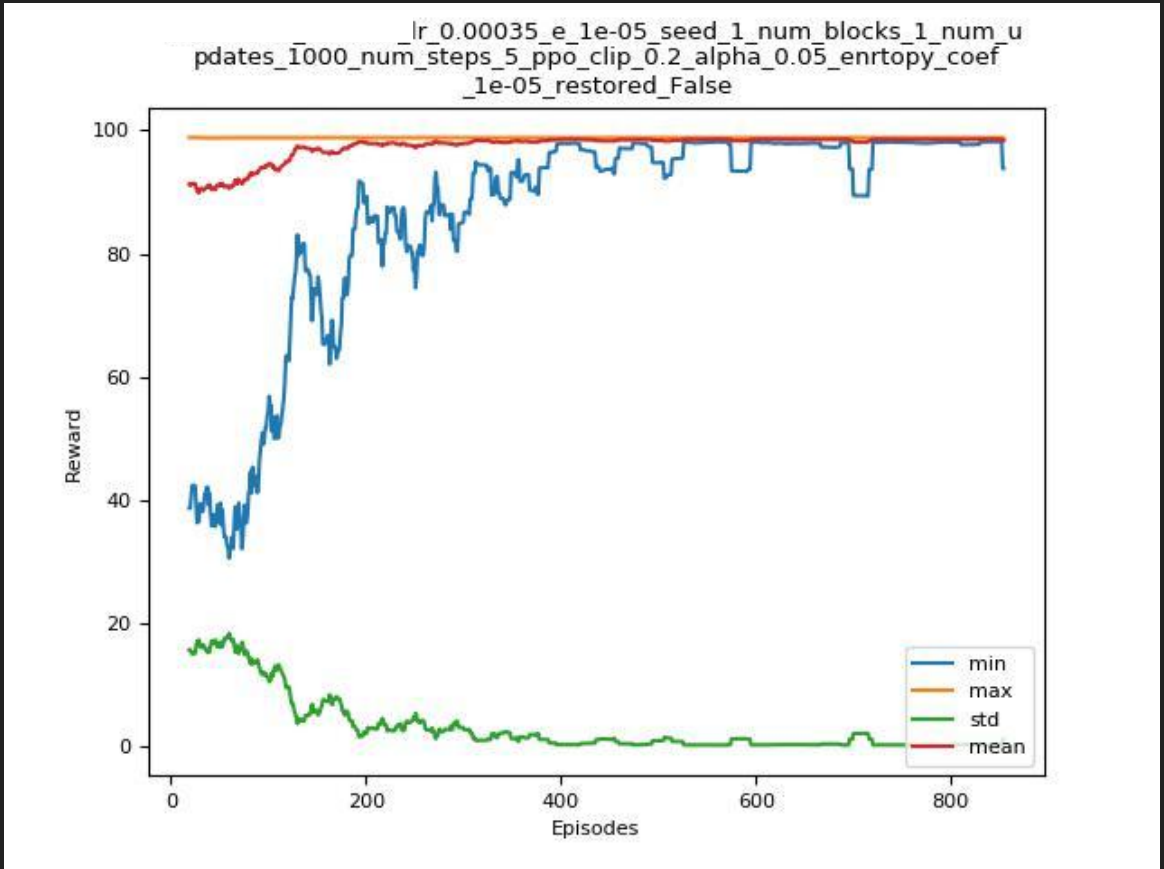

Small search space results (pre-computed reward)

Two best “cell” architectures:

Cell Architecture	Val Accuracy	Test Accuracy
5x5 conv, 1x3 then 3x1 conv	98.85	98.88
3x3 sep conv, 7x7 conv	98.85	98.2

Large search space results

Best “cell” architecture thus far:

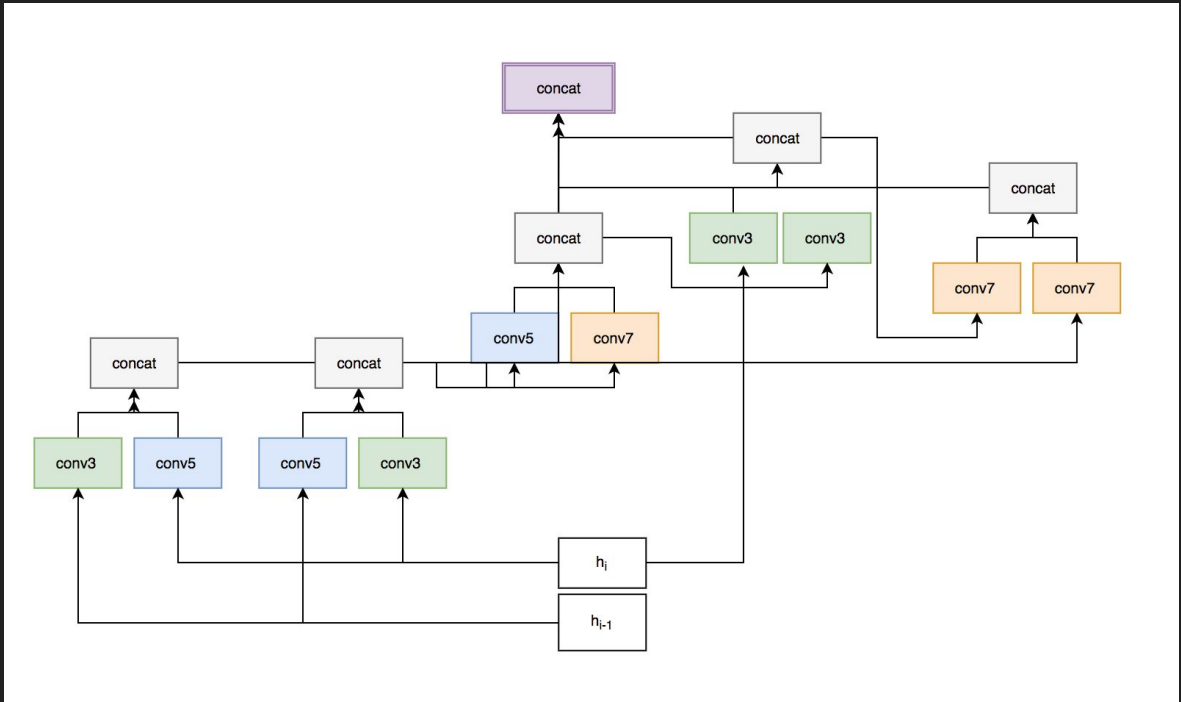

Val Accuracy	Test Accuracy
41.41	42.97

Related work on architecture search

1. Neural architecture search with reinforcement learning. ICLR, 2017.
2. Learning transferable architectures for scalable image recognition. arXiv 2017.
3. Progressive Neural Architecture Search. arXiv 2017