Recent Advancements!

Lecture # 10

Hassan Sajjad and Fahim Dalvi
Qatar Computing Research Institute, HBKU

Today

- Domain adaptation
- Multilingual systems - Zero shot
- Multi-task learning
- Multi-modal learning
- Reinforcement learning
- Adversarial training

Recent Advancements

Domain Adaptation

Domain Adaptation

- Training data comes in various styles, genre and domains
  - news, TED talks, Europarl
- Translation systems built on one domain may not perform well on other domains

Domain adaptation aims to use all the available data for the benefit of the in-domain data

Domain Adaptation

- Suppose, we want to train an NMT model for TED talks and there are three corpora available for training

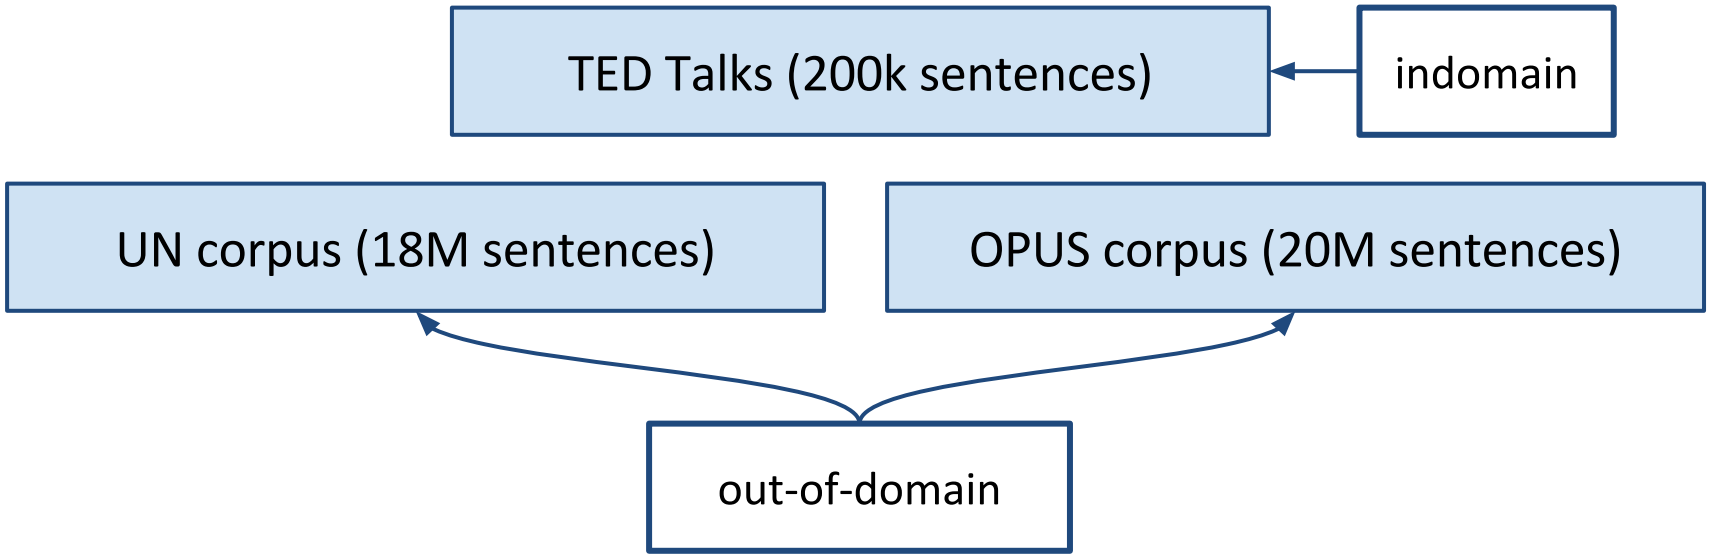

Domain Adaptation

- Suppose, we want to train an NMT model for TED talks and there are three corpora available for training

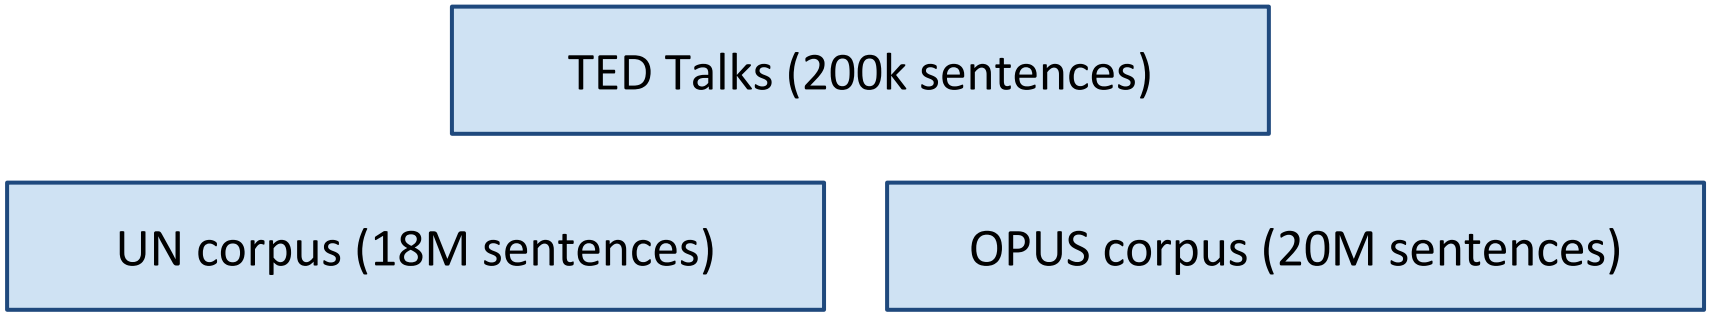

- What would be a good training strategy in order to built a good domain specific system?

Domain Adaptation

- What would be a good training strategy in order to built a good TED talks system?
- Use only TED talks because that's what we care about!

TED Talks (200k sentences)

UN corpus (18M sentences)

OPUS corpus (20M sentences)

Domain Adaptation

- What would be a good training strategy in order to built a good TED talks system?
- Use only TED talks because that's what we care about!

TED data is small. It would result in overfitting

TED Talks (200k sentences)

UN corpus (18M sentences)

OPUS corpus (20M sentences)

Domain Adaptation

- What would be a good training strategy in order to built a good TED talks system?
- Combine all three corpora together!

TED Talks (200k sentences)

UN corpus (18M sentences)

OPUS corpus (20M sentences)

Domain Adaptation

- What would be a good training strategy in order to built a good TED talks system?
- Combine all three corpora together!

TED is comparatively small. It will be lost in the concatenation of three corpora

TED Talks (200k sentences)

UN corpus (18M sentences)

OPUS corpus (20M sentences)

Domain Adaptation

- What would be a good training strategy in order to built a good TED talks system?

- **Solution:** Train the model first on the large out-of-domain corpora to learn a generic model
- Later, **fine-tune** it on the in-domain data

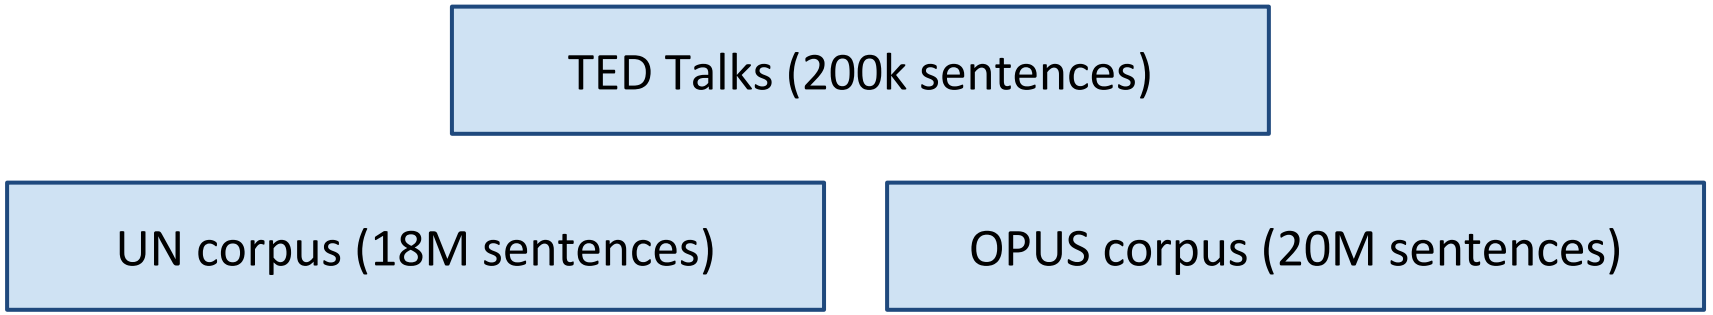

Domain Adaptation - Fine-tuning

Consider an NMT training scenario that takes an input corpus and trains a model

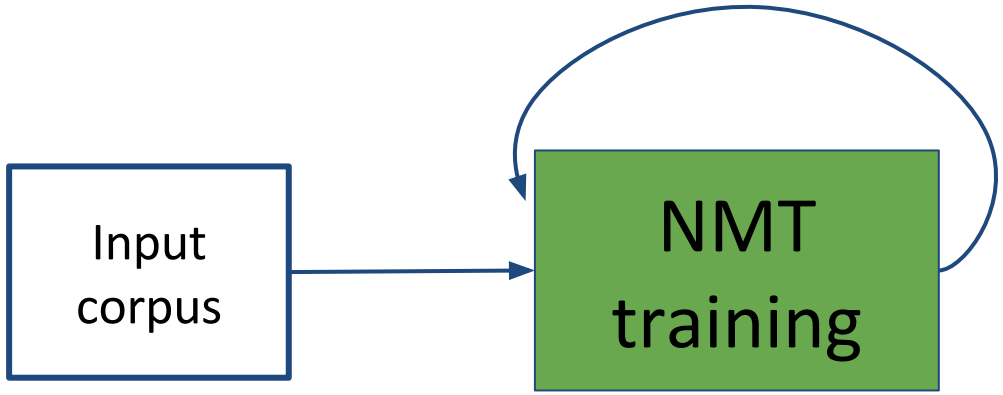

Domain Adaptation - Fine-tuning

First we train the model on the out-of-domain data for 20 epochs

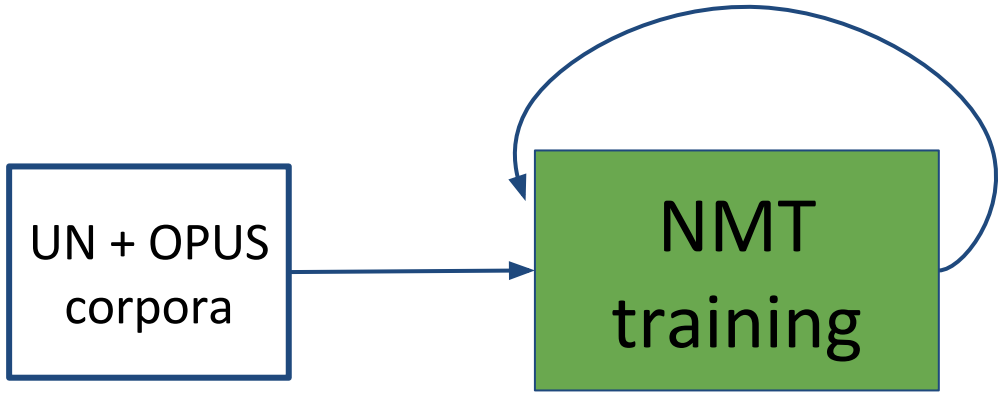

Domain Adaptation - Fine-tuning

Now model has learned to translate language from a generic data

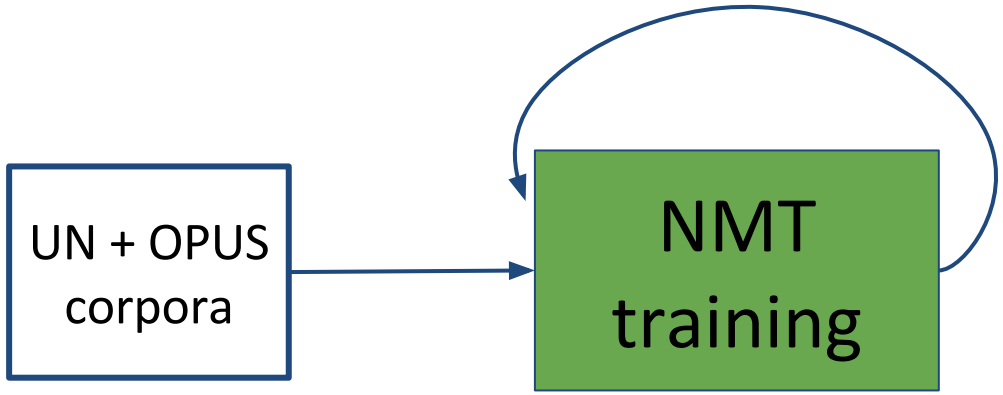

Domain Adaptation - Fine-tuning

Let’s fine-tune the model to our domain of interest which is TED talks!

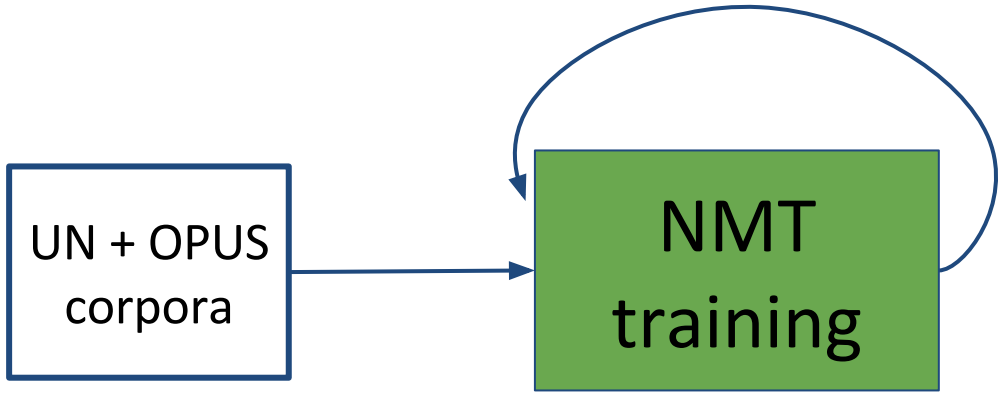

Domain Adaptation - Fine-tuning

Remove out-of-domain corpora and input TED talks to the already trained model

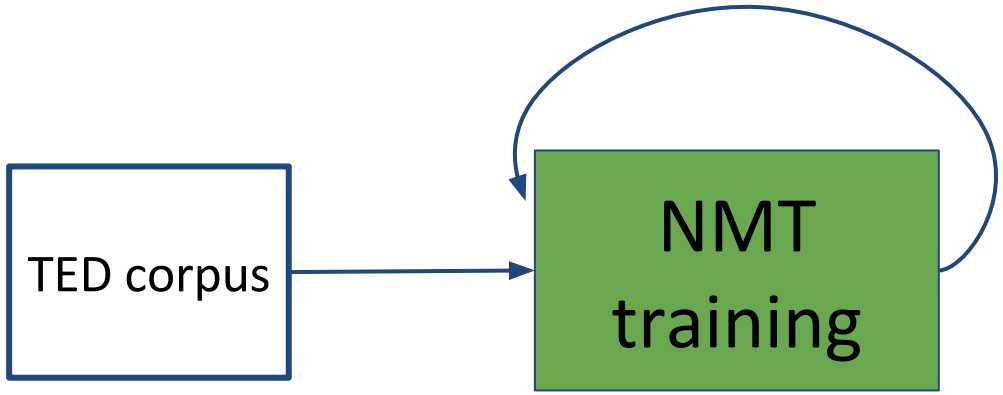

Domain Adaptation - Fine-tuning

Remember, NMT runs in mini-batches. Essentially we are simply training an NMT system on a few more mini-batches but of TED corpus

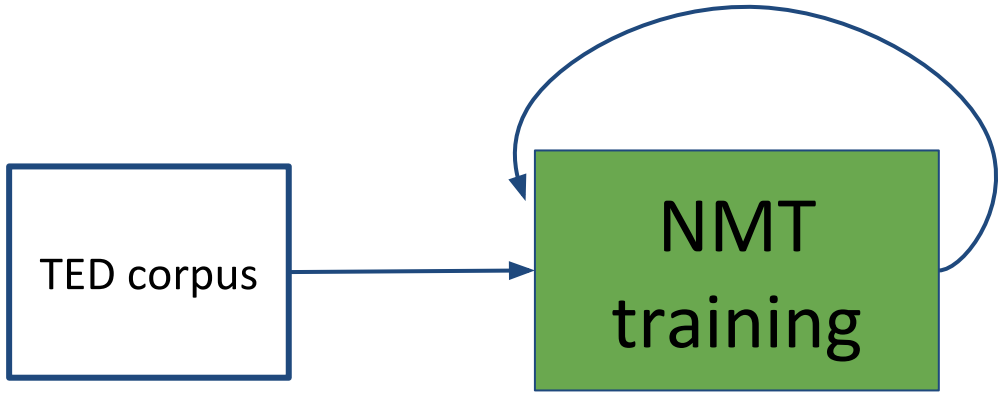

Domain Adaptation - Fine-tuning

System learned initial parameters from a large out-of-domain corpora and later we fine-tuned the trained parameters in favor of our in-domain corpus

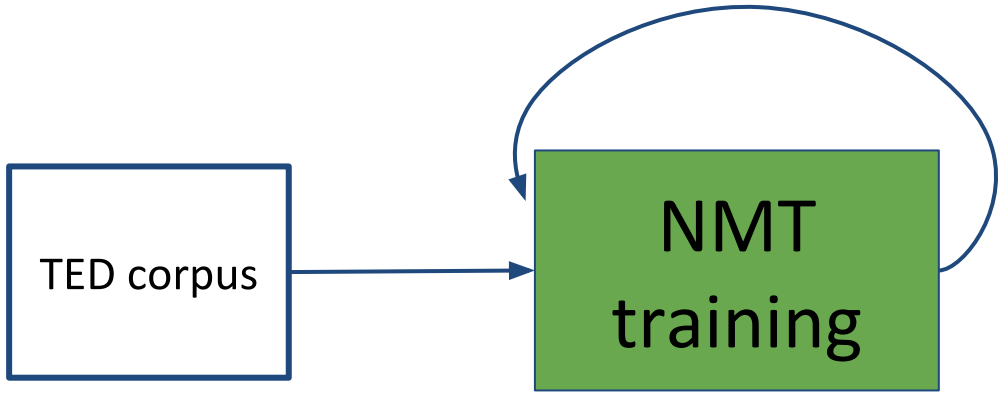

Domain Adaptation - Fine-tuning

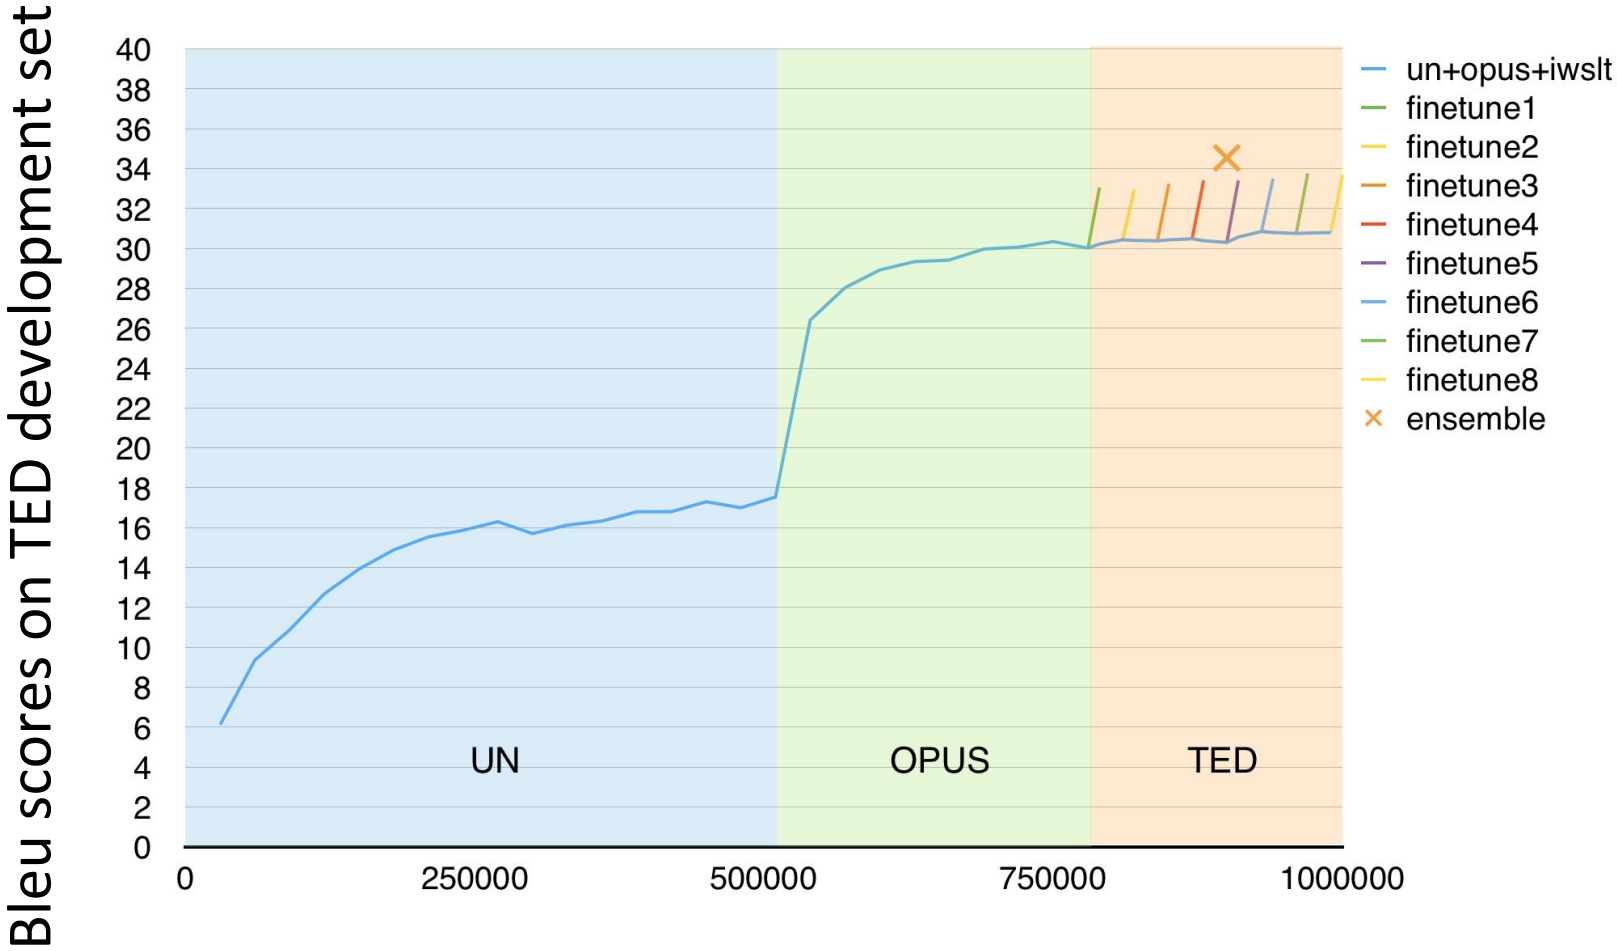

https://arxiv.org/pdf/1708.08712.pdf

Domain Adaptation - Fine-tuning

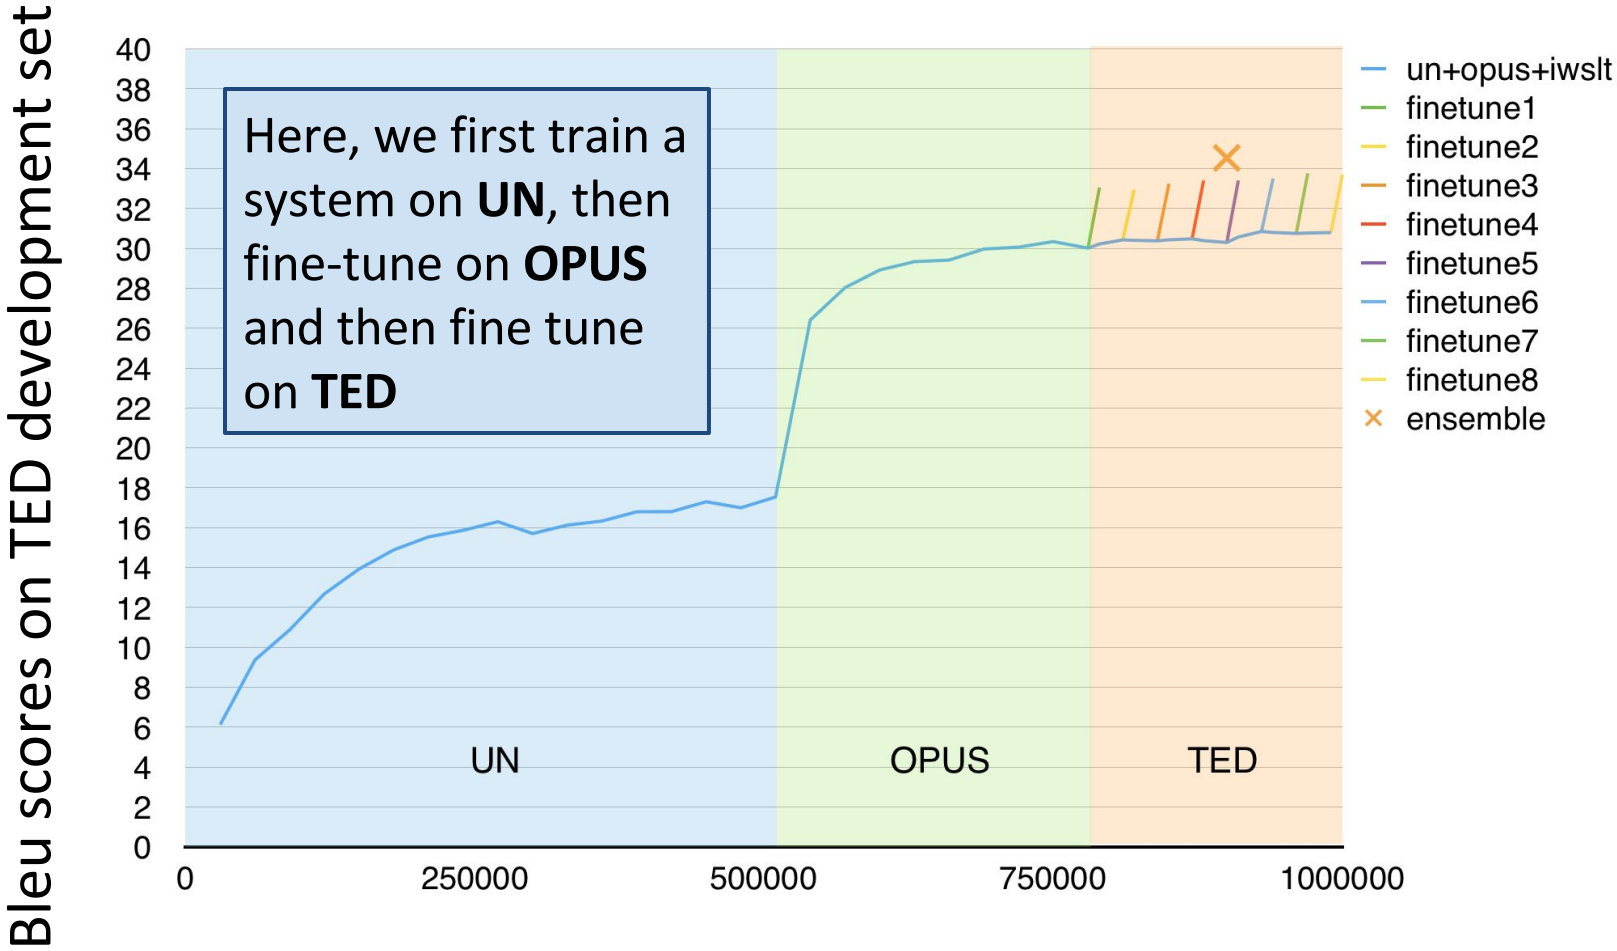

Domain Adaptation - Fine-tuning

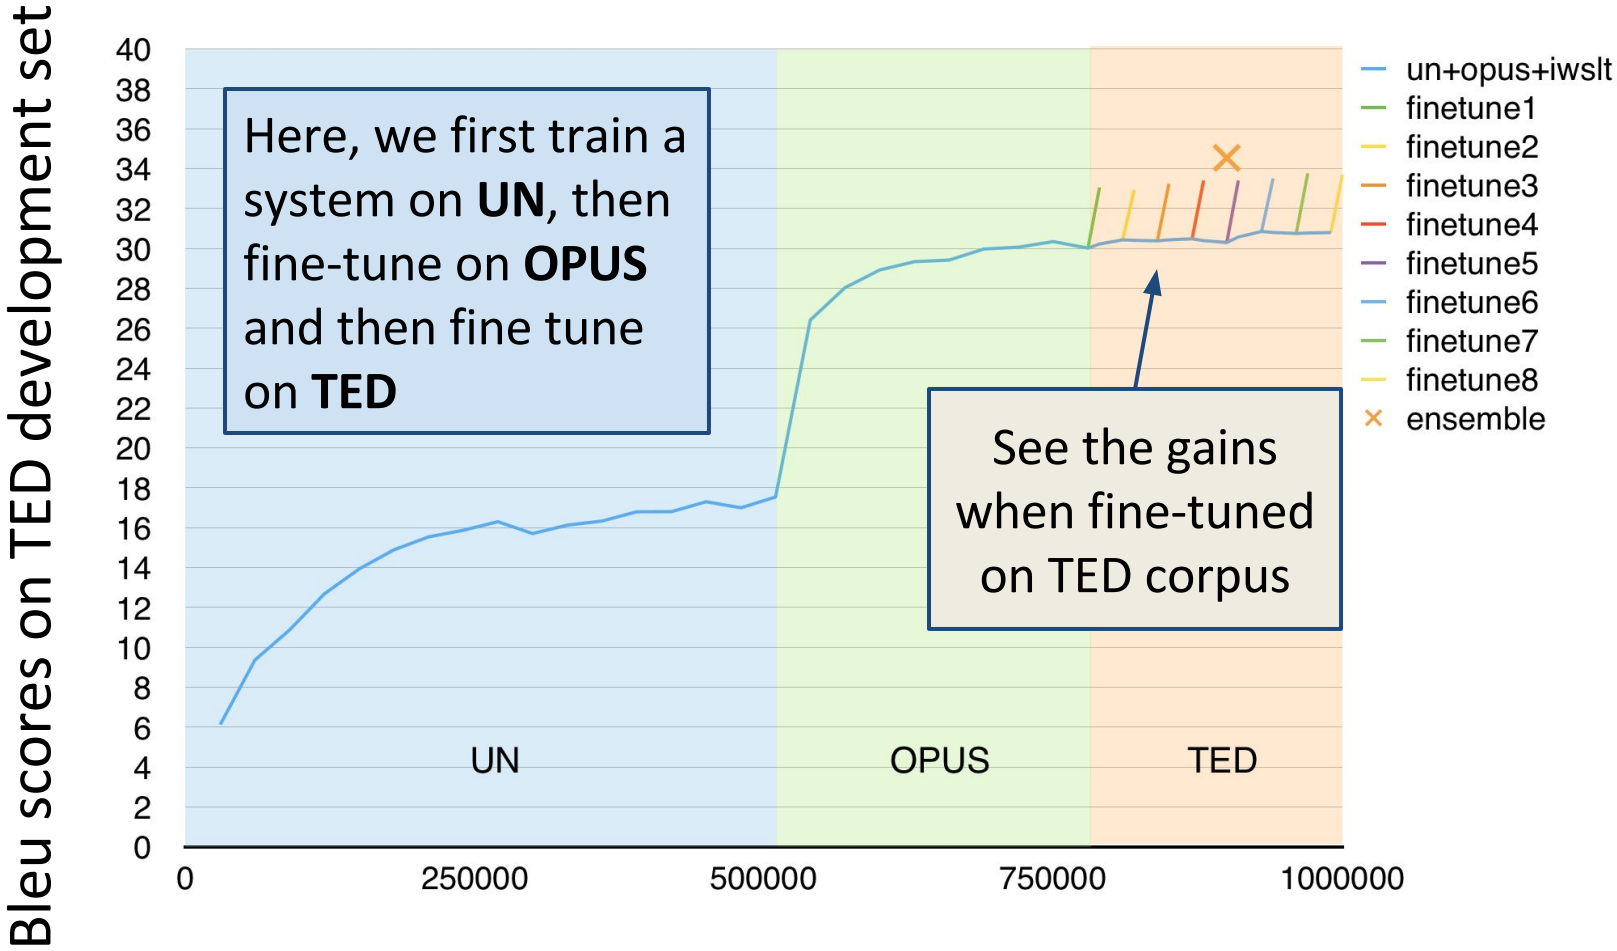

Domain Adaptation - Fine-tuning

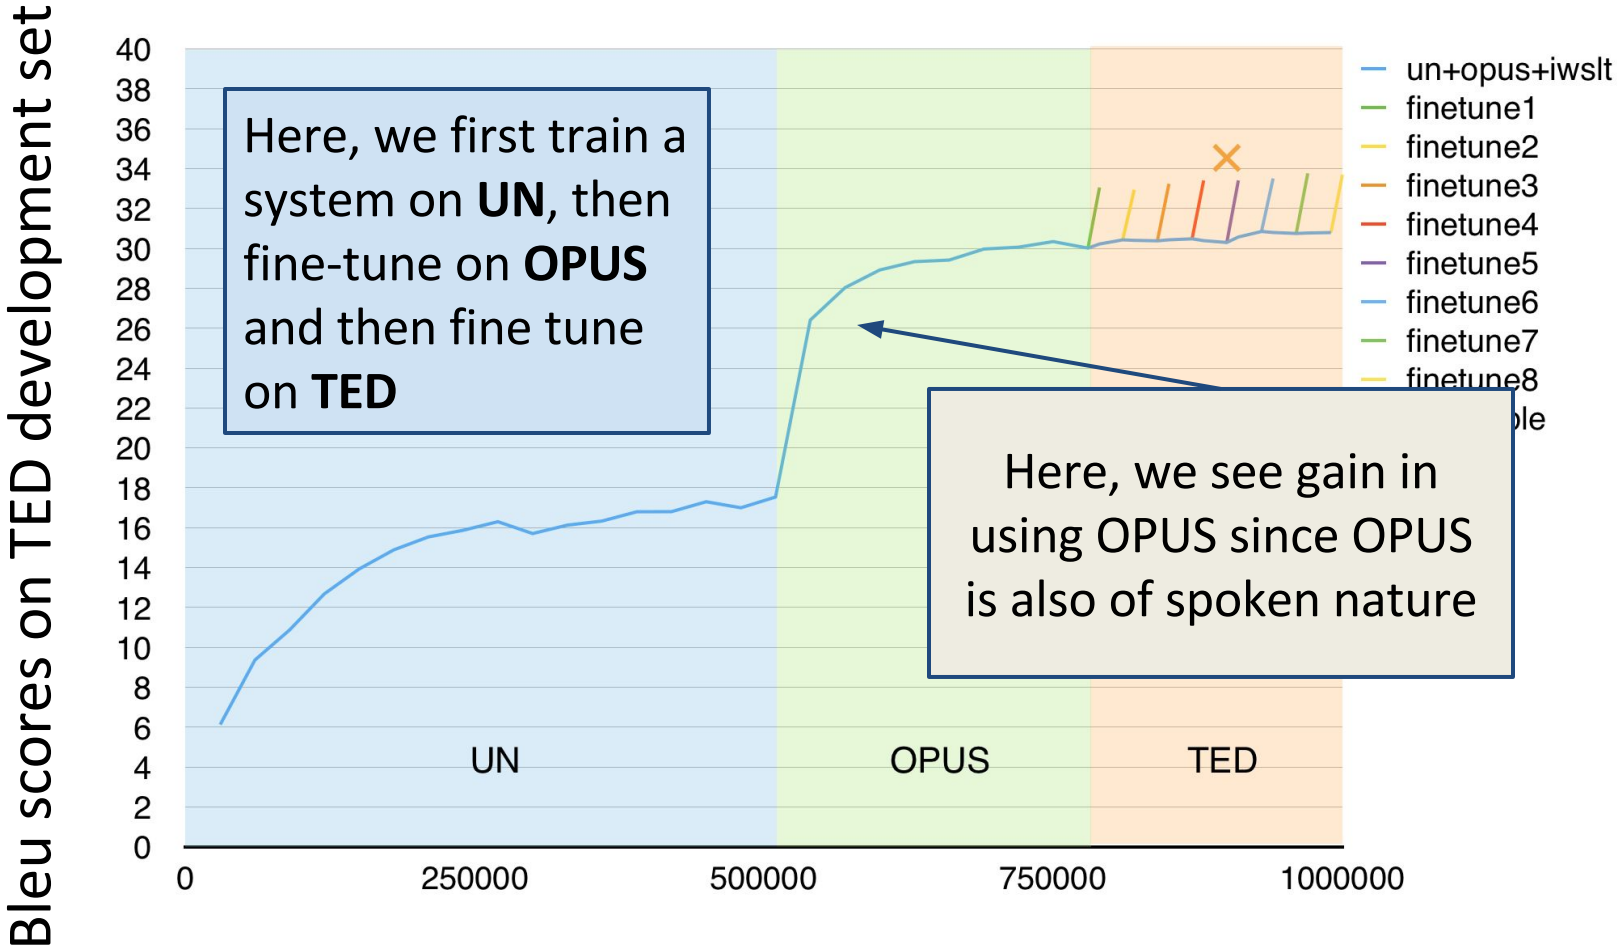

Domain Adaptation

- Fine-tuning is an effective and commonly used method to bias model parameters towards a specific scenario

- However, domain adaptation is a fairly unexplored research area in neural machine translation
  - a lot of roam for better methods for various scenarios

Recent Advancements

Multilingual Translation Systems

Multilingual Translation Systems

We briefly saw yesterday that the encoder and the decoder are not inherently limited to a single language!

Multilingual Translation Systems

Idea: Instead of training individual systems that translate between German, Arabic, English etc. train just **one** system!

Multilingual Translation Systems

Q: How do we do this practically?

Hello, how are you?	Hallo wie geht's dir?
Where is the nearest restaurant?	Wo ist das nächste Restaurant?
It is a good day today.	Es ist ein guter Tag heute.
⋮	⋮
Thank you!	Vielen Dank!

English → German system

English → Arabic system

it was a pleasure to meet you.	كان لقائك من دواعي سروري
The world is not flat	العالم ليست مسطحة
Good morning	صباح الخير.
⋮	⋮
Thank you!	!شكرا

Keine problem	لا مشكلة
Einmal wasser	مرة واحدة الماء
Ich habe meine auto	لدي سيارتي
⋮	⋮
Bist du verrückt?	هل انت مجنون؟

German → Arabic system

Multilingual Translation Systems

Q: How do we do this practically?

<de> Hello, how are you? <de> Where is the nearest restauran <de> It is a good day today. ⋮ <de> Thank you!	Hallo wie geht's dir? Wo ist das nächste Restaurant? Es ist ein guter Tag heute. ⋮ Vielen Dank!
<ar> it was a pleasure to meet you. <ar> The world is not flat <ar> Good morning ⋮ <ar> Thank you!	كان لقائك من دواعي سروري العالم ليست مسطحة .صباح الخير ⋮ !شكرا
<ar> Keine problem <ar> Einmal wasser <ar> Ich habe meine auto ⋮ <ar> Bist du verrückt?	لا مشكلة مرة واحدة الماء لدي سيارتي ⋮ هل انت مجنون؟

Single English, Arabic → Arabic, German system!

Multilingual Translation Systems

Q: How do we do this practically?

<de> Hello, how are you? <de> Where is the nearest restauran <de> It is a good day today. ⋮ <de> Thank you!	Hallo wie geht's dir? Wo ist das nächste Restaurant? Es ist ein guter Tag heute. ⋮ Vielen Dank!
<ar> it was a pleasure to meet you. <ar> The world is not flat <ar> Good morning ⋮ <ar> Thank you!	كان لقائك من دواعي سروري العالم ليست مسطحة .صباح الخير ⋮ !شكرا
<ar> Keine problem <ar> Einmal wasser <ar> Ich habe meine auto ⋮ <ar> Bist du verrückt?	لا مشكلة مرة واحدة الماء لدي سيارتي ⋮ هل انت مجنون؟

Single English, Arabic → Arabic, German system!

The NMT system learns that the first token in the source sentence is what defines the target language!

Multilingual Translation Systems

Q: Why does this work?

A: Because we are forcing the network to learn so many languages, the **path of least resistance** for the network is to learn **an intermediate language**

Multilingual Translation Systems

Q: Why does this work?

A: Because we are forcing the network to learn so many languages, the **path of least resistance** for the network is to learn **an intermediate language**

The network encodes the source language to this intermediate language, and then decodes from the intermediate language into the target language

Multilingual Translation Systems

Slight tangent...

Forbes: Facebook AI Creates Its Own Language In Creepy Preview Of Our Potential Future

The Atlantic: An Artificial Intelligence Developed Its Own Non-Human Language

Independent: Facebook's artificial intelligence robots shut down after they start talking to each other in their own language

Multilingual Translation Systems

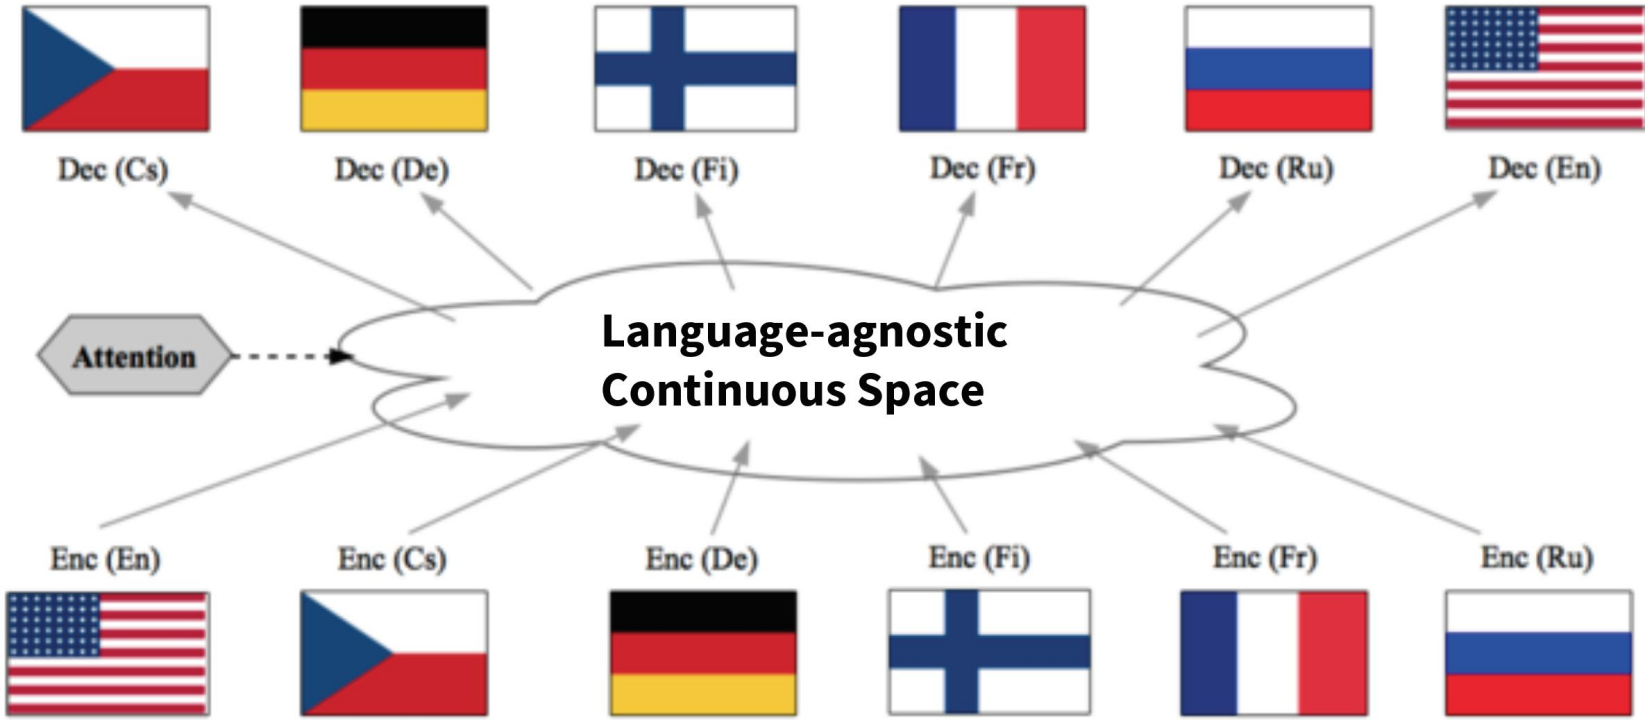

Dong at al., 2015, Luong at al. ICLR 2016, Firat et al. NAACL 2016

Zero-Shot Translation

Q: Can an NMT system learn to produce translations between languages that it has not seen any explicit data for?

Zero-Shot Translation

Q: Can an NMT system learn to produce translations between languages that it has not seen any explicit data for?

A: Yes of course! The only thing the network needs to learn is to convert the source language to the intermediate language, and convert the intermediate language to the target language!

Zero-Shot Translation

Consider the example from earlier, but with only two language pairs:

<de> Hello, how are you? <de> Where is the nearest restauran <de> It is a good day today. ⋮ <de> Thank you!	Hallo wie geht's dir? Wo ist das nächste Restaurant? Es ist ein guter Tag heute. ⋮ Vielen Dank!
<ar> Keine problem <ar> Einmal wasser <ar> Ich habe meine auto ⋮ <ar> Bist du verrückt?	لا مشكلة مرة واحدة الماء لدي سيارتي ⋮ هل انت مجنون؟

Zero-Shot Translation

<de> Hello, how are you? <de> Where is the nearest restauran <de> It is a good day today. ⋮ <de> Thank you!	Hallo wie geht's dir? Wo ist das nächste Restaurant? Es ist ein guter Tag heute. ⋮ Vielen Dank!
<ar> Keine problem <ar> Einmal wasser <ar> Ich habe meine auto ⋮ <ar> Bist du verrückt?	لا مشكلة مرة واحدة الماء لدي سيارتي ⋮ هل انت مجنون؟

If our system works well, it has learned four language pairs:

English ⟶ Intermediate

Intermediate ⟶ German

German ⟶ Intermediate

Intermediate ⟶ Arabic

Zero-Shot Translation

<de> Hello, how are you? <de> Where is the nearest restauran <de> It is a good day today. ⋮ <de> Thank you!	Hallo wie geht's dir? Wo ist das nächste Restaurant? Es ist ein guter Tag heute. ⋮ Vielen Dank!
<ar> Keine problem <ar> Einmal wasser <ar> Ich habe meine auto ⋮ <ar> Bist du verrückt?	لا مشكلة مرة واحدة الماء لدي سيارتي ⋮ هل انت مجنون؟

Now, if we want to translate from English → Arabic, we can!

English → Intermediate

Intermediate → German

German → Intermediate

Intermediate → Arabic

Zero-Shot Translation Systems

- These systems are very useful in scenarios where we have a limited amount of parallel text for some language pairs
  - Especially true for resource poor languages

Recent Advancements

Multi-task learning

Multi-task Learning

- Learn multiple tasks together
- For example, learn to translate and produce POS tags together
- Implicitly introducing an inductive bias
  - model has to consider more than one tasks to reduce loss
- Benefits
  - learning various tasks together helps to achieve better generalization

Multi-task Learning

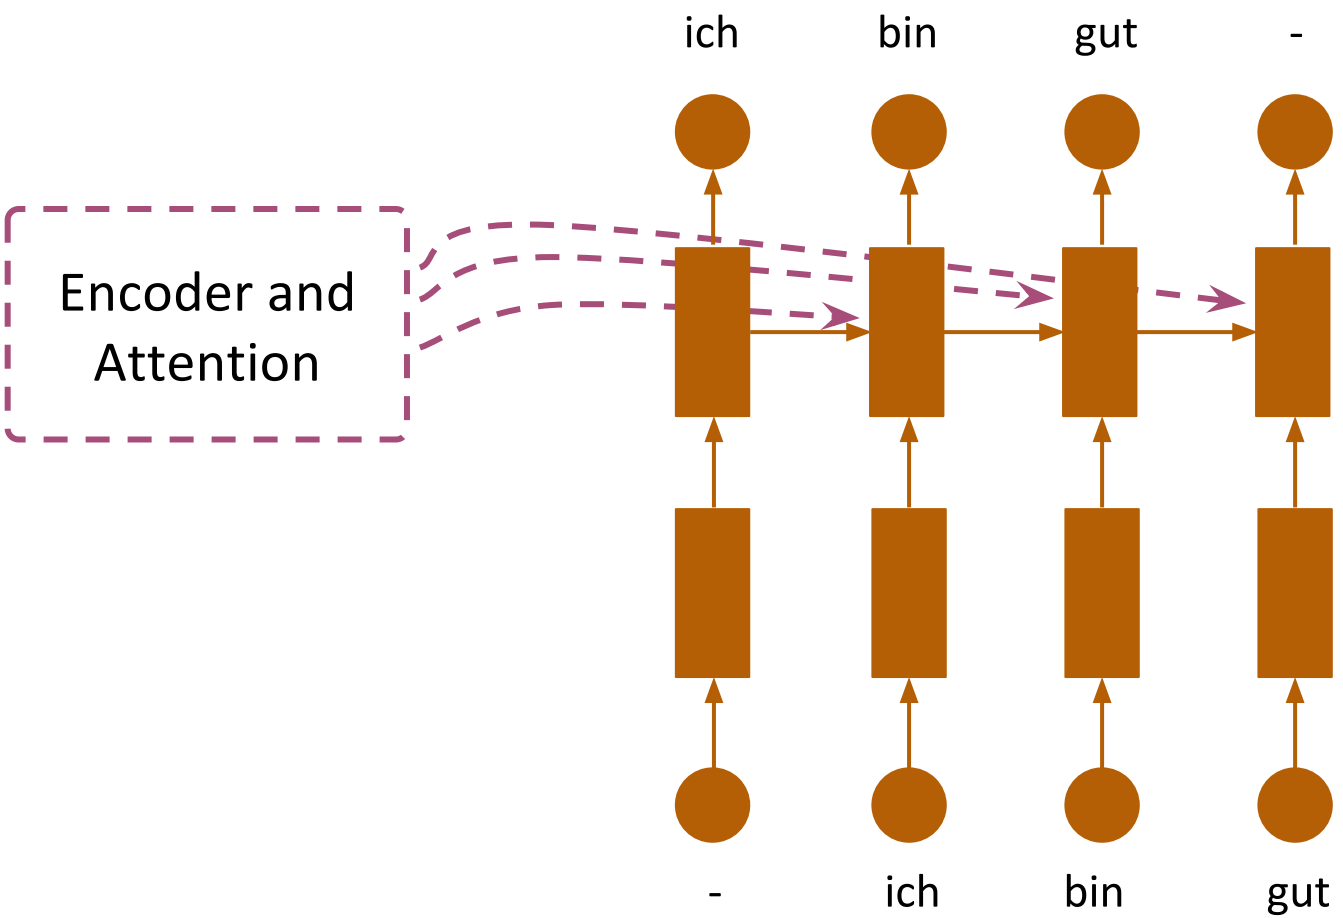

Multi-task Learning

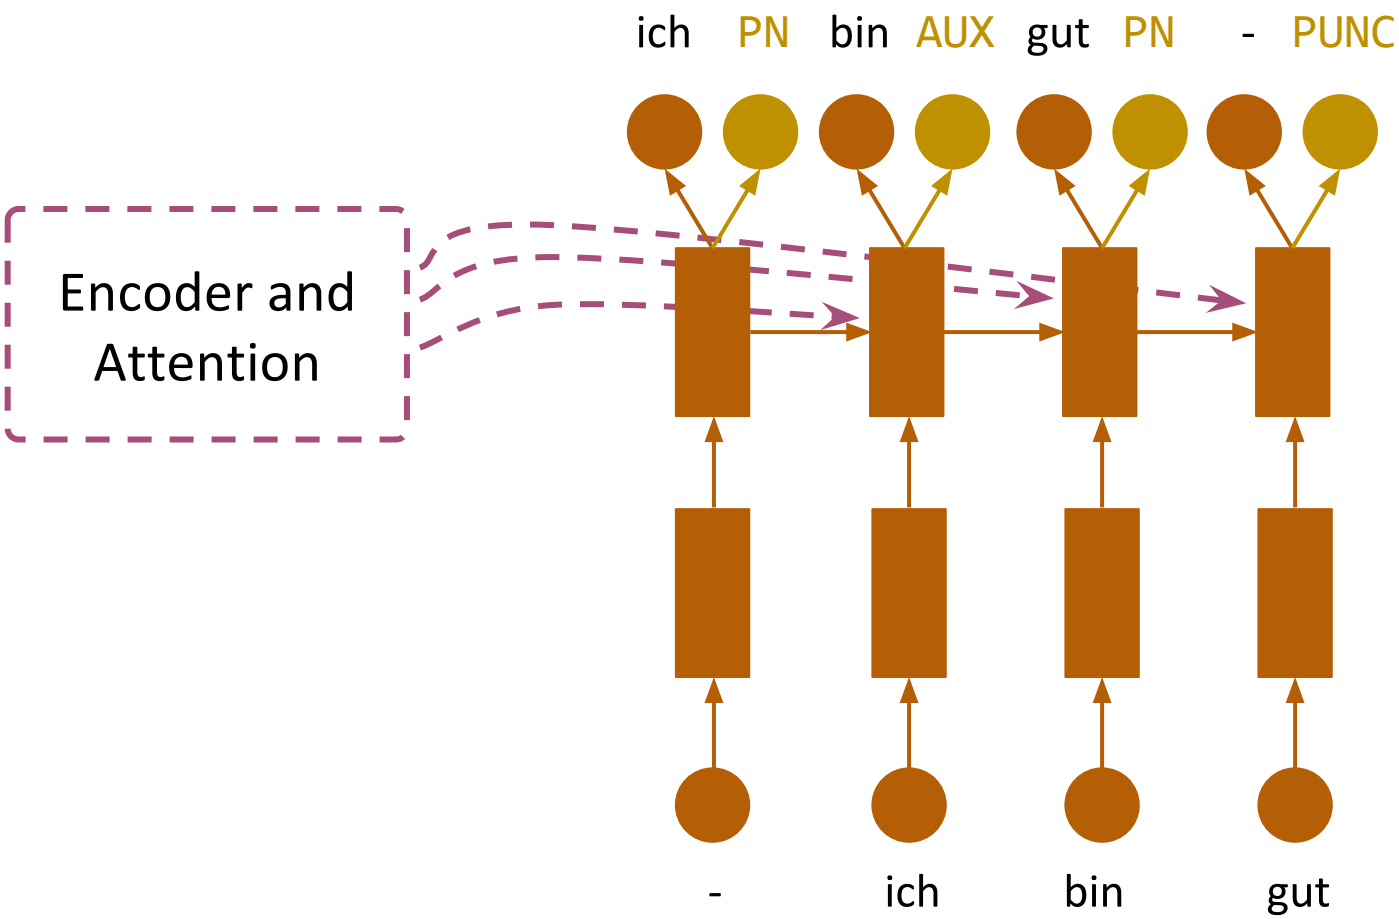

Multi-task Learning

- POS tagging and translation tasks are learned together
- Input
  - source sentence - target sentence
  - source sentence - target POS-tagged sequence
- Shared embeddings, shared attention
- Loss function minimizes the combined loss of both tasks

Multi-task Learning

- Since we are learning weights of multiple tasks simultaneously, chances of overfitting are reduced
- As we increase the number of tasks, model has to find weights that minimize the overall loss
- Thus, the resulting model generally has better generalization capability

Multi-task Learning

- Multi-task learning can also be seen as a way to inject linguistic knowledge into the network!
- One idea is to inject target morphology in a multitask setting
  - improves translation quality in translating into morphologically rich languages, such as Arabic

Recent Advancements

Multi-modal learning

Multi-modal Learning

Humans tend to learn using various senses e.g. babies use vision and speech together to learn communication

Multi-modal Learning

Humans tend to learn using various senses e.g. babies use vision and speech together to learn communication

We have seen the we can use the same encoder or decoder with more than one language

Multi-modal Learning

Q: As there was no inherent limitation of a single language for the encoder/decoder, is there any inherent limitation for the source/target to be a language at all?

Multi-modal Learning

Q: As there was no inherent limitation of a single language for the encoder/decoder, is there any inherent limitation for the source/target to be a language at all?

A: No! As long as we can teach the machine to convert from any arbitrary input to its intermediate language, we should be able to work with that input!

Multi-modal Learning

Image

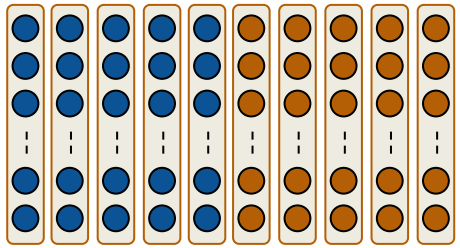

Encoder-Decoder Model

A grumpy cat

Caption

Multi-modal Learning

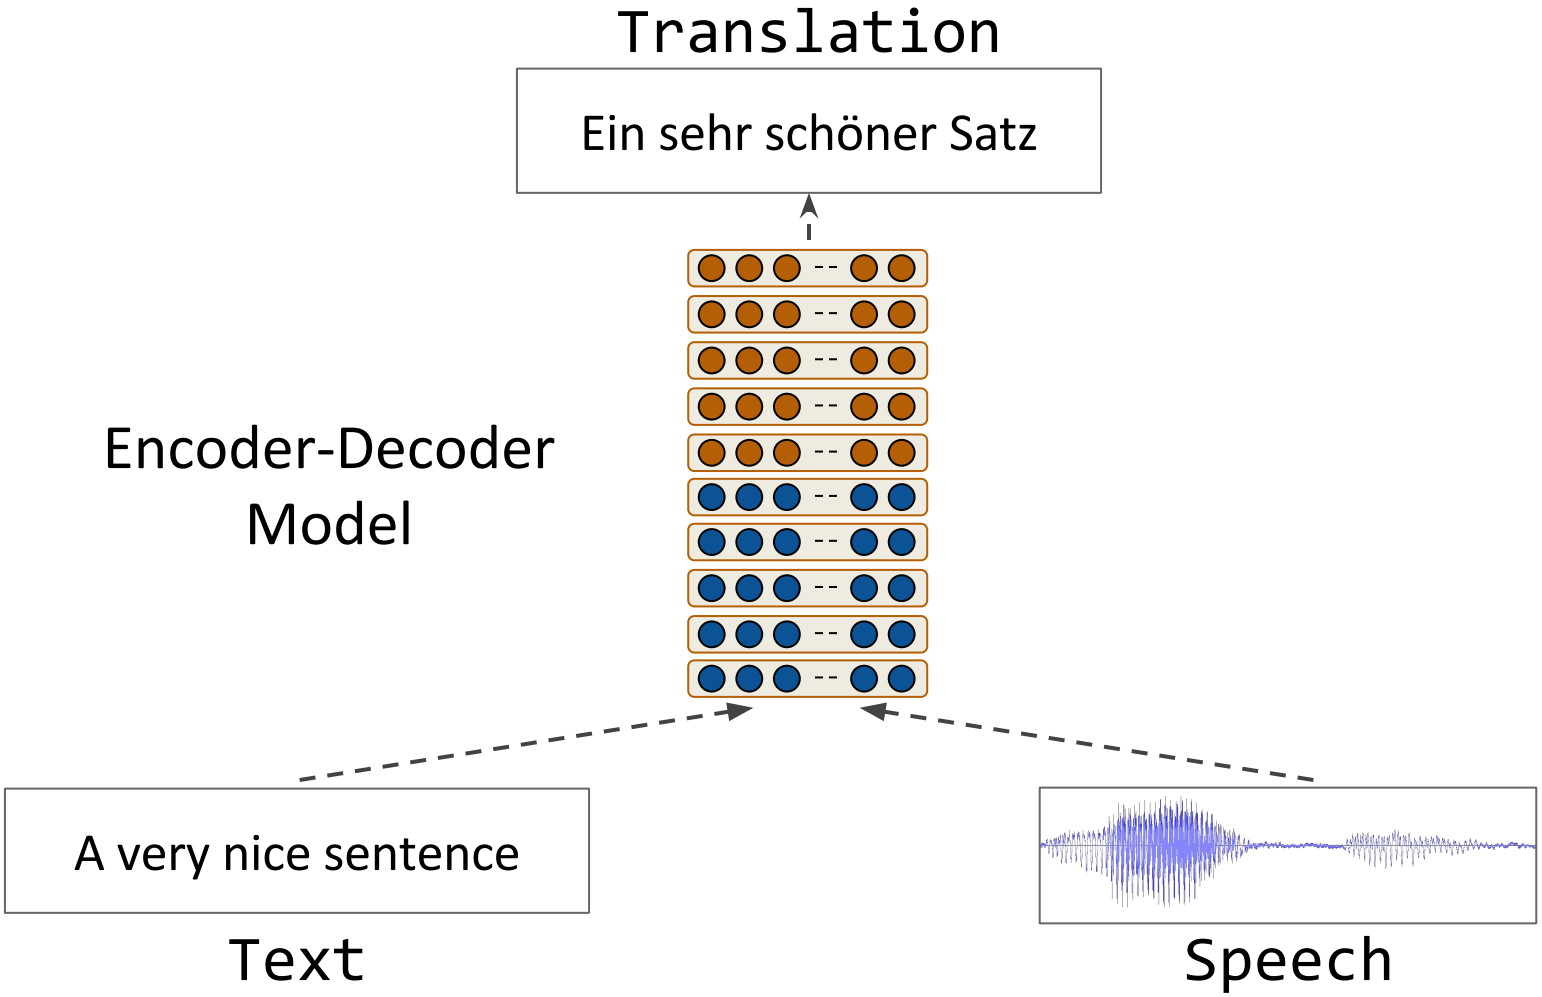

Multi-modal Learning

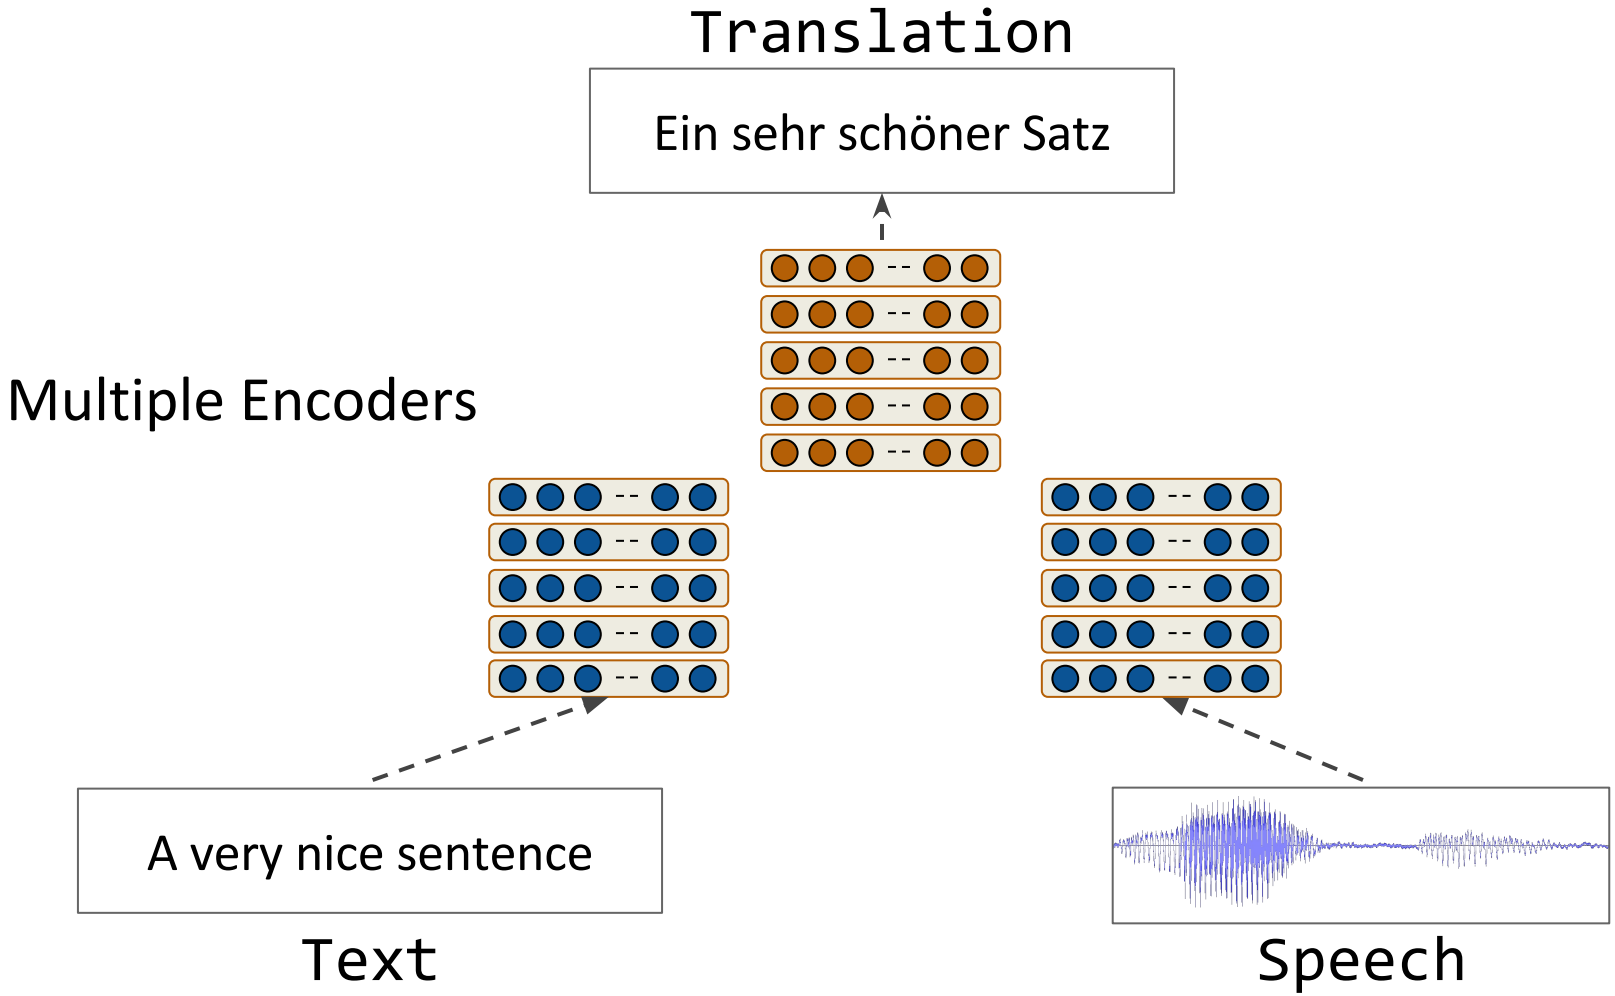

Multimodal Multitask Learning

Each encoder is usually designed for a specific type of input (text, speech, images)

The challenging part is to bring all senses to one space. Jointly training all these tasks with combined losses is one way to do this - still an open problem!

Recent Advancements

Generative Adversarial Networks

Generative Adversarial Networks

So far, we have seen only one class of machine learning problems: Classification problems

Generative Adversarial Networks

So far, we have seen only one class of machine learning problems: Classification problems

Another class of problems is generating outputs - say generating speech, drawing pictures, producing “handwritten” texts...

Generative Adversarial Networks

We can do this using **Generative Adversarial Networks**!

Humans have evolved over time, and a fundamental motivator for a lot of the progress has been *competition*

Generative Adversarial Networks

We can do this using **Generative Adversarial Networks**!

Humans have evolved over time, and a fundamental motivator for a lot of the progress has been *competition*

Idea: If humans (generally) get better because of *competition*, why not make machines compete with each other as well!

Generative Adversarial Networks

Intuition: We will have two neural networks - one will try and generate something, while the other will try to distinguish if its input is from the real world, or is generated by the first neural network!

Generative Adversarial Networks

Intuition: We will have two neural networks - one will try and generate something, while the other will try to distinguish if its input is from the real world, or is generated by the first neural network!

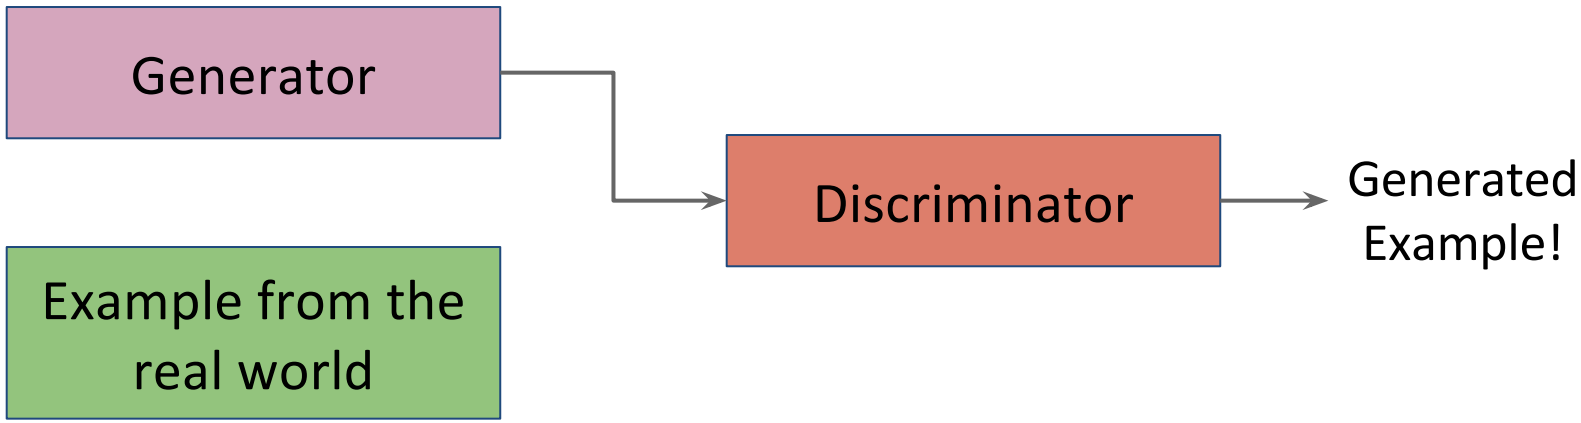

Generative Adversarial Networks

Intuition: We will have two neural networks - one will try and generate something, while the other will try to distinguish if its input is from the real world, or is generated by the first neural network!

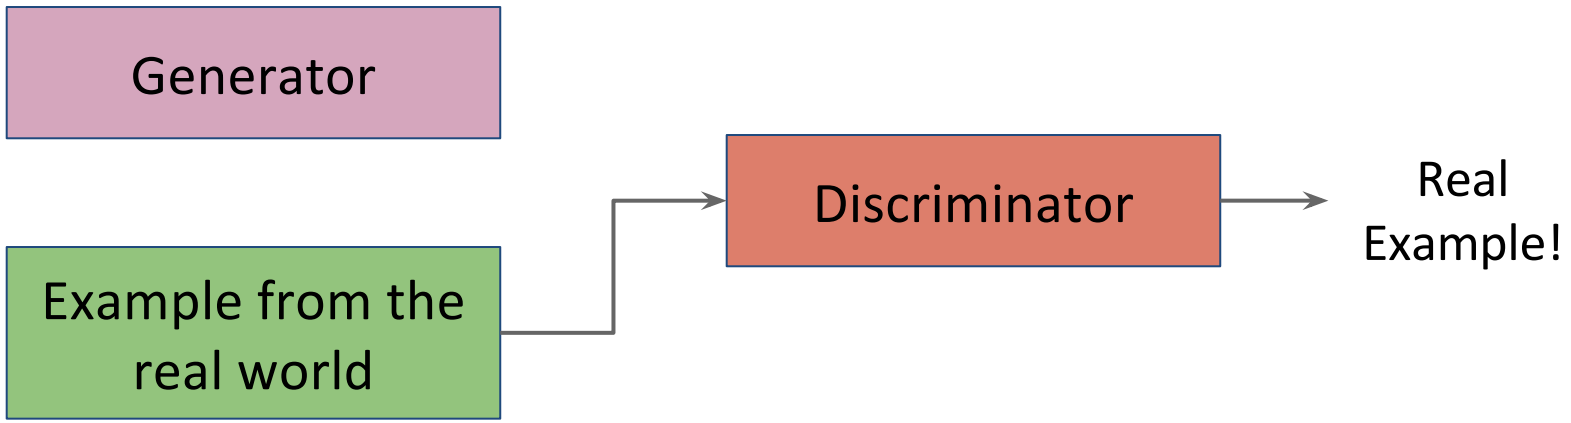

Generative Adversarial Networks

Training process:

1. Freeze Generator and draw outputs from it
2. Draw equal number of outputs from the real world
3. Train the Discriminator on the mix of this data

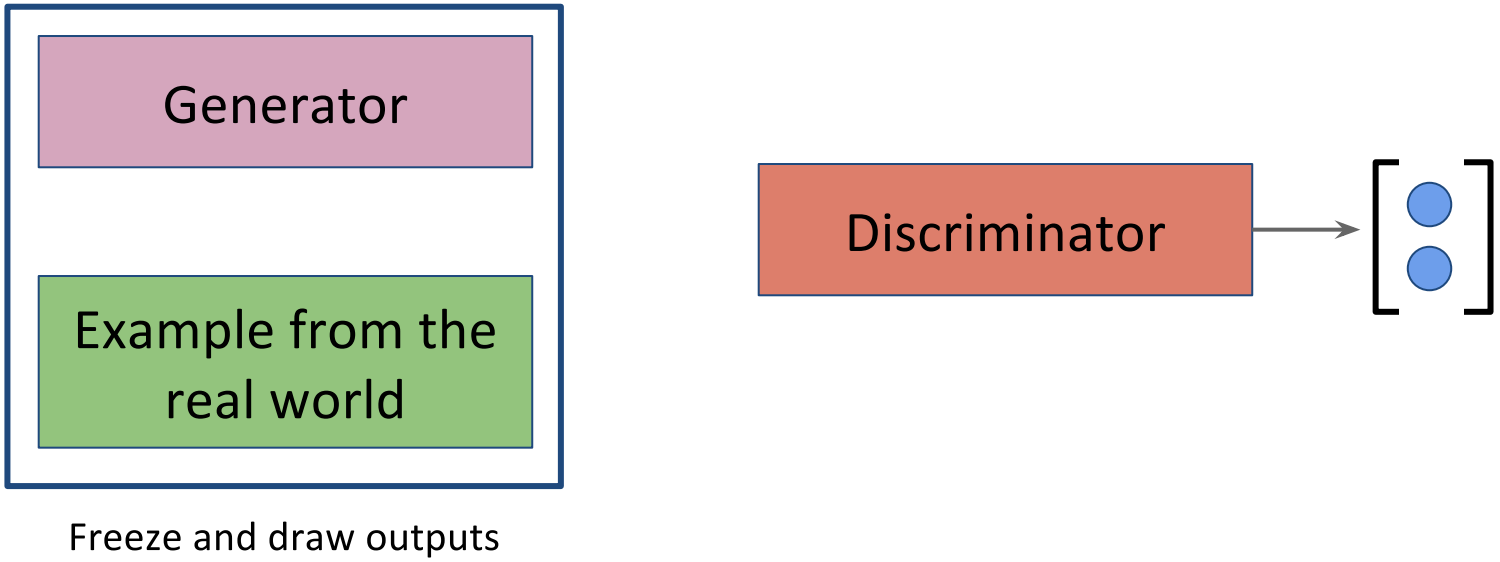

Generative Adversarial Networks

Training process:

1. Freeze Generator and draw outputs from it
2. Draw equal number of outputs from the real world
3. Train the Discriminator on the mix of this data

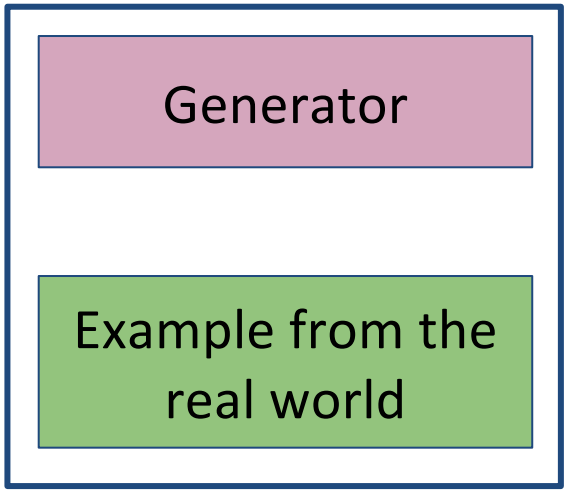

Freeze and draw outputs

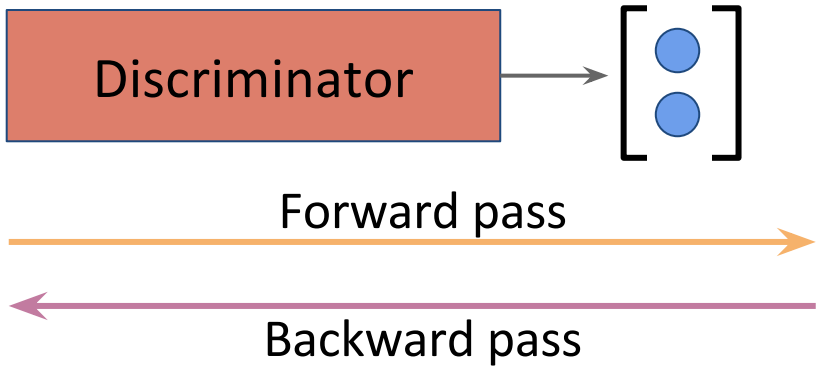

Generative Adversarial Networks

Training process:

4. Freeze Discriminator
5. Draw an output from the generator
6. Pass the output through the frozen discriminator, and backpropagate the loss!

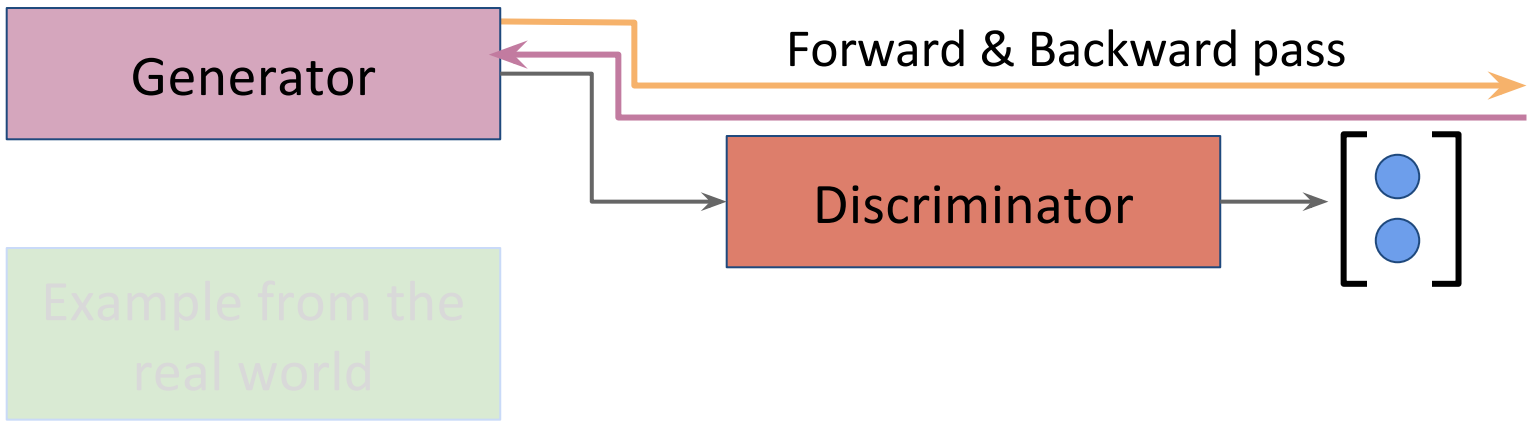

Generative Adversarial Networks

Training process:

7. Rinse and Repeat steps 1-6
8. Use the generator to generate real-world like examples!

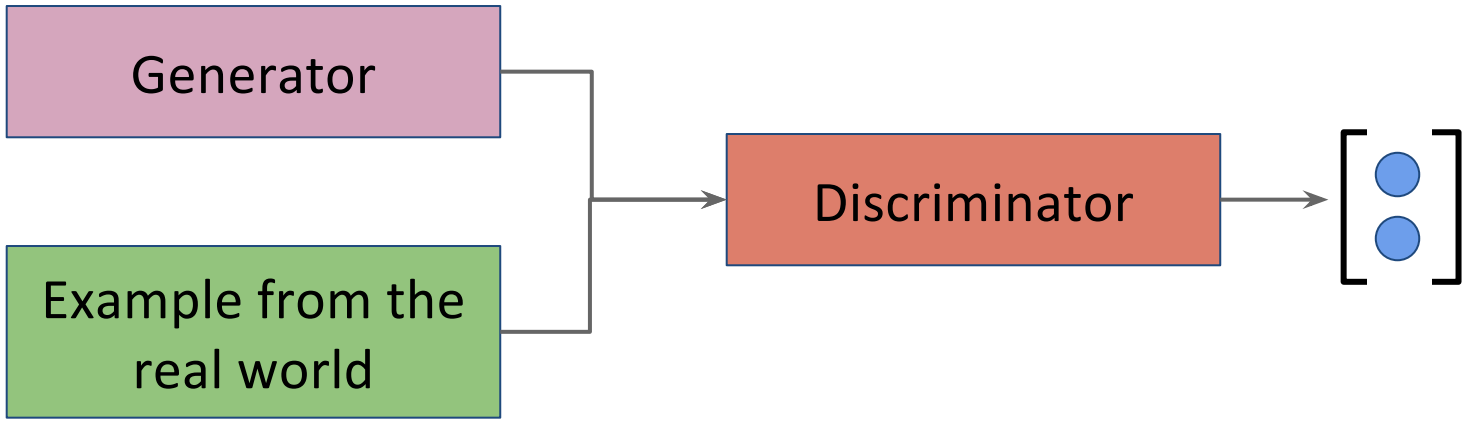

Generative Adversarial Networks

https://www.reddit.com/r/MachineLearning/comments/49wrt4/adversarial_images_for_deep_learning/

Generative Adversarial Networks

Learning what bedrooms look like:

"Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks"
https://arxiv.org/abs/1511.06434v2

Generative Adversarial Networks

Prisma App: Art from pictures!

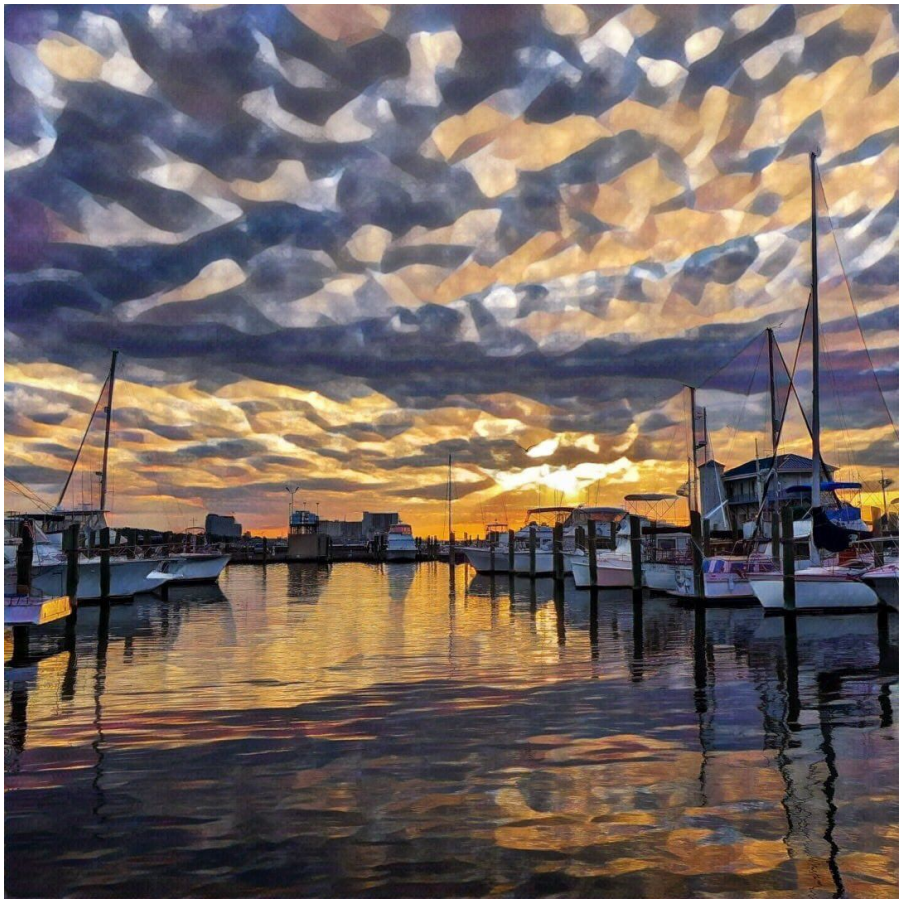

https://twitter.com/lapdog99/status/760639333812490240

Generative Adversarial Networks

NIPS 2016 Tutorial: Generative Adversarial Networks
https://arxiv.org/pdf/1701.00160.pdf

Captions to text!

Generative Adversarial Networks

Ofcourse, we also have issues...

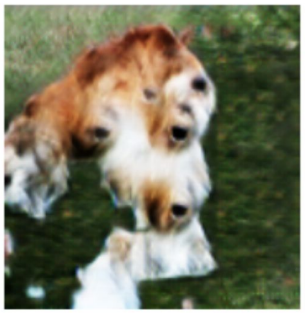




NIPS 2016 Tutorial: Generative Adversarial Networks
https://arxiv.org/pdf/1701.00160.pdf

Recent Advancements

Reinforcement Learning

Reinforcement Learning

Another class of problems is teaching a machine to perform some **sequence of actions** to **reach an eventual goal**!

Reinforcement Learning

Another class of problems is teaching a machine to perform some **sequence of actions** to **reach an eventual goal**!

Imagine a child's mind when it's learning to walk:

1. The child would notice how adults around it walk
2. It will try to first stand and balance itself - and will fall repeatedly before finding the right "parameters"
3. It will then learn to take small steps. Again, tuning its "parameters" to avoid falling

Reinforcement Learning

This is kind of what reinforcement learning algorithms do!

The task at hand is to train an “agent” who will perform some “actions” and eventually either succeed or fail.

Reinforcement Learning

This is kind of what reinforcement learning algorithms do!

The task at hand is to train an “agent” who will perform some “actions” and eventually either succeed or fail.

Researchers have come up with nice algorithms so that the “agent” can learn from its failures - and eventually succeed

Reinforcement Learning

But first, an example!

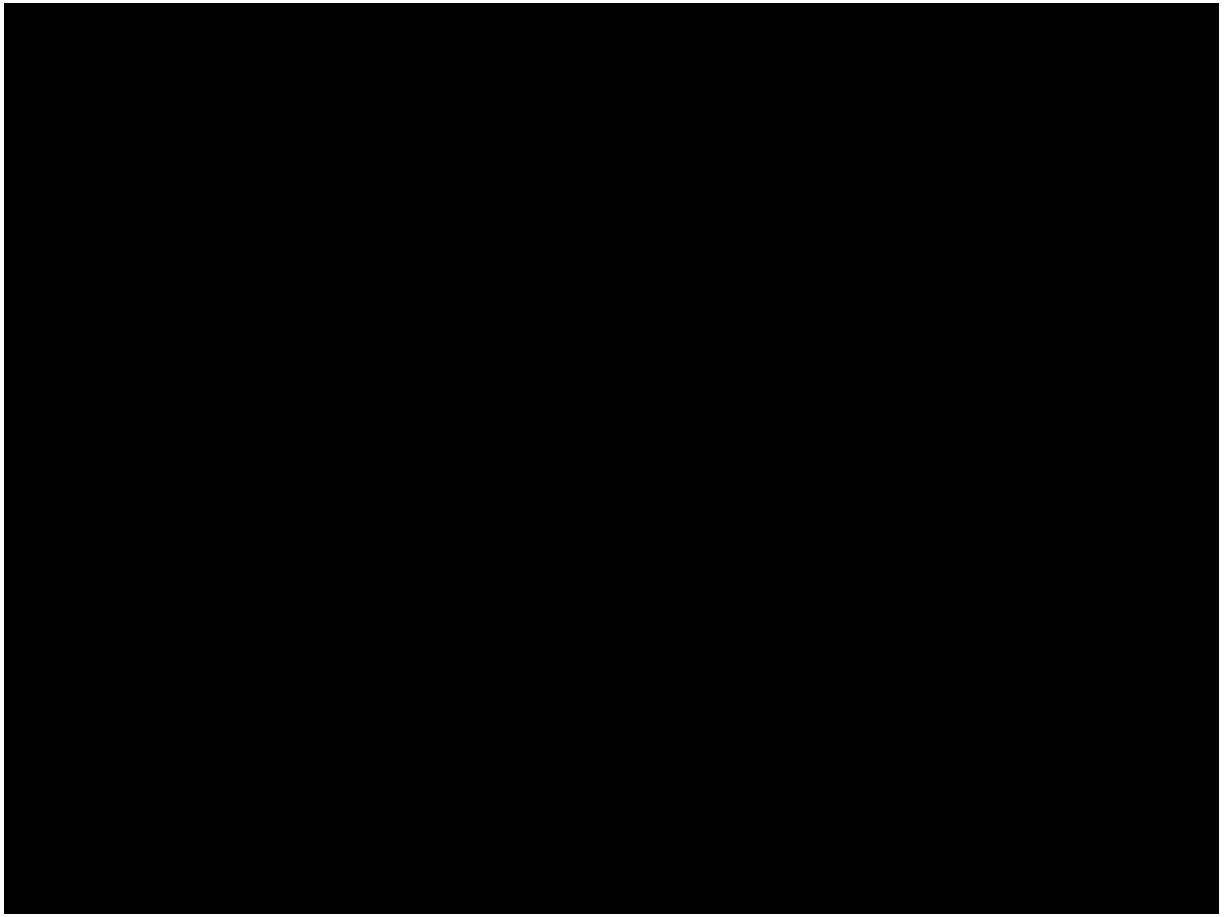

http://heli.stanford.edu/

Reinforcement Learning

A simplified walkthrough: Playing Pong!

Reinforcement Learning

Goal: We want to teach our computer to play the game of Pong

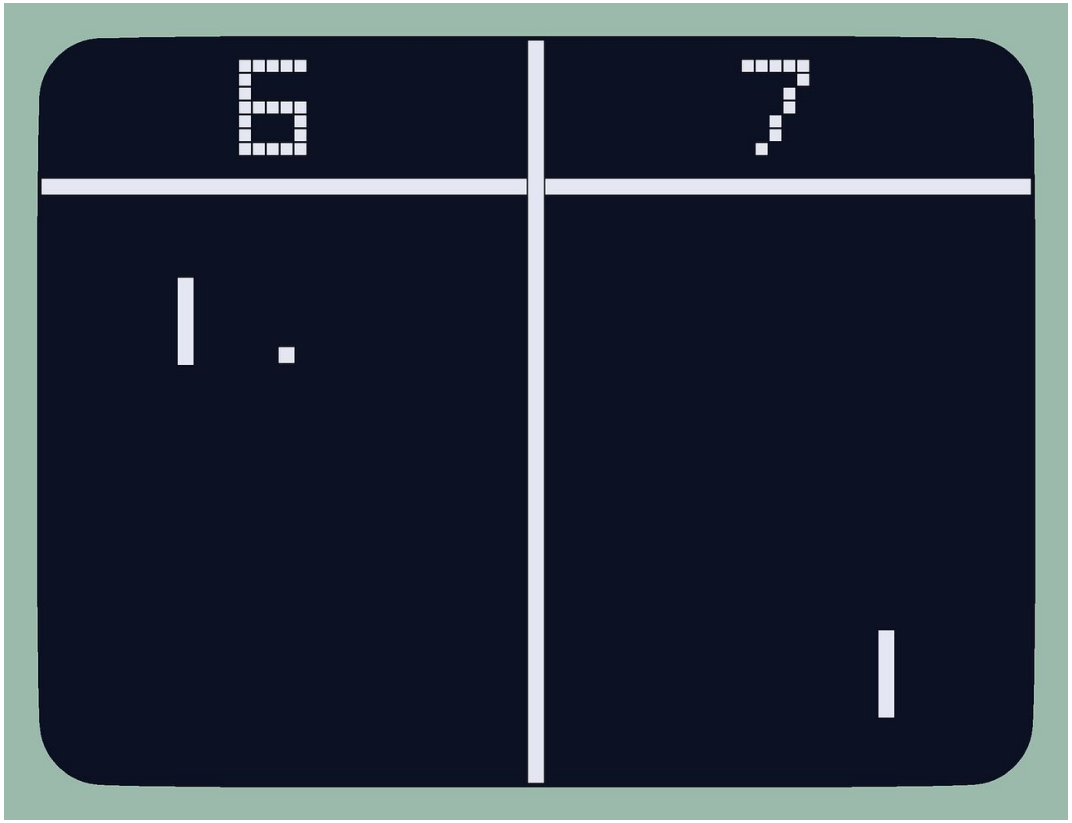

Reinforcement Learning

Goal: We want to teach our computer to play the game of Pong

General Idea: At each timestep (say every 100ms), we want to see where the ball is, and decide if we want to move our paddle up or down

Reinforcement Learning

Training process:

1. Initially, we will move our paddle up or down randomly, and most probably fail

Reinforcement Learning

Training process:

1. Initially, we will move our paddle up or down randomly, and most probably fail

2. Before failing, we have taken a set of actions

```
[up, down, up, up, up, down, down, up, down ...]
```

Reinforcement Learning

Training process:

1. Initially, we will move our paddle up or down randomly, and most probably fail

2. Before failing, we have taken a set of actions

`[up, down, up, up, up, down, down, up, down ...]`

3. Some action(s) caused us to eventually fail, so let us penalize all these actions a little bit (In some variations, penalize recent actions more)

Reinforcement Learning

Training process:

1. Initially, we will move our paddle up or down randomly, and most probably fail

2. Before failing, we have taken a set of actions

`[up, down, up, up, up, down, down, up, down ...]`

3. Some action(s) caused us to eventually fail, so let us penalize all these actions a little bit (In some variations, penalize recent actions more)

4. Every time we succeed, boost all the actions a little bit

Reinforcement Learning

Training process:

1. Initially, we will move our paddle up or down randomly, and most probably fail

2. Before failing, we have taken a set of actions

`[up, down, up, up, up, down, down, up, down ...]`

3. Some action(s) caused us to eventually fail, so let us penalize all these actions a little bit (In some variations, penalize recent actions more)

4. Every time we succeed, boost all the actions a little bit

5. Over time, “good actions at the right timesteps” will be boosted, and incorrect actions will be penalized!

Reinforcement Learning

Now make your computer play a few million games of Pong.

Overtime, it will learn to perform the right actions depending on the location of the ball!

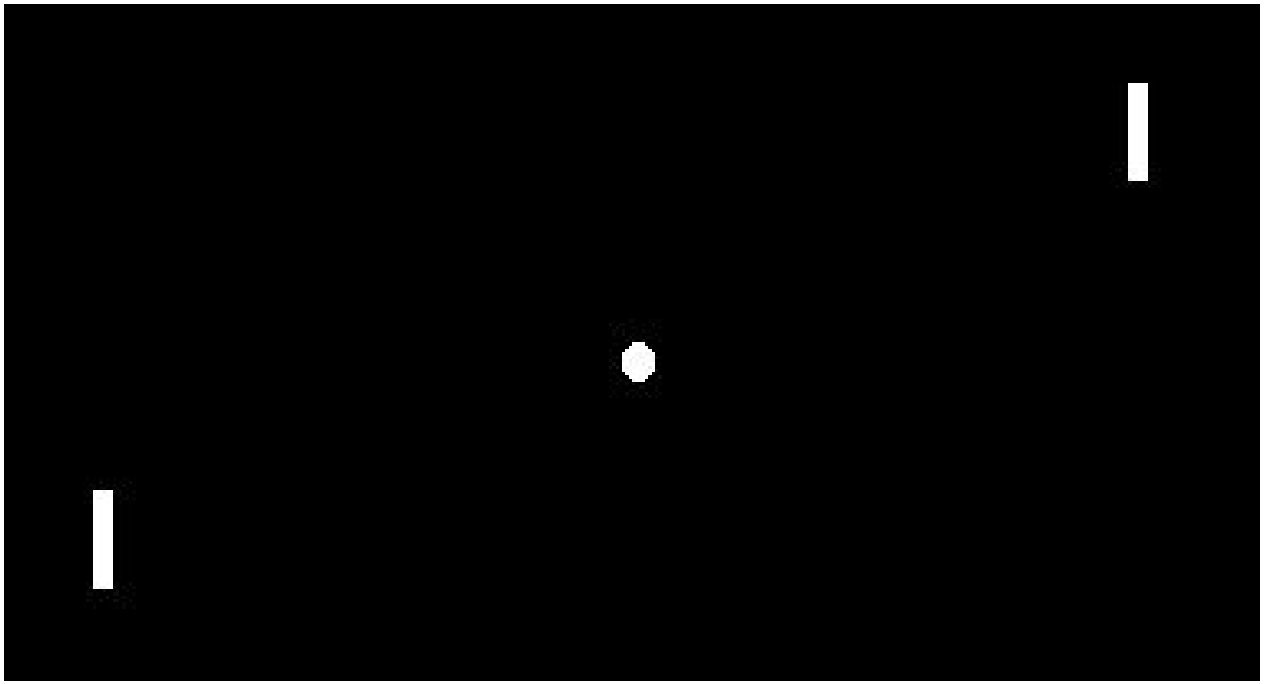

http://karpathy.github.io/2016/05/31/rl/

Reinforcement Learning

More examples: AlphaGo

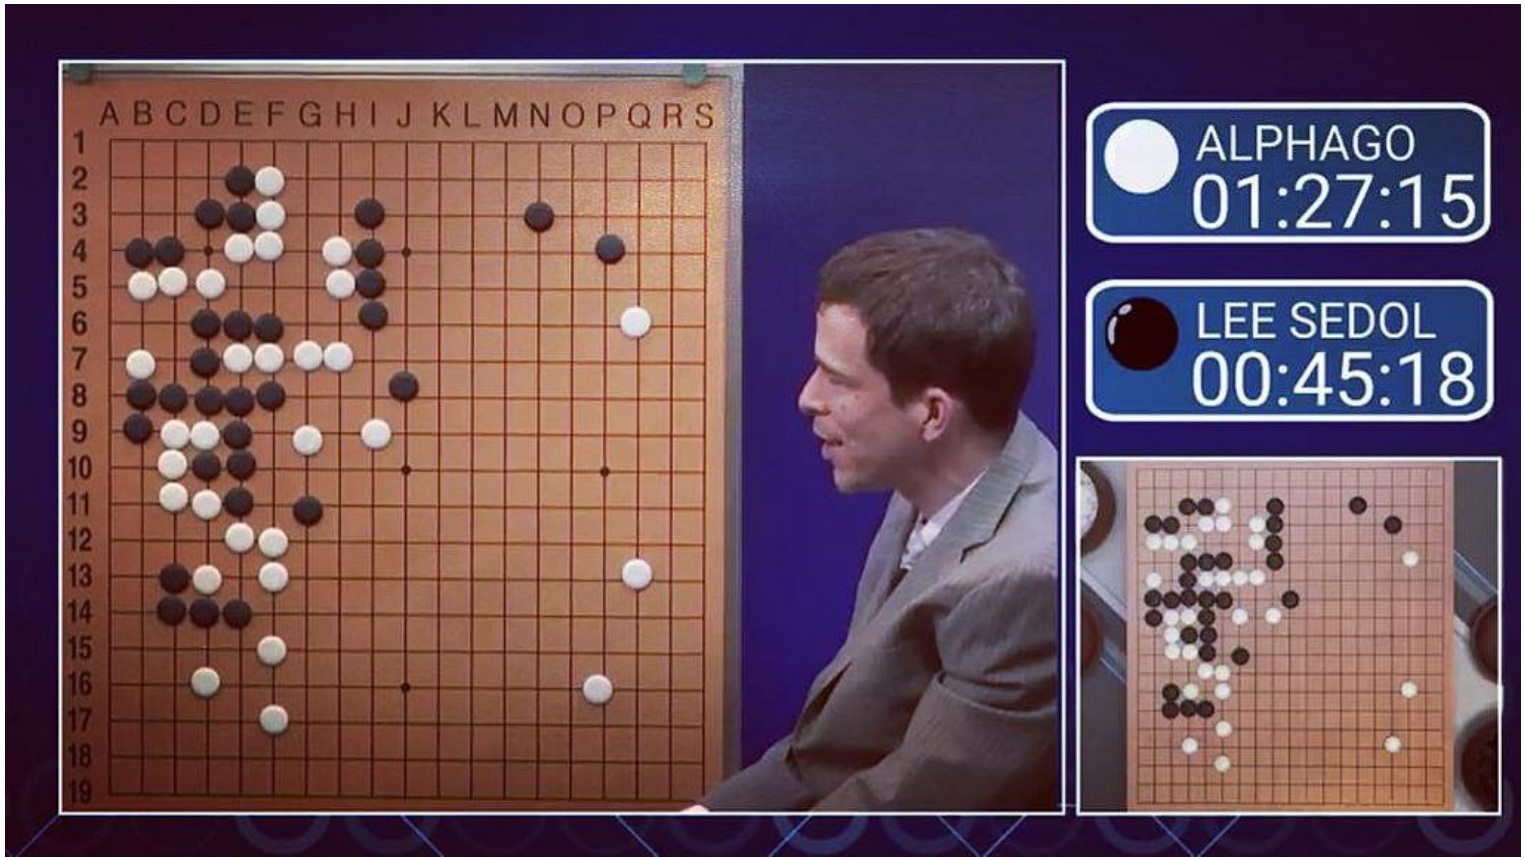

Reinforcement Learning

More examples: DOTA

Reinforcement Learning

More examples: Not just games...

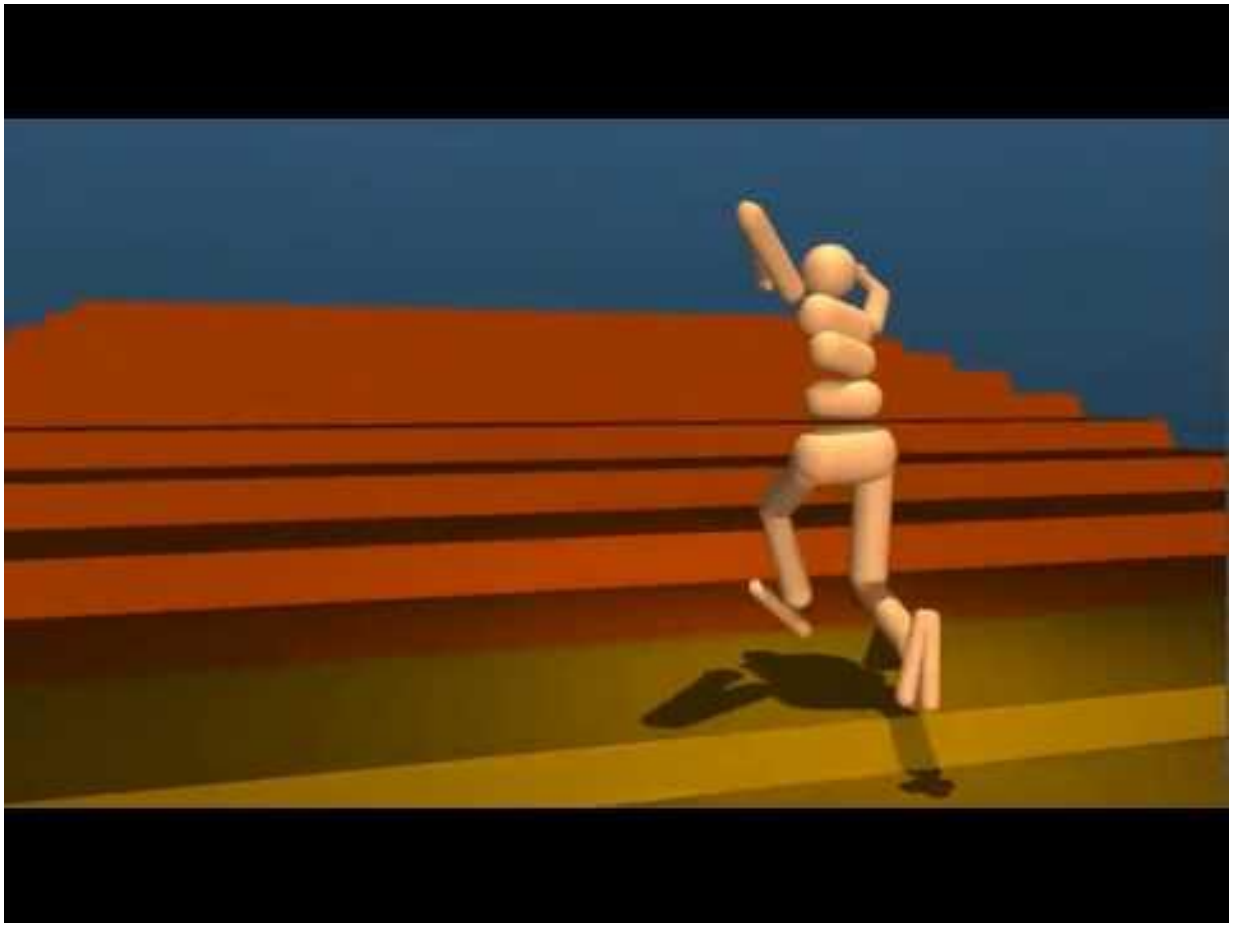

Summary

- Domain adaptation aims to use all available data in favor of the in-domain data
- Multilingual systems enable Zero-shot translation
- Multi-task learning improves generalization capability of the model
- Multi-modal learning is the way forward to build general AI's that understand the world like humans do
- Generative Adversarial Networks aim to solve the inverse problem of generating instead of just classifying
- Reinforcement Learning is a super-general framework that can help us teach agents how to act and react in the real world