Discreteness in Generative Models

-- *Discrete Sequence Generation*

Hao Dong

Peking University

Discreteness in Generative Models
– Discrete Sequence Generation

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

PEKING UNIVERSITY

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

RNN Language Modelling

- **Recap: Recurrent Neural Network**

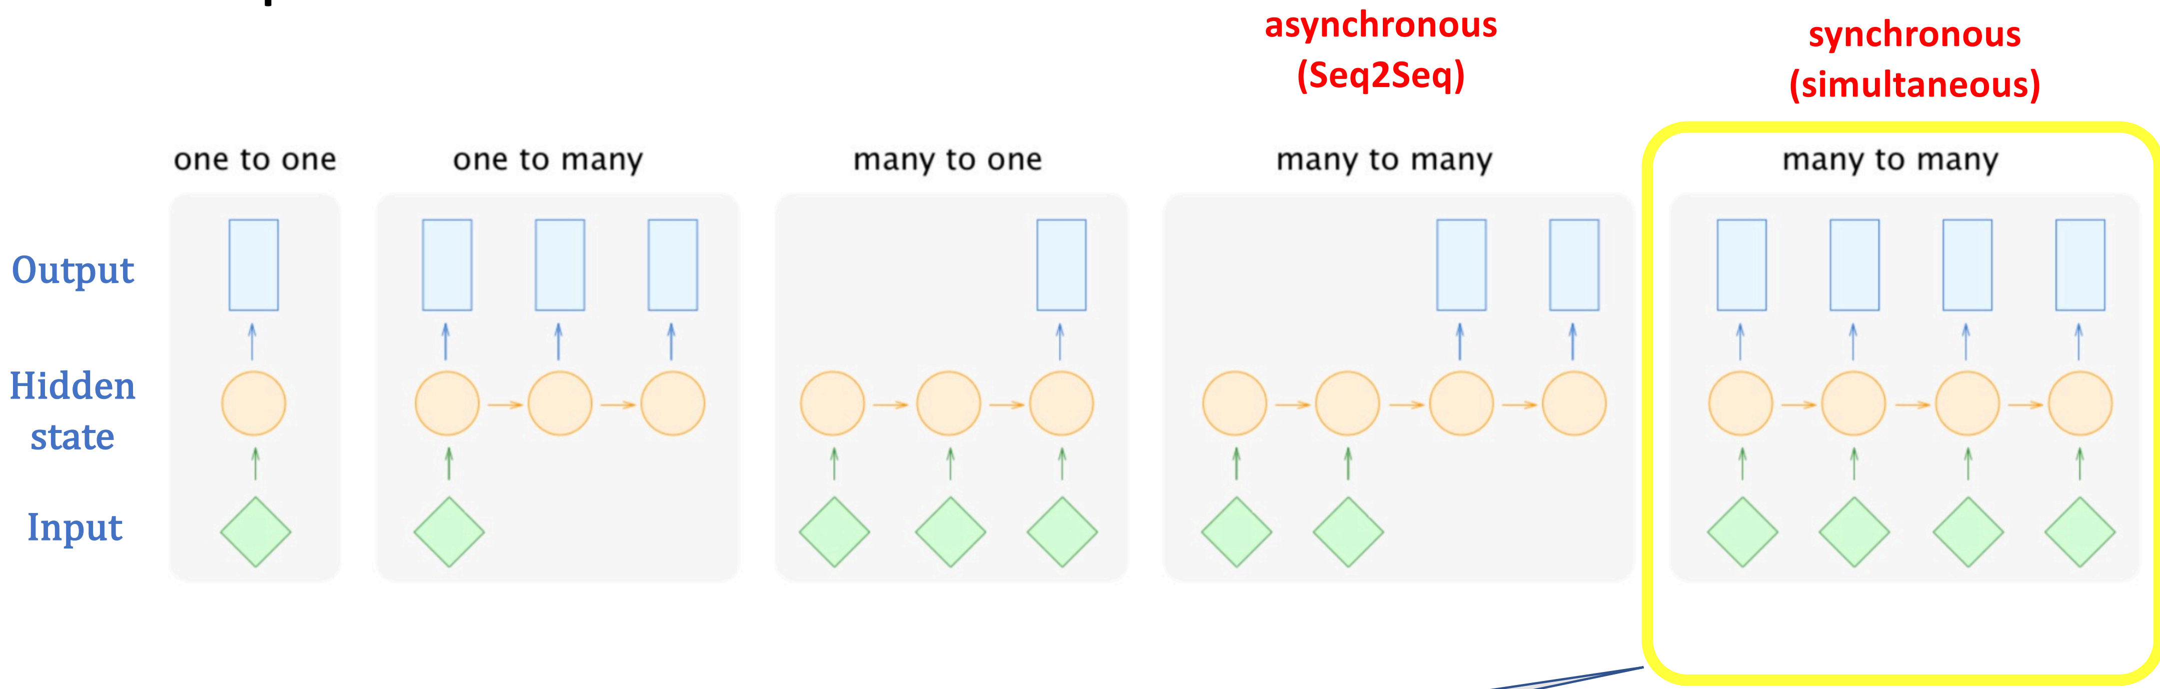

Multiple inputs and multiple outputs

RNN Language Modelling

- **Synchronous Many-to-Many**

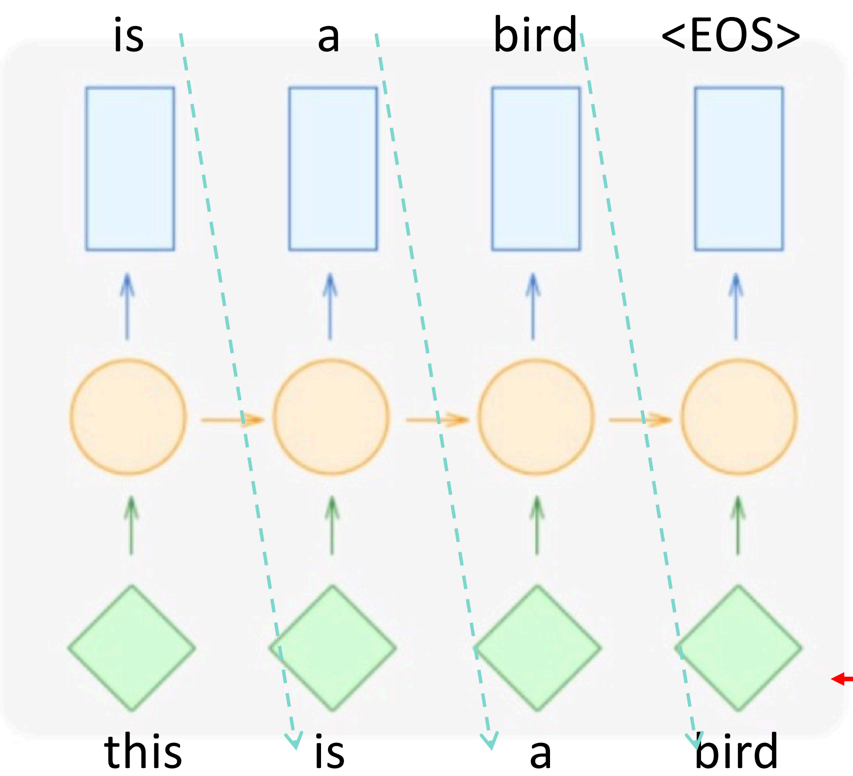

- MLE training
- The output of each step is equal to the input of its next step.

For testing, input “this,” output the entire sentence

RNN Language Modelling

- **Limitations**
  - 1. No latent space – Representation learning
  - 2. Exposure bias – Long sentence generation problem

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inversed RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

Generating Sentences from a Continuous Space

- **Why Continuous Space**

i went to the store to buy some groceries .
i store to buy some groceries .
i were to buy any groceries .
horses are to buy any groceries .
horses are to buy any animal .
horses the favorite any animal .
horses the favorite favorite animal .
horses are my favorite animal .

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **VAE Approach**

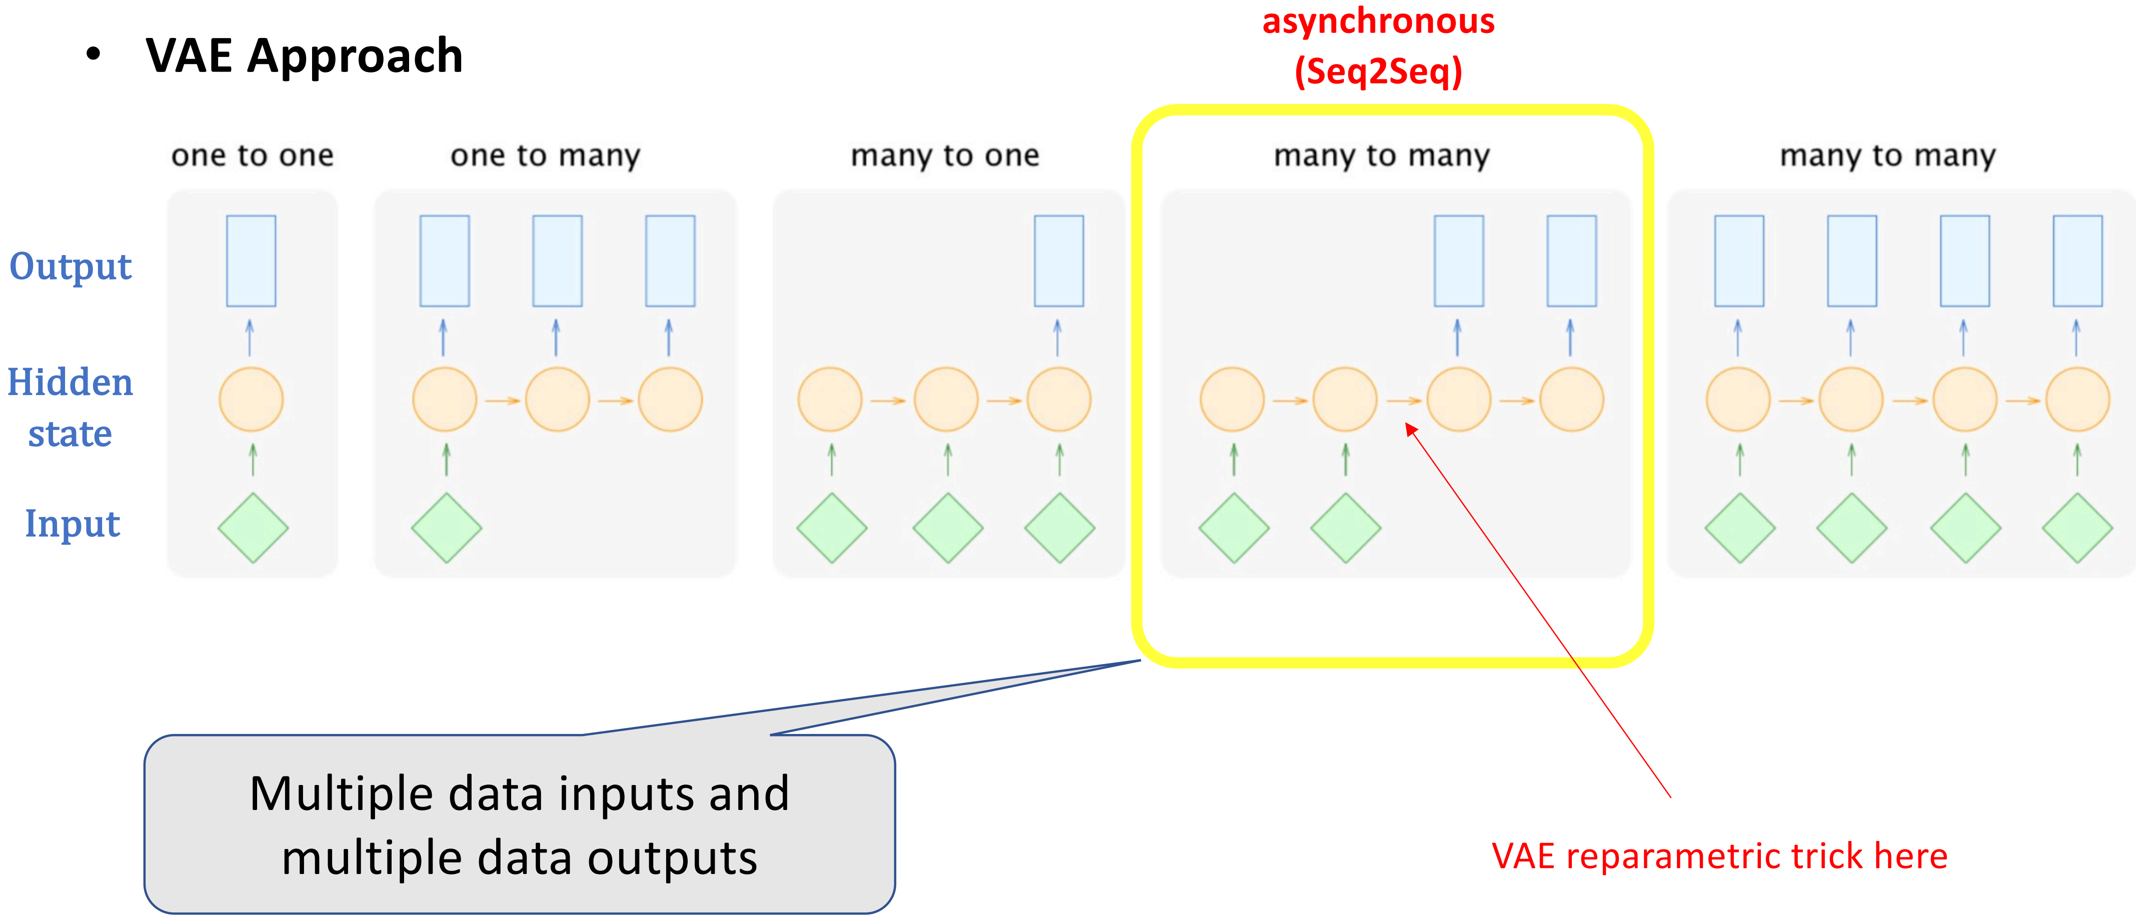

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **VAE Approach: Training**

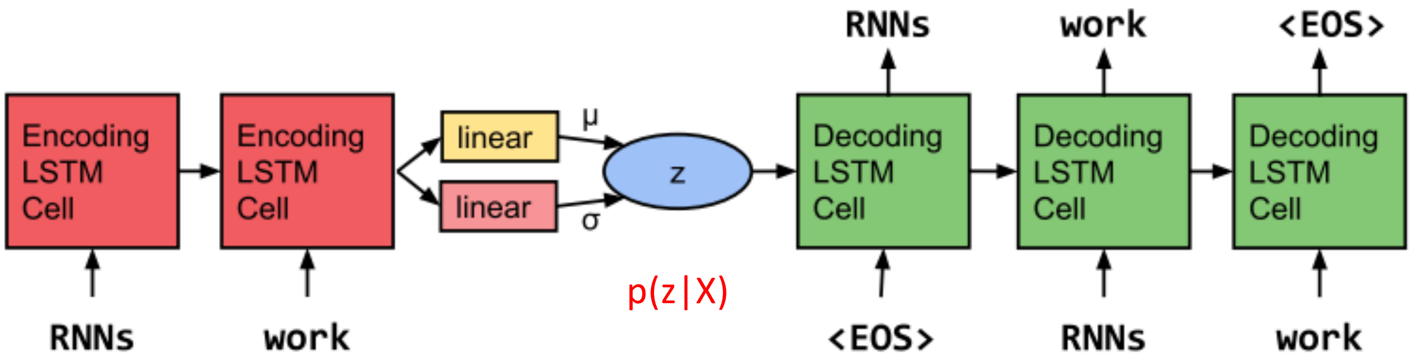

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **VAE Approach: Testing**

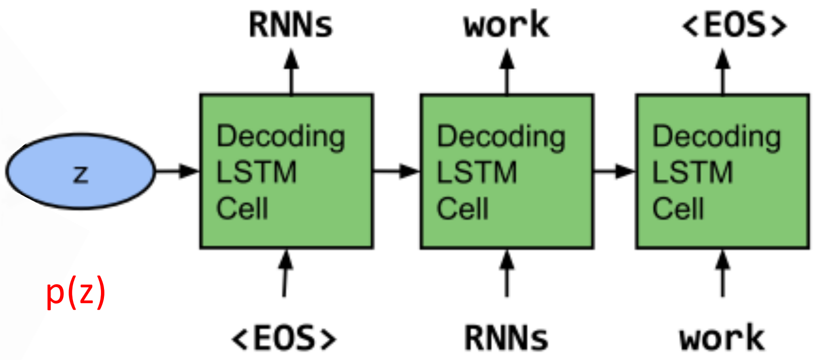

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **Application**

Impute missing words within sentences.

but now , as they parked out front and owen stepped out of the car , he could see _ _ _ _ _ _		
True: *that the transition was complete .*	**RNNLM:** *it , " i said .*	**VAE:** *through the driver 's door .*
you kill him and his _ _		
True: *men .*	**RNNLM:** *. "*	**VAE:** *brother .*
not surprising , the mothers dont exactly see eye to eye with me _ _ _ _		
True: *on this matter .*	**RNNLM:** *, i said .*	**VAE:** *, right now .*

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **Application**

Results from the mean of the posterior distribution and samples from that distribution.

INPUT	**we looked out at the setting sun .**	**i went to the kitchen .**	**how are you doing ?**
MEAN	*they were laughing at the same time .*	*i went to the kitchen .*	*what are you doing ?*
SAMP. 1	*ill see you in the early morning .*	*i went to my apartment .*	*" are you sure ?*
SAMP. 2	*i looked up at the blue sky .*	*i looked around the room .*	*what are you doing ?*
SAMP. 3	*it was down on the dance floor .*	*i turned back to the table .*	*what are you doing ?*

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **Application**

Interpolation between two points in the VAE manifold.

" i want to talk to you . "
"i want to be with you . "
"i do n't want to be with you . "
i do n't want to be with you .
she did n't want to be with him .
he was silent for a long moment .
he was silent for a moment .
it was quiet for a moment .
it was dark and cold .
there was a pause .
it was my turn .

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

Generating Sentences from a Continuous Space

- **Limitation**

  Difficult to generate long sentences

Generating Sentences from a Continuous Space.
Bowman, Samuel R. Vilnis, Luke. Vinyals, Oriol. Dai, Andrew. Jozefowicz, Rafal. Bengio, Samy. Arxiv 2015.

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

Recap: Inverse RL vs. GAN

- **Reinforcement Learning** and Imitation Learning

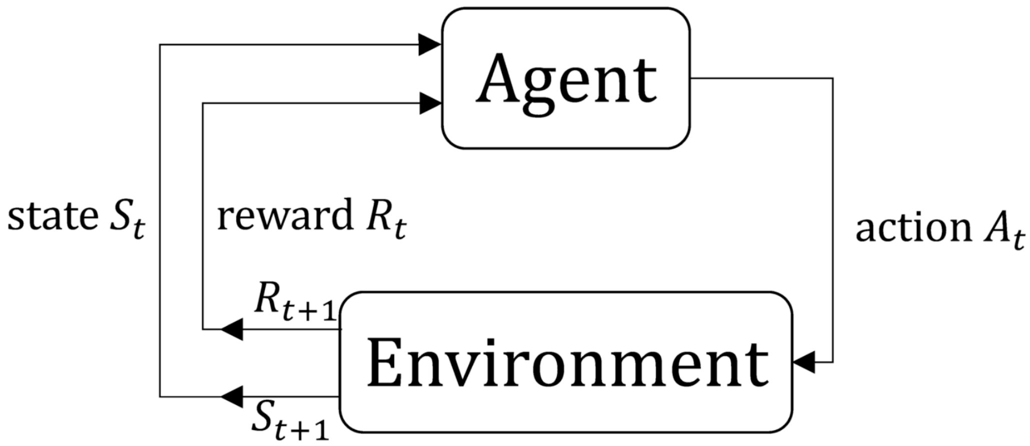

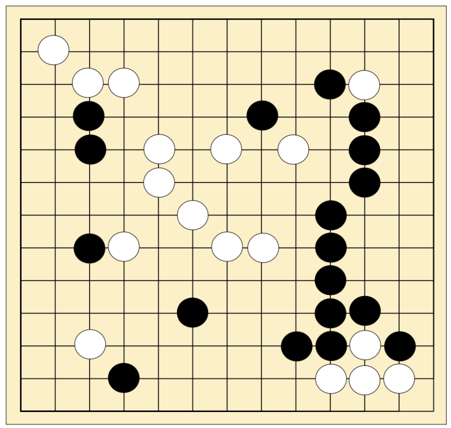

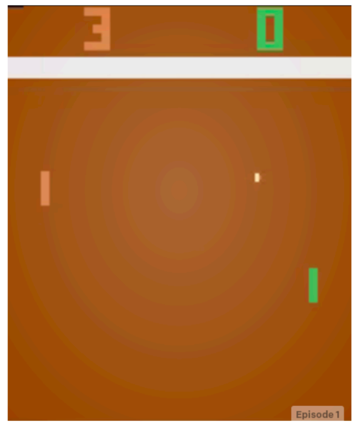

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**

  - **Why?**

    Difficult and even impossible to define reward functions for many environments

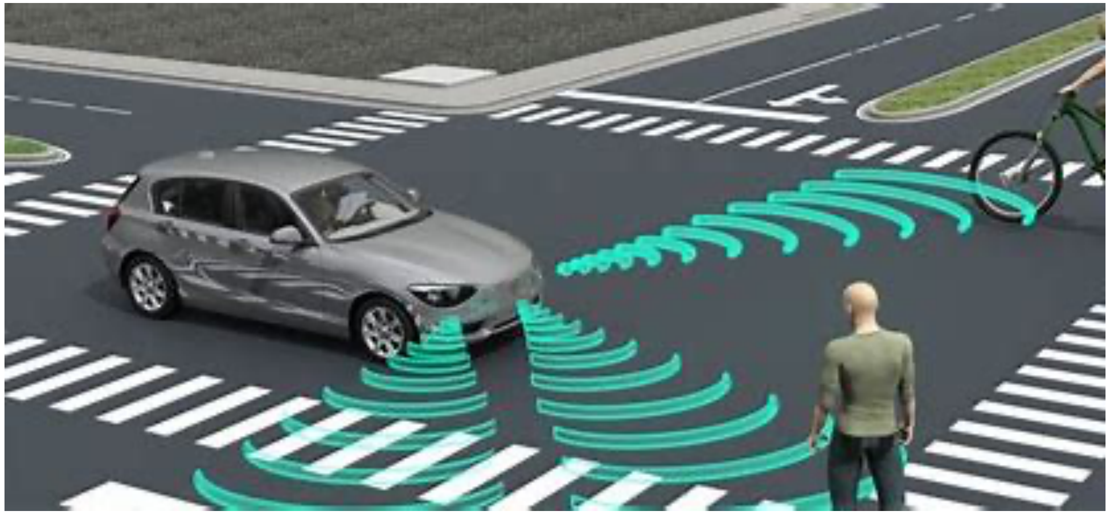

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**
  - **Behaviour Cloning**
  - **Inverse RL**

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**
  - **Behaviour Cloning == Supervised Learning**
  - Given expert's demonstrations: (s_1, a_1), (s_2, a_2)..... (s_N, a_N)
  - Supervised training:

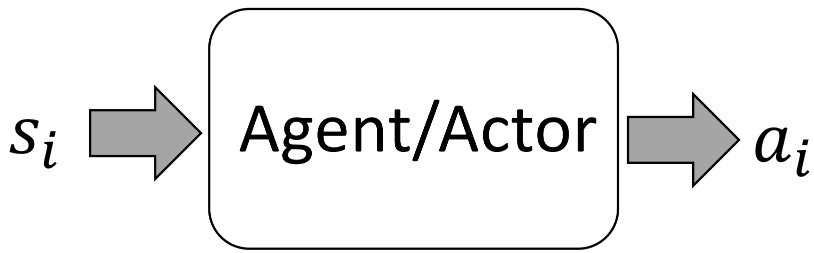

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**
  - **Behaviour Cloning == Supervised Learning**
  - Problem
    - Expert only samples from limited observation (states)
      (Solution: Dataset Aggregation)
    - If machine has limited capacity, it may choose the wrong behavior to copy
    - Assume the training and testing data distributions are the same

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**
    - **Inverse Reinforcement Learning**

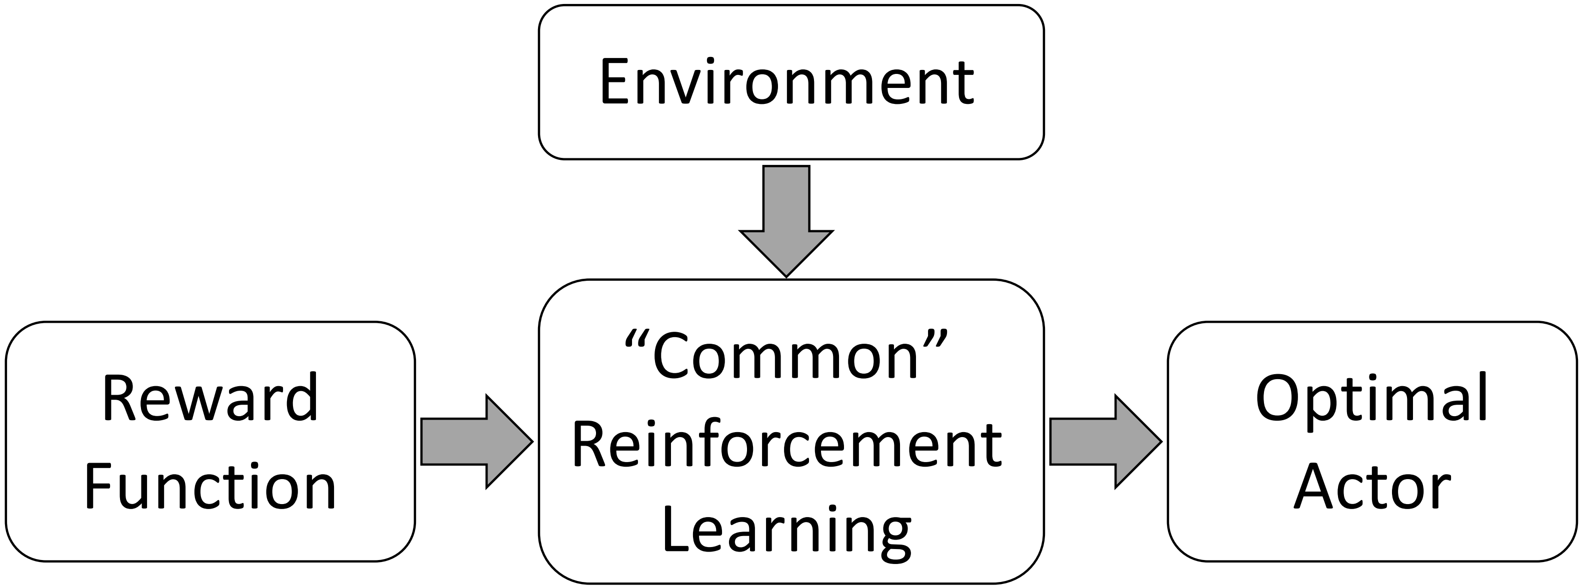

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**
  - **Inverse Reinforcement Learning**

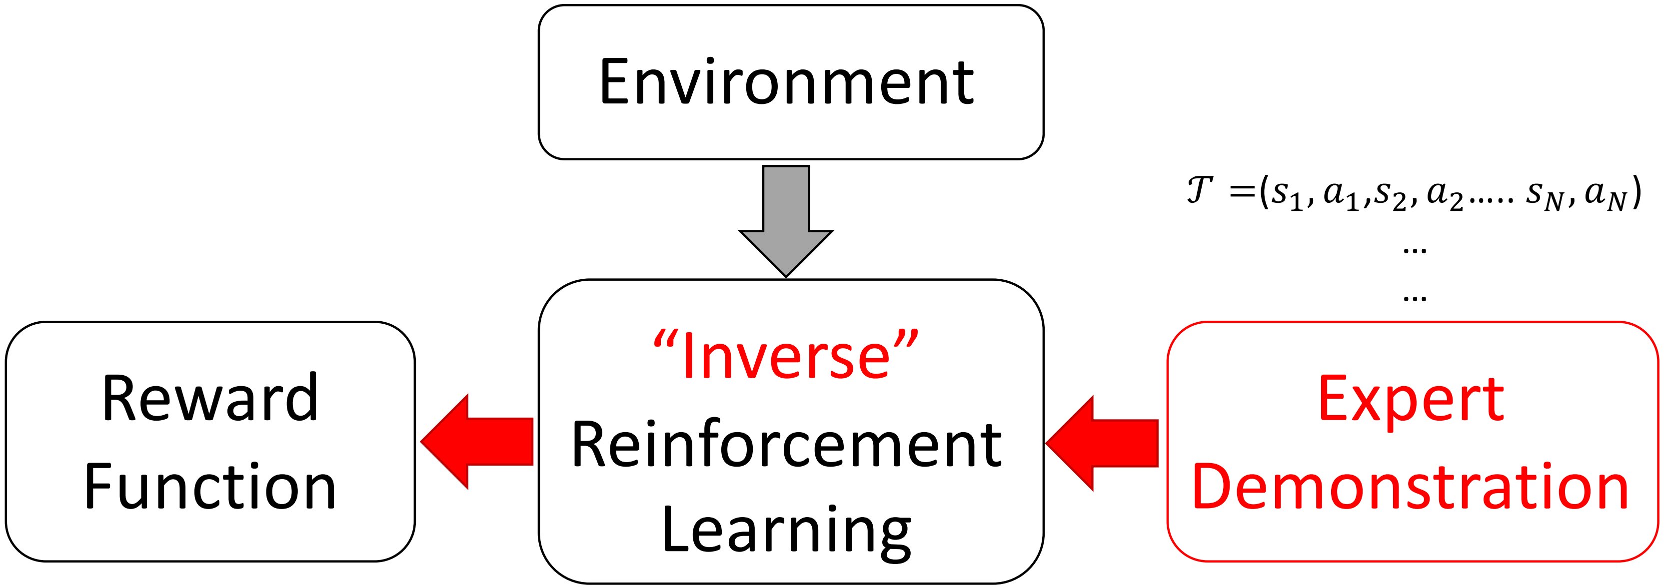

Recap: Inverse RL vs. GAN

- Reinforcement Learning and **Imitation Learning**

  - **Inverse Reinforcement Learning**

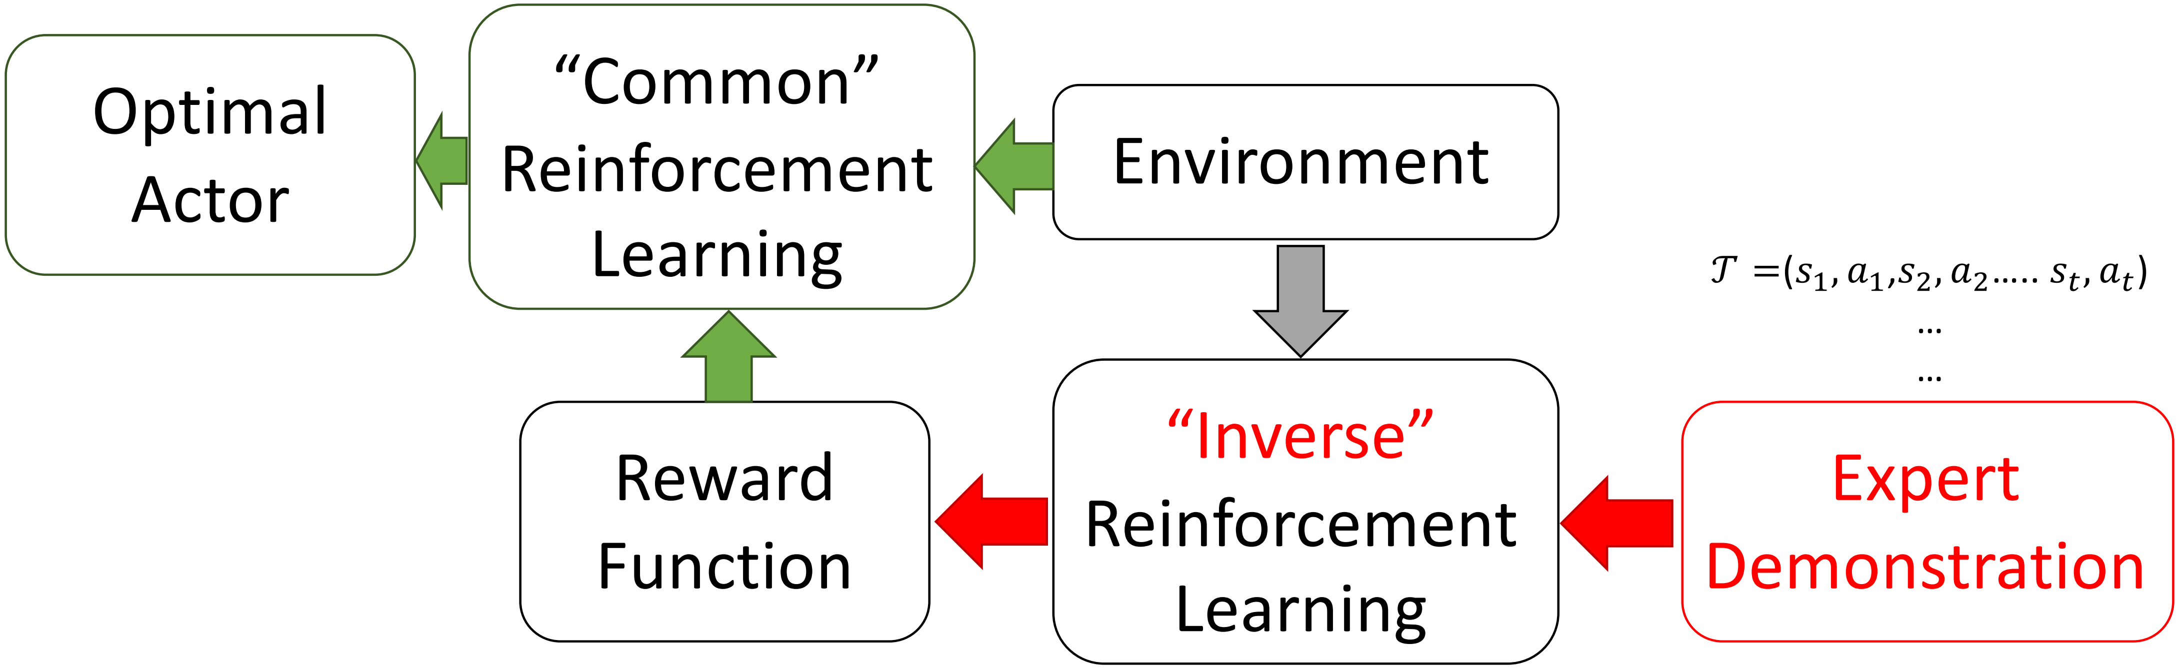

Recap: Inverse RL vs. GAN

- **Inverse RL: Training**

$$\sum_{n=1}^{N} R(\mathcal{T}_n) > \sum_{n=1}^{N} R(\hat{\mathcal{T}}_n)$$

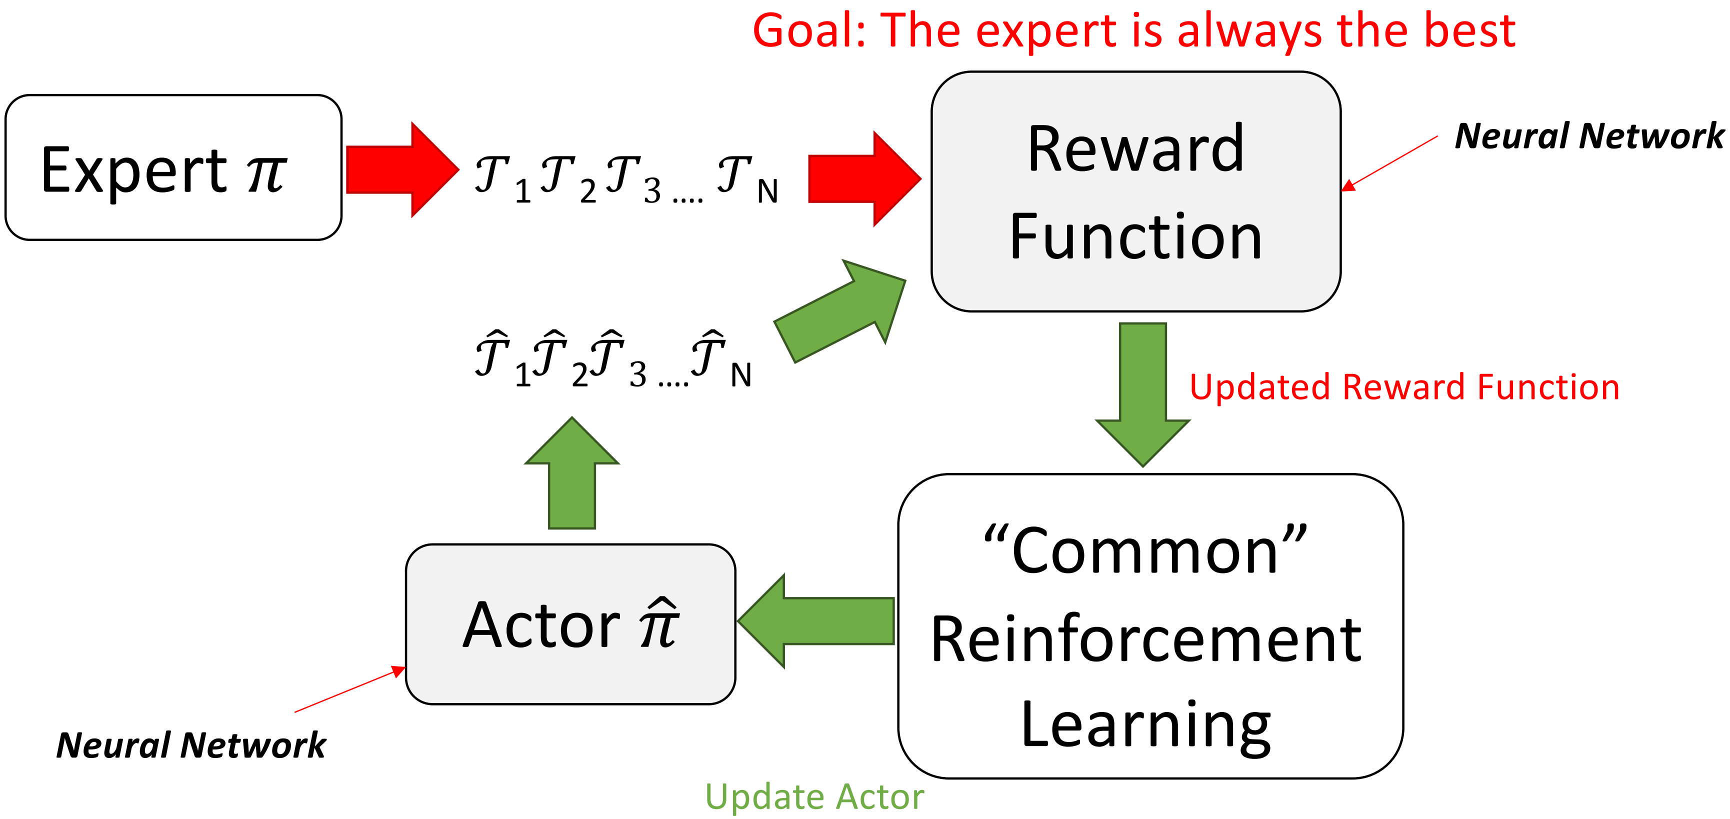

$\mathcal{T} = (s_1, a_1, s_2, a_2 \ldots\ldots s_t, a_t)$

Recap: Inverse RL vs. GAN

- **GAN**

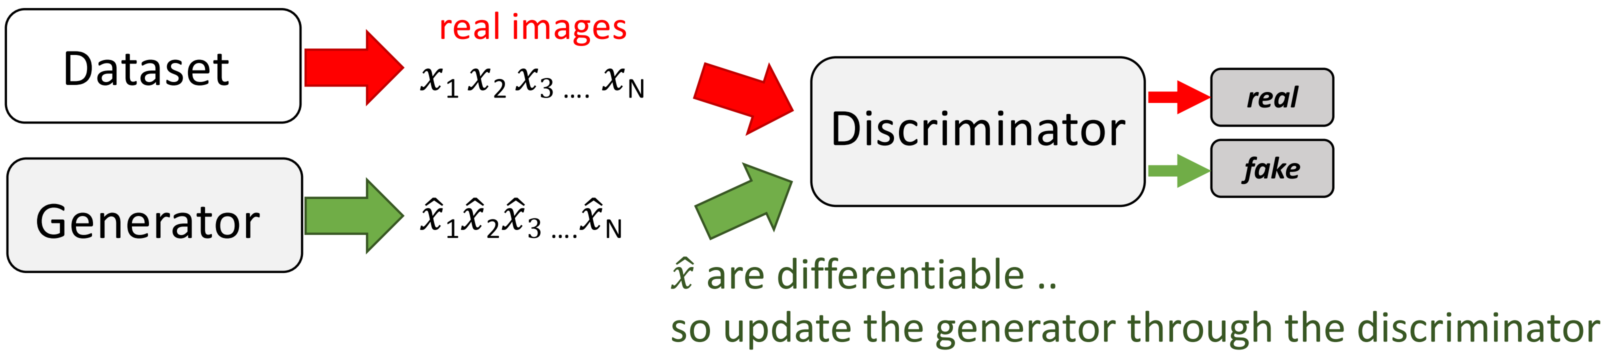

- **Inverse RL**

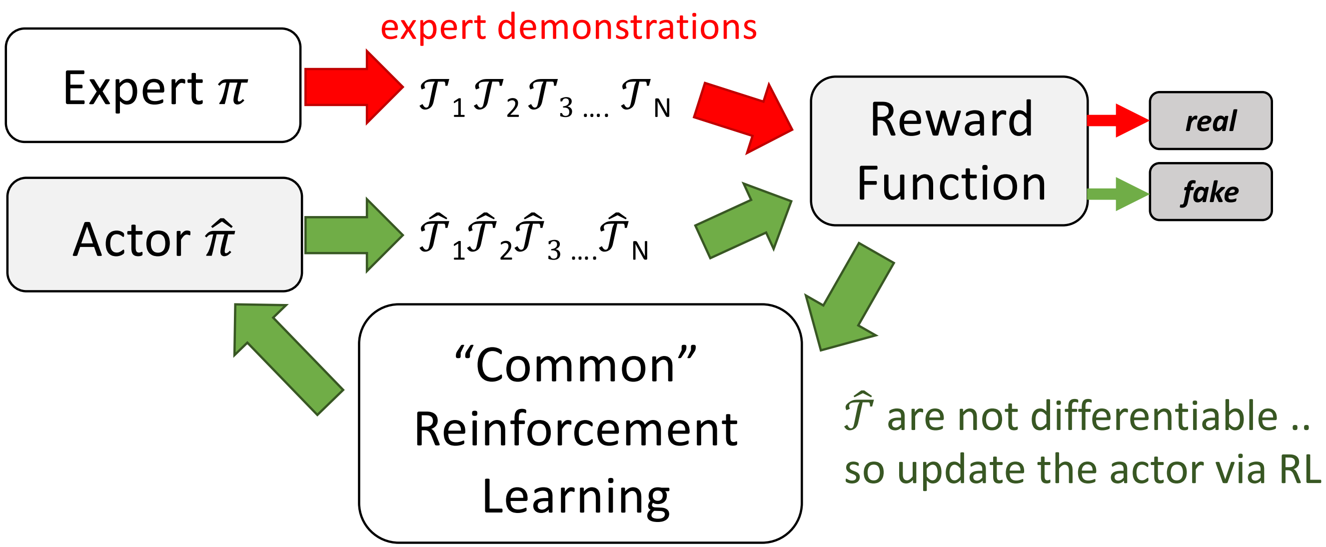

PEKING UNIVERSITY

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient

- **Motivation**

    MLE has exposure bias

    MLE tends to move to the mean when there is uncertainty

    “this is a …” “this is one” ← “a” or “one” which one is correct?

    MLE-free sentence generation using adversarial networks could be better

- **Challenge**

    Sentences are discrete

    Sentence generation is not differentiable

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. Lantao Yu, Weinan Zhang, Jun Wang, Yong Yu. AAAI. 2016.

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient

- **Inverse RL**

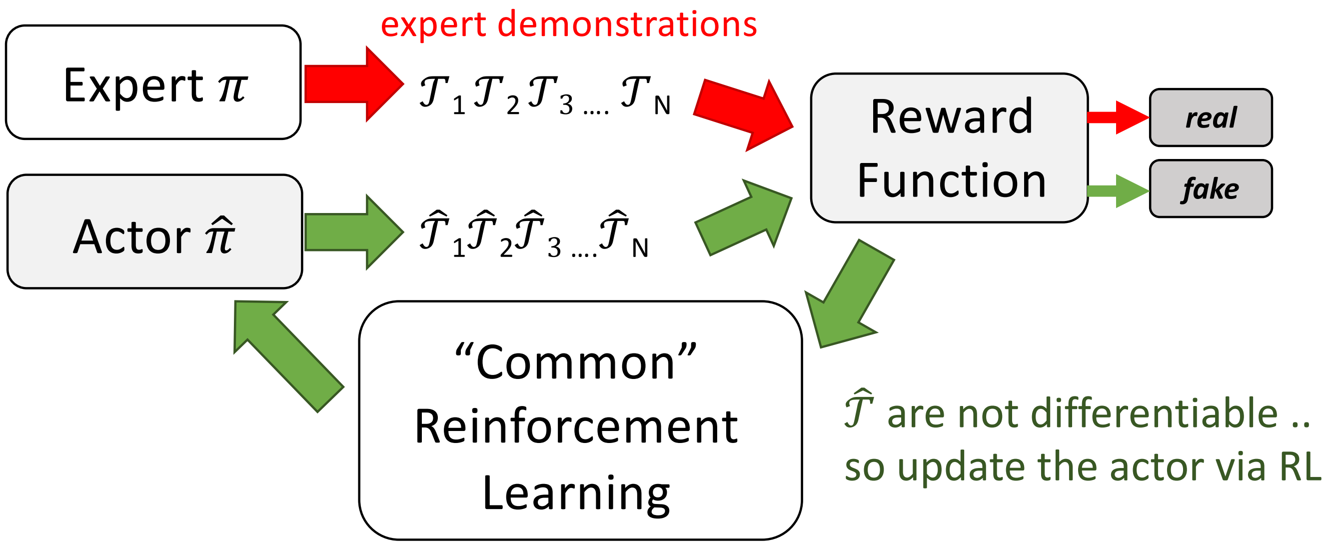

- **SeqGAN**

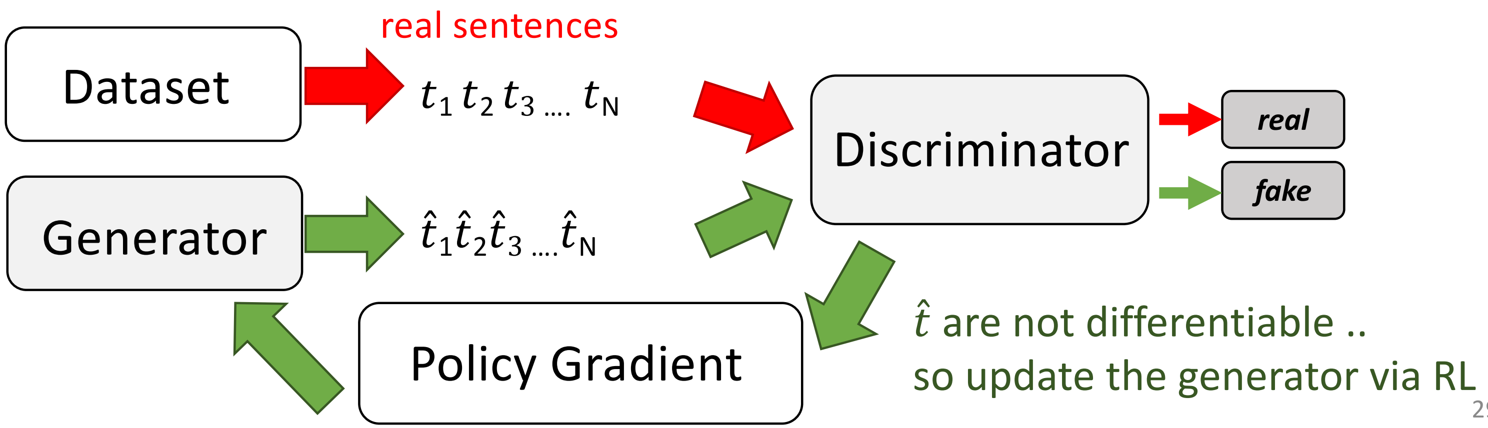

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient

- **Method Details**

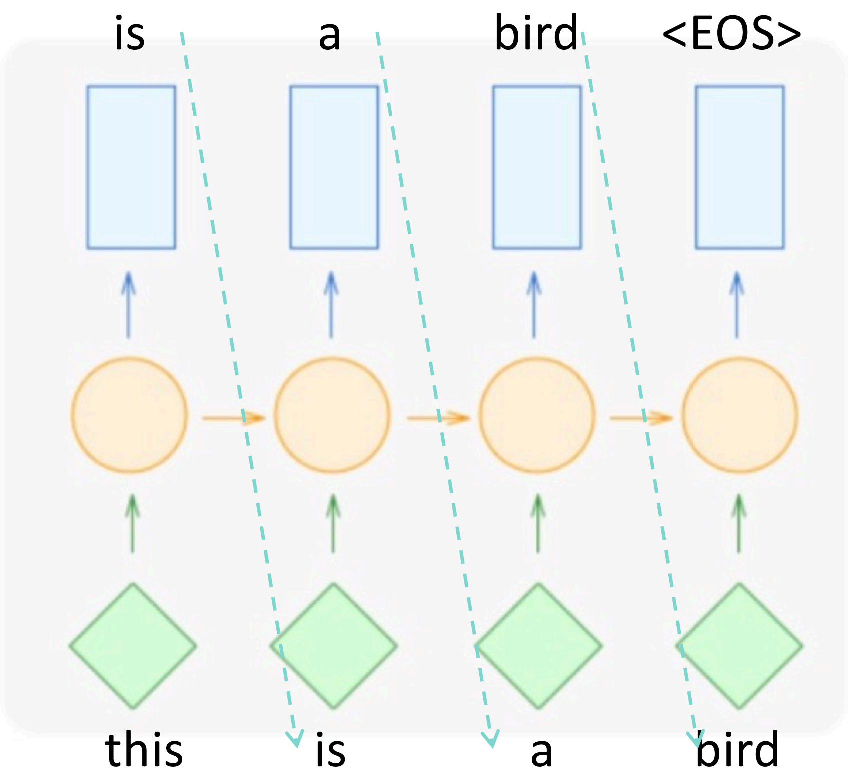

- The generator is a LSTM or GRU language model
- Pretrain the generator using MLE before RL training for better initialization
- The discriminator is a CNN network

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. Lantao Yu, Weinan Zhang, Jun Wang, Yong Yu. AAAI. 2016.

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient

- **Method Details**

Algorithm 1 Sequence Generative Adversarial Nets

Require: generator policy G_θ; roll-out policy G_β; discriminator D_ϕ; a sequence dataset $\mathcal{S} = \{X_{1:T}\}$

Initialize G_θ, D_ϕ with random weights θ, ϕ.
Pre-train G_θ using MLE on $\mathcal{S}$
$\beta \leftarrow \theta$
Generate negative samples using G_θ for training D_ϕ
Pre-train D_ϕ via minimizing the cross entropy
repeat
    for g-steps **do**
        Generate a sequence $Y_{1:T} = (y_1, \ldots, y_T) \sim G_\theta$
        for t in $1 : T$ **do**
            Compute $Q(a = y_t; s = Y_{1:t-1})$ by Eq. (4)
        end for
        Update generator parameters via policy gradient Eq. (8)
    end for
    for d-steps **do**
        Use current G_θ to generate negative examples and combine with given positive examples $\mathcal{S}$
        Train discriminator D_ϕ for k epochs by Eq. (5)
    end for
    $\beta \leftarrow \theta$
until SeqGAN converges

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. Lantao Yu, Weinan Zhang, Jun Wang, Yong Yu. AAAI. 2016.

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient

- **Results**

Table 2: Chinese poem generation performance comparison.

Algorithm	Human score	p-value	BLEU-2	p-value
MLE	0.4165	0.0034	0.6670	$< 10^{-6}$
SeqGAN	**0.5356**		**0.7389**	
Real data	0.6011		0.746	

Table 3: Obama political speech generation performance.

Algorithm	BLEU-3	p-value	BLEU-4	p-value
MLE	0.519	$< 10^{-6}$	0.416	0.00014
SeqGAN	**0.556**		**0.427**	

Table 4: Music generation performance comparison.

Algorithm	BLEU-4	p-value	MSE	p-value
MLE	0.9210	$< 10^{-6}$	22.38	0.00034
SeqGAN	**0.9406**		**20.62**	

SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. Lantao Yu, Weinan Zhang, Jun Wang, Yong Yu. AAAI. 2016.

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

RankGAN: Adversarial Ranking for Language Generation

- **Motivation**

  the existing GANs restrict the discriminator to be a binary classifier,
  limiting their learning capacity for tasks that need to synthesize output with rich structures
  such as natural language descriptions

  so … we use a Ranker as the discriminator … improved SeqGAN

RankGAN: Adversarial Ranking for Language Generation.
Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, Ming-Ting Sun. NeurIPS. 2017

RankGAN: Adversarial Ranking for Language Generation

- **Method**

Generator hopes the Ranker rank the fake sentence to the top
… a better reward function for RL training

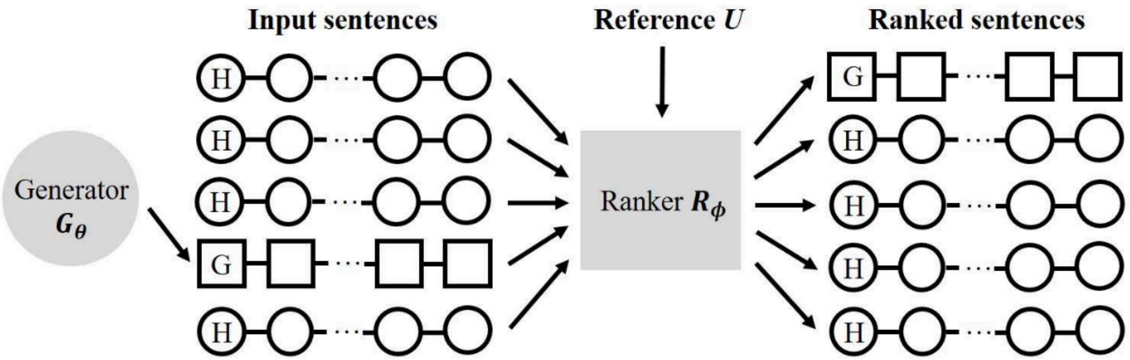

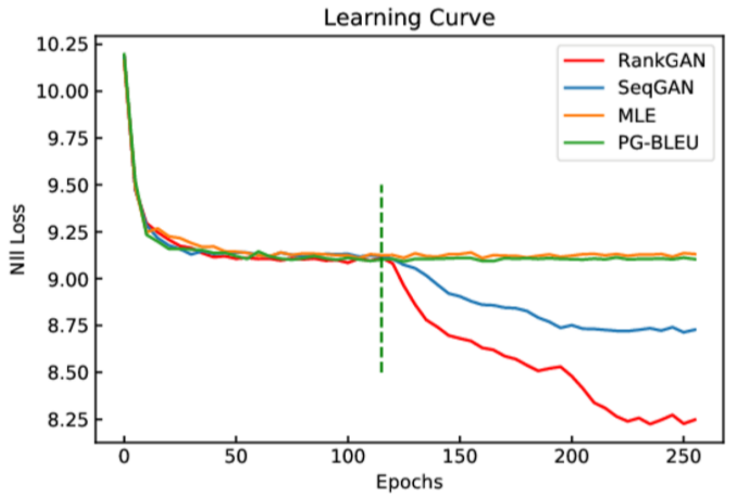

RankGAN: Adversarial Ranking for Language Generation.
Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, Ming-Ting Sun. NeurIPS. 2017

RankGAN: Adversarial Ranking for Language Generation

北京大學 PEKING UNIVERSITY

- **Method**

$$\min_{\theta} \max_{\phi} \mathfrak{L}(G_\theta, R_\phi) = \mathop{\mathbb{E}}_{s \sim \mathcal{P}_h} \left[\log R_\phi(s|U, \mathcal{C}^-) \right] + \mathop{\mathbb{E}}_{s \sim G_\theta} \left[\log(1 - R_\phi(s|U, \mathcal{C}^+)) \right]$$

U is a set of real data for reference
C^- is a set of randomly sampled fake data
C^+ is a set of randomly sampled real data
p_h and G_θ are real and fake distribution
The "Rank" of a data is the similarity with the reference U
y_s and y_u are the feature vector of U and s (s is a singe data point)
The similarity score of the input sequence s given a reference u:

$$\alpha(s|u) = cosine(y_s, y_u) = \frac{y_s \cdot y_u}{||y_s|| ||y_u||}$$

The ranking score for a certain sequence s given a comparison set C:

$$P(s|u, \mathcal{C}) = \frac{exp(\gamma \alpha(s|u))}{\sum_{s' \in \mathcal{C}'} exp(\gamma \alpha(s'|u))}$$

The ranking score for the input sentence s :

$$R_\phi(s|U, \mathcal{C}) = \mathop{\mathbb{E}}_{u \in U} \left[P(s|u, \mathcal{C}) \right]$$

RankGAN: Adversarial Ranking for Language Generation.
Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, Ming-Ting Sun. NeurIPS. 2017

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

MaskGAN: Better Text Generation via Filling in the ____

- **Motivation**

SeqGAN and RankGAN use Policy Gradient for RL training,

… thus we can only have a single reward for the whole sentence

Could we have a reward for each word?

MaskGAN uses Actor Critic for RL training.

MaskGAN: Better Text Generation via Filling in the ____ . William Fedus, Ian Goodfellow, Andrew M. Dai. ICLR 2018

MaskGAN: Better Text Generation via Filling in the ____

- **Method**

Both Generator and Discriminator are Seq2Seq

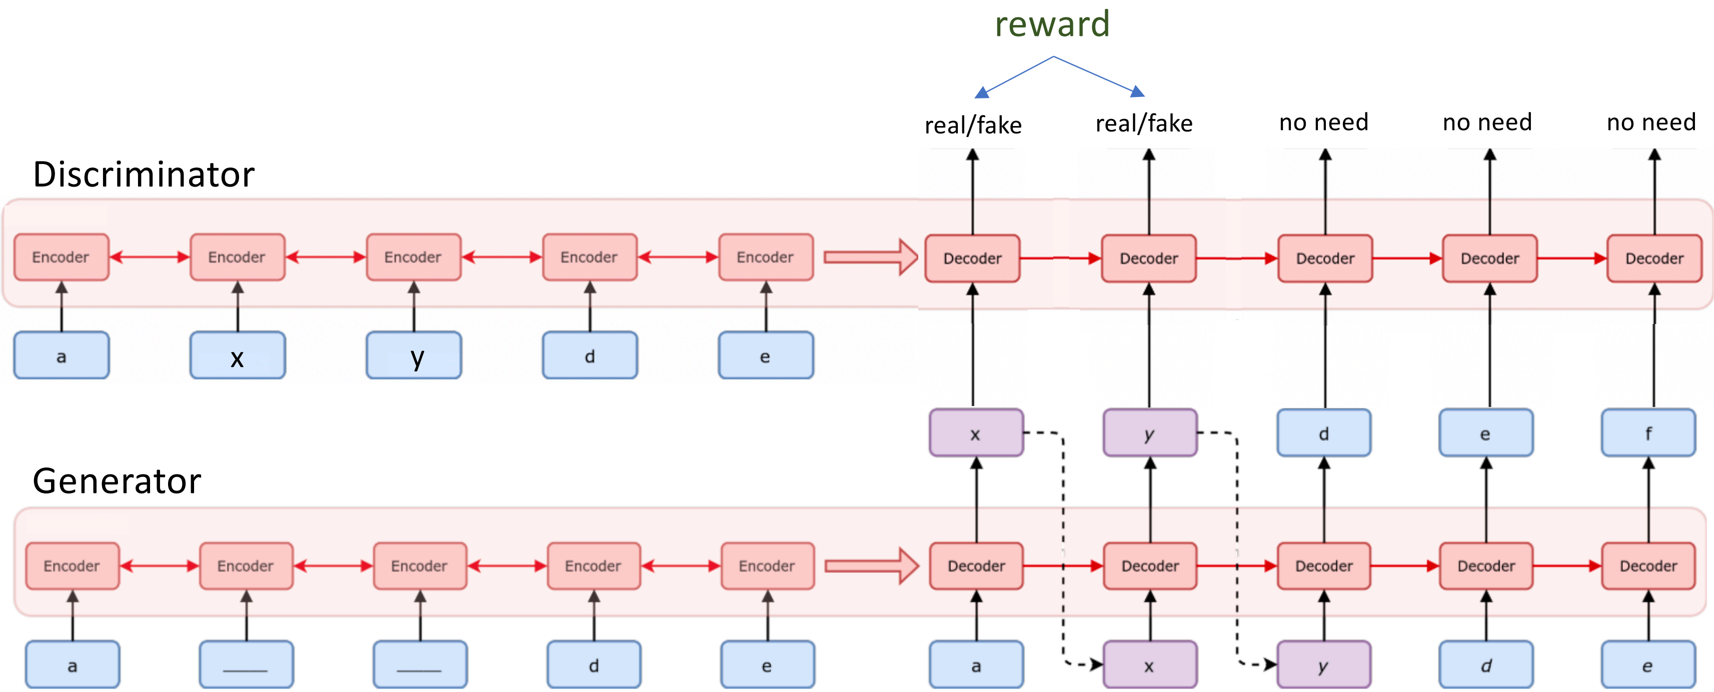

MaskGAN: Better Text Generation via Filling in the ____ . William Fedus, Ian Goodfellow, Andrew M. Dai. ICLR 2018

MaskGAN: Better Text Generation via Filling in the _____

- **Results**
  - Conditional Sentence Generation == Inpainting

Ground Truth	**Pitch Black was a complete shock to me when I first saw it back in 2000 In the previous years I**
MaskGAN	Pitch Black was a complete shock to me when I first saw it back in 1979 I was really looking forward
MaskMLE	Black was a complete shock to me when I first saw it back in 1969 I live in New Zealand

  - Unconditional Sentence Generation == All words are masked

MaskGAN	**Positive**: Follow the Good Earth movie linked Vacation is a comedy that credited against the modern day era yarns which has helpful something to the modern day s best It is an interesting drama based on a story of the famed

MaskGAN: Better Text Generation via Filling in the ____ . William Fedus, Ian Goodfellow, Andrew M. Dai. ICLR 2018

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

北京大學 PEKING UNIVERSITY

Adversarial Text Generation without Reinforcement Learning

- **Motivation**

  RL training is too slow …

- **Method: Stage 1**

  The generator is a common Seq2Seq

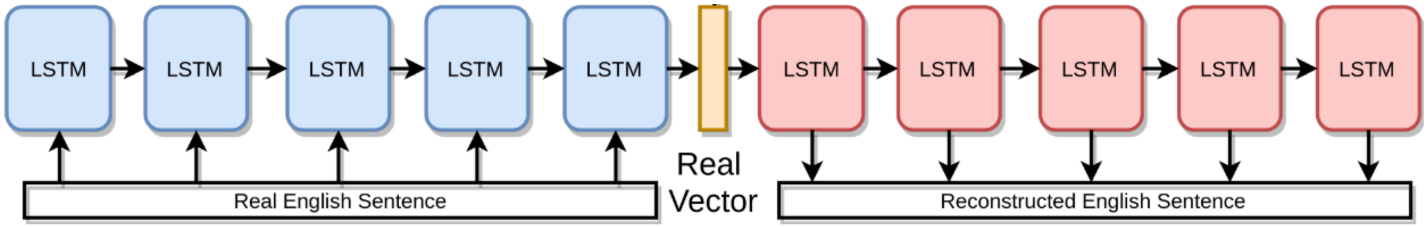

Adversarial Text Generation without Reinforcement Learning. Donahue, David. Rumshisky, Anna. 2018

Adversarial Text Generation without Reinforcement Learning

- **Method: Stage 2**

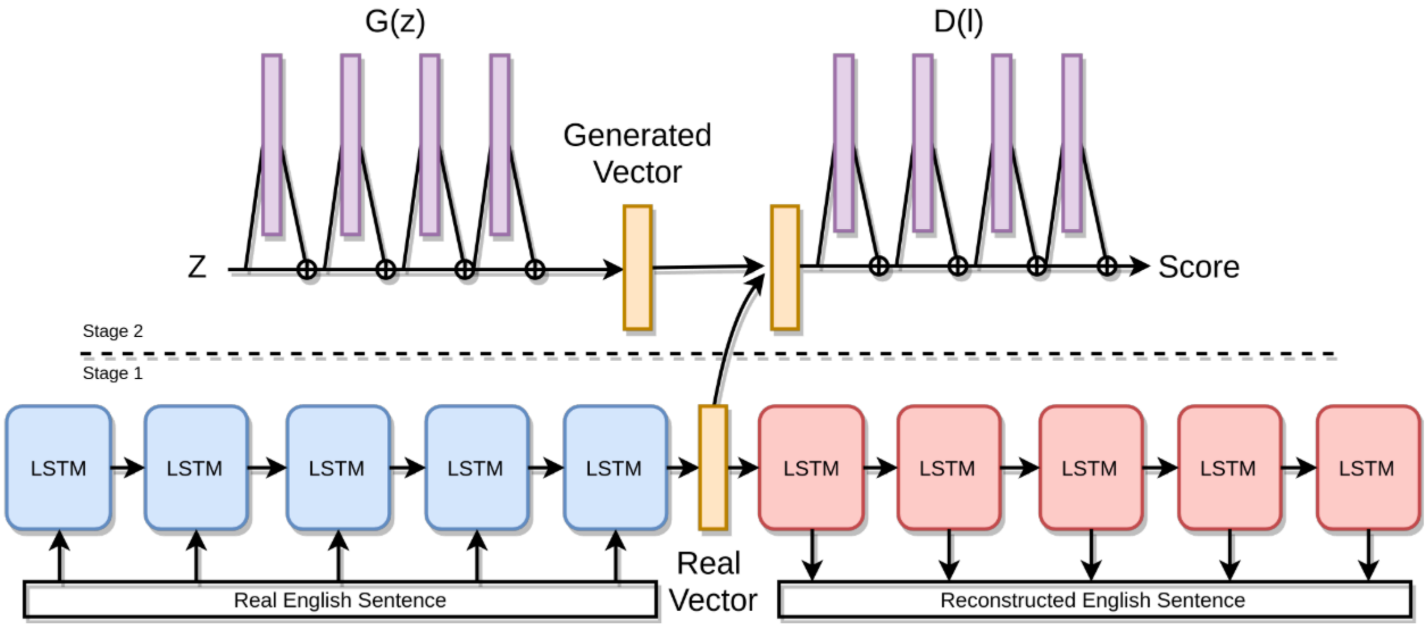

Adversarial Text Generation without Reinforcement Learning. Donahue, David. Rumshisky, Anna. 2018

- Introduction
  - RNN Language Modelling
  - Generating Sentences from a Continuous Space
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN
  - RankGAN
  - MaskGAN
- GANs
  - Adversarial Text Generation without Reinforcement Learning
  - GANs for Sequences Generation with the Gumbel-softmax

GANs for Sequence Generation with Gumbel-softmax

- **Motivation**
  Train like a normal GAN

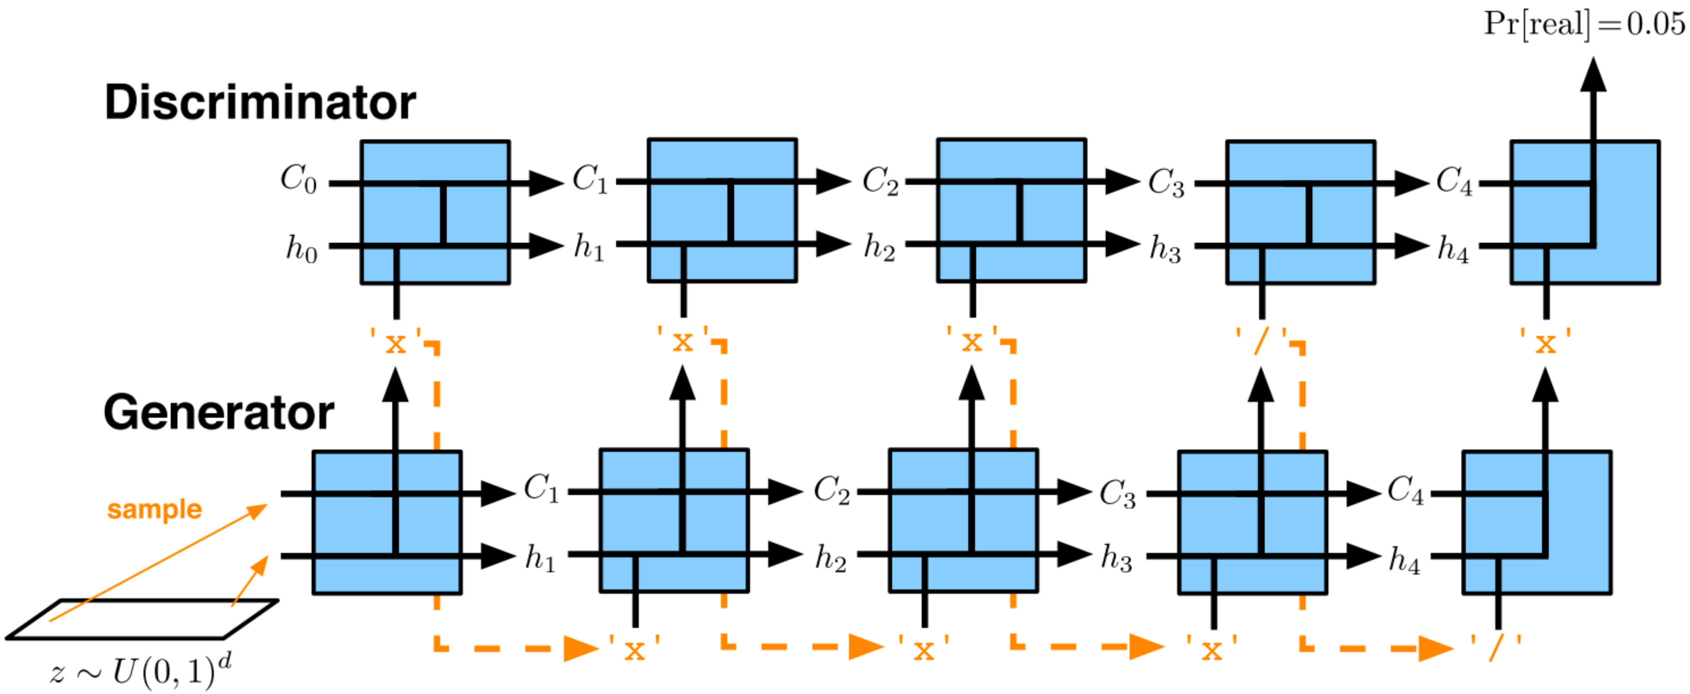

GANs for Sequences of Discrete Elements with the Gumbel-softmax Distribution.
Matt Kusner, Jose Miguel Hernandez-Lobato. NeurIPS. 2016.

GANs for Sequence Generation with Gumbel-softmax

PEKING UNIVERSITY

- **SeqGAN**

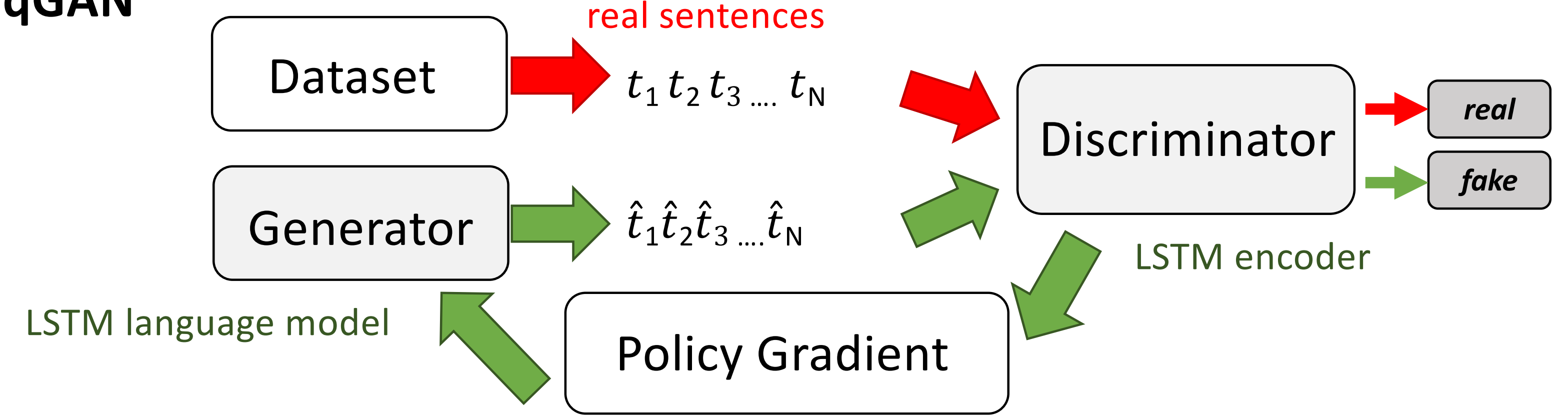

- **Gumbel-softmax GAN**

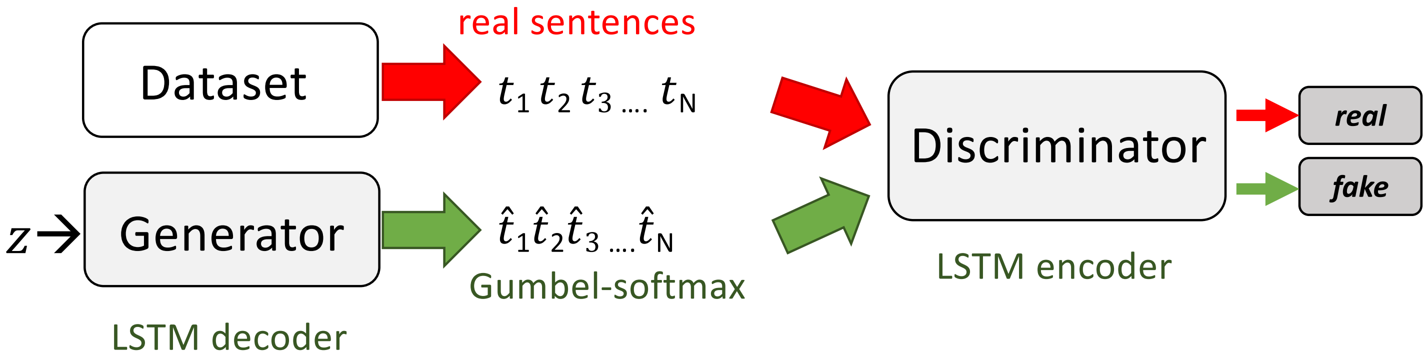

GANs for Sequences of Discrete Elements with the Gumbel-softmax Distribution.
Matt Kusner, Jose Miguel Hernandez-Lobato. NeurIPS. 2016.

GANs for Sequence Generation with Gumbel-softmax

- **Limitation**
  - Gumbel-softmax’s gradient is too large … poor performance

GANs for Sequences of Discrete Elements with the Gumbel-softmax Distribution.
Matt Kusner, Jose Miguel Hernandez-Lobato. NeurIPS. 2016.

Discreteness in Generative Models
– Discrete Sequence Generation

- Introduction
  - RNN Language Modelling --
  - Generating Sentences from a Continuous Space 2016
  - Recap: Inverse RL vs. GAN
- GAN+RL
  - SeqGAN 2017
  - RankGAN 2017
  - MaskGAN 2018
- GANs
  - Adversarial Text Generation without Reinforcement Learning 2018
  - GANs for Sequences Generation with the Gumbel-softmax 2016

Thanks