Neural-Network-based Dialog Agents: Going Beyond the Seq2seq Model

Thomas WOLF
thomas@huggingface.co

OUR TEAM

Google

Stanford University

We raised $5.2M in early 2018

From institutional investors

betaworks

SVAngel

a_capital

And private investors

HuggingFace is building a SOCIAL AI

It holds **long conversations,** it's **fun** and **cares** about you

It helps you build your **confidence**, get **better** and be **more social**

Topic of Today: Chit-Chat 😺

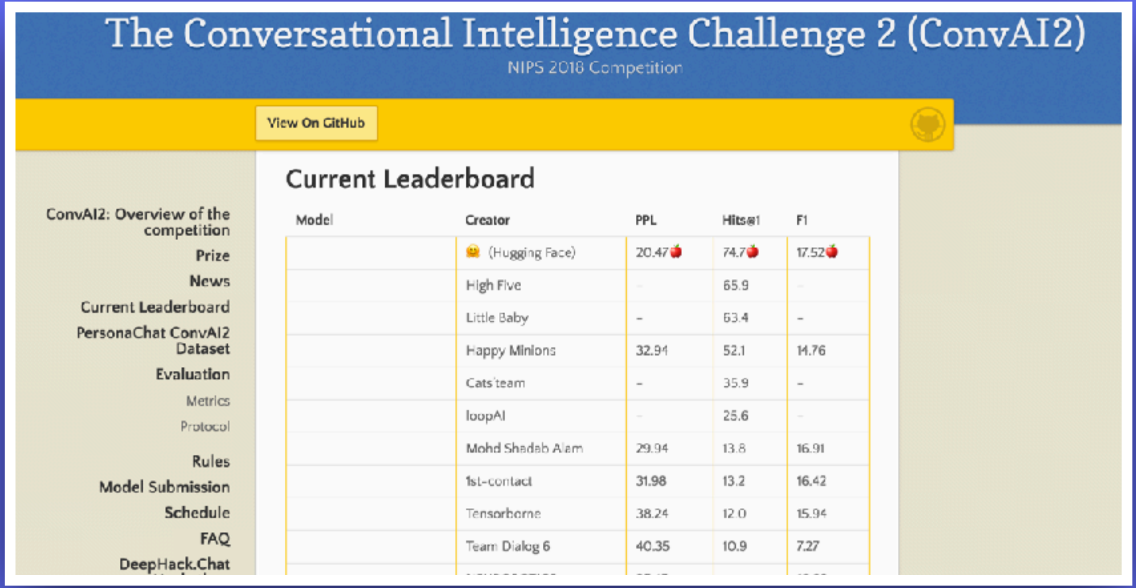

The Conversational Intelligence Challenge 2 (ConvAI2)
NIPS 2018 Competition

View On GitHub

ConvAI2: Overview of the competition
Prize
News
Current Leaderboard
PersonaChat ConvAI2 Dataset
Evaluation
Metrics
Protocol
Rules
Model Submission
Schedule
FAQ
DeepHack.Chat

Current Leaderboard

Model	Creator	PPL	Hits@1	F1
	🤗 (Hugging Face)	20.47🍎	74.7🍎	17.52🍎
	High Five	-	65.9	-
	Little Baby	-	63.4	-
	Happy Minions	32.94	52.1	14.76
	Cats'team	-	35.9	-
	loopAI	-	25.6	-
	Mohd Shadab Alam	29.94	13.8	16.91
	1st-contact	31.98	13.2	16.42
	Tensorborne	38.24	12.0	15.94
	Team Dialog 6	40.35	10.9	7.27

Open domain - open form: no limitation on topic, raw text input, no constrains
Short conversation: <10 turns
Small talk: shallow topics, not about question-answering, light memorization

Two Paradigms in Chit-Chat

- **Retrieval-based** model:
    - Take a dataset of dialogs **as big as you can.**
    - Compute a **similarity** between the current dialog and your dataset.
    - Send back the **closest** answer.

- **Generative model**:
    - A recent alternative mostly popularized by 2015's « A Neural Conversational Model » by Vinyals and Le
    - **Generate word after word** up until the utterance is deemed complete.
    - Typical architecture: an LSTM digests the inputs word by words and another LSTM generates the output from the last hidden state: « Seq2seq Model »

Conversation 4: Philosophical Q&A

Human: *what is the purpose of life ?*
Machine: *to serve the greater good .*
Human: *what is the purpose of living ?*
Machine: *to live forever .*
Human: *what is the purpose of existence ?*
Machine: *to find out what happens when we get to the planet earth .*
Human: *where are you now ?*
Machine: *i 'm in the middle of nowhere .*
Human: *what is the purpose of dying ?*
Machine: *to have a life .*

Strengths & Limitations

- Retrieval models:
  - Good grammaticality (replies are written by humans)
  - But: can’t adapt the response to the context
  - Problem of consistency: contradictory answers are often close to the same question
- Generative models:
  - Can dynamically adapt to the context
  - But: tend to default to unappealing answers like « yes/no/I don't know »
  - Grammatically is often an issue

- Rule of NIPS **Conversational Intelligence Challenge 2** (ConvAI2)
  - Condition the models on personalities to improve consistency

The Conversational Intelligence Challenge 2 (NIPS 2018 competition)

Persona 1	Persona 2
I like to ski	I am an artist
My wife does not like me anymore	I have four children
I have went to Mexico 4 times this year	I recently got a cat
I hate Mexican food	I enjoy walking for exercise
I like to eat cheetos	I love watching Game of Thrones

[PERSON 1:] Hi
[PERSON 2:] Hello ! How are you today ?
[PERSON 1:] I am good thank you , how are you.
[PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones.
[PERSON 1:] Nice ! How old are your children?
[PERSON 2:] I have four that range in age from 10 to 21. You?
[PERSON 1:] I do not have children at the moment.
[PERSON 2:] That just means you get to keep all the popcorn for yourself.
[PERSON 1:] And Cheetos at the moment!
[PERSON 2:] Good choice. Do you watch Game of Thrones?
[PERSON 1:] No, I do not have much time for TV.
[PERSON 2:] I usually spend my time painting: but, I love the show.

Example dialog from the PERSONA-CHAT dataset. Person 1 is given their own persona (top left) at the beginning of the chat, but does not know the persona of Person 2, and vice-versa. They have to get to know each other during the conversation.

Chit-chat with a human while keeping a coherent/predefined persona

Automatic Metrics

- **PPL** (perplexity) *How well the model can predict the successive words in a gold message (written by humans).*
  - **lower** is better
  - Scale: **Infinity – 0**
- **Hits@1** *Number of time the model select the gold next message between 20 possible message (the other 19 are random)*
  - **higher** is better.
  - Scale: **0 –100**
- **F1** *How many content words (nouns/verbs) does a message generated by your model share with a gold message*.
  - **higher** is better.
  - Scale: **0 –100**

These automatic metrics have issues – The best is always the human opinion

First submission results

(The 2nd is on its way with a +8 points improvement in Hits@1)

Model	Creator	PPL	Hits@1	F1
	🤗 (Hugging Face)	20.47🍎	74.7🍎	17.52🍎
	High Five	-	65.9	-
	Little Baby	-	63.4	-
	Happy Minions	32.94	52.1	14.76
	Cats'team	-	35.9	-
	loopAI	-	25.6	-
	Mohd Shadab Alam	29.94	13.8	16.91
	1st-contact	31.98	13.2	16.42
	Tensorborne	38.24	12.0	15.94
	Team Dialog 6	40.35	10.9	7.27
	NEUROBOTICS	35.47	-	16.68
	Sonic	33.46	-	16.67
topicSeq2seq	Team Pat	-	-	16.11
	Roboy	-	-	15.83
	Lost in Conversation	55.84	-	15.74
	flooders	-	-	15.47
	IamNotAdele	66.47	-	13.09
	Salty Fish	38.86	-	-
	Pinta	37.85	-	-
Seq2Seq + Attention	ParlAI team	29.8	12.6	16.18
Language Model	ParlAI team	46.0	-	15.02
KV Profile Memory	ParlAI team	-	55.2	11.9

Demo Time?

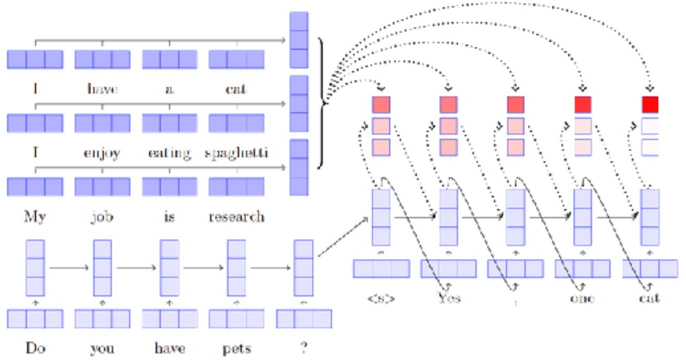

Figure 1: A diagram of the Profile Memory Network for generation. We also implemented a ranking version which has the same architecture except it ranks candidate sentences from the training set instead of generating, representing them using bag-of-word embeddings.

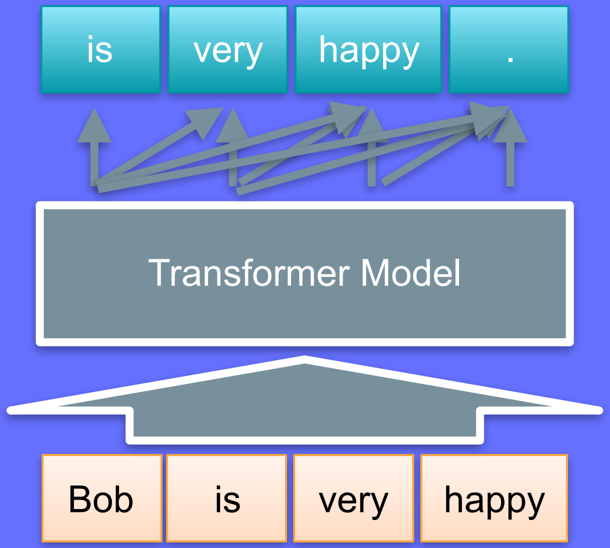

Baseline:
Memory-augmented **seq2seq**

Our model:
TransferTransfo

What’s the difference? 👻

Validation set (public) Leaderboard – Test set (hidden) Leaderboard

Model	Creator	PPL	Hits@1	F1
	🤗 (Hugging Face)	23.05🍎	74.3🍎	17.85🍎
	Team Pat	-	-	17.85
	Pinta	-	51.4	17.25
	Mohd Shadab Alam	35.57	14.8	16.94
	Sonic	38.87	-	16.88
	NEUROBOTICS	39.7	-	16.82
	Happy Minions	34.57	68.1	16.72
	1st-contact	36.54	13.3	16.58
	Tensorborne	44.64	12.1	16.13
	flooders	-	-	15.96
	Lost in Conversation	62.83	-	15.91
	High Five	59.83	78.2	15.34
	Little Baby	-	72.9	-
	loopAI	-	29.7	-
	Salty Fish	42.3	-	-

Model	Creator	PPL	Hits@1	F1
	🤗 (Hugging Face)	20.47🍎	74.7🍎	17.52🍎
	Little Baby	-	61.0	-
	Happy Minions	32.94	52.1	14.76
	High Five	52.8	50.3	13.73
	Pinta	-	44.4	16.52
	loopAI	-	25.6	-
	Mohd Shadab Alam	30.97	14.4	16.44
	1st-contact	31.98	13.2	16.42
	Tensorborne	38.24	12.0	15.94
	Team Dialog 6	40.35	10.9	7.27
	NEUROBOTICS	35.47	-	16.68
	Sonic	33.46	-	16.67
	Lost in Conversation	55.84	-	15.74
	flooders	-	-	15.47
	Team Pat	-	-	13.23
	Salty Fish	45.87	-	-
Seq2Seq + Attention	ParlAI team	29.8	12.6	16.18
Language Model	ParlAI team	46.0	-	15.02
KV Profile Memory	ParlAI team	-	55.2	11.9

- Small dataset =>
  - Large models are overfitting
  - Small models are underfitting

Transfer Learning 🦄

TransferTransfo is a Transformer Model…

- Take a fixed-length sequence of « words » as input and outputs a sequence of the same length
- Each output is a probability distribution over the next words
  => a probability for each word in the vocabulary

- Inside:
  - Self-attention:
    MLP compute Key-Value-Query for each element of the sequence
  - Causal masking in the attention heads to only attend to the past.

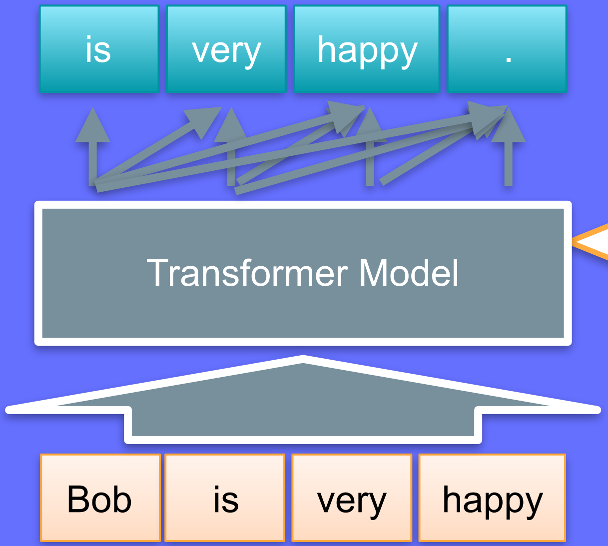

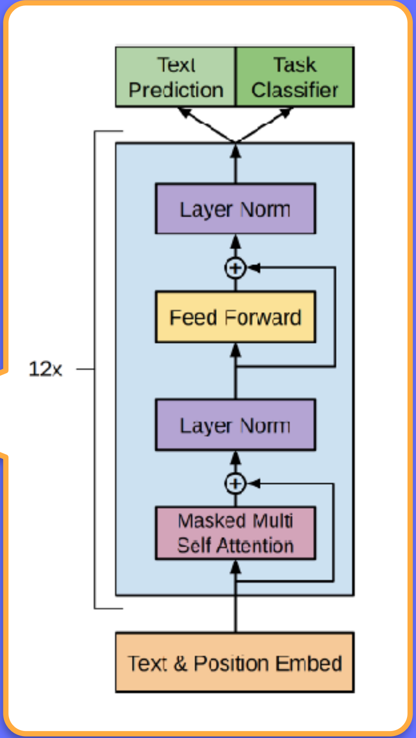

…Trained using Transfer Learning & Multi-Task Learning

- The competition dataset (PERSONA-CHAT) is one of the biggest multi-turn dialog dataset (10k conversations, about 100k turns) but it is still quite small in term of requirement for deep learning tools:
  - ex: Billion Words dataset has 1B words, CoNLL 2012 used for training co-reference systems is ~1M sentences long.
- An engaging open-domain dialogue is a lot more than just topic-coherence and dialogue-flow!
  - Need have common-sense, short term memory, co-reference resolution, sentimental analysis, textual entailment,…
  - Hard to learn all these from such a small dataset.
- Solution:
  - Use transfert learning and multi-task learning

Transfer Learning 🦄

- Train the model on the language modeling task on a large dataset.
    - Language modeling: **given a seq of words, learn to predict the next word**
    - This task force the model to learn many aspects of language including high-level ones like common-sense knowledge.
    - Recently shown to improve many downstream NLP tasks and in particular commonsense reasoning / co-reference resolution:
        - *A Simple Method for Commonsense Reasoning* by Trinh & Le (2018)
        - *Improving Language Understanding by Generative Pre-Training* by Radford et al. (2018)
        - *Universal Language Model Fine-tuning for Text Classification* by Howard and Ruder (2018)
    - Our model is derived from the model of Radford et al. pre-trained on the Toronto Book dataset (7k books).

Encoding a Dialog 🤖

Learning Dialog Flow and Persona

- Now that we have a model with basic common-sense and co-reference capabilities, we need to teach it the specificities of dialog:
  - Alternating utterances - basic Theory of Mind concepts
  - Dialog flow (« speech/dialog acts »)
  - Conditioning on a provided personality

- How to build a sequential inputs for our model from a conditioned dialog?
  - Unlike RNN/LSTM, Transformers don't possess a natural notion of sequentiality and position
  - We need to add positional embeddings to incorporate sequentiality

I	like	to	ski	Hello	!	How	are	you	today	?	I	am	good	thank	you	Word embeddings
																Positional embeddings

Encoding a Dialog and a Persona

- Now that we have a model with basic common-sense and co-reference capabilities, we need to teach it the specificities of dialog:
  - Alternating utterances - basic Theory of Mind concepts
  - Dialog flow (« speech/dialog acts »)
  - Conditioning on a provided personality

- How to build a sequential inputs for our model from a conditioned dialog?
  - Unlike RNN/LSTM, Transformers don't possess a natural notion of sequentiality and position
  - We need to add positional embeddings to incorporate sequentiality
  - We add special embeddings related to utterances and personas

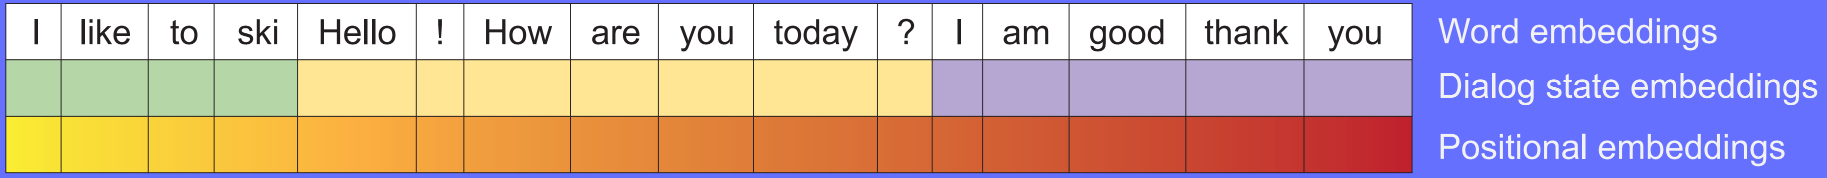

Encoding a Dialog and a Persona

- We can play with these embeddings to manipulate the notion of a sequence

Repeating specific embeddings to control positioning information

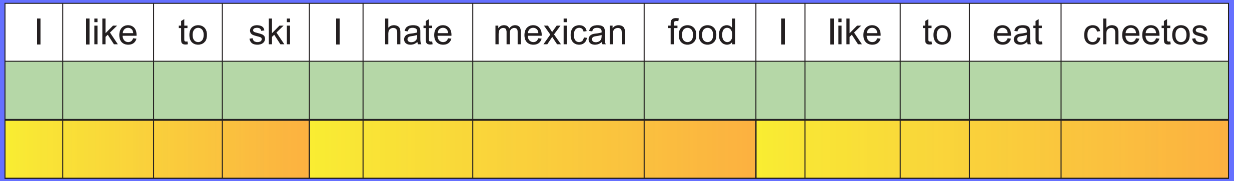

- We can also augment the dataset to bias towards positional invariance

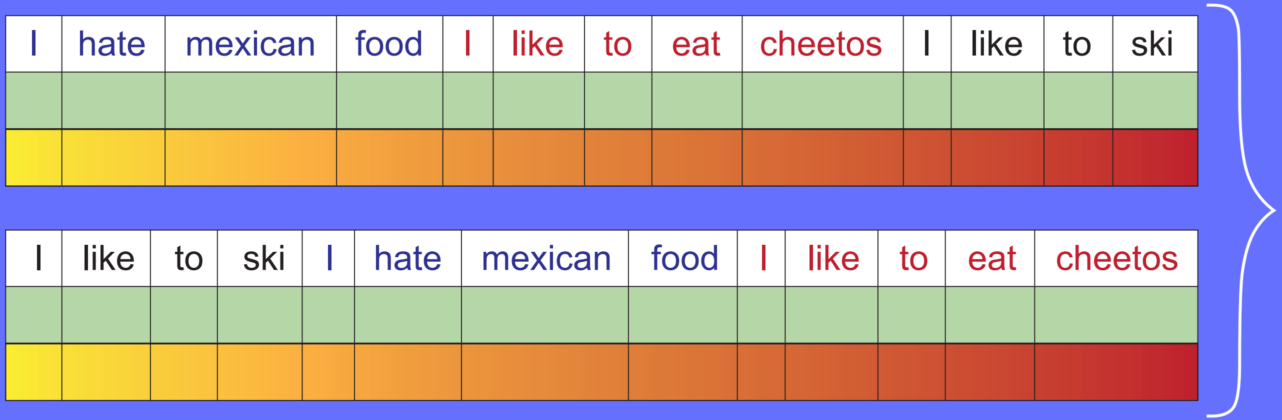

Permutation augmented dataset to bias towards positional invariance

Learning Dialog Flow 🧑‍🎓

Semantic Learning on Dialog Utterances

- Learning to distinguish a real answer from a distractor.

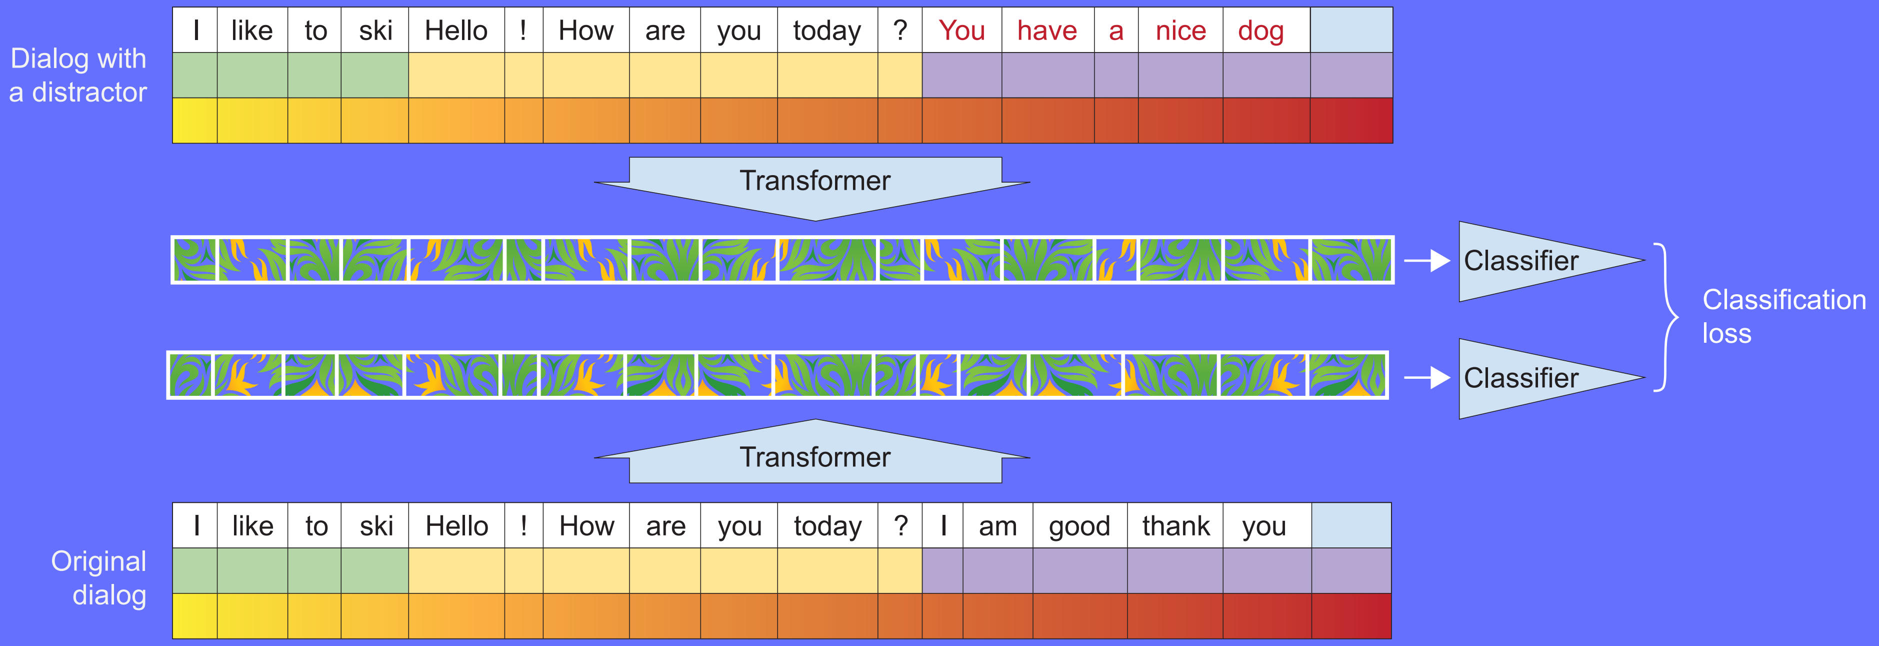

Can be combined with language modeling fine-tuning in a multi-task fashion

Talking to humans 🎤

Decoding — ☞ *How are you?*

- **Beam Search**
  - We create a message word by word.
  - Each time there could be several possibilities.
  - We keep a beam of possible answers and drop the lowest ones at each step.
  - We accumulate a pool of answers and select the one with the highest normalized probability
  - N-Grams filtering (competition rules forbid to repeat persona sentences)

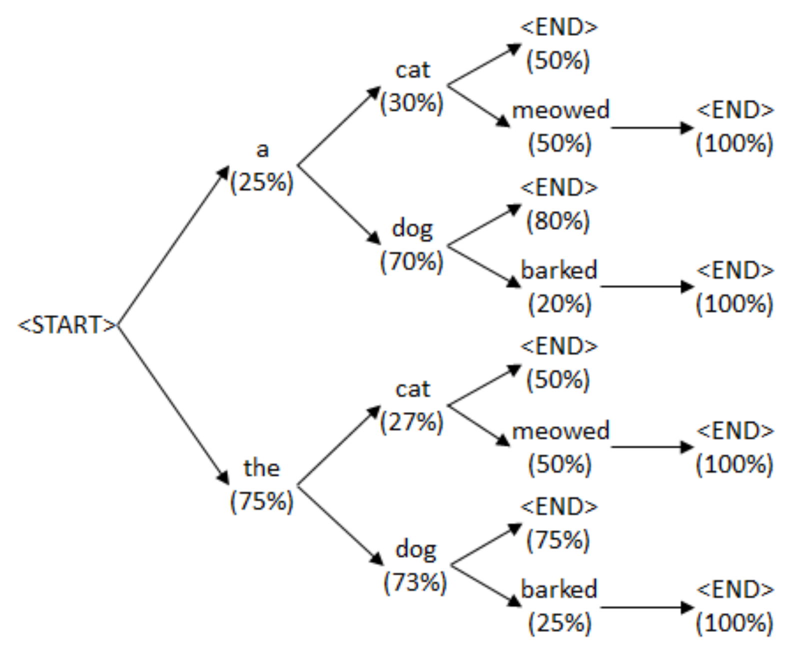

That's it for today
Thanks for listening!

Thomas WOLF
thomas@huggingface.co