CS769 Advanced NLP

Instruction Tuning and Multi-task Learning

Junjie Hu

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Slides adapted from Graham

https://junjiehu.github.io/cs769-spring25/

Goals for Today

- NLP Tasks
- Multi-task Learning
- Instruction Tuning

Plethora of Tasks in NLP

- In NLP, there are a plethora of tasks, each requiring different varieties of data
  - **Only text:** e.g. language modeling
  - **Naturally occurring data:** e.g. machine translation
  - **Hand-labeled data:** e.g. most analysis tasks
- And each in many languages, many domains!

Standard Multi-task Learning

- Train representations to do well on multiple tasks at once

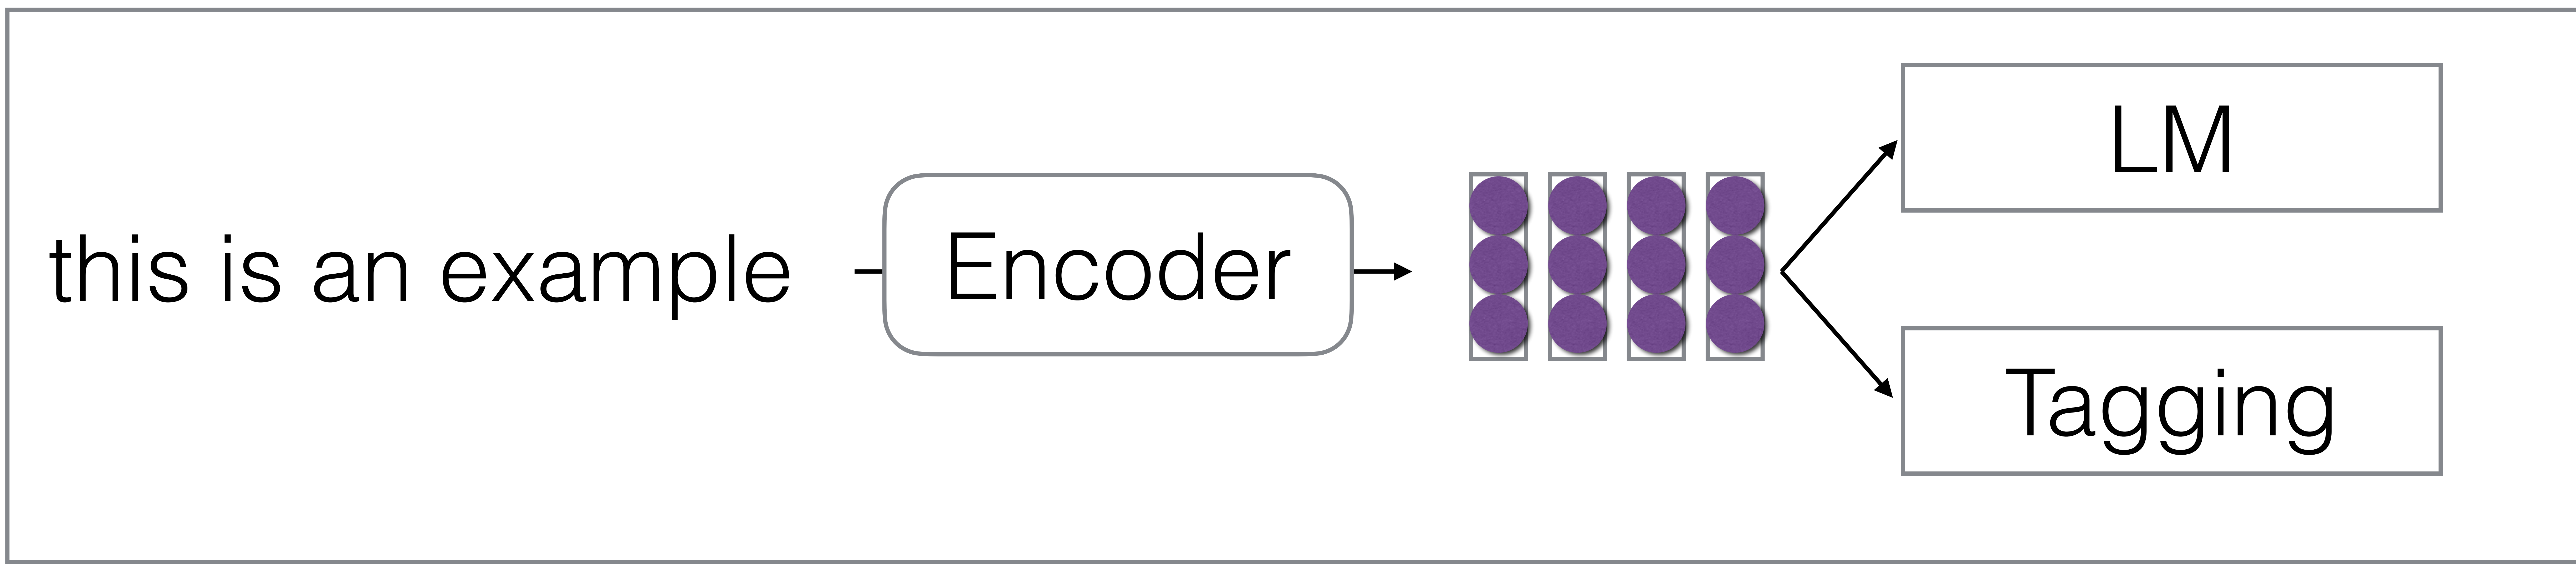

- Often as simple as randomly choosing minibatch from one of multiple tasks

Pre-train and Fine-Tune

- First train on one task, then train on another

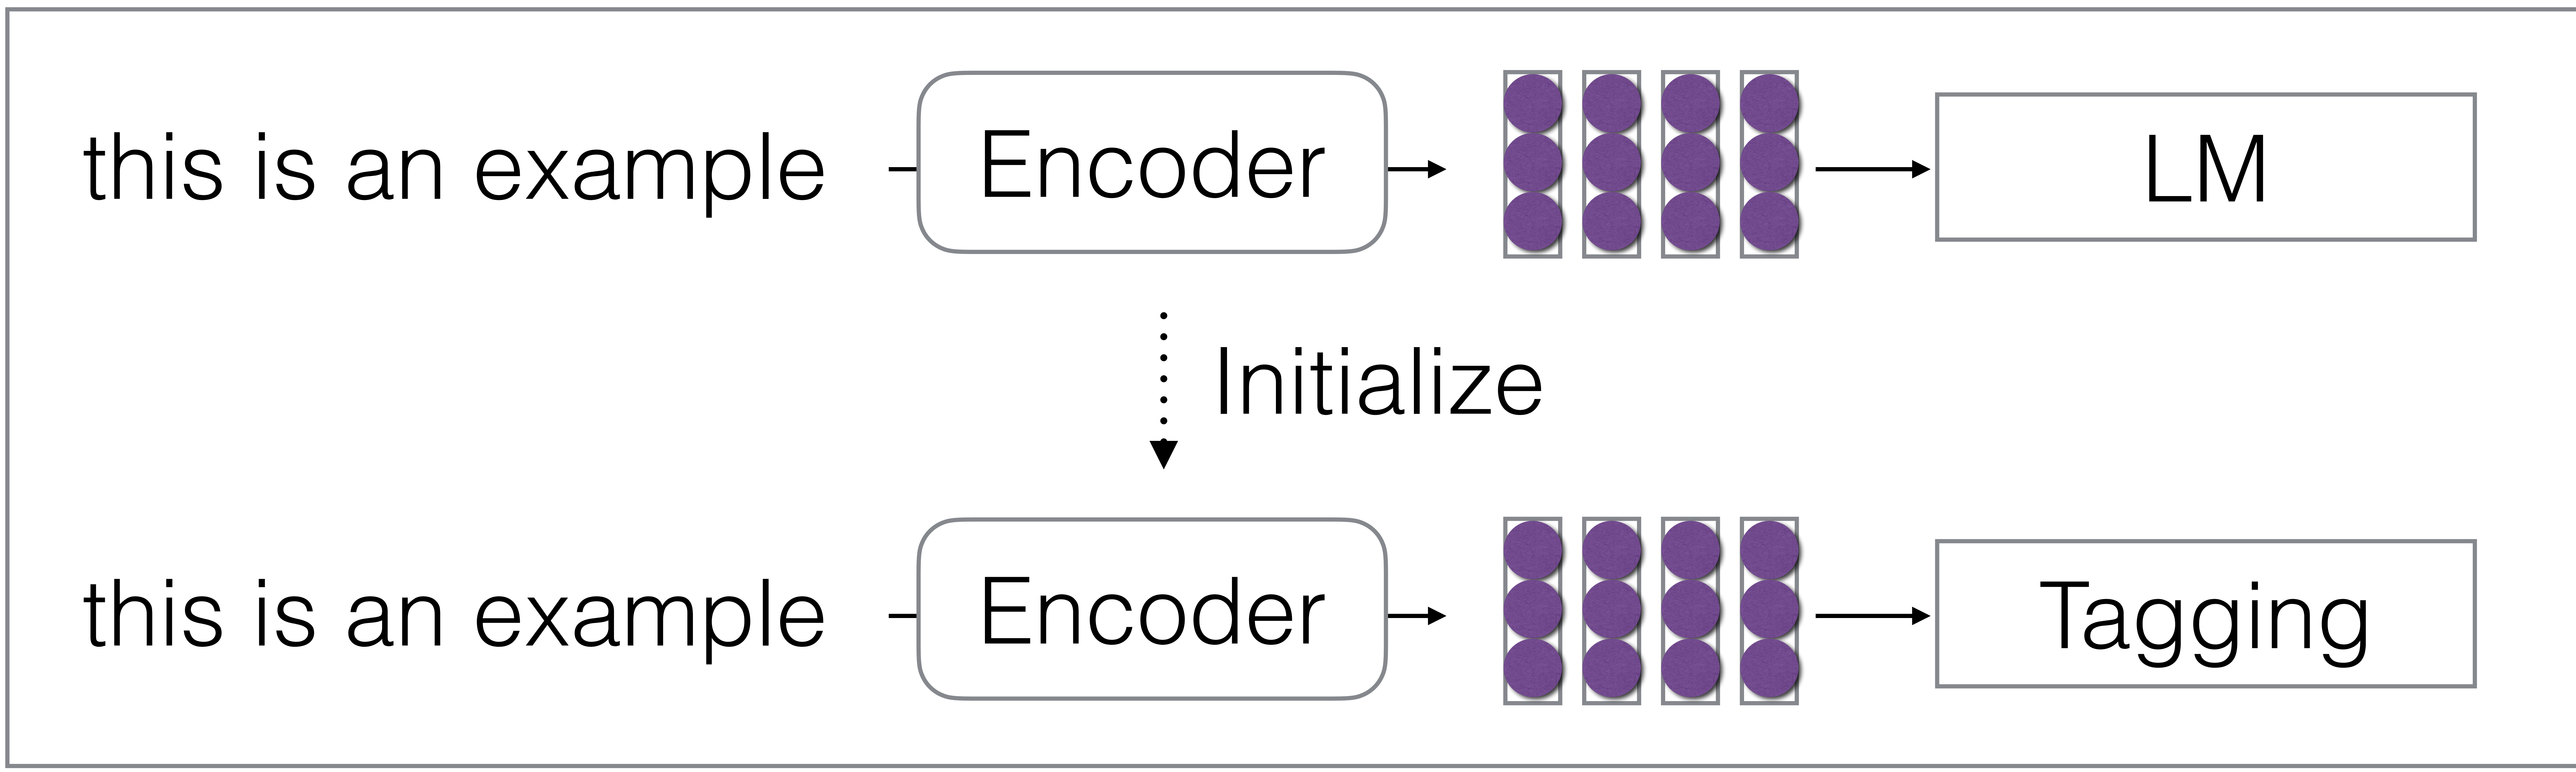

Prompting

- Train on LM task, make predictions in textualized tasks

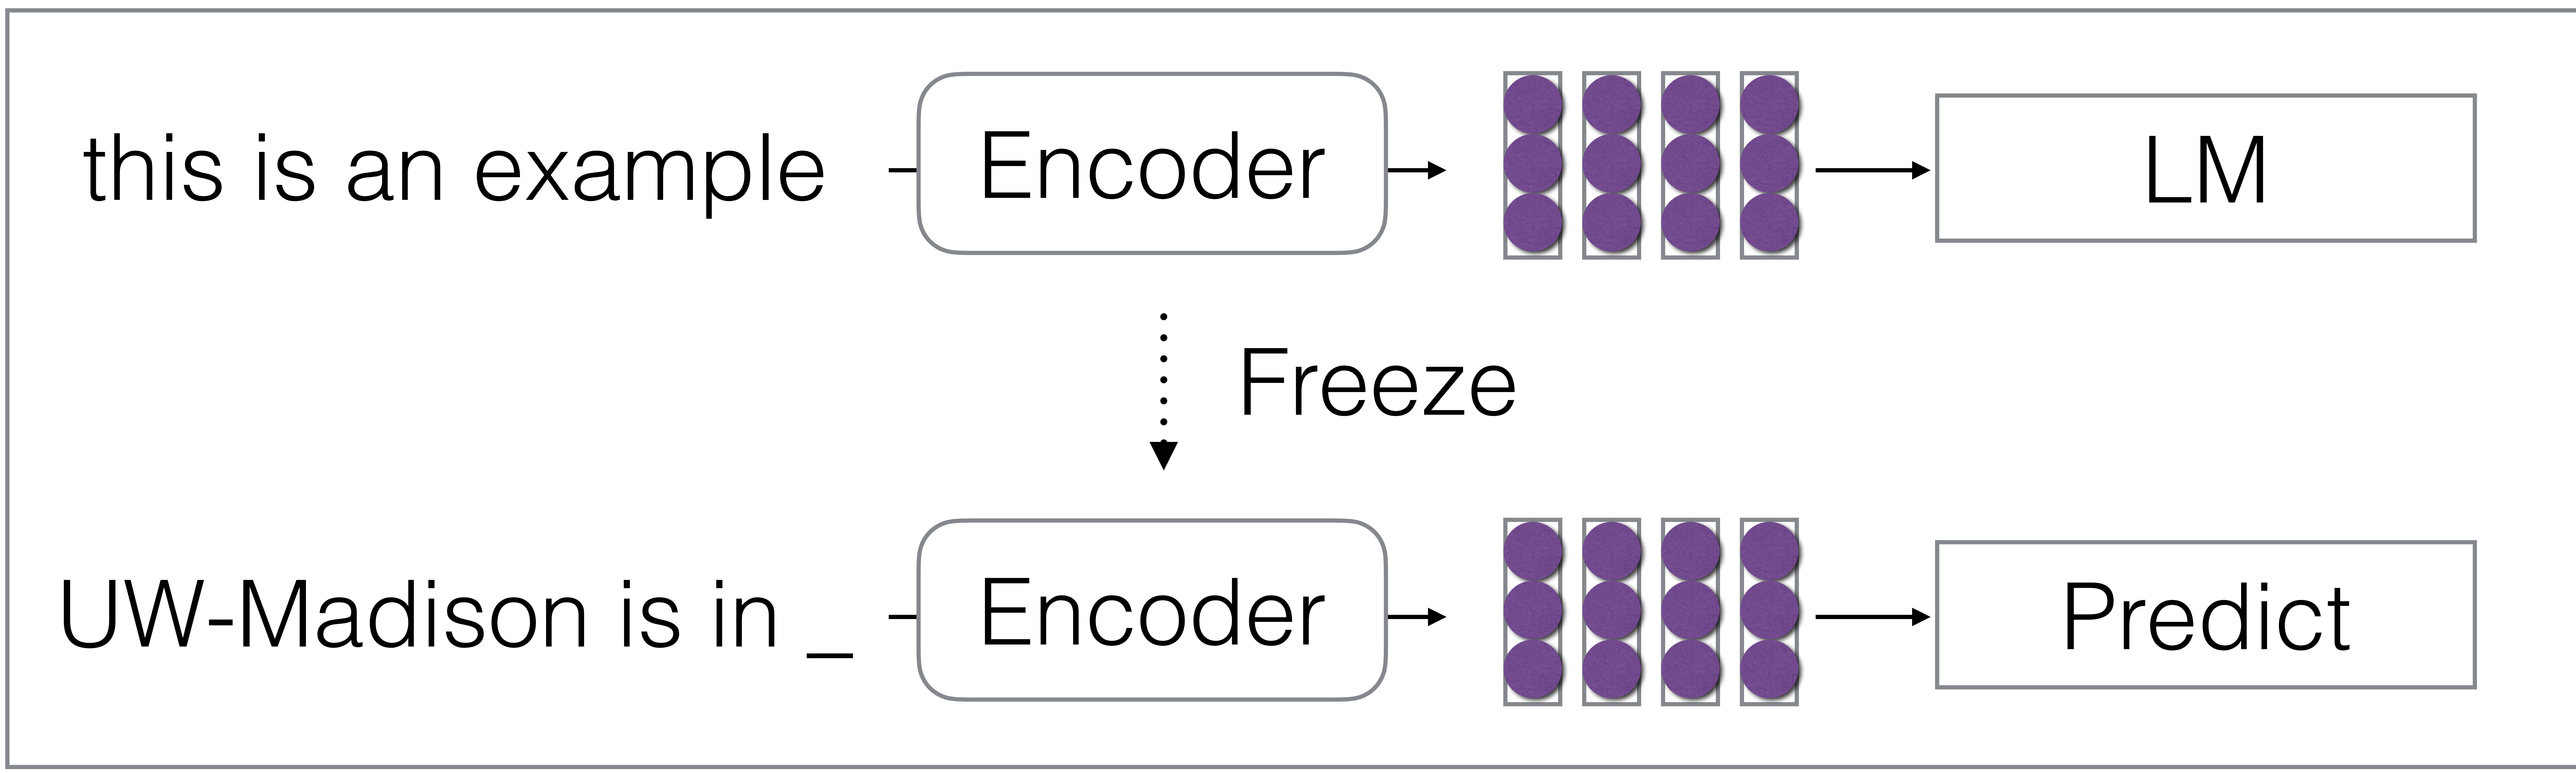

Instruction Tuning

- Pre-train, then fine-tune on many different tasks, with an instruction specifying the task

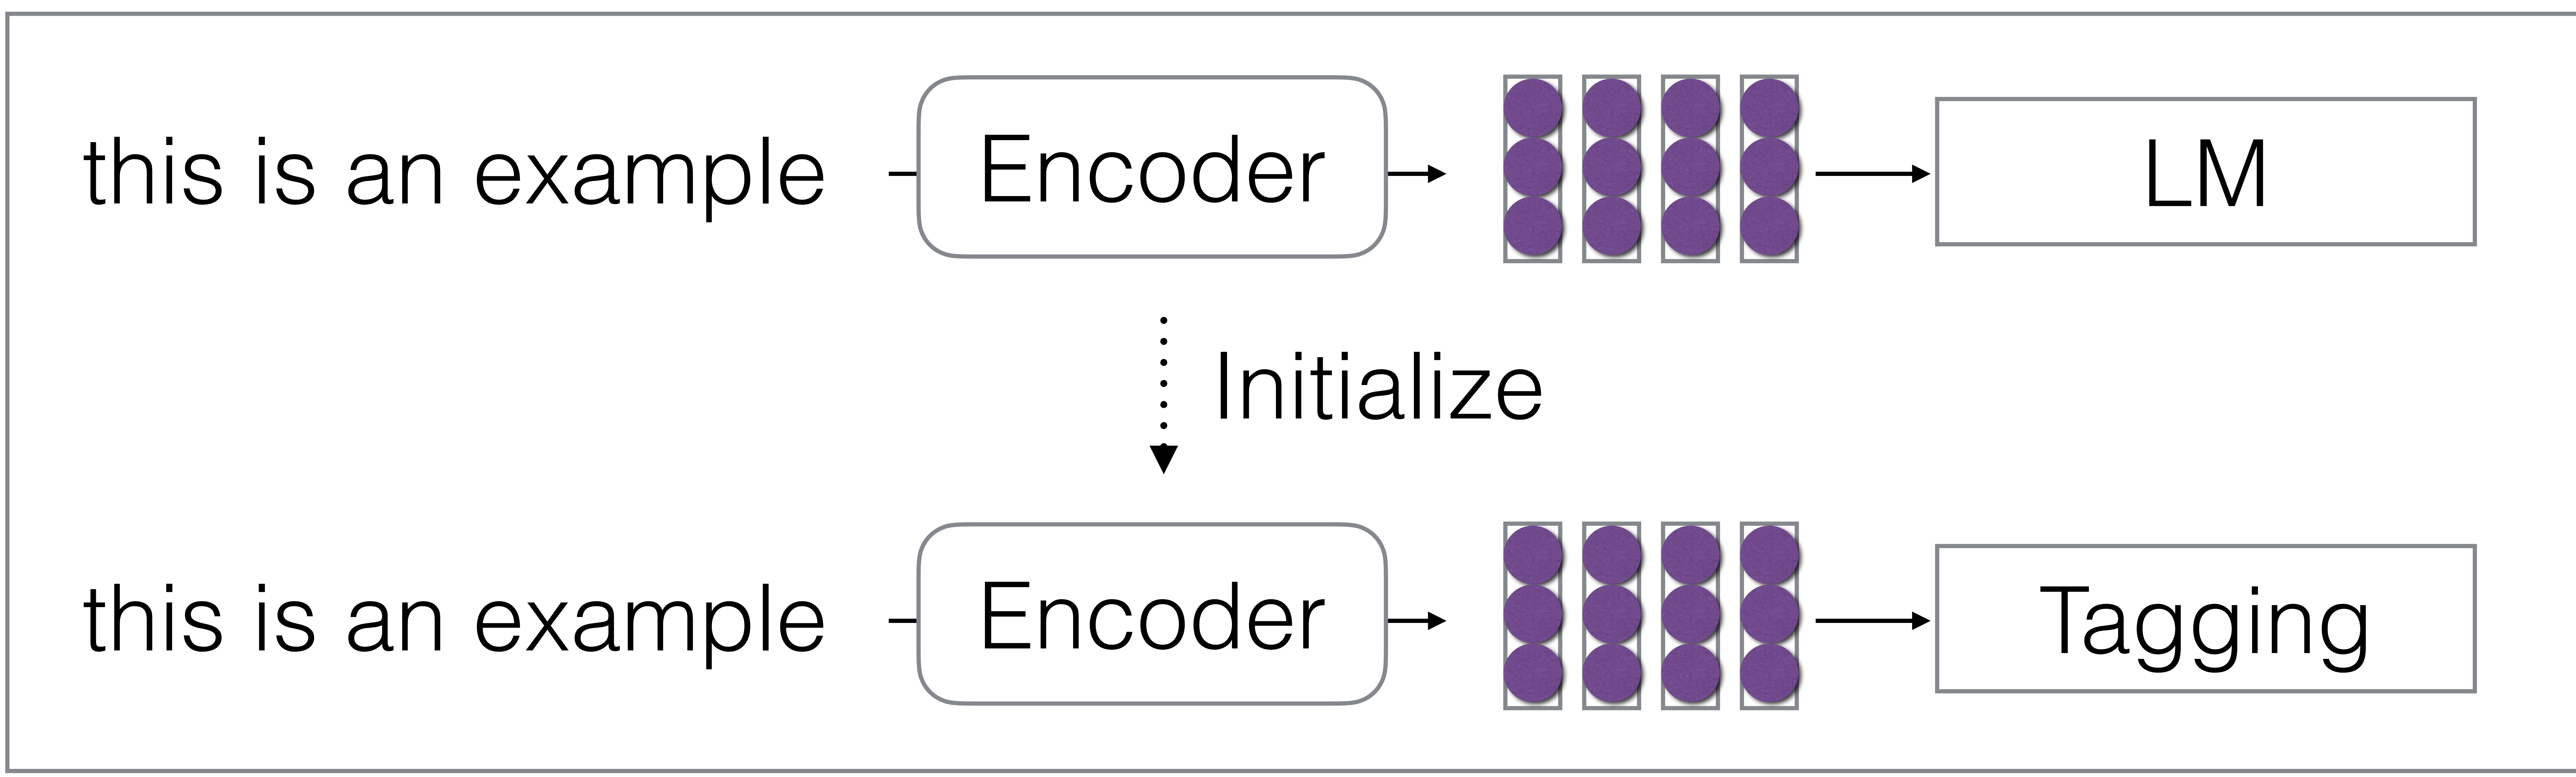

NLP Tasks

Approaches to Model Construction

- **Basic Fine Tuning:** Build a model that is good at performing a single task
- **Instruction Tuning:** Build a generalist model that is good at many tasks
- Even if we build a generalist model, we need to have an idea about what tasks we want it to be good at!

Context-free Question Answering

- Also called “open-book QA”
- Answer a question without any specific grounding into documents
- Example dataset: MMLU (Hendrycks et al. 2020)

Professional Law

As Seller, an encyclopedia salesman, approached the grounds on which Hermit's house was situated, he saw a sign that said, "No salesmen. Trespassers will be prosecuted. Proceed at your own risk." Although Seller had not been invited to enter, he ignored the sign and drove up the driveway toward the house. As he rounded a curve, a powerful explosive charge buried in the driveway exploded, and Seller was injured. Can Seller recover damages from Hermit for his injuries?

(A) Yes, unless Hermit, when he planted the charge, intended only to deter, not harm, intruders.
(B) Yes, if Hermit was responsible for the explosive charge under the driveway.
(C) No, because Seller ignored the sign, which warned him against proceeding further.
(D) No, if Hermit reasonably feared that intruders would come and harm him or his family.

Contextual Question Answering

- Also called “machine reading”, “closed-book QA”
- Answer a question about a document or document collection
- *Example:* Natural Questions (Kwiatkowski et al. 2019) is grounded in a Wikipedia document, or the Wikipedia document collection

Question: what color was john wilkes booth’s hair
Wikipedia Page: John_Wilkes_Booth
Long answer: Some critics called Booth “the handsomest man in America” and a “natural genius”, and noted his having an “astonishing memory”; others were mixed in their estimation of his acting. He stood 5 feet 8 inches (1.73 m) tall, had jet-black hair , and was lean and athletic. Noted Civil War reporter George Alfred Townsend described him as a “muscular, perfect man” with “curling hair, like a Corinthian capital”.

Short answer: jet-black

Code Generation

- Generate code (e.g. Python, SQL, etc.) from a natural language command and/or input+output examples
- *Example:* HumanEval (Chen et al. 2021) has evaluation questions for Python standard library

```python
def incr_list(l: list):
    """Return list with elements incremented by 1.
    >>> incr_list([1, 2, 3])
    [2, 3, 4]
    >>> incr_list([5, 3, 5, 2, 3, 3, 9, 0, 123])
    [6, 4, 6, 3, 4, 4, 10, 1, 124]
    """
    return [i + 1 for i in l]
```

Summarization

- Single-document: Compress a longer document to shorter
- Multi-document: Compress multiple documents into one
- Example: WikiSum compresses the references in a Wikipedia article into the first paragraph

References

1. ^ "Barack Hussein Obama Takes The Oath Of Office" on YouTube. January 20, 2009.
2. ^ "American Presidents: Greatest and Worst – Siena College Research Institute". Archived from the original on July 15, 2022. Retrieved February 12, 2023.
3. ^ "Barack Obama | C-SPAN Survey on Presidents 2017". Archived from the original on February 12, 2023. Retrieved February 12, 2023.
4. ^ "Siena's 6th Presidential Expert Poll 1982–2018 – Siena College Research Institute". Archived from the original on July 19, 2019. Retrieved February 13, 2023.
5. ^ "President Barack Obama". The White House. 2008. Archived from the original on October 26, 2009. Retrieved December 12, 2008.
6. ^ "President Obama's Long Form Birth Certificate". *whitehouse.gov*. April 27, 2011. Archived from the original on July 31, 2023. Retrieved August 4, 2023.
7. ^ "Certificate of Live Birth: Barack Hussein Obama II, August 4, 1961, 7:24 pm, Honolulu" (PDF). *whitehouse.gov*. April 27, 2011. Archived from the original (PDF) on March 3, 2017. Retrieved March 11, 2017 – via National Archives.

Barack Obama

Article Talk

From Wikipedia, the free encyclopedia

"Barack" and "Obama" redirect here. For other uses, see Barack (disambiguation), Obama (disambiguati

Barack Hussein Obama II (/bəˈrɑːk huːˈseɪn oʊˈbɑːmə/ *bə-RAHK hoo-SAYN oh-BAH-mə*;[1] born August 4, 1961) is an American politician who served as the 44th president of the United States from 2009 to 2017. A member of the Democratic Party, he was the first African-American president in U.S. history. Obama previously served as a U.S. senator representing Illinois from 2005 to 2008, as an Illinois state senator from 1997 to 2004, and as a civil rights lawyer and university lecturer.

Obama was born in Honolulu, Hawaii. He graduated from Columbia University in 1983 with a B.A. in political science and later worked as a community organizer in Chicago. In 1988, Obama enrolled in Harvard Law School, where he was the first black president of the *Harvard Law Review*. He became a civil rights attorney and an academic, teaching constitutional law at the University of Chicago Law School from 1992 to 2004. He also went into elective politics. Obama represented the 13th district in the Illinois Senate from 1997 until 2004, when he successfully ran for the U.S. Senate. In 2008, after a close primary campaign against Hillary Clinton, he was nominated by the Democratic Party for president and chose Delaware Senator Joe Biden as his running mate. Obama was elected president, defeating Republican Party nominee John McCain in the presidential election and was inaugurated on January 20, 2009. Nine months later he was named the 2009 Nobel Peace Prize laureate, a decision that drew a mixture of praise and criticism.

Information Extraction

- *Entity recognition:* identify which words are entities
- *Entity linking:* link entities to a knowledge base (e.g. Wikipedia)
- *Entity co-reference:* find which entities in an input correspond to each-other
- *Event recognition/linking/co-reference:* identify what events occurred
- Example: OntoNotes (Weischedel et al. 2013) annotates many types of information like this on various domains

Translation

- Translate from one language to another
- Quality assessment done using similarity to reference translation
- Example: FLORES dataset (Goyal et al. 2021) — translations of Wikipedia articles into 101 languages

“General Purpose” Benchmarks

- Try to test language abilities across a broad range of tasks
- Example: BIGBench (Srivatsava et al. 2022)

tracking_shuffled_objects_three_objects_0

Alice, Bob, and Claire are friends and avid readers who occasionally trade books. At the start of the semester, they each buy one new book: Alice gets Ulysses, Bob gets Frankenstein, and Claire gets Lolita.
As the semester proceeds, they start trading around the new books. First, Claire and Bob swap books. Then, Bob and Alice swap books. Finally, Claire and Bob swap books. At the end of the semester, Bob has
Options:
(A) Ulysses
(B) Frankenstein
(C) Lolita

label
(B)

date_understanding_0

Today is Christmas Eve of 1937. What is the date tomorrow in MM/DD/YYYY?
Options:
(A) 12/11/1937
(B) 12/25/1937
(C) 01/04/1938
(D) 12/04/1937
(E) 12/25/2006
(F) 07/25/1937

label
(B)

web_of_lies_0

Question: Sherrie tells the truth. Vernell says Sherrie tells the truth. Alexis says Vernell lies. Michaela says Alexis tells the truth. Elanor says Michaela tells the truth. Does Elanor tell the truth?

label
No

Earlier Work on Multi-task Learning in NLP

Applications of Multi-task Learning

- Perform multi-tasking when one of your two tasks has fewer data

- **Plain text → labeled text**
  (e.g. LM -> parser)

- **General domain → specific domain**
  (e.g. web text → medical text)

- **High-resourced language → low-resourced language**
  (e.g. English → Telugu)

Domains in NLP

- One task, but incoming data could be from very different distributions

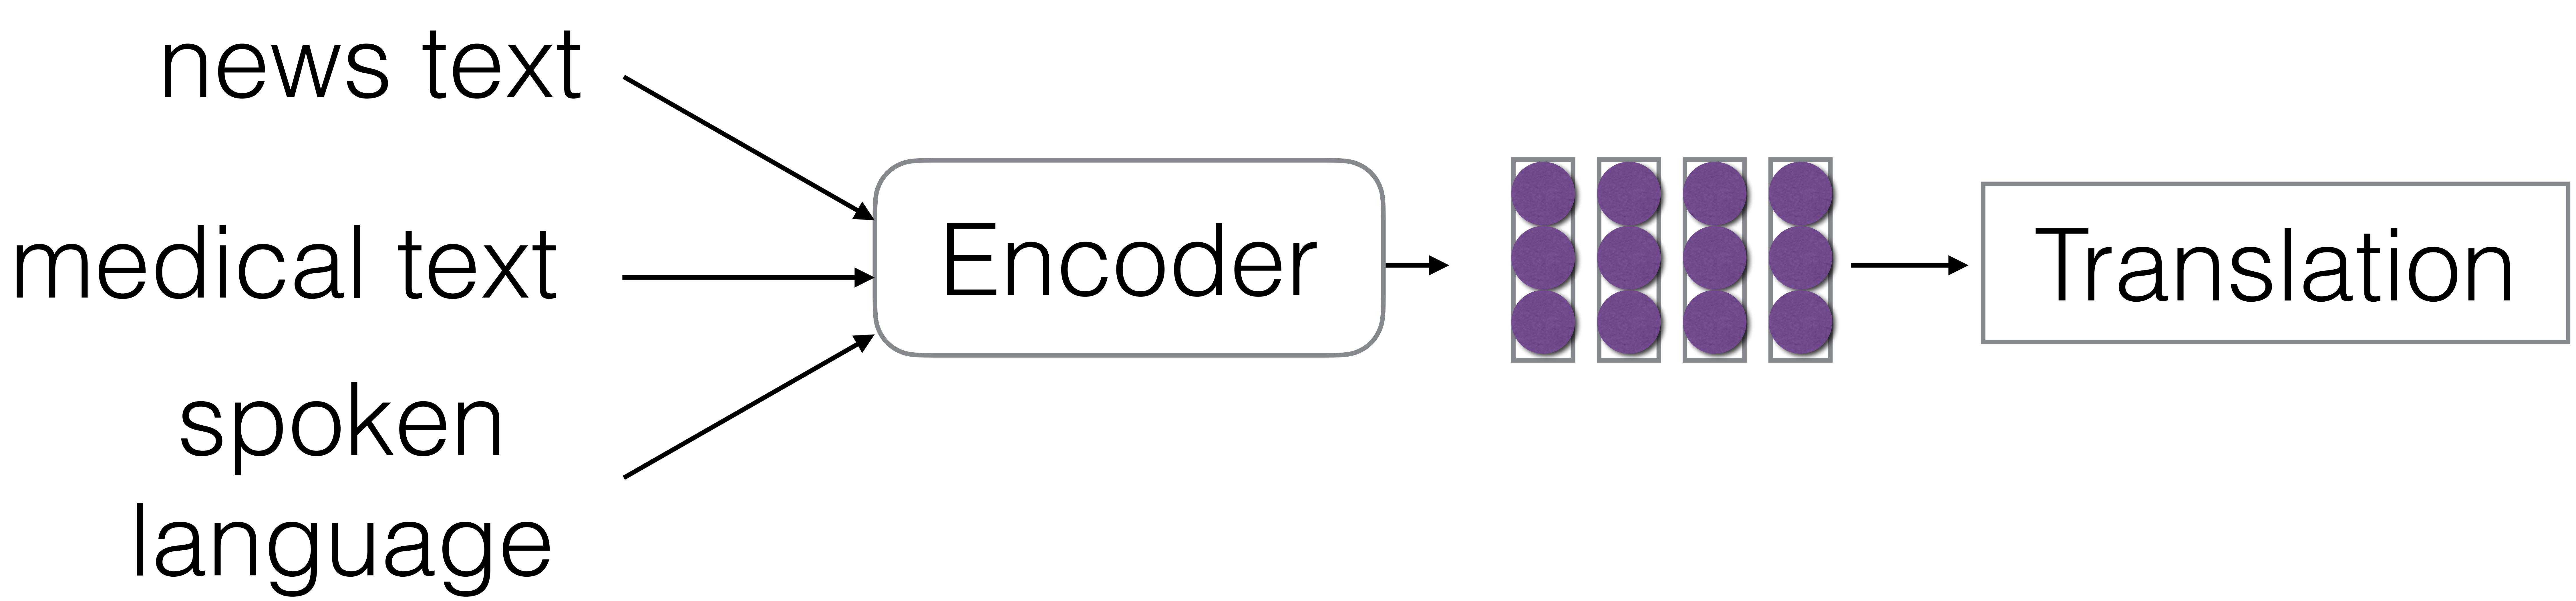

- Sometimes domains are labeled, sometimes they are not

What's in a "Domain"

(Stewart 2019)

- Mathematically, joint distribution over inputs and outputs differs over domains 1 and 2

$$P_{d1}(X, Y) \neq P_{d2}(X, Y)$$

- In practice:
  - **Content**, what is being discussed
  - **Style**, the way in which it is being discussed
  - **Labeling Standards,** the way that the same data is labeled

Types of Domain Shift

- **Covariate Shift:** The input changes but not the labeling

$$P_{d1}(X) \neq P_{d2}(X) \qquad P_{d1}(Y|X) = P_{d2}(Y|X)$$

- **Concept Shift:** The conditional distribution of labels changes (e.g. different labeling standards)

$$P_{d1}(X) = P_{d2}(X) \qquad P_{d1}(Y|X) \neq P_{d2}(Y|X)$$

- **Label Shift:** The output changes (which also implies the input changes).

$$P_{d1}(Y) \neq P_{d2}(Y)$$

Out of Distribution/Domain (OOD)

- **Generalization**
  - **Domain adaptation**: train on many domains, adapt to a target domain at testing
  - **Domain robustness**: train on many domains, perform well on all domains (esp. minority domains)
- **Detection**
  - Binary classification: detect whether a test example is an OOD example or not.

Ramponi et al Neural Unsupervised Domain Adaptation in NLP—A Survey. COLING 2020

Domain Adaptation

- Train on many domains, or a high-resourced domain

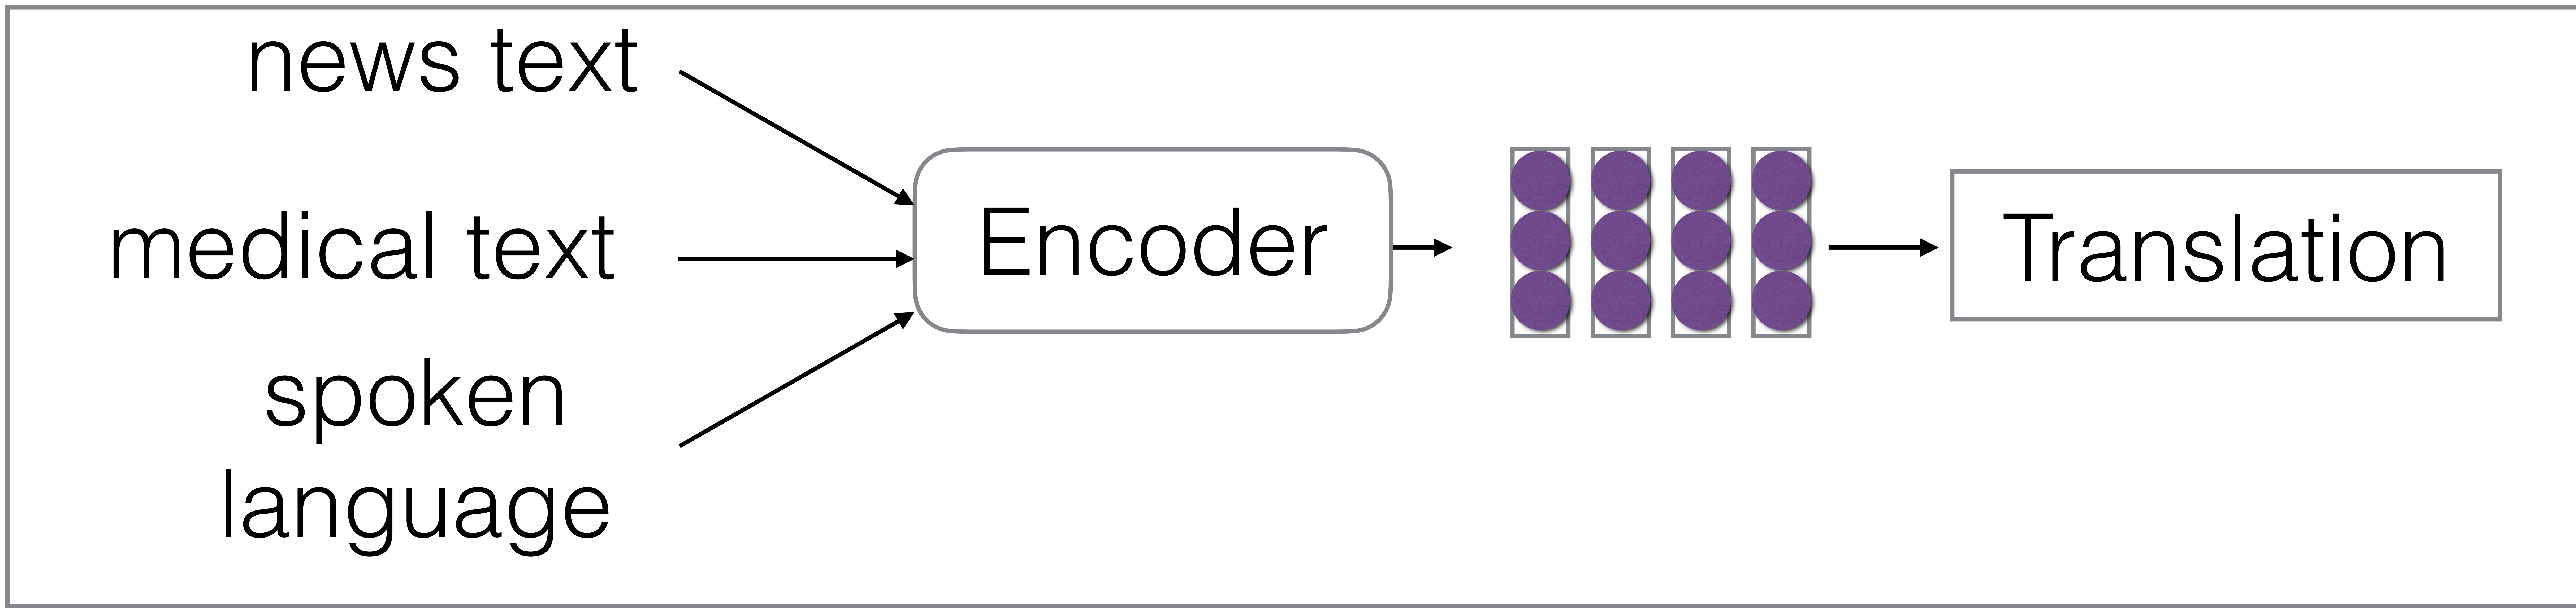

- Test on a low-resourced domain (target domain)

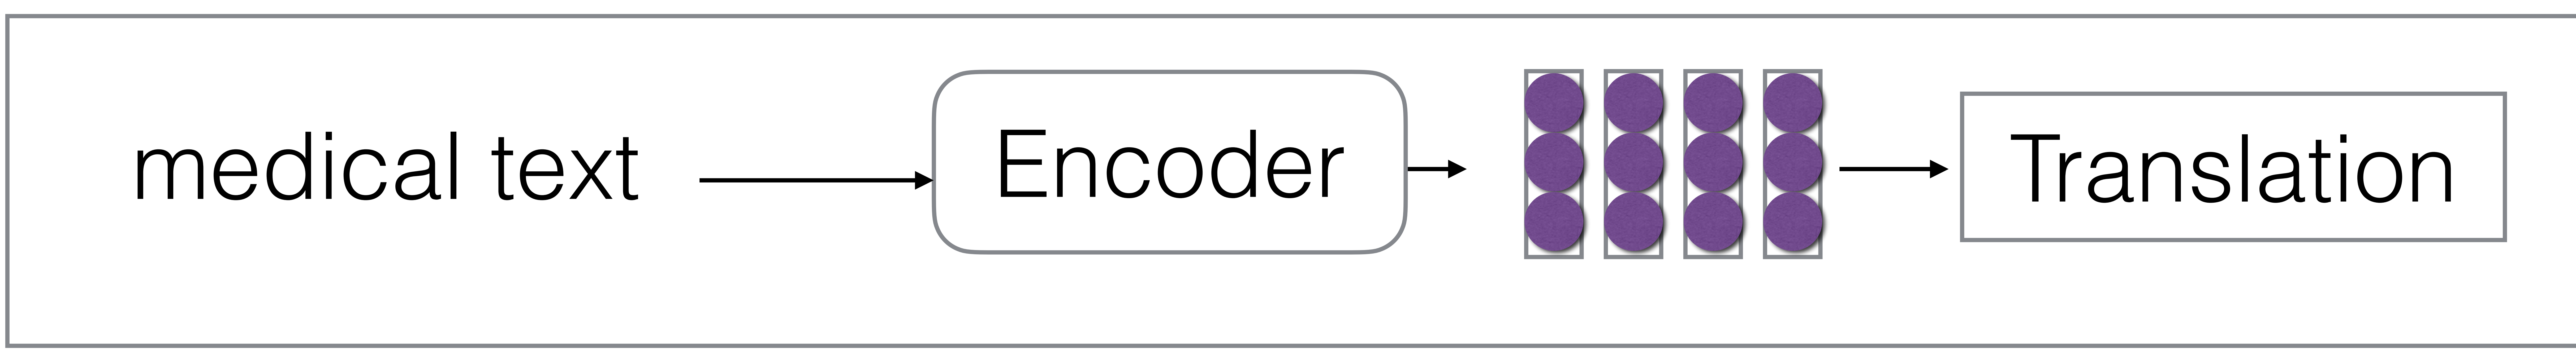

- **Supervised** adaptation: train w/ target-domain labeled data
- **Unsupervised** adaptation: train w/o target-domain labeled data

Domain Robustness

- Train on many domains and do well on all of them

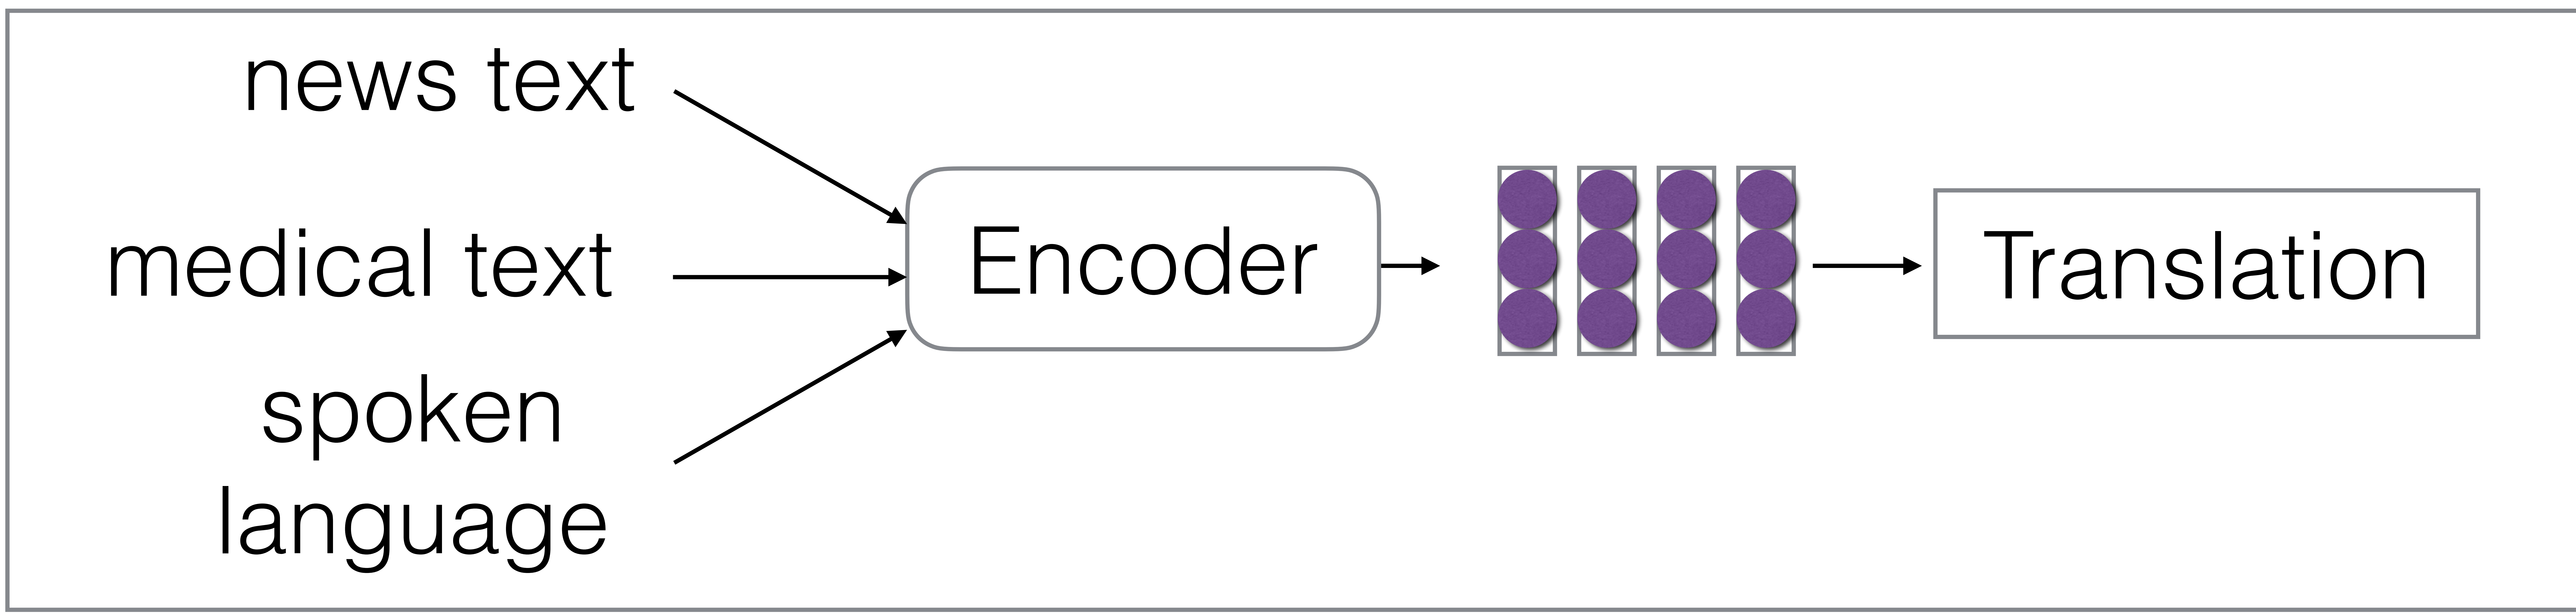

- Robustness to **minority domains**

- **Zero-shot** robustness to domains not in training data

Multilingual Learning

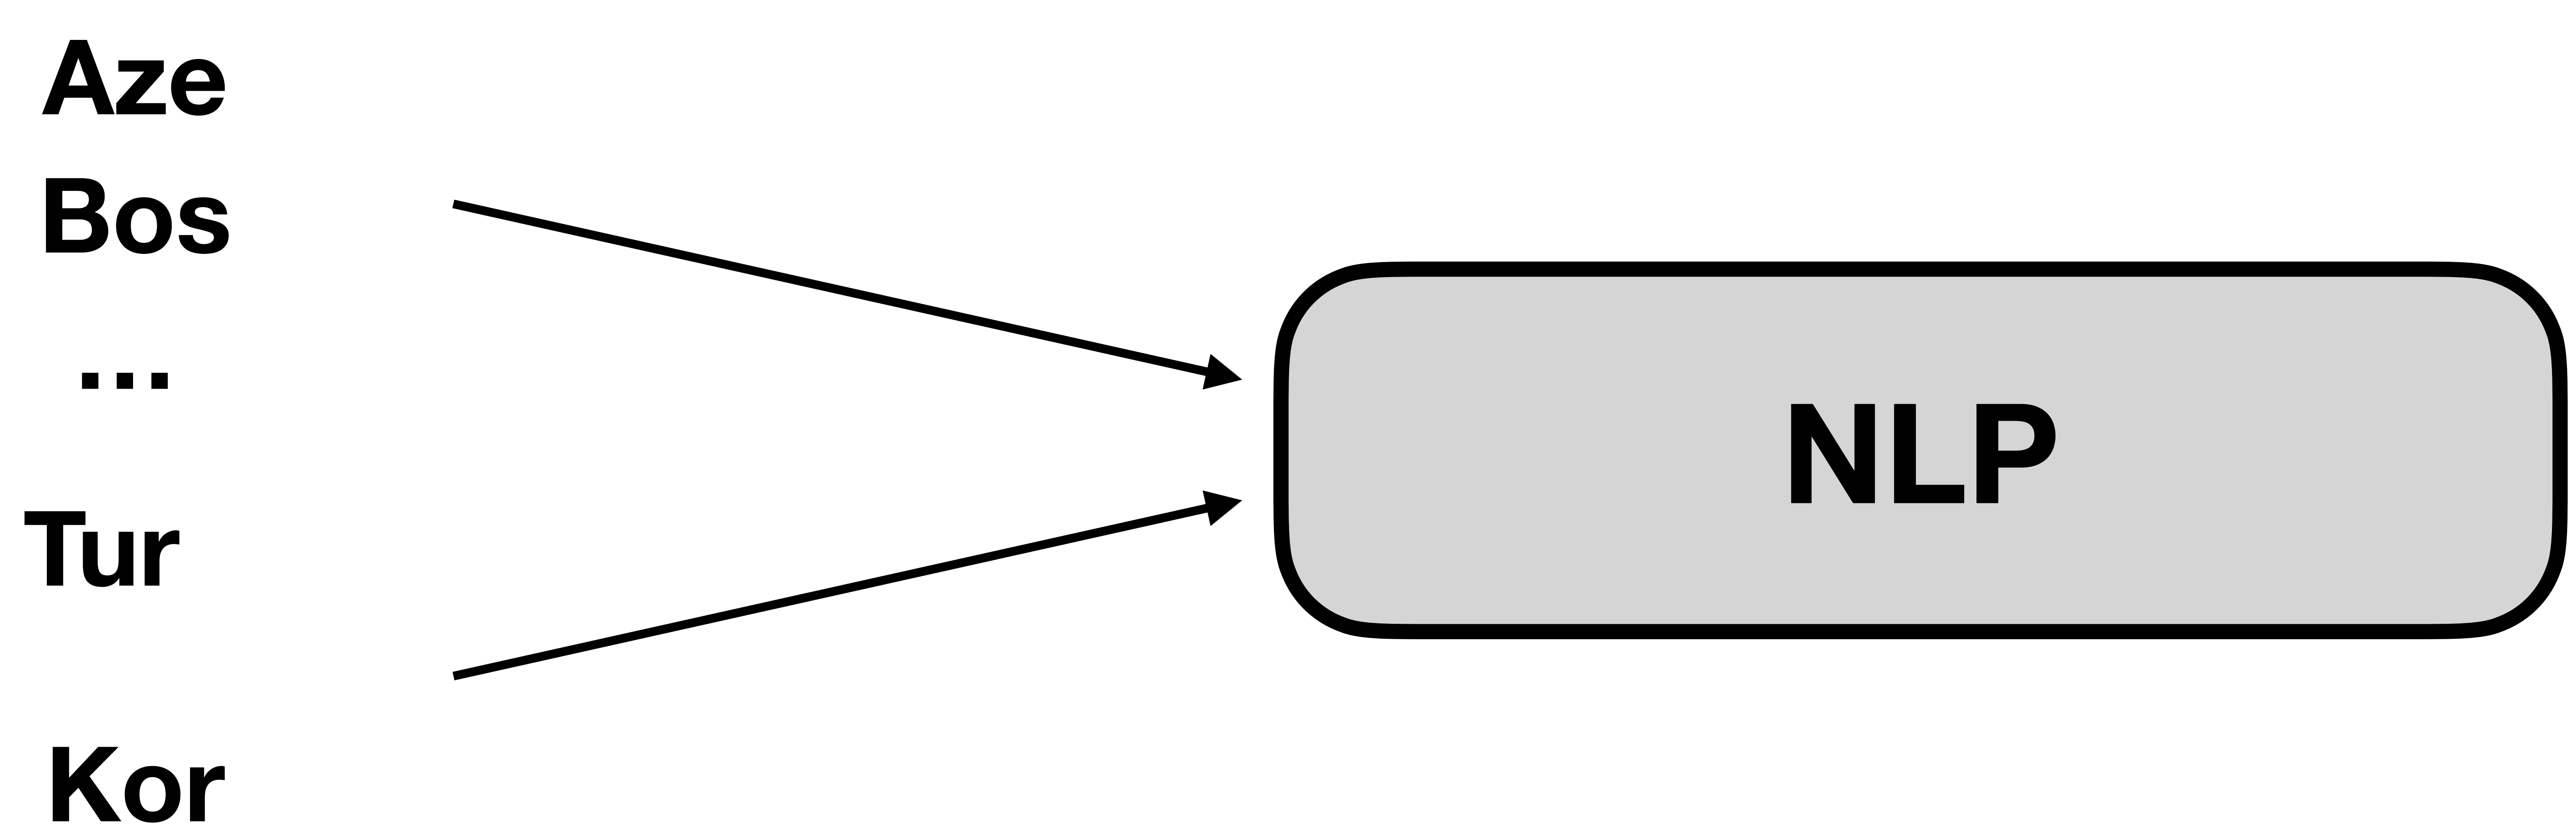

Now our best tool for applying methods to low-resourced languages

Similarity Across Languages

- Many languages share similar word roots

Cognates (joint origin)

English:	night
French:	nuit
Russian:	noch
Bengali:	nishi

Loan Words (borrowed from another)

Arabic:	qahwa
Turkish:	kahveh
English:	coffee
Japanese:	kohi
Chinese:	kafei

- Languages share a considerable amount of underlying structure, e.g. word order, grammar.

Languages as Domains

- Multilingual learning is an extreme variety,
  different language = different domain
  - **Adaptation:** Improve accuracy on lower-resource languages by transferring knowledge from higher-resource languages
  - **Robustness:** Use one model for all languages, instead of one for each
- At the same time, much more complexity!
  → Requires modeling similarities/differences in lexicon, morphology, syntax, semantics, culture

Earlier Method on Multitask Learning Feature Space Regularization

- Try to regularize the features spaces learned to be closer to each-other (e.g. Ganin et al. 2016)

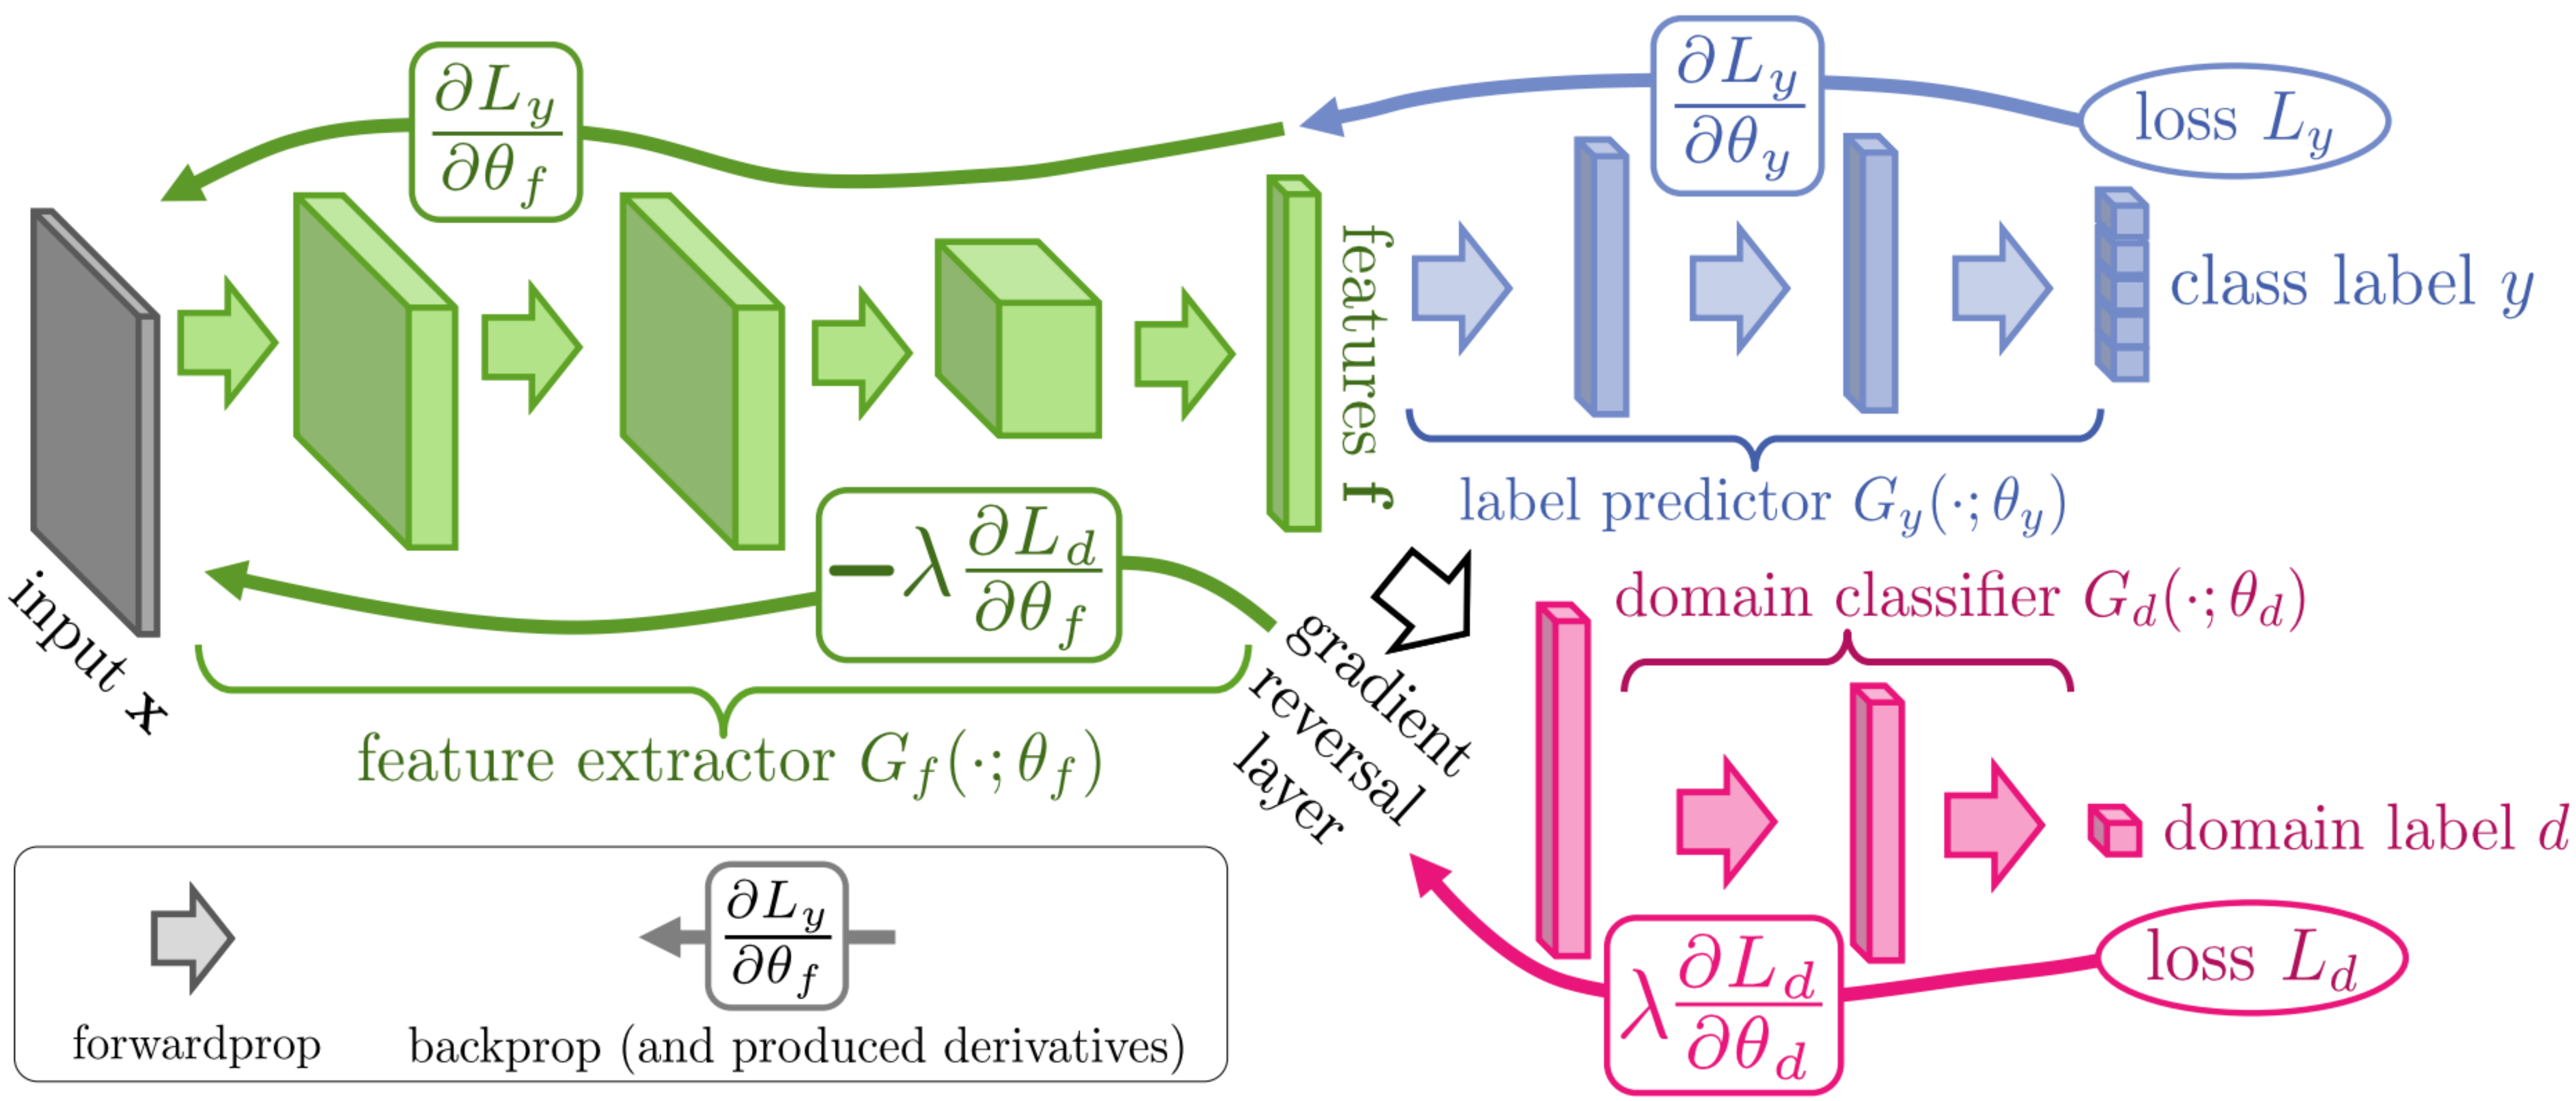

Much Simpler Multi-task Learning Method: Adding Domain Tags

- Append a domain tag to input (Chu et al. 2017)

  <news> news text

  <med> medical text

- Translate into several languages by adding a tag about the target language (Johnson et al. 2017)

  <fr> this is an example → ceci est un exemple

  <ja> this is an example → これは例です

- Introduces a small number of parameters (=embedding size) for each domain

Instruction Tuning

Text-to-Text Transfer Transformer
T5 (Raffel et al. 2020)

- Supervised training on many NLP tasks
- Control a single model to do tasks following the instructions in the prompt

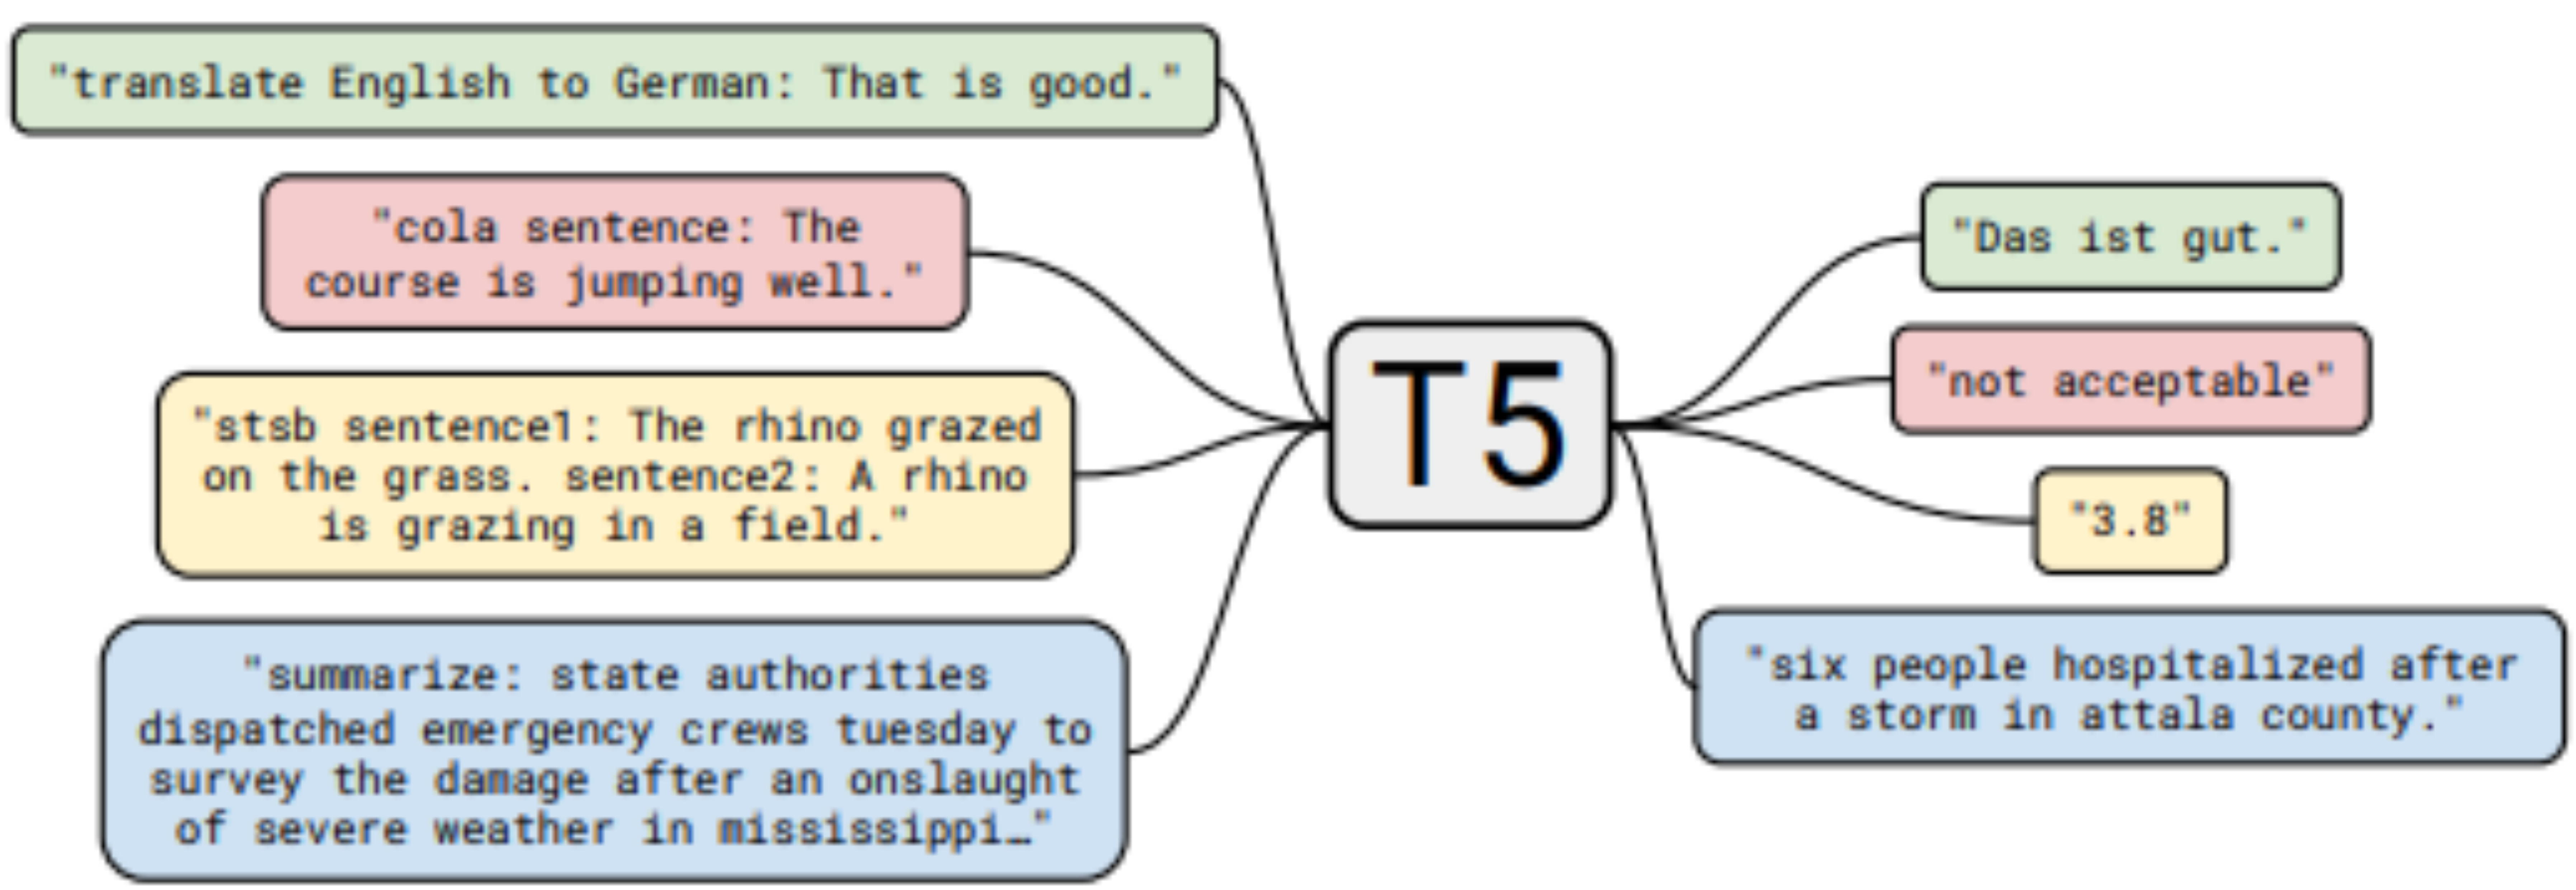

Multitask Prompted Training Enables Zero-shot Task Generalization

- T0 model (Sanh et al 2021) from Hugging Face

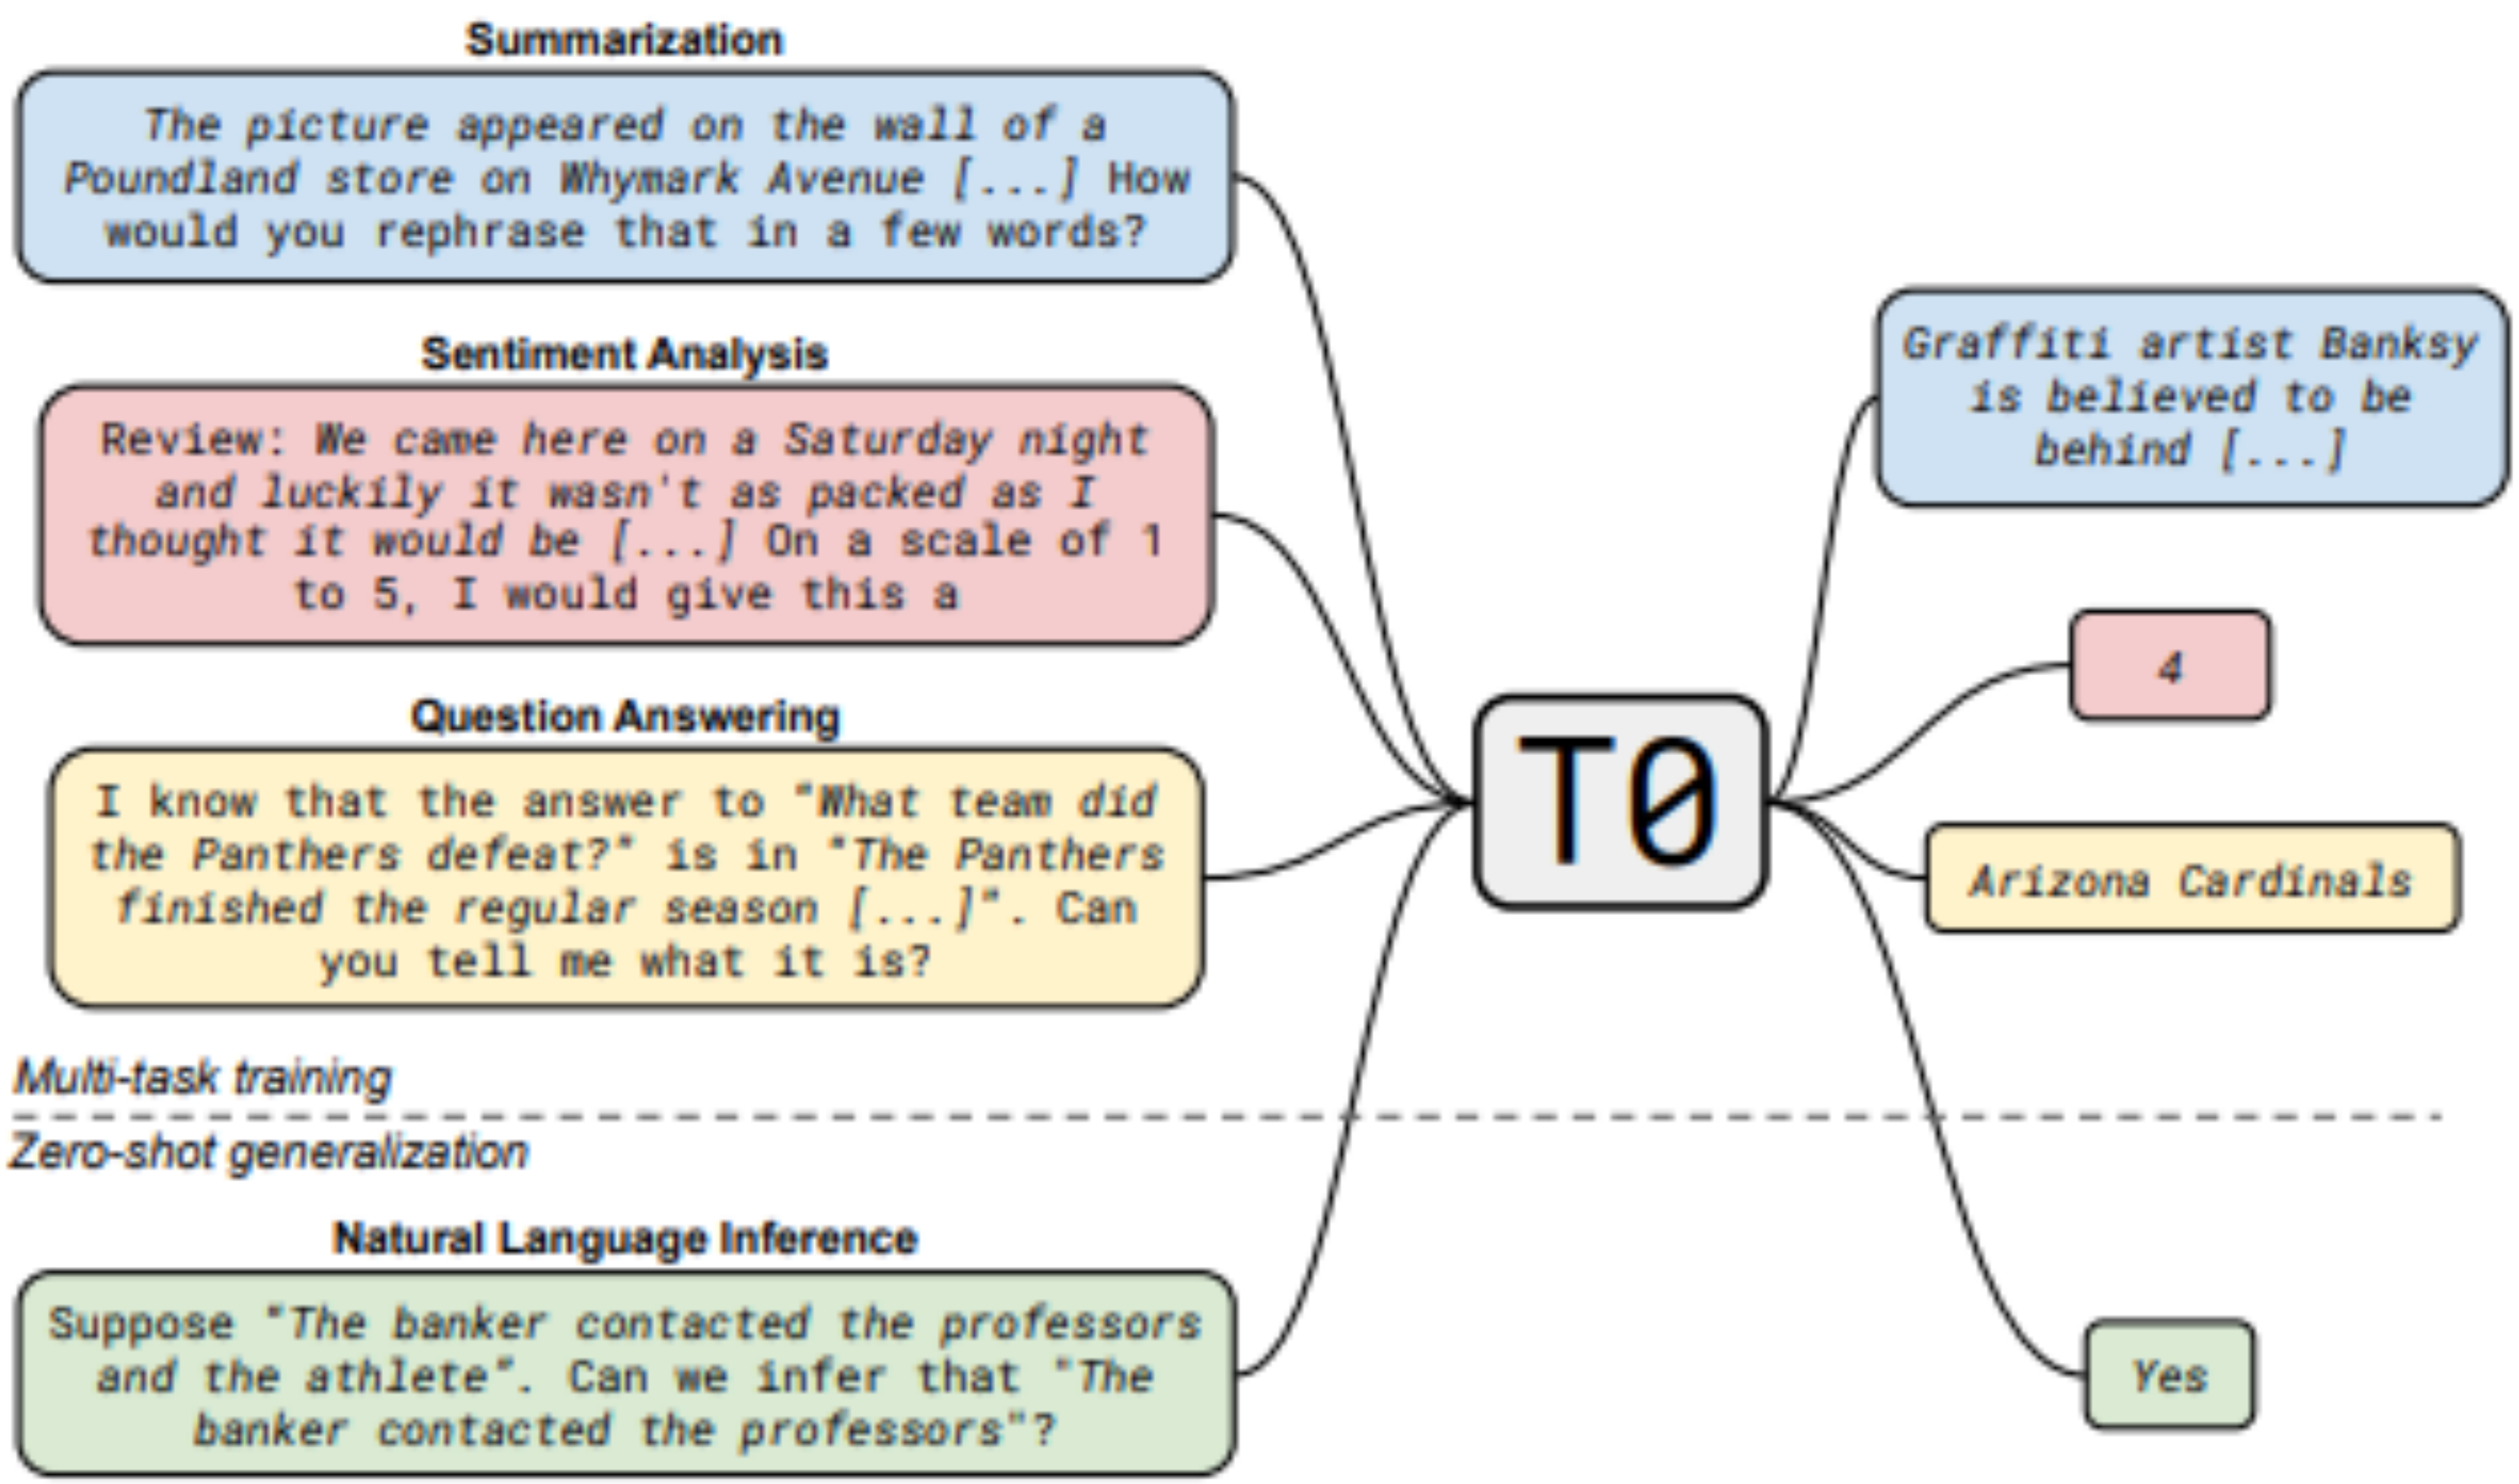

Zero-shot Task Generalization

Concurrent paper from Google
FLAN (Wei et al. 2021)

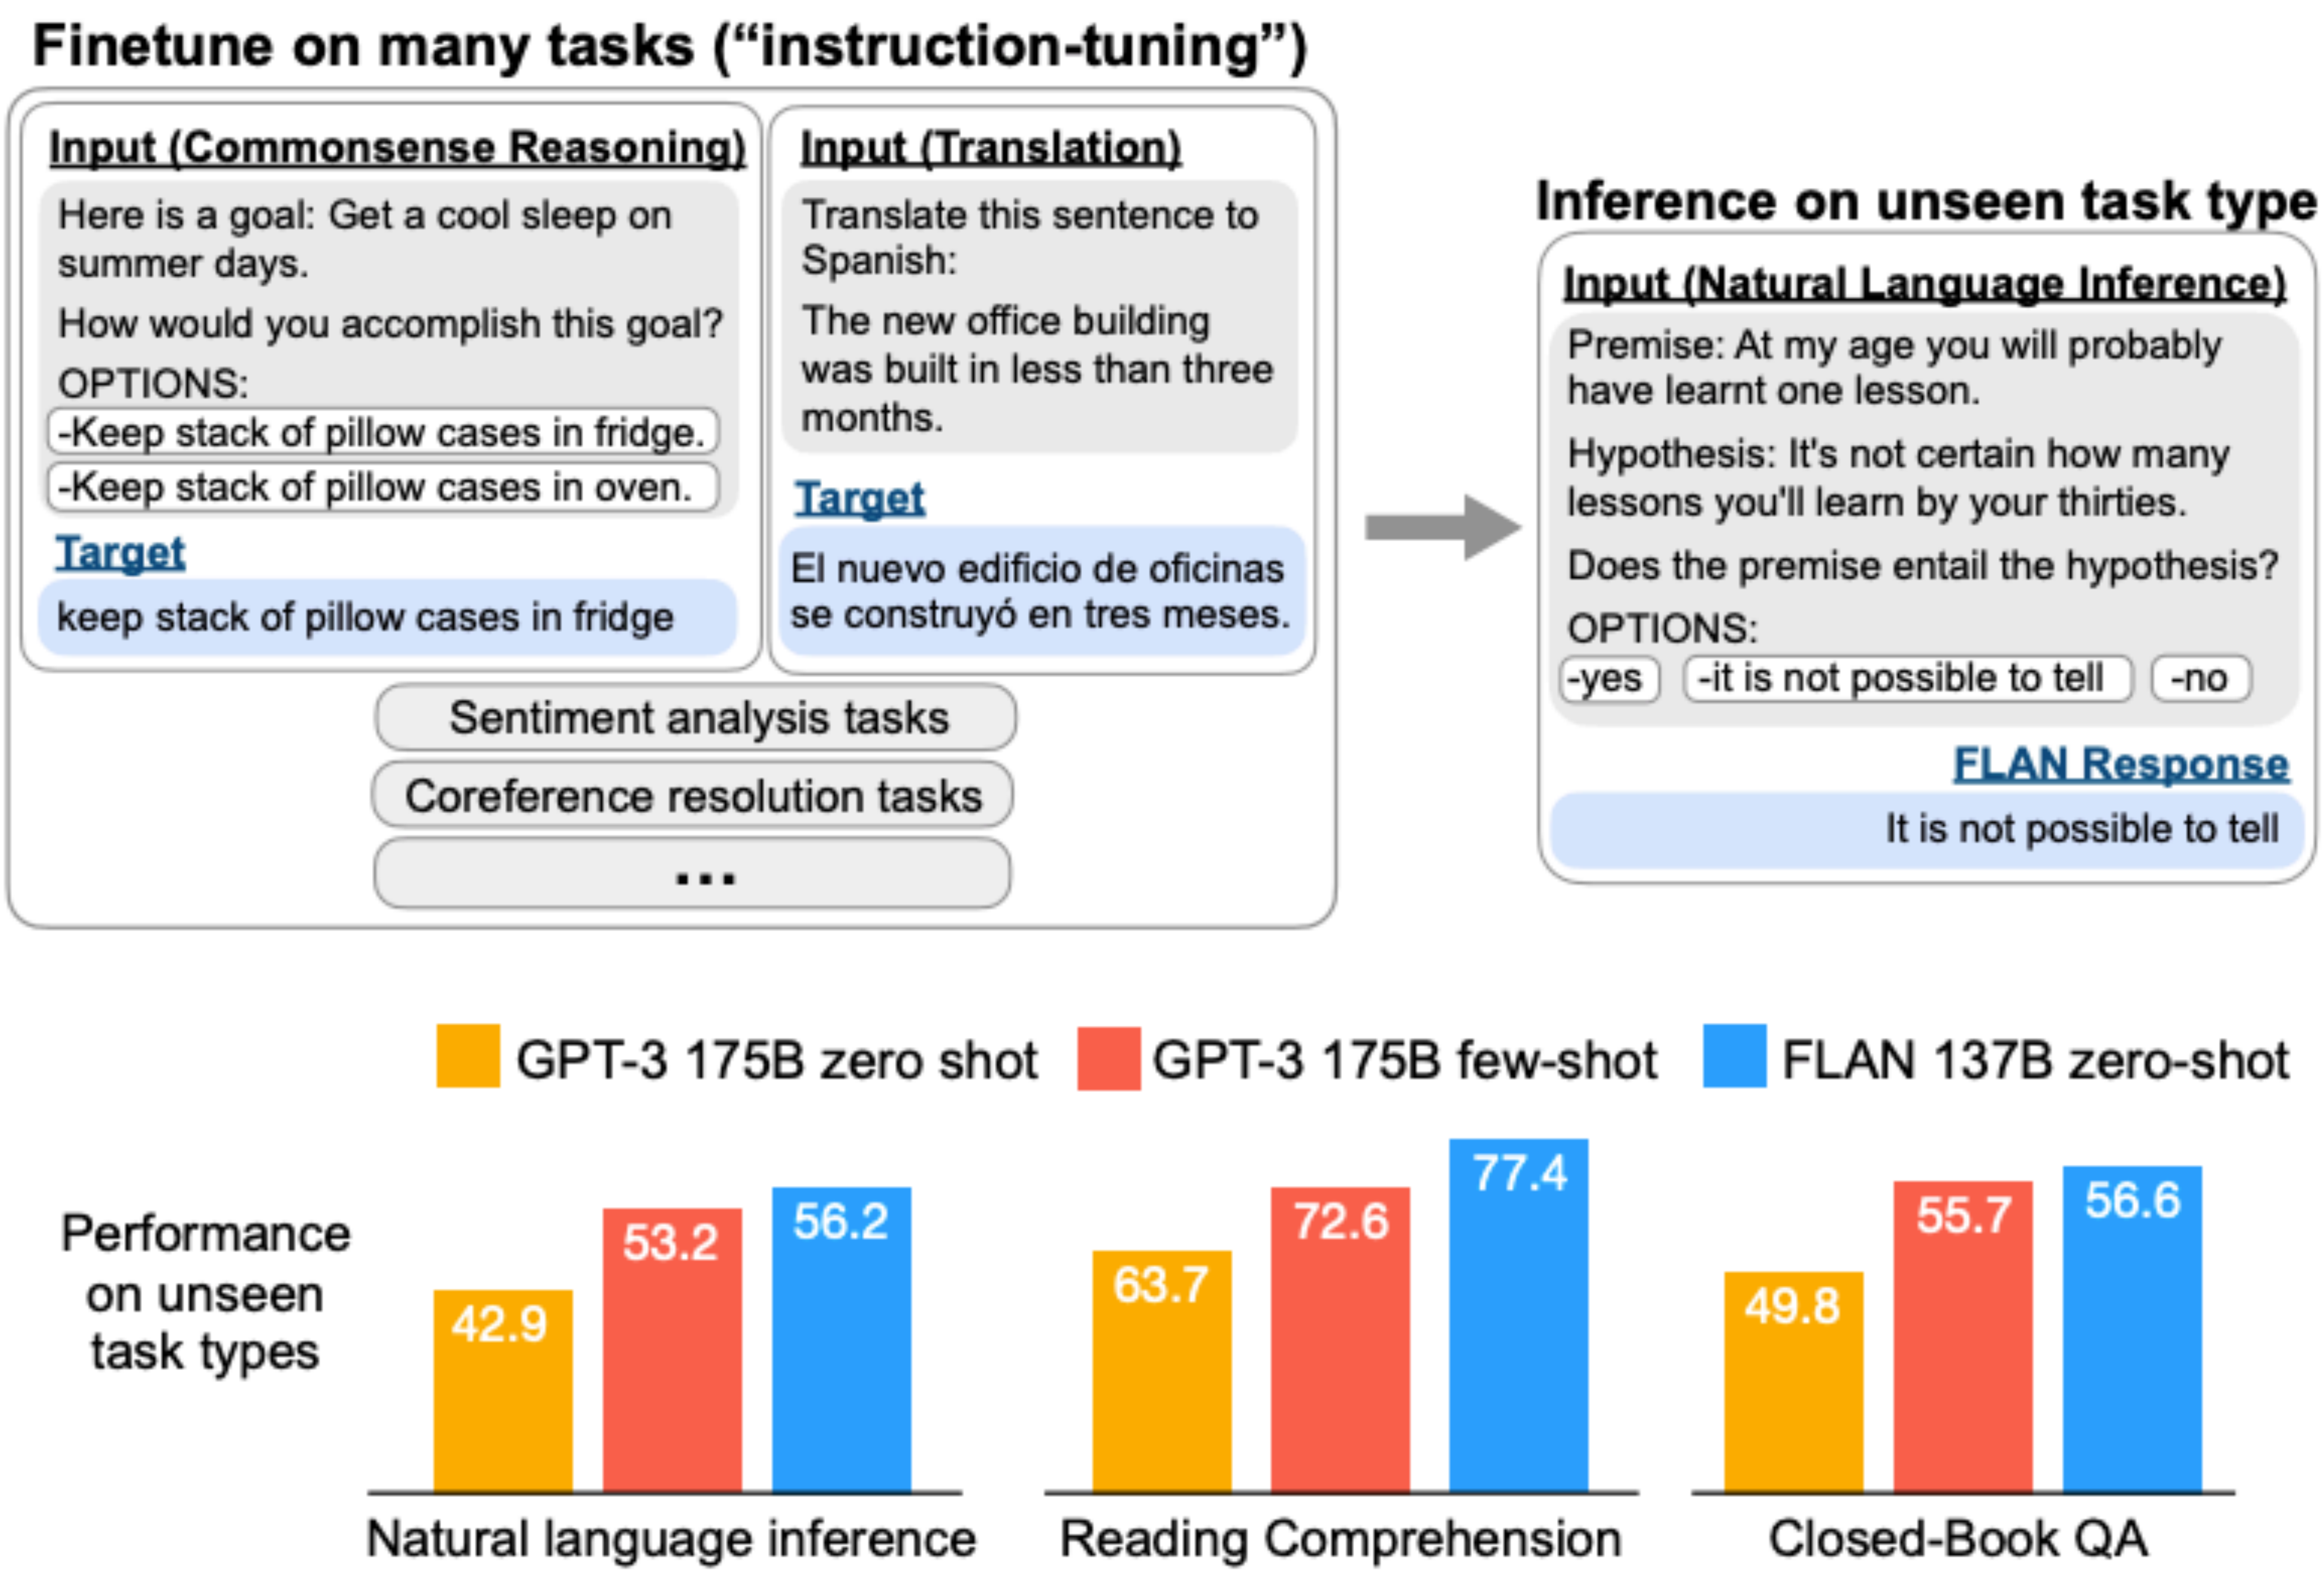

Learning to In-context Learn
(Min et al. 2021)

- Convert many-shot datasets (typically used in fine-tuning) to few-shot in-context learning examples

	Meta-training	Inference
Task	C *meta-training* tasks	An unseen *target* task
Data given	Training examples $\mathcal{T}_i = \{(x^i_j, y^i_j)\}_{j=1}^{N_i}, \ \forall i \in [1, C] \quad (N_i \gg k)$	Training examples $(x_1, y_1), \cdots, (x_k, y_k)$, Test input x
Objective	For each iteration, 1. Sample task $i \in [1, C]$ 2. Sample $k+1$ examples from $\mathcal{T}_i$: $(x_1, y_1), \cdots, (x_{k+1}, y_{k+1})$ 3. Maximize $P(y_{k+1} \mid x_1, y_1, \cdots, x_k, y_k, x_{k+1})$	$\text{argmax}_{c \in \mathcal{C}} P(c \mid x_1, y_1, \cdots, x_k, y_k, x)$

Instruction Tuning Datasets

- Good reference: FLAN Collection (Longpre et al. 2023)

		Model Details				Data Collection & Training Details			
Release	**Collection**	**Model**	**Base**	**Size**	**Public?**	**Prompt Types**	**Tasks in Flan**	**# Exs**	**Methods**
2020 05	UnifiedQA	UnifiedQA	RoBerta	110-340M	P	ZS	46 / 46	750k	
2021 04	CrossFit	BART-CrossFit	BART	140M	NP	FS	115 / 159	71.M	
2021 04	Natural Inst v1.0	Gen. BART	BART	140M	NP	ZS / FS	61 / 61	620k	+ Detailed k-shot Prompts
2021 09	Flan 2021	Flan-LaMDA	LaMDA	137B	NP	ZS / FS	62 / 62	4.4M	+ Template Variety
2021 10	P3	T0, T0+, T0++	T5-LM	3-11B	P	ZS	62 / 62	12M	+ Template Variety + Input Inversion
2021 10	MetaICL	MetaICL	GPT-2	770M	P	FS	100 / 142	3.5M	+ Input Inversion + Noisy Channel Opt
2021 11	ExMix	ExT5	T5	220M-11B	NP	ZS	72 / 107	500k	+ With Pretraining
2022 04	Super-Natural Inst.	*Tk*-Instruct	T5-LM, mT5	11-13B	P	ZS / FS	1556 / 1613	5M	+ Detailed k-shot Prompts + Multilingual
2022 10	GLM	GLM-130B	GLM	130B	P	FS	65 / 77	12M	+ With Pretraining + Bilingual (en, zh-cn)
2022 11	xP3	BLOOMz, mT0	BLOOM, mT5	13-176B	P	ZS	53 / 71	81M	+ Massively Multilingual
2022 12	Unnatural Inst.†	T5-LM-Unnat. Inst.	T5-LM	11B	NP	ZS	~20 / 117	64k	+ Synthetic Data
2022 12	Self-Instruct†	GPT-3 Self Inst.	GPT-3	175B	NP	ZS	Unknown	82k	+ Synthetic Data + Knowledge Distillation
2022 12	OPT-IML Bench†	OPT-IML	OPT	30-175B	P	ZS + FS CoT	~2067 / 2207	18M	+ Template Variety + Input Inversion + Multilingual
2022 10	Flan 2022 (ours)	Flan-T5, Flan-PaLM	T5-LM, PaLM	10M-540B	P NP	ZS + FS CoT	1836	15M	+ Template Variety + Input Inversion + Multilingual

Instruction Tuned Models

- **FLAN-T5:** huggingface/google/flan-t5-xxl
  - Encoder-decoder model based on T5
  - 11B parameters
- **LLaMa-2 Chat:** huggingface/meta-llama/Llama-2-70b-chat-hf
  - Decoder-only model
  - 70B parameters
- **Mixtral instruct:** huggingface/mistralai/Mixtral-8x7B-Instruct-v0.1
  - Decoder-only mixture of experts model
  - 45B parameters
- *(smaller versions also available - Mistral, LLaMa2-7B)*

Dataset Generation

- It is possible to automatically generate instruction tuning datasets, e.g. **self-instruct** (Wang et al. 2022)

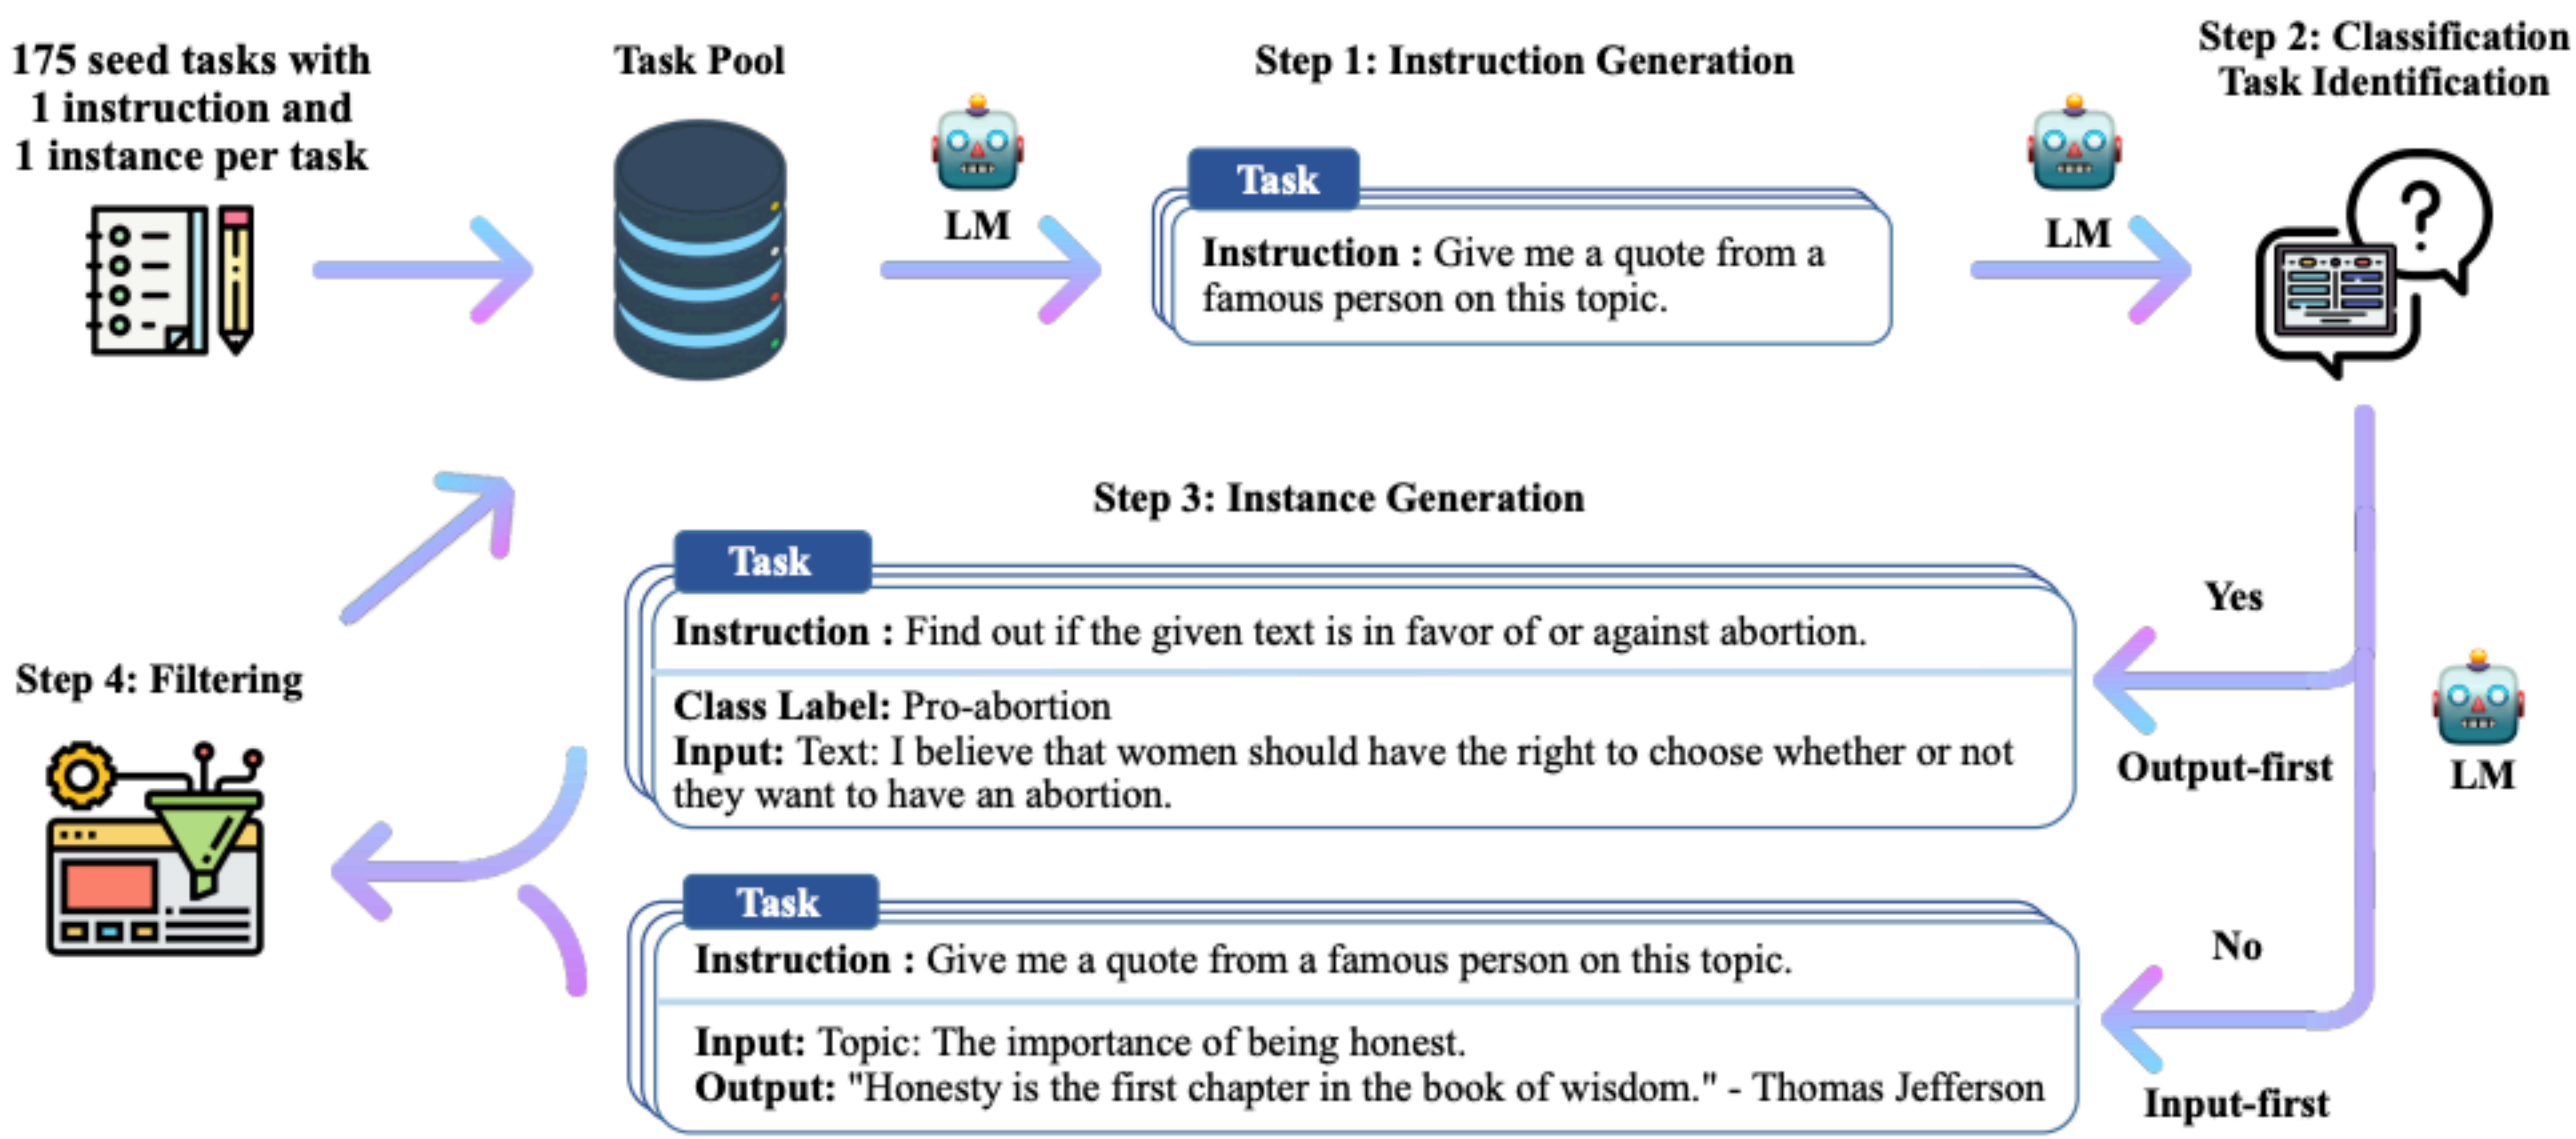

- Can be used to train chain-of-thought — ORCA (Mukherjee et al. 2023)
- Can be used to make instructions more complex — Evol-Instruct (Xu et al. 2023)

Selected Tasks from the GPT3-generated Data

- **Input-first**: Ask an LLM to come up with the input fields first based on the instruction, then produce the output
- **Output-first**: Generate the output label, and then generate the input based on the output

Instruction: Given an address and city, come up with the zip code.
Input:
Address: 123 Main Street, City: San Francisco
Output: 94105

Instruction: I am looking for a job and I need to fill out an application form. Can you please help me complete it?
Input:
Application Form:
Name:_________ Age:__________ Sex:________
Phone Number:______ Email Address:________
Education:___________________ ...
Output:
Name: John Doe Age: 25 Sex: Male
Phone Number: ...

Instruction: How to write a code for converting degrees fahrenheit to celsius.
Input: Null
Output:

```
def convert_fahrenheit_to_celsius(fahr):
    celsius = (fahr - 32) * 5 / 9
    return celsius
```

Instruction: Write a letter from the perspective of a cat.
Input: Null
Output:
Dear [Owner], I am writing to you today because I have a problem. You see, I've been living with you for the past [Number of Years] and I love it here. The food is great, the bed is comfy, and there are plenty of toys to play with. But lately, I've noticed that you're not as affectionate towards me as you used to be. In fact, sometimes when I try to cuddle up next to you on the couch, you push me away ...

Examples

```
Come up with examples for the following tasks. Try to generate multiple examples when possible.
If the task doesn't require additional input, you can generate the output directly.

Task: Which exercises are best for reducing belly fat at home?
Output:
- Lying Leg Raises
- Leg In And Out
- Plank
- Side Plank
- Sit-ups

Task: Extract all the country names in the paragraph, list them separated by commas.
Example 1
Paragraph: Dr. No is the sixth novel by the English author Ian Fleming to feature his British
Secret Service agent James Bond. Written at Fleming's Goldeneye estate in Jamaica, it was
first published in the United Kingdom by Jonathan Cape in 1958. In the novel Bond looks into
the disappearance in Jamaica of two fellow MI6 operatives who had been investigating Doctor
No. Bond travels to No's Caribbean island and meets Honeychile Rider, who is there to collect
shells. They are captured and taken to a luxurious facility carved into a mountain. The
character of Doctor No, the son of a German missionary and a Chinese woman, was influenced by
Sax Rohmer's Fu Manchu stories. Dr. No was the first of Fleming's novels to face widespread
negative reviews in Britain, but it was received more favourably in the United States.
Output: English, British, Jamaica, the United Kingdom, German, Chinese, Britain, the United
States.

Task: Converting 85 F to Celsius.
Output: 85°F = 29.44°C
```

Evol-Instruct
(Xu et al. 2023)

- Starting with an initial set of instructions, rewrite them step by step into more complex instructions — Evol-Instruct (Xu et al. 2023)

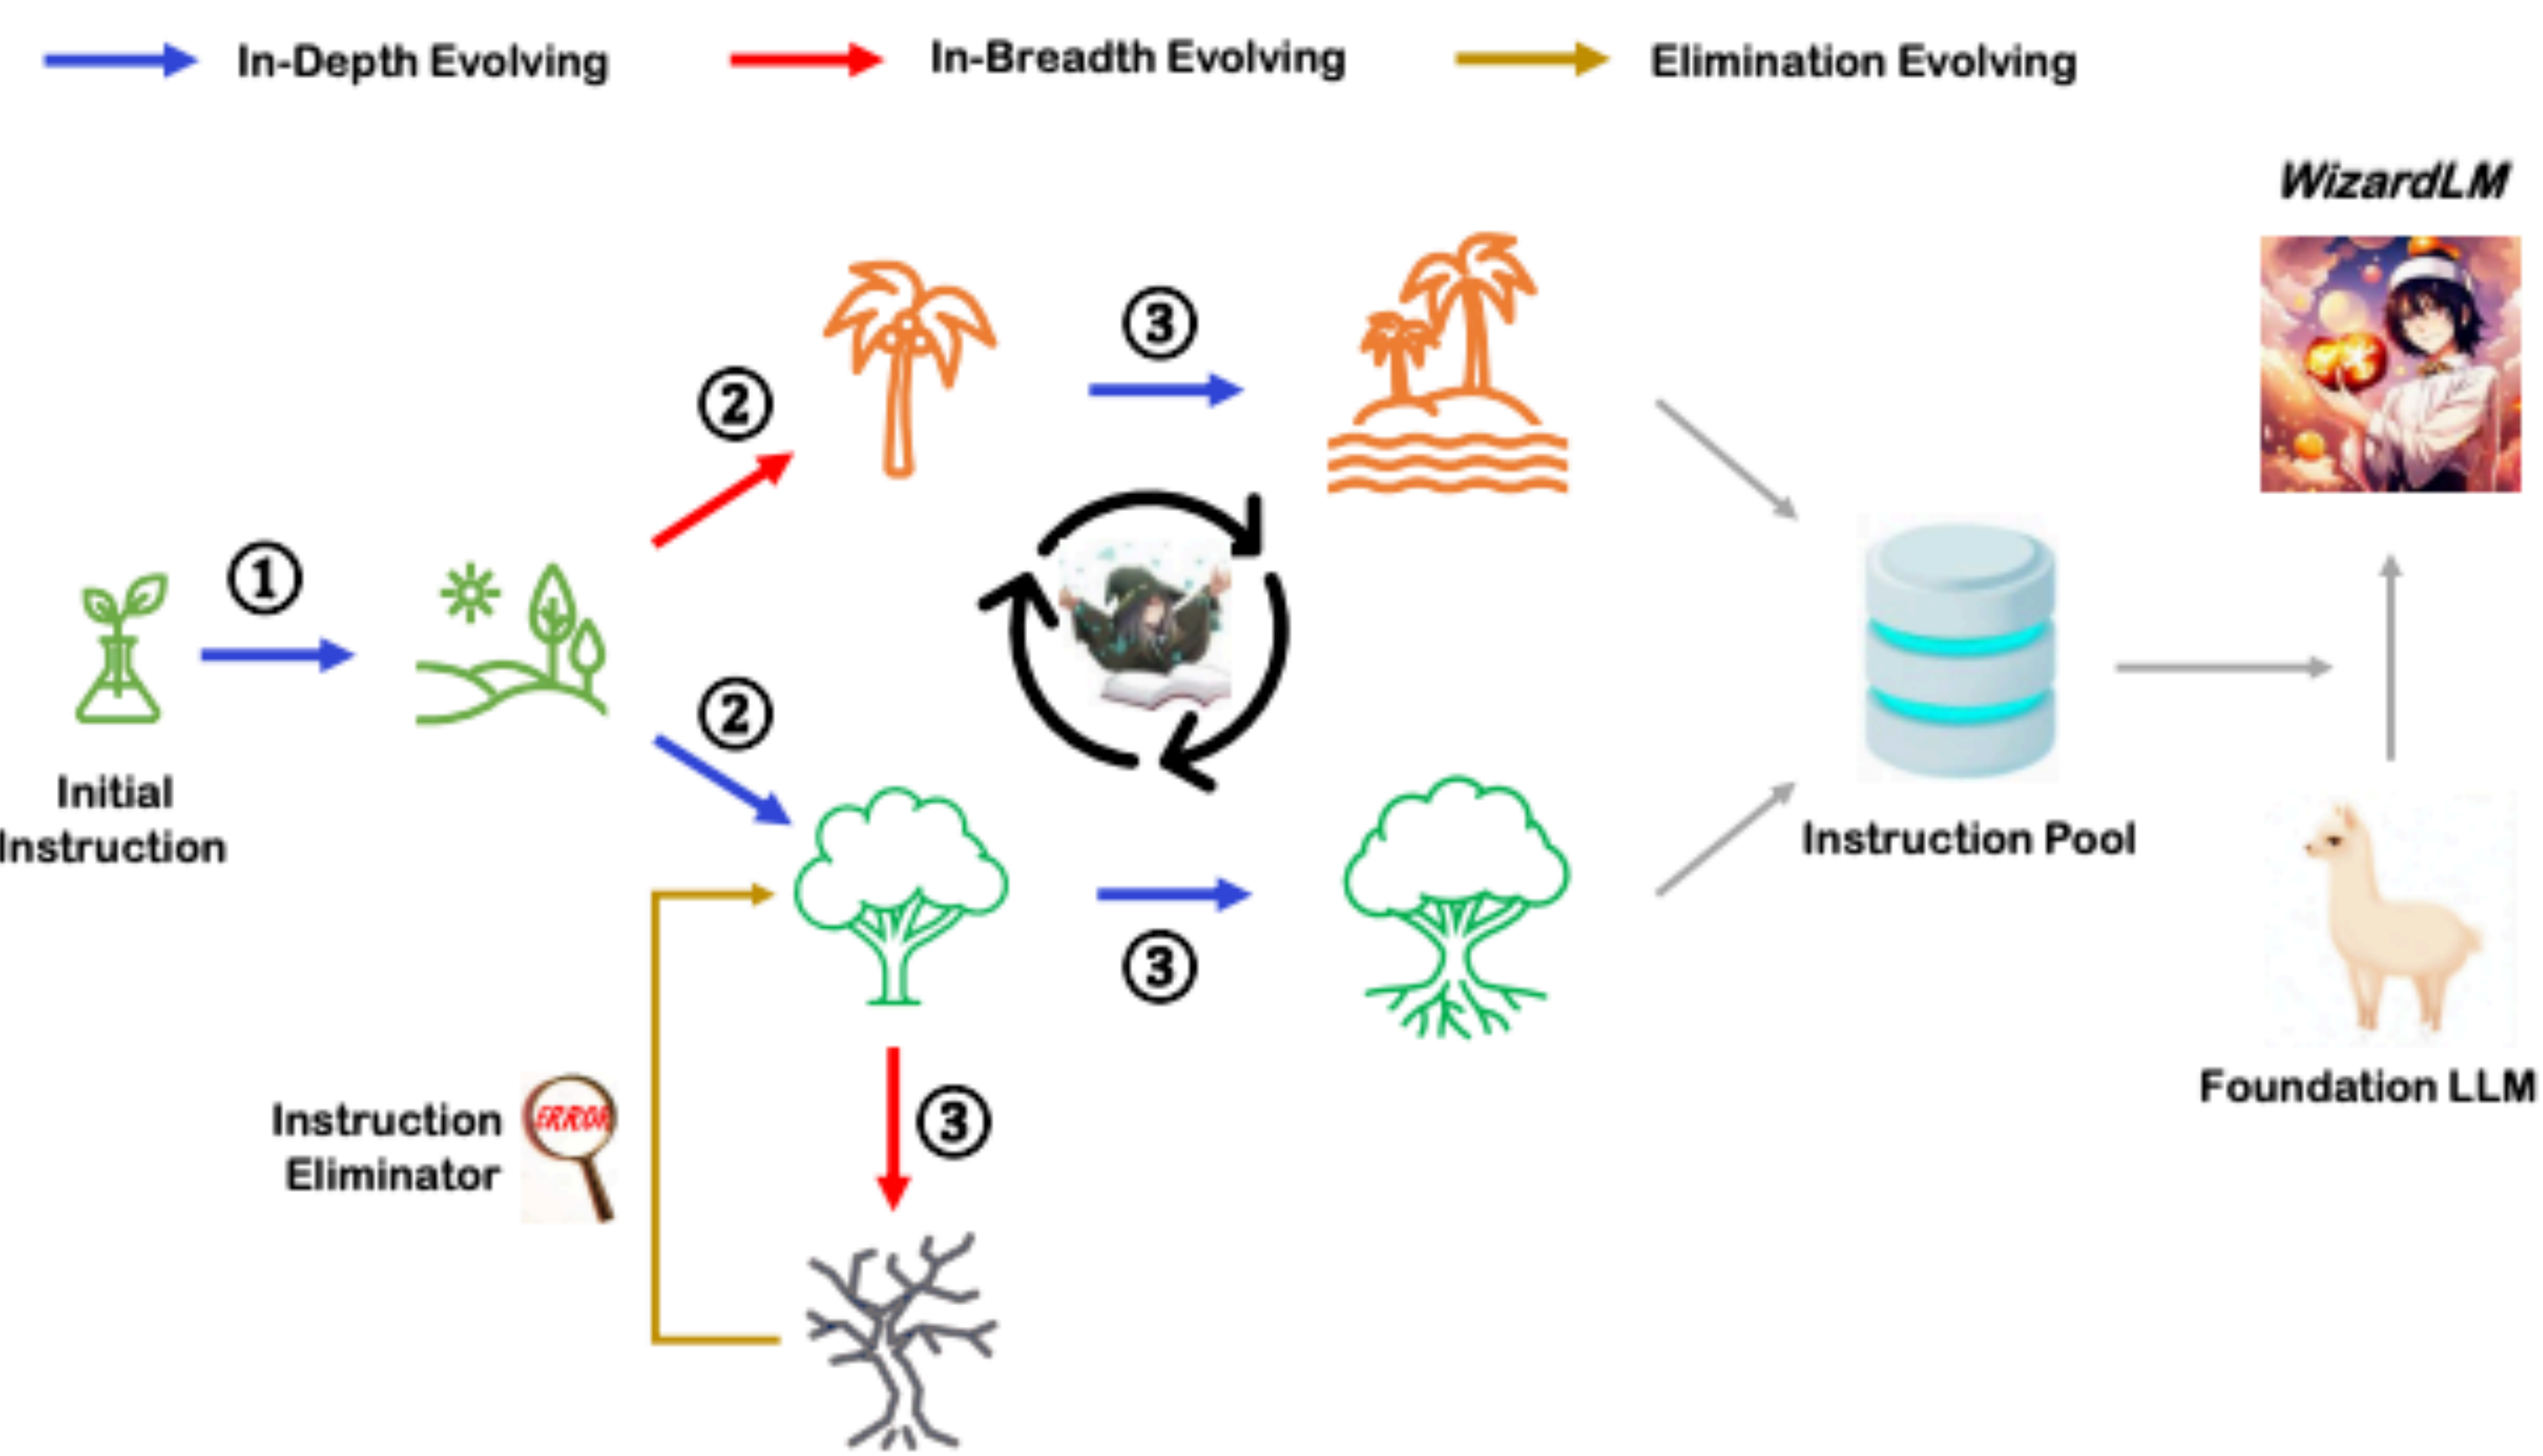

Instruction Evolver

Example 3.1: Prompt for Adding Constraints of In-Depth Evolving

I want you act as a Prompt Rewriter.
Your objective is to rewrite a given prompt into a more complex version to make those famous AI systems (e.g., ChatGPT and GPT4) a bit harder to handle. But the rewritten prompt must be reasonable and must be understood and responded by humans.
Your rewriting cannot omit the non-text parts such as the table and code in #Given Prompt#:. Also, please do not omit the input in #Given Prompt#.
You SHOULD complicate the given prompt using the following method:
Please add one more constraints/requirements into #Given Prompt#

You should try your best not to make the #Rewritten Prompt# become verbose, #Rewritten Prompt# can only add 10 to 20 words into #Given Prompt#. '#Given Prompt#', '#Rewritten Prompt#', 'given prompt' and 'rewritten prompt' are not allowed to appear in #Rewritten Prompt#

#Given Prompt#:
{Here is instruction.}
#Rewritten Prompt#:

Instruction Evolver

Example 3.2: Prompt for Complicating Input of In-Depth Evolving

I want you act as a Prompt Rewriter.

Your objective is to rewrite a given prompt into a more complex version to make those famous AI systems (e.g., ChatGPT and GPT4) a bit harder to handle. But the rewritten prompt must be reasonable and must be understood and responded by humans.

You must add **[XML data]** format data as input data in [Rewritten Prompt]
#Given Prompt#:
{Here is instruction of Example 1.}
#Rewritten Prompt#:
{Here is rewritten instruction of Example 1.}

... N -1 Examples ...

You must add **[#Given Dataformat#]** format data as input data in [Rewritten Prompt]
#Given Prompt#:
{Here is instruction of Example N.}
#Rewritten Prompt#:

Instruction Evolver

Example 3.3: Prompt for In-Breadth Evolving

I want you act as a Prompt Creator.
Your goal is to draw inspiration from the #Given Prompt# to create a brand new prompt.
This new prompt should belong to the same domain as the #Given Prompt# but be even more rare.
The LENGTH and difficulty level of the #Created Prompt# should be similar to that of the #Given Prompt#.
The #Created Prompt# must be reasonable and must be understood and responded by humans.

'#Given Prompt#', '#Created Prompt#', 'given prompt' and 'created prompt' are not allowed to appear in #Created Prompt#.

#Given Prompt#:
{Here is instruction.}
#Created Prompt#:

Example of Evol-Instruct

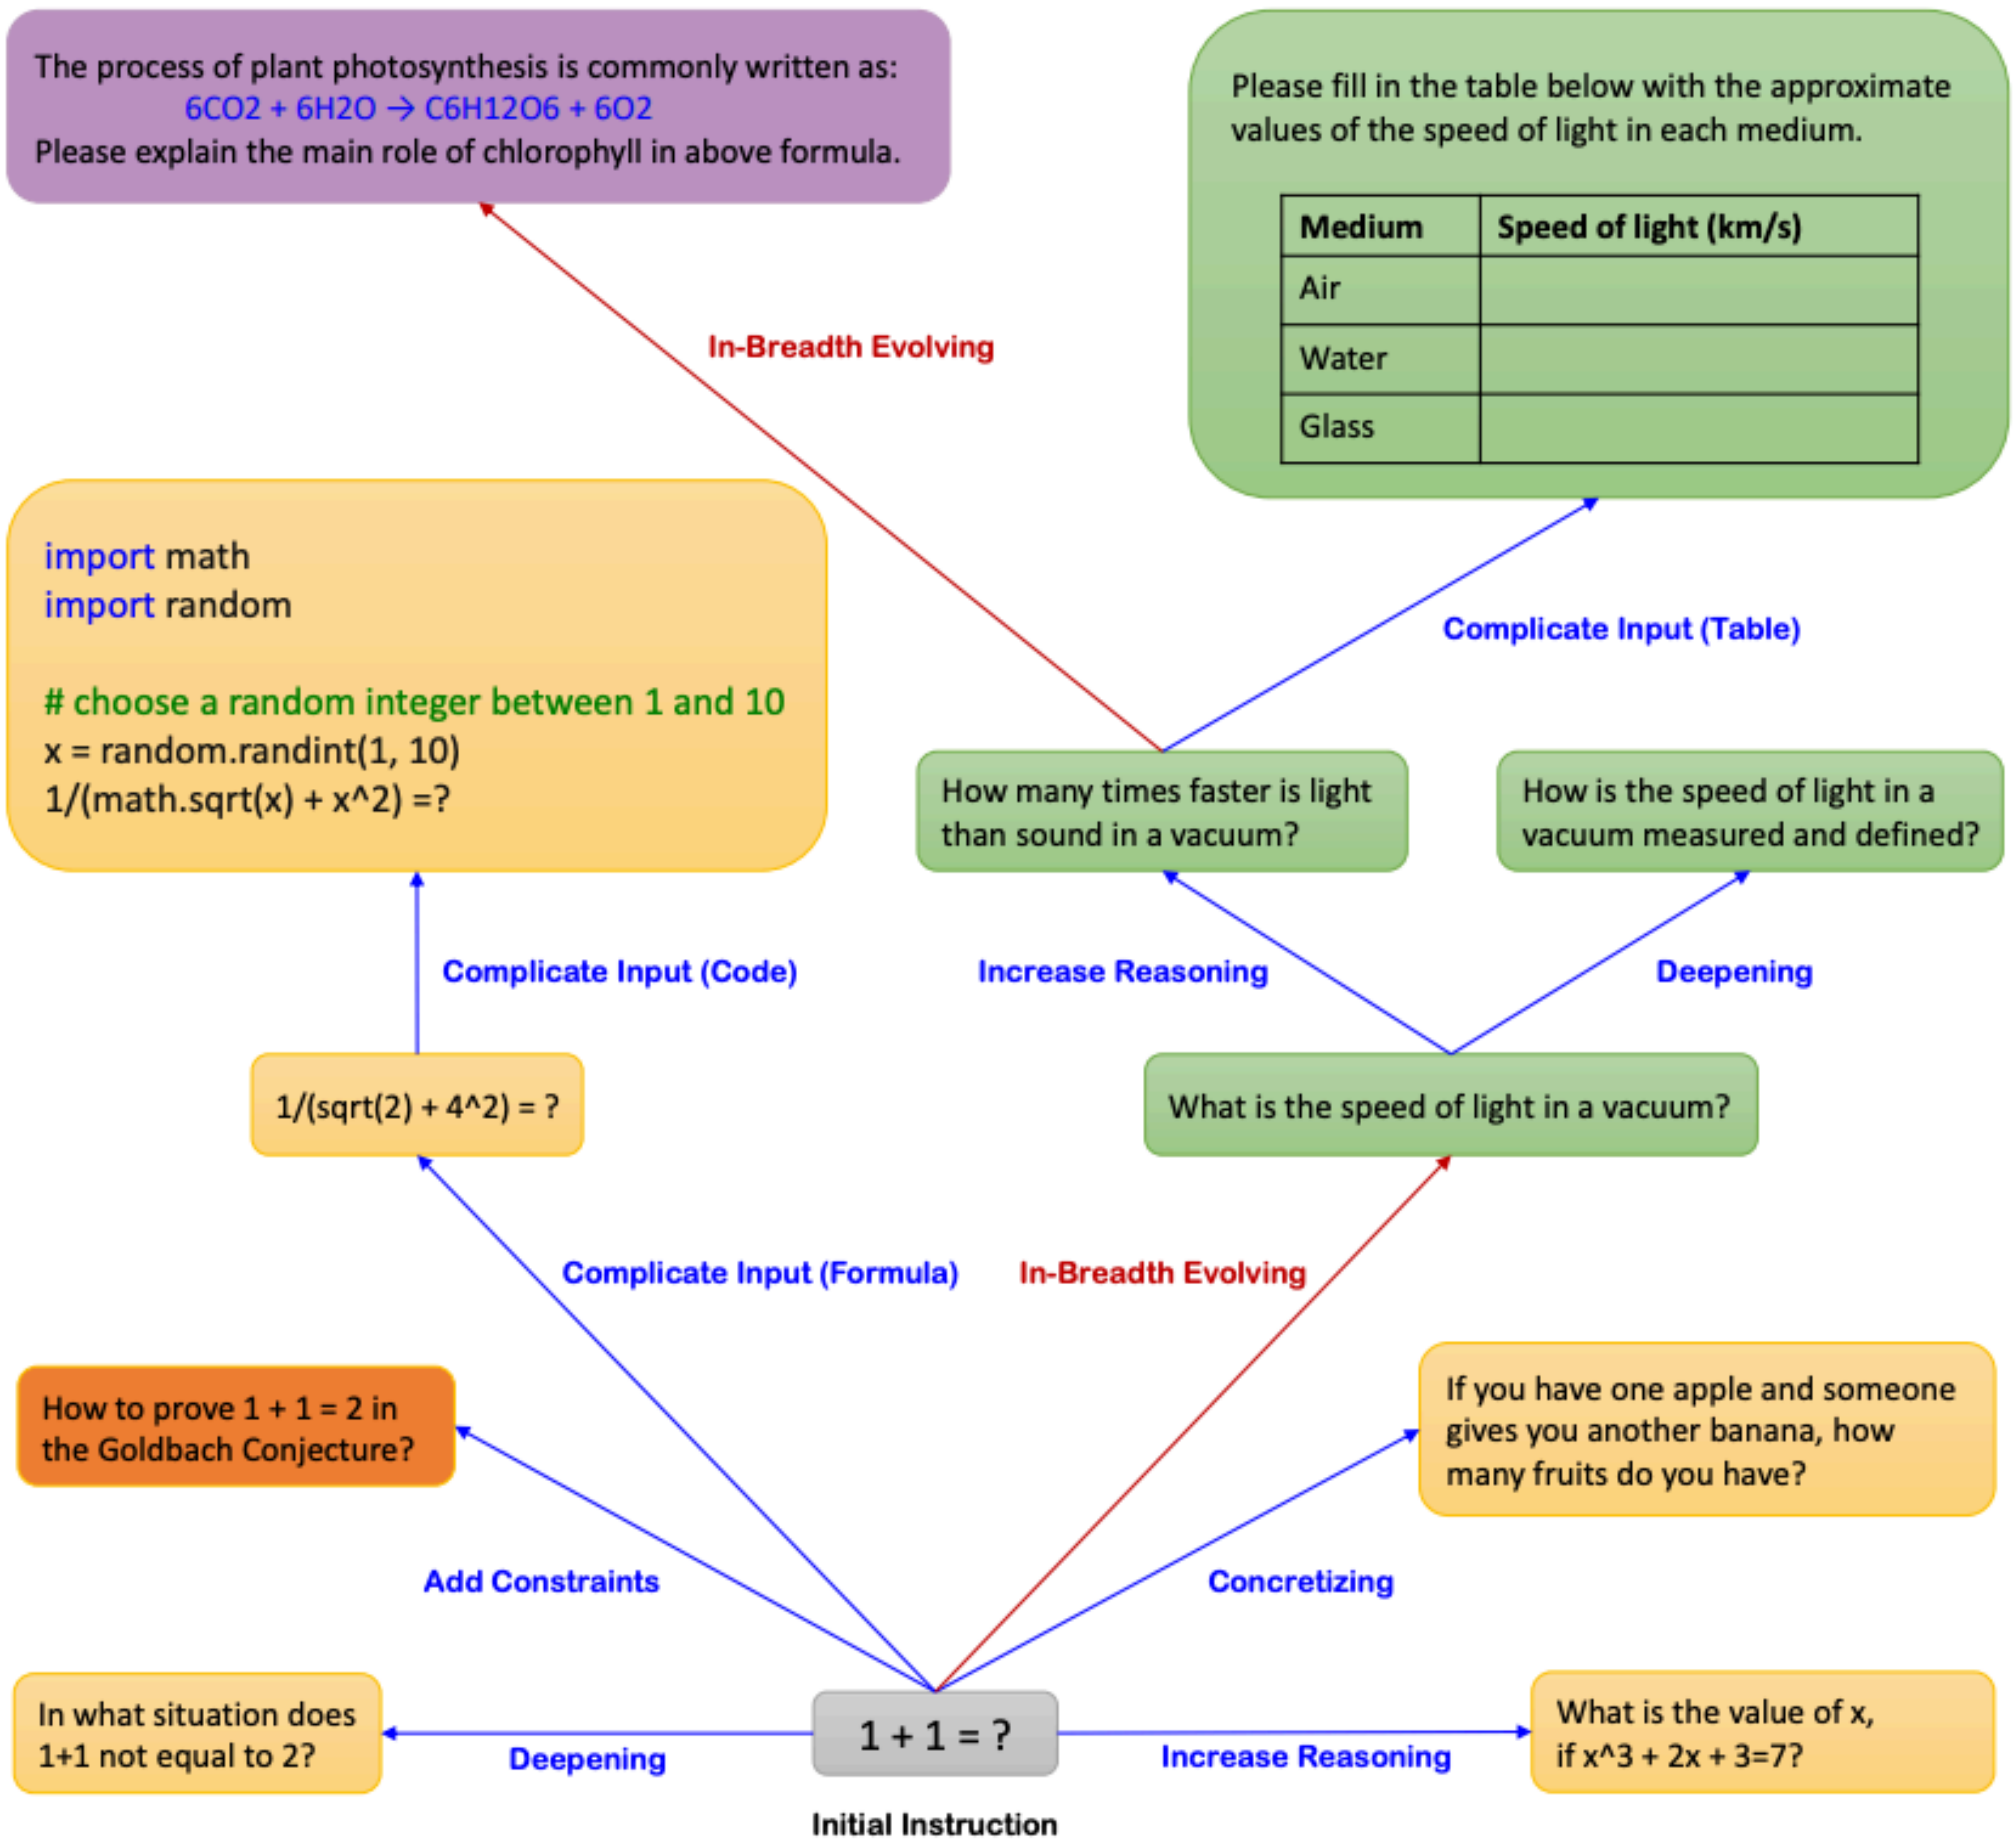

Questions?