Human-In-The-Loop for Machine Learning

Yasmine Nahal (ESR9)

AIDD Seminar - **16.03.2022**

What I will talk about

- ➢ What is “Human-In-The-Loop”?
- ➢ Motivation
- ➢ Human interaction
  - Data processing
  - Model training and inference
- ➢ Pros and cons
- ➢ Future directions & open challenges

What is "Human-In-The-Loop"

Good at making optimal decisions **when there is large and high-quality data**

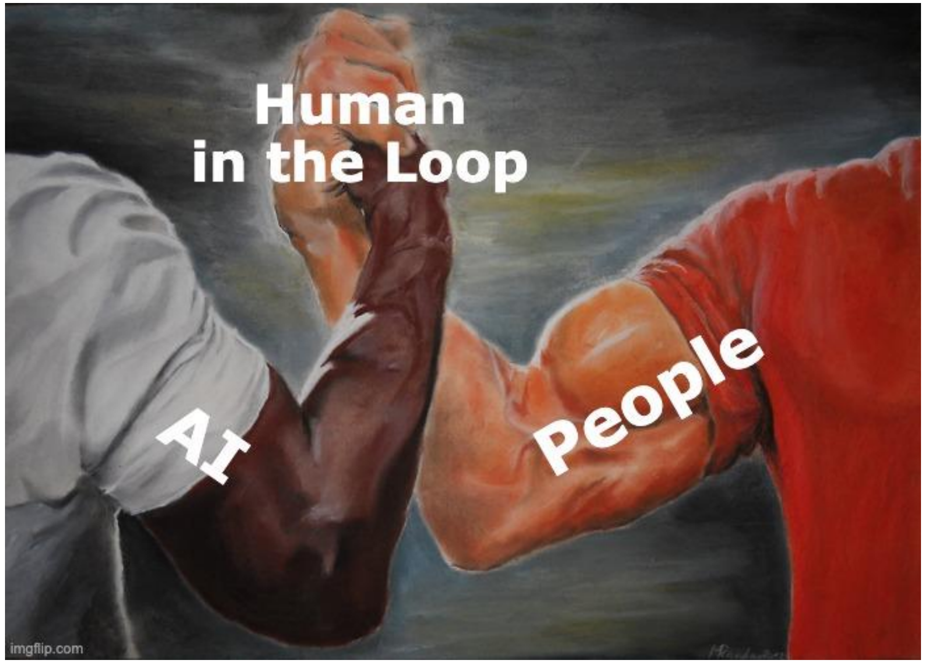

Good at recognizing patterns **within small and low-quality data**

- Human-In-The-Loop (**HITL**) is a ML method that combines **human** and **AI** to build effective ML models.

Background and motivation

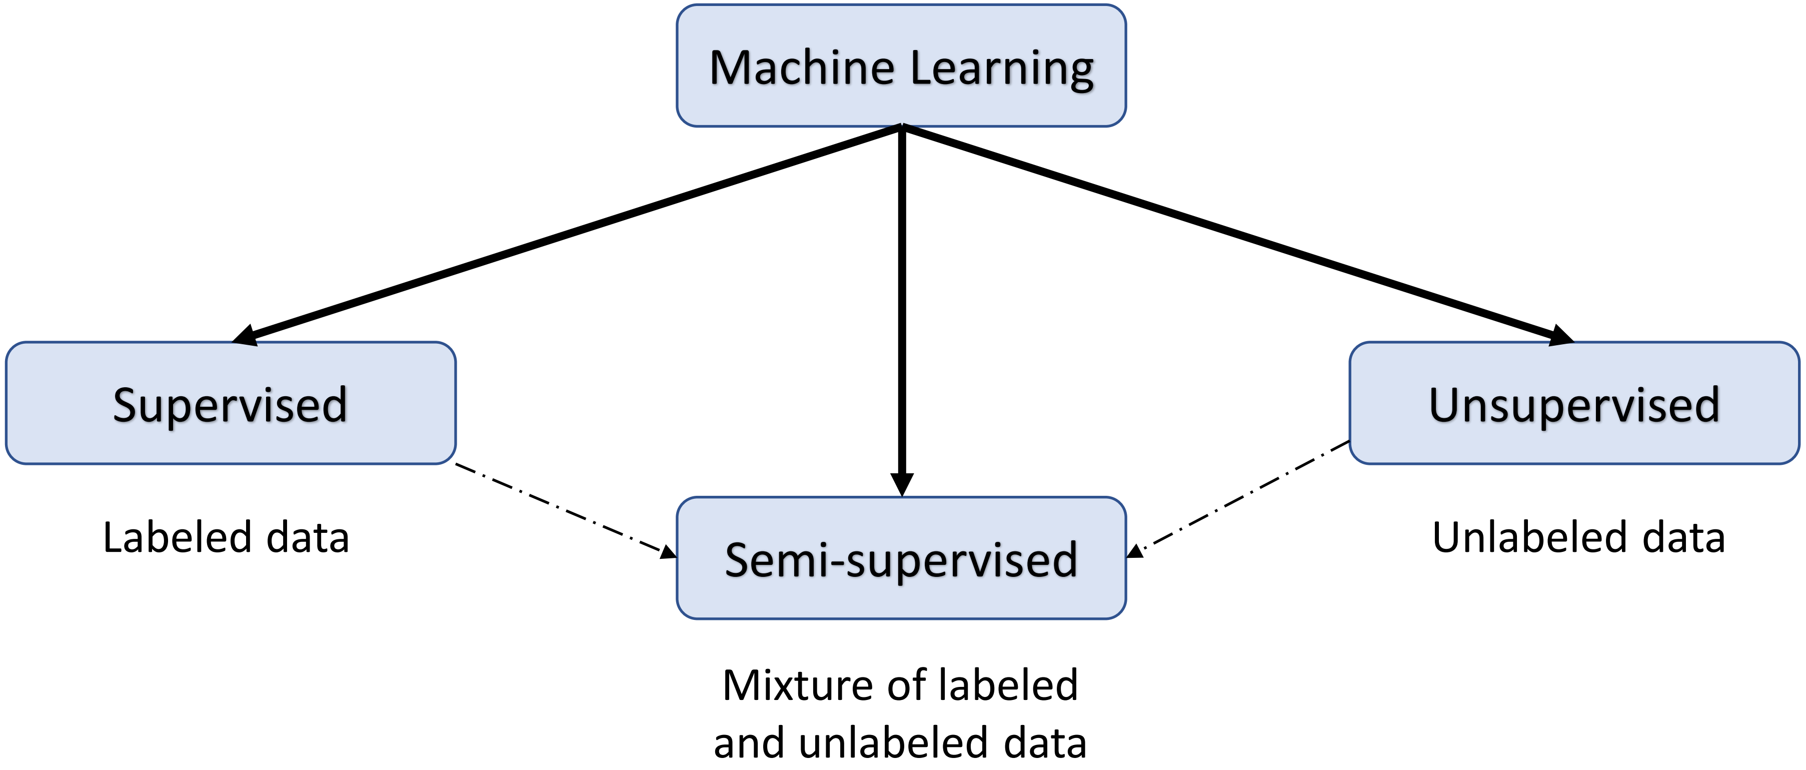

Background and motivation

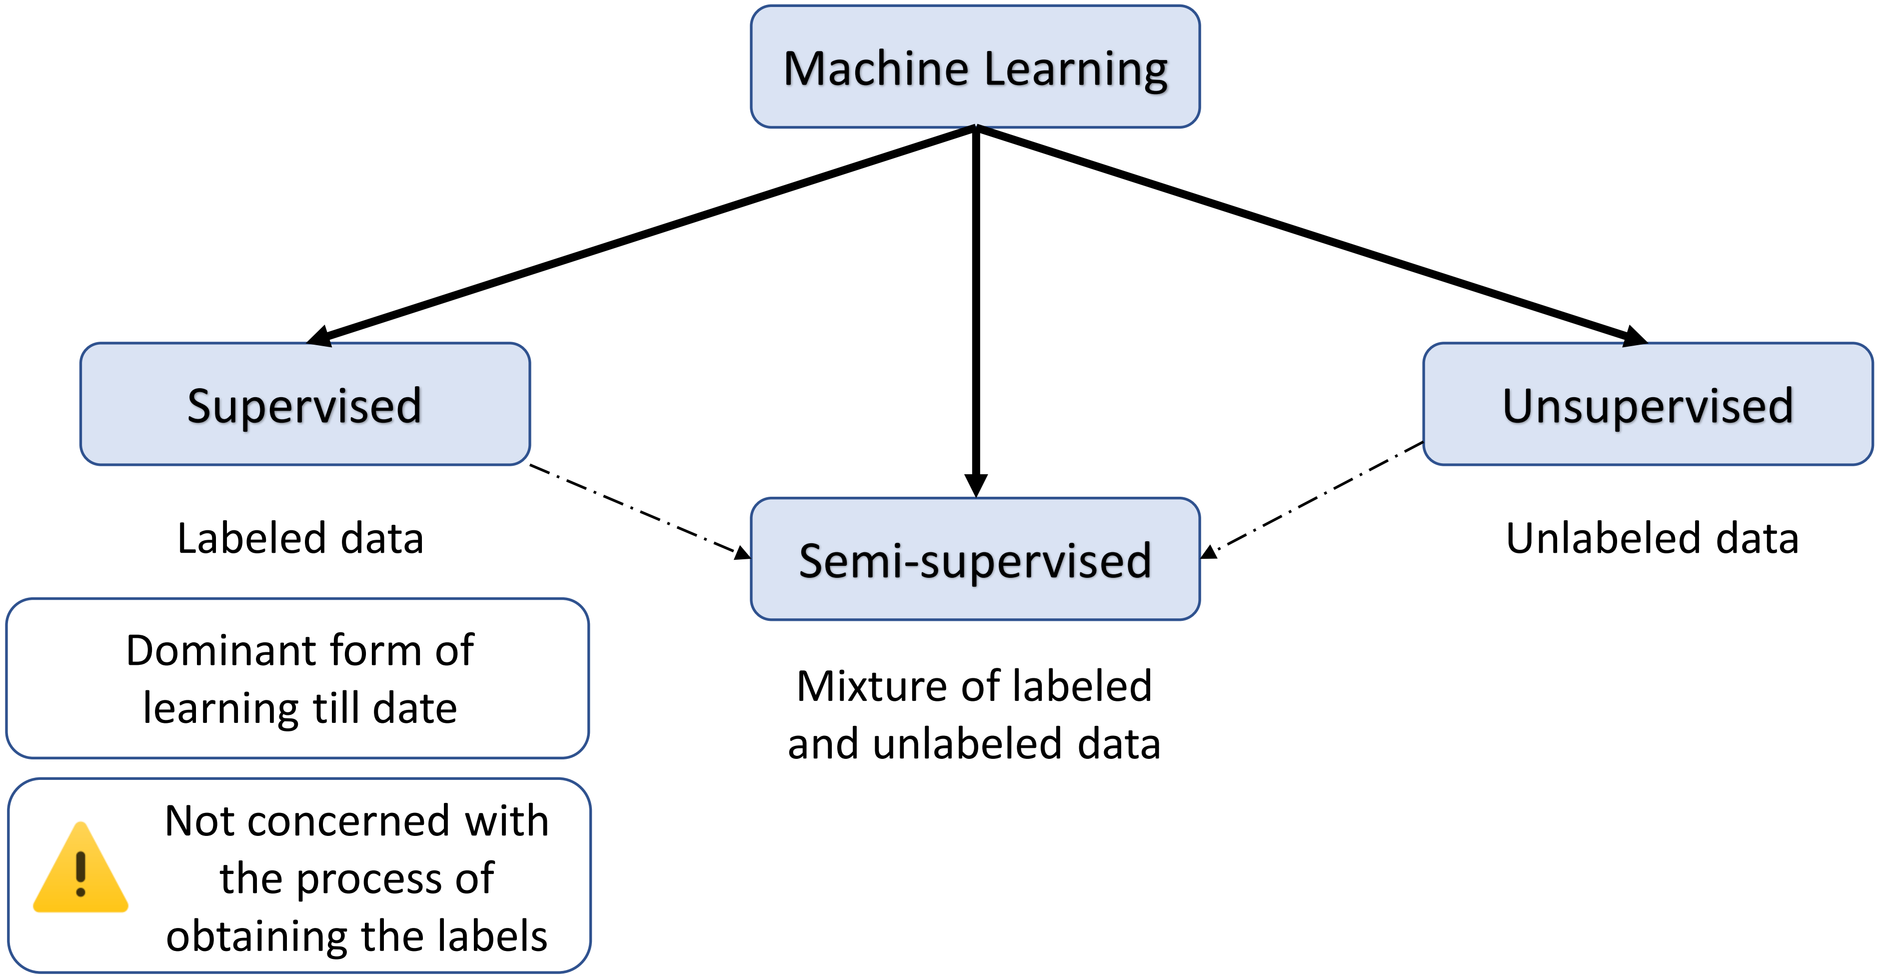

Background and motivation

- ML models perform well with **sufficient amount of training data with labels**
  - Available datasets becoming outdated in size and density

- Need for new data annotation
  - Data growing constantly
  - Laborious

Background and motivation

- ML models perform well with **sufficient amount of training data with labels**
  - Available datasets becoming outdated in size and density
- Need for new data annotation
  - Data growing constantly
  - Laborious

Query for human feedback on key samples

If the machine can learn human knowledge, it will help deal with the lack of training data.

Background and motivation

- ML models perform well with **sufficient amount of training data with labels**
  - Available datasets becoming outdated in size and density

- Need for new data annotation
  - Data growing constantly
  - Laborious

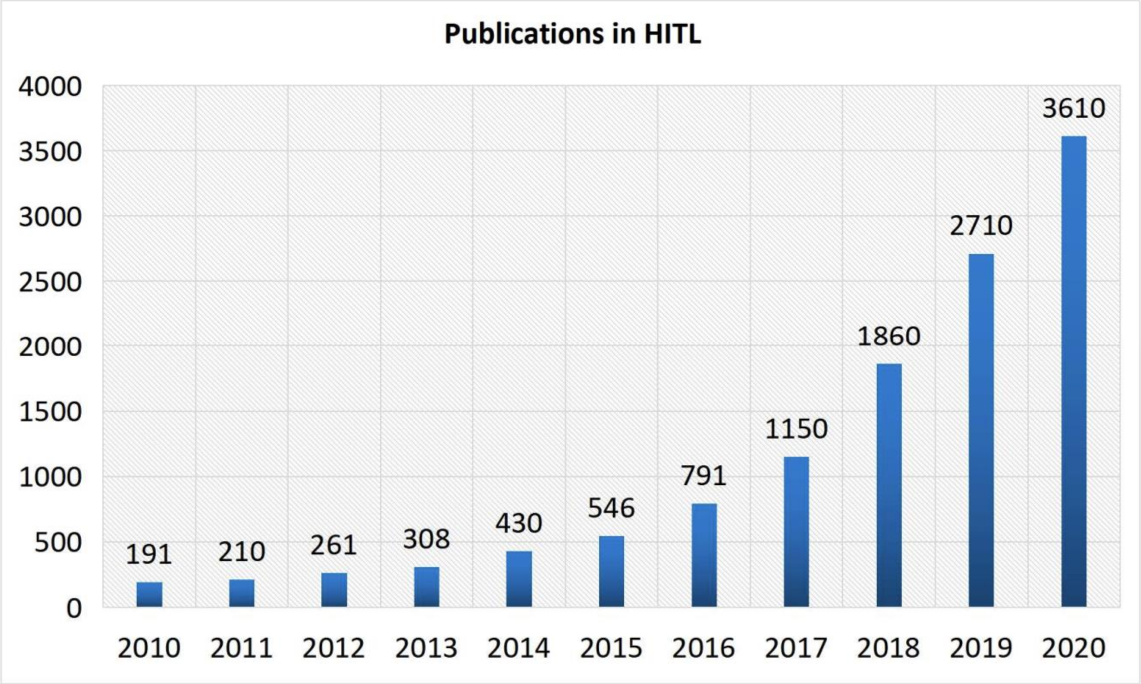

Query for human feedback on key samples

If the machine can learn human knowledge, it will help deal with the lack of training data.

Where can humans intervene?

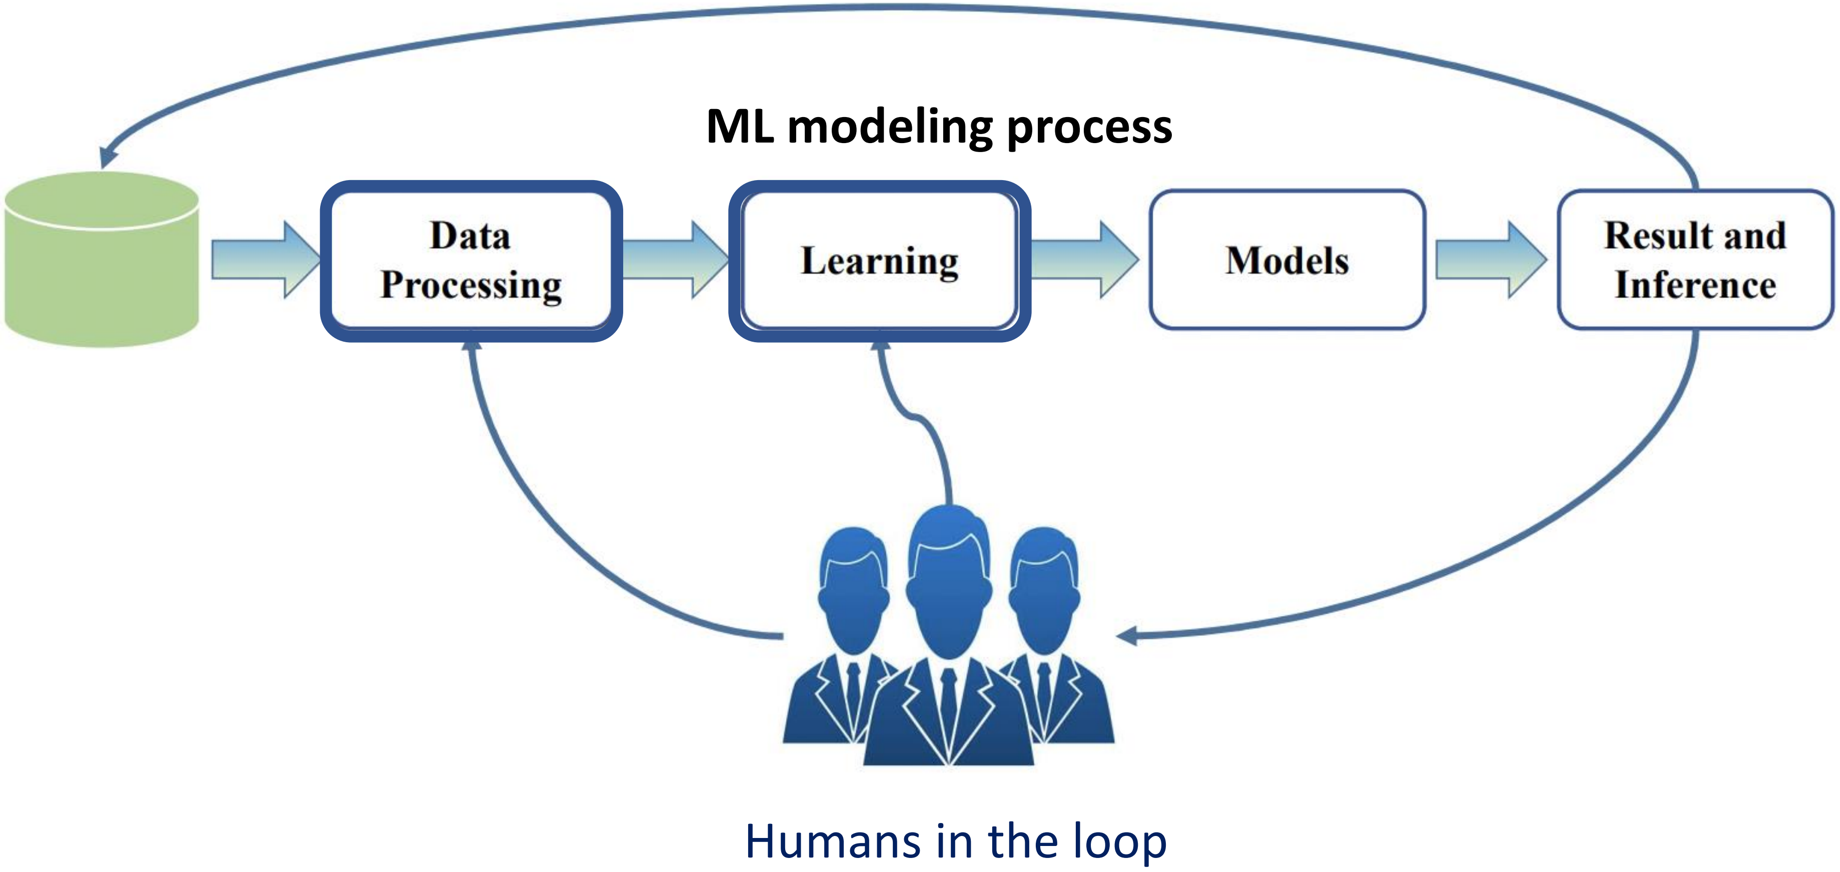

Data Processing

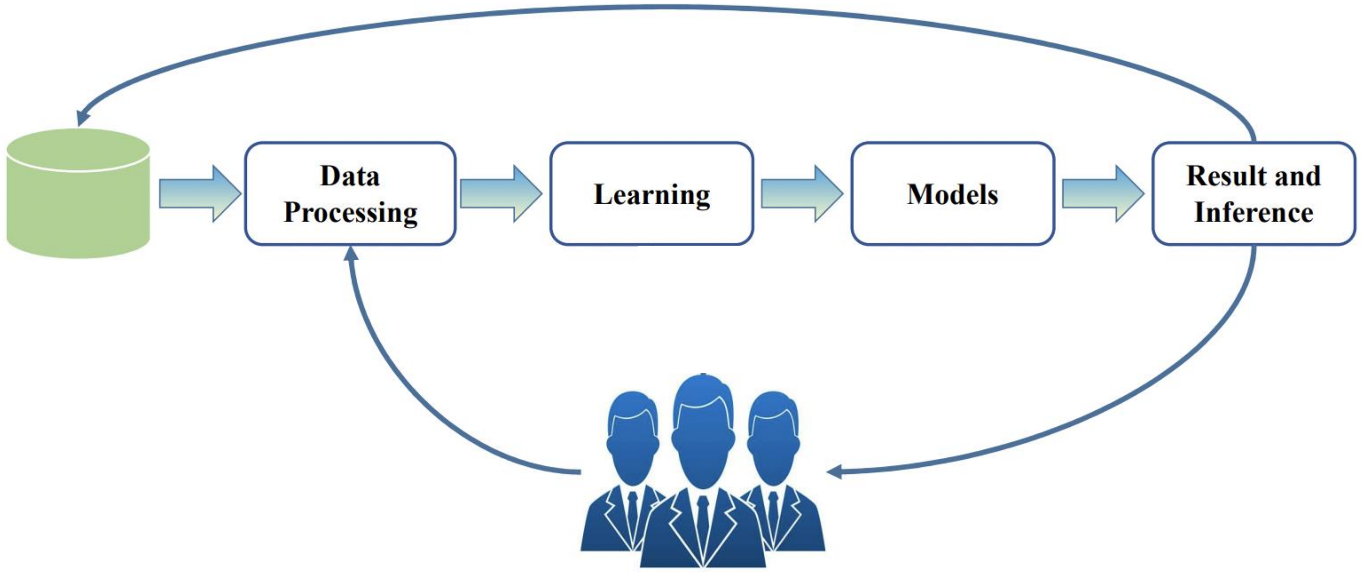

Data pre-processing

- Data scientist spend about 80% of their time on data processing
- Challenge: orchestration of data pre-processing tasks is **difficult**, **unguaranteed to be optimal**

Goal: propose an adaptive data preparation approach to **learn from humans the optimal sequence of data preparation tasks**.

Method: Learn2Clean. Active reinforcement learning-based approach where humans are introduced to adapt/select the sequence of data pre-processing tasks.

Application field: Web data mining

Laure Berti-Équille, CIDR 2020

Data pre-processing

- Challenge: identify incorrectly labelled training data

Goal: identify incorrect labels in influential corpora used to train state-of-the-art models for NER

Method: Semi supervised approach to flag potentially-incorrect labels in the corpus then manual review of these labels by humans in the loop

Application field: NLP, Name entity recognition (NER)

Cutler et al. 2020

2. Data annotation

- Challenge: annotation of **new collected data**

Goal: release the limitations of pre-labeling and upgrades the model with continuously collected data

Method: Deep reinforcement active learning to guide agents (models) to dynamically select new training samples annotated by humans

Application field: CNNs for pedestrian re-identification

Liu et al. 2019

2. Data annotation

- Challenge: **incorporate human intelligence** through data annotation

Goal: incorporate human intelligence to generate high-quality comics

Method: Data-driven generation of comics from digital illustration components (line drawings, irregularities, texture..) annotated by an artist

Application field: Comics generation

Zhang et al. 2021

3. Iterative labeling

- Challenge : improve user experience from direct to dynamic data labelling

Goal: improve user experience in labelling new data

Method: partially automated labeling scheme for data annotation with deep learning and HITL

Application field: CV, visual recognition

Yu et al. 2015

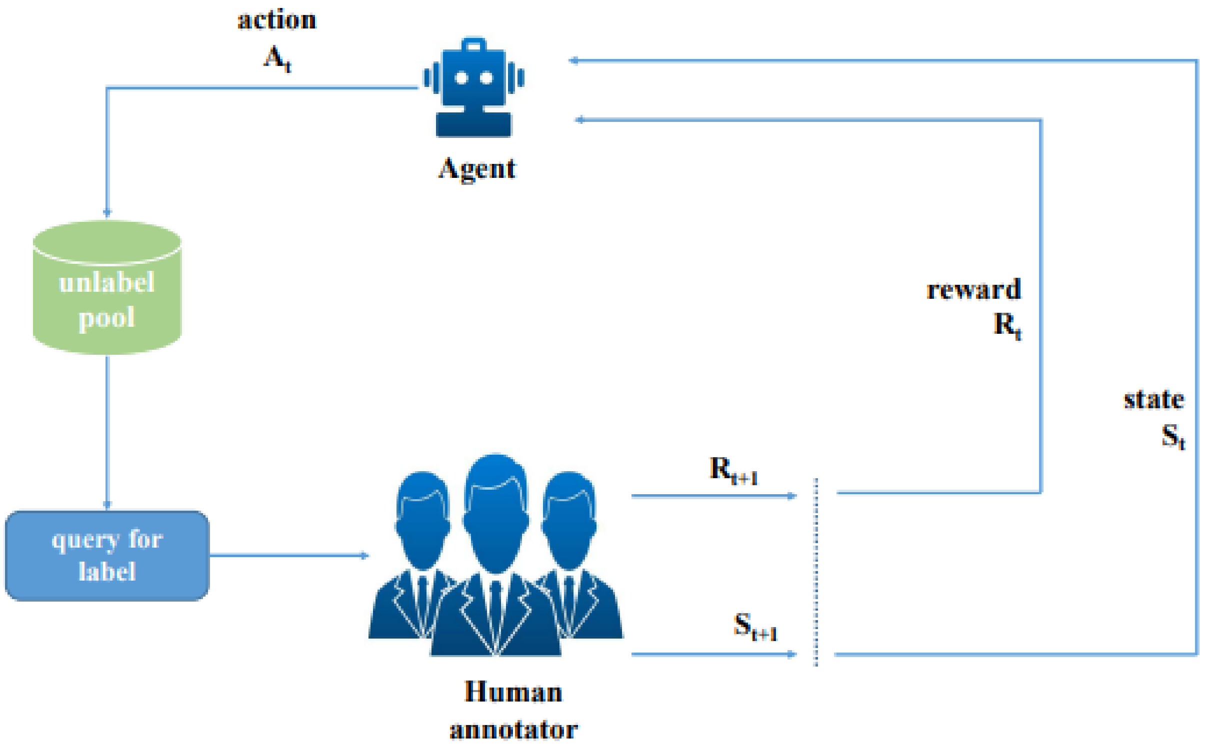

HITL framework based on reinforcement learning

Model training and inference

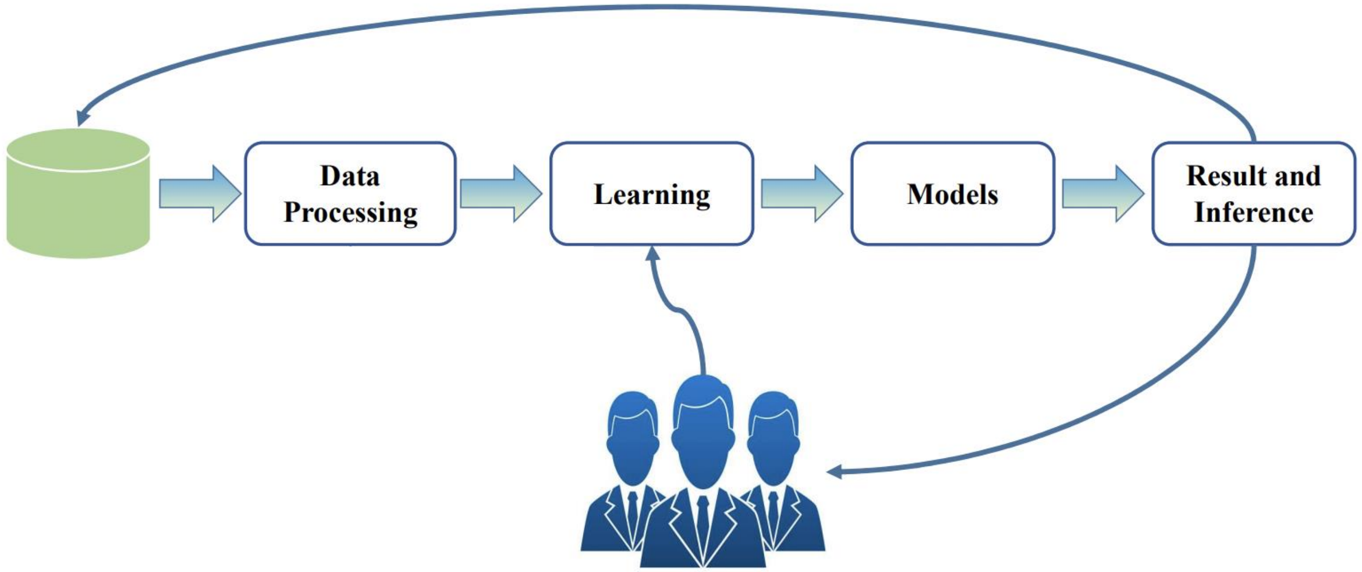

Natural Language Processing

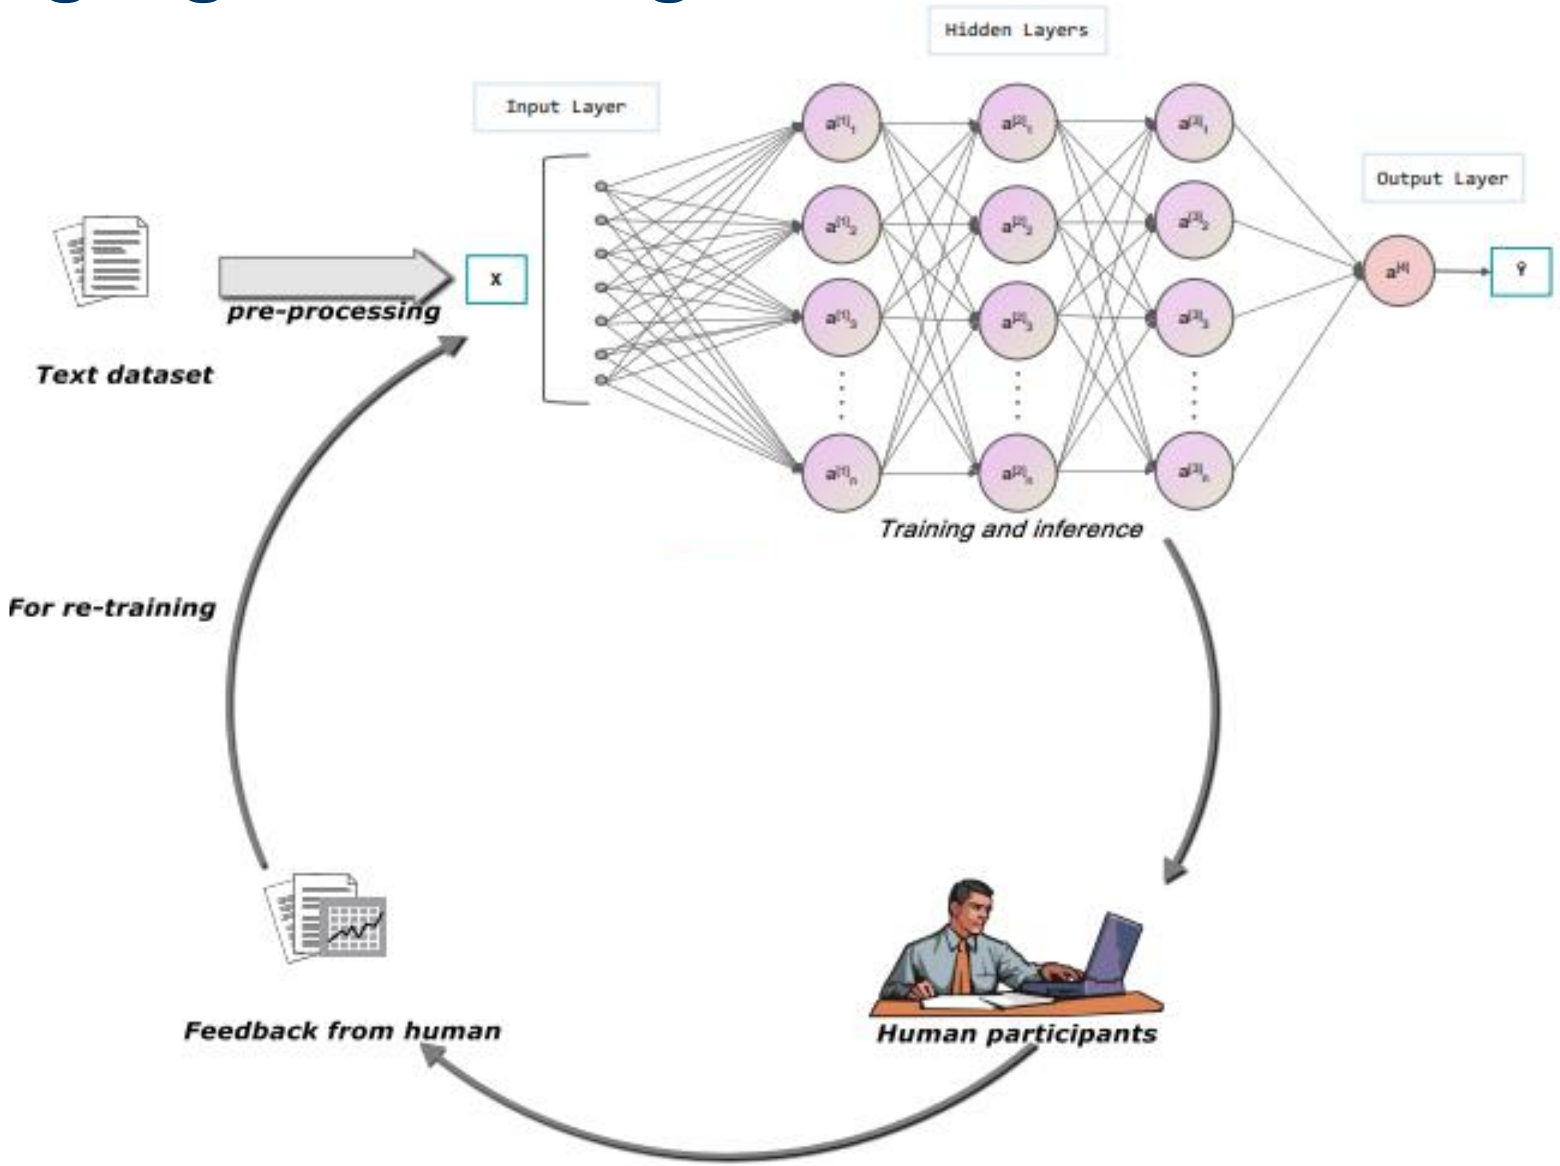

Natural Language Processing

Work	Task						Motivation		
	TC	SSP	TM	TS	QA	SA	Performance	Interpretability	Usability
Zaidan et al. (2007) [77]	✓						✓	✓	
Zhang et al. (2016) [78]	✓						✓	✓	
Arous et al. (2021) [79]	✓						✓	✓	
Karmakharm et al. (2019) [76]	✓						✓		✓
He et al. (2016) [48]		✓					✓		
Su et al. (2018) [84]		✓					✓	✓	✓
Yao et al. (2019) [83]		✓					✓	✓	✓
Yao ZiYu et al. (2019) [85]		✓					✓	✓	
Hu et al. (2014) [87]			✓				✓		
Smith et al. (2018) [88]			✓				✓	✓	✓
Kim et al. (2019) [89]			✓				✓		✓
Ziegler et al. (2019) [90]				✓			✓		
Stiennon et al. (2020) [91]				✓			✓		
Hancock et al. (2019) [92]					✓		✓		
Wallace et al. (2019) [61]					✓		✓		✓
Liu et al. (2021) [93]						✓		✓	

Natural Language Processing

Text classification

Goal: learn from additional feedback from journalists

Method: retrain a rumour classification system with an updated dataset with user-provided annotations

Application field: Rumour Classification

Karmakharm et al. 2019

Goal: improve interpretability

Method: incorporate constraints defined by humans into the loss function of the classifier

Zaidan et al. 2007

Question-Answering

Goal: learn from continuous additional feedback from users

Method: self-feeding mechanism to generate new examples in the train set

Application field: Chatbots

Hancock et al. 2019

Computer vision

Work	Task					Motivation		
	OD	IR	IS	IE	VOS	Performance	Interpretability	Usability
Yao et al. (2012) [114]	✓					✓		
Roy et al. (2018) [115]	✓					✓		✓
Madono et al. (2020) [116]	✓					✓		
Roels et al. (2019) [127]		✓				✓	✓	
Weber et al. (2020) [125]		✓				✓		
Wang et al. (2020) [133]			✓			✓		
Ravanbakhsh et al. (2020) [135]			✓			✓		
Kapoor et al. (2014) [136]				✓		✓		
Murata et al. (2019) [137]				✓		✓		
Fischer et al. (2020) [138]				✓		✓		
Benard et al. (2017) [139]					✓	✓		✓
Caelles et al. (2018) [140]					✓	✓		✓
Oh et al. (2019) [141]					✓	✓		✓

OD: Object Detection. IR: Image Restoration. IS: Image Segmentation. IE: Image Enhancement. VOS: Video Object Segmentation

Computer Vision

Object Detection

Goal: learn from human knowledge by correcting a few annotations

Method: interactive object detection tool to ask humans to correct a few annotations for fine-tuning

Application field: CNNs, pedestrian detection

Yao et al. 2012

Image Restoration

Goal: learn to restore incomplete images with the help of human experts

Method: interactive ML system based on Deep Image Prior

Application field: CNNs

Weber et al. 2020

Other fields

- Recommender systems

Goal: extend state-of-the-art recommenders to feedback on pairs of recommendations and explanations to improve future recommendations

Method: ELIXIR incorporate user latent preference vectors

Ghazimatin et al. 2021

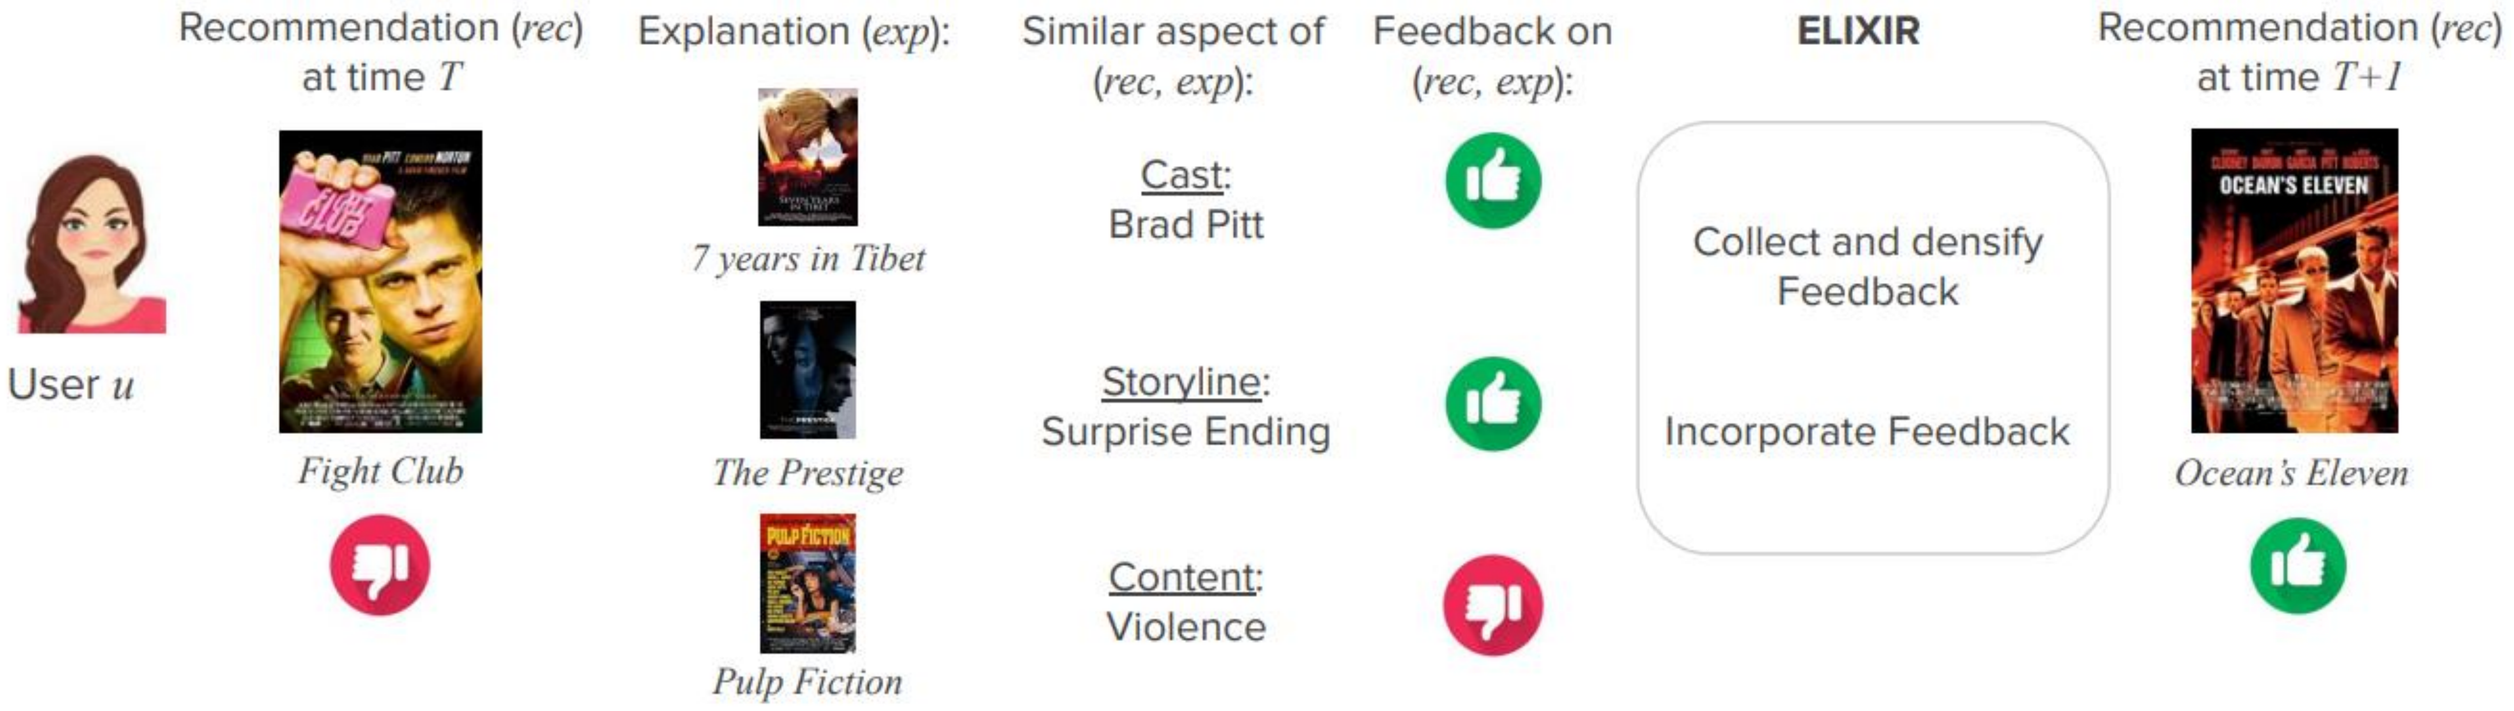

Other fields

- Collaboration between AstraZeneca and Aalto University : **Human-in-the-Loop** for **molecular design**

- REINVENT: *de novo* drug design tool developed by AstraZeneca which trains an agent **to generate new molecules** with specific properties by using **a scoring function F(x).**

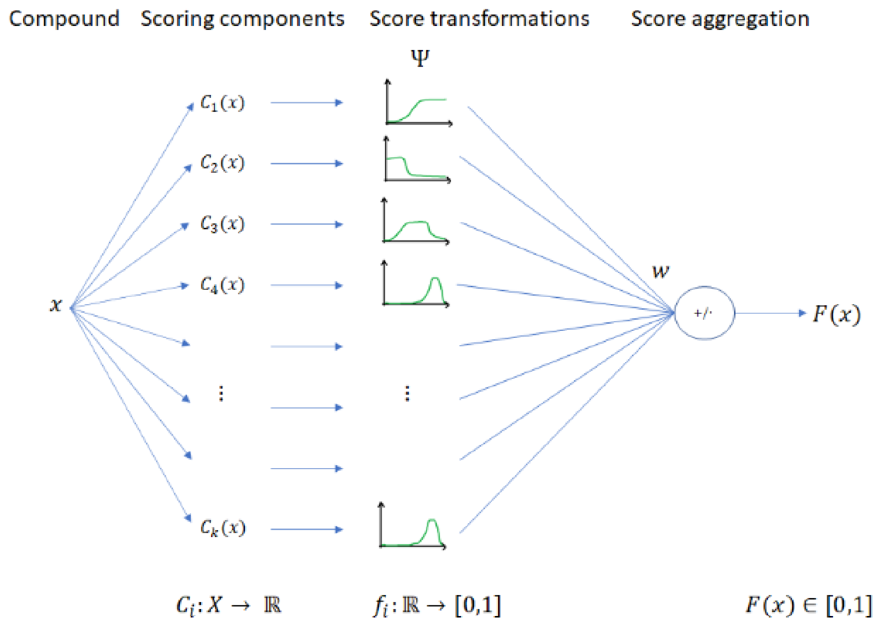

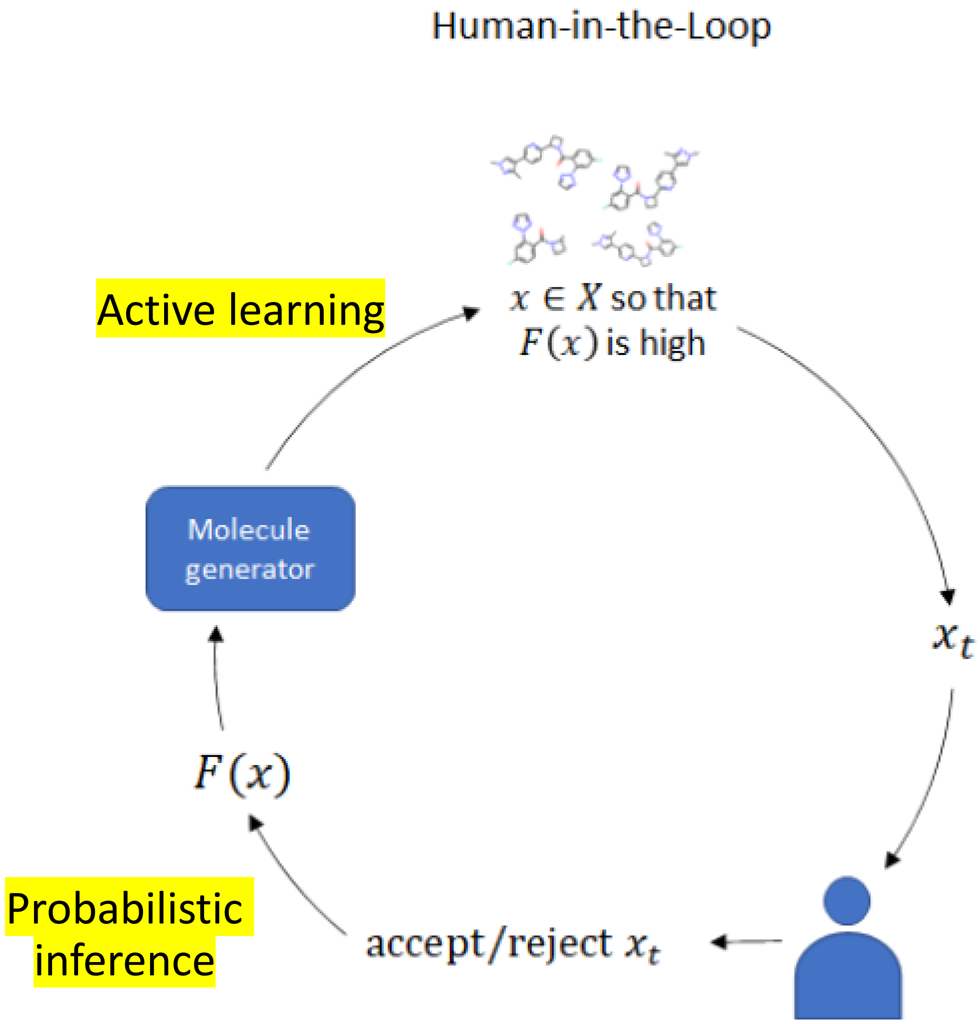

Other fields

- Collaboration between AstraZeneca and Aalto University :
  Human-in-the-Loop for **molecular design**

<u>Goal 1 :</u> adapt the parameters of the scoring function to generate molecules with high QED score

<u>Goal 2 :</u> learn human knowledge as a separate scoring component to generate molecules with high activity on DRD2 receptor

Sundin et al.

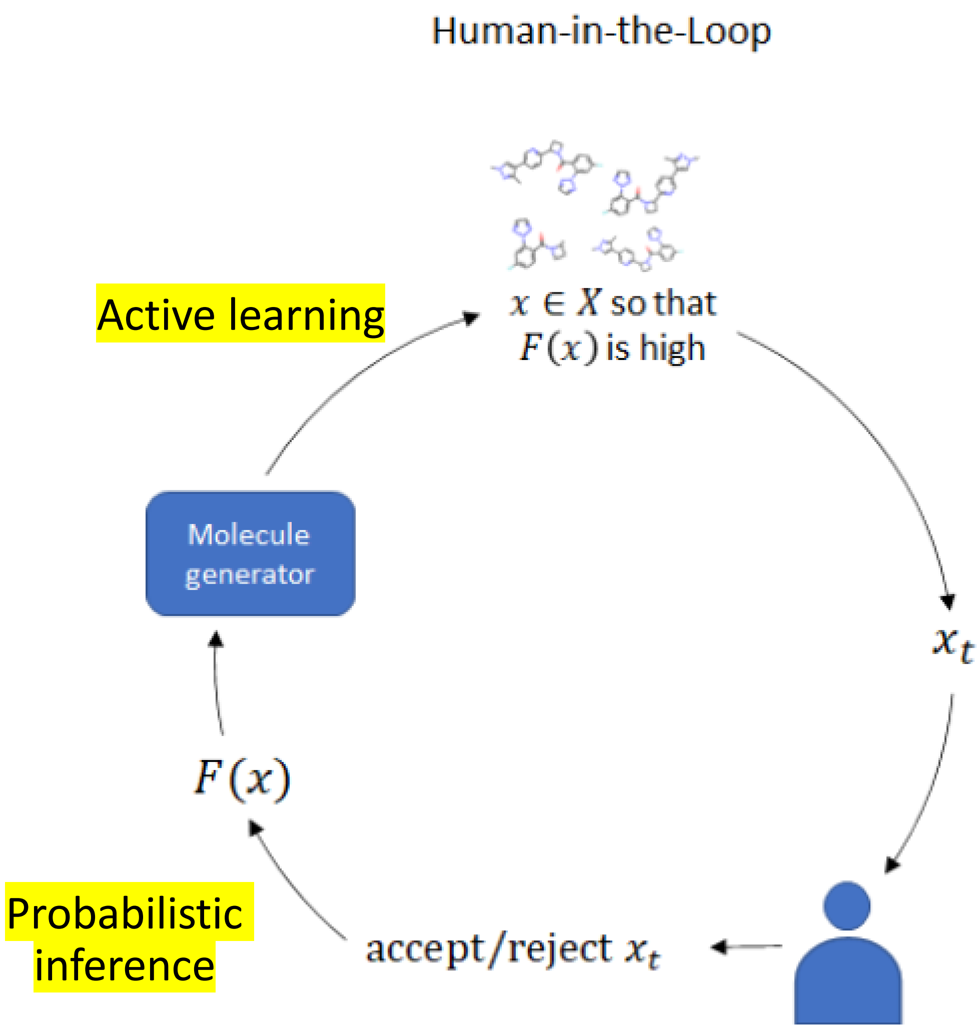

Pros and cons

- Enable to label new data points
- Human feedback can be a valuable input

- Data labeling and constant feedback are costly

Example where human input can be valuable

Open questions and challenges

- How to analyse the quality of provided feedback?
- How to pick the most representative feedback?
- How to display what the model learned from this feedback?
- How to choose an appropriate human intervention time?
- Computational challenges

Future directions

- ➢Besides data annotation, how can the model learn from user experience?
- ➢Explore other methods for selecting key samples
- ➢No uniform standard for HITL benchmarks
- ➢Hope of achieving a universal model through HITL fine-tuning

Thank you!