Lifelong Learning and Continuous Knowledge Learning in Chatbots

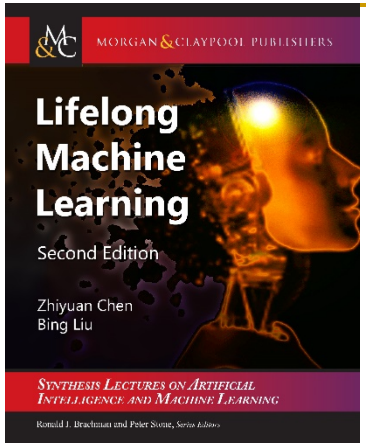

Bing Liu
Department of Computer Science
University of Illinois at Chicago
liub@uic.edu

Introduction

- Learning is key to all human and AI agent activities
- Classic machine learning
  - Isolated single-task learning

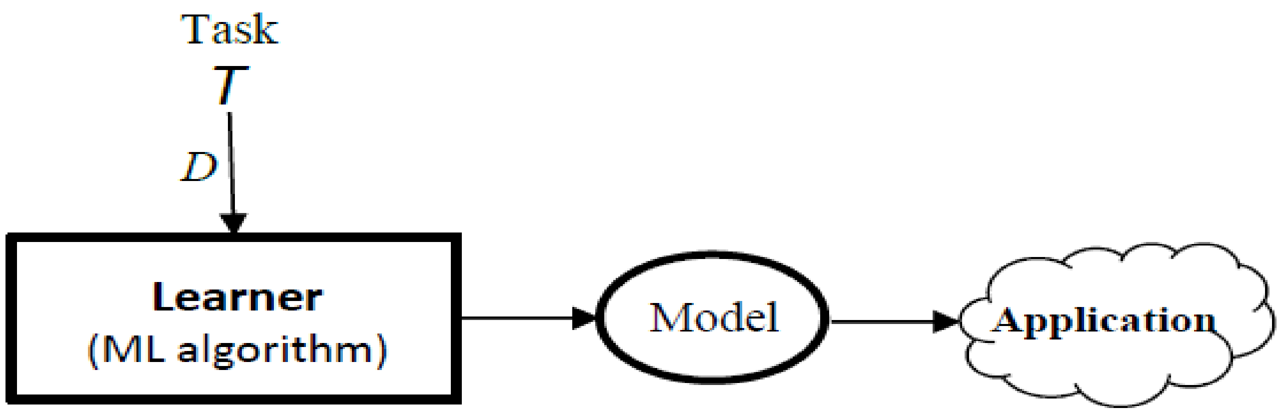

- Knowledge learned not accumulated or used in new learning
  - Needs a large number of training examples: cannot expect humans to label everything in the world.
    - Manual labeling of data has become an industry! No education bg needed.
  - Suitable for well-defined tasks in restricted and closed environment

Lifelong learning
(Chen and Liu, 2016-book)

- Humans never learn in isolation or from scratch. We
  - learn continuously and accumulate knowledge learned in the past and use it to learn more & to learn better.
  - learn in open environments in a self-supervised manner

- *Lifelong Learning* (LL): imitate this human learning capability
- **Goal:** Create a machine that learns like humans
  - Without it, a system will never achieve
    - human-level intelligence or Artificial General Intelligence (AGI)

Practical applications **need** LL

- Chatbots, personal assistants, self-driving cars, and other physical robots, working in real-life environments need LL.
  - Chatbots will not be intelligent if they cannot learn more knowledge after they are deployed
    - Impossible to know what people may say.
    - Little of our ability to converse about any subject matter is learned formally.
  - Self-driving cars are not going to fly with only rules and off-line training
- They face the real open world. No closed-world assumption.
  - They have to continuously learn and accumulate knowledge and adapt to new situations in a self-supervised manner.

Outline

- What is lifelong learning?
- Continual learning and meta-learning
- Open world learning
  - Its use in Intelligent Personal Assistant
- Learning on the job in the open world
  - Continuous knowledge learning in chatbot
- Summary

Definition of lifelong learning (LL) (existing)

(Thrun 1995, Silver et al 2013; Chen and Liu, 2016 –book)

- The learner has performed learning on a sequence of tasks, from 1 to *N*.
- When faced with the (*N*+1)th task, it uses the relevant knowledge in its *knowledge base* (KB) to help learn the (*N*+1)th task.
- After learning (*N*+1)th task, KB is updated with learned results from (*N*+1)th task.

Definition of lifelong learning (LL) (new)

(Thrun 1995, Silver et al 2013; Ruvolo & Eaton, 2013, Chen & Liu 2014, Fei et al 2016, Shu et al 2017a, 2017, **Chen & Liu, 2018)**

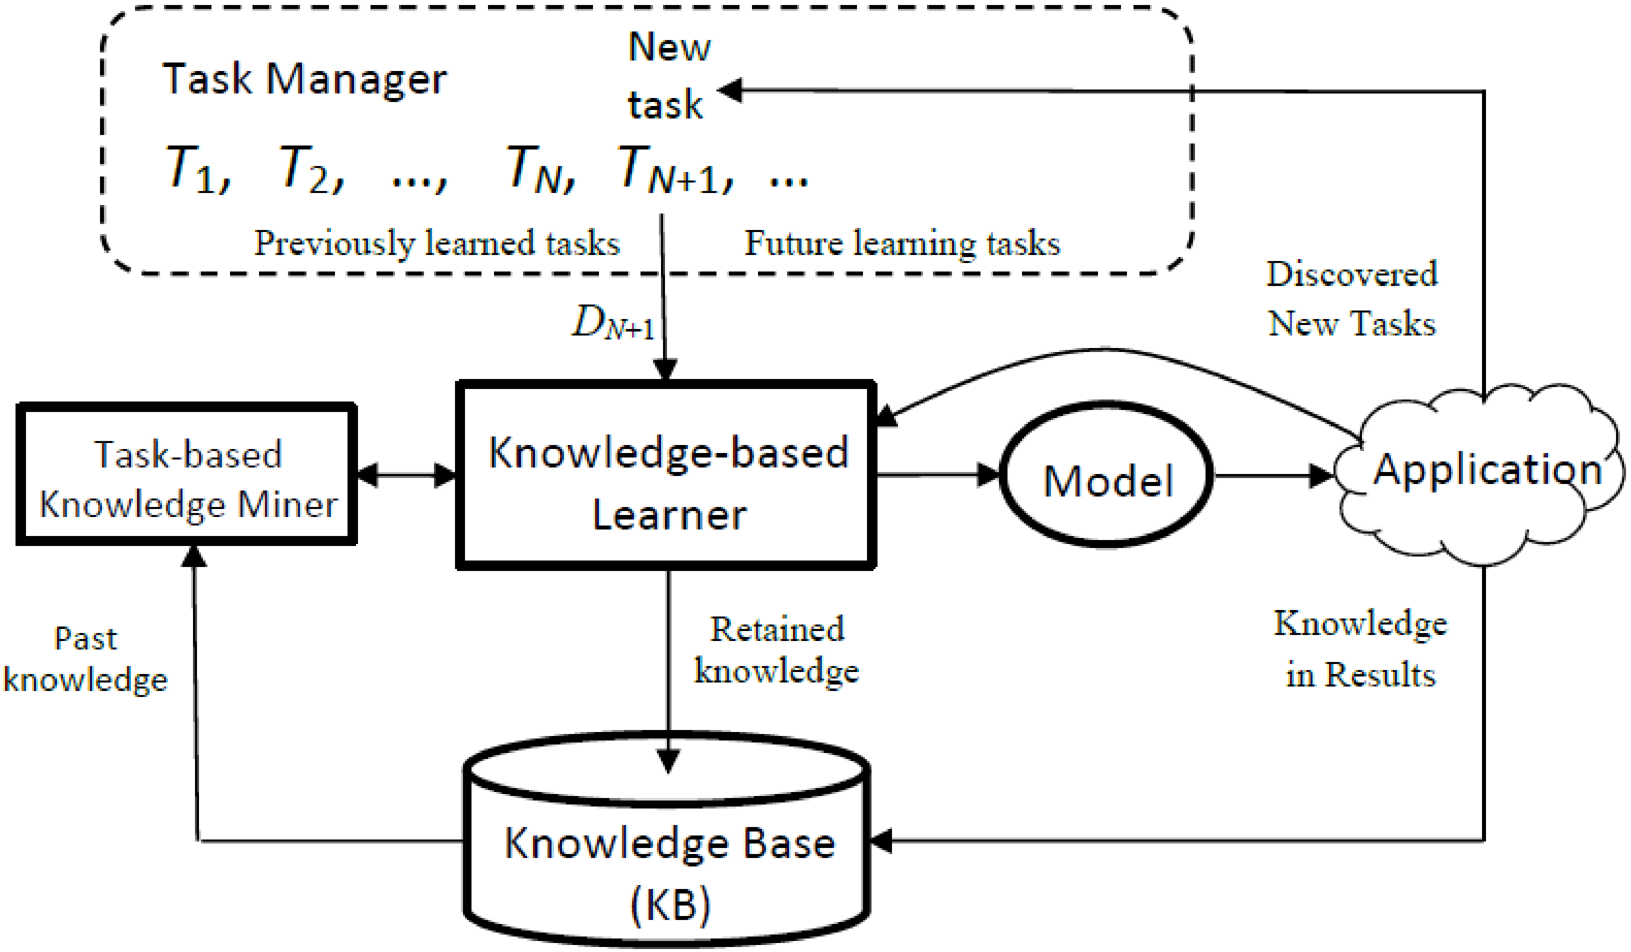

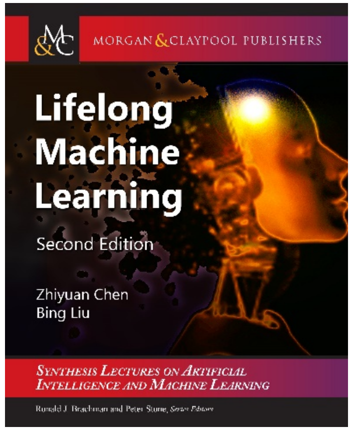

2nd edition

Key characteristics of LL
(Chen and Liu, 2018-book)

1. **Continuous learning process** (w/o supervision)
2. **Knowledge accumulation in KB** (long-term memory)
3. **Using/adapting the past knowledge** to help future learning
4. **Learning in the open world**, discovering new tasks & learning them incrementally in an *interactive environment* with self-motivation
5. **Learning on the job** or learning while working, during model application or testing

- Both (4) and (5) need some form of self-supervision using agents' own knowledge and environmental feedback

Two Types of Shared Knowledge

- *Global knowledge*: These methods assume a *global latent structure* among tasks that are shared (Thrun, 1996, Ruvolo and Eaton, 2013, Bou Ammar et al., 2014, …)
  - Global structure **L**: learned and leveraged in the new task learning.

$$\boldsymbol{\theta}^t = \mathbf{L}\mathbf{s}^t$$

- *Local knowledge:* do not assume a *global latent structure* among tasks (Chen and Liu, 2014a,b, Fei et al., 2016, Liu et al., 2016, Shu et al., 2016, 2017).
  - During the new task learning, they select or meta-mine those pieces of prior knowledge to use based on need.

Approach: shared global knowledge

(Ruvolo & Eaton, 2013)

- Each model's parameter vector θ^t is a linear combination of the weight vector s^t and the basis model components **L** (Kumar et al., 2012).

$$\boldsymbol{\theta}^t = \mathbf{L}\mathbf{s}^t$$

- Initial objective function

$$e_T(\mathbf{L}) = \frac{1}{T}\sum_{t=1}^{T}\min_{\mathbf{s}^{(t)}}\left\{\frac{1}{n_t}\sum_{i=1}^{n_t}\mathcal{L}\left(f\left(\mathbf{x}_i^{(t)};\mathbf{L}\mathbf{s}^{(t)}\right),y_i^{(t)}\right) + \mu\|\mathbf{s}^{(t)}\|_1\right\} + \lambda\|\mathbf{L}\|_{\mathsf{F}}^2 , \qquad (1)$$

Lifelong topic modeling

(Chen and Liu 2014)

- Use results from past tasks to find some topical knowledge
  - e.g., {*price*, *cost*} & {*price*, *expensive*}
  - Use it to guide new task modeling
- Graphical model: same as LDA, but different inference
  - Generalized Pólya Urn Model (GPU)
- Idea: When assigning a topic *t* to a word *w*, also assign *a fraction of t* to words in must-links sharing with *w*.

$$P(z_i = t | \boldsymbol{z}^{-i}, \boldsymbol{w}, \alpha, \beta, \mathbf{A}') \propto$$
$$\frac{n_{d,t}^{-i} + \alpha}{\sum_{t'=1}^{T} (n_{d,t'}^{-i} + \alpha)} \times \frac{\sum_{w'=1}^{V} \mathbf{A}'_{t,w',w_i} \times n_{t,w'}^{-i} + \beta}{\sum_{v=1}^{V} (\sum_{w'=1}^{V} \mathbf{A}'_{t,w',v} \times n_{t,w'}^{-i} + \beta)}$$

Open problems/challenges

- Is the past knowledge actually correct?
- Is the past knowledge actually applicable?
- What is past knowledge and how to represent it?
- …

- LL forces us to think about the issue of knowledge and the role it plays in learning.
  - knowledge representation, acquisition, reasoning, maintenance, etc.

LL research is ramping up

- Many related topics and names
  - Lifelong learning
  - Never-ending learning
  - Continual learning (continuous learning)
  - Open-world learning
  - Meta-learning
  - Developmental learning (in robotics)
- DARPA program (2018): Lifelong learning machines
- DARPA program (new 2019): Open-world learning (SAIL-ON)

Transfer, Multitask → Lifelong

- Transfer learning: using source domain to help target domain,
  - Learning is not continuous
  - No accumulation of knowledge except data
  - Only one directional: help target domain
- Multitask learning: Jointly optimize multi. tasks
  - No accumulation of knowledge except data
  - Hard to re-learn all when tasks are numerous
- Both no discovery of new problems or learning in testing

Outline

- What is lifelong learning?
- Continual learning and meta-learning
- Open world learning
  - Its use in Intelligent Personal Assistant
- Learning on the job in the open world
  - Continuous knowledge learning in chatbot
- Summary

Continual learning and catastrophic forgetting

- Continual learning (CL) is mainly for solving the *catastrophic forgetting* problem in neural networks.
- Catastrophic forgetting (CF): learning a new task is likely change the weights that have been learned for past tasks,
  - degrading the models for past tasks.
  - like a human brain, ideally we want a network (like a brain) to learn many tasks with little interference (or little forgetting of the past).
- CL mainly aims to solve CF. Unlike lifelong learning, CL usually does not emphasize leveraging the past knowledge.

Meta-learning

- Meta-learning, also called *learning to learn*, is often used in one-shot or few-shot learning.
- It trains a meta-model with a large set of tasks. Each task has a set of labeled examples.
  - The learned model can quickly adapt to a new task using only a few examples (few-shot learning).
  - Basically, meta-learning treats all these tasks as training "instances."
- Assume: training tasks and test tasks - same distribution
  - But in real life, new tasks fundamentally different in some aspects.

Outline

- What is lifelong learning?
- Continual learning and meta-learning
- Open world learning
  - Application in Intelligent Personal Assistant
- Learning on the job in the open world
  - Continuous knowledge learning in chatbot
- Summary

Open-world learning

(Fei et al 2016, Chen and Liu 2016, 2018)

- Traditional learning makes the closed world assumption:
  - What the agent sees in training is what it will see in testing.
  - Nothing new or nothing unexpected.
- Clearly not true in the real-world (**open**) environment.
  - Self-driving cars will see things it has never been trained for.
  - A chatbat will definitely hear something that it does not know.
- Can the agent detects unknowns and learn them?
  - If so, it will become more and more knowledgeable.

Open-world learning
(Fei et al., 2016; Shu et al., 2017)

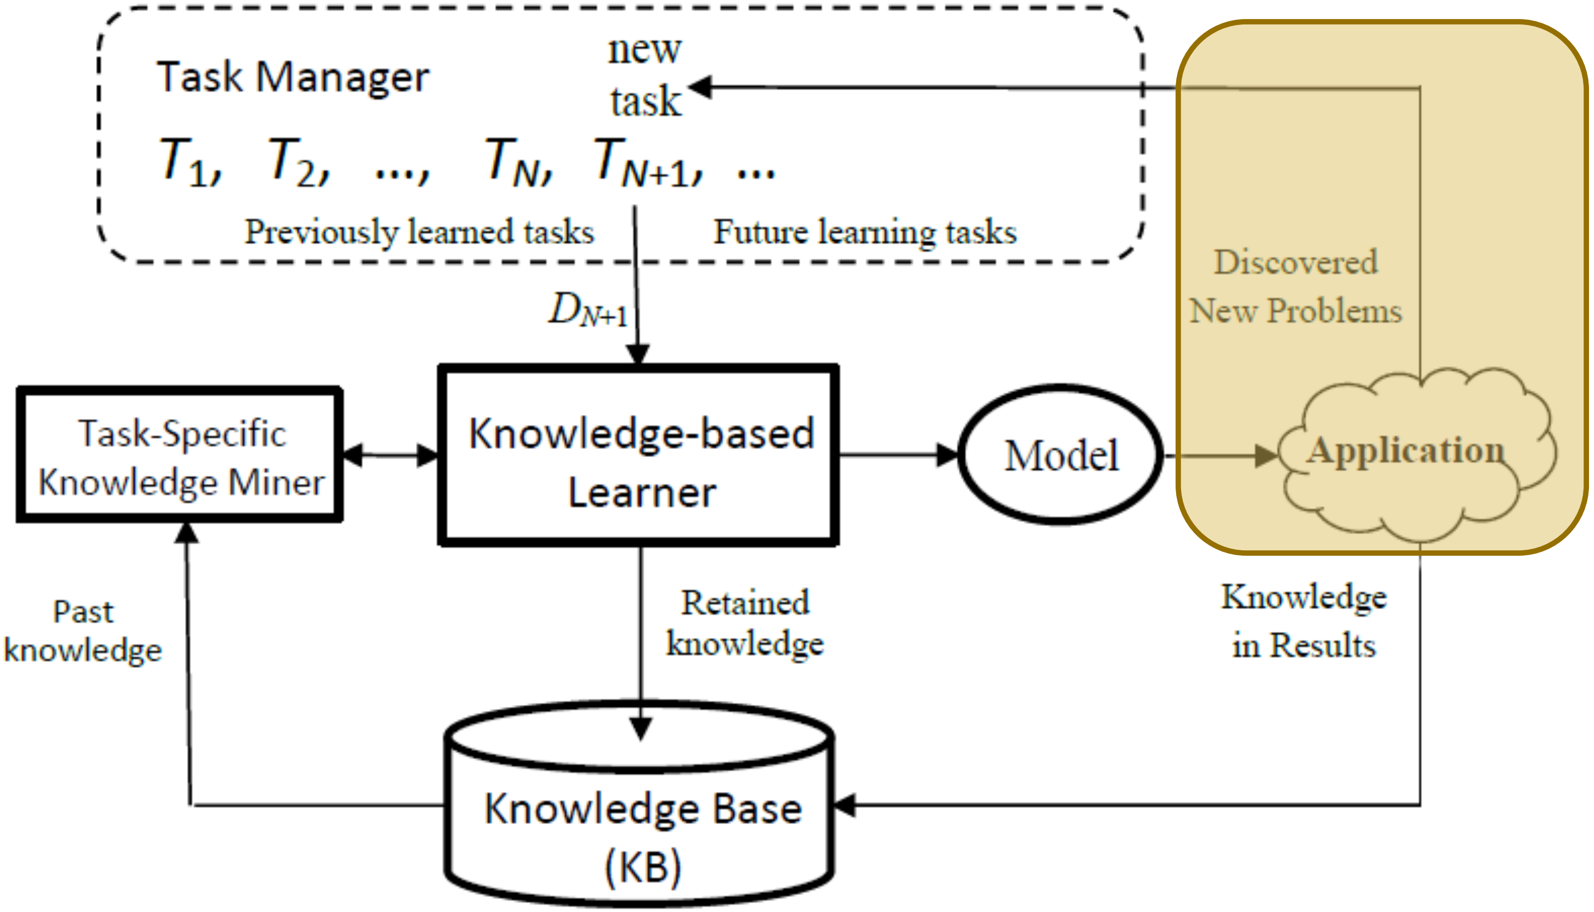

Open-world Learning (OWL)

(Shu, Xu and Liu 2018)

- At any point in time, the learning system is aware of a set of seen classes $S = \{c_1, \ldots, c_m\}$ and has an OWL model or classifier for S but is unaware of a set of unseen classes $U = \{c_{m+1}, \ldots\}$ (any class not in S can be in U) that the model may encounter.
- Goal:
  - Classify seen class instances and detect unseen class instances
  - Group these instances into classes
  - Learn the new classes incrementally

Intelligent Personal Assistant needs OWL

(Xu et al, 2019)

- Intelligent Personal Assistants (IPA) (e.g., Amazon Alexa, Google Assistant, and Microsoft Cortana)
  - The first task is to classify user utterances into existing known intent classes (e.g., Alexa's skills) and
  - detect utterances from unknown intent classes (not supported)
- But, with the support to allow the 3rd-party to develop new skills (Apps),
  - IPAs must recognize new/unseen intent classes and include/learn them in the classification model.

Learning to Accept Classes (L2AC)
(Xu et al, 2019)

- Goal of L2AC is two-fold:
  - (1) classifying examples from classes in S and reject examples from classes in U, and
  - (2) when a new class c_{m+1} (without loss of generality) is removed from U (now U = $\{c_{m+2}, \dots\}$) and added to S (now S = $\{c_1, \dots, c_m, c_{m+1}\}$, still being able to perform (1) without retraining the model.
- L2AC maintains a dynamic set S of seen classes that allows new classes to be added or deleted with no model re-training needed.

L2AC architecture

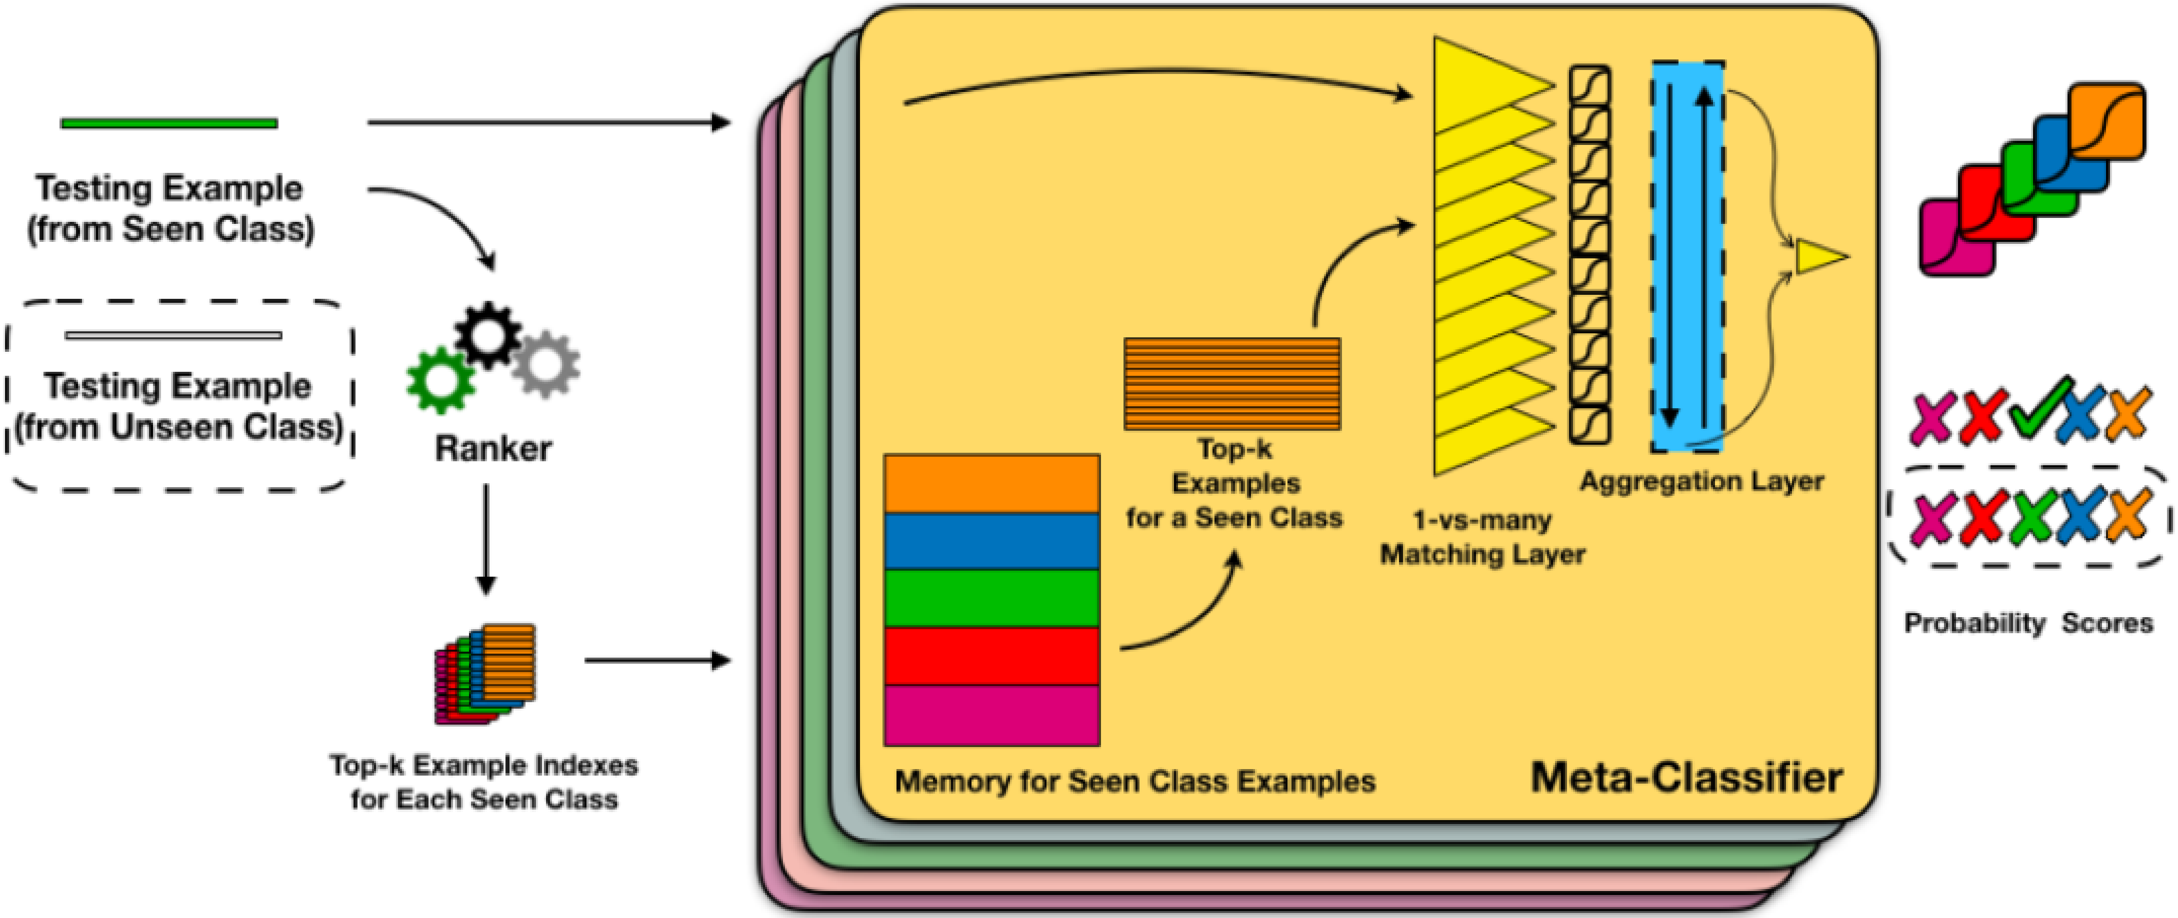

Training the meta-classifier

- Since the meta-classifier is a general classifier that is supposed to work for any class, training the meta-classifier

$$p_\theta(c|x_t, x_{a_{1:k}}|x_t, c)$$

- It requires examples from another set M of classes called meta-training classes.
- A large |M| is desirable to cover of features for seen and unseen classes in testing

DOC: Deep Open Classification

(Shu et al. 2017)

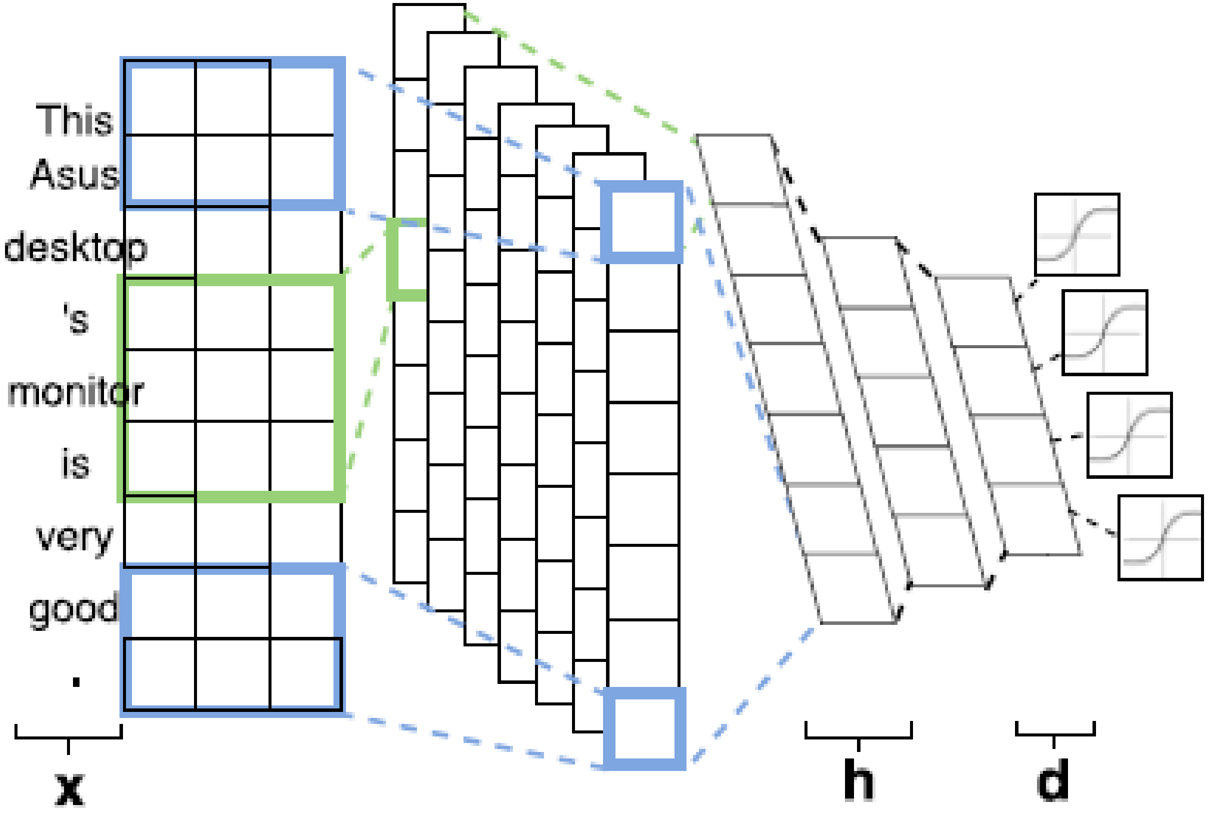

Figure 1: Overall Network of DOC

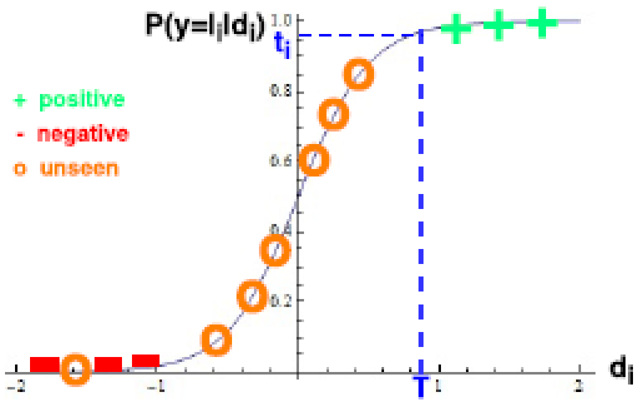

Figure 2: Open space risk of sigmoid function and desired decision boundary $d_i = T$ and probability threshold t_i.

Learning incrementally (cumulatively)

- How to discover classes (or characterize) in the unseen data?
- Incrementally add a class without retraining from scratch
- "Human learning": uses the past knowledge F_t to help learn the new class l_{t+1}.
  - Find similar classes SC from known classes Y^t.
    - Old classes: Y^t = {movie, cat, politics, soccer}.
    - New class: l_{t+1}= basketball
    - SC = {soccer}
  - Building F_{t+1} by focusing on separating l_{t+1} and SC.

Outline

- What is lifelong learning?
- Continual learning and meta-learning
- Open world learning
  - Its use in Intelligent Personal Assistant
- Learning on the job in the open world
  - Continuous knowledge learning in chatbot
- Summary

Learning on the job (while working)
(Shu et al 2017)

- It is known in learning science that about 70% of our human knowledge comes from 'on-the-job' learning.
  - Only about 10% through formal training
  - The rest 20% through observation of others
- An AI agent should be able to learn on the job too as
  - The world is constantly changing.
  - There is always new knowledge, e.g., learning during human-machine conversation (Mazumder et al., 2018)

Lifelong learning extended

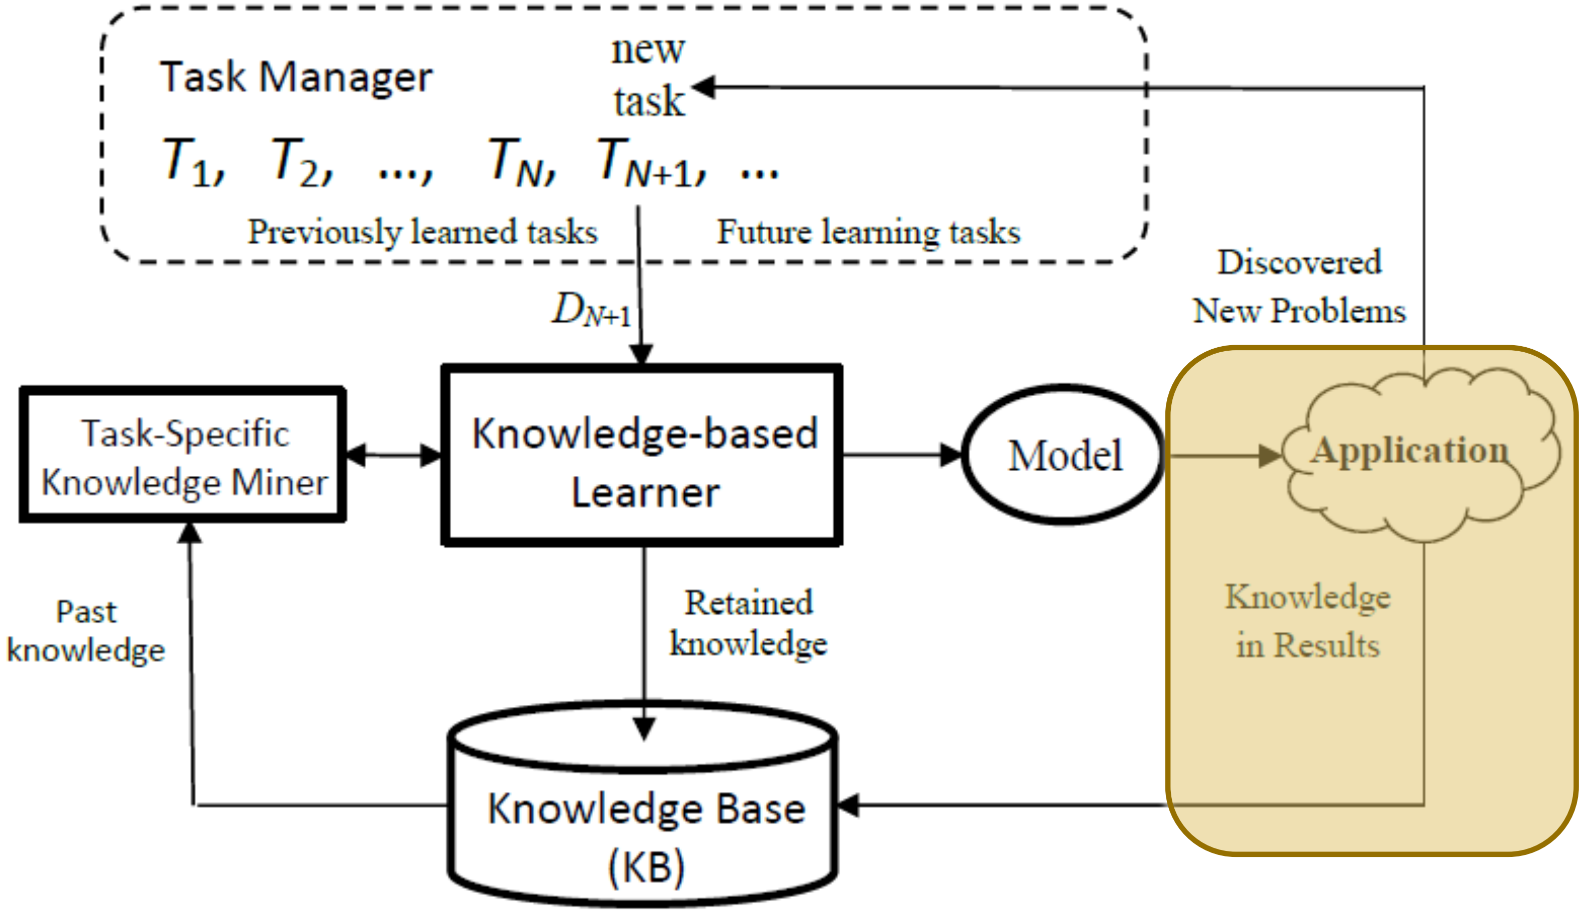

Improving model in testing or application
(Shu et al 2017)

- Can a model's performance be improved after training?
- This paper proposes a technique to do so in the context of CRF for information extraction.
- Idea: connect features with extraction results
  - More results → better features
  - It exploits dependency features
    - As the model sees more data, more features are identified
    - The features help produce better results in the new domain using same model.

Chatbots: Learning during conversation

(Mazumder et al., 2018)

- Chatbots: Early chatbots were mostly built using
  - markup languages, e.g., AIML
  - handcrafted conversation generation rules, and/or
  - information retrieval techniques
- Recently deep learning based approaches are popular
  - They may or may not use explicit knowledge bases (KB)
    - Without using explicit KBs, produce generic/dull responses
    - Use KBs to generate knowledge grounded responses.

Major shortcoming

- One major weakness of existing chat systems is that
  - They do not explicitly or implicitly learn new knowledge in the conversation process.
  - This limits the scope of their applications.
- Learning during chatting enables chatbots to
  - learn new knowledge to expand its KB to become more and more knowledgeable.
  - improve its conversation ability.

Continuous knowledge learning engine

(Sahisnu at al., 2018)

- lifelong interactive learning and inference (LiLi).
- We assume relation extraction, etc, are done by other systems
- Given a query (s, r?, t) from the user, the system has the following 4 capabilities
  - 1. to formulate an inference strategy for a given query that embeds *processing* and *interactive* actions.
    - Processing action: selection of related facts, deriving inference chain, etc., that advances the inference process.
    - Interactive action: deciding what to ask, formulating a suitable question, etc., that enable interaction with the user.

Lifelong interactive learning (contd.)

- 2. to learn interaction behaviors (deciding what to ask and when to ask the user).
- 3. to leverage the acquired knowledge in the current and future inference process.
- 4. to perform 1, 2 and 3 in a lifelong manner for continuous knowledge learning.

- We envision that the proposed system is incorporated in a chatbot in a multi-user environment.
  - The system can perform cross-verification with other users by asking them whether facts) from the particular user is correct.

LiLi: Modeled as Knowledge Base Completion

- We focus on fact triples: e.g., Obama is a citizen of USA
  - (*Obama*, *CitizenOf*, *USA*)
- KB completion (KBC): a binary classification problem
  - Given a query triple, (s, r, t), predict whether the source entity s and target entity t can be linked by the relation r.
  - This problem is solved under the *closed-world* assumption, i.e., s, r and t are all *known* to exist in the KB
- For conversation, we need to remove this closed-world assumption

Open-world knowledge base completion (OKBC)

- The new problem *open-world knowledge base completion*:
  - s, r, or t may not be in the KB.
- 1. Mapping open-world to closed-world.
  - By asking the user questions to convert the open-world problem to a closed-world problem.
  - This acquires some new knowledge and will be added to KB
    - The updated with KB (more knowledgeable)
- 2. Solving the problem in the closed-world (C-PR).
  - Prediction – the prediction may not be correct, but not a problem.

System actions

- If the relation r is unknown, ask for a clue. ask the user to provide a clue (+ve example) c+.
- If an entity s/t is unknown, ask for connecting link (CLQ). Ask the user to connect unknown entities s and/or t with the KB.
  - System will suggest (the most contextually relevant node/entity)
- Ask missing link query (MLQ). Formulate a missing link query (MLQ) to ask the user
- Enumerate path features (for predictor).
- Invoke Predictor.

An example

User:	(Obama, CitizenOf, USA)? "Is Obama a citizen of USA?" [**Query**]
LiLi:	I do not know what "**CitizenOf**" mean? Can you provide me an example? [*Ask for Clue*]
User:	(David Cameron, CitizenOf, UK). "David Cameron is a citizen of UK." [**SF1**]
LiLi:	Got it. Can you tell me how "**Obama**" and "**Honolulu**" are related? [*CLQ*]
User:	(Obama, BornIn, Honolulu). "Obama was born in Honolulu." [**SF2**]
LiLi:	Got it. Can you tell also me how "**Honolulu**" and "**Hawaii**" are related? [*MLQ*]
User:	(Honolulu, CapitalOfState, Hawaii). "Honolulu is the state capital of Hawaii." [**SF3**]
LiLi:	(Obama, CitizenOf, USA) is correct. [**Answer**]

Experiments

Table 4: Dataset statistics [*kwn* = known, *unk* = unknown]

KB Statistics	Freebase (FB)	WordNet (WN)
# Relations ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	1,345 / 1,273	18 / 12
# Entities ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	13, 871 / 13, 223	13, 595 / 13, 150
# Triples ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	854, 362 / 652, 790	107, 146 / 66, 338
# Test relations (*kwn* / *unk*)	25 (17 / 8)	18 (12 / 6)
# Train / valid / test instances	11,260 / 1223 / 7083	6628 / 711 / 3500
Entity statistics in Test Data query triples [s = source entity ; t = target entity]		
Avg. % triples per relation with only s *unk*	16.28	13.29
Avg. % triples per relation with only t *unk*	16.29	10.94
Avg. % triples per relation both s and t *unk*	4.69	4.17

Prediction results

Table 5: Comparison of predictive performance of *various versions* of CKLE [% TTO = % of Test Triples Observed].

% TTO	Models	Freebase									WordNet								
		known Relations			*unknown Relations*			*Overall*			*known Relations*			*unknown Relations*			*Overall*		
		P(+)	F(+)	B-Acc	P(+)	F(+)	B-Acc	P(+)	F(+)	B-Acc	P(+)	F(+)	B-Acc	P(+)	F(+)	B-Acc	P(+)	F(+)	B-Acc
UN-FILTERED [evaluation on all test queries, irrespective of whether user can answer any CLQs or MLQs]																			
10%	F-th	**0.54**	0.40	0.58	**0.48**	0.22	0.52	**0.53**	0.34	0.56	0.75	0.35	0.59	0.52	0.58	0.64	0.59	0.47	0.62
	BG	0.48	0.52	0.62	0.42	**0.50**	**0.53**	0.46	**0.51**	0.59	**0.64**	**0.66**	**0.74**	0.44	0.54	0.57	0.54	**0.61**	0.69
	CKLE	0.51	**0.55**	**0.64**	0.45	0.38	**0.53**	0.49	0.50	**0.60**	0.63	0.58	0.69	**0.61**	**0.6**	**0.68**	**0.62**	0.59	**0.69**
50%	F-th	**0.49**	0.36	0.56	**0.57**	0.41	0.59	**0.52**	0.37	0.57	0.72	0.41	0.61	0.41	0.47	0.57	0.52	0.44	0.60
	BG	0.45	0.48	0.59	0.43	**0.55**	0.58	0.44	0.51	0.59	0.62	**0.64**	**0.73**	0.37	0.51	0.56	0.48	0.58	0.67
	CKLE	**0.49**	**0.52**	**0.62**	0.50	0.52	**0.61**	0.49	**0.52**	**0.61**	**0.63**	0.62	0.71	**0.57**	**0.54**	**0.66**	**0.61**	**0.59**	**0.69**
100%	F-th	**0.50**	0.36	0.56	**0.53**	0.39	0.58	**0.51**	0.37	0.57	**0.73**	0.40	0.61	0.39	0.48	0.56	0.49	0.44	0.59
	BG	0.46	0.50	0.60	0.42	**0.55**	0.60	0.44	**0.52**	0.60	0.62	0.62	0.71	0.36	0.50	0.55	0.48	0.57	0.65
	CKLE	0.49	**0.51**	**0.62**	0.48	0.52	**0.61**	0.49	0.51	**0.62**	0.64	**0.63**	**0.72**	**0.59**	**0.55**	**0.66**	**0.62**	**0.60**	**0.70**
FILTERED [evaluation on test queries for which user has answered atleast one CLQs or MLQs (when asked)]																			
10%	F-th	0.53	0.49	0.55	0.52	0.32	0.51	0.53	0.43	0.53	1.0	0.24	0.56	0.81	0.85	0.7	0.85	0.52	0.58
	BG	**0.58**	**0.62**	**0.62**	**0.56**	**0.65**	**0.58**	0.57	**0.64**	**0.61**	0.86	**0.88**	0.57	**0.90**	**0.95**	**0.83**	0.88	**0.90**	0.68
	CKLE	0.56	0.61	0.60	0.5	0.47	0.51	**0.54**	0.57	0.57	**1.0**	0.70	**0.77**	0.8	0.8	0.65	**0.90**	0.74	**0.71**
50%	F-th	0.49	0.47	0.54	**0.59**	0.46	0.52	0.53	0.46	0.53	0.86	0.43	0.34	0.78	0.75	0.64	0.81	0.57	0.48
	BG	0.51	**0.58**	**0.58**	0.56	**0.68**	0.52	**0.54**	0.57	**0.62**	0.91	**0.92**	0.60	**0.81**	**0.86**	**0.63**	0.87	**0.90**	**0.62**
	CKLE	**0.52**	**0.58**	**0.58**	**0.59**	0.6	**0.53**	**0.54**	**0.59**	0.56	**0.93**	0.78	**0.68**	0.80	0.76	0.58	**0.88**	0.77	**0.62**
100%	F-th	0.52	0.49	0.54	**0.54**	0.45	0.50	0.52	0.48	0.53	0.88	0.45	0.44	0.83	0.82	0.53	0.85	0.61	0.44
	BG	0.51	0.58	0.56	0.53	**0.66**	**0.51**	0.52	**0.61**	0.55	0.91	**0.92**	0.59	**0.87**	**0.91**	**0.66**	0.89	**0.92**	0.62
	CKLE	**0.53**	**0.60**	**0.58**	**0.54**	0.59	**0.51**	**0.53**	0.60	**0.56**	**0.94**	0.83	**0.66**	**0.87**	0.82	**0.66**	**0.91**	0.82	**0.64**

Outline

- What is lifelong learning?
- Continual learning and meta-learning
- Open world learning
  - Its use in Intelligent Personal Assistant
- Learning on the job in the open world
  - Continuous knowledge learning in chatbot
- Summary

Summary

- LL wants to enable an AI agent to learn continuously and to make use of knowledge learned in the past to help future learning and problem solving.
- A large space with huge challenges (Chen & Liu, 2016, 2018 book):
  - Correctness and applicability of knowledge, knowledge representation and reasoning, composition, interactive learning, etc.
- Learning in conversation interactively is promising.
  - Without it, a chatbot will not be intelligent.
  - Next generation of chatbots should have the self-learning capability.

Thank You

Q&A

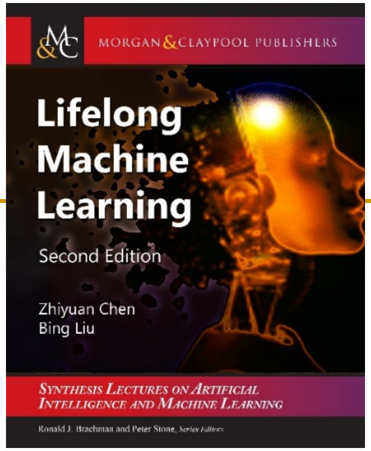

2nd edition

References: https://www.cs.uic.edu/~liub/lifelong-learning.html