MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS
Universität Zürich UZH
Caltech
THE UNIVERSITY OF CHICAGO
ETH zürich

Teaching Multiple Concepts to Forgetful Learners

Yuxin Chen
chenyuxin@uchicago.edu


Anette Hunziker

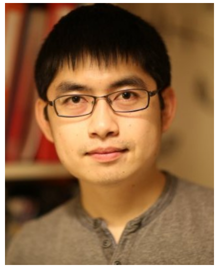
Yuxin Chen


Oisin Mac Aodha


Manuel Gomez Rodriguez


Andreas Krause


Pietro Perona


Yisong Yue


Adish Singla

Applications: Language learning

- Over 300+ million students
- Based on **spaced repetition** of flash cards
- Can we compute **optimal personalized schedule** of repetition?

Teaching Interaction Using Flashcards

Interaction at time $t = 1, 2, \ldots T$

1. Teacher displays a flashcard $x_t \in \{1,2,..,n\}$
2. Learner's recall is $y_t \in \{0, 1\}$
3. Teacher provides the correct answer

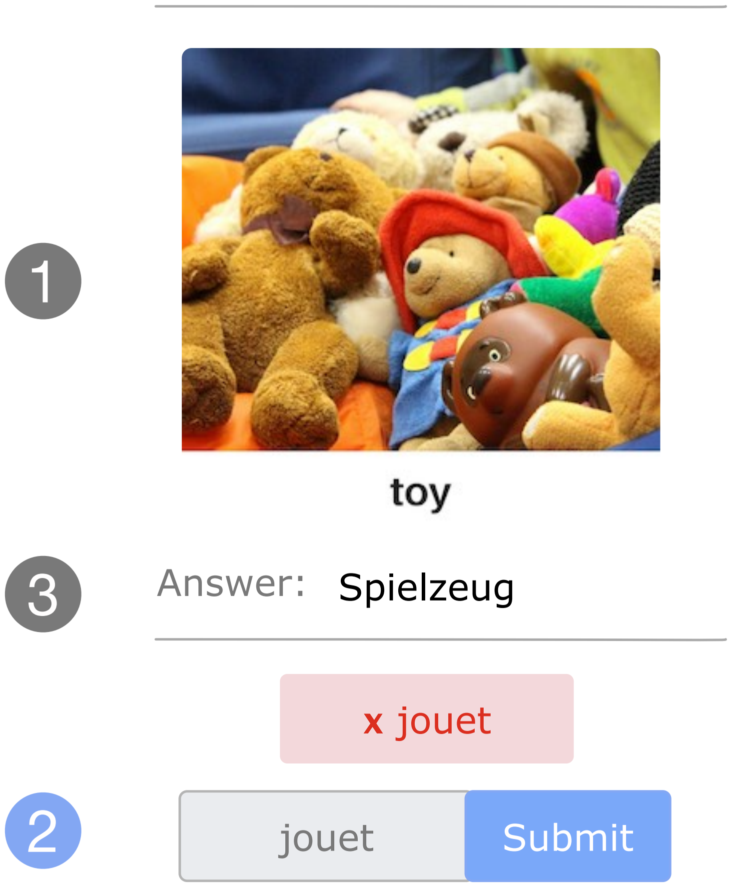

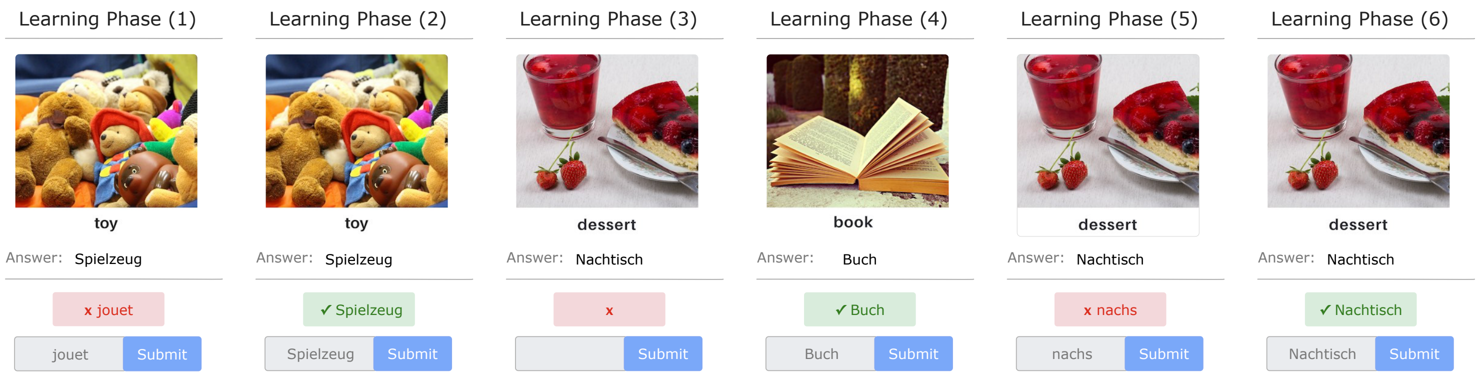

Background on Teaching Policies

Example setup

- $T = 20$ and $n = 5$ concepts given by $\{a, b, c, d, e\}$

Naïve teaching policies

- Random: $a \to b \to a \to e \to c \to d \to a \to d \to c \to a \to b \to e \to a \to b \to d \to e \to$
- Round-robin: $a \to b \to c \to d \to e \to a \to b \to c \to d \to e \to a \to b \to c \to d \to e \to a \to$

Key limitation: Schedule agnostic to learning process

Background: Pimsieur Method (1967)

Used in mainstream language learning platforms

Based on spaced repetition ideas

$$a \to b \to a \to b \to c \to a \to c \to b \to d \to c \to d \to a \to b \to d \to c \to e \to$$

Background: Leitner System (1972)

Adaptive spacing intervals

Key limitation: No guarantees on the optimality of the schedule

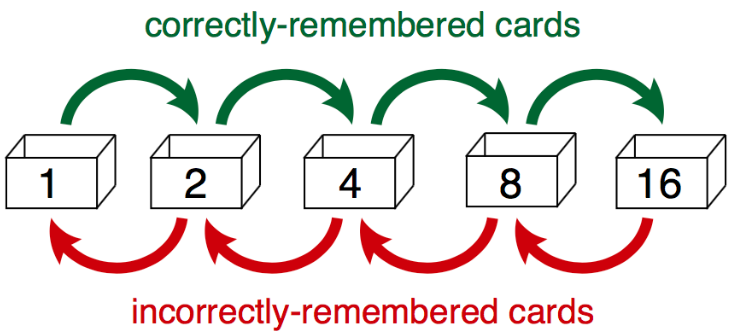

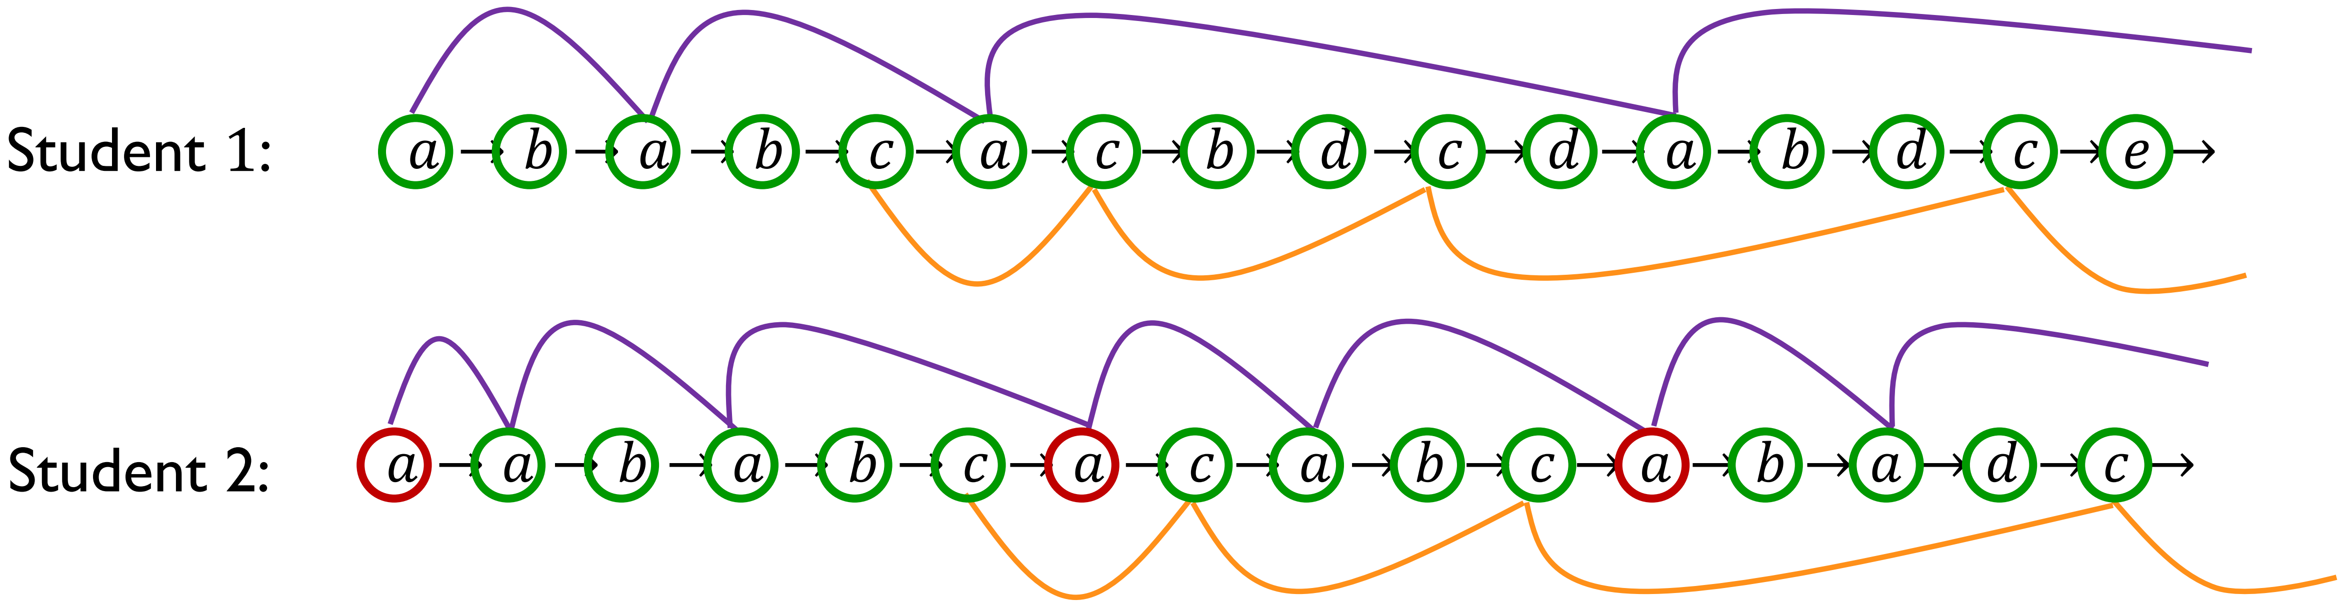

Modeling Forgetfulness

Half-life Regression (HRL) model [Settles & Meeder, ACL 2016]

Time since last teaching concept

Time t

Recall Probability of Concept i:

$$p_i(t \mid \text{history}) = 2^{-\frac{\Delta t_i}{h_i}}$$

Half-life estimate (depends on feedback)

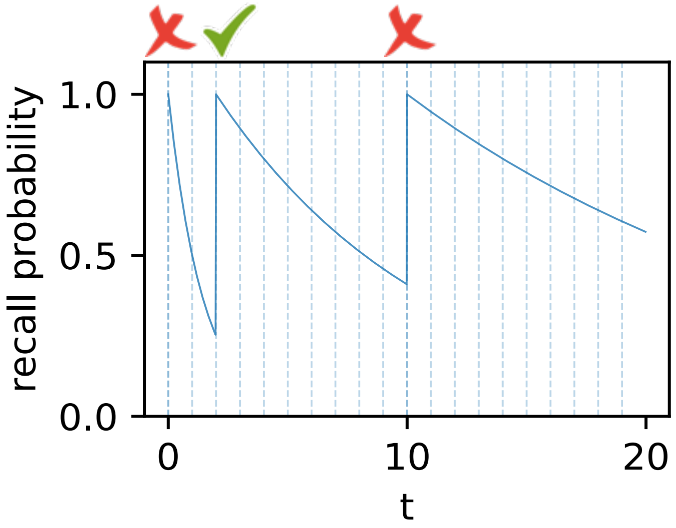

✓ $h_i \mathrel{+}= a_i$

✗ $h_i \mathrel{+}= b_i$

Interactive Teaching Protocol

- For t = 1…T
  - Teacher chooses concept $i \epsilon \{1, \ldots, m\}$ (e.g., a flashcard)
  - Learner tries to recall concept (success or fail)
  - Teacher reveals answer (e.g., "Spielzug")

toy

- Goal: maximize

$$f(\text{history}) = \frac{1}{m}\frac{1}{T}\sum_{i=1}^{m}\sum_{t=1}^{T} p_i(t \mid \text{history}_{1:t-1})$$

"Area Under Curve"

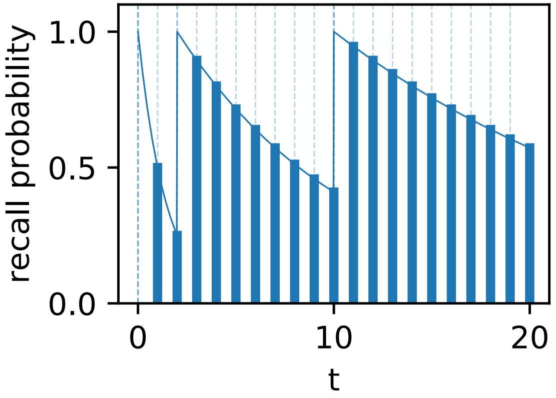

Naive Approaches

- Round Robin
  - Doesn’t adapt to new estimates of learner recall probabilities
  - Over-teaches easy concepts
  - Under-teaches hard concepts

- Lowest Recall Probability
  - Generalization of Pimsleur method and Leitner system
  - Doesn’t consider change to recall probability

Greedy Teaching Algorithm (interactive)

- Choose concept i to maximize

$$\Delta(i|\text{history}) = E_{y_t}\left[f\left(\text{history} \oplus (i, y_t)\right) - f(\text{history})\right]$$

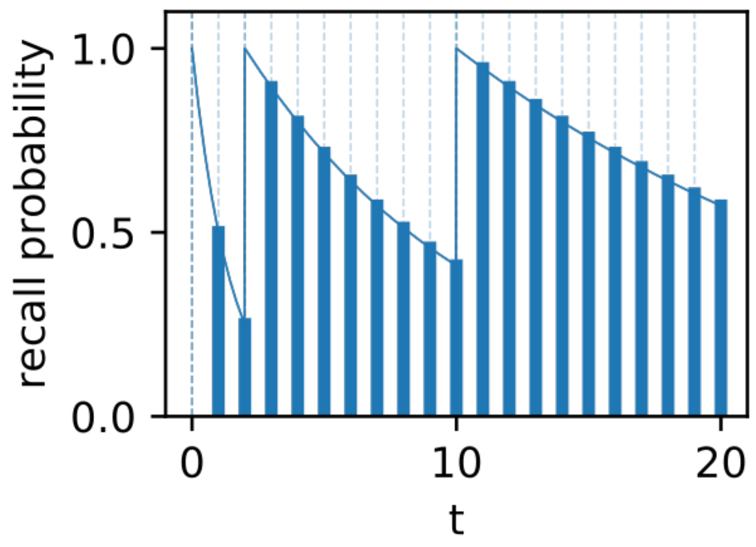

y_t: success or failure of recall at time t

(randomness over model estimate)

$$p_i(t \mid \text{history}) = 2^{-\frac{\Delta t_i}{h_i}}$$

(h_i updated after observing y_t)

Characteristics of the Optimization Problem

- Non-submodular
  - Gain of a concept x can increase given longer history
  - Captured by submodularity ratio γ over sequences

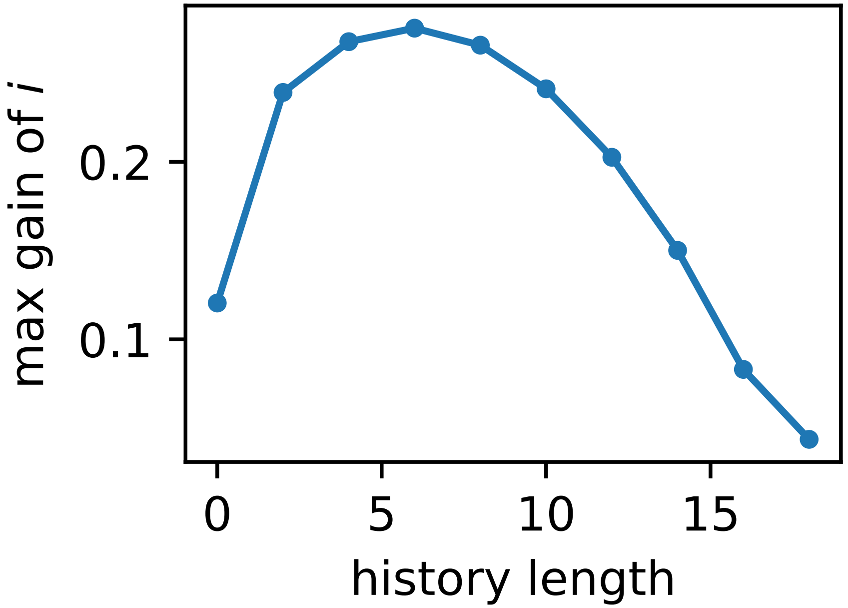

Characteristics of the Optimization Problem (cont.)

- Post-fix non-monotone
  - $f(\text{orange} \oplus \text{blue}) < f(\text{blue})$
  - Captured by curvature ω

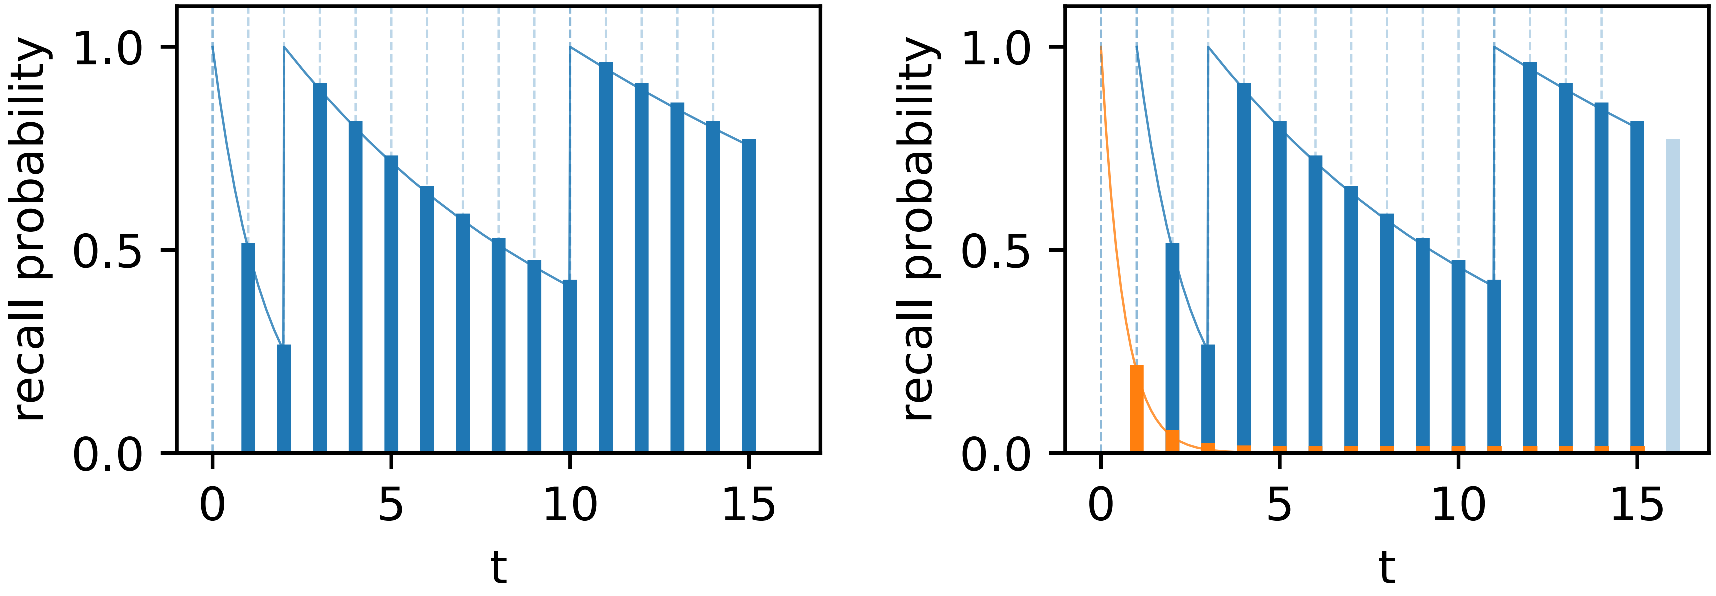

Theoretical Guarantees: General Case

- Guarantees for the general case (**any memory model**)
- Utility of π^{gr} (greedy policy) compared to π^{opt} is given by

$$F(\pi^{\mathrm{gr}}) \geq F(\pi^{\mathrm{opt}}) \sum_{t=1}^{T} \left(\frac{\gamma_{T-t}}{T} \prod_{\tau=0}^{t-1} \left(1 - \frac{\omega_\tau \cdot \gamma_\tau}{T}\right) \right) \geq F(\pi^{\mathrm{opt}}) \frac{1}{\omega_{\max}} (1 - e^{-\omega_{\max} \cdot \gamma_{\min}})$$

Theorem 1

Corollary 2

Theoretical Guarantees: HLR Model

- Consider the task of teaching n concepts where each concept is following an independent **HLR model** with the same parameters $(a^x = z, b^x = z)\ \forall\ x\ \in \{1,2,..,n\}$. A sufficient condition for the algorithm to achieve $(1-\epsilon)$ high utility is

$$\mathrm{z} \geq \max\left\{\log T, \log(3n), \log\left(\frac{2n^2}{\epsilon T}\right)\right\}$$

Illustration: Simulation Results

Greedy

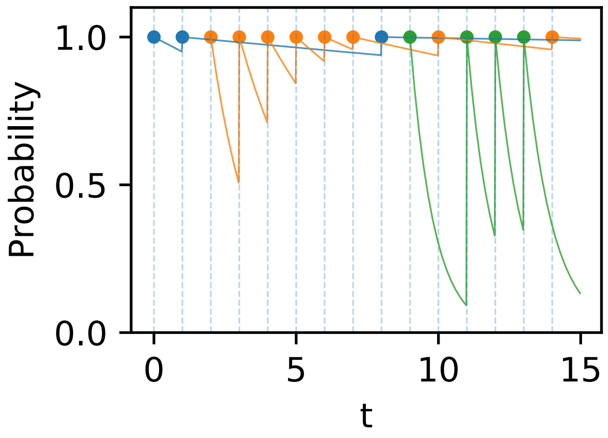

Round Robin

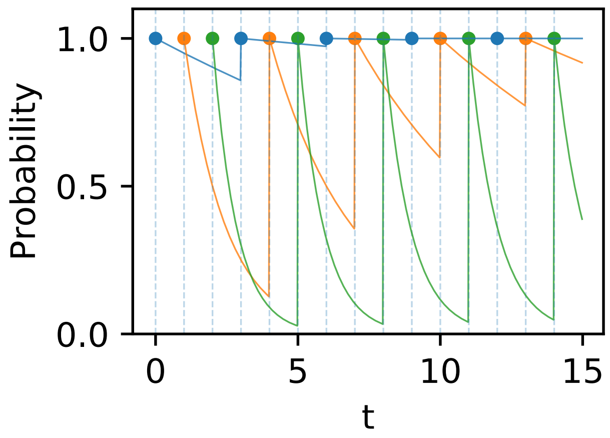

Optimal

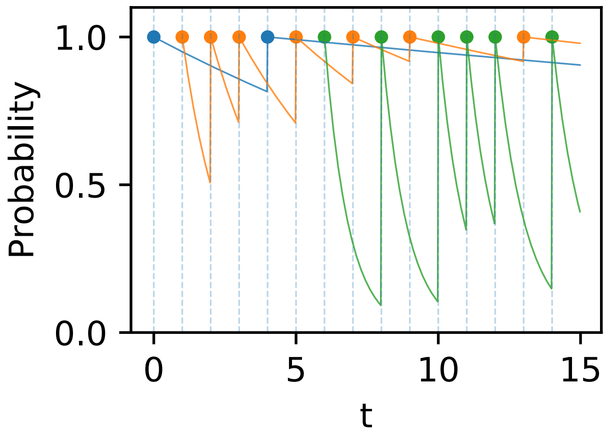

Objective

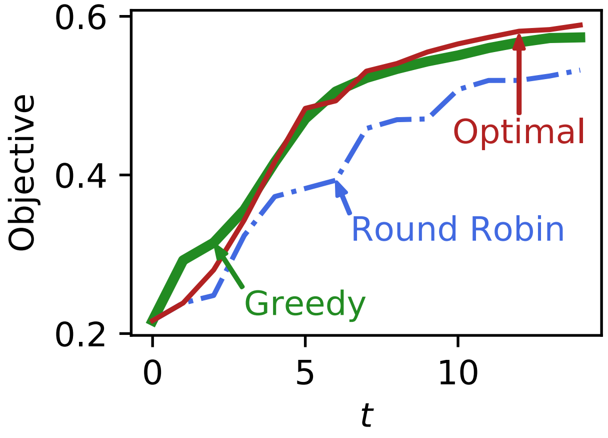

User Study

150 participants from Mechanical Turk platform

T=40, m=15, total study time is about 25 mins

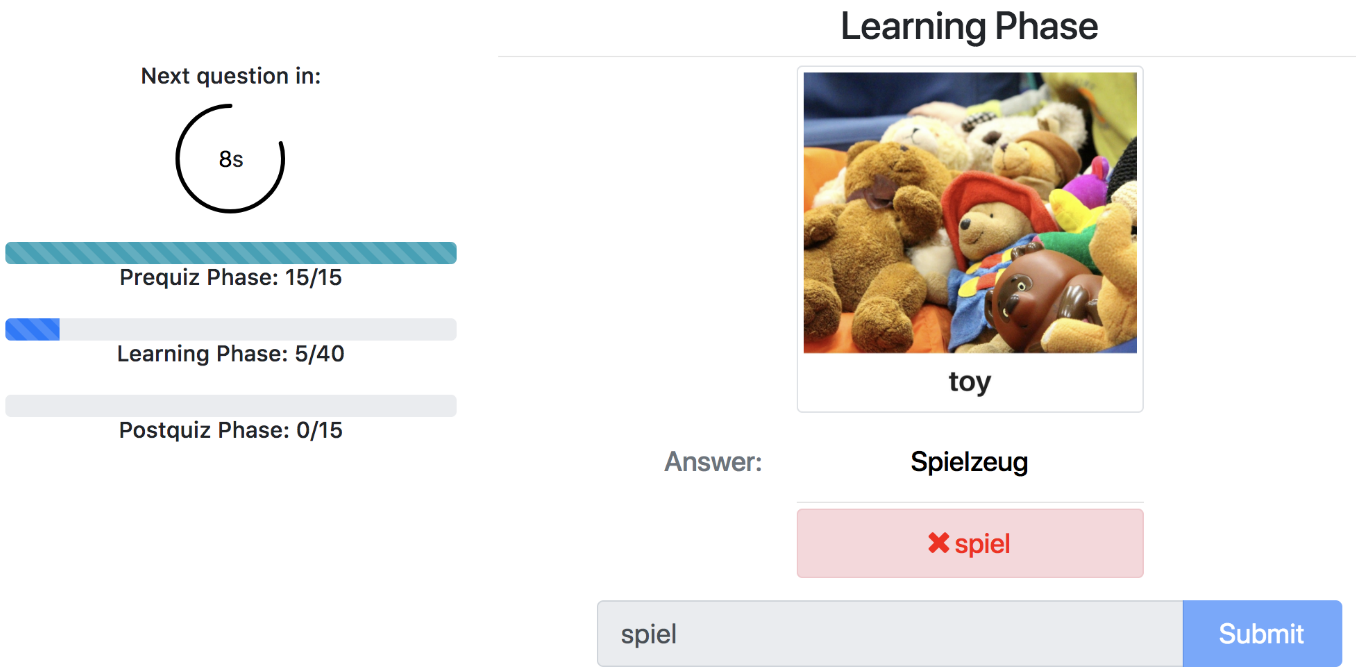

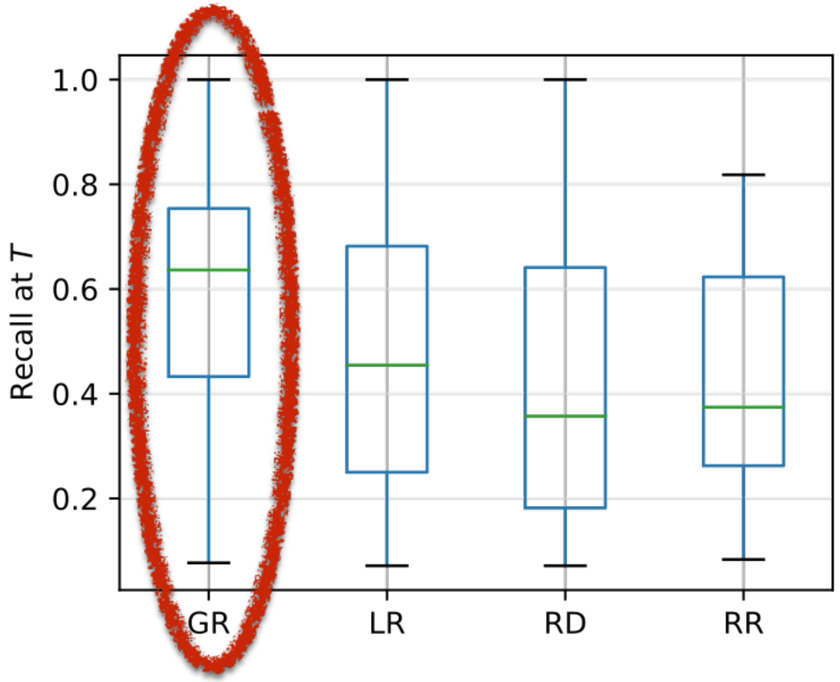

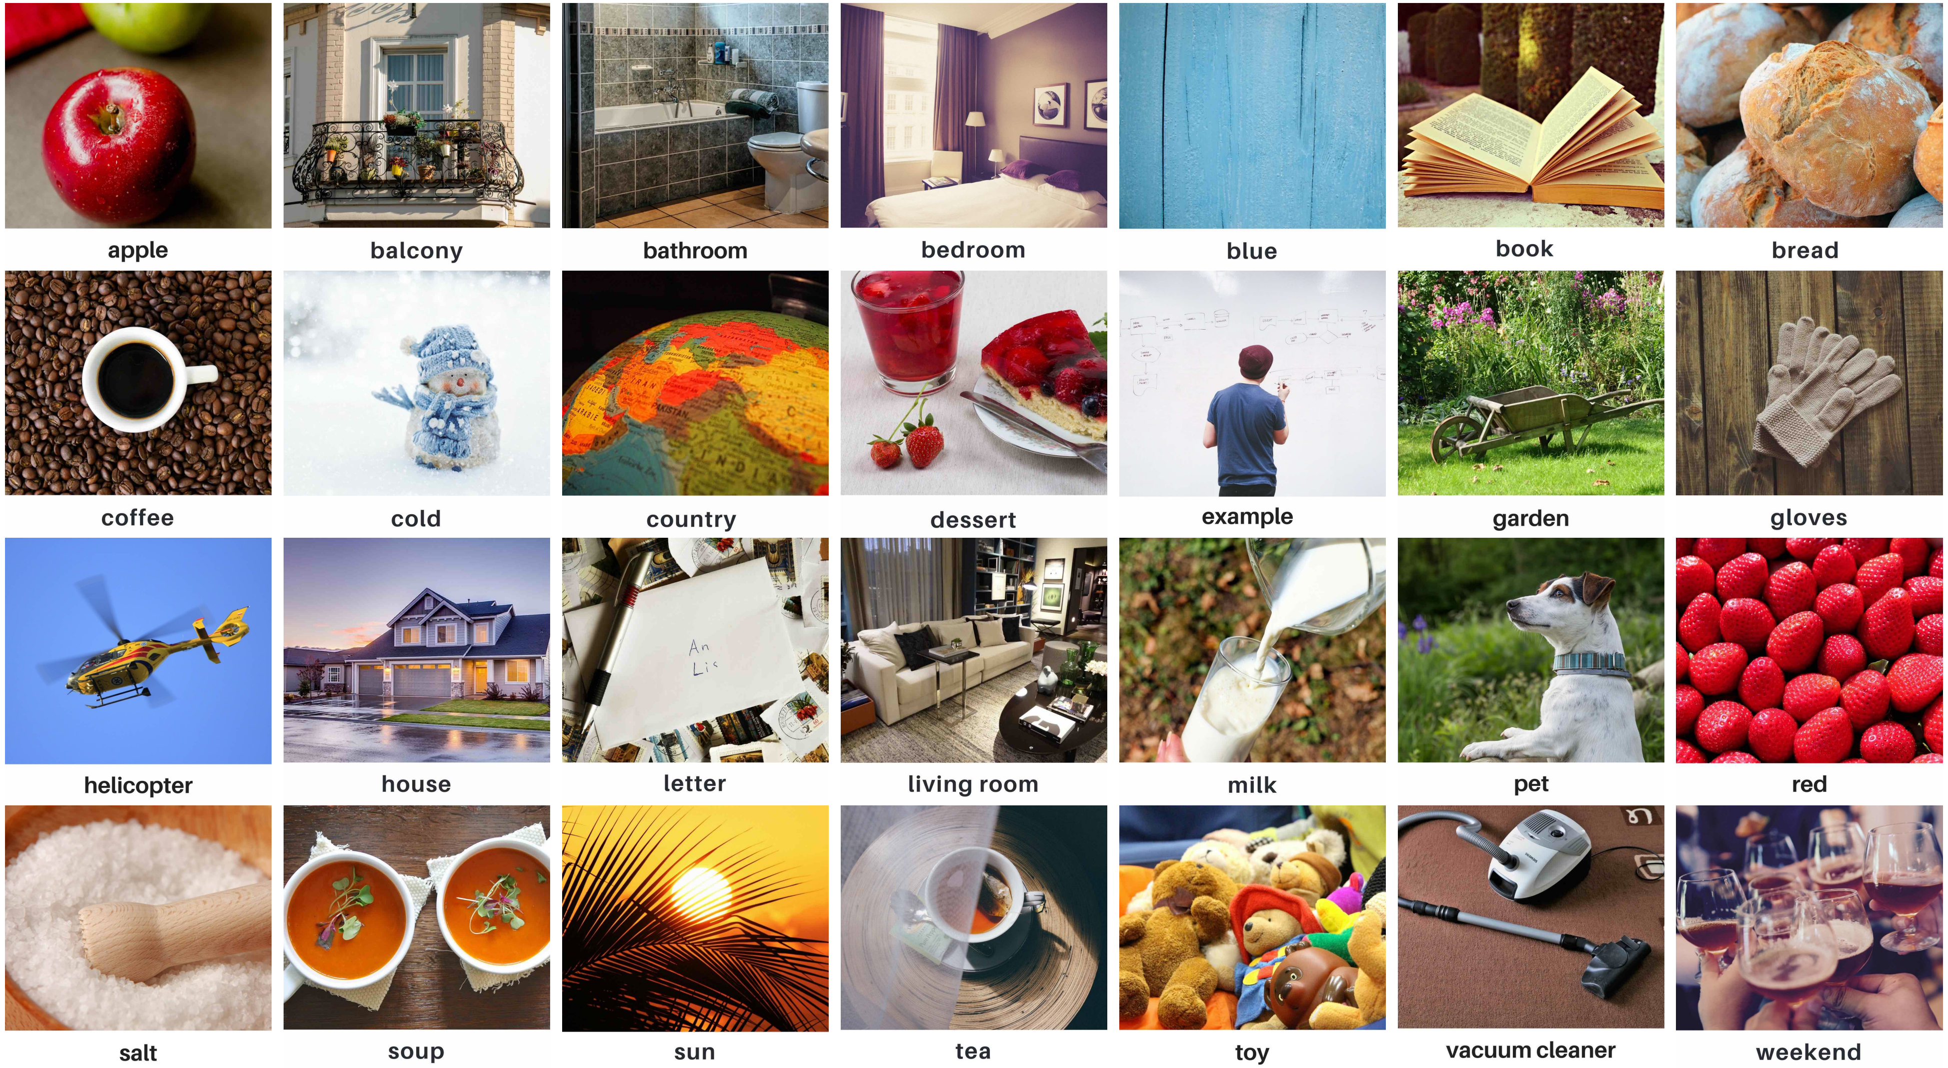

German

	GR	LR	RR	RD
Avg. gain	**0.572**	0.487	0.462	0.467
p-value	-	0.0652	0.0197	0.0151

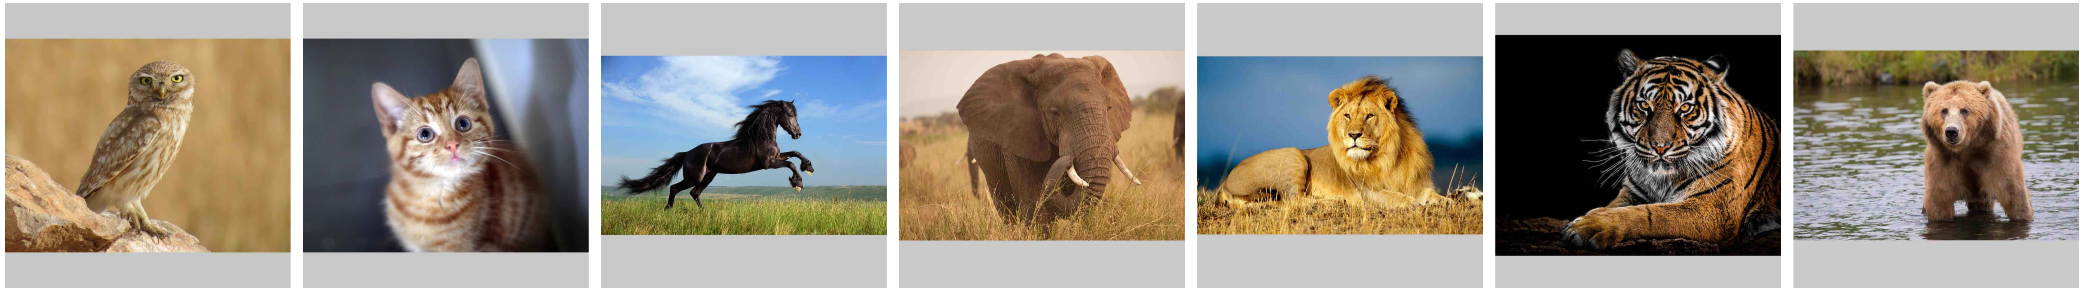

(a) Common: Owl, Cat, Horse, Elephant, Lion, Tiger, Bear

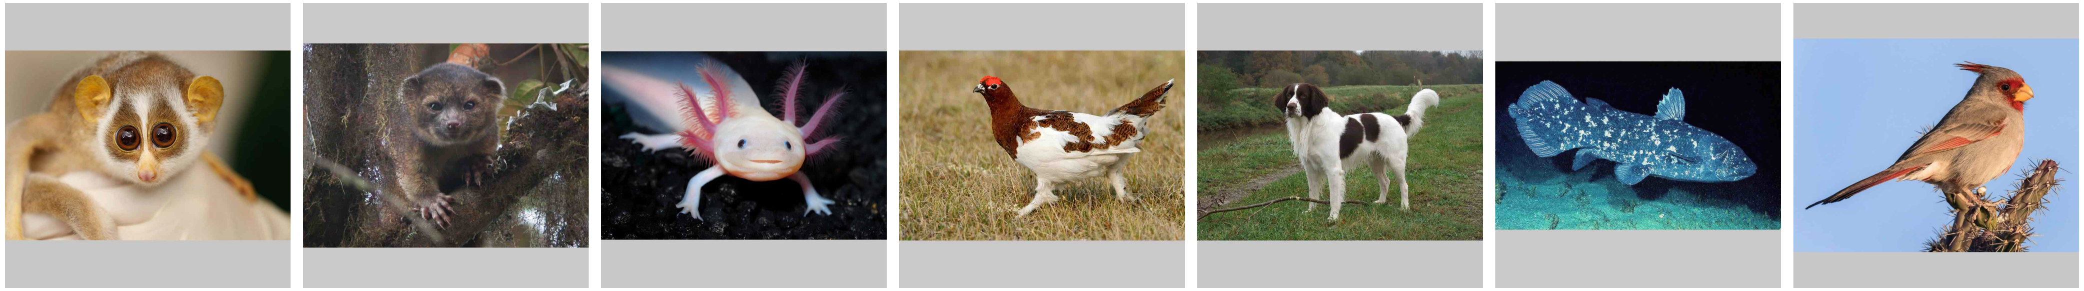

(b) Rare: Angwantibo, Olinguito, Axolotl, Ptarmigan, Patrijshond, Coelacanth, Pyrrhuloxia

Biodiversity (all species)

	GR	LR	RR	RD
Avg. gain	**0.475**	0.411	0.390	0.251
p-value	-	0.0017	0.0001	0.0001

Biodiversity (rare species)

	GR	LR	RR	RD
Avg. gain	**0.766**	0.668	0.601	0.396
p-value	-	0.0001	0.0001	0.0001

Summary: Teaching Concepts to People

- Teaching forgetful learners
  - Limited memory (modeling forgetfulness)
  - Engagement (interface design)

- Challenges not covered in this talk:
  - Limited inference power and noise
  - Mismatch in representation
  - Interpretability (e.g., teaching via labels vs. rich feedback)
  - Safety (e.g., when teaching physical tasks)
  - Fairness (e.g., when teaching a class)
  - . . .

Questions?