# Teaching Multiple Concepts to a Forgetful Learner

University of Zurich[UZH]

Anette Hunziker, Yuxin Chen, Oisin Mac Aodha, Manuel Gomez Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, Adish Singla

ETH zürich

MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS

## 1 Motivating Applications

duolingo memrise

- Language learning apps used by over 300+ million students
- Based on spaced repetition technique
  - Spacing effect: practice should spread out over time
  - Lag effect: spacing between practices should gradually increase
- No known guarantees on scheduling multiple concepts over fixed horizon
- Key research problem that we tackle in this paper is:

Can we compute near-optimal personalized schedule of repetition?

## 2 Teaching Interaction using Flashcards

**Interaction at time $t = 1, 2, \dots T$**

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

## 3 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

**Pimsleur method (1967)**

- Used in mainstream language learning platforms
- Based on spaced repetition ideas

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

Key limitation: Non-adaptive schedule ignores learner's responses

**Leitner system (1972)**

- Adaptive spacing intervals

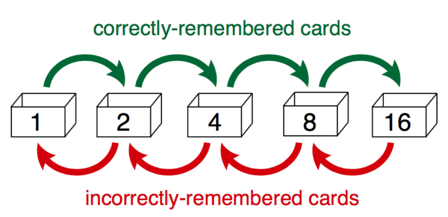


Student 1: a – b – a – b – c – a – c – b – d – c – d – a – b – d – c – e –

Student 2: a – a – b – a – b – c – a – c – a – b – c – a – b – a – d – c –

Key limitation: No guarantees on the optimality of the schedule

## 4 Learner: Memory Model and Responses

- Half-life regression (HLR) model [Settles, Meeder'16]
- Denote history up to time $t$ as $(x_{1:t}, y_{1:t})$
  - Last time step when concept $x$ was taught is $l_t^x \in \{1,..,t\}$
  - $\Delta_{t,\tau}^x = (\tau - l_t^x)$ is time past for $\tau \in \{t+1,..,T\}$
  - Learner's mastery for concept $x$ at time $t$ is $h_t^x$
- Recall probability based on exponential forgetting: $g^x(\tau, (x_{1:t}, y_{1:t})) = 2^{-\left(\Delta_{t,\tau}^x / h_t^x\right)}$
- Changes in half-life $h^x$ parameterized by $(a^x, b^x)$

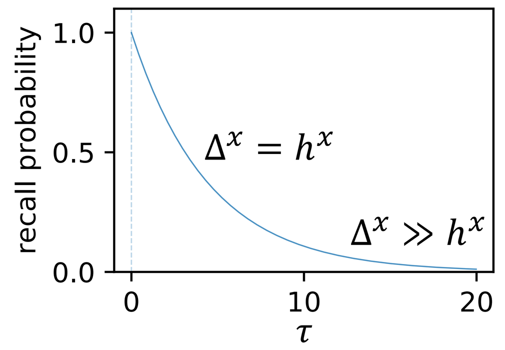


✓ $h^x$ += $a^x$

✗ $h^x$ += $b^x$

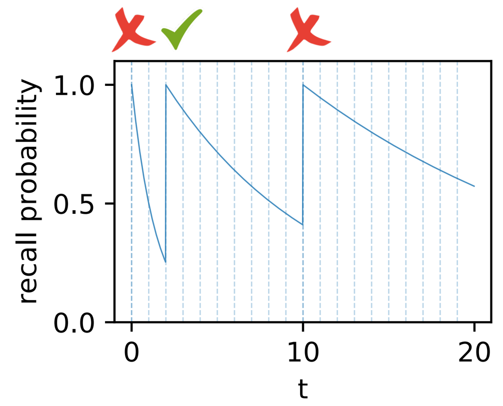


## 5 Teacher: Scheduling as Optimization

**Teacher's objective function**

- Given a sequence of concepts and observations $x_{1:T}, y_{1:T}$, we define

$$f(x_{1:T}, y_{1:T}) = \frac{1}{nT} \underbrace{\sum_{x=1}^{n} \sum_{t=1}^{T} g^x(t+1, (x_{1:t}, y_{1:t}))}_{\text{Area under the curve}}$$

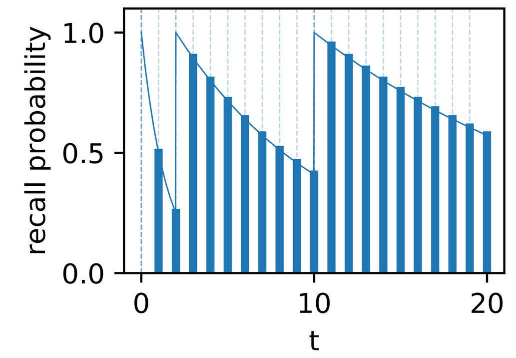


**Optimization problem**

- Teaching policy is given by $\pi: (x_{1:t-1}, y_{1:t-1}) \to \{1,2,..,n\}$
- Average utility of a policy $\pi$ is $F(\pi) = \mathbb{E}_{(x,y)}[f(x_{1:T}^\pi, y_{1:T}^\pi)]$
- Optimal policy is given by $\pi^* = \mathrm{argmax}_\pi F(\pi)$

**Adaptive greedy algorithm**

- for $t = 1, 2, \dots T$:
  - Select $x_t \leftarrow \mathrm{argmax}_x \mathbb{E}_{(y)}[f(x_{1:t-1} \oplus x, y_{1:t-1} \oplus y)] - f(x_{1:t-1}, y_{1:t-1})$
  - Observe learner's recall $y_t \in \{0, 1\}$
  - Update $x_{1:t} \leftarrow x_{1:t-1} \oplus x_t$; $y_{1:t} \leftarrow y_{1:t-1} \oplus y_t$

**Characteristics of the problem**

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

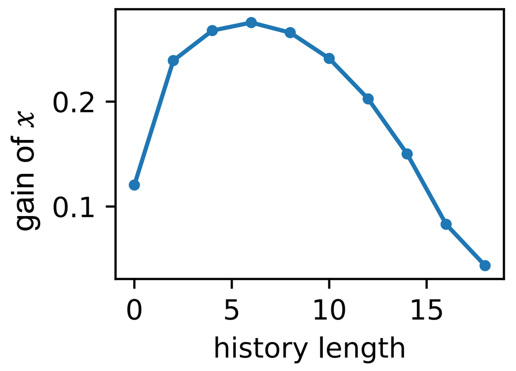


- Post-fix non-monotone
  - $f(\text{orange} \oplus \text{blue}) < f(\text{blue})$
  - Captured by curvature $\omega$

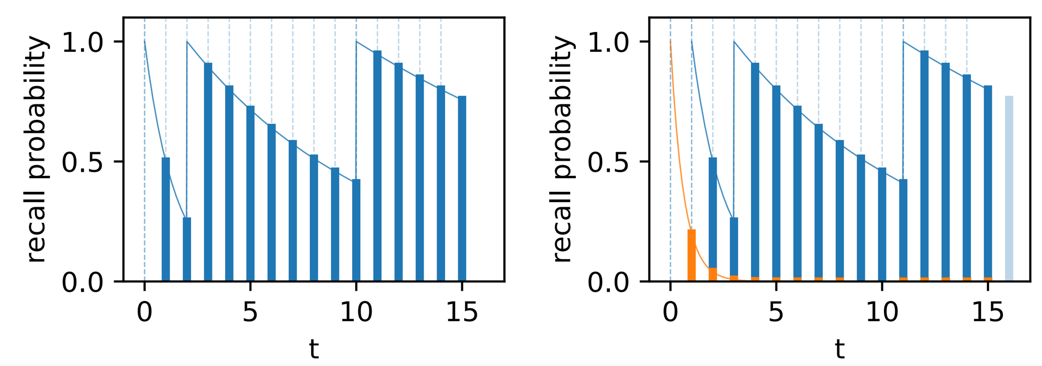


## 6 Theoretical Guarantees

**Guarantees for general case (any memory model)**

- Utility of $\pi^{gr}$ (greedy policy) compared to $\pi^{opt}$ is given by

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

Theorem 1 Corollary 2

**Guarantees for the HLR model**

- Theorem 5. 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 $z \geq \max\{\log T, \log(3n), \log\left(\frac{2n^2}{\epsilon T}\right)\}$.

**Illustration**

- $T$=15 and $n$=3 concepts using HLR model with different parameters

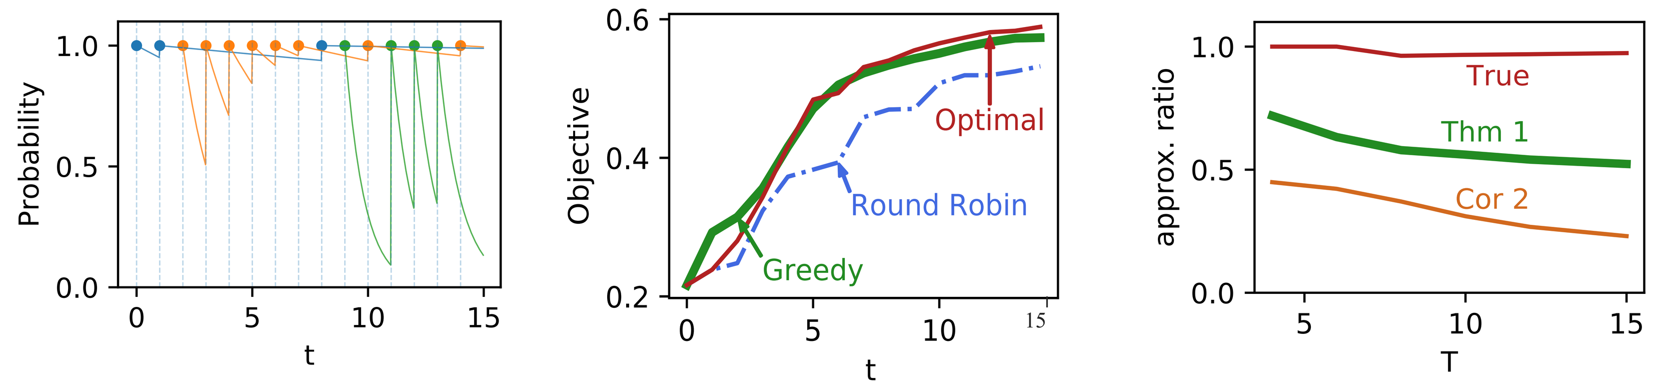


## 7 Results on Human Participants

**Online learning platforms**

- German vocabulary for language learning: https://www.teaching-german.cc/
- Recognizing animal species from images: https://www.teaching-biodiversity.cc/

**Experimental setup**

- Performance measured by gain in knowledge: *postquiz score – prequiz score*
- $T = 40, n = 15$; participants from a crowdsourcing platform (80 and 320)
- Dataset of 100 English-German word pairs
- Dataset of 50 animal images of common and rare species

**Algorithms**

- GR: Our algorithm; RD: Random; RR: Round-robin
- LR: Least-recall (generalization of Pimsleur method and Leitner system)

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

| | | GR | LR | RR | RD |
|---|---|---|---|---|---|
| Biodiversity (all species) | Avg. gain | 0.475 | 0.411 | 0.390 | 0.251 |
| | p-value | - | 0.0017 | 0.0001 | 0.0001 |

| | | GR | LR | RR | RD |
|---|---|---|---|---|---|
| Biodiversity (rare species) | Avg. gain | 0.766 | 0.668 | 0.601 | 0.396 |
| | p-value | - | 0.0001 | 0.0001 | 0.0001 |