Assistive Teaching of Motor Control Tasks to Humans

Megha Srivastava, Erdem Biyik, Suvir Mirchandani, Noah Goodman, Dorsa Sadigh

Motor control tasks are everywhere…

Motor control tasks are everywhere…

and are challenging to learn!

Motor control tasks are everywhere…

and are challenging to **teach others**!

There will always be new motor control tasks to teach


CorPath

What makes teaching motor control tasks challenging?

What makes teaching motor control tasks challenging?

Requires specialized instructors

What makes teaching motor control tasks challenging?

Requires specialized instructors

Individual student variations

What makes teaching motor control tasks challenging?

Requires specialized instructors

Individual student variations

Diverse physical conditions

Can AI-assistance help teach humans motor control tasks?

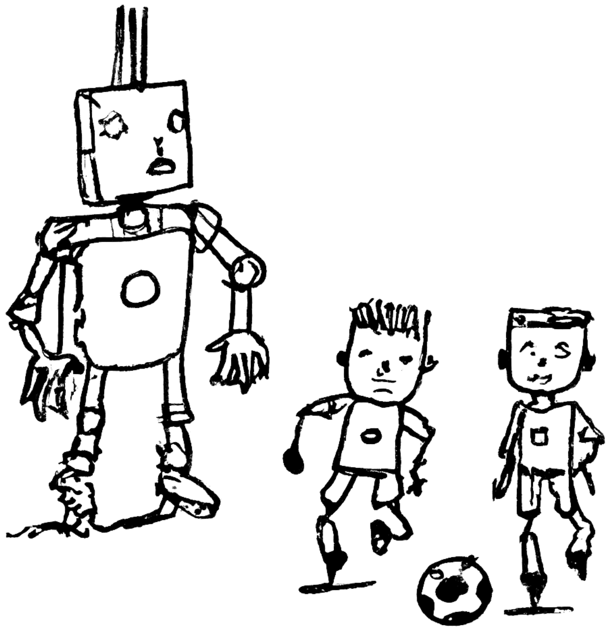

Can we leverage expert knowledge of a motor control task to help any human learn the task themselves?

AI-Assistance has helped bring more accessible, uniform teaching for simpler domains

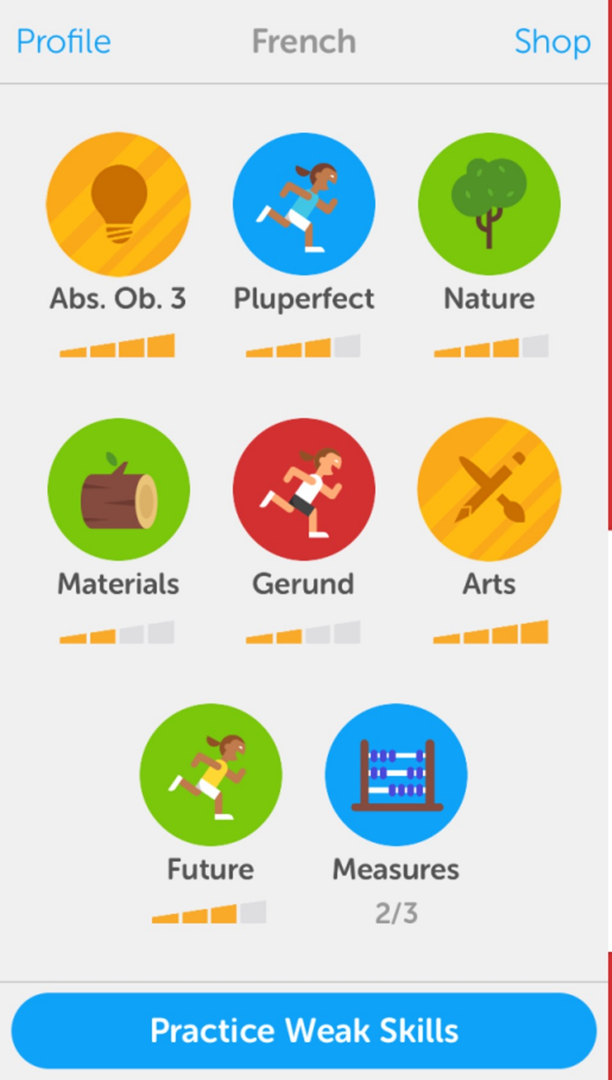

Profile
French
Shop
Abs. Ob. 3
Pluperfect
Nature
Materials
Gerund
Arts
Future
Measures
2/3
Practice Weak Skills

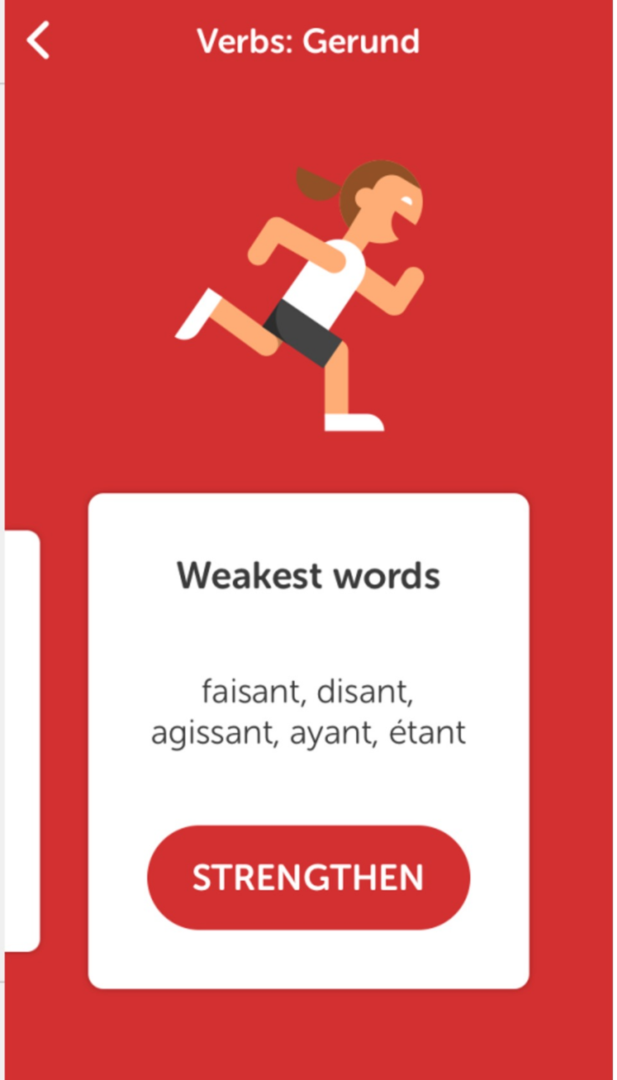

Verbs: Gerund
Weakest words
faisant, disant,
agissant, ayant, étant
STRENGTHEN

[Duolingo]

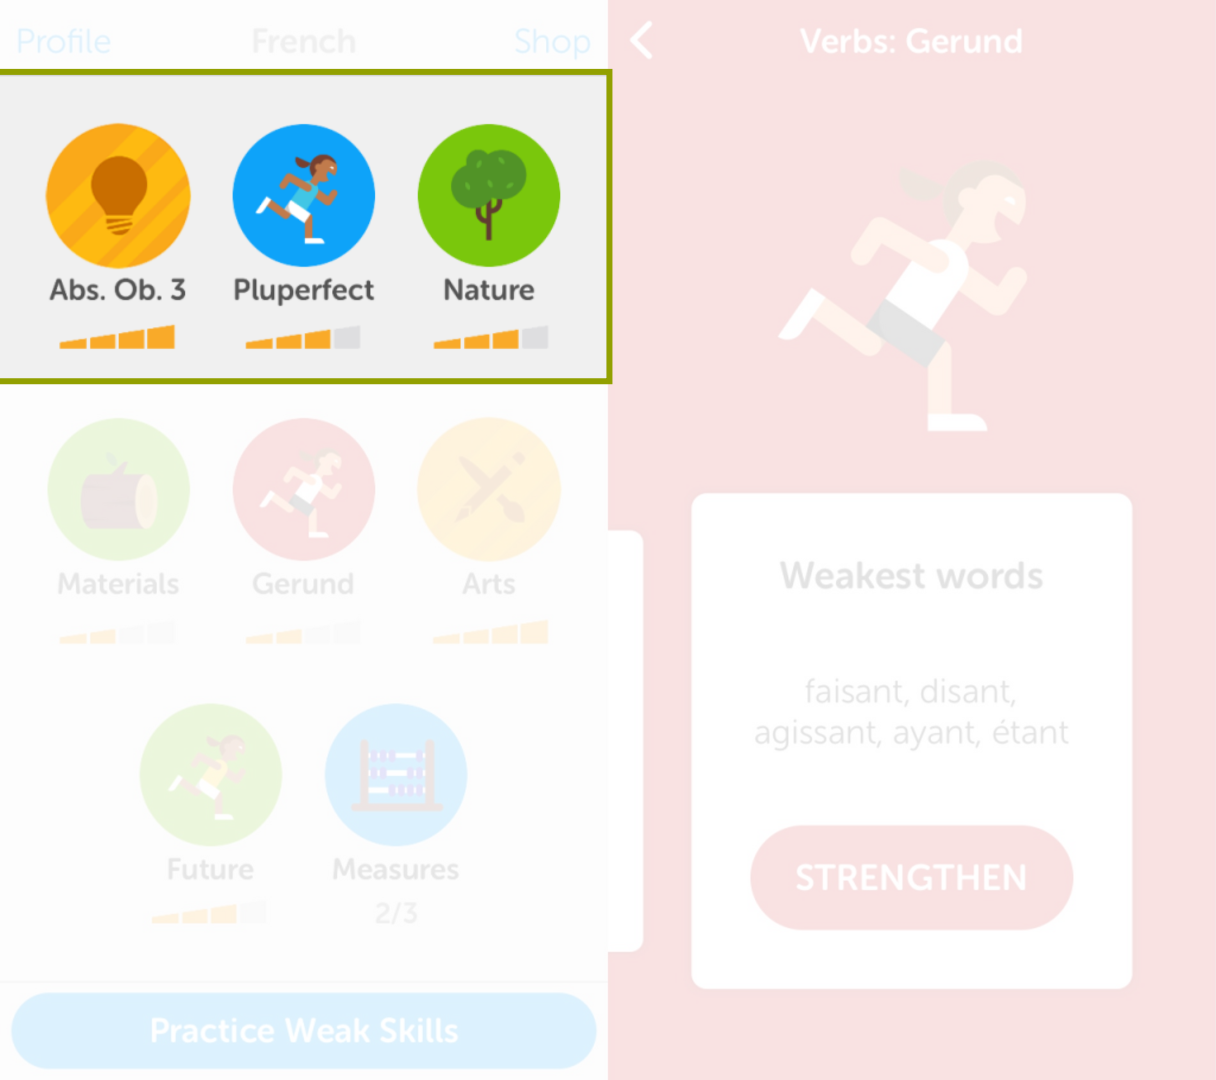

Skill identification

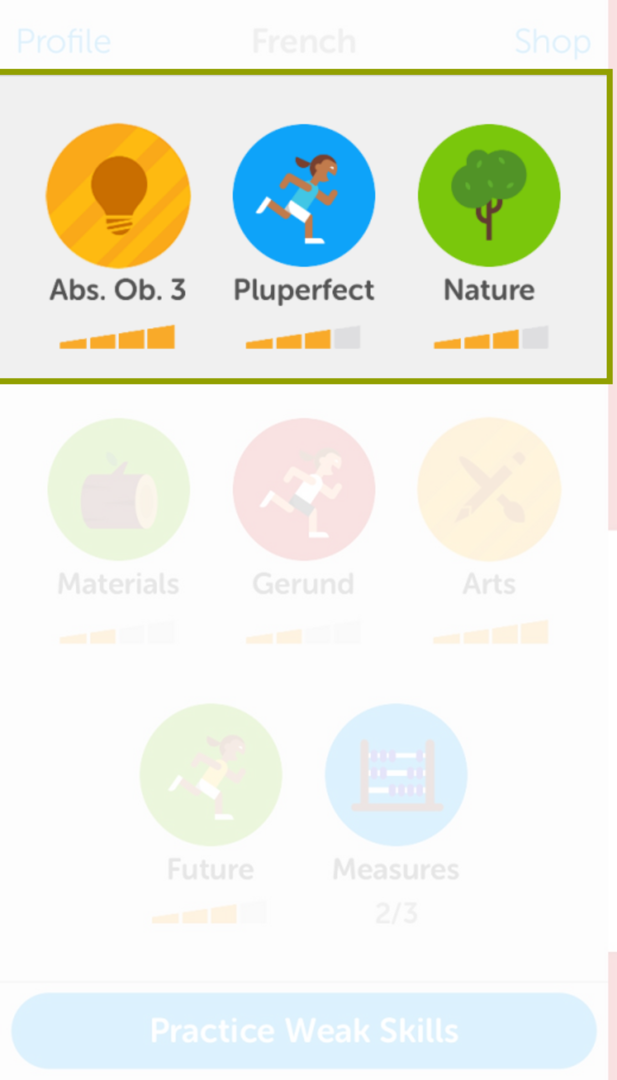

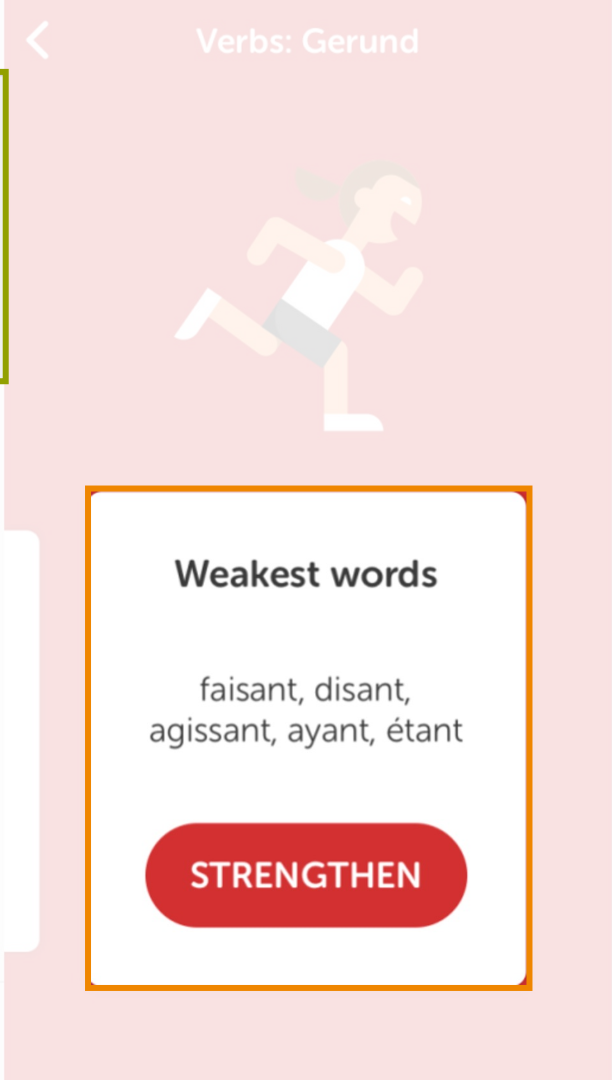

Skill identification

Individualization

[Srivastava & Goodman (2021)]

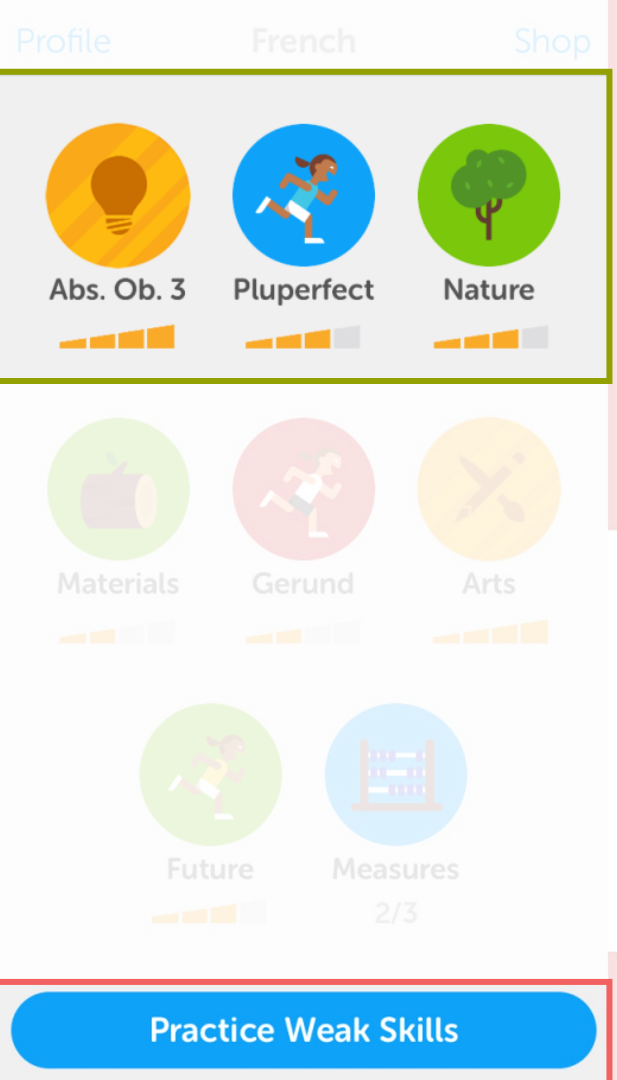

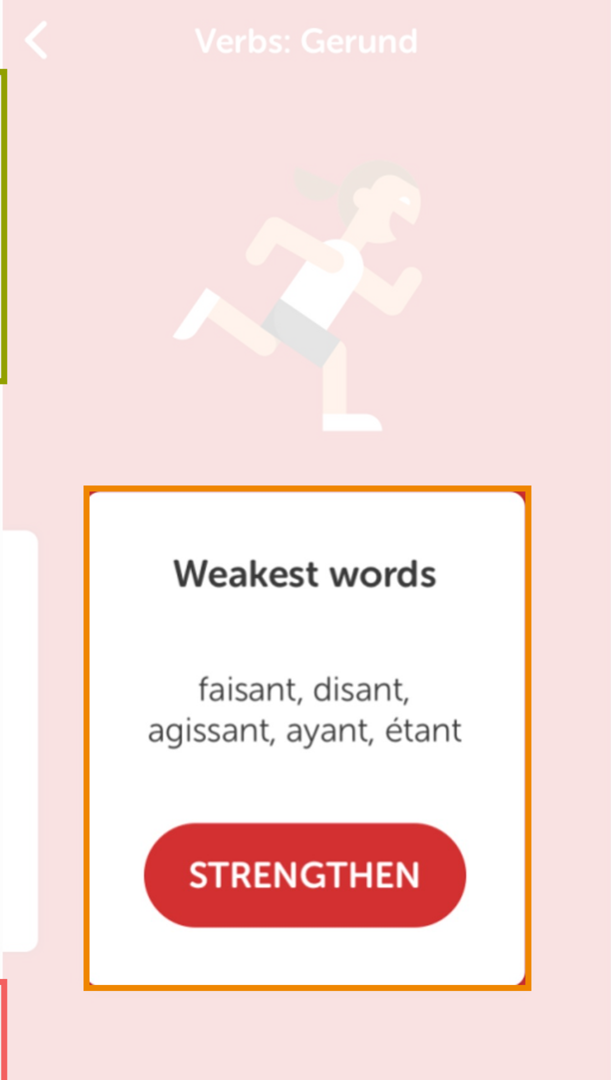

Skill identification

Individualization

Curricula creation (“drills”)

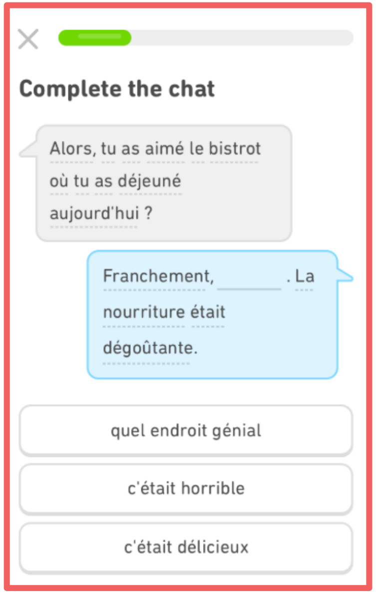

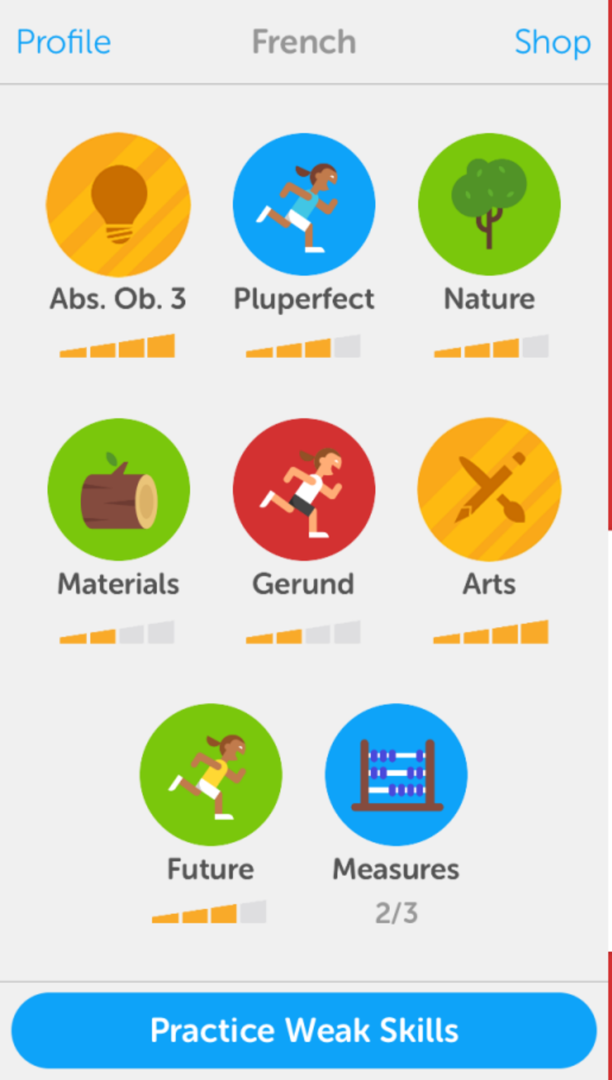

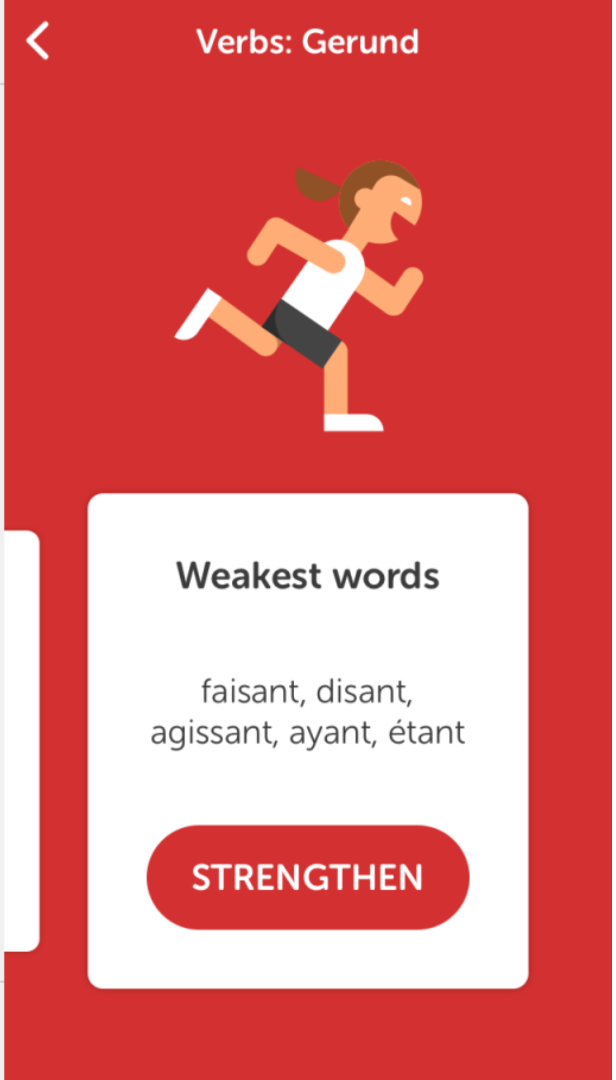

Prior work:
common education domains
(e.g. math, language learning)

Skills in these domains are
standardized & easy to detect!

Key complexity of motor
control tasks:
trajectories over time

Key complexity of motor control tasks: **trajectories over time**

WRITING

Key complexity of motor control tasks: trajectories over time

Initial State s_0 :
(x,y)

Action: (x,y)

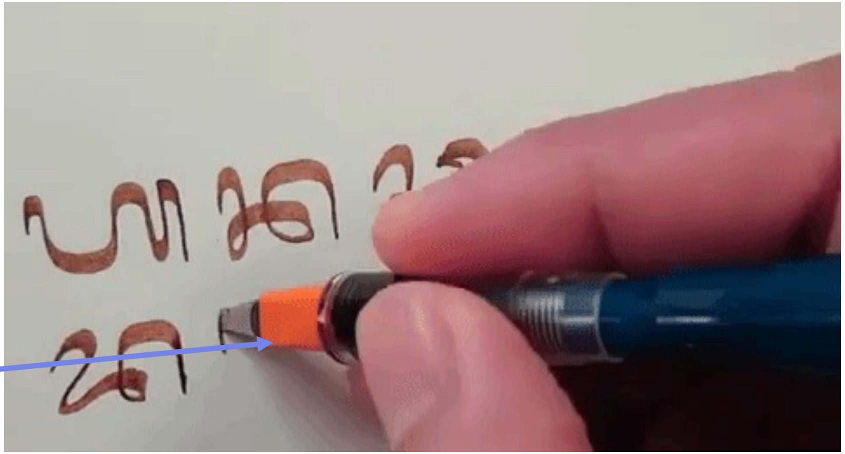

Reward:
Overlap w/ Goal Sequence

WRITING

Scenario: ξ: (s_0, r) initial state and reward pair (e.g. goal character sequence)

Trajectory: τ: $(s_0, a_0)\ldots(s_T, a_T)$ sequence for a particular scenario ξ: (s_0, r)

Key complexity of motor control tasks: trajectories over time

WRITING

Skill identification

Individualization

Curricula creation ("drills")

Key complexity of motor control tasks: trajectories over time

WRITING

Skill identification

How do we identify motor control skills from motion trajectories?

Unsupervised Skill Discovery: CompILE [Kipf et. al. '19]

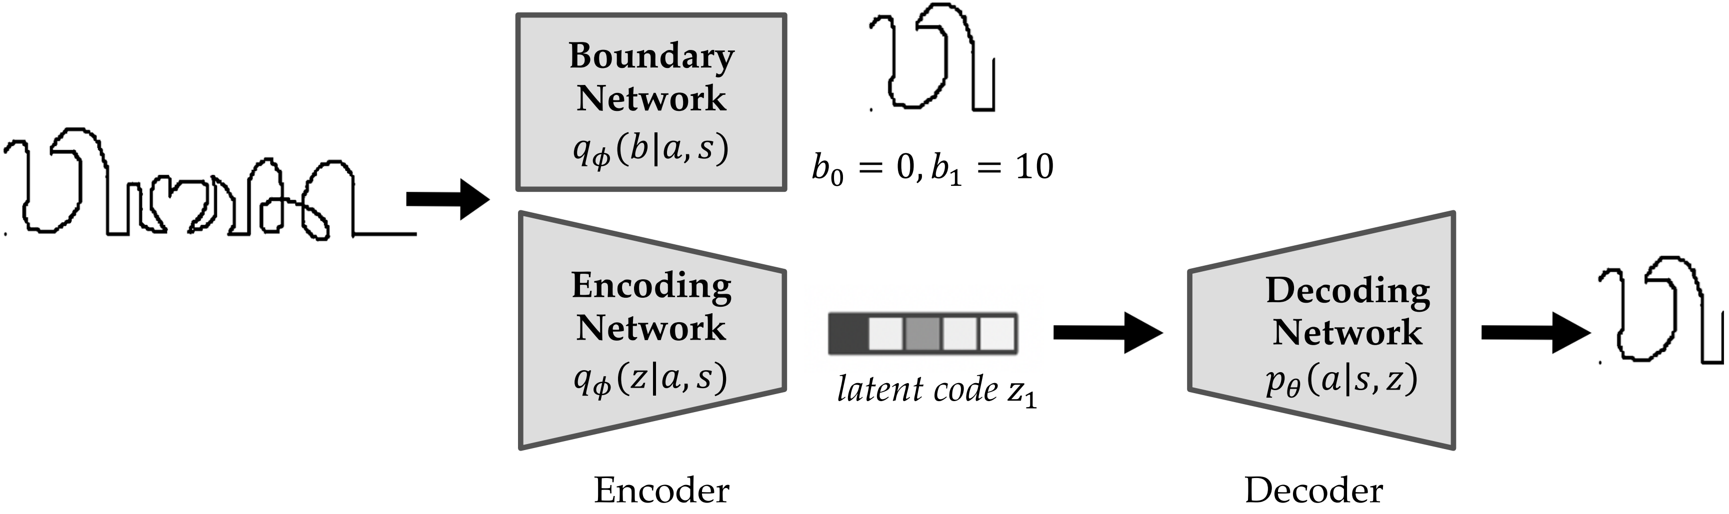

Re-construction Loss Function (across set C of segments in trajectory):

$$-\mathbb{E}_{q_\phi(b,z\,|a,s)} \sum_{i=1:|C|} [P(t \in C_i) * \log p_\theta(a|s, z_i)]$$

Unsupervised Skill Discovery: CompILE [Kipf et. al. '19]

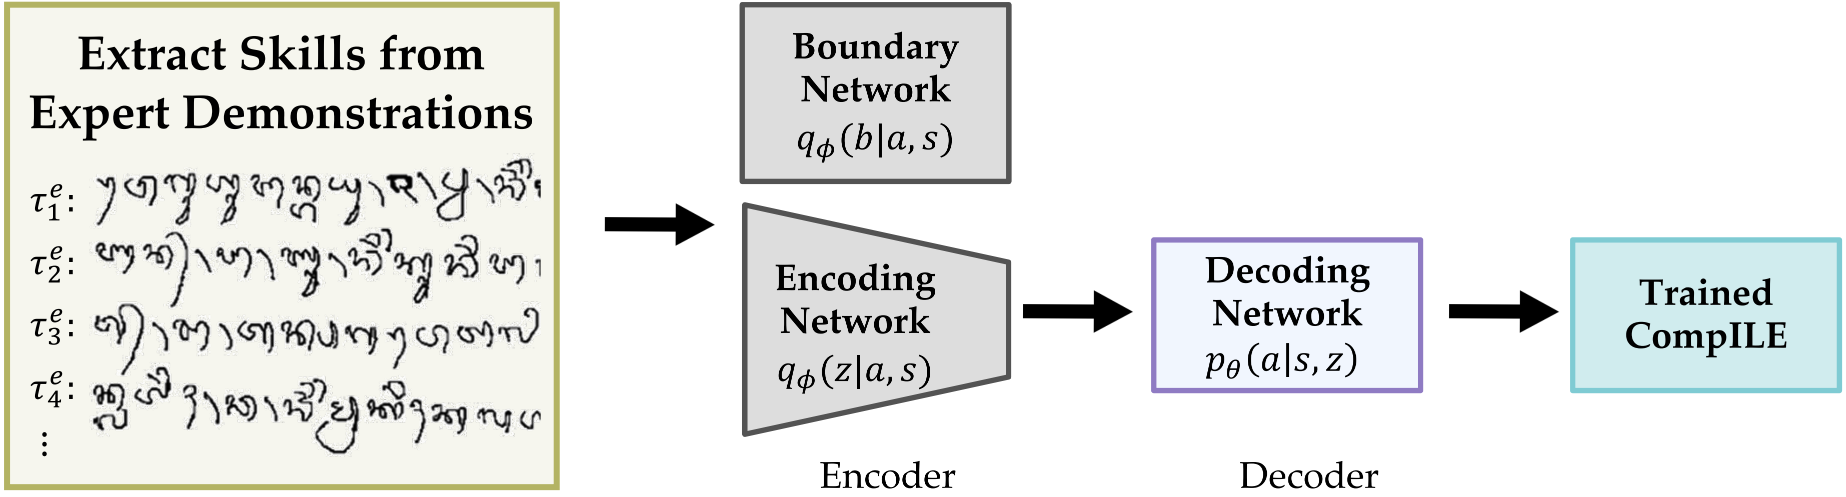

Unsupervised Skill Discovery: CompILE [Kipf et. al. '19]

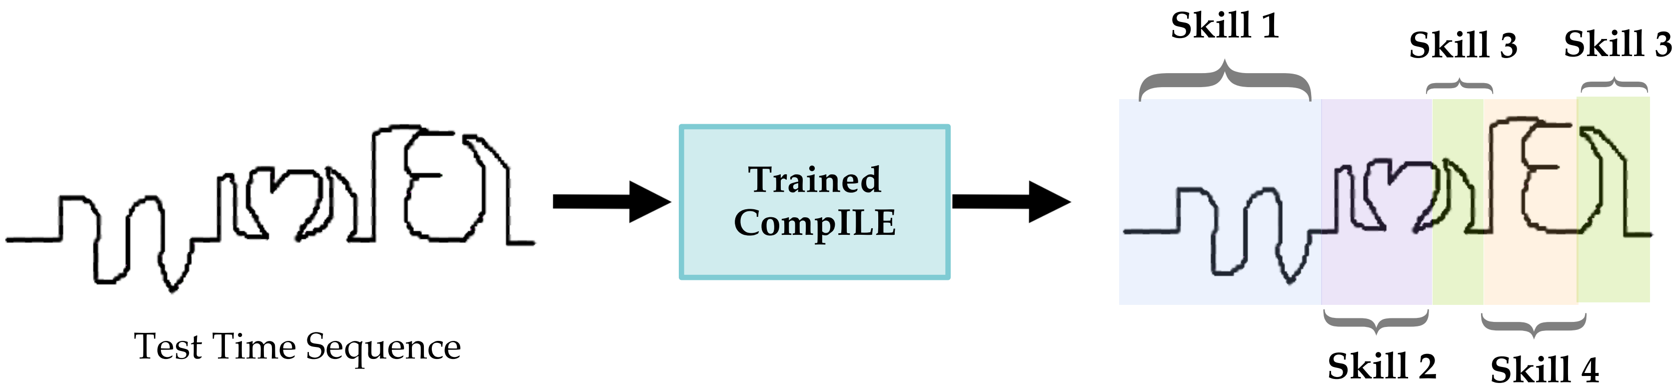

Key idea: Use expert demos + unsupervised skill discovery

Key complexity of motor control tasks: trajectories over time

Writing

Individualization

How do we identify individual expertise from *student* motion trajectories?

Identifying Individual Student Skill Expertise

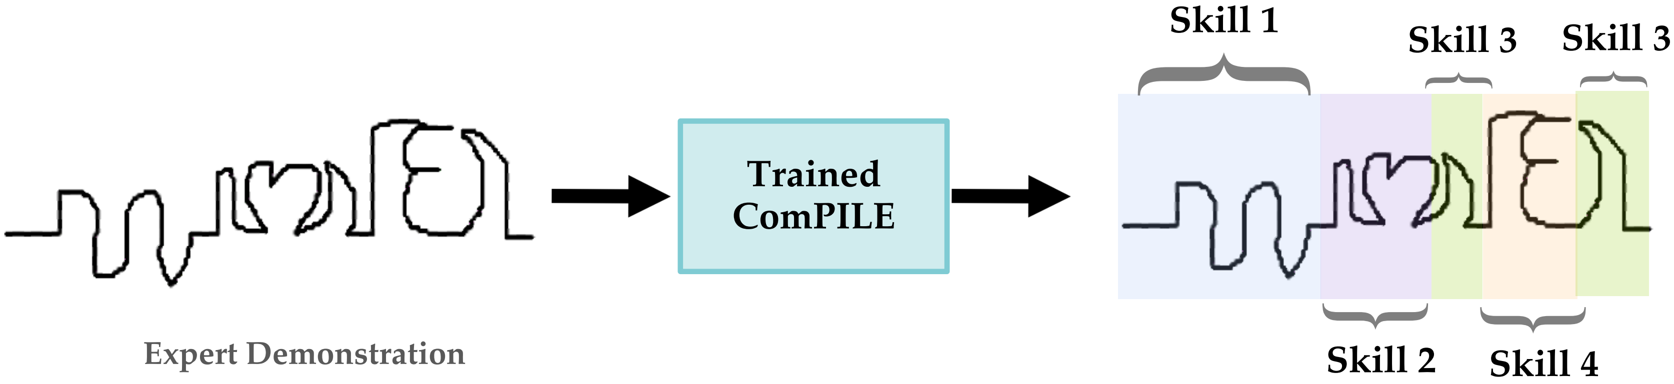

Identifying Individual Student Skill Expertise

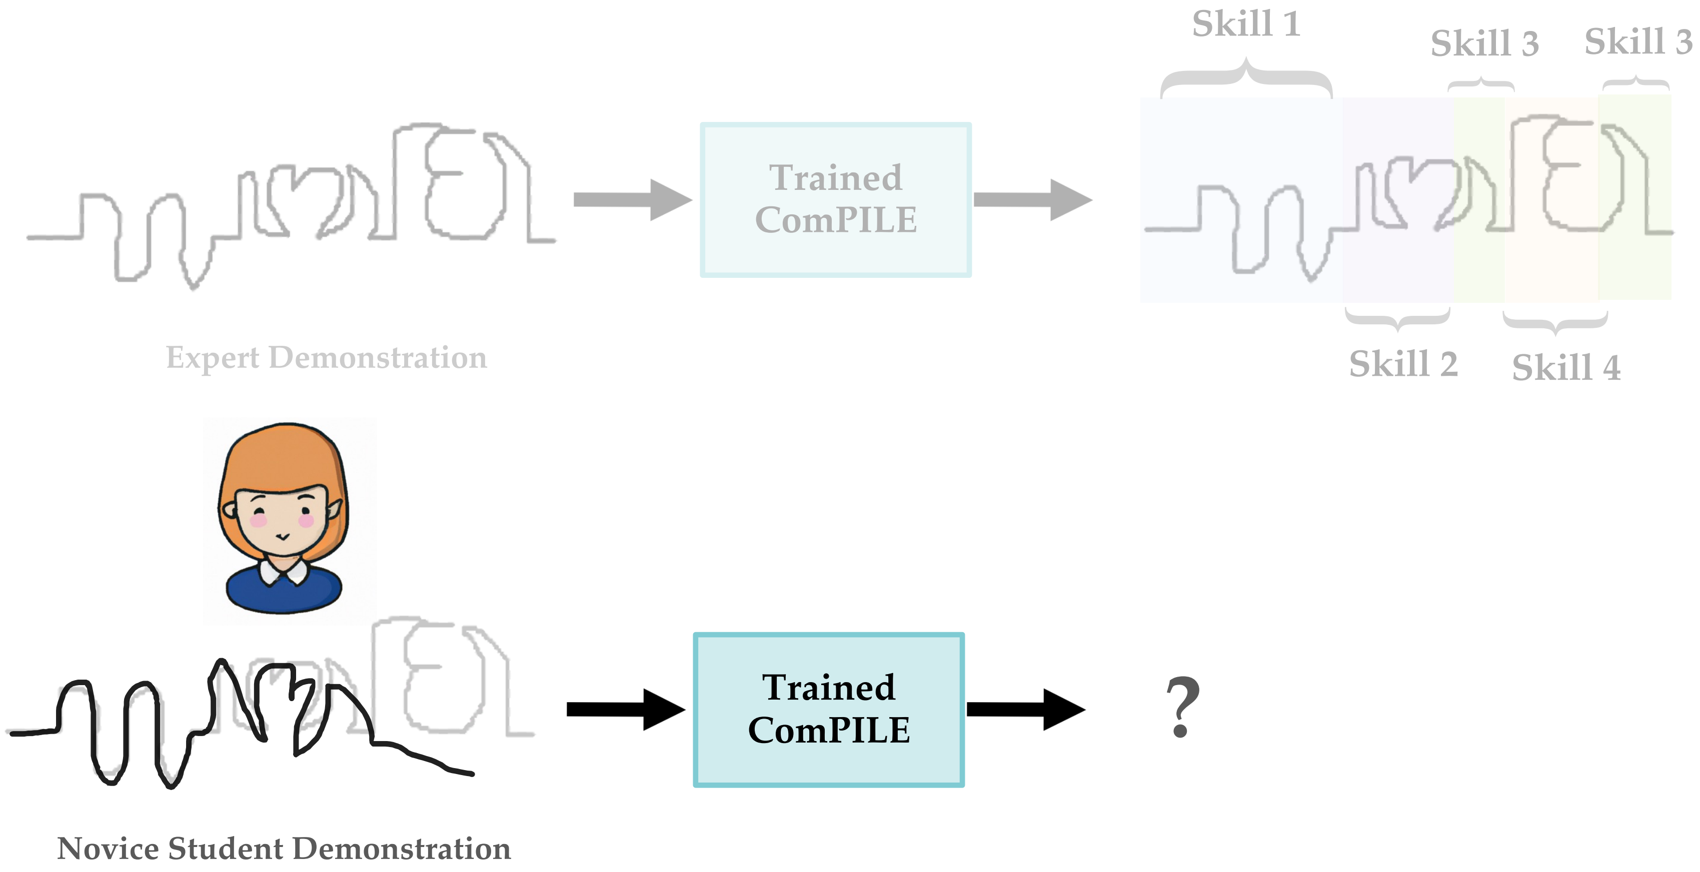

Identifying Individual Student Skill Expertise

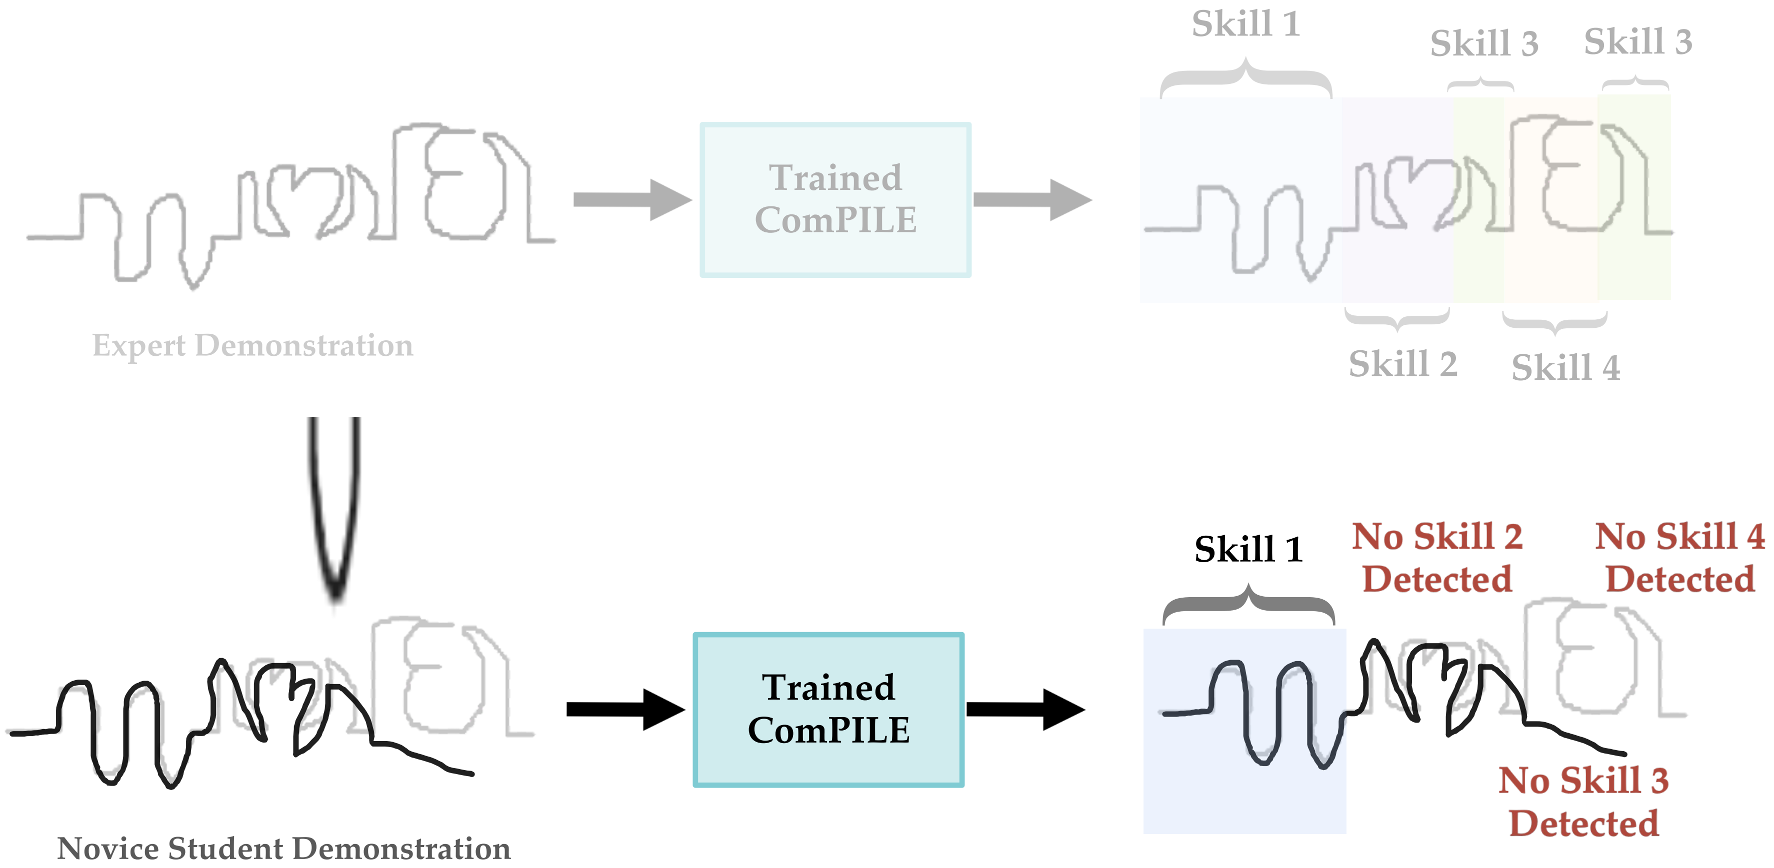

Identifying Individual Student Skill Expertise

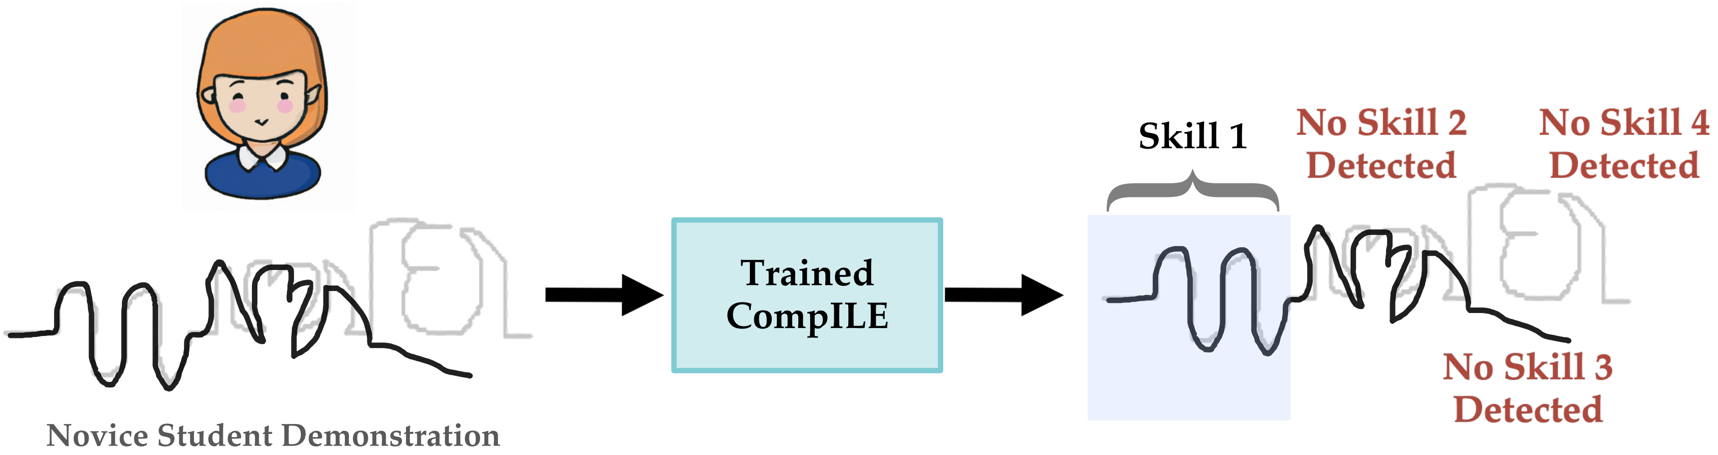

Penalize Skill 2 more than Skill 3 and Skill 4?
temporal decay term j

Identifying Individual Student Skill Expertise

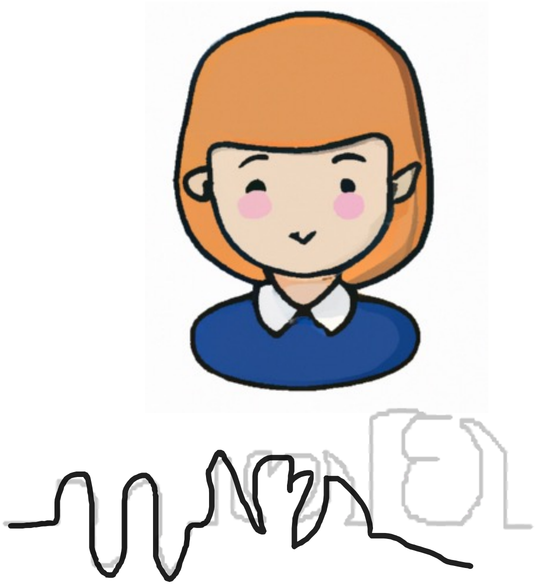

Reward $r =$

$- \alpha *$ (pixel dist. between student and expert)

$+\beta *$ (highest x value of trajectory)

Skill 1: 0
Skill 2: r/j
Skill 3: r/2j
Skill 4: r/3j

Skill Scores

Identifying Individual Student Skill Expertise

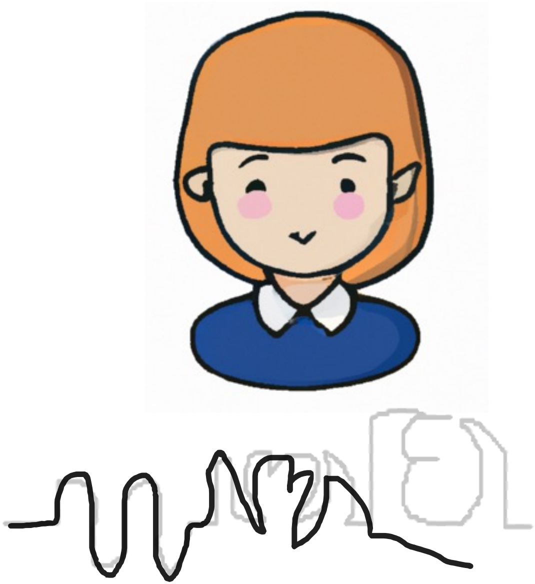

Reward $r = -12$

Penalty $j = 1$

Skill 1: 0
Skill 2: -12
Skill 3: -6
Skill 4: -4

Skill Scores

Identifying Individual Student Skill Expertise

Skill 1: 0
Skill 2: -12
Skill 3: -6
Skill 4: -4

Skill 1: -12
Skill 2: -6
Skill 3: -4
Skill 4: -3

…

Skill 1: -9
Skill 2: -8
Skill 3: -5
Skill 4: -3

Diverse Scenarios

Maximum Set-Coverage
over expert demonstrations / skills

Individual's Most Challenging Skills

Identifying Individual Student Skill Expertise

Skill 1: -1
Skill 2: -4
Skill 3: -1
Skill 4: 0

Skill 1: -1
Skill 2: -10
Skill 3: -5
Skill 4: -7

Skill 1: -9
Skill 2: -8
Skill 3: -5
Skill 4: -3

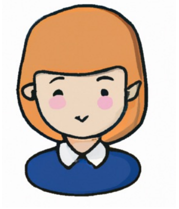

Key complexity of motor control tasks: trajectories over time

WRITING

Curricula creation ("drills")

How do we create novel drills that improve learning from motion trajectories?

Drills: **Repetitive** sequences that targets skills in their **most common contexts**

Drills: Repetitive sequences that targets skills in their most common contexts

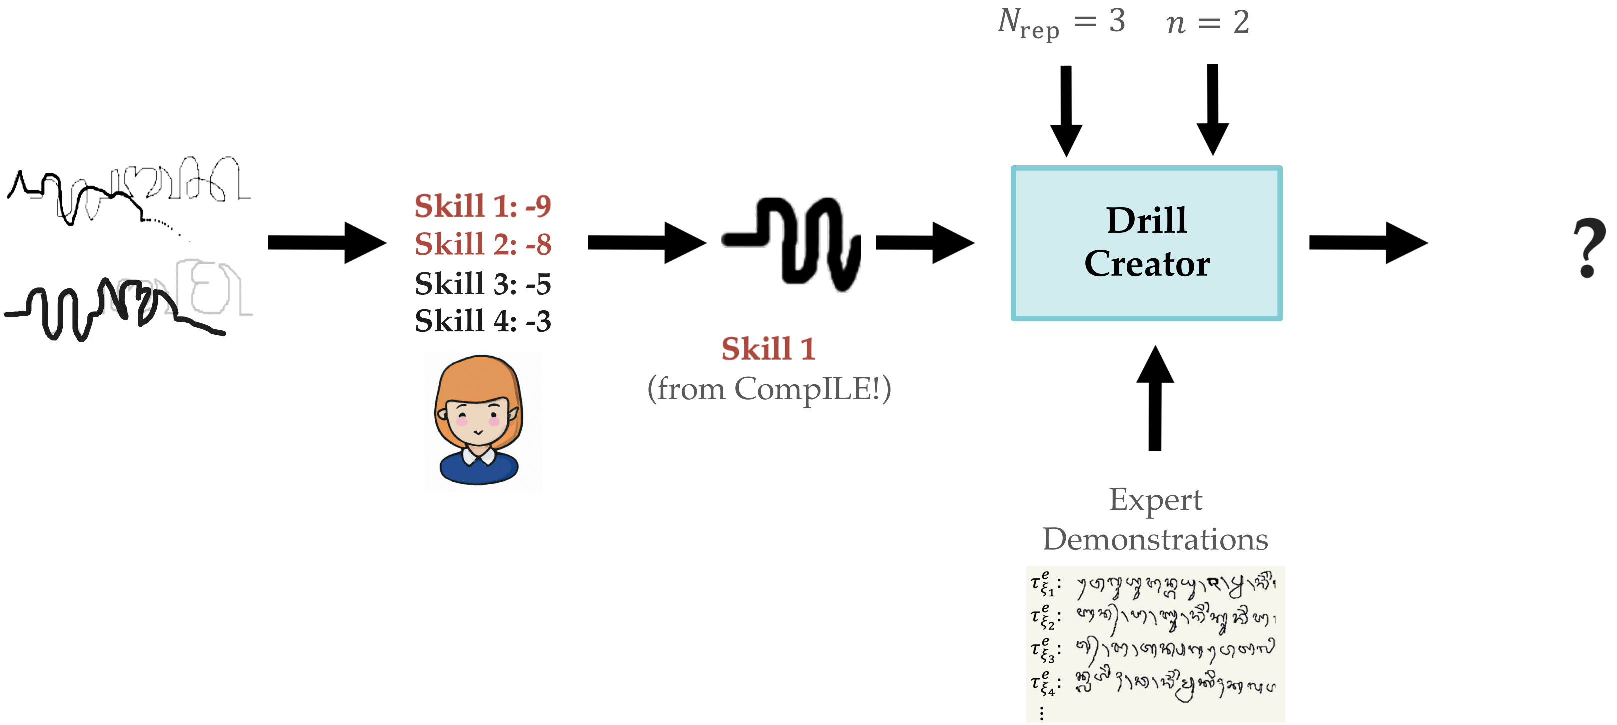

Drills: Repetitive sequences that targets skills in their most common contexts

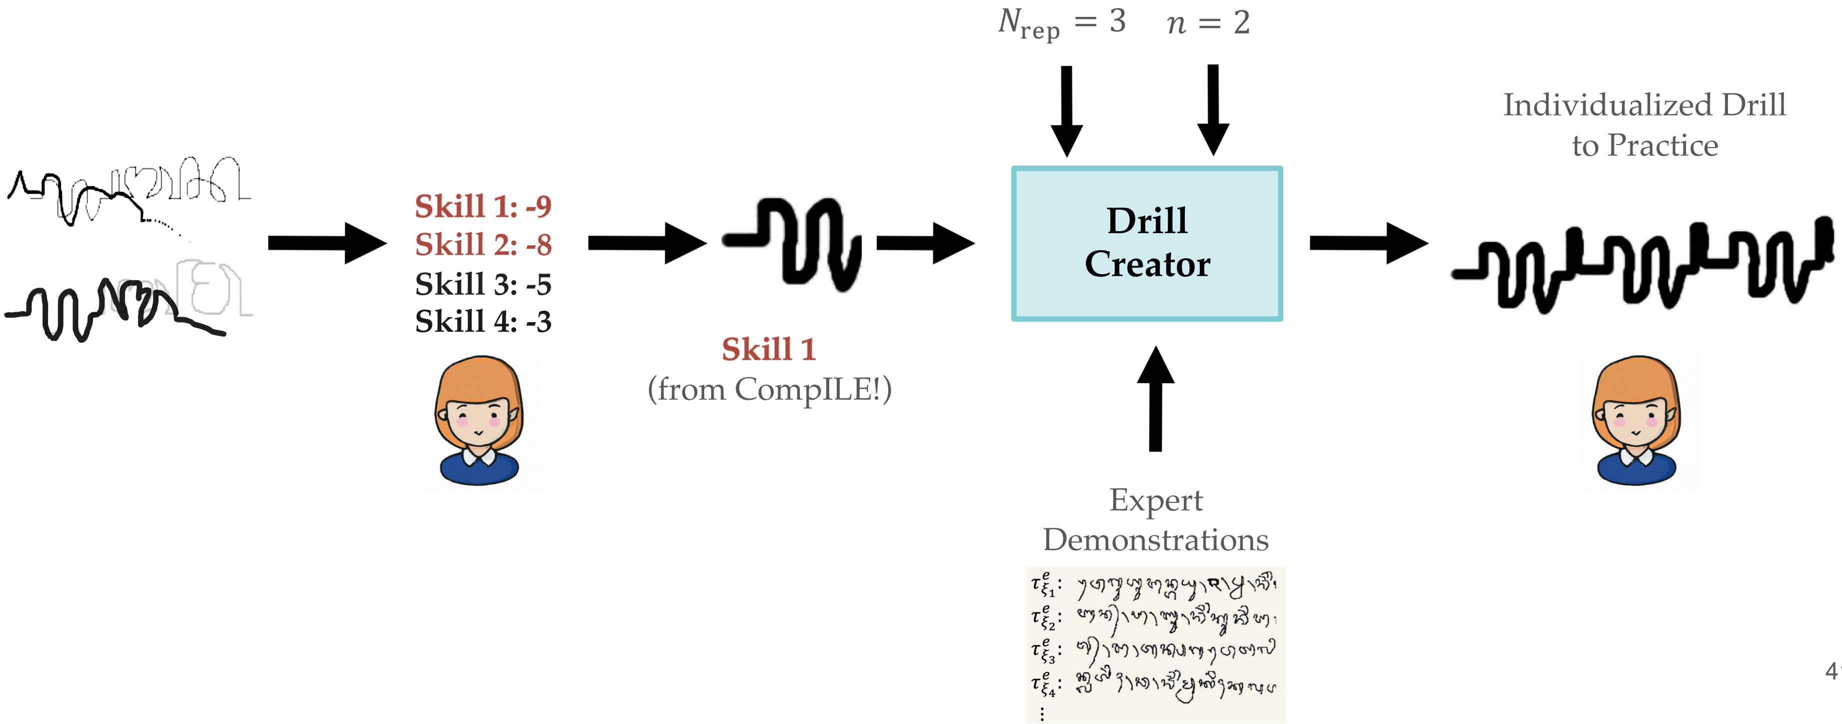

1. Extract Skills from Expert Demonstrations

τ_1^e:
τ_2^e:
τ_3^e:
τ_4^e:
$\vdots$

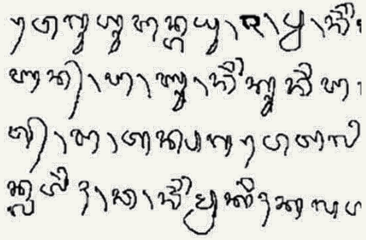

2. Select Scenarios with Diverse Skills

τ_1^e:
τ_2^e:

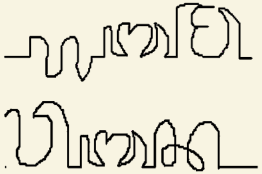

3. Identify Individual Student Skill Expertise

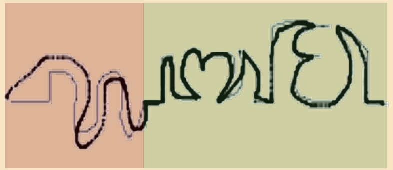

4. Automatically Create Individualized Drills

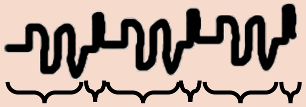

Experiments: Parking & Writing

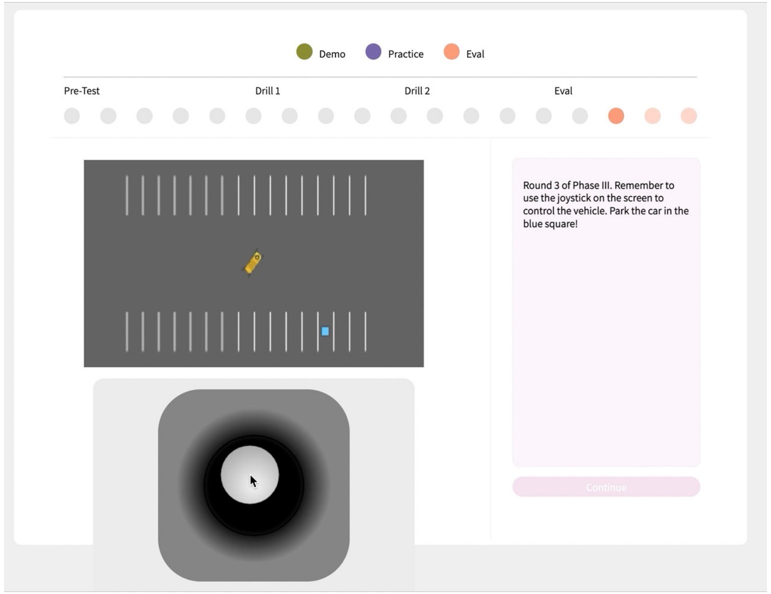

PARKING

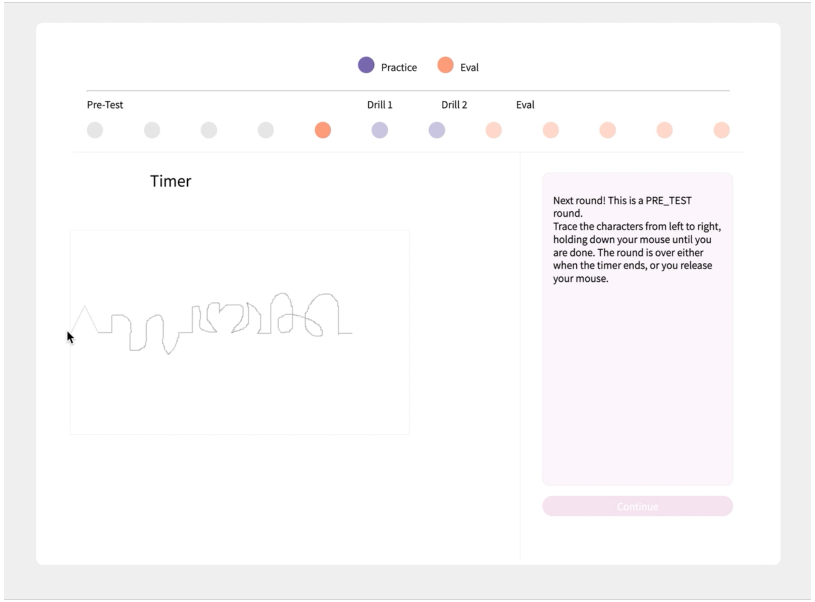

WRITING

Writing Task

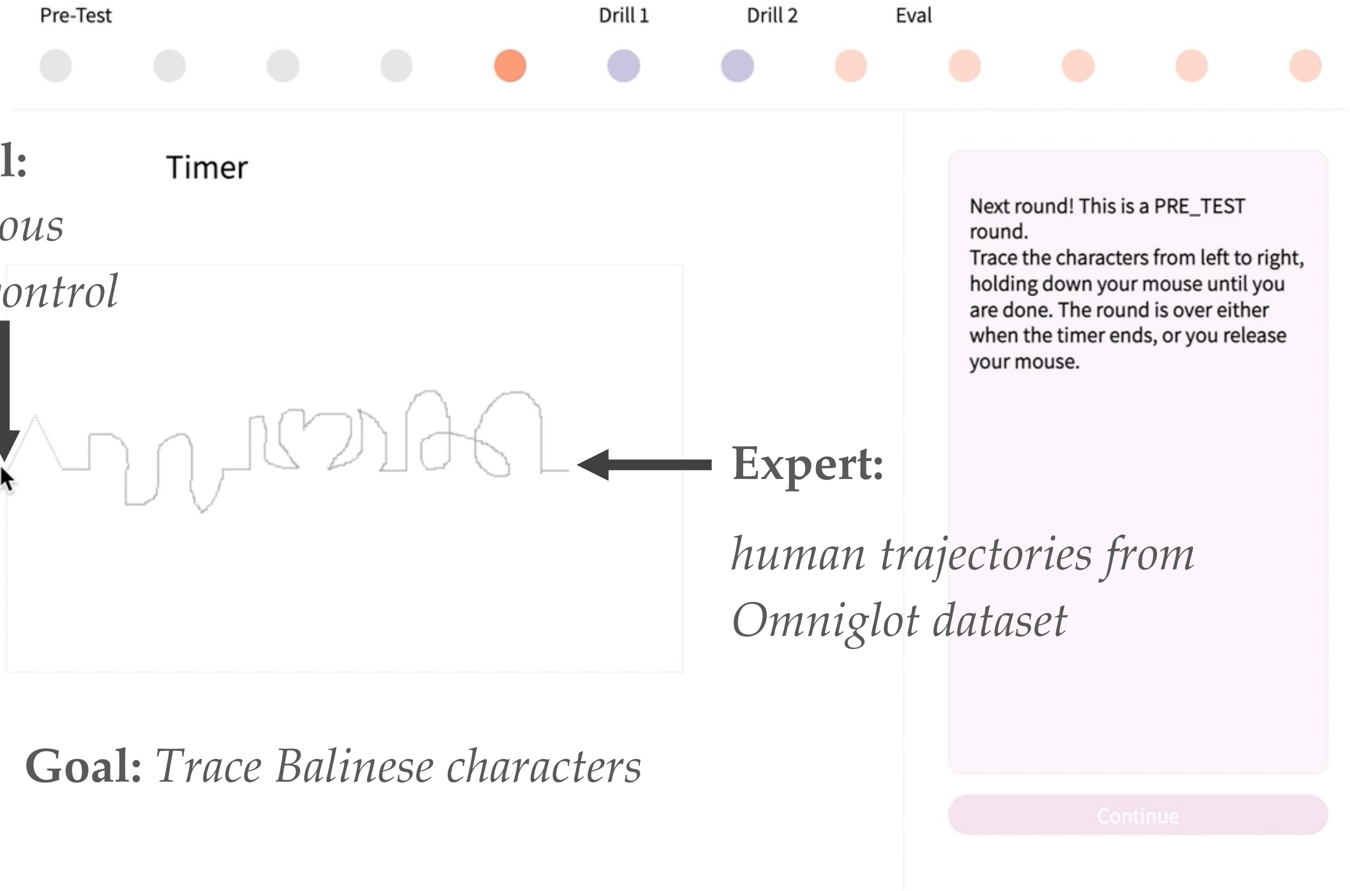

Control: *continuous mouse control*

Expert: *human trajectories from Omniglot dataset*

Goal: *Trace Balinese characters*

Writing Task

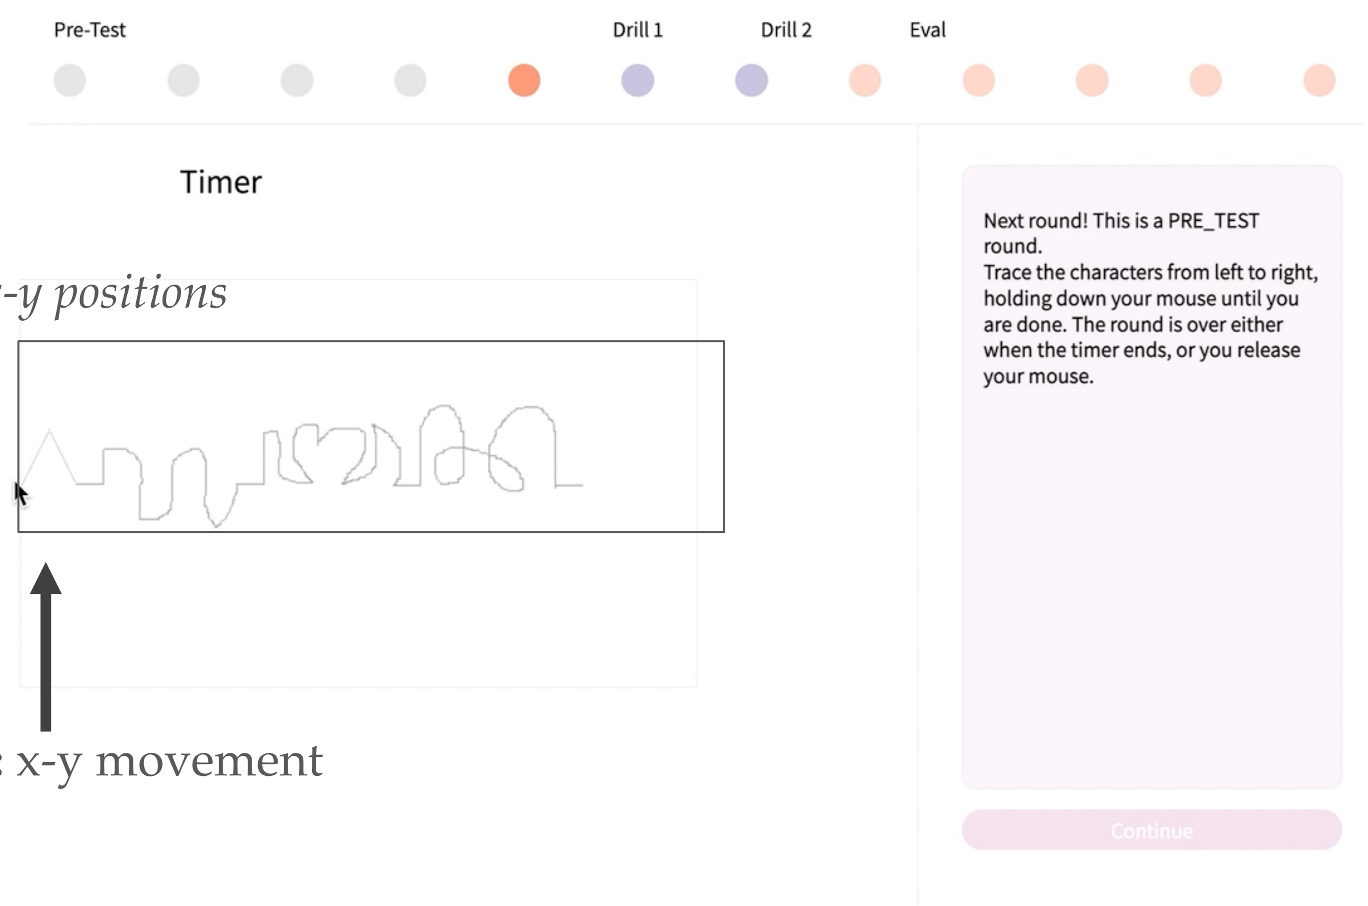

State (2-d): *x-y positions*

Action (2-d): x-y movement

Parking Task

Expert: *Optimal Soft-Actor Critic Agent*

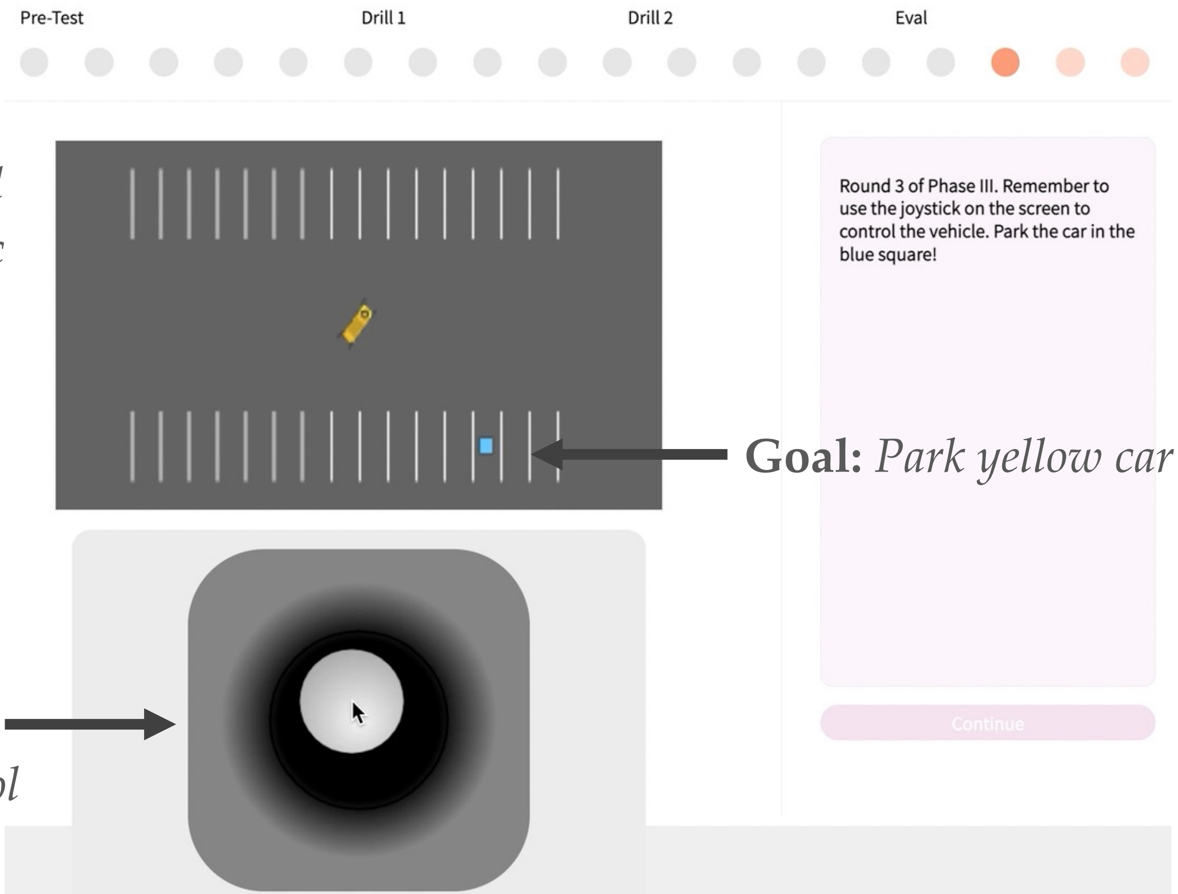

Goal: *Park yellow car on blue spot*

Control: *continuous mouse control*

Parking Task

State (6-d):
position,
velocity,
heading

Action (2-d):
acceleration,
heading

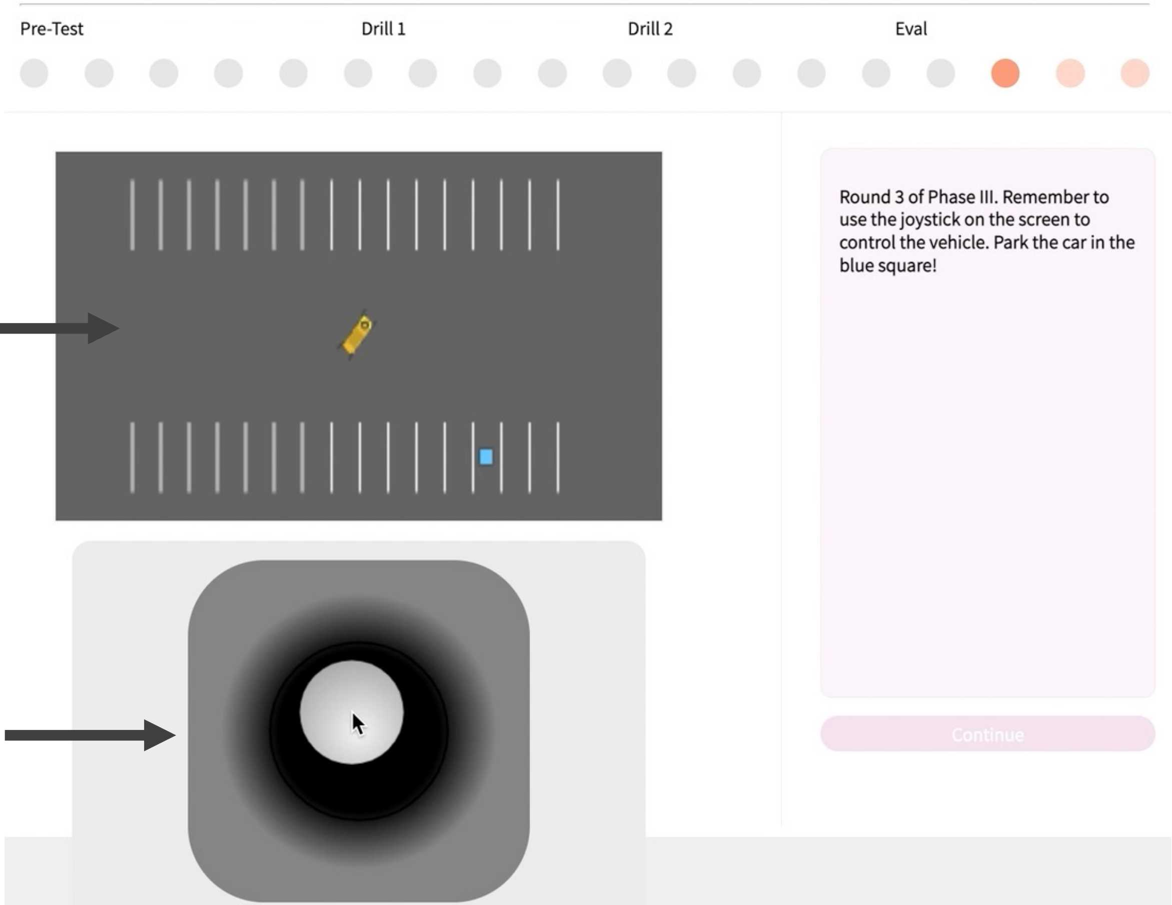

Are skills returned from CompILE useful for learning?

Are skills returned from CompILE useful for learning?

PARKING

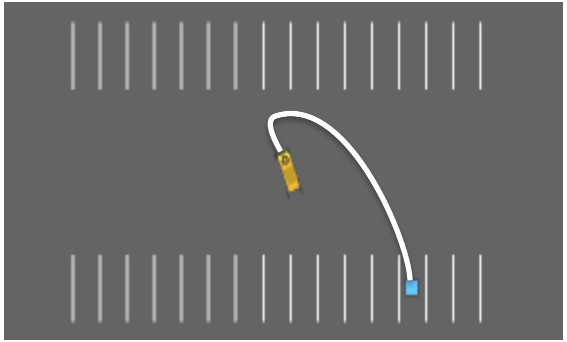

Full Trajectory

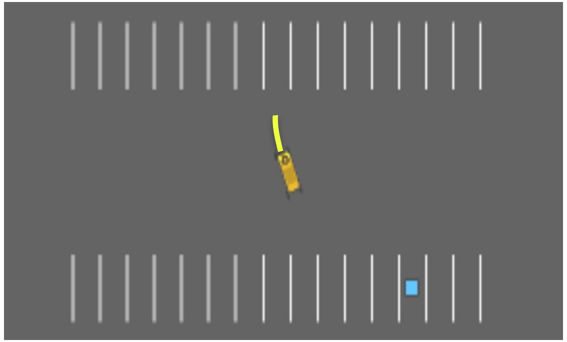

Time Heuristic Skills

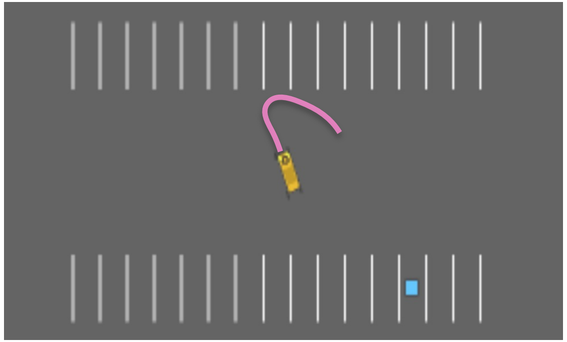

CompILE Skills

WRITING

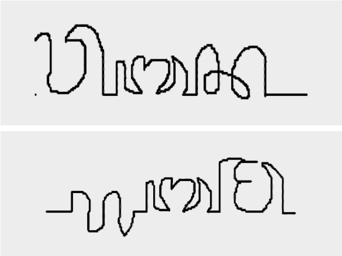

Full Trajectory

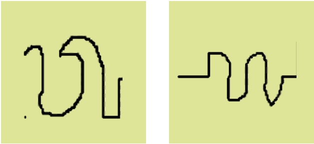

Time Heuristic Skills

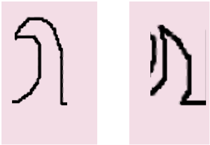

CompILE Skills

Are skills returned from CompILE useful for learning?

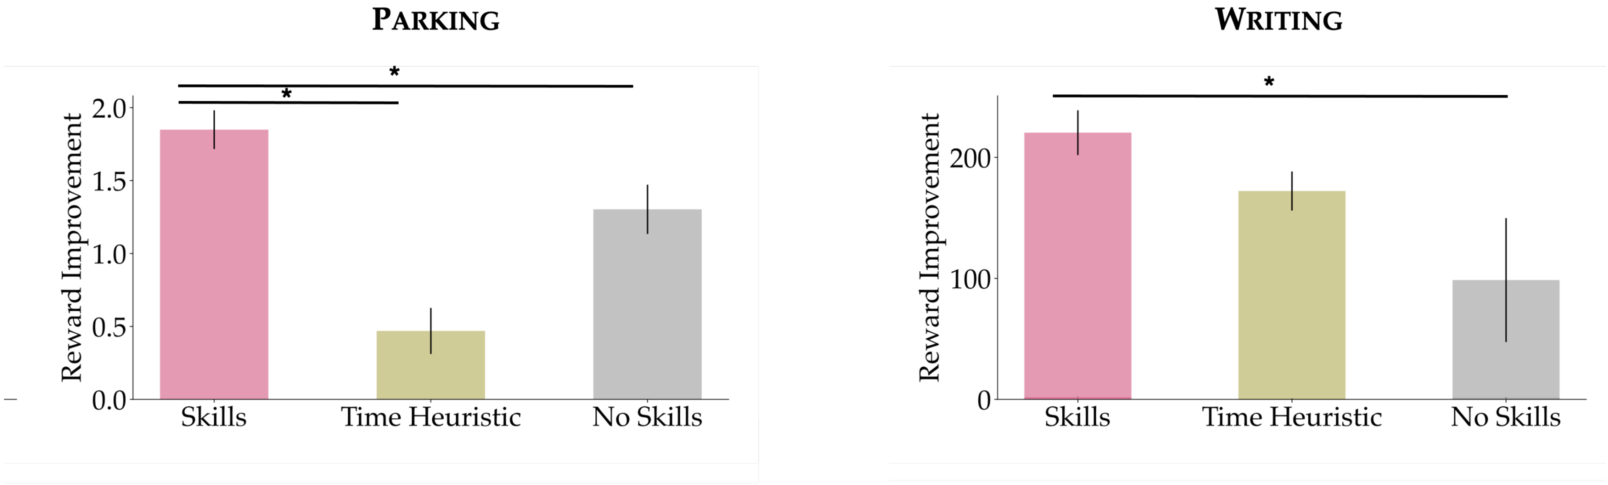

- Prolific user study (n=20 parking, n=25 writing)

- **Reward Improvement:** $\frac{\sum_{\xi_i^{\text{eval}}} r_i}{n_{\text{eval}}} - \frac{\sum_{\xi_i^{\text{pretest}}} r_i}{n_{\text{pretest}}}$

- **CompILE Skills** outperform **Full Trajectory**, **Time Heuristic** inconsistent!

Do individualized drills help students learn?

Do individualized drills help students learn?

PARKING

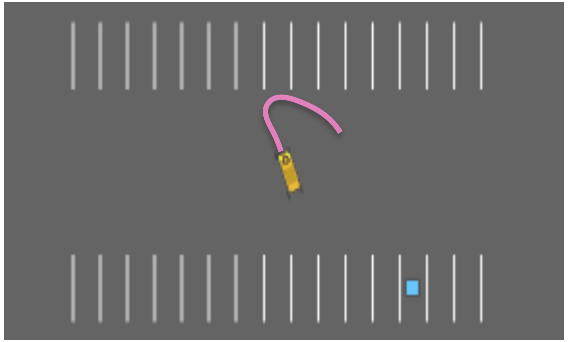

Skills Only

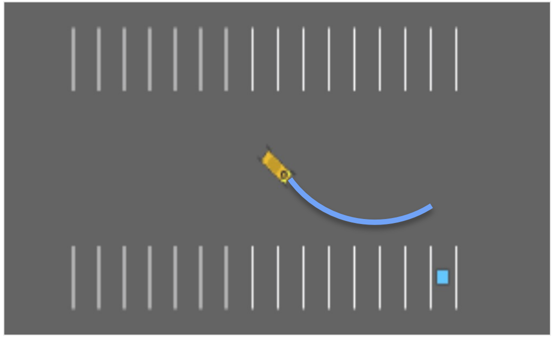

Randomly Chosen Drills

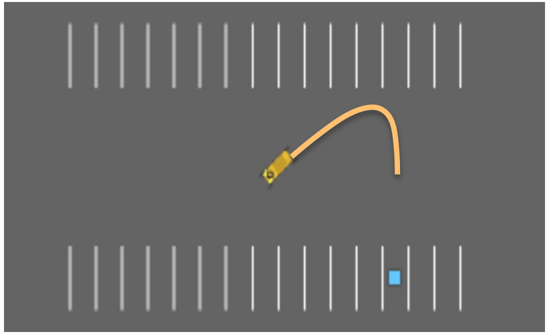

Individualized Drills

WRITING

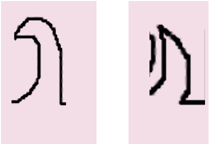

Skills Only

Randomly Chosen Drills

Individualized Drills

Distribution of hardest skills across individuals

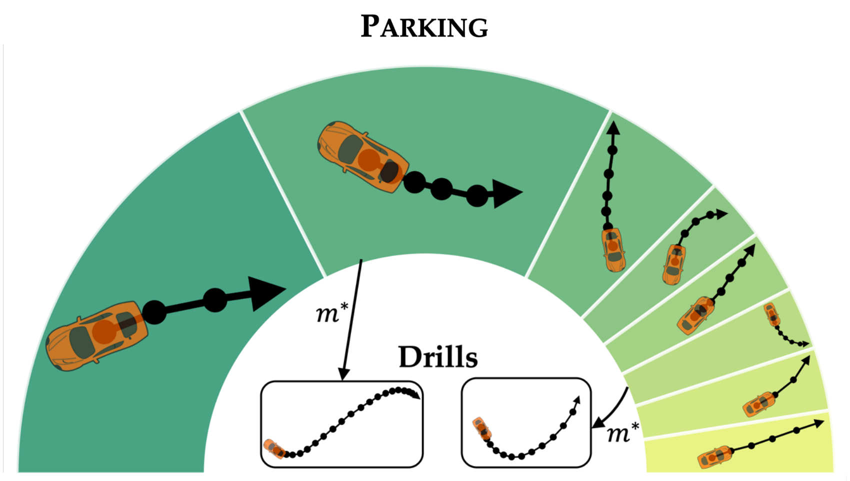

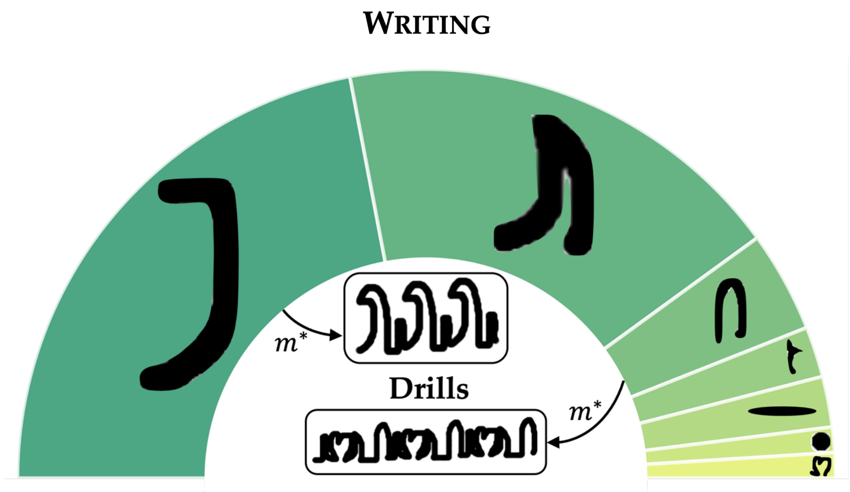

Do individualized drills help students learn?

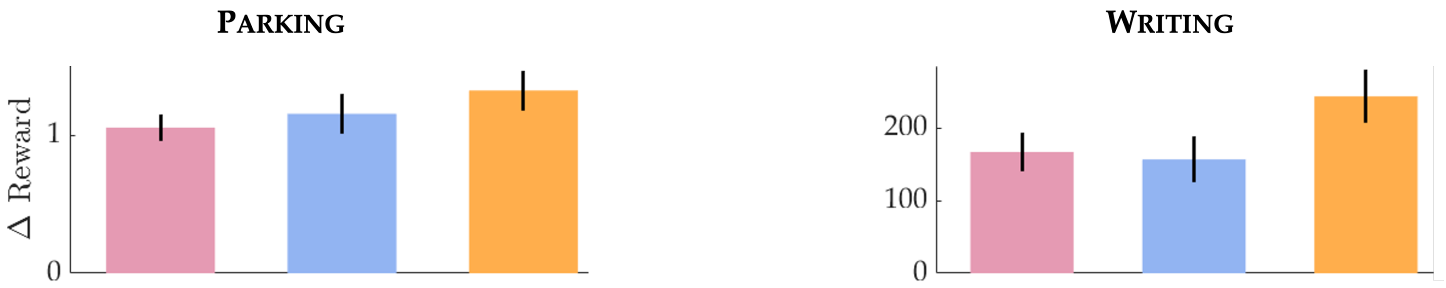

Individualized drills generally improves student performance

Do individualized drills help students learn?

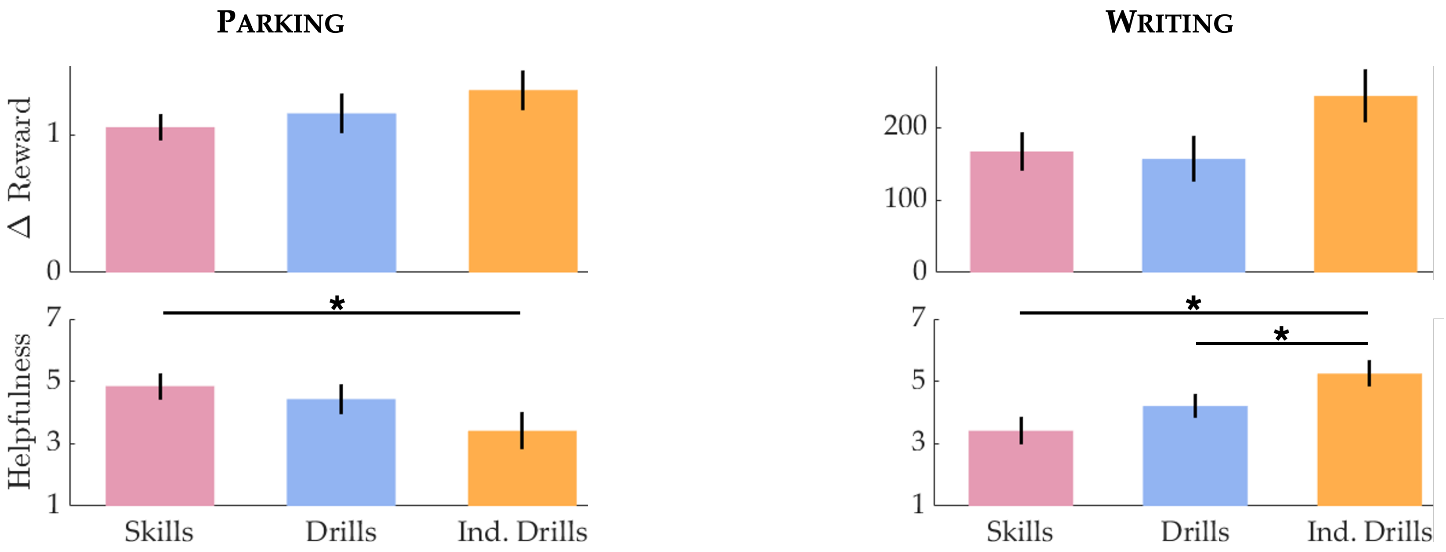

Do individualized drills help students learn?

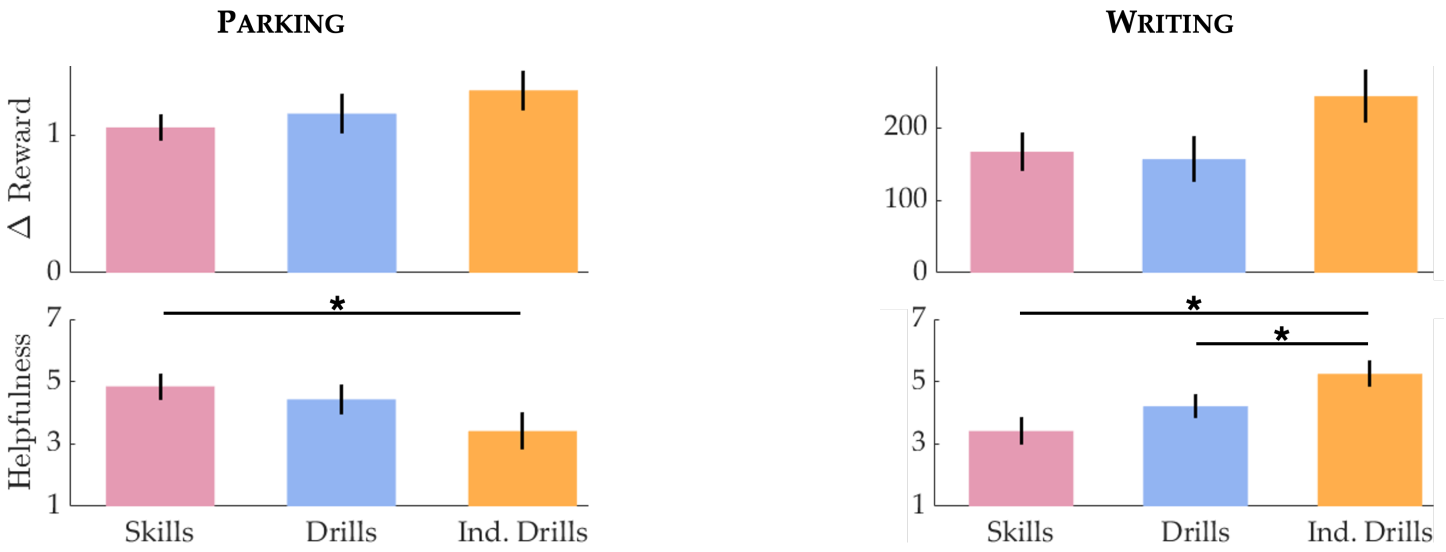

Participants significantly prefer ind. drills for Writing

Do individualized drills help students learn?

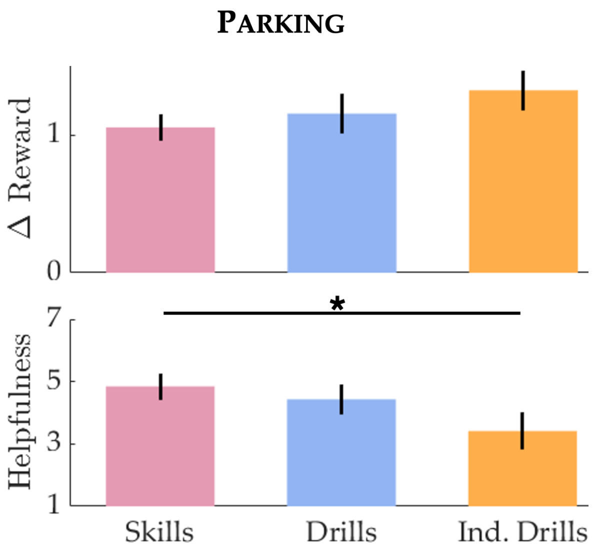

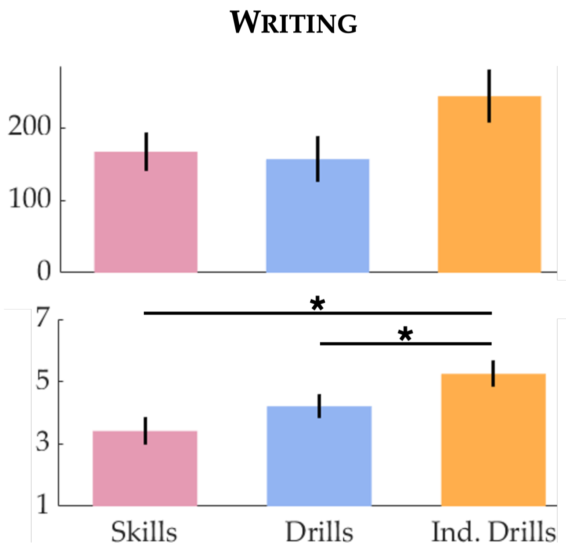

Participants significantly prefer skills over ind. drills for Parking *despite performing better w/ drills!*

Participants significantly prefer ind. drills for Writing

Optimal Expert Action:
Reverse (tricky!)

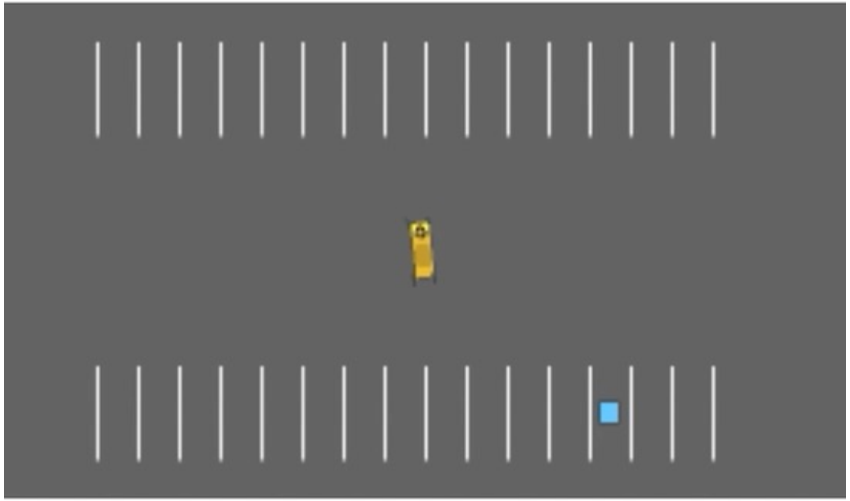

Students learn to more closely follow expert

No Individualization:
27% of students try to reverse

Individualization:
53% of students try to reverse
but find it hard!

Key Take-Aways

- AI-Assistance for skill discovery, individualization, and drill-creation
- Easier to “do” than “teach” → expert demonstrations can come from anyone!
- Participants benefited from AI-Assistance for two different control tasks

Future Directions

Stronger models of student motor learning

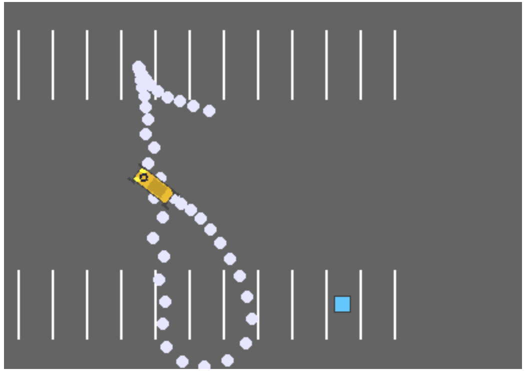

Half-Trained Student

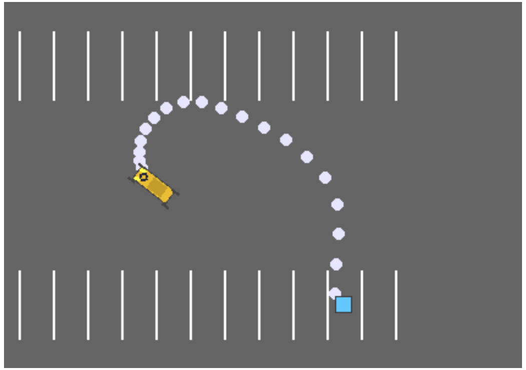

Reversing Difficulty Student

Future Directions

Risks of expert-student mismatch & accounting for preferences

"I hope this becomes a learning tool for writing new scripts.
Really cool concept just hard to get to grips with when I have a disability."
- User Study Participant