UNIVERSITY OF CAMBRIDGE
Department of Computer Science and Technology

IEEE RO-MAN 2020 VIRTUAL CONFERENCE

ORTA DOĞU TEKNİK ÜNİVERSİTESİ
MIDDLE EAST TECHNICAL UNIVERSITY

Continual Learning for Affective Robotics

Why, What and How?

Nikhil Churamani, Sinan Kalkan and Hatice Gunes

Affective Robots

Affective Robots

Healthcare

(1)

Affective Robots

Healthcare


(1)

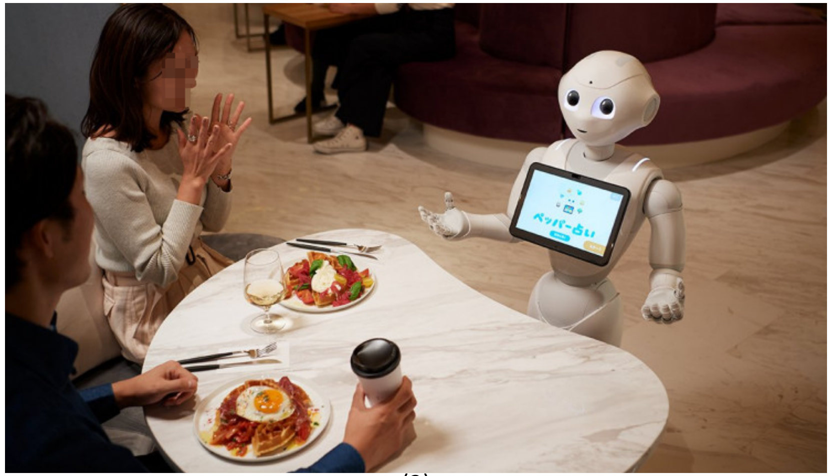
(2)

Service

Affective Robots

Healthcare

(1)

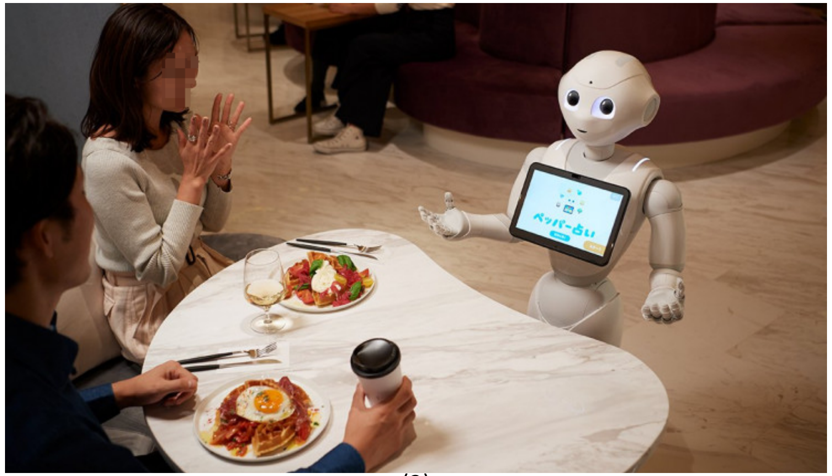

(2)

Service

Elderly care

(3)

Affective Robots

Healthcare

(1)

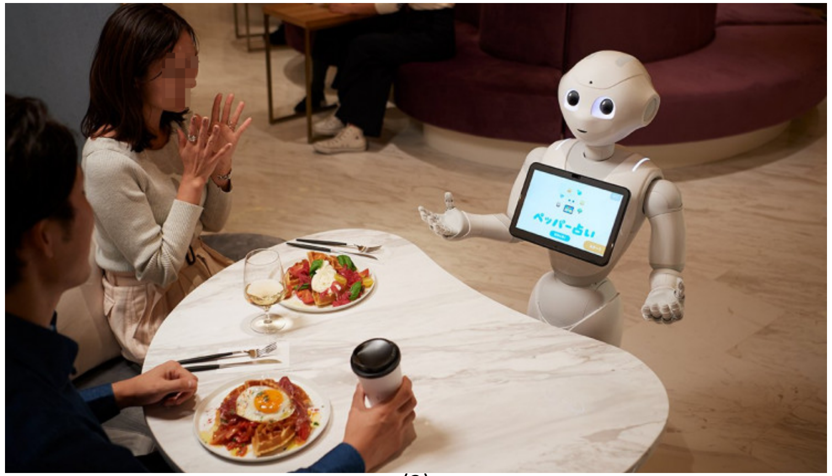

(2)

Service

Elderly care

(3)

(4)

Companion

UNIVERSITY OF CAMBRIDGE

(1) http://abovewhispers.com/2016/06/14/robot-receptionists-introduced-at-hospitals-in-belgium/
(2) https://interestingengineering.com/softbank-is-opening-a-cafe-where-pepper-robots-will-work-alongside-humans
(3) https://www.thetimes.co.uk/article/robot-carers-for-the-elderly-are-now-a-reality-in-japan-but-do-we-want-them-here-mw8zpw0zd
(4) https://customerthink.com/4-ways-social-robots-improve-customer-experience-in-retail-stores/

Need for Adaptation

Need for Adaptation

Equipped with
Learning Models

Need for Adaptation

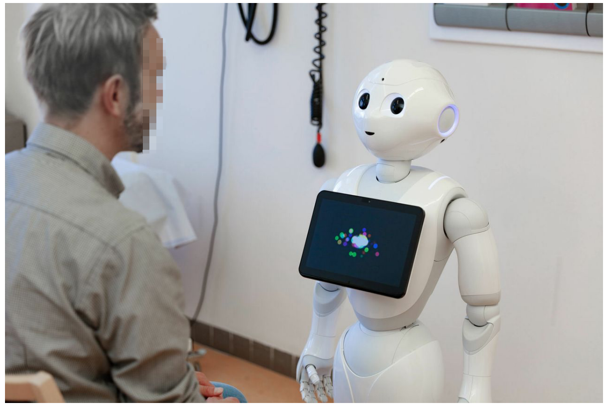

(a)

Equipped with
Learning Models

Interact with First User

Need for Adaptation

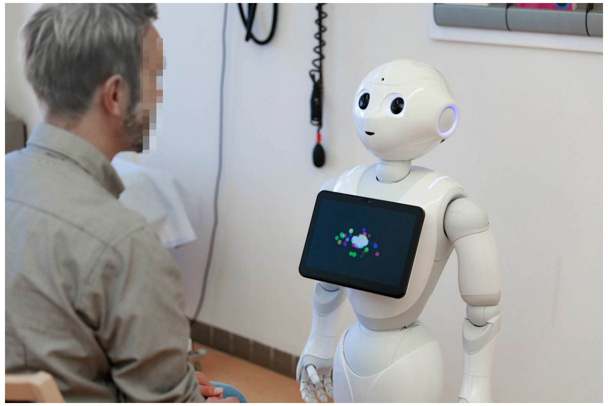

(a)

(b)

Equipped with
Learning Models

Interact with First User

Adapt with Second User.

Need for Adaptation

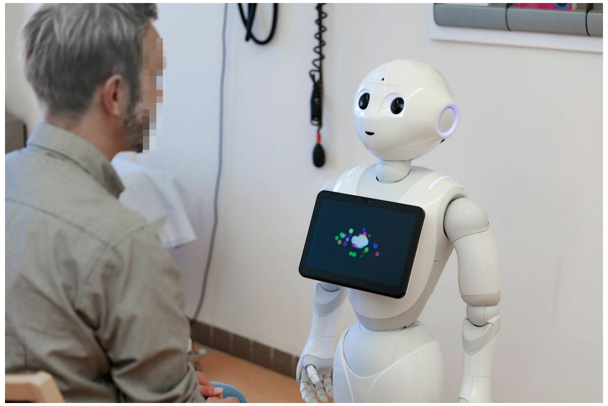
(a)


(b)

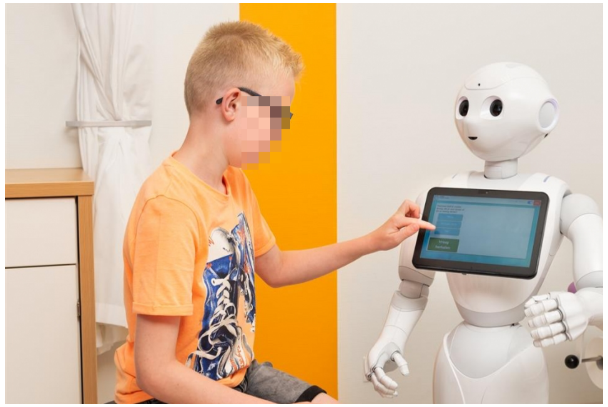
(c)

Equipped with Learning Models

Interact with First User

Adapt with Second User.

Adapt with Third User.

UNIVERSITY OF CAMBRIDGE

(a) Boumans R, van Meulen F, Hindriks K, *et al* Robot for health data acquisition among older adults: a pilot randomised controlled cross-over trial *BMJ Quality & Safety* 2019;**28**:793-799.
(b) https://www.scmp.com/lifestyle/health-wellness/article/3024028/how-robot-nurses-could-help-care-worlds-elderly-and
(c) https://eindhovennews.com/news/2018/06/robot-pepper-helps-children-hospital-visits/

Need for Adaptation

Equipped with Learning Models

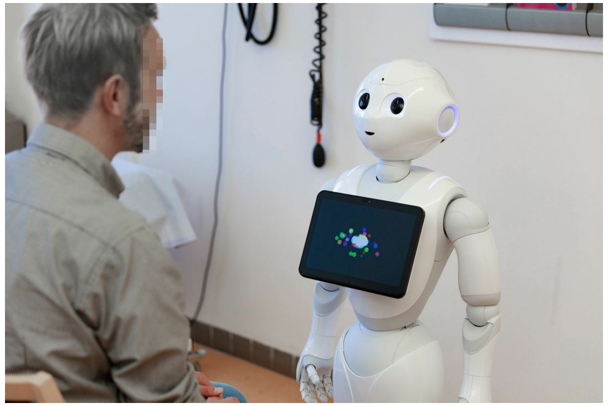

(a)

Interact with First User

(b)

Adapt with Second User.

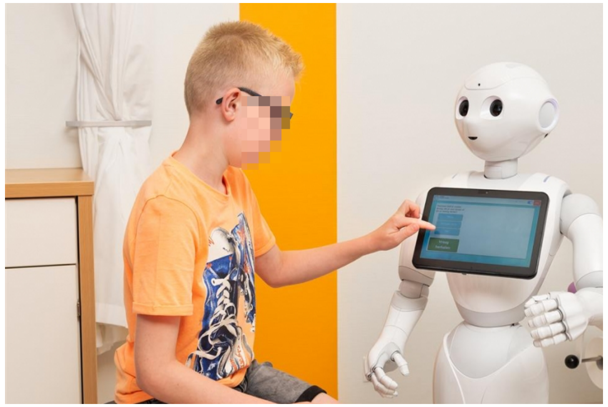

(c)

Adapt with Third User.

...

(a) Boumans R, van Meulen F, Hindriks K, *et al* Robot for health data acquisition among older adults: a pilot randomised controlled cross-over trial *BMJ Quality & Safety* 2019;**28**:793-799.
(b) https://www.scmp.com/lifestyle/health-wellness/article/3024028/how-robot-nurses-could-help-care-worlds-elderly-and
(c) https://eindhovennews.com/news/2018/06/robot-pepper-helps-children-hospital-visits/

Traditional Approaches vs. Continual Learning

Traditional Approaches | **Continual Learning**

T. Lesort *et al.*, “Continual learning for robotics: Definition, frame- work, learning strategies, opportunities and challenges,” *Information Fusion*, vol. 58, pp. 52–68, 2020.

Traditional Approaches vs. Continual Learning

Traditional Approaches

- Models **trained in isolation** on benchmark datasets.
- Large datasets enable **generalisation** across contexts.
- Training data might be very **different** from application scenarios.
- Generalisation comes **at the cost** of learning individual differences.
- **Cumbersome** to retrain and update models.

Continual Learning

UNIVERSITY OF CAMBRIDGE

T. Lesort *et al.*, "Continual learning for robotics: Definition, frame- work, learning strategies, opportunities and challenges," *Information Fusion*, vol. 58, pp. 52–68, 2020.

Traditional Approaches vs. Continual Learning

Traditional Approaches

- Models **trained in isolation** on benchmark datasets.
- Large datasets enable **generalisation** across contexts.
- Training data might be very **different** from application scenarios.
- Generalisation comes **at the cost** of learning individual differences.
- **Cumbersome** to retrain and update models.

Continual Learning

- Agents **acquire** and **integrate** knowledge **incrementally** about changing environments.
- Data only made available **sequentially**.
- Highly sensitive towards changing data conditions.
- **Adaptations** in learning to avoid forgetting.
- CL Problem Formulation:

T. Lesort *et al.*, "Continual learning for robotics: Definition, frame- work, learning strategies, opportunities and challenges," *Information Fusion*, vol. 58, pp. 52–68, 2020.

Traditional Approaches vs. Continual Learning

Traditional Approaches

- Models **trained in isolation** on benchmark datasets.
- Large datasets enable **generalisation** across contexts.
- Training data might be very **different** from application scenarios.
- Generalisation comes **at the cost** of learning individual differences.
- **Cumbersome** to retrain and update models.

Continual Learning

- Agents **acquire** and **integrate** knowledge **incrementally** about changing environments.
- Data only made available **sequentially**.
- Highly sensitive towards changing data conditions.
- **Adaptations** in learning to avoid forgetting.
- CL Problem Formulation:

$$A_i^{CL} : \langle \underbrace{h_{i-1}}_{\text{Model}}, \overbrace{Tr_i}^{\text{New Data}}, \underbrace{M_{i-1}}_{\text{Experience}}, \overbrace{t_i}^{\text{Task}} \rangle \rightarrow \langle \underbrace{h_i}_{\text{Improved Model}}, \overbrace{M_i}^{\text{Updated Experience}} \rangle$$

T. Lesort *et al.*, "Continual learning for robotics: Definition, frame- work, learning strategies, opportunities and challenges," *Information Fusion*, vol. 58, pp. 52–68, 2020.

Personalised Affect Perception

Personalised Affect Perception

User (C_u) :

- Age: 30-40
- Gender: Female
- Personality: Extrovert

Context (T_i) :

- Environment: Living_Room
- Objective: Home_Assist

Task (t_i) :

- Task: Engagement

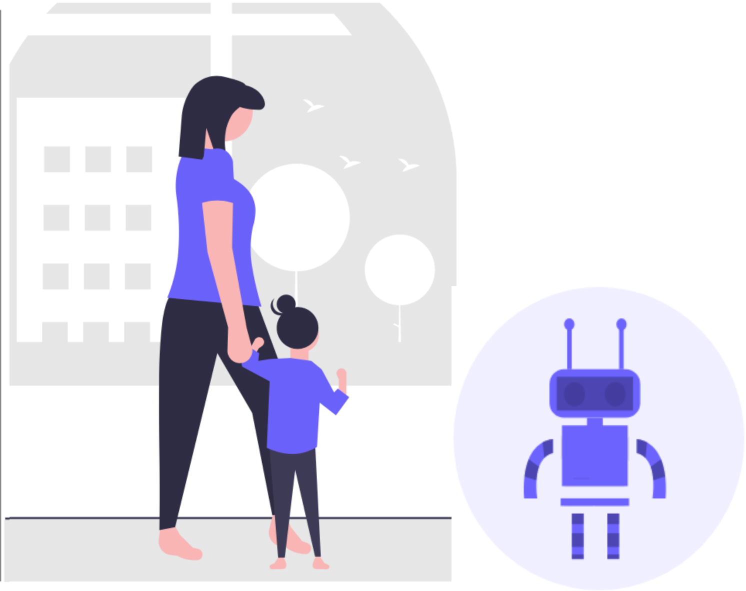

Personalised Affect Perception

Traditional ML

User (C_u) :
- Age: 30-40
- Gender: Female
- Personality: Extrovert

Context (T_i) :
- Environment: Living_Room
- Objective: Home_Assist

Task (t_i) :
- Task: Engagement

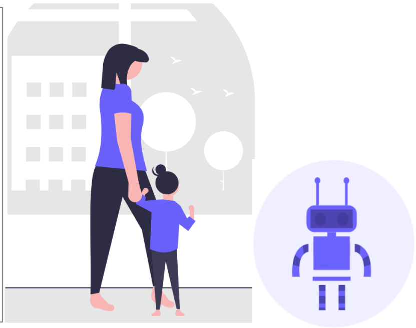

Continual Learning

Personalised Affect Perception

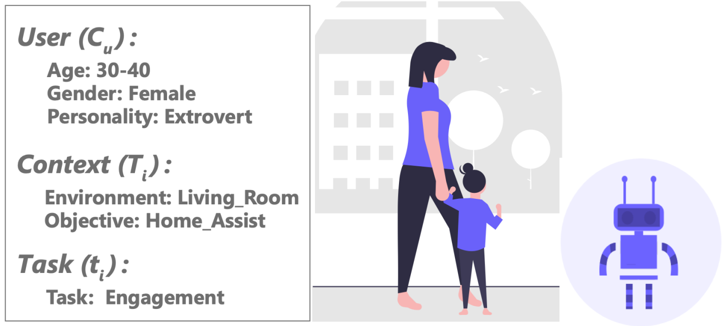

Traditional ML

- Robot perception is pre-trained and does **not** adapt.
- Pre-training on '*large-enough*' datasets to ensure generalisation.
- Individual differences in expression are **ignored**.

Continual Learning

Personalised Affect Perception

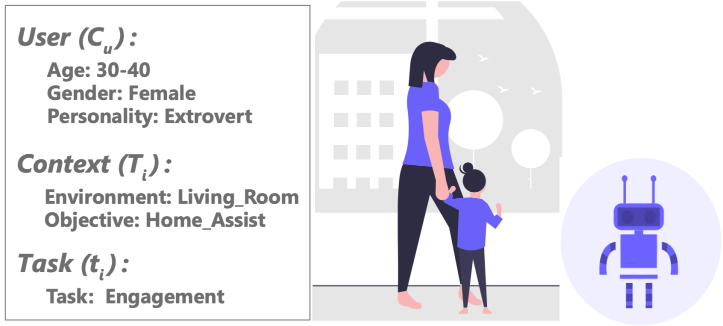

Traditional ML

- Robot perception is pre-trained and does **not** adapt.
- Pre-training on '*large-enough*' datasets to ensure generalisation.
- Individual differences in expression are **ignored**.

Continual Learning

- Robot uses pre-trained perception and **adapts** it with each interaction.
- Interactions help to gather more **data**.
- User-specific attributes provide **context**.
- Agent *continually* updates its perception.
  - With each interaction.
  - Personalising towards each user.

$$\forall u \in \mathcal{U}, \forall i \in \mathcal{I}, P^{CL}_{u,i} : \langle h_{i-1}, Tr_i, M_{i-1}, C_u, t_i \rangle \rightarrow \langle h_i, M_i \rangle$$

Context-specific Robot Behaviour Generation

Context-specific Robot Behaviour Generation

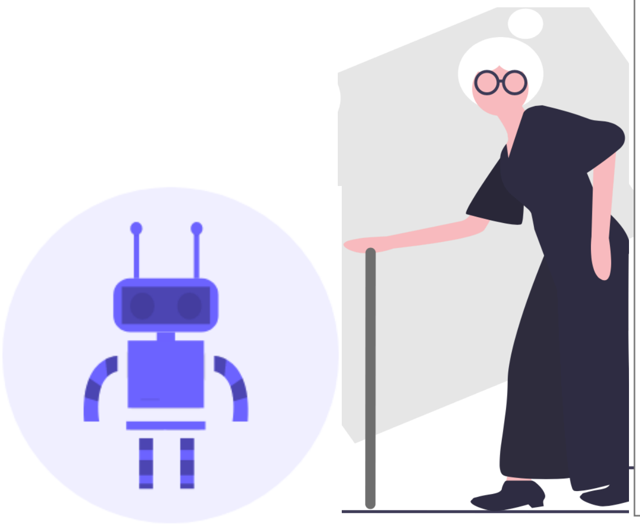

User (C_u) :

Age: 60-70
Gender: Female
Personality: Introvert
Condition: Arthritis

Context (T_i) :

Environment: Care_Home
Objective: Care_Assist

Task (t_i) :

Task: Medicine_Reminder

Context-specific Robot Behaviour Generation

Traditional ML

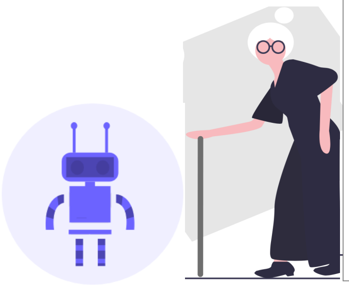

User (C_u) :

Age: 60-70
Gender: Female
Personality: Introvert
Condition: Arthritis

Context (T_i) :

Environment: Care_Home
Objective: Care_Assist

Task (t_i) :

Task: Medicine_Reminder

Continual Learning

Context-specific Robot Behaviour Generation

Traditional ML

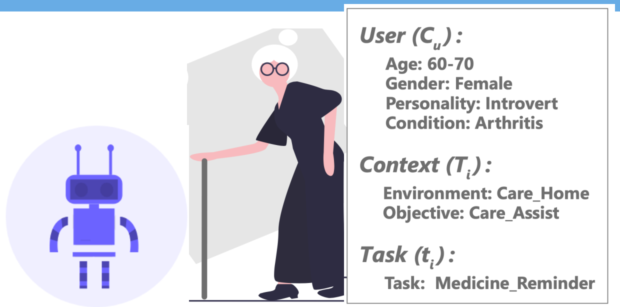

Continual Learning

- Robot interacts in **restricted settings**.
- RL models perfect **task** behaviours.
- Context and task-switching is **intractable**.
- **Large amount** of interactions needed to learn.
- Interactive RL with human in the loop helps *speed up* convergence.

Context-specific Robot Behaviour Generation

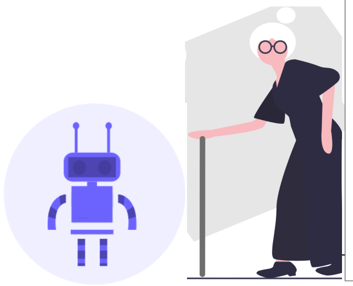

User (C_u) :
Age: 60-70
Gender: Female
Personality: Introvert
Condition: Arthritis

Context (T_i) :
Environment: Care_Home
Objective: Care_Assist

Task (t_i) :
Task: Medicine_Reminder

Traditional ML

- Robot interacts in **restricted settings**.
- RL models perfect **task** behaviours.
- Context and task-switching is **intractable**.
- **Large amount** of interactions needed to learn.
- Interactive RL with human in the loop helps *speed up* convergence.

Continual Learning

- Robot tries to **interact** with the user.
- User feedback provides **data** for learning.
- User-specific attributes provide **affordances**.
- Task-specific attributes provide **context**.
- Agent *continually* improves its behaviour at the task.
  - Improving its behaviour with each interaction.
  - Extending this learning to different users.

$$\forall u \in \mathcal{U}, \forall i \in \mathcal{I}, B^{CL}_{u,i} : \langle h_{i-1}, Tr_i, M_{i-1}, C_u, T_i, t_i \rangle \rightarrow \langle h_i, M_i \rangle$$

UNIVERSITY OF CAMBRIDGE

Continual Learning for Affective Robotics

Continual Learning for Affective Robotics

Continual Learning for Affective Robotics

User (C_u) :
- Age: 30-40
- Gender: Female
- Personality: Extrovert

Context (T_i) :
- Environment: Living_Room
- Objective: Home_Assist

Task (t_i) :
- Task: Engagement

NI

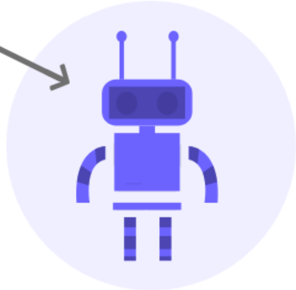

NI: New Instances

Continual Learning for Affective Robotics

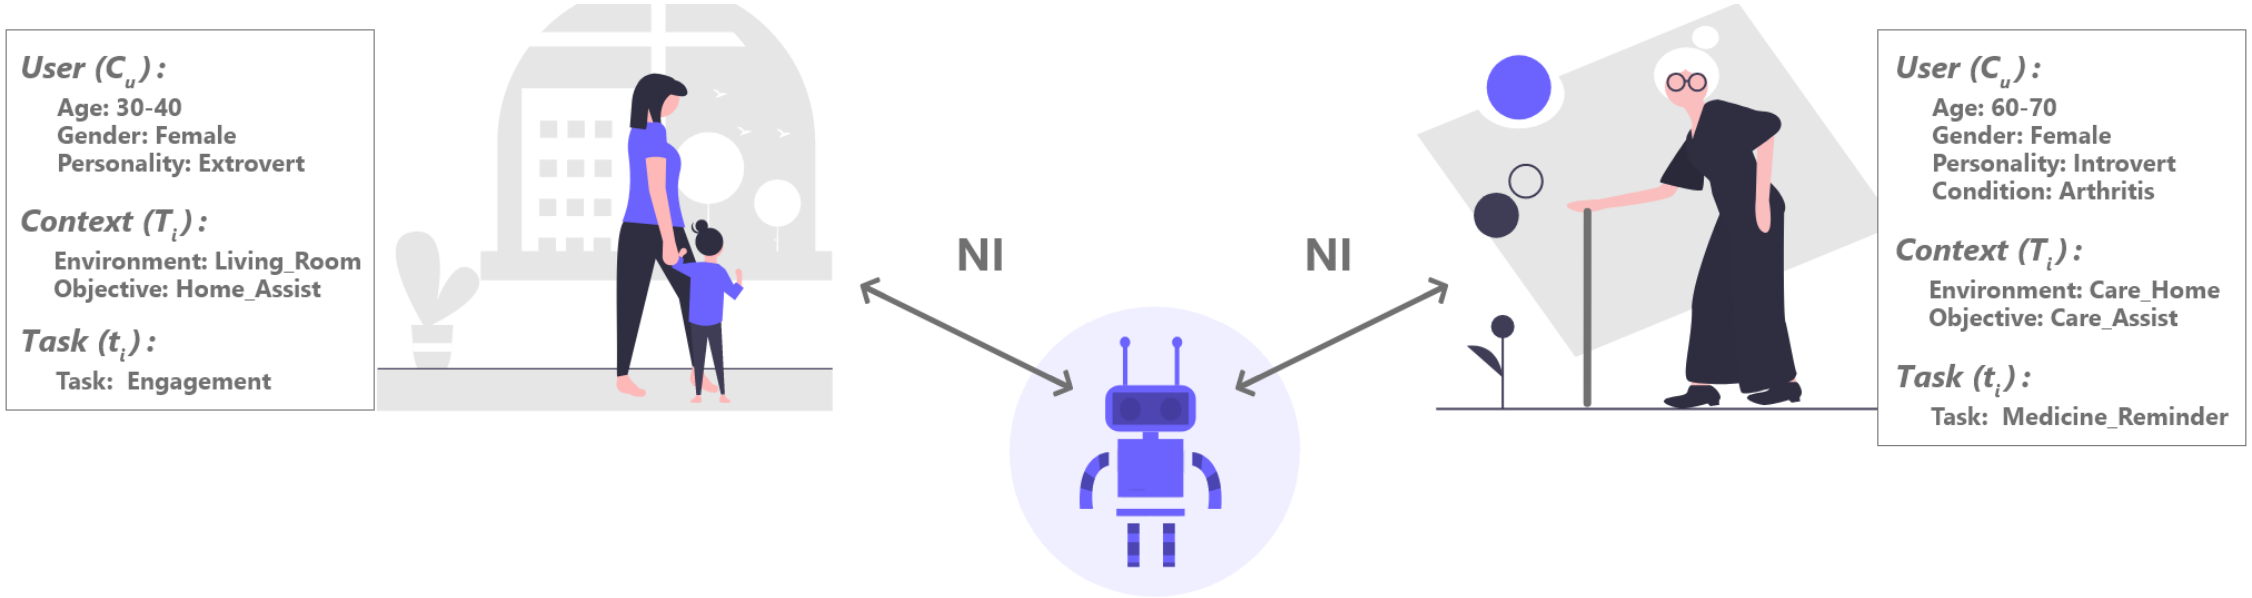

NI: New Instances

Continual Learning for Affective Robotics

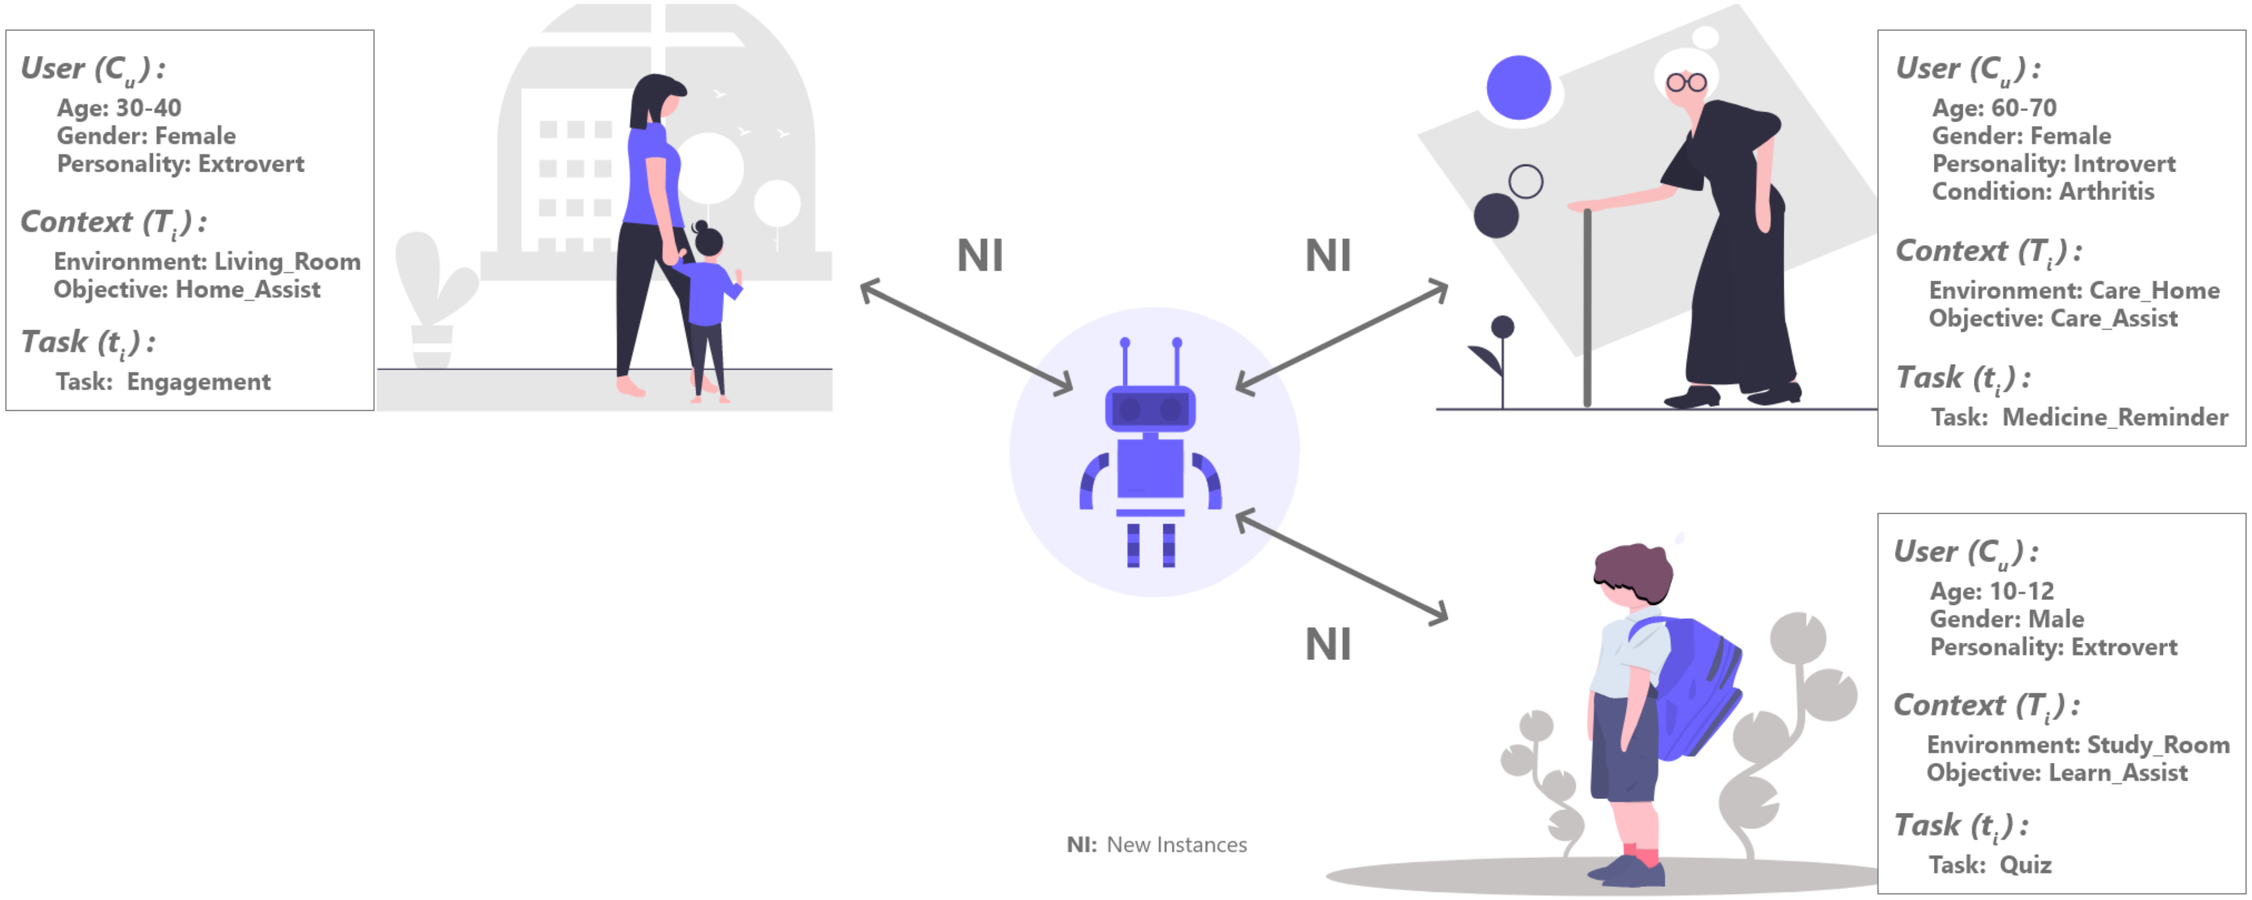

Continual Learning for Affective Robotics

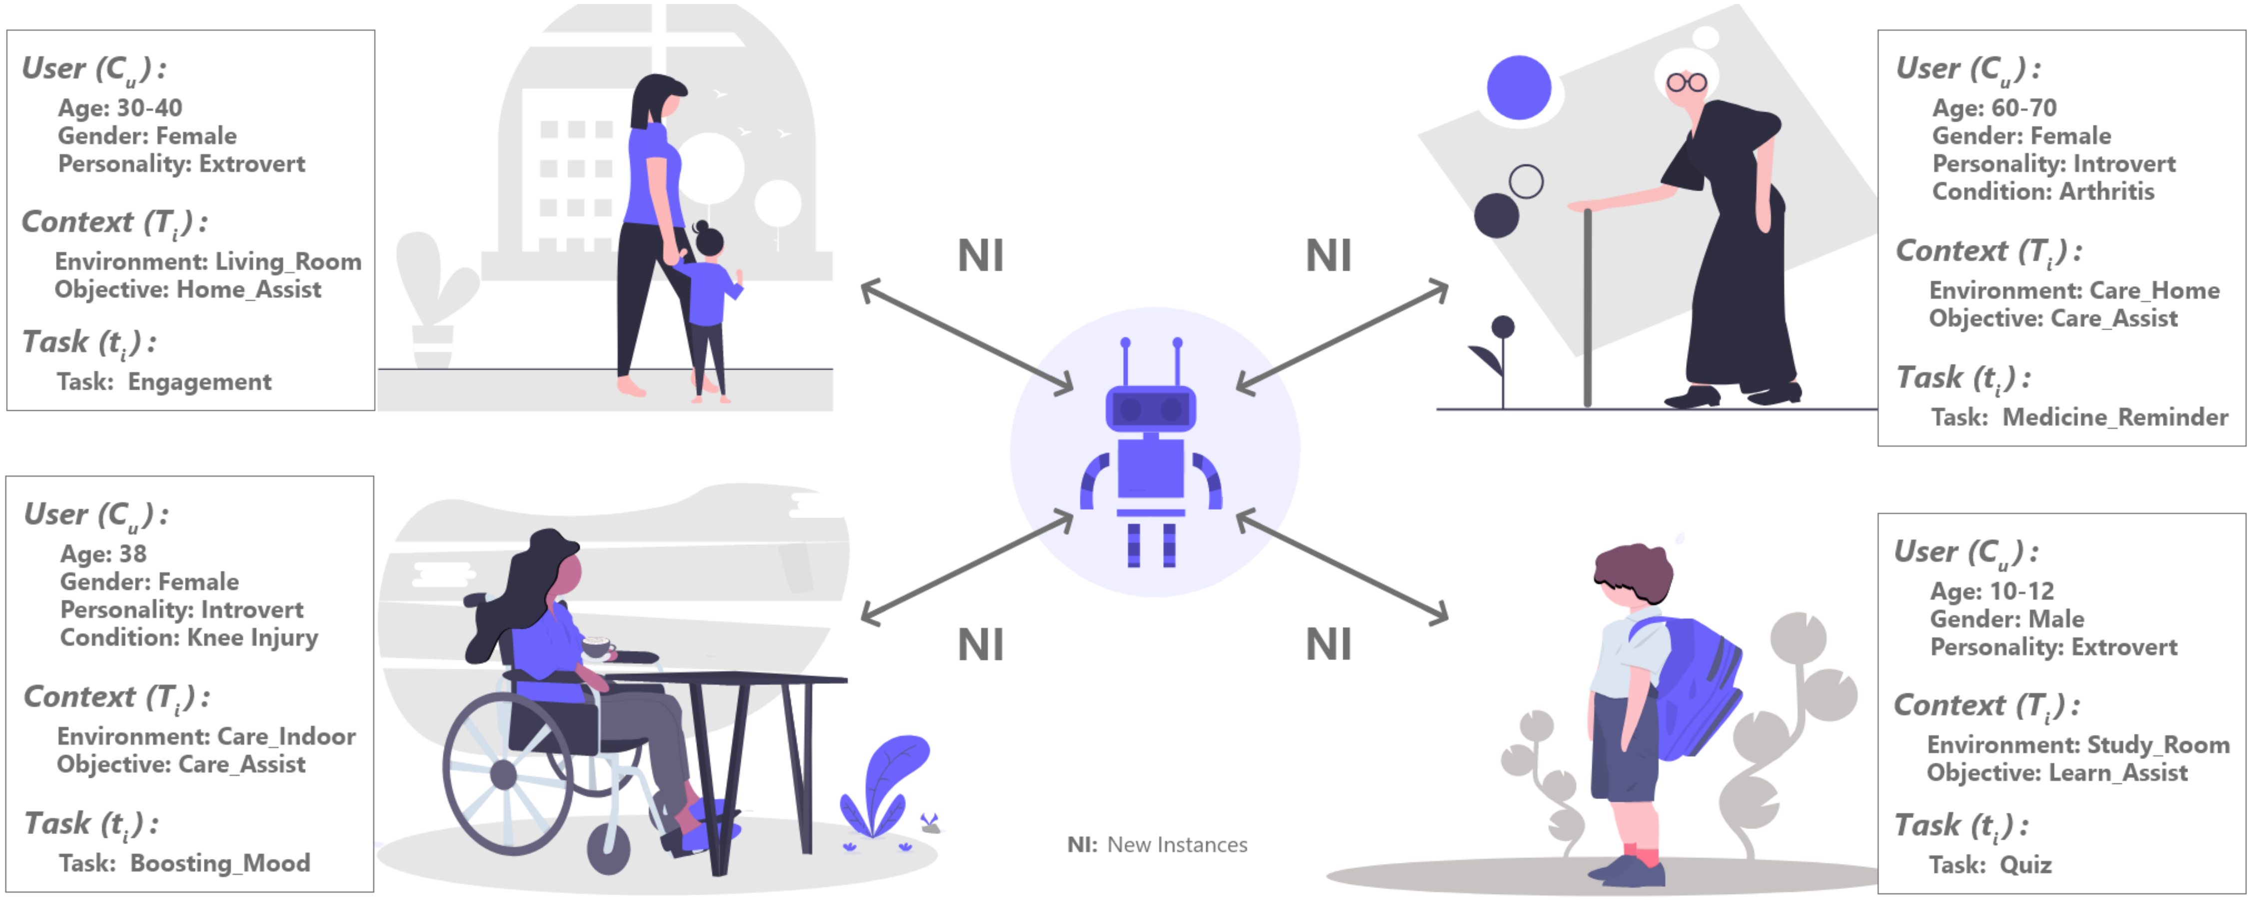

Continual Learning for Affective Robotics

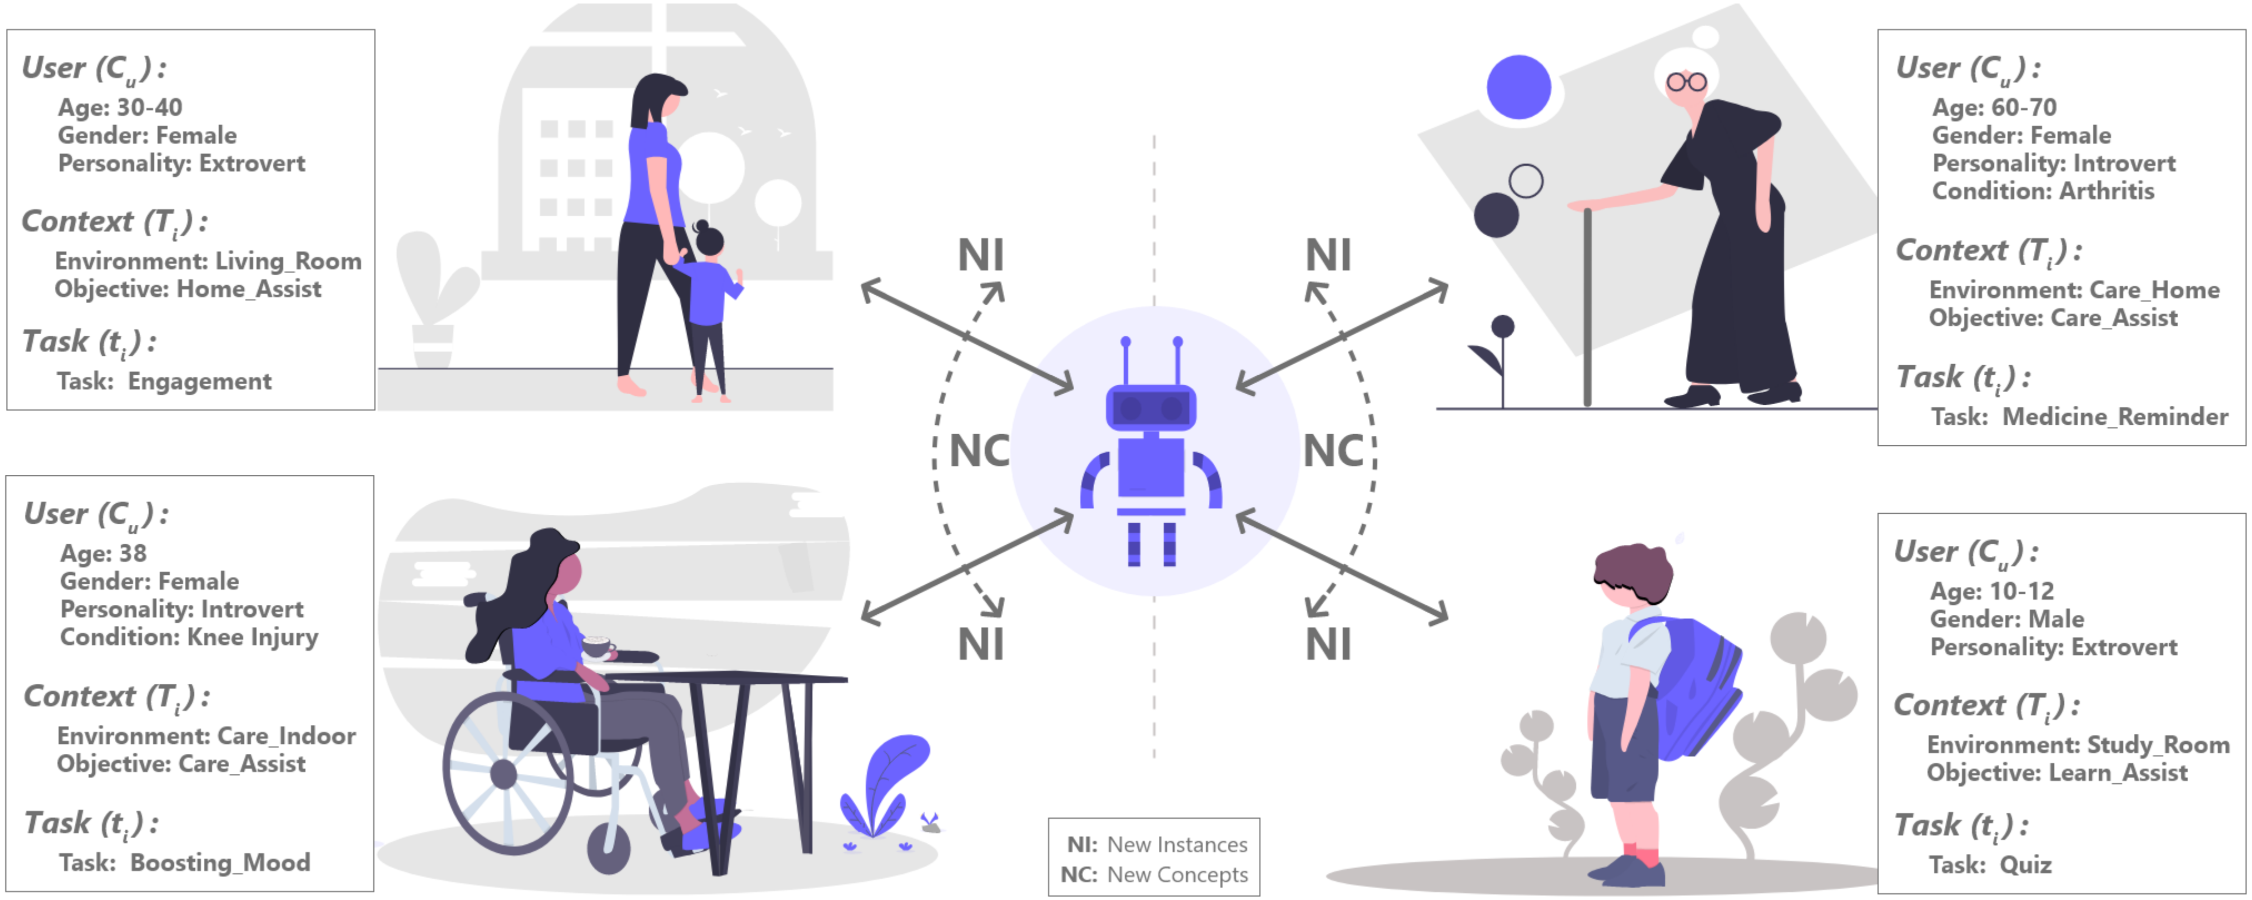

Challenges and Recommendations

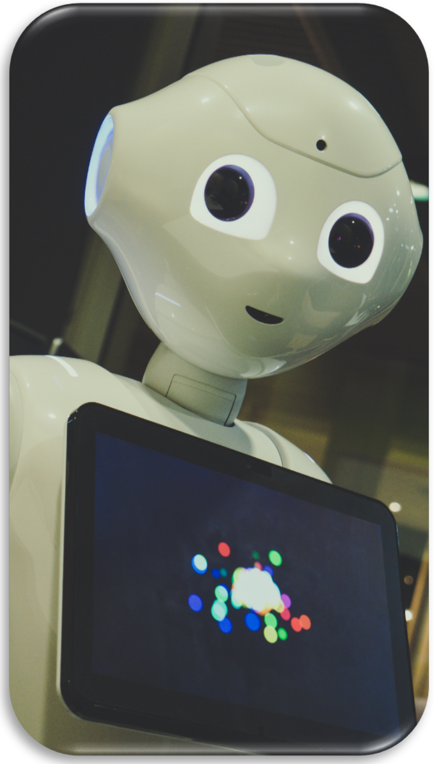

Challenges and Recommendations

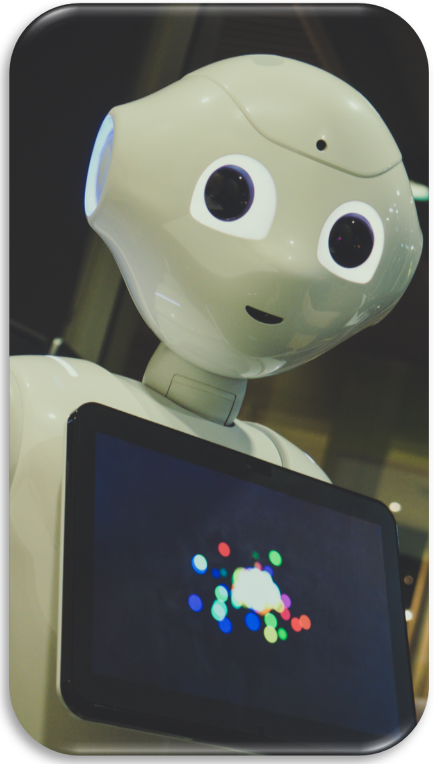

Gathering Person-specific Data

UNIVERSITY OF CAMBRIDGE

Challenges and Recommendations

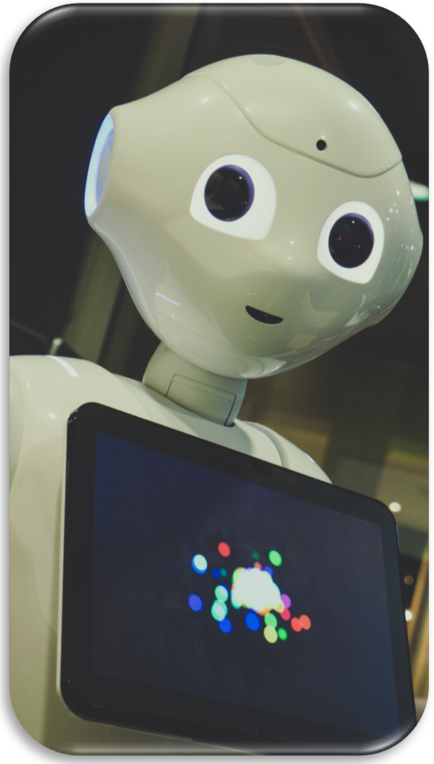

Gathering Person-specific Data

Why?

- **Interactions** are the only source of data.
- Initial interactions impacted due to **slow learning**.
- User specific **data unavailable** before any interaction.

Challenges and Recommendations

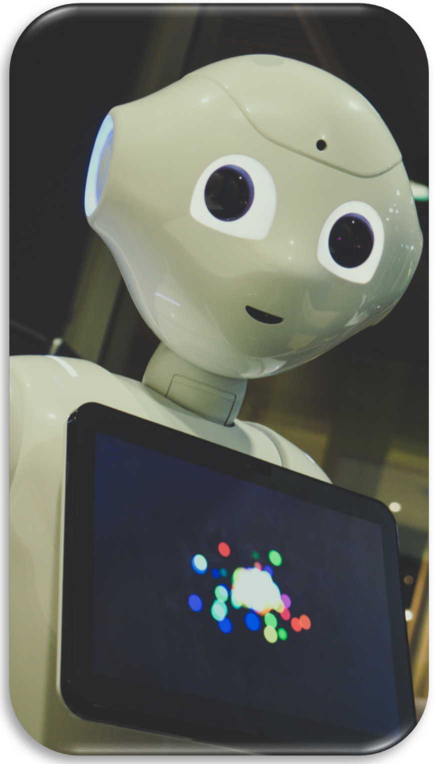

Gathering Person-specific Data

Why?

- **Interactions** are the only source of data.
- Initial interactions impacted due to **slow learning**.
- User specific **data unavailable** before any interaction.

How?

- Conduct **introductory interaction rounds** enable collecting additional data.
- Use a generative model to **simulate** additional person-specific data for **augmenting** learning.

Challenges and Recommendations

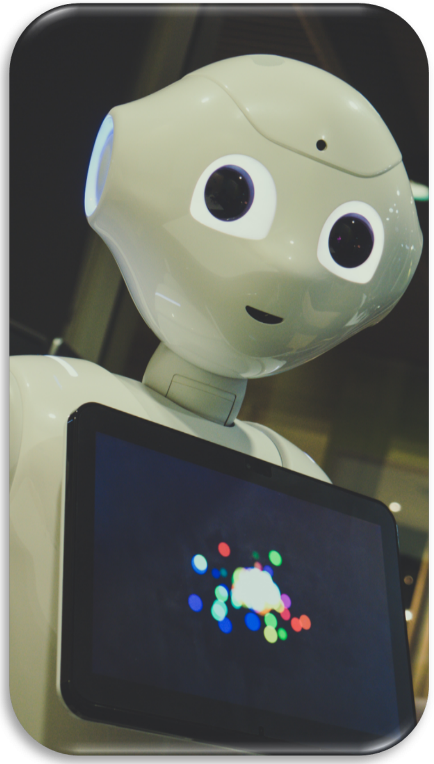

Obtaining Ground Truth

Challenges and Recommendations

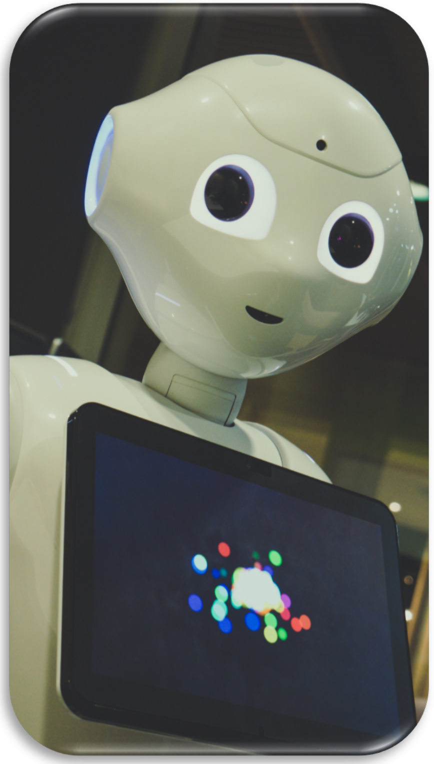

Obtaining Ground Truth

Why?

- Human affect is **subjective**.
- Ground truth **changes** with users and contexts.
- Unsupervised learning may be **intractable** in real-time.

Challenges and Recommendations

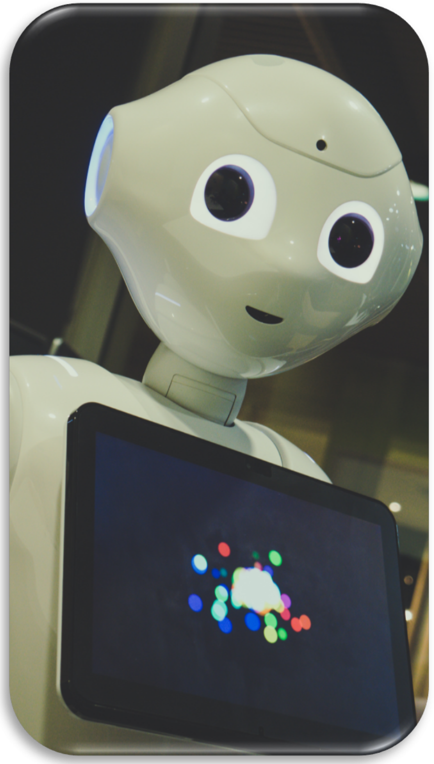

Obtaining Ground Truth

Why?

- Human affect is **subjective**.
- Ground truth **changes** with users and contexts.
- Unsupervised learning may be **intractable** in real-time.

How?

Learn Normative Baselines

- Contextually neutral interactions provide a **baseline** for measuring human behaviour.

Learn Semantic Associations

- Group users based on person-specific attributes to **speed up** learning.

Challenges and Recommendations

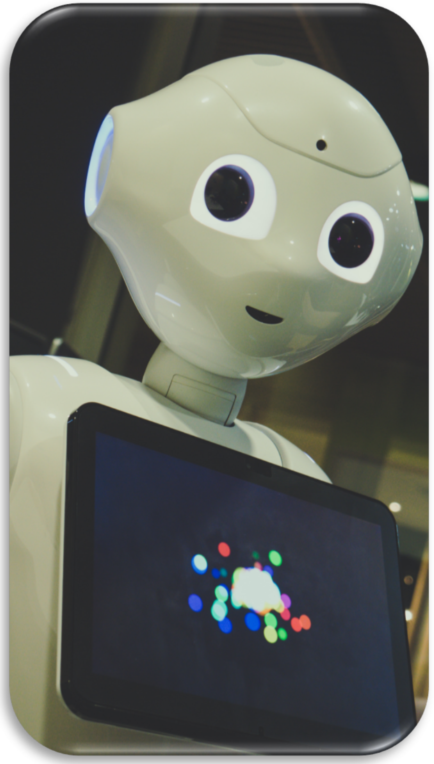

Learning without Task Boundaries

Challenges and Recommendations

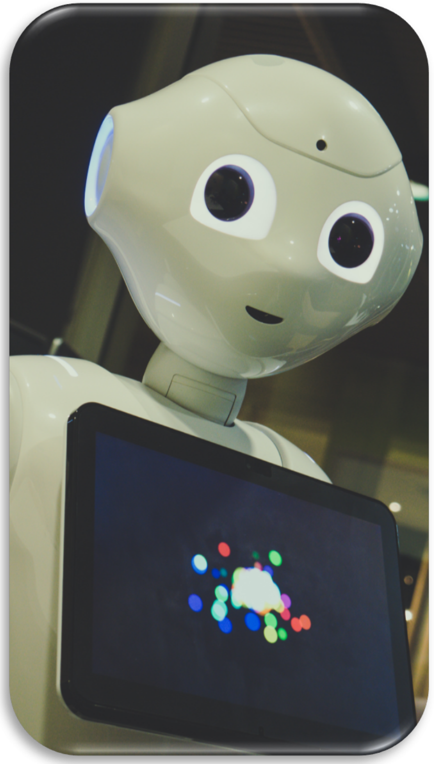

Learning without Task Boundaries

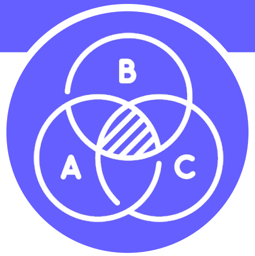

Why?

- Human interactions are fluid and **toggle** between contexts.
- Robots need robust and quick **context-switching**.
- Contextual attributions required for learning may **overlap**.

Challenges and Recommendations

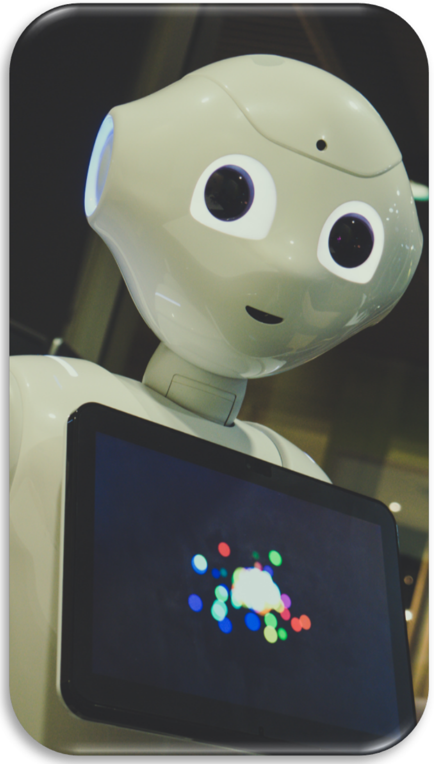

Learning without Task Boundaries

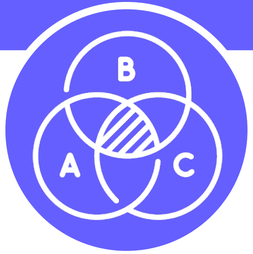

Why?

- Human interactions are fluid and **toggle** between contexts.
- Robots need robust and quick **context-switching**.
- Contextual attributions required for learning may **overlap**.

How?

Learn Contextual Attributions

- Context-aware embeddings enable **distinguishing** between task boundaries.
- Context attributes (e.g. environment or objective) facilitate **context-switching**.

Challenges and Recommendations

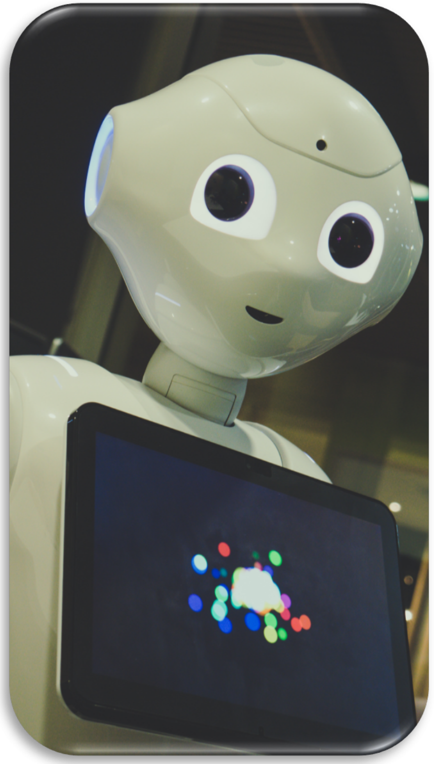

Hardware and Memory

Challenges and Recommendations

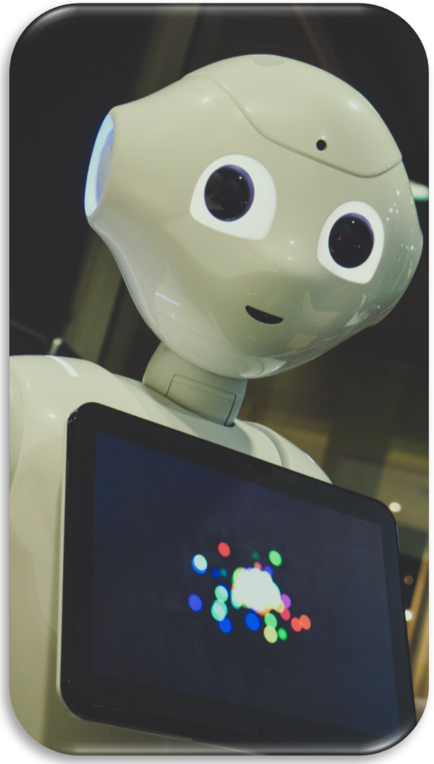

Hardware and Memory

Why?

- Hardware capabilities on robots are **limited** for online learning.
- Memory-Computation **trade-off** needs to be considered.
- Online-adaptation becomes **challenging**.

Challenges and Recommendations

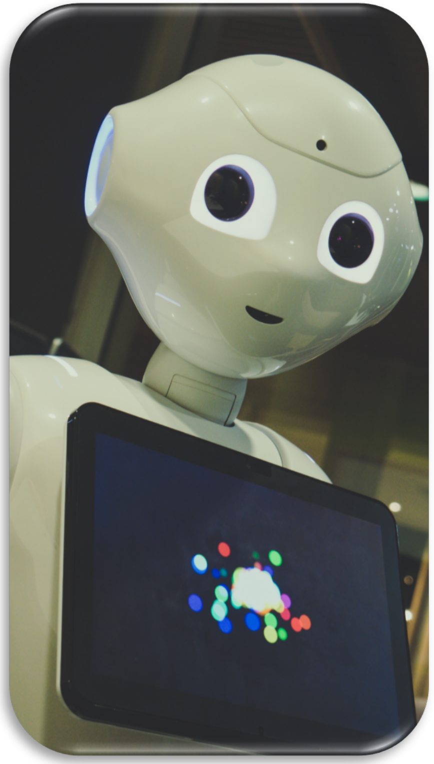

Hardware and Memory

Why?

- Hardware capabilities on robots are **limited** for online learning.
- Memory-Computation **trade-off** needs to be considered.
- Online-adaptation becomes **challenging**.

How?

Allow Controlled Forgetting

- Biologically-inspired forgetting makes way for **new knowledge**.

Balance Memory with Computation

- Generative models for **pseudo-rehearsal**.
- Robotics as Service (**RaaS**).

Challenges and Recommendations

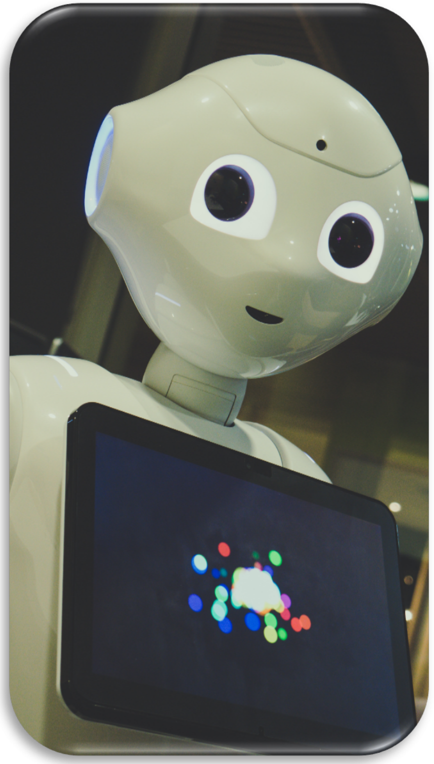

Evaluations

Challenges and Recommendations

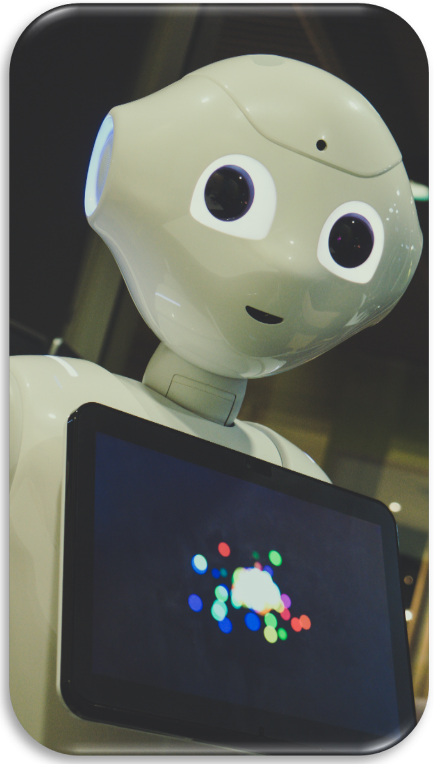

Evaluations

Why?

- Evaluations provide guarantees for **Fairness** and **Reproducibility**.
- Continual Learning metrics evaluate **robustness** towards dynamic shifts in data distributions.

Challenges and Recommendations

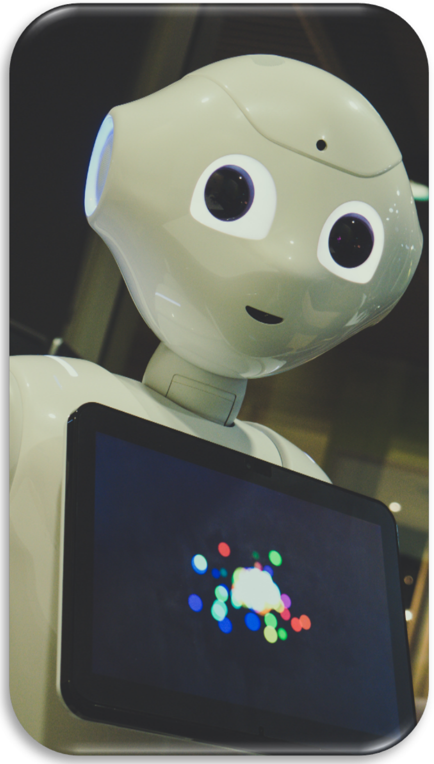

Evaluations

Why?

- Evaluations provide guarantees for **Fairness** and **Reproducibility**.
- Continual Learning metrics evaluate **robustness** towards dynamic shifts in data distributions.

How?

- **CL metrics** (e.g. *Forward* and *Backward Transfer* or *Catastrophic Forgetting rate*) measure **robustness**.
- **Traditional Metrics** (e.g. F1 or AUC-ROC scores) measure task **performance**.

Conclusion

- **Real-world** interactions are **complex** and **unpredictable**.
- Affective Robots need to **adapt on-the-fly**.
- Desiderata from Affective Robots; **personalisation** and behaviour **adaptation**.
- *Continual Learning* enables **perpetual evolution** of robot capabilities.

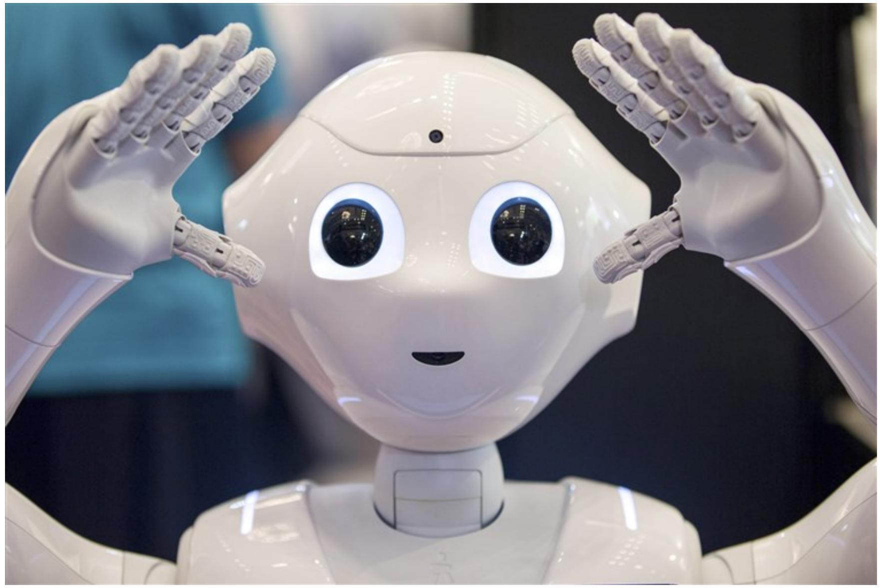

UNIVERSITY OF CAMBRIDGE

https://www.nbcnews.com/tech/tech-news/we-re-entering-age-friendly-robots-n703336

Acknowledgement

Nikhil Churamani

Sinan Kalkan

Hatice Gunes

Nikhil Churamani is supported by the EPSRC (grant EP/R513180/1 ref. 2107412).

Hatice Gunes is supported by the EPSRC (grant ref. EP/R030782/1) and the Alan Turing Institute Faculty Fellowship (G102185).

S. Kalkan is supported by Scientific and Technological Research Council of Turkey (TUBITAK) through BIDEB 2219 International Postdoctoral Research Scholarship Program.

Engineering and Physical Sciences Research Council

UNIVERSITY OF CAMBRIDGE

UNIVERSITY OF CAMBRIDGE
Department of Computer Science and Technology

IEEE RO-MAN
2020 VIRTUAL CONFERENCE

ORTA DOĞU TEKNİK ÜNİVERSİTESİ
MIDDLE EAST TECHNICAL UNIVERSITY

Continual Learning for Affective Robotics

Why, What and How?

Nikhil Churamani, Sinan Kalkan and Hatice Gunes