Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

b-it Bonn-Aachen International Center for Information Technology

Institute for AI and Autonomous Systems

Robot Personalisation

Autonomously Adapting to an Individual's Needs

Dr. Alex Mitrevski
Master of Autonomous Systems

Structure

- Personalisation preliminaries
- Learning for personalisation

Personalized Behaviour Models: A Survey Focusing on Autism Therapy Applications

Michał Stolarz, Hochschule Bonn-Rhein-Sieg, Sankt Augustin, Germany, michal.stolarz@smail.inf.h-brs.de
Alex Mitrevski, Hochschule Bonn-Rhein-Sieg, Sankt Augustin, Germany, aleksandar.mitrevski@h-brs.de
Mohammad Wasil, Hochschule Bonn-Rhein-Sieg, Sankt Augustin, Germany, mohammad.wasil@h-brs.de
Paul G. Plöger, Hochschule Bonn-Rhein-Sieg, Sankt Augustin, Germany, paul.ploeger@h-brs.de

Personalised Robot Behaviour Modelling for Robot-Assisted Therapy in the Context of Autism Spectrum Disorder

Michał Stolarz[†§], Alex Mitrevski[†§], Mohammad Wasil[†], and Paul G. Plöger[†]

Proceedings of the 2004 IEEE International Workshop on Robot and Human Interactive Communication
Kurashiki, Okayama Japan September 20-22, 2004

Robots We Like to Live With?! – A Developmental Perspective on a Personalized, Life-Long Robot Companion

Kerstin Dautenhahn

International Journal of Social Robotics (2021) 13:169–185
https://doi.org/10.1007/s12369-020-00629-w

Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider[1,2] · Franz Kummert[1]

Personalisation Preliminaries

What is Personalisation?

- Personalisation is the **adaptation of a robot's behaviour to the preferences and needs of individual users**
- Depending on the application, personalisation can be performed **for each individual user** or **for groups of comparable users**

Robots can be referred to as personalised if "their individuality reflects the needs and requirements of the (social) environment where the robot is operating in" (Dautenhahn, 2004)

Why Personalisation?

- Robot personalisation can be desirable for various reasons:

Why Personalisation?

- Robot personalisation can be desirable for various reasons:
  - Different users may have **different expectations from a robot** (e.g. collaboration with an expert should differ from collaboration with a novice)

Why Personalisation?

- Robot personalisation can be desirable for various reasons:
  - Different users may have **different expectations from a robot** (e.g. collaboration with an expert should differ from collaboration with a novice)
  - Individual differences between users may be **too large for a generic robot to be useful** (e.g. a therapy robot should adapt to the individual needs of the treated person)

Why Personalisation?

- Robot personalisation can be desirable for various reasons:
  - Different users may have **different expectations from a robot** (e.g. collaboration with an expert should differ from collaboration with a novice)
  - Individual differences between users may be **too large for a generic robot to be useful** (e.g. a therapy robot should adapt to the individual needs of the treated person)
- Personalisation may **increase the acceptance of robots in everyday applications** — users are likelier to use robots **over a prolonged period of time** if they can be adapted to their needs

Reconfigurability vs. Adaptation

International Journal of Social Robotics (2021) 13:169–185
https://doi.org/10.1007/s12369-020-00629-w

Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider[1,2] · Franz Kummert[1]

- In principle, there are two ways in which robots can be personalised:

Reconfigurability

The robot's behaviour can be described by **a collection of settings** that a user can manually adjust based on their needs

Adaptation

The behaviour is represented through **a learnable model** of the user's preferences that is updated throughout the robot's operation

Reconfigurability vs. Adaptation

International Journal of Social Robotics (2021) 13:169–185
https://doi.org/10.1007/s12369-020-00629-w

Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider[1,2] · Franz Kummert[1]

- In principle, there are two ways in which robots can be personalised:

Reconfigurability

The robot's behaviour can be described by **a collection of settings** that a user can manually adjust based on their needs

Adaptation

The behaviour is represented through **a learnable model** of the user's preferences that is updated throughout the robot's operation

- Both techniques have their own merits — and most devices we use daily are manually adjustable anyway — but **adaptation can reduce the burden on the user to find appropriate settings**

Reconfigurability vs. Adaptation

International Journal of Social Robotics (2021) 13:169–185
https://doi.org/10.1007/s12369-020-00629-w

Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider[1,2] · Franz Kummert[1]

- In principle, there are two ways in which robots can be personalised:

Reconfigurability

The robot's behaviour can be described by **a collection of settings** that a user can manually adjust based on their needs

Adaptation

The behaviour is represented through **a learnable model** of the user's preferences that is updated throughout the robot's operation

- Both techniques have their own merits — and most devices we use daily are manually adjustable anyway — but **adaptation can reduce the burden on the user to find appropriate settings**
  - Manual reconfigurability can be particularly difficult to achieve if the user's behaviour cannot be described by well-defined options (e.g. the manner in which a robot moves in certain cases)

Reconfigurability vs. Adaptation

International Journal of Social Robotics (2021) 13:169–185
https://doi.org/10.1007/s12369-020-00629-w

Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider[1,2] · Franz Kummert[1]

- In principle, there are two ways in which robots can be personalised:

Reconfigurability

The robot's behaviour can be described by **a collection of settings** that a user can manually adjust based on their needs

Adaptation

The behaviour is represented through **a learnable model** of the user's preferences that is updated throughout the robot's operation

- Both techniques have their own merits — and most devices we use daily are manually adjustable anyway — but **adaptation can reduce the burden on the user to find appropriate settings**
  - Manual reconfigurability can be particularly difficult to achieve if the user's behaviour cannot be described by well-defined options (e.g. the manner in which a robot moves in certain cases)
  - Reconfigurability also requires some level of system expertise, which most users may lack

Personalisable Aspects

Personalized Behaviour Models: A Survey Focusing on Autism Therapy Applications

Michał Stolarz, *Hochschule Bonn-Rhein-Sieg*, Sankt Augustin, Germany, michal.stolarz@smail.inf.h-brs.de
Alex Mitrevski, *Hochschule Bonn-Rhein-Sieg*, Sankt Augustin, Germany, aleksandar.mitrevski@h-brs.de
Mohammad Wasil, *Hochschule Bonn-Rhein-Sieg*, Sankt Augustin, Germany, mohammad.wasil@h-brs.de
Paul G. Plöger, *Hochschule Bonn-Rhein-Sieg*, Sankt Augustin, Germany, paul.ploeger@h-brs.de

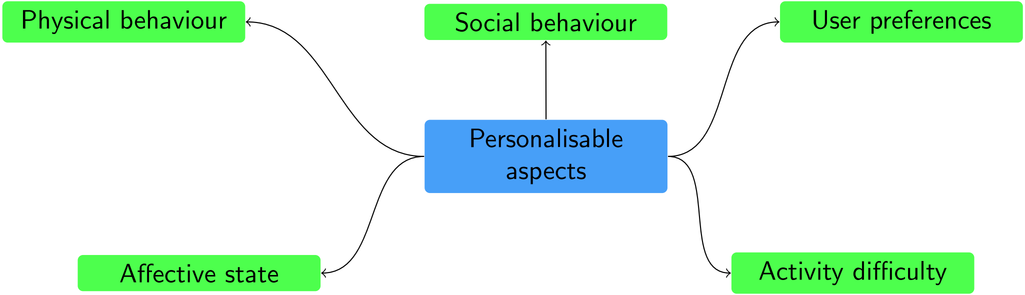

Personalisation: Safety, Ethics, and Transparency

The personalisation aspect is **closely related to robot safety, ethics, and transparency**

Personalisation: Safety, Ethics, and Transparency

The personalisation aspect is **closely related to robot safety, ethics, and transparency**

Safe adaptation

Personalisation should be done so that **the robot's behaviour complies with predefined safety expectations** (i.e. the personalisation process should not make a robot unsafe)

Personalisation: Safety, Ethics, and Transparency

The personalisation aspect is **closely related to robot safety, ethics, and transparency**

Safe adaptation

Personalisation should be done so that **the robot's behaviour complies with predefined safety expectations** (i.e. the personalisation process should not make a robot unsafe)

Preventing unethical adaptation

Personalisation **must be performed in accordance with ethical standards** (e.g. it should not be possible to personalise a robot to perform unethical actions)

Personalisation: Safety, Ethics, and Transparency

The personalisation aspect is **closely related to robot safety, ethics, and transparency**

Safe adaptation

Personalisation should be done so that **the robot's behaviour complies with predefined safety expectations** (i.e. the personalisation process should not make a robot unsafe)

Preventing unethical adaptation

Personalisation **must be performed in accordance with ethical standards** (e.g. it should not be possible to personalise a robot to perform unethical actions)

Transparent adaptation

Users should be able to **comprehend the aspects that a robot has personalised** (particularly if the personalised behaviour does not correspond to their expectations)

Learning for Personalisation

Types of Learning for Personalisation

- Learning for personalisation is typically performed with some form of **interactive machine learning (IML)**[1]
  - IML is a user-in-the-loop learning paradigm, where a robot learns through interaction with a user

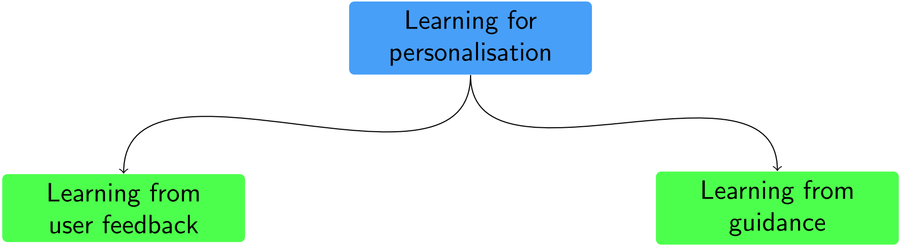

- We can particularly distinguish between IML methods that depend on who provides feedback — the robot user itself (learning from feedback) or a supervisor (learning from guidance)

[1] M. Chetouani, "Interactive Robot Learning: An Overview,", *ECCAI Advanced Course on Artificial Intelligence*, pp. 140–172, 2021.

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**
- The policy $\pi : S \rightarrow A$ is then **a model of which action $\boldsymbol{a} \in A$ to select based on the user's state $\boldsymbol{s} \in S$**

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**
- The policy $\pi : S \rightarrow A$ is then **a model of which action $a \in A$ to select based on the user's state $s \in S$**
- Generally speaking, learning from feedback is challenging because it requires **a large number of interactions with a user** — S needs to be explored by the robot as well as possible

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**
- The policy $\pi : S \rightarrow A$ is then **a model of which action $a \in A$ to select based on the user's state $s \in S$**
- Generally speaking, learning from feedback is challenging because it requires **a large number of interactions with a user** — S needs to be explored by the robot as well as possible

User reaction

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**
- The policy $\pi : S \rightarrow A$ is then **a model of which action $a \in A$ to select based on the user's state $s \in S$**
- Generally speaking, learning from feedback is challenging because it requires **a large number of interactions with a user** — S needs to be explored by the robot as well as possible

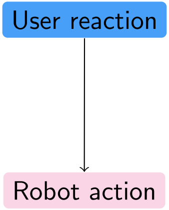

Learning From Feedback

- In learning from feedback, **a robot learns a behaviour policy based on interactions with and observations of a user**
- The policy $\pi : S \rightarrow A$ is then **a model of which action $a \in A$ to select based on the user's state $s \in S$**
- Generally speaking, learning from feedback is challenging because it requires **a large number of interactions with a user** — S needs to be explored by the robot as well as possible

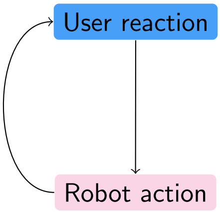

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**

- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**
- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**
- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**
- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**
- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]
- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

User reaction

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**
- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**
- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]
- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**

- In this case, the objective of the policy $\pi : S \to A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**

- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]

- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

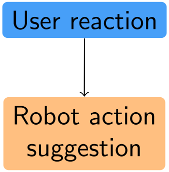

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**

- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**

- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]

- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

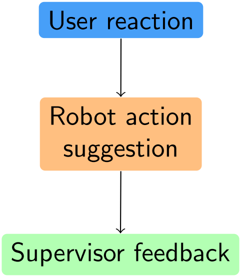

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**
- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**
- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]
- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

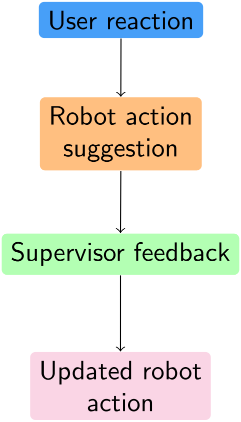

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

Guidance-Based Learning

- In learning from guidance, the user also interacts with a user, but **the selection of actions is guided by a supervisor**
- In this case, the objective of the policy $\pi : S \rightarrow A$ is to **select an action $a \in A$ that would correspond to how the supervisor would act in a state $s \in S$**
- Learning from guidance **reduces the learning data requirements** — supervisor feedback serves as **a ground-truth action label**
  - But the possibility that a supervisor makes a mistake should be taken into account as well[2]
- Particularly in domains such as robot-assisted therapy, learning from guidance is more suitable than learning from feedback, as **robot mistakes can have more serious consequences**

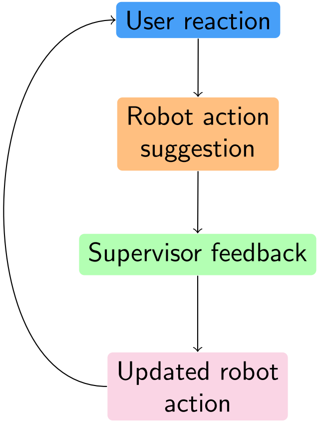

[2] For instance by learning a model of the feedback as in W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," in *Proc. 5th Int. Conf. Knowledge capture*, 2009, pp. 9–16.

User Modelling

- Personalisation is typically associated with **user models**

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture
- The main use of user models is to **simulate user interactions**
  - Just as learning from guidance, user models can reduce the real interaction data requirements

User Modelling

Personalised Robot Behaviour Modelling for Robot-Assisted Therapy in the Context of Autism Spectrum Disorder

Michał Stolarz[†§], Alex Mitrevski[†§], Mohammad Wasil[†], and Paul G. Plöger[†]

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture
- The main use of user models is to **simulate user interactions**
  - Just as learning from guidance, user models can reduce the real interaction data requirements

User reaction

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture
- The main use of user models is to **simulate user interactions**
  - Just as learning from guidance, user models can reduce the real interaction data requirements

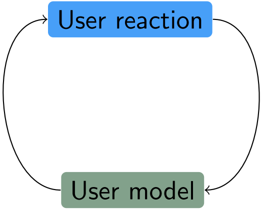

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture
- The main use of user models is to **simulate user interactions**
  - Just as learning from guidance, user models can reduce the real interaction data requirements

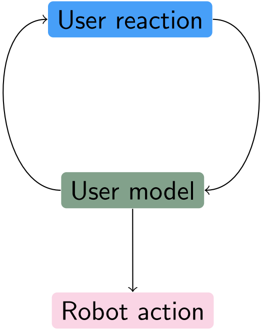

User Modelling

- Personalisation is typically associated with **user models**
- A user model **is a predictive model of certain aspects of the user's behaviour** (e.g. given a task and the user's affective state, the model can be used to predict whether the user will succeed at the task)
  - Simplified user models may be defined manually, but learning is typically a more general approach
- The aspects that should be captured by a user model **depend on the personalisation objective**
  - A robot for social interaction may need to attend to the user's affective state, while a collaborative mobile manipulator may need to consider the user's posture
- The main use of user models is to **simulate user interactions**
  - Just as learning from guidance, user models can reduce the real interaction data requirements

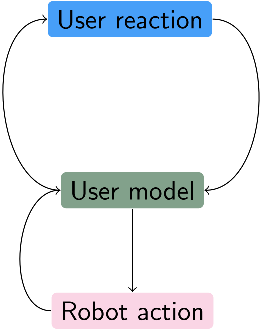

User Grouping

- In some cases, it may be sensible to **consider groups of related users instead of individual users** to learn personalised policies
  - Users can be grouped into categories based on particular criteria, such as similar preferences

User Grouping

- In some cases, it may be sensible to **consider groups of related users instead of individual users** to learn personalised policies
  - Users can be grouped into categories based on particular criteria, such as similar preferences
- User grouping is also a measure that can reduce the data requirements for personalisation — instead of collecting data only from a given user, the data from all users in a group are aggregated into a **group-specific user model**

User Grouping

- In some cases, it may be sensible to **consider groups of related users instead of individual users** to learn personalised policies
  - Users can be grouped into categories based on particular criteria, such as similar preferences
- User grouping is also a measure that can reduce the data requirements for personalisation — instead of collecting data only from a given user, the data from all users in a group are aggregated into a **group-specific user model**
- Categorising users into groups is usually **a task for unsupervised learning**, as the groups and group assignments are unknown in advance
  - To generalise a policy to a new user, the user should be categorised into one of the existing groups
  - If applied without care, group categorisation may lead to inappropriate user stereotyping!

Personalised Policies

- Users add a complexity layer to the overall policy learning problem that we have been considering in the course so far
- Let U represent a space of users; we can then define **a context- and user-specific policy as**

$$\pi : S \times C \times U \rightarrow A$$

  This is an extension of the definition of a policy for a task family
- In principle, we could observe users as part of the context C, but the dependence on the user is best treated separately, particularly due to the user association with user models

Example: Personalisation for Robot-Assisted Therapy

The MigrAVE System

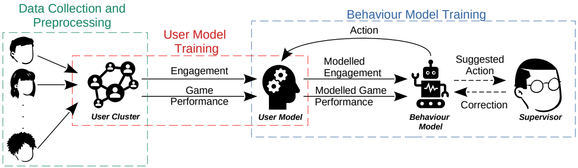

- In our personalisation system developed for the MigrAVE project, we combined various of the aspects discussed on the previous slides
- The objective here was to enable a robot to personalise (i) the difficulty of activities, e.g. games, and (ii) the type of feedback provided to the user (encouraging, challenging, or no feedback)
- User models that estimate the user's engagement and performance in the activity were learned for multiple user groups
- We used the user model to simulate learning from feedback, but also used learning from guidance
  - Learning from guidance is particularly suitable in this case, as we need therapists to have control of the robot's behaviour

Robot Personality Modelling

- When considering social robot interaction, the main aspect that needs to be personalised is **the robot's personality** — but what is a robot's personality?

[3] Details about the model can be found in R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *Journal of personality* vol. 60, no. 2, pp. 175–215, 1992.

[4] H. -L. Cao et al., "A Personalized and Platform-Independent Behavior Control System for Social Robots in Therapy: Development and Applications," in *IEEE Trans. Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 334–346, Sept. 2019

Robot Personality Modelling

- When considering social robot interaction, the main aspect that needs to be personalised is **the robot's personality** — but what is a robot's personality?
- Robot personality is generally modelled **based on models about human personality**

[3] Details about the model can be found in R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *Journal of personality* vol. 60, no. 2, pp. 175–215, 1992.

[4] H. -L. Cao et al., "A Personalized and Platform-Independent Behavior Control System for Social Robots in Therapy: Development and Applications," in *IEEE Trans. Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 334–346, Sept. 2019

Robot Personality Modelling

- When considering social robot interaction, the main aspect that needs to be personalised is **the robot's personality** — but what is a robot's personality?
- Robot personality is generally modelled **based on models about human personality**
- One such model is the **five-factor model (Big Five)**[3], according to which there are five main personality traits that can be observed: agreeableness, conscientiousness, extraversion, openness to experience, and neuroticism
  - For each individual, these are present on a scale

[3] Details about the model can be found in R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *Journal of personality* vol. 60, no. 2, pp. 175–215, 1992.

[4] H. -L. Cao et al., "A Personalized and Platform-Independent Behavior Control System for Social Robots in Therapy: Development and Applications," in *IEEE Trans. Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 334–346, Sept. 2019

Robot Personality Modelling

- When considering social robot interaction, the main aspect that needs to be personalised is **the robot's personality** — but what is a robot's personality?
- Robot personality is generally modelled **based on models about human personality**
- One such model is the **five-factor model (Big Five)**[3], according to which there are five main personality traits that can be observed: agreeableness, conscientiousness, extraversion, openness to experience, and neuroticism
  - For each individual, these are present on a scale
- There are attempts to impose these personality types on a robot through the **affective state**[4]
  - The affective state is, in turn, often modelled using the valence-arousal-dominance model, where events can result in different values for these variables

[3] Details about the model can be found in R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *Journal of personality* vol. 60, no. 2, pp. 175–215, 1992.

[4] H. -L. Cao et al., "A Personalized and Platform-Independent Behavior Control System for Social Robots in Therapy: Development and Applications," in *IEEE Trans. Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 334–346, Sept. 2019

Personalisation Over a Long Term

- **Personalisation cannot be observed as a one-time learning process** — people's skills and preferences generally evolve over time, so **personalised policies should ideally be adapted over time as well**

Personalisation Over a Long Term

- **Personalisation cannot be observed as a one-time learning process** — people's skills and preferences generally evolve over time, so **personalised policies should ideally be adapted over time as well**
- This long-term aspect poses challenges on various levels:

Personalisation Over a Long Term

- **Personalisation cannot be observed as a one-time learning process** — people's skills and preferences generally evolve over time, so **personalised policies should ideally be adapted over time as well**
- This long-term aspect poses challenges on various levels:
  - **Learning**: The policy representation should have a degree of plasticity so that it can combine new with previously known information

Personalisation Over a Long Term

- **Personalisation cannot be observed as a one-time learning process** — people's skills and preferences generally evolve over time, so **personalised policies should ideally be adapted over time as well**
- This long-term aspect poses challenges on various levels:
  - **Learning**: The policy representation should have a degree of plasticity so that it can combine new with previously known information
  - **Privacy**: Long-term personalisation requires data from a user to be collected potentially continuously — this needs to be done without violating the user's privacy

Personalisation Over a Long Term

- **Personalisation cannot be observed as a one-time learning process** — people's skills and preferences generally evolve over time, so **personalised policies should ideally be adapted over time as well**
- This long-term aspect poses challenges on various levels:
  - **Learning**: The policy representation should have a degree of plasticity so that it can combine new with previously known information
  - **Privacy**: Long-term personalisation requires data from a user to be collected potentially continuously — this needs to be done without violating the user's privacy
- In a sense, **long-term personalisation is the holy grail of service robotics** — service robots can only truly be useful once this problem is solved satisfactorily

Summary

- Personalisation is the adaptation of a robot's behaviour to the needs and preferences of individual users
- Various aspects of a robot's behaviour can be personalised, such as its physical or social behaviour, its affective state, the difficulty of activities that it performs with people, as well as general configuration preferences
- Personalised robot behaviour policies are typically learned using interactive machine learning, using learning from feedback or learning from guidance
- User models are often learned as an intermediate representation of users that enables interaction experiences to be simulated; users may also be categorised into groups of similar users so that personalisation is applied over a group
- Personalisation should be seen as a long-term learning aspect to address changing needs and preferences of users