Observing and Clustering Coaching Behaviours to Inform the Design of a Personalised Robotic Coach

Ross, Martin; Broz, Frank; Baillie, Lynne

DOI
10.1145/3447526.3472043

Publication date
2021

Document Version
Final published version

Published in
Proceedings of MobileHCI 2021 - ACM International Conference on Mobile Human-Computer Interaction

Observing and Clustering Coaching Behaviours to Inform the Design of a Personalised Robotic Coach

Martin K. Ross
Heriot-Watt University, Edinburgh, UK
mkr30@hw.ac.uk

Frank Broz
Delft University of Technology, Rotterdam, Netherlands
f.broz@tudelft.nl

Lynne Baillie
Heriot-Watt University, Edinburgh, UK
l.baillie@hw.ac.uk

ABSTRACT

Adherence to repetitive rehabilitation exercises is important in motor recovery after stroke. Similarly, repetitive solo practice exercises can improve the skill level of sports players. In both of these scenarios, regular human coaching has benefits, but in practice, the required training is often carried out alone, resulting in lowered adherence. This work presents a mixed methodology approach, novel in the context of designing for HRI, towards informing the design of a personalised robotic coach for stroke rehabilitation and squash. Using observations of human-human interactions, we first obtained action sequences of behaviours exhibited by coaches and physiotherapists. We then clustered these action sequences into behaviour graphs, with each graph representing a coaching policy usable for robotic control. Next we obtained coaches' and physiotherapists' reflections on the graphs' applicability to the real world. Finally, we provide an explanation of how the policies visualised in these graphs could be used for robotic control.

CCS CONCEPTS

• **Human-centered computing** → Interaction design; Interaction design process and methods; User centered design; • **Computer systems organization** → Embedded and cyber-physical systems; Robotics;

KEYWORDS

Systematic observations, Stroke, Coaching, HCI, Human Robot Interaction (HRI)

ACM Reference Format:
Martin K. Ross, Frank Broz, and Lynne Baillie. 2021. Observing and Clustering Coaching Behaviours to Inform the Design of a Personalised Robotic Coach. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction (MobileHCI '21), September 27–October 01, 2021, Toulouse & Virtual, France.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3447526.3472043

MobileHCI '21, September 27–October 01, 2021, Toulouse & Virtual, France

ACM ISBN 978-1-4503-8328-8/21/09…$15.00
https://doi.org/10.1145/3447526.3472043

1 INTRODUCTION

Stroke is one of the leading causes of acquired adult disability with survivors commonly suffering permanent impairments such as fatigue, weakness in the arms and legs, aphasia and forgetfulness [1]. Although research in rehabilitation techniques strongly suggest that home-based rehabilitation (i.e. without the supervision of a rehabilitation therapist) is beneficial to the patient [2], it is often not adhered to due to (among other reasons) a lack of motivation [3]. Squash, on the other hand, is an intermittent, high-intensity racket sport in which repetitive, solo drills are used frequently by many of the top professionals either with or without the presence of a coach. We have chosen to consider both of these case studies, from different domains, in the same body of work due to the similarities in the individual, often unsupervised and repetitive nature of practice which helps in making long term functional improvements after stroke [2] and helps high performance sports players improve their skill level [4]. Additionally, in both rehabilitation and sports coaching the relationship built between the physiotherapist/coach and their client is an important part of making functional improvements. The intrinsic motivation of the client (which can lead to increased adherence to exercise [5]) can be affected by the actions of the physiotherapist or coach [6], [7]. Therefore there is significant overlap between these two cases providing evidence of the potential to use a similar HCI system to good effect in both areas.

This paper presents a mixed methodology approach towards achieving an understanding of the behaviours used by stroke rehabilitation physiotherapists and squash coaches, with the goal of informing the design of a personalised robotic coach which could provide motivation in solo squash practice and rehabilitation after stroke. Our approach combines data collection methods adapted from sports coaching literature with computational techniques and mathematical modelling, which is a unique way of approaching designing for HCI. The methodology we used consisted of systematically observing practicing physiotherapists and professional squash coaches, clustering the obtained data into behaviour graphs (each graph is a visual representation of a policy usable for robotic control), and finally obtaining professionals' reflections on the applicability of these graphs through semi-structured interviews. A policy, as commonly defined in machine learning literature (e.g. [8]) expresses an agent's way of behaving at a given time. It is a mapping from perceived states of the environment to actions to be taken when in those states.

We believe the results presented in this paper contribute to the MobileHCI community by providing a foundation for designing a data-driven Human-Robot Interaction (HRI) system to assist professionals in these areas. The envisioned personalised robotic coach, for example, capable of increasing motivation and adherence to individual exercise has potential benefits in both squash and rehabilitation after stroke. Moreover, the unique methodology presented could be applied to any number of other cases to inform the design of HCI systems. In particular the contribution of behaviour graphs as a method of representing a policy in a usable and readable way to stakeholders could be replicated by other researchers in designing HCI systems. The aims of these studies were to:

A1. Gather and analyse data about the behaviours of stroke physiotherapists and squash coaches to discover the behaviours which would be required in a robotic coach capable of similar behaviour.
A2. Discover differences and commonalities in the behaviours of physiotherapists and squash coaches to predict whether a robotic coach implementing similar behavioural policies could be used in both areas.
A3. Obtain physiotherapists' and squash coaches' reflections on the behaviour graphs generated by the clustering algorithm.
A4. Discover how physiotherapists and squash coaches personalise their coaching approach and gauge which factors would be most useful to take into account when deciding on a coaching policy to be used by our robotic system.

2 BACKGROUND

2.1 Robotics in Sport and Rehabilitation

More than half of the 80 million stroke survivors worldwide [9] suffer permanent impairments [1]. Survivors frequently feel abandoned after leaving hospital because the necessary follow-up assessments and support for them and their family are often not delivered [10]. A variety of HCI devices have the potential to benefit stroke survivors including haptic devices [11] and exergames [12]. However robots have perhaps the biggest potential to fill this gap and have already been considered as rehabilitation coaches for stroke survivors [13], [14]. For example Wade, Parnandi and Matarić showed that it is possible for a fully autonomous robotic coach to lead a stroke survivor through a rehabilitation session, without the presence of a trained therapist [13]. Feingold Polak & Levy Tzedek [14] conducted a medium-term study involving 4 stroke patients over 15 interaction sessions (2-3 sessions per week over 5-7 weeks), in which the gamified robotic system was rated very highly by all participants. Both of these studies show the potential of using a system of this kind to engage stroke survivors in rehabilitation, but the area requires further investigation. Both studies were conducted in a lab or rehabilitation facility and the limited number of participants in Feingold Polak & Levy Tzedek's study, combined the short-term nature of Wade's study (each participant interacted with the system for 2 sessions of 10 minutes each), means the use of robots for long-term rehabilitation and for use in an unsupervised setting e.g. the home, remains uninvestigated.

In a sports coaching context, ongoing work is exploring the use of a robot to coach users through the couch to 5km running program [15]. Using a method of a domain expert manually correcting the behaviour of the robotic coach during sessions, it has been shown that a robot can learn to replicate this behaviour when acting autonomously. The user experience of this system is yet to be explored, but participants used it for an average of 15.4 hours over 3 months, suggesting that a system like this could be used in long-term recovery after stroke and to increase adherence and engagement in repetitive solo practice sessions in sport. However, all of these systems lack the personalisation aspect which is an aim of the current work and is becoming a requirement of such a robotic coaching system [14], [16], [17].

By conducting focus groups and interviews with physiotherapists, occupational therapists, speech therapists and sports rehabilitation therapists, Winkle et al. [17] presented a set of design implications/guidelines which focused on using a robot as a practical tool and interaction partner for facilitating self-practice exercise sessions for adults in their own homes. This included a list of key patient traits, a subset of which we presented to coaches and physiotherapists during semi-structured interviews in the second study of the current work (see Section 6.1.2).

2.2 Data Based Approaches to Personalisation

Personalisation has become a key factor in many HRI systems, with robotic systems employing strategies intended to build the relationship between the user and the robot (e.g. by using continuity behaviours, and the user's name) being preferred by users over systems that are purely functional [18], [19].

A number of different methods have been explored to achieve personalisation of a robot's actions. Using Reinforcement Learning (RL) to personalise the teaching behaviours of a robot has been shown to result in higher levels of positive valence towards the robot [20] and more effective teaching [21]. This is a promising method of low-level adaption to individuals within sessions. However, the current work focusses on strategies for high-level personalisation to groups of users.

Other techniques focus on learning sets of policies from task demonstrations. For example, Chen et al. proposed Multi-Style Reward Distillation (MSRD) [22] which takes a group of demonstrations and learns methods to achieve the same outcome using different strategies. This is similar to the role of a coach, the difference being that the task reward was clear and fixed in Chen et al.'s work: to return a table tennis ball. In the domain of coaching, the task reward could be any number of things depending on the type of session and the user involved.

Nikolaidis et al. [23] proposed a method of clustering human demonstrations into similar styles and applying inverse reinforcement learning over these clusters. They showed that, using their method, it was possible to learn a reward function that was representative of each user type in a collaborative packing task. Yin et al. [24] also partitions demonstrations into different styles and employs inverse reinforcement learning, but uses the learned ensemble as the mode observation model, which is user-specific in Nikolaidis' work. The clustering method presented by Nikolaidis is used in this study and described in Section 5.

2.3 Techniques for Data Collection

Systematic observations are a validated and well-used approach in sports science to gain a better understanding of the behaviours used by coaches [25]. Previous research using this technique in the coaching domain has primarily focused on team sports, finding that expert coaches will commonly exhibit a large number of instructional behaviours [26], constructing their practice sessions in a manner that maximises the transfer of information to their players [27]. However, coaches also attempt to create a social bond between themselves and their players using the coaching behaviours they employ [28]. The relationship formed between the two parties is likely to be even more important in an individual sport, as evidenced by a workshop study concluding that personal trainers favoured technology which would enhance the relationship between themselves and their clients [29].

Systematic observations using similar instruments have also been used during stroke rehabilitation sessions. For example, Van Vilet et al. used this method to compare the differences in two different approaches (Bobath-based and movement-science based) to physiotherapy after stroke [30]. This study put heavy emphasis on the physical activities used in sessions over the communication behaviours exhibited by the physiotherapists. This means that the behavioural data is unfortunately not detailed enough on which to base the design of a robotic coach or any other HCI system. A similar problem can be seen in the field of sports coaching. For example, Sussenbach et al. created a motivational model based on HHI observations which was implemented in a robotic cycling instructor [31], but the observations focused more on the structure of the interaction than on the behavioural action sequences which would be needed to create a data-driven system. However, these study do show the applicability of using systematic observations in the domains of sport and rehabilitation.

The use of additional qualitative methods in conjunction with quantitative data collection has become favoured in a variety of research contexts [7], [25], [32]. In the field of sports coaching, many systematic observation studies (14 of the studies found in Cope et al.'s review [25] of systematic observations of sports coaches between 1997 and 2016), have also used interviews to gain an understanding of why coaches use the behaviours they do [33]. However, as demonstrated by Harvey et al. [34], if coaches are not reflective about their practices in any meaningful way, there can be discrepancy between what they were observed doing, and their understanding and explanations of why they used the behaviours they did. In rehabilitation after stroke, semi-structured interviews with 32 professionals (including 10 physiotherapists) indicated that patient motivation can be affected by the behaviours of the health professionals who interact with them [7].

Despite the previous research in these areas, no studies that the authors are aware of have used semi-structured interviews to explore the behaviours of coaches in individual sports, or to explore with domain professionals the application of robotics in the context of sports coaching. Qualitative methods such as workshops, focus groups and interviews are frequently used during the design process of HCI systems (e.g. [17], [35]) but typically focus more on idea generation and high-level interactions. Likewise, the authors are not aware of any systematic observation studies which analyse the behaviours of stroke physiotherapists to an extent which would provide usable data in the context of designing for HCI/HRI.

3 OVERVIEW OF STUDIES

Figure 1 shows an overview of the process used in this body of work and the following sections describe in detail the rigorous and well defined procedure for mapping observations of HHI to an HRI system capable of similar behaviour. To discover how human practitioners in both stroke rehabilitation and squash deliver their sessions, we first conducted systematic observations of one-to-one sessions in both domains (detailed in Section 4). This provided us with action sequences of coaching behaviours which we clustered into 12 different coaching policies (see Section 5) each representing a different style of coaching which could be implemented in a robotic coach. However we did not know at this stage when it would be most appropriate to apply each policy. Therefore, we went back to the coaches and physiotherapists and presented to them the coaching policies, visualised as behaviour graphs (see Section 5.1). By doing this during online semi-structured interviews, we were able to deduce the situations and types of users each of the graphs were likely to be most applicable to. The semi-structured interview process is described in Section 6.

4 STUDY 1 – OBSERVATIONS OF PROFESSIONAL SQUASH COACHES AND STROKE PHYSIOTHERAPISTS

4.1 Observation Instrument

A variety of observation instruments have been used in sports coaching observation studies, the most popular of which (between 1997 and 2016 [25]) has been the Arizona State University Observation Instrument (ASUOI) [36]. The ASUOI was developed to create a sensitive tool, capable of collecting highly specific data on coaching behaviours, expanding on and modifying several behavioural categories in other observation instruments available at the time [36]. It is a pen and paper based instrument which allows the recording of 13 distinct behavioural categories exhibited by the observed coach. Some researchers have chosen to modify the original, validated ASUOI to better suit the needs of their study [37], [38] - an approach also taken in the current work. Using a more complex, computer-based instrument (such as the CAIS [39]) would not have allowed for easy adaption and would have been overly complex for one-to-one interactions.

The instrument contributed by this work is adapted from the validated and widely used ASUOI. The behavioural categories of "management" (used more often in group coaching), "silence" (only applicable when using interval recording, see Section 4.1.1) and "other" (an additional box for field notes was provided instead) were removed. "Positive reinforcement" and "punishment" were added as new behavioural categories. Research has shown that both are techniques which could be used to strengthen a particular action by providing a consequence that an individual finds either rewarding or penalising [40]. "Console" was also added as a new behavioural category during the coder training process described in Section 4.1.2. Finally, the "concurrent instruction" and "post-instruction" categories were split into positive and negative versions as is the case in other, more recent observation instruments (e.g. the CAIS [39]). The final instrument (digital recreation shown in Figure 2) contains 16 behavioural categories (see Table 1) which can be recorded in each observation session, with field notes indicating any additional behaviours displayed by the coach.

4.1.1 Recording Behaviours Using the Instrument. Figure 2 shows an example of the observation instrument (digitally recreated for

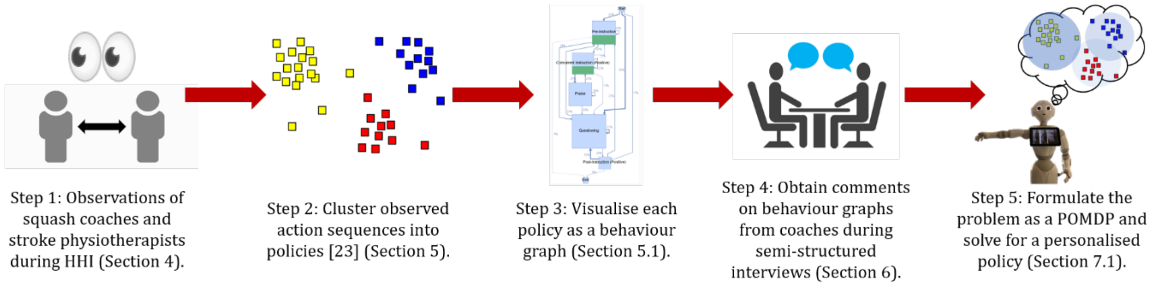

Figure 1: : The process used in this body of work beginning with observations of HHI through to using the obtained data to personalise the behaviour of a robotic coach for squash or stroke rehabilitation.

legibility) after it had been completed using pen and paper during one of the sessions. Traditionally, either interval recording (writing down a number corresponding to the behaviour being exhibited by the coach every 3-5 seconds) or event recording (placing a tally in the corresponding column on the observation instrument each time the coach exhibits a behaviour) is used during systematic observations with the ASUOI [36]. For the current study, an adapted version of event recording which involved writing a number in sequential order in the corresponding row of the observation instrument was used. If multiple behaviours occurred concurrently they were noted after a decimal point in the original number. For example, if the order of behaviours observed was: questioning, followed by pre-instruction and positive modelling concurrently, then praise; the observer would record "1" in the questioning row, "2" in the pre-instruction row, "2.1" in the positive modelling row, and finally "3" in the praise row. The original version of event recording [36] only gave the frequency of each behaviour and not the sequential order, while interval recording only allows for one behaviour to be recorded every 3-5 seconds, and in practice behaviours can happen more frequently than this. The adapted method used in the current work combined the benefits of interval and event recording by producing the total number of each behaviour observed (including behaviours which were used concurrently) and the order of the observed behaviours.

4.1.2 Coder Training. Before a study such as this is conducted, it is recommended that the coder receive training on the chosen observation instrument [25]. Since we used an adapted observation instrument for this study, we followed a similar process to that used in [41]. During this process, the observer (lead author in our case) familiarised themself with the ASUOI through careful reading of the relevant literature. Next, we validated the developed instrument with a top-level squash coach (currently coaching a national team) and an experienced coder. Then our coder practiced with the developed instrument, under the guidance of the experienced coder, using both video footage and live sessions of professional squash coaches, with gaps of 24 hours, 7 days, and 14 days to allow for memory lapse [42]. A pilot study was undertaken consisting of observation sessions with three squash coaches. Following this and discussion with the experienced coder, the instrument as presented in the current work was finalised.

4.2 Method

4.2.1 Participants. Côté and Gilbert define coaching effectiveness as: "The consistent application of integrated professional, interpersonal, and intrapersonal knowledge to improve athletes' competence, confidence, connection, and character in specific coaching contexts" [43]. They add that coaches who demonstrate coaching effectiveness over an extended period of time may then be considered expert coaches. Given this definition, the inclusion criteria used in this study were:

1.1 Qualification: minimum of level 2 coaching certification for squash coaches or BSc physiotherapy or related subject for physiotherapists.
OR
1.2 Experience: minimum of 5 years coaching/physiotherapy experience.
2. Currently coach squash or administer stroke rehabilitation on at least a weekly basis, and have done for at least the last year.
3. Squash coaches only: have worked with both senior and junior players, and either international or developmental players, in the last year.

For this observation study, we recruited 10 practicing stroke physiotherapists and 8 professional squash coaches (demographic information shown in Table 2). The physiotherapists were recruited through local rehabilitation centres and physiotherapy companies, while the squash coaches were recruited through personal contacts of the researcher, with the help of the sport's National Governing Body and through advertisement at a coach development session. Stroke survivors (N=18) and squash players (N=15) were also involved in the study, although they were not the focus of any observations. For this reason, it was deemed most appropriate to restrict the information provided to the therapists and coaches, rather than their clients. Full ethical approval for all observations was obtained from our university and participants received no remuneration for their time.

4.2.2 Procedure. Participants were asked to carry out a physiotherapy/squash coaching session as normal while they were observed by the researcher. Each participant (with one exception) was observed during two separate sessions. This is in line with recommendations in the sports coaching literature that a single observation session of a coach cannot be deemed an example of how a coach behaves due to the contextual and situational nature of coaching and should

Table 1: Observation instrument behaviours and definitions. All definitions have been taken from [31] and adapted slightly for use in stroke rehabilitation sessions unless otherwise stated.

Behaviour	Definition	Examples Squash Coaching	Stroke Physiotherapy
Pre-instruction	Initial information given to the player or patient preceding the desired action to be executed, explaining the execution of an exercise or technique.	"Keep an open racket face so the ball stays above the tin."	"Hold the putty in your fingertips and roll it into a ball."
Concurrent Instruction (Positive)[1]	Cues or reminders given during the actual execution of the skill, play or exercise, framed in a positive or supportive way.	"Racket up"	"Keep your body still"
Concurrent Instruction (Negative)	Cues or reminders given during the actual execution of the skill, play or exercise, framed in a negative or unsupportive way.	"Don't let the racket drop"	"Don't let your body rotate"
Post-instruction (Positive)	Correction, re-explanation, or instructional feedback given after the actual execution of the skill, play or exercise, framed in a positive or supportive way.	"I liked the extension of the follow through on that last shot."	"Excellent stable body position."
Post-instruction (Negative)	Correction, re-explanation, or instructional feedback given after the actual execution of the skill, play or exercise, framed in a negative or unsupportive way.	"Don't let your follow through come so far round your body."	"You didn't keep your body straight there."
Manual Manipulation	Physically moving the player or patient's body to a proper position or through a correct range of motion.	Positioning the player's racket preparation.	Guiding the patient's arm through an exercise.
Questioning	A question to the player or patient concerning goals, strategies, techniques etc. associated with the activity.	"What is the proper grip on the forehand?"	"Where should your left arm be for this exercise?"
Positive Modelling	A demonstration of correct performance of a skill, playing technique or exercise.	Correctly executing a drop shot.	Rolling a piece of putty into a ball.
Negative Modelling	A demonstration of incorrect performance of a skill, playing technique, or exercise.	Hitting a shot into the tin with a closed racket face.	Using the other hand to position the putty.
First Name	Using the first name or nickname when speaking directly to the player or patient.	"Nice shot, Tank!"	"That's good Sarah."
Hustle	Verbal statements intended to intensify the efforts of the player or patient or affect their desire to improve.	"Be quick, be quick"	"Challenge yourself with the range of movement"
Praise	Verbal or nonverbal compliments, statements, or signs of acceptance.	"Nice going"	Smiles and pats on the back.
Scold	Verbal or nonverbal behaviours of displeasure.	Kicking the ground	"That was a bad effort"
Console[2]	A verbal statement intended to acknowledge an undesirable outcome, without scolding.	"Unlucky"	"It's alright, don't worry"
Positive Reinforcement2	A physical reward given when a player or patient does something considered good by the coach/therapist.	Rewarding a win with a 2 point lead next time.	Giving the patient a chocolate for doing well.
Punishment2	The infliction of a physical penalty as retribution for the player or patient's bad behaviour or poor performance.	Making a player do 10 court sprints for losing a conditioned game.	Making a patient do the exercise again because it was not done well.

1. In the original ASUOI [36] there were only three instructional categories. However, after consultation with an experienced coder, it was decided to split the definitions of concurrent and post-instruction to differentiate between positive and negative behaviours.
2. The final 3 behaviours were not included in the original ASUOI. Console was added after an initial pilot study squash session because it was observed a number of times throughout as a way of acknowledging the good effort of the player without the desired outcome. Positive Reinforcement and Punishment were added due to their links to providing motivation for individuals – which will be a goal of the envisioned robotic system.

be avoided [25]. Equally, if limited time is spent observing a coach, that person may act or behave in certain ways to satisfy the observation period [44]. One physiotherapist was only observed once as the COVID-19 pandemic posed unexpected restrictions to the completion of this study (only the last session with one physiotherapist was affected). This participant's data is included in all analysis because the length of time the participant was observed for (56:30 mins) was greater than 4 of the other physiotherapists who were

Coach ID C4.2

Categories	**Seg:** Technical **Time:** 00:00 – 17:30	
Pre-Instruction	2 6 7 54 57	5
Concurrent Instruction (Positive)	3 5 8 12 14 18 19 24 27 28 29 30 33 41 45 47 53 55 58 59 62 69 71 73 78 80 82 84 87 91 93 94 95 97	34
Concurrent Instruction (Negative)	13 43 56	3
Post Instruction (Positive)	9 11 20 22 32 35 40 46 49 66 74 76 89 101	14
Post Instruction (Negative)	10 21 50 75	4
Manual Manipulation	22.1	1
Questioning	1 15 16 66	4
Positive Modelling	2.1 2.3 2.4 2.5 7.1 9.1 11.1 20.1 22.1 32.1 32.2 35.1 54.1 57.1 76.1	15
Negative Modelling	2.2 10.1 21.1 50.1 75.1	5
First Name		
Hustle	4	1
Praise	17 23 25 26 31 34 36 37 38 39 42 44 48 51 52 60 61 63 64 65 67 68 70 72 77 79 81 83 85 86 88 90 92 96 98 99 100	37
Scold		
Console		
Positive Reinforcement		
Punishment		

Additional Comments/Observed Behaviours:
Discussion at start about what going to do rather than coach telling.
Post instr: pos -> neg -> pos pattern
Lots of silence
Some players give opinion without coach having to ask Q's

Figure 2: The observation instrument (digitally recreated for legibility) after it has been used for a session. The numbers in the table represent the order in which behaviours were used by the observed coach, and the numbers on the right are a count of the total uses of the corresponding behaviour.

Table 2: Demographic information of coaches and physiotherapists observed in Study 1.

	Squash Coaches	**Stroke Physiotherapists**
Participants (N)	8	10
Gender	6M, 1F, 1 preferred not to say	2M, 8F
Age Range	25-63 (mean = 41 ± 12.3)	28-53 (mean = 41.6 ± 8.4)
Qualification	3 level 2, 3 level 3, 2 level 4 coaching qualifications from the sport's National Governing Body	8 BSc Physiotherapy, 1 MSc Physiotherapy (Pre-Reg), 1 Chartered Physiotherapist
Experience	10-30 years (mean = 18.4 ± 6.5)	4.5 – 31 years (mean = 17.3 ± 8.6)

observed for 2 shorter sessions, meaning sufficient time was spent with that participant.

There was no upper or lower limit on length of session because it is important to observe the session just as it would normally be, without imposing any restrictions on the participants. All sessions were one-to-one physiotherapy or coaching sessions (i.e. the practitioner interacting with one client, not a group), with one exception. Individual sessions were chosen because the robotic coach to be developed is intended to act as a coach for one person at a time. The one exception was a session in which the physiotherapist alternated between attending to two patients, only directing minimal instructions to both stroke survivors. This data is also included in the analysis because it was still deemed that the physiotherapist was able to give specific feedback at all necessary times to each individual, which would not be possible in a group session with more than two participants.

The main method for gathering data was using the Observation Instrument (Figure 2) described in Section 4.1 to code the behaviours exhibited by the domain professionals. The coding was performed live during the session for all but the first 5 observed squash sessions. These 5 sessions were both video and audio recorded for coder training purposes (see Section 4.1.2) with live

Table 3: Behaviours used by the 8 observed squash coaches and 10 observed physiotherapists.

Behaviour	Occurrences		Frequency(behaviours per minute)		Percentage	
	Coaches	Physios	Coaches	Physios	Coaches	Physios
Pre-instruction	241	526	0.39	0.72	5.45%	11.41%
Concurrent Instruction (Positive)	1032	1314	1.66	1.80	23.35%	28.49%
Concurrent Instruction (Negative)	83	130	0.13	0.18	1.88%	2.82%
Post-instruction (Positive)	429	234	0.69	0.32	9.71%	5.07%
Post-instruction (Negative)	230	62	0.37	0.08	5.20%	1.34%
Manual Manipulation	12	323	0.02	0.44	0.27%	7.00%
Questioning	249	591	0.40	0.81	5.63%	12.81%
Positive Modelling	385	275	0.62	0.38	8.71%	5.96%
Negative Modelling	234	44	0.38	0.06	5.29%	0.95%
First Name	66	140	0.11	0.19	1.49%	3.04%
Hustle	73	87	0.12	0.12	1.65%	1.89%
Praise	1276	864	2.05	1.18	28.87%	18.73%
Scold	23	3	0.04	0.00	0.52%	0.07%
Console	87	19	0.14	0.03	1.97%	0.41%
Positive Reinforcement	0	0	0.00	0.00	0.00%	0.00%
Punishment	0	0	0.00	0.00	0.00%	0.00%
Total	4420	4612	7.11	6.31	100.00%	100.00%

coding and video coding having mean retest agreements of 95.8%. Each session was timed to the nearest 5 seconds so that the frequency of behaviours could be calculated. The coder also took field notes of any other behaviours or points of interest they noticed throughout the session and felt may be useful in achieving the aims of this study.

The physiotherapy sessions took place in a variety of locations including the homes of stroke survivors, physiotherapy gyms and rehabilitation centres. During the physiotherapy sessions, the observer's position varied depending on the space available but was always within earshot of the physiotherapist. Considering the squash sessions, the observer was positioned outside the court (either on the balcony or behind the glass back wall) and used a Bb Talkin' Advance microphone attached to the coach's clothing to allow them to hear what the coach was saying.

4.3 Results

The observed squash coaching sessions ranged in length (excluding breaks) from 22 to 55 minutes (mean = 38:51 ± 8:49 minutes). The observed physiotherapy sessions ranged in length from 20 to 59 minutes (mean = 36:32 ± 14:15 minutes). In total, 9032 behaviours were used across all 36 sessions. As per common practice in the sports coaching literature when using such a systematic observation instrument, a full breakdown of the behaviours used by all participants can be found in Table 3

All observed coaches used more positive behaviours (positive concurrent instruction, positive post instruction, positive modelling, and praise) than negative (negative concurrent instruction, negative post instruction, negative modelling, and scold) or neutral (pre-instruction, manual manipulation, questioning, first name, hustle, console) behaviours. However, there were noticeable differences in coaching styles, as shown in Figure 3 a. In particular, there was a wide variety in the amount of questioning, positive concurrent instruction, and positive and negative modelling used by the coaches.

For each physiotherapist, positive behaviours accounted for more than half of the total behaviours used, and for all but one, negative behaviours accounted for less than 10% of all behaviours. Figure 3 b shows that the use of pre-instruction was consistent between physiotherapists. However, there were big differences in many of the other behaviours, such as concurrent instruction (both positive and negative), questioning, manual manipulation, positive modelling, and first name.

There were other similarities in the data obtained between the two groups of participants. For example, praise was the most frequently used behaviour for 7 of the 8 observed coaches and was in the top 2 most frequently used behaviours for 7 of the 10 physiotherapists. Concurrent instruction (positive) was the most frequently used behaviour for 8 of the 10 observed physiotherapists and was in the top 2 most frequently used behaviours for 7 of the 8 squash coaches.

5 CLUSTERING

Through the systematic observation process detailed in Section 4.2, action sequences of coaching behaviours were produced. The similarities of the behaviours used by squash coaches and stroke physiotherapists indicate that a robotic coach implementing very similar behaviours could be used to good effect in both cases. However, the action sequences obtained through the observations were each only a single example of coaching and we didn't yet know how they would apply to different situations. To effectively use the observation data collected to guide the behaviour of a robotic coach, the next step in this work was to group the observed sessions into coaching policies using a clustering algorithm. This would produce a relatively small set of distinct policies based on our data that correspond meaningfully to different ways of coaching or personalising coaching for different types of people. In this context a policy is defined as a mapping between perceived states of the environment (i.e. the previous actions taken by the robotic coach) and the likelihood of each action being selected by the robotic coach when in that state.

A number of different clustering techniques could have been used at this stage (some of which are discussed in Section 2.2). However, given the similarity of the data obtained and the personal nature of the coaching process, it was decided to use Nikolaidis' expectation maximisation based algorithm [23] in the current work. This algorithm works by repeatedly executing an E-step (assigning each data point to the nearest cluster based on the sum of squared distance) and an M-step (updating each cluster's centroid using the assigned data points) until the assignments of each data point to a cluster do not change. The data points used in this case were transition matrices calculated from the action sequences obtained through the HHI observations. The transition matrices represented the likelihood of each behaviour occurring based on the previously observed behaviour. To obtain the number of clusters, and therefore the number of coaching policies generated, the mean gap value [45] over 5 runs was used. The gap value is a commonly used method of estimating the optimal number of clusters n for a given set of data points. It attempts to find the value for n where the decrease of the pooled within-cluster sum-of-squared error flattens markedly. This resulted in 6 clusters/policies for both the squash coaches' and stroke physiotherapists' data.

An adaption to our data was required to fit the algorithm because concurrent behaviours had been observed which were not a feature in Nikolaidis' work. Concurrent behaviours were encoded into compound behaviours and decoded again once clustering had occurred[3]. The data from each participant group was clustered separately to identify from the interviews if domain professionals would be able to provide confirmation of policies from the other case study. This would allow a very similar robotic coach to be used in a wide variety of cases.

Clustering the data produced policies which could be implemented in a robotic system to choose which action (behaviour) to take at each timestep of a coaching session. Learning the reward that corresponds to each of the clustered policies will also be required but has not been done yet because it was unnecessary for the current body of work. Due to the observation instrument used to obtain the original action sequences, we only have the behaviours used by the coach, not the joint actions of both parties involved as in Nikolaidis' work. The novelty of the current work comes from the open, real-world scenario of coaching from which our data was obtained and will be used, and in the use of behaviour graphs as a visual representation for the clustered policies. For a brief explanation of how the policies could be implemented to control the behaviour of a robotic coach and how to integrate the behaviours of the robot's interaction partner see Section 7.1.

[3]An exception to this was the concurrent behaviours which happened alongside manual manipulation (physically moving the client's body into the correct position) which were extensive in the physiotherapy observations. The envisioned robotic system will not be capable of performing manual manipulation so it was deemed most appropriate to remove these concurrent behaviours when clustering for simplicity in presenting the clustered data to participants during the interview process.

5.1 Behaviour Graphs

We wanted to obtain reflections on the policies created through clustering from squash coaches and stroke physiotherapists to confirm their applicability to both areas. However, the original dataset contained over 9000 behaviours used by 18 participants over 36 sessions and was therefore too complex to show in a meaningful way to coaches and physiotherapists. Therefore we created behaviour graphs (such as those in Figure 4), each of which is a visualisation of one of the coaching policies obtained through clustering that could be implemented in a roboic coach. These graphs provided a way to show real data to participants during semi-structured interviews in an anonymised manner that still kept the characteristics of the original data. In this way we could obtain the reflections of squash coaches and stroke physiotherapists on complex, real-world interaction data.

The graphs are simplified data visualisations of the underlying policies they represent, allowing us to visualise the temporal relationships between behaviours over a coaching session, as well as their relative frequencies. In the graphs, the size of each node represents the frequency of that behaviour and the transitions represent the likelihood of one action being followed by another. For example, the graph in Figure 4 a shows that 73% of the time the session will start with a pre-instruction, and following this the most likely next behaviour would be praise. Only behaviours which accounted for more than 5% of the total behaviours used by the coach, and only transitions with a probability greater than 0.1 are displayed[4]. The coloured areas within the nodes represent concurrent behaviours used. For example, in the graph shown in Figure 4 a the green areas within the pre-instruction and post-instruction (positive) boxes mean that some of the time the coach would use a demonstration of correct technique while giving a verbal instruction.

There were noticeable differences between the graphs produced from the clustered policies. Some of these differences can be seen in Figure 4 and all 12 graphs have been made available on GitHub[5] . For example, the graphs in Figures 4 a and 4b show lots of praise being used, whereas in Figure 4 d the praise box is very small in comparison to the amount of questioning used. Other big differences between the graphs were in the amount of manual manipulation, modelling and use of first name. Furthermore, many of the graphs tended to begin with pre-instruction (e.g. Figures 4 a, 4b and 4c) but others would start with a question (e.g. Figures 4 d) opening up the possibility for a robotic coach implementing these policies to be adaptive or adaptable depending on the user. For a description of how the underlying policies can be used as a starting point for high level personalisation during interactions between the planned robotic coaching system and it's user, see Section 7.1.

[4]Exceptions are the "Start" and "End" nodes whose size is consistent across all graphs and whose transitions are displayed if the probability exceeds 0.03. This was done to avoid situations where "Start" or "End" had no transitions coming in or out.
[5]The behaviour graphs representing the 12 clustered policies, as they were presented to the domain professionals during the interview process, can be viewed here: https://github.com/M4rtinR/BehaviourGraphVisualisations.

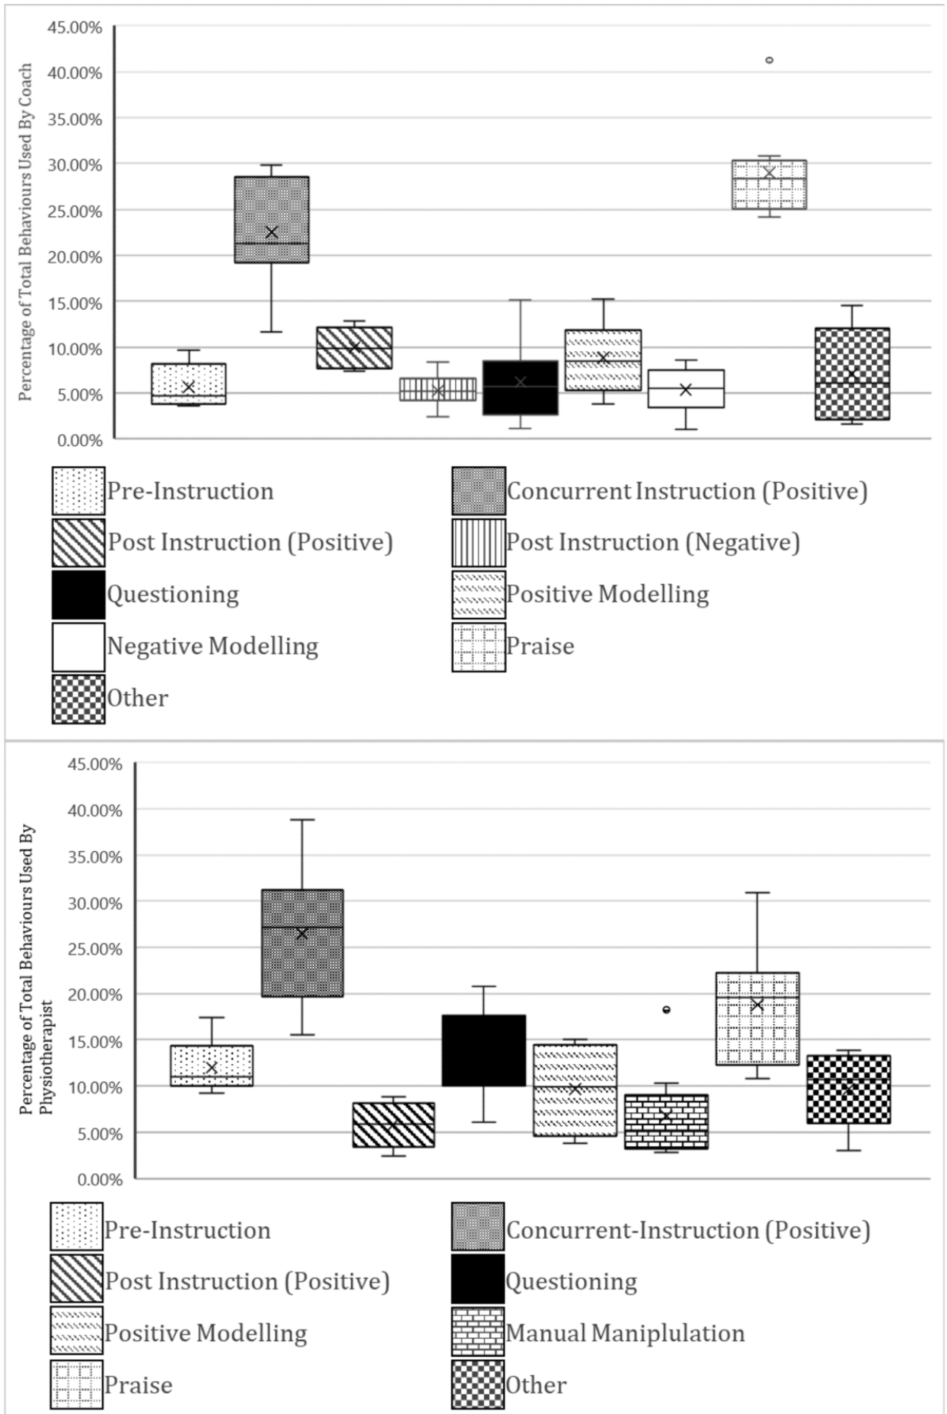

Figure 3: Box plots where each box shows the middle 50% of coaches' use of each behaviour and the whiskers show the top and bottom 25%. Only behaviours which accounted for more than 5% of total behaviours used are shown. (a) shows the behaviours used by squash coaches and (b) shows the behaviours used by stroke physiotherapists.

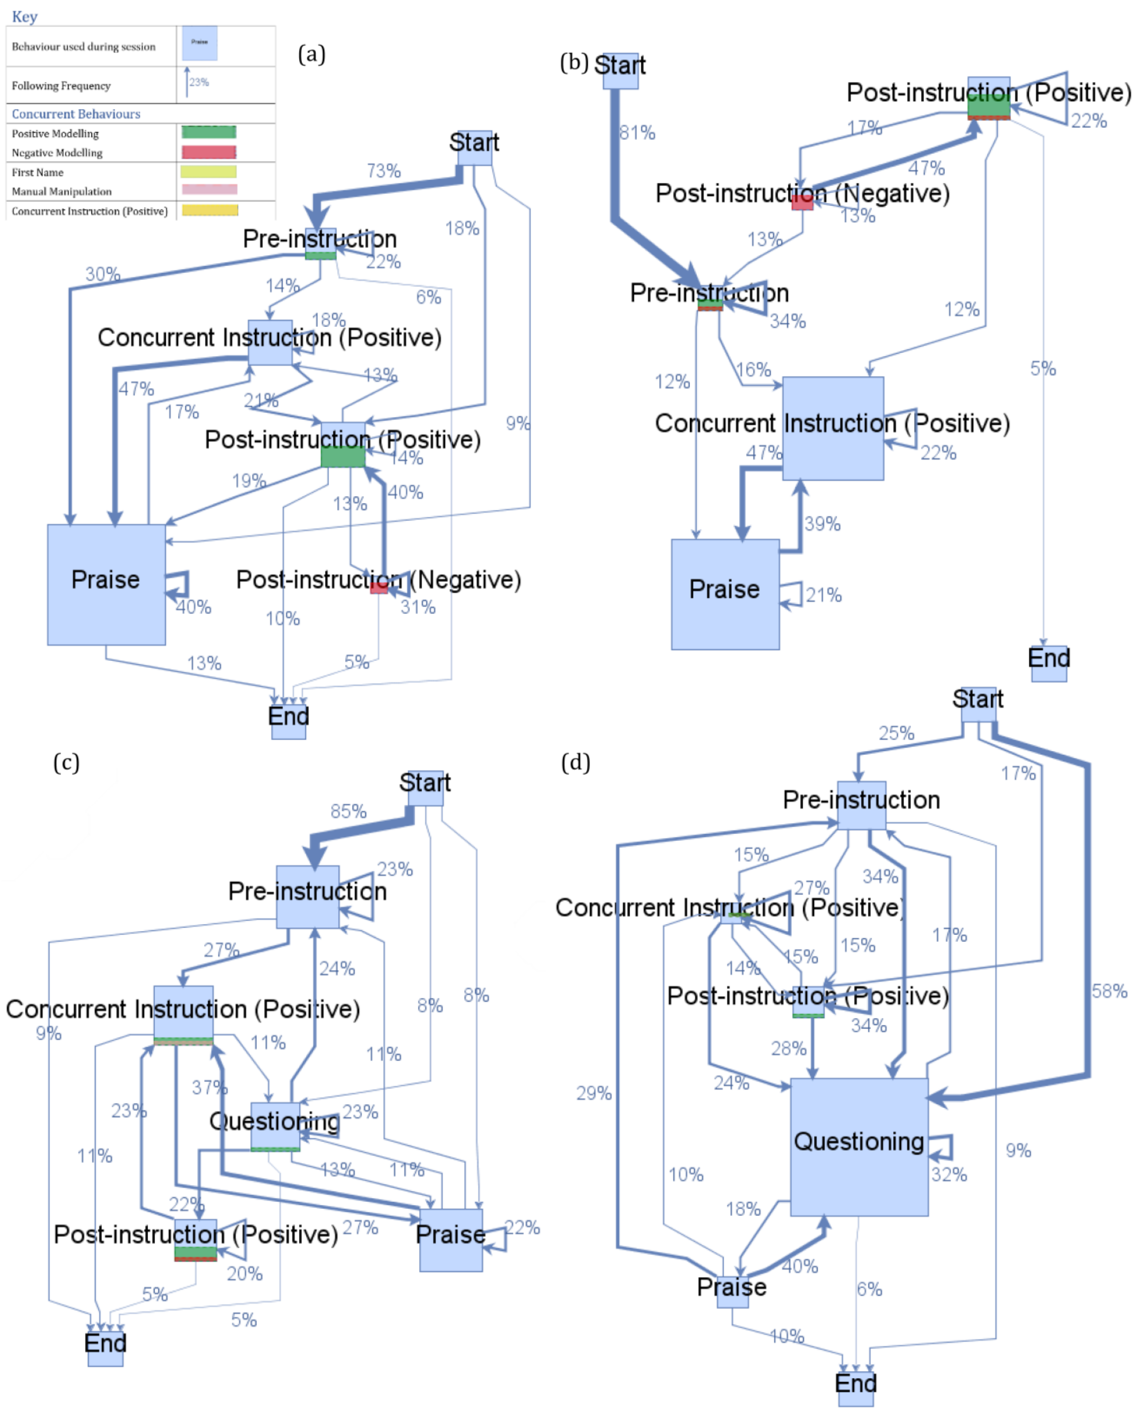

Figure 4: Four of the twelve behaviour graphs obtained through clustering the observation data. (a) and (b) come from the squash coaches' clustered data, while (c) and (d) come from the stroke physiotherapists' clustered data. Each graph is a visual representation of the underlying coaching policy which could be used for robotic control.

Table 4: Demographic Information of the coaches and physiotherapists interviewed in Study 2.

	Squash Coaches	**Stroke Physiotherapists**
Participants	9	8
Gender	8M, 1F	7F, 1M
Age Range	26-63 (Mean: 39.1 ± 11.6)	28-55 (Mean: 40.1 ± 8.8)
Qualification	3 Level 2, 4 Level 3, 1 Level 4, 1 Level 4 + MSc in Mental Toughness for Squash	6 BSc Physiotherapy, 1 Chartered Physiotherapist, 1 Grad Dip Phys, M.C.S.P
Experience	10-30 years (Mean: 16.7 ± 6.7)	4.5-31 years (Mean 20.9 ± 8.5)

Table 5: Segments of the semi-structured Interviews conducte

Phase	**Activities**	**Description of Questions**
Segment 1: Use of an Autonomous Robotic Coach	- Explanation and reminder of research project. - Video of a Pepper robot leading a rehabilitation exercise.	Broad, open-ended questions related to how a robotic coaching system could help in the daily lives of the participant/their clients to make the interviewee feel at ease with the situation [46]. It was found in [17] that participants in focus groups reacted quite negatively to the idea of a robotic coach before seeing a demonstration of its behaviour, so a decision was taken to begin with a video demonstration.
Segment 2: Discussion of Behaviour Graphs	- Introduce and describe example behaviour graph. - Explain Mural and share behaviour graphs.	The online whiteboard tool Mural [47] (overview shown in Figure 5) was used to share the 12 behaviour graphs (examples given in Figure 4) with participants. Each group of 6 graphs (grouped by domain) was shared separately and the order randomised. Each graph was numbered to facilitate discussion and a key similar to the one shown in Figure 4 was positioned in the centre of the 6 graphs. Questions were asked which encouraged the participant to explore the graphs and discuss the behaviours they noticed and scenarios in which they were likely to take place. Throughout this segment the facilitator annotated the graphs to ensure they had understood the participant fully.
Segment 3: Personalisation of Coaching	- Introduce and describe user traits.	Participants were asked how they would personalise their coaching behaviour and asked to identify any traits they would look for in their clients to inform their style of interaction. They were then introduced to a list of traits, shortened from those suggested by Winkle et al. [17] to include only the traits which could inform the behaviour of a robot during a coaching or rehabilitation session. These traits were: previous activity levels / engagement in sport; motivation / self-efficacy; functional goal(s) and/or interests. Participants were finally invited to relate these and their own traits back to the behaviour graphs seen during Segment 2.

6 STUDY 2 – SEMI-STRUCTURED INTERVIEWS AND BEHAVIOUR GRAPH DISCUSSION WITH COACHES AND PHYSIOTHERAPISTS

Following the clustering process, we used semi-structured interviews to explore and discuss with domain professionals the behaviour graphs and when it would be appropriate to use each of them. An autonomous robotic coach which was able to choose which of the confirmed behaviour graphs to use has the potential to personalise its behaviour in an attempt to increase motivation and adherence to an individual exercise routine.

6.1 Method

6.1.1 Participants. Table 4 shows the demographic information of all 17 interview participants. The majority of participants took part in the initial observation study, but 2 additional participants from each domain were interviewed.

6.1.2 Procedure. Qualitative data was gathered using the semi-structured interview process detailed in Table 5. The interviews were one-to-one, conducted via video call and recorded for later transcription. Consent was obtained through electronic signature form each interviewee to participate in the interview and to publicly release their anonymised data.

At the start of each interview, an overview of the work was given by the interviewer and a video was shown which displayed a Pepper robot leading an able-bodied person through a stroke rehabilitation exercise. In the video, the robot explains and demonstrates an "arm-to-side" exercise, displaying images of the exercise on it's tablet computer screen, and then exercises along with the user. This led into a discussion around how such a robot could be used to assist the coach or physiotherapist with their clients. During segment 2 discussions of the behaviour graphs took place. All 12 behaviour graphs were shown to participants, a selection of which can be seen in Figure 4. Of these 12 graphs, 6 came from clustering the observed squash coaches' data and 6 from the stroke physiotherapists' data,

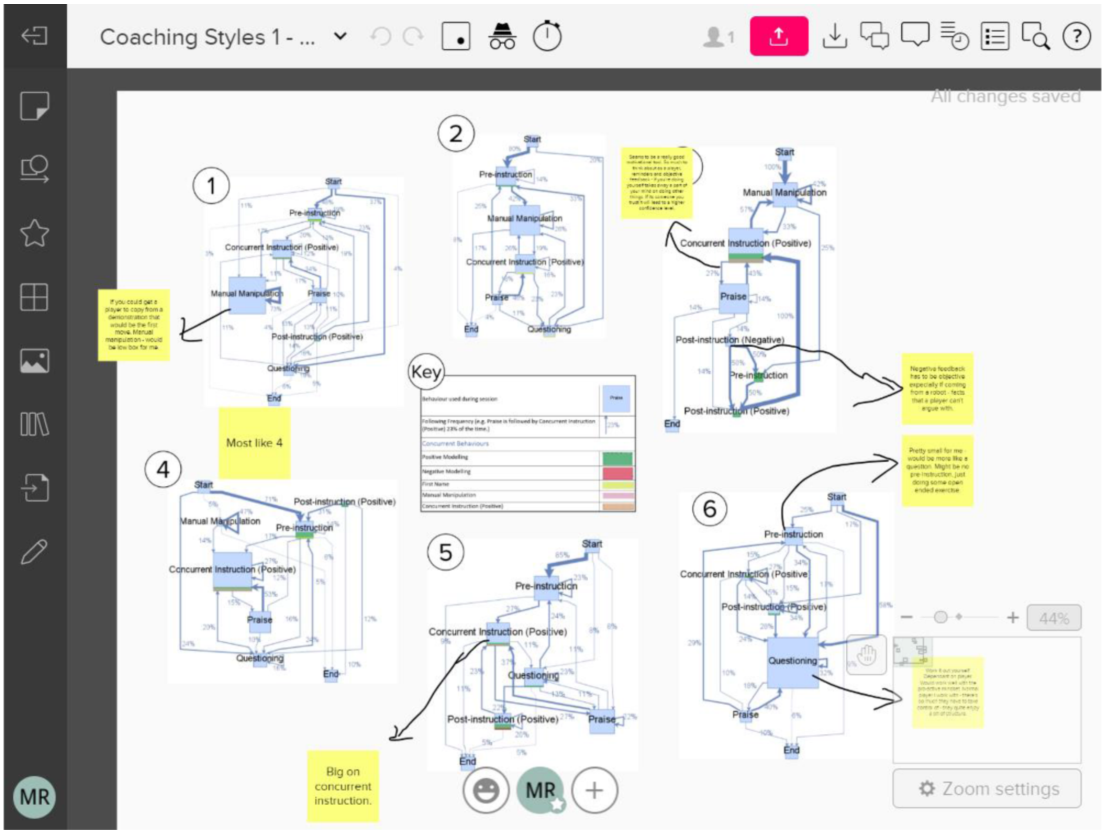

Figure 5: The online collaborative whiteboard tool Mural being used for annotations during one of the interviews, shown here as an overview to clarify how the tool was used. 6 of the behaviour graphs were displayed and discussed at a time with participants able to zoom in on particular graphs or graph elements.

therefore coaches viewed and commented on data obtained through observations of physiotherapists as well as squash coaches, and vice versa. Typically, in the domain of sports, coaches are presented with their own data during interviews [25]. However, discrepancy is often found between what they were observed doing, and their understanding and explanations of why they used the behaviours they did [34]. A novelty of this study is that the data they are presented is not just their own, so this discrepancy should be avoided. Additionally, if participants confirmed the applicability of all the graphs to their area, this would further emphasise that a system implementing very similar behaviour could be used in both sports and rehabilitation. The online whiteboard tool Mural [47] was used to share the graphs with participants. Mural allowed collaborative annotating of the graphs as discussions took place and a screenshot showing an overview of the tool can be seen in Figure 5

6.1.3 Data Analysis. The transcripts were analyzed using the Constant Comparative Method (CCM) [48], a data analysis method of Grounded Theory. This is a method of qualitative data analysis that aims to allow the generation of themes using specific coding and analytic procedures which has been used successfully in past HCI/HRI studies [49], [50]. The analysis was verified by an experienced HCI researcher external to the project. The annotated murals were also analysed to ensure they correlated with the interview transcriptions.

6.2 Results

The interviews lasted between 36 and 68 minutes. The results presented in this section have been structured to provide useful insights into how an autonomous robotic coach for use in both squash coaching and stroke physiotherapy could be designed.

6.2.1 Robot Features and Uses. All of the physiotherapists and six of the nine squash coaches agreed that the robot could be used effectively in helping someone with individual training. Two squash coaches also suggested its use as an assistant during group training sessions led by the coach.

Eight of the squash coaches identified the robot as being a useful source of information either with built in drills and exercises which would help the user improve towards their goals or with feedback it could give on technical aspects of the game. The majority of stroke physiotherapists (six out of eight) on the other hand saw the robot being used as more of a motivating tool and often referred to the idea of the user building a relationship with it: "*I think the robot, maybe this is just me, but I think for some people almost feels like you've got a little buddy.*"

Both sets of participants (five physiotherapists and six squash coaches) viewed the way the robot interacts as a very important part of the system. They were enthusiastic about the idea of the robot being able to use multiple interaction modalities, expressing the importance of this to work with a wide variety of people: "*You've got demonstration. You've got the video screen you could be putting diagrams up and you could get a video and they, and then the robot can verbally prompt them through things.*" A key difference between the groups was the heavy emphasis of the use of the Pepper robot's screen by squash coaches, particularly for demonstrations: "*Obviously if you've got like a video of I don't know, James Willstrop [a former world number 1 squash player renowned for his accuracy] showing you how to do it. You know obviously kind of gravitate towards that more than, than a robot showing you that same kind of you know swing or whatever it was.*" The physiotherapists saw more use in the robot providing physical demonstrations rather than the focus being mainly on the screen.

6.2.2 Behaviour Graph Selection and High-Level Personalisation. Annotating the behaviour graphs through Mural during discussions with participants in Segment 2 of the interviews helped confirm the appropriateness of the coaching policies which the graphs represent, for both domains. During the conversation, participants were specifically asked if they would use the particular graph or behaviour currently being discussed in their own coaching. All participants responded that they would, with one exception. One physiotherapist and two squash coaches also stated outright, without prompt, that all of the graphs presented were valid: "*Well yeah, I can see how all the different graphs could work in, in different scenarios.*" The only exception to this was in the amount of manual manipulation used by squash coaches. However, through further discussion of the given graphs, coaches were able to describe situations in which manual manipulation could be useful – particularly with beginner players.

Participants were also invited to compare and contrast the different graphs. These discussions helped to identify situations which each graph would be most suited to. Table 6 summarises this information by grouping coaching/physiotherapy sessions into categories and gives a list of the graphs which would be most suited to a session in the given category. A selection of the behaviour graphs is given in Figure 4 and the full set can be found by following the link in the footnote of Section 5.1. These example categories were the most frequently cited use of each of the presented graphs by both groups of participants and are based on the traits identified by participants in Segment 3 of the interviews. The end goal of this work is to produce a data-driven controller for a robotic coach and the information presented in Table 6 will enable this. For a full explanation of how an HRI system could use this information to personalise its behaviour see Section 7.1.

In Segment 3 of the interview process, discussions centred around the identification of traits which could be used to inform personalisation of interaction style used by our participants. All but one participant either agreed with the three given traits (previous activity levels/engagement in sport, motivation/self-efficacy, and functional goals) or identified them themselves as things that would affect their behaviours within a session. The one exception was a squash coach who felt that: "*You can eliminate a few possibilities based upon the information that you get within this, but I wouldn't necessarily say it helped you to pick the optimum approach to take from my coaching point of view.*" This coach felt that getting to know a particular player was the best way to decide on a method of coaching. This emphasises suggestions in the literature [14], [16], [17] that both high level personalisation based on user traits and low level adaption to individuals is required for a robotic coaching system: "*so you're gonna have to actually build sort of character coaches as such you know or, build up certain personality traits that if they've got this sort of trait, this is the program that we've got to deliver.*" A description of the user traits commonly identified by participants can be found in Table 7

Five of the physiotherapists also stressed the importance of including family in the rehabilitation process as early as possible. However, this was more likely to impact the exercises prescribed to patients or the quantity of information delivered rather than particular coaching behaviours used within a session.

Additionally, a big focus was put on the relationship which is built up with the client over a period of time. It is an important skill of a physiotherapist or coach to be able to read someone's personality and figure out the learning style of each individual they interact with: "*I think that's what makes you know the really good therapists really good, 'cause they can change it so much. . . So much of what we have to do is we have to connect with the client.*" This further indicates the need for both high level and low level personalisation in a robotic coach.

7 DISCUSSION

7.1 Implementation Process

The main contribution of this work is the data obtained in the described observation and interview studies which can form a data-driven controller for an HRI system. Therefore, instead of being presented with high-level design implications, it is most useful for the reader to understand exactly how the results presented could be used for implementation. In this work we plan to use a Pepper robot to autonomously guide a user through an exercise routine. However using a very similar process it would be possible for other HCI applications/systems to be created. For example a mobile application or smart speaker could perform a subset of the coaching behaviours given or mixed reality devices could be used to implement virtual coaches. Although a full description of the planned implementation process is out with the scope of the current work, it can be summarised as follows.

Through clustering, we have obtained twelve different coaching policies, which we have visually represented as behaviour graphs. Using a method similar to [23], we will first learn a reward function from the underlying policies through inverse reinforcement learning (we plan to use Maximum Entropy IRL [51]). The learned reward structure can feed into a Partially Observable Markov Decision Process (POMDP) model of the coaching scenario. If it was known which policy would work best for a user, that policy could simply be chosen and executed. One of the benefits of modelling the interaction as a POMDP is that it allows the system to reason about which policy (or mix of policies) will work best and handle that uncertainty. The partially observable variable in our case will be the "type" of the user, each type corresponding to a particular

Table 6: Categories of session that each behaviour graph would be most likely to suit. The categories are taken directly from the discussions and annotations of the behaviour graphs during Segment 2 of the interviews. The numbers in the "graphs" column correspond to the 12 graphs presented to interview participants and can be seen by following the link in the footnote of Section 5.1.

Category	Graphs	Description
Low motivation	2 (Fig.4 a), 6, 9	Someone low on motivation who needs encouragement (this might involve an open, continuous session in squash or trying new exercises in stroke rehabilitation).
High motivation	3 (Fig.4 b)	Someone high on motivation looking to make a change in technique for an exercise.
Early in relationship	1, 4, 9	One of the first sessions with a player or stroke survivor when the coach or physiotherapist is looking to form a connection with them.
Late in relationship	12 (Fig.4 d)	An experienced squash player or a stroke survivor quite far on in their rehabilitation who the coach/physiotherapist is trying to encourage to become independent.
Beginner player	7, 8	A squash player just starting out.
Experienced player	4, 11 (Fig.4 c)	A squash player who has been playing the sport for a long time and at a high level.
Repetitive exercises	3 (Fig.4 b), 4, 6, 8, 9, 10	A session involving lots of repetitive exercises (e.g. high-level balance in stroke rehabilitation or technical practice in squash).
Complex exercises	2 (Fig4 a), 5, 11 (Fig.4 c)	A session involving complex exercises which require a lot of explanation (e.g. working on a new tactic in squash or a new movement in stroke rehabilitation).
Challenging exercises	7, 8, 10	An exercise which the player/stroke survivor finds very difficult at first (this could involve a new technique in squash or progressing a physical movement in stroke rehabilitation).
High cognitive impairment	1, 5, 8	Interacting with a stroke survivor who e.g. gets confused easily or suffers emotional lability.
High physical impairment	7	Interacting with a stroke survivor who has very limited movement in the target limb.

Table 7: Traits commonly identified by squash coaches and stroke physiotherapists as things which would help them personalise their coaching behaviour.

Trait	Description	Identified By
Motivation/Self-efficacy	Motivation effects how much a person would do on their own. There was a heavy emphasis with physiotherapists to find out why somebody is not motivated. Whereas squash coaches would try different methods to improve motivation such as making the session more enjoyable or making sure progress was being made by the player.	7 Coaches 7 Physiotherapists
Goals	A collaborative approach to goal setting was taken by seven of the squash coaches and all eight physiotherapists. Physiotherapists strongly linked goals to motivation, pointing out the need for both parties to be on the same page and breaking down long-term goals into short-term ones. These short-term goals were more frequently discussed by both groups, highlighting the need to revisit goals continually as a measure of progress and a tool for motivation.	9 Coaches 8 Physiotherapists
Previous activity levels/ engagement in sport	The playing experience and ability of the player, particularly in squash but also in other sports, would dictate the amount of praise (lower for experienced players) and questioning (higher for experienced players) used in sessions.	9 Coaches 4 Physiotherapists
Type of session	If a session was technically- or tactically-based, coaches would expect to use more questioning. However in a session which was focused more on fitness or execution under pressure, more concurrent instruction and negative behaviours would occur to make sure the play didn't break down.	4 Coaches 0 Physiotherapists
Effects of condition	A stroke can have different effects on different people so it is important to be aware of this. If the patient has language difficulties, more visual cues might be necessary, some cognitive issues would require simplification of the message, and some people might need pushed a bit harder than others who were more in tune with their bodies.	0 Coaches 8 Physiotherapists

coaching policy. During execution the POMDP will maintain a belief distribution over the 12 policies. To form this, it will incorporate the recommendations in Table 6 and user-provided information on the traits given in Table 7. For example, when interacting with a user who has reported high levels of motivation that day and is conducting repetitive exercises, the system's belief distribution will make it more likely to choose an action based on the policy represented by graph 3 (shown in Figure 4 b). If a particular session falls into more than one of the categories given in Table 6, the belief distribution will be split across these categories. The policy which solves this POMDP will provide short-term, high-level personalisation when implemented alongside interactional behaviour in a robotic coach.

Domain specific coaching behaviours will comprise the robot's actions associated with the policy during the coaching session. The majority of a user's actions will be the user performing a particular exercise. In this case, data obtained from sensors attached to the user's body/equipment would give an indication of the stage of the interaction and therefore which of the coaching behaviours would be appropriate at that time. Additional user actions during policy execution will include selecting which exercise they would like to practice during the session, how a particular exercise felt in terms of difficulty or pain level, or asking for clarification of an instruction. All of these actions will result in a different subset of the coaching behaviours being available to the robot at a given timestep. In this way, the system can choose appropriate actions based on the stage of the interaction while still being data-driven using the POMDP policy.

7.2 Similarities in Coaching Behaviours Across Domains

The above description of how the envisioned robotic coaching system could be implemented was made possible by fulfilling the aims of the current work. The first aim (**A1**) was to gather and analyse data about the behaviours of stroke physiotherapists and squash coaches. The data presented in Table 3 satisfies this aim. The systematic observation process we used, detailed in Sections 4.1 & 4.2, allowed in-depth action sequences of behaviours used in the real world by professionals in both areas to be obtained. This allowed direct comparisons to be made between the two domains, satisfying **A2**: to discover differences and commonalities in the behaviours of physiotherapists and squash coaches to predict whether a robotic coach implementing similar behavioural policies could be used in both fields. This was a novelty of our first study and the second contribution of our work. No previous work has been identified which compares the behaviours of sports coaches with rehabilitation physiotherapists in one-to-one sessions. Additionally, this work investigates the behaviours of coaches in individual sports and physiotherapists in stroke rehabilitation in a level of detail which has not been done before.

We found striking similarities in terms of the frequency of praise and positive concurrent instruction (e.g. Figures 4 a, b & c), and the lack of negative behaviours used by both groups (e.g. negative post-instruction was the only negative behaviour above the 5% threshold in all of the clustered graphs given in Figure 4). This leads us to believe that an autonomous robotic coach implementing very similar behaviours could be used in either area. Previous research using systematic observations in coaching for team sports indicates that coaches structure sessions and use behaviours in a way that maximises the transfer of information to players [26], [27]. A similar trend was seen in our work, with instructional behaviours accounting for 45.59% of the total behaviours used by squash coaches and 49.13% of behaviours used by physiotherapists.

Despite this evidence showing the similarities in the behaviours used by squash coaches and stroke physiotherapists, there were some small differences as well. Sports coaches have previously been shown to adapt their behaviour to form a relationship with their players [28]. It seems this was also the case for our participants, with coaches and physiotherapists bringing their own style of interaction to each session. The biggest difference between the domains examined in this work was in manual manipulation which was used more frequently by stroke physiotherapists than by squash coaches. However, this is a behaviour that a social robotic coach could not replicate effectively given the capabilities of today's technology and could be effective without it [13], [14], [52]. This, and the other subtle differences identified, highlight the differences in application area of these two groups of professionals and the need for such systematic observations in any domain. No positive reinforcement (physical reward) or punishment (physical retribution) was used by any of the observed coaches or physiotherapists (see Table 1 for full definitions of these behaviours) so we recommend these categories be removed from the instrument for any future studies of this nature.

7.3 Behaviour Graphs

A third contribution of this work is the representation of coaching policies as behaviour graphs. Each of the 12 graphs created (examples shown in Figure 4) is a visual representation of one of the policies obtained through clustering the observational data. These were presented to squash coaches and stroke physiotherapists to achieve **A3** of this work: to obtain professionals' reflections of the behaviour graphs generated by the clustering algorithm. It is common to obtain reflections of coaches on their own coaching practice [25]. Our novel strategy, which combines data-driven and qualitative methods, allowed discovery of situations and types of users for which specific coaching policies would be most appropriate without bias of the interviewee viewing their own data.

Participants were able to read and understand the technical detail given in the presented graphs and they effectively facilitated discussions on coaching behaviours which helped to produce the recommended uses of each graph given in Table 6. This is a vital step in achieving high level personalisation of a robotic system. Knowing when, and with which type of person, to apply each policy has been demonstrated to produce a more responsive system which improved team efficiency in a collaborative packing task [23]. This work is a starting point of applying a similar technique to the more open, real-world task of coaching. A robotic coaching system which could choose which of the behaviour graphs to use could help combat staffing difficulties which are common in this area [53] and might allow rehabilitation to become closer to the ideal.

The fourth and final aim of this work (**A4**) was to discover how domain professionals personalise their coaching approach and gauge which factors would be most useful to take into account when deciding on a behaviour graph to be used by our robotic system. This aim originates from the recommendations in literature that personalisation is becoming a requirement of such a robotic coach [14], [16], [17]. The results presented in Section 6.2.2 identify user traits which the interviewees commonly identified as using to personalise their own interaction styles. The list of user traits discovered in this work differs slightly from those presented by Winkle et al. [17]. The coaches and physiotherapists interviewed in Study 2 also identified motivation and individual goals as factors which would affect their behaviour. However, the physiotherapists in the current work put less focus on the previous activity levels/engagement in sport of their clients than the squash coaches did. An additional trait identified in this work was the effect of a patient's condition. This is similar to "cognition" defined by Winkle but encompasses physical, as well as cognitive, effects. Physiotherapists in particular advocated the use of all of the possible modalities of interaction available on a robot to allow the system to be effectively used by the widest range of people. In addition to motivation and goals, the squash coaches also identified the type of session as something which would inform their coaching behaviours.

8 CONCLUSION

This paper has presented a novel mixed methodology approach to informing the design of a personalised robotic coach for sports and rehabilitation coaching. By first performing systematic observations of domain professionals in squash coaching and stroke physiotherapy, we have been able to cluster the obtained action sequences of coaching behaviours into coaching policies usable for robotic control. Then, by producing behaviour graph visualisations of these policies and presenting them to professional squash coaches and stroke physiotherapists during semi-structured interviews, we have gained an understanding of the needs of an HRI system to be used in each case. We believe our method of data collection could be applied to a wide variety of cases to enhance future design in HCI for interaction with data-driven systems. Furthermore, the similarities found and the confirmation of them by coaches and physiotherapists as regards to coaching policies which could be used for squash coaching and stroke physiotherapy, indicate that a robotic coach implemented as described in Section 7.1 could be effective in increasing adherence and performance.

ACKNOWLEDGMENTS

The authors would like to thank Scottish Squash for their help with participant recruitment. Also, thanks go to Edward Hall for lending his help in the coder training phase. This work was funded by Engineering and Physical Sciences Research Council (EPSRC) Grant ID: EPSRC DTP18.

REFERENCES

[1] World Health Organisation, "WHO EMRO | Stroke, Cerebrovascular accident | Health topics," 2018. http://www.emro.who.int/health-topics/stroke-cerebrovascular-accident/index.html
[2] S. Hillier and G. Inglis-Jassiem, "Rehabilitation for community-dwelling people with stroke: home or centre based? A systematic review," *International Journal of Stroke*, vol. 5, pp. 178–186, 2010, doi: 10.1111/j.1747-4949.2010.00427.x.
[3] K. K. Miller, R. E. Porter, E. DeBaun-Sprague, M. V. Puymbroeck, and A. A. Schmid, "Exercise after Stroke: Patient Adherence and Beliefs after Discharge from Rehabilitation," *Topics in Stroke Rehabilitation*, vol. 24, no. 2, pp. 142–148, Feb. 2017, doi: 10.1080/10749357.2016.1200292.
[4] J. Baker, J. Côté, and B. Abernethy, "Learning from the Experts: Practice Activities of Expert Decision Makers in Sport," *Research Quarterly for Exercise and Sport*, vol. 74, no. 3, pp. 342–347, Sep. 2003, doi: 10.1080/02701367.2003.10609101.
[5] M. Kilpatrick, E. Hebert, and J. Bartholomew, "College Students' Motivation for Physical Activity: Differentiating Men's and Women's Motives for Sport Participation," *Journal of American College Health*, vol. 54, no. 2, pp. 87–94, 2014, doi: 10.3200/JACH.54.2.87-94.
[6] B. J. Almagro, P. Sáenz-López, and J. A. Moreno, "Prediction of sport adherence through the influence of autonomy-supportive coaching among spanish adolescent athletes," *Journal of Sports Science and Medicine*, vol. 9, no. 1, pp. 8–14, 2010.
[7] N. Maclean, P. Pound, C. Wolfe, and A. Rudd, "The Concept of Patient Motivation: A Qualitative Analysis of Stroke Professionals' Attitudes," *Stroke*, vol. 33, no. 2, pp. 444–448, Feb. 2002, doi: 10.1161/hs0202.102367.
[8] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 1st ed. Cambridge, MA: MIT Press, 1998.
[9] P. B. Gorelick, "The global burden of stroke: persistent and disabling," *The Lancet Neurology*, vol. 18, no. 5, pp. 417–418, May 2019, doi: 10.1016/S1474-4422(19)30030-4.
[10] Stroke Association, "State of the nation Stroke statistics - February 2018," 2018. doi: 10.1080/07393149508429756.
[11] C. O'Neil, M. D. Dunlop, and A. Kerr, "Supporting Sit-To-Stand Rehabilitation Using Smartphone Sensors and Arduino Haptic Feedback Modules," in *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*, Copenhagen Denmark, Aug. 2015, pp. 811–818. doi: 10.1145/2786567.2793705.
[12] G. Alankus, A. Lazar, M. May, and C. Kelleher, "Towards customizable games for stroke rehabilitation," in *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, New York, New York, USA, 2010, pp. 2113–2122. doi: 10.1145/1753326.1753649.
[13] E. Wade, A. R. Parnandi, and M. J. Matarić, "Using Socially Assistive Robotics to Augment Motor Task Performance in Individuals Post – Stroke," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 2403–2408.
[14] R. Feingold Polak and S. L. Tzedek, "Social Robot for Rehabilitation: Expert Clinicians and Post-Stroke Patients' Evaluation Following a Long-Term Intervention," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, United Kingdom, Mar. 2020, pp. 151–160. doi: 10.1145/3319502.3374797.
[15] K. Winkle, S. Lemaignan, P. Caleb-Solly, U. Leonards, A. Turton, and P. Bremner, "Couch to 5km Robot Coach: An Autonomous, Human-Trained Socially Assistive Robot," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, United Kingdom, Mar. 2020, pp. 520–522. doi: 10.1145/3371382.3378337.
[16] A. Kubota, E. I. C. Peterson, V. Rajendren, H. Kress-Gazit, and L. D. Riek, "JESSIE: Synthesizing Social Robot Behaviors for Personalized Neurorehabilitation and Beyond," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, United Kingdom, Mar. 2020, pp. 121–130. doi: 10.1145/3319502.3374836.
[17] K. Winkle, P. Caleb-Solly, A. Turton, and P. Bremner, "Social Robots for Engagement in Rehabilitative Therapies: Design Implications from a Study with Therapists," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18*, 2018, pp. 289–297. doi: 10.1145/3171221.3171273.
[18] J. Fasola and M. J. Matarić, "Using Socially Assistive Human-Robot Interaction to Motivate Physical Exercise for Older Adults," in *Household Service Robotics*, vol. 100, no. 8, 2014, pp. 487–516. doi: 10.1016/B978-0-12-800881-2.00022-0.
[19] B. Irfan, M. Hellou, A. Mazel, and T. Belpaeme, "Challenges of a Real-World HRI Study with Non-Native English Speakers: Can Personalisation Save the Day?," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, United Kingdom, Mar. 2020, pp. 272–274. doi: 10.1145/3371382.3378278.
[20] G. Gordon *et al.*, "Affective personalization of a social robot tutor for children's second language skills," in *Proceedings of the 30th Conference on Artificial Intelligence (AAAI 2016)*, 2016, pp. 3951–3957. doi: 10.1016/j.bbmt.2012.12.001.
[21] S. Roy, E. Kieson, C. Abramson, and C. Crick, "A reinforcement learning model for robots as teachers," in *27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2018, pp. 294–299.
[22] L. Chen, R. Paleja, M. Ghuy, and M. Gombolay, "Joint Goal and Strategy Inference across Heterogeneous Demonstrators via Reward Network Distillation," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, United Kingdom, Mar. 2020, pp. 659–668. doi: 10.1145/3319502.3374791.
[23] S. Nikolaidis, R. Ramakrishnan, K. Gu, and J. Shah, "Efficient Model Learning from Joint-Action Demonstrations for Human-Robot Collaborative Tasks," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*, Mar. 2015, pp. 189–196. Accessed: Mar. 02, 2020. [Online]. Available: https://dl.acm.org/doi/10.1145/2696454.2696455

[24] H. Yin, F. S. Melo, A. Paiva, and A. Billard, "An ensemble inverse optimal control approach for robotic task learning and adaptation," *Auton Robot*, vol. 43, no. 4, pp. 875–896, Apr. 2019, doi: 10.1007/s10514-018-9757-y.

[25] E. Cope, M. Partington, and S. Harvey, "A review of the use of a systematic observation method in coaching research between 1997 and 2016," *Journal of Sports Sciences*, vol. 35, no. 20, pp. 2042–2050, 2017, doi: 10.1080/02640414.2016.1252463.

[26] P. Potrac, R. Jones, and C. Cushion, "Understanding Power and the Coach's Role in Professional English Soccer: A Preliminary Investigation of Coach Behaviour," *Soccer & Society*, vol. 8, no. 1, pp. 33–49, 2007, doi: 10.1080/14660970600989509.

[27] S. Horton, J. Baker, and J. Deakin, "Expert in action: A systematic observation of 5 national team coaches," *International Journal of Sport Psychology*, vol. 36, no. 4, pp. 299–319, 2005.

[28] P. Potrac, R. Jones, and K. Armour, "'It's All About Getting Respect': The Coaching Behaviors of an Expert English Soccer Coach," *Sport, Education and Society*, vol. 7, no. 2, pp. 183–202, 2002, doi: 10.1080/1357332022000018869.

[29] H. Rutjes, M. C. Willemsen, and W. A. IJsselsteijn, "Beyond Behavior: The Coach's Perspective on Technology in Health Coaching," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, Glasgow, Scotland Uk, 2019, pp. 1–14. doi: 10.1145/3290605.3300900.

[30] P. M. van Vilet, N. B. Lincoln, and E. Robinson, "Comparison of the content of two physiotherapy approaches for stroke," *Clinical Rehabilitation*, vol. 15, pp. 398–414, 2001.

[31] L. Sussenbach *et al.*, "A robot as fitness companion: Towards an interactive action-based motivation model," *IEEE RO-MAN 2014 - 23rd IEEE International Symposium on Robot and Human Interactive Communication: Human-Robot Co-Existence: Adaptive Interfaces and Systems for Daily Life, Therapy, Assistance and Socially Engaging Interactions*, pp. 286–293, 2014, doi: 10.1109/ROMAN.2014.6926267.

[32] D. Hebesberger, T. Koertner, C. Gisinger, J. Pripfl, and C. Dondrup, "Lessons learned from the deployment of a long-term autonomous robot as companion in physical therapy for older adults with dementia: A mixed methods study," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2016, vol. 2016-April, pp. 27–34. doi: 10.1109/HRI.2016.7451730.

[33] M. Partington, C. Cushion, and S. Harvey, "An investigation of the effect of athletes' age on the coaching behaviours of professional top-level youth soccer coaches," *Journal of Sports Sciences*, vol. 32, no. 5, pp. 403–414, 2014, doi: 10.1080/02640414.2013.835063.

[34] S. Harvey, C. J. Cushion, E. Cope, and B. Muir, "A season long investigation into coaching behaviours as a function of practice state: the case of three collegiate coaches," *Sports Coaching Review*, vol. 2, no. 1, pp. 13–32, May 2013, doi: 10.1080/21640629.2013.837238.

[35] S. Uzor and L. Baillie, "Exploring & designing tools to enhance falls rehabilitation in the home," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, pp. 1233–1242, 2013, doi: 10.1145/2470654.2466159.

[36] A. C. Lacy and P. W. Darst, "Evolution of a Systematic Observation System: The ASU Coaching Observation Instrument," *Journal of Teaching in Physical Education*, vol. 3, pp. 59–66, 1984, doi: 10.1123/jtpe.3.3.59.

[37] P. R. Ford, I. Yates, and A. M. Williams, "An analysis of practice activities and instructional behaviours used by youth soccer coaches during practice: Exploring the link between science and application," *Journal of Sports Sciences*, vol. 28, no. 5, pp. 483–495, 2010, doi: 10.1080/02640410903582750.

[38] M. Smith and C. J. Cushion, "An investigation of the in-game behaviours of professional, top-level youth soccer coaches," *Journal of Sports Sciences*, vol. 24, no. 4, pp. 355–366, Apr. 2006, doi: 10.1080/02640410500131944.

[39] C. Cushion, S. Harvey, B. Muir, and L. Nelson, "Developing the Coach Analysis and Intervention System (CAIS): Establishing validity and reliability of a computerised systematic observation instrument," *Journal of Sports Sciences*, vol. 30, no. 2, pp. 201–216, 2012, doi: 10.1080/02640414.2011.635310.

[40] J. E. R. Staddon and D. T. Cerutti, "Operant Conditioning," *Annu Rev Psychol*, vol. 54, pp. 115–144, 2003, doi: 10.1146/annurev.psych.54.101601.145124.

[41] E. T. Hall, S. Gray, and J. Sproule, "The microstructure of coaching practice: behaviours and activities of an elite rugby union head coach during preparation and competition," *Journal of Sports Sciences*, vol. 34, no. 10, pp. 896–905, May 2016, doi: 10.1080/02640414.2015.1076571.

[42] A. C. Lacy and P. W. Darst, "The Arizona State University observation instrument (ASUOI)," in *Analyzing Physical Education and Sport Instruction*, Human Kinetics Publications, Box 5076, Champaign, IL 61820 ($35, 1989.

[43] J. Côté and W. Gilbert, "An Integrative Definition of Coaching Effectiveness and Expertise," *International Journal of Sports Science & Coaching*, vol. 4, no. 3, pp. 307–323, 2009.

[44] M. Partington and C. J. Cushion, "Performance during performance: using Goffman to understand the behaviours of elite youth football coaches during games," *Sports Coaching Review*, vol. 1, no. 2, pp. 93–105, Oct. 2012, doi: 10.1080/21640629.2013.790167.

[45] R. Tibshirani, G. Walther, and T. Hastie, "Estimating the number of clusters in a data set via the gap statistic," *J Royal Statistical Soc B*, vol. 63, no. 2, pp. 411–423, May 2001, doi: 10.1111/1467-9868.00293.

[46] A. Galletta, Mastering the Semi-Structured Interview and Beyond: From Research Design to Analysis and Publication. NYU Press, 2013.

[47] "MURAL." https://www.mural.co/ (accessed Oct. 03, 2020).

[48] B. G. Glaser and A. L. Strauss, "The Constant Comparative Method of Qualitative Analysis," in *The discovery of grounded theory: strategies for qualitative research*, 4. paperback printing., New Brunswick: Aldine, 1967, pp. 101–115.

[49] T. Georgiou, L. Baillie, M. K. Ross, and F. Broz, "Applying the Participatory Design Workshop Method to Explore how Socially Assistive Robots Could Assist Stroke Survivors," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge United Kingdom, Mar. 2020, pp. 203–205. doi: 10.1145/3371382.3378232.

[50] M. Akazue, M. Halvey, and L. Baillie, "Using thermal stimuli to influence affect in different picture display sizes," *Pers Ubiquit Comput*, vol. 21, no. 4, pp. 739–759, Aug. 2017, doi: 10.1007/s00779-017-1018-0.

[51] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum Entropy Inverse Reinforcement Learning," in *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, 2008, p. 6.

[52] N. A. Malik, F. A. Hanapiah, R. A. A. Rahman, and H. Yussof, "Emergence of Socially Assistive Robotics in Rehabilitation for Children with Cerebral Palsy: A Review," *International Journal of Advanced Robotic Systems*, vol. 13, no. 3, 2016, doi: 10.5772/64163.

[53] M. P. McGlinchey and S. Davenport, "Exploring the decision-making process in the delivery of physiotherapy in a stroke unit," *Disability and Rehabilitation*, vol. 37, no. 14, pp. 1277–1284, Jul. 2015, doi: 10.3109/09638288.2014.962106.