

WIRATUNGA, N., CORSAR, D., MARTIN, K., WIJEKOON, A., ELYAN, E., COOPER, K., IBRAHIM, Z., CELIKTUTAN, O., DOBSON, R.J., MCKENNA, S., MORRIS, J., WALLER, A., ABD-ALHAMMED, R., QAHWAJI, R. and CHAUDHURI, R. 2019. Learning to self-manage by intelligent monitoring, prediction and intervention. In Wiratunga, N., Coenen, F. and Sani, S. (eds.) *Proceedings of the 4th International workshop on knowledge discovery in healthcare data (KDH 2019), co-located with the 28th International joint conference on artificial intelligence (IJCAI-19)*, 10-11 August 2019, Macao, China. CEUR workshop proceedings, 2429. Aachen: CEUR-WS [online], pages 60-67. Available from: <http://ceur-ws.org/Vol-2429/paper10.pdf>

Learning to self-manage by intelligent monitoring, prediction and intervention.

WIRATUNGA, N., CORSAR, D., MARTIN, K., WIJEKOON, A., ELYAN, E., COOPER, K., IBRAHIM, Z., CELIKTUTAN, O., DOBSON, R.J., MCKENNA, S., MORRIS, J., WALLER, A., ABD-ALHAMMED, R., QAHWAJI, R. and CHAUDHURI, R.

2019

Learning to Self-Manage by Intelligent Monitoring, Prediction and Intervention

Nirmalie Wiratunga¹, David Corsar¹, Kyle Martin¹, Anjana Wijekoon¹, Eyad Elyan¹, Kay Cooper¹,
Zina Ibrahim^{2,3}, Oya Celiktutan², Richard J. Dobson^{2,3},
Stephen McKenna⁴, Jacqui Morris⁴, Annalu Waller⁴,
Raed Abd-Alhammed⁵, Rami Qahwaji⁵,
Ray Chaudhuri⁶

¹ Robert Gordon University, Aberdeen, UK

² King's College London, London, UK

³ University College London, London, UK

⁴ University of Dundee, Dundee, UK

⁵ University of Bradford, Bradford, UK

⁶ King's College Hospital, NHS Foundation Trust, UK

{n.wiratunga, d.corsar, k.martin, a.wijekoon, e.elyan, k.cooper}@rgu.ac.uk,

{zina.ibrahim, oya.celiktutan, richard.j.dobson}@kcl.ac.uk,

{s.j.z.mckenna, j.y.morris, a.waller}@dundee.ac.uk,

{r.a.a.abd, r.s.r.qahwaji}@bradford.ac.uk,

ray.chaudhuri@nhs.net

Abstract

Despite the growing prevalence of multimorbidities, current digital self-management approaches still prioritise single conditions. The future of out-of-hospital care requires researchers to expand their horizons; integrated assistive technologies should enable people to live their life well regardless of their chronic conditions. Yet, many of the current digital self-management technologies are not equipped to handle this problem. In this position paper, we suggest the solution for these issues is a model-aware and data-agnostic platform formed on the basis of a tailored self-management plan and three integral concepts - Monitoring (M) multiple information sources to empower Predictions (P) and trigger intelligent Interventions (I). Here we present our ideas for the formation of such a platform, and its potential impact on quality of life for sufferers of chronic conditions.

Although the number of individuals living with multiple morbidities is predicted to increase significantly over the coming years [Barnett *et al.*, 2012], current self-management solutions prioritise single conditions. A 2018 study from the UK National Institute of Health Research (NIHR) predicted that two-thirds of people aged 65 and over will have multiple morbidities by 2035, and 17% with four or more conditions. One third of these people will have a mental illness (e.g. dementia or depression). Increased life expectancy for both men and women means people will spend a longer time living with multiple morbidities, placing increased demand on the healthcare system.

Integrated assistive technologies promises a to enable people to live their life well regardless of their chronic conditions. Advances in telecommunications and Artificial Intelligence (AI) technologies paves the way for personalised virtual health companions that provide a intelligently-on connection between the patient and those providing their care. Such companions should be an intermediary, supporting patients with confidently managing their condition(s), proactively engaging with health care professions only when necessary, reducing their workload and replacing the current reactive patient-clinician interaction. This can be achieved by reasoning with data from automated observations of patient's progress along their healthcare plan using real-time predictions to trigger appropriate proactive interventions. Chronic patients have to live with their conditions 24/7, so it is natural that their care should reflect that.

This position paper presents our framework for multimorbidity virtual health companions, each tailored to the unique health needs of individuals, assisting them to take

1 Introduction

Chronic health conditions currently incur over 80% of all healthcare spending in the United Kingdom. Living with one or more chronic illnesses almost certainly means major changes in one's life; the latter can be minimised with effective self-management. Studies show that the capability to self-manage (chronic) health conditions effectively promises lower associated healthcare costs and more efficient use of primary and secondary care [Wolff *et al.*, 2002].

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

an assured, active role in managing their health. The framework is based on a configurable architecture comprising multiple reasoning components for *Monitoring, Prediction and Intervention* (MPI). This will provide a plug and play model enabling the bespoke integration of existing and yet-to-be-created devices, along with modes of reasoning as necessary in a library of AI skills. For example, humanoid-robot driven conversational dialogue systems, autonomous image and video analysis, analysis of real-time wearable and implant-generated data, and advanced telecommunication networks (e.g. 5G) for remote interactions.

This paper is structured as follows. In Section 2 we discuss related work within digital self-management. In Section 3 we present our concepts on requirements for a generic framework with emphasis on reasoning centered around Monitoring, Prediction and Intervention (MPI) components, and detail how these can be expanded to cover multiple morbidities. In Section 4 we describe several technologies which we expect to be key players in the future and explore their integration within a data agnostic framework. Finally, in Section 5 we provide some conclusions.

2 Related Work

Self-management is a set of approaches which aim to enable people living with long-term conditions to take control of their care and manage their own health. Assistive technology can support self-management on several levels:

1. Problem-solving (e.g. coping with flare-ups or adapting plan activities);
2. Decision-making (e.g. when to seek support or help with decisions around positive behaviour changes like improving diet, reducing alcohol consumption, quitting smoking or increasing physical activity and social interaction);
3. Resource utilisation (e.g. making best use of healthcare and other resources, including 3rd sector, peer-support, web-based resources and other sources of information or advice);
4. Forming patient-healthcare provider relationships and encouraging patients to interact with their healthcare provider appropriately. This would ideally occur before emergency or crisis situations arise to prevent decline in health and/or hospital admission.
5. Action planning and self-tailoring (e.g. encouraging patient participation in creating their own self-management plan (which might include physical activity, specific exercises, relaxation) and tailoring it to their specific needs, improving patient's knowledge of their conditions).

The goal of self-management is to encourage behaviour change in sufferers of chronic conditions. A systematic review of interventions to promote physical activity [Morris *et al.*, 2014] illustrated that interventions involving behaviour change strategies are more effective for sustaining longer-term physically active lifestyles than time-limited interventions involving structured exercises alone.

These interventions are commonly delivered face to face by healthcare practitioners. However current studies indicate that healthcare time is extremely limited and of short duration. Without ongoing support, patient physical activity levels decline as maintaining motivation is difficult [Morris *et al.*, 2012]. Innovative, person-centred strategies to monitor and predict physical activity and exercise behaviours, to scan and anticipate environmental barriers to activity, and to provide social and motivation support are required. These must support evidence-based, personally tailored behaviour change strategies by monitoring and providing feedback on performance; provide virtual real-time social support for activity; provide feedback on physical performance and evaluation of environmental barriers to physical activity.

2.1 Case-based reasoning for self-management

Previous work has demonstrated the effectiveness of applying decision support and reasoning systems to the management of a specific chronic disease. For instance Case-based reasoning (CBR) which is an AI approach that solves new problems using specific knowledge extracted from previously solved problems, has been successfully used to incorporate evidence-base practices. Here, reasoning is facilitated by a collection of cases, a unique set of past experiences stored in a case base. However to the best of our knowledge CBR has only been applied in self-management of single chronic diseases.

CBR has been applied to managing diabetes types 1 and 2, using records that provide details about periodical visits with a physician in a case consisting of features that represent a problem (e.g. weight, blood glucose level), its solution (e.g. levels of insulin) and the outcome (e.g. hyper/hypo(glycemia)) observed after applying the solution [Marling *et al.*, 2012; Montani *et al.*, 2000]. More recent work [Chen *et al.*, 2017], explored the management of diabetes type 1 to support monitoring of blood glucose levels before, during and after exercises. Interventions recommend carbohydrate intake based on similar cases retrieved from the case base. In related work on self-management of low-back pain (LBP) [Bach *et al.*, 2016], CBR recommends care plans from similar patients. Management involves a human activity recognition (HAR) component to monitor the patient activity using sensor data that is continuously polled from a wearable device. Patient reported monitoring is used by the SelfBACK system to manage exercise adherence. Monitoring allows the system to detect periods of low activity behaviour, at which point a notification is generated to nudge the user to be more active - the intervention. An important contribution of this work is the integration of behaviour change techniques such as goal setting to focus the expected level of activity. Thereafter comparison of expected and actual behaviours analyse goal achievement.

Evidence for self-management of patients with multimorbidities is limited, despite the prevalence of co-occurring conditions and its impact on patients and healthcare systems [Smith *et al.*, 2012]. Interventions tend to have mixed effects requiring careful design underpinned by evidence-based practice. Personalisation is important to ensure that care plans are tailored to the needs of the individual. Al-

though there has been recent work on personalised learning using state-of-the-art learning architectures (e.g. matching networks) more work is needed when applying them to individuals with multimorbidities [Sani *et al.*, 2018].

2.2 Pervasive and ubiquitous self-management

Pervasive and ubiquitous AI enabled devices are arguably best placed to continuously monitor a person's adherence to self-management plans, make real-time predictions about the likelihood of adherence and the impact of that. However intervention requires a good understanding of human behaviours and direct inspection by health providers, which although valuable, cannot be scaled to large and diverse groups of people. Wearables, such as smart watches or phones, are the most common form of physical activity monitoring devices and sources of delivering digital interventions. These are embedded with inertial measurement devices (e.g. accelerometers or gyroscopes) that generate time-series data which can be exploited for human activity recognition of ambulatory activities, activities of daily living, gait analysis and pose recognition [Sani *et al.*, 2018; Reiss and Stricker, 2012; Chavarriaga *et al.*, 2013].

Although commercial wearable activity trackers such as Fit-Bit are increasingly being used to monitor levels of physical activity and provide feedback to users, their utility is limited. Accuracy in determining activity in people who walk at slow ambulatory speeds in free living conditions is low [Feehan *et al.*, 2018], and evidence of effects on physical activity levels are uncertain [Lynch *et al.*, 2018]. They have limited interactive and personalisation options, due in part to a reliance upon text notifications as the intervention method. As a direct result, current wearable technologies struggle to understand and address individual barriers to promoting behaviour change in individuals with complex disabilities, provide insufficient information to determine specific rehabilitation activities (such as exercises), and are not adapted to people with communication impairments. Importantly, provision of real-time information on outdoor environmental hazards (e.g. stairways) and mental health issues (such as anxiety or lack of confidence) which are likely to impact physical activity behaviours in patients is limited.

Another promising solution is to develop and employ social robots that can be designed to offer psychological support. In this context, Socially Assistive Robotics (SAR) is a rapidly growing domain that aims to enhance psychological well-being through human interactions with a robot. Robots can be programmed to perceive and interpret human actions and nonverbal cues, and provide assistance both at the level of goal setting and tracking and socio-emotional communication through personalised conversational dialogues. The usefulness of social robots in mental healthcare contexts has been investigated by a number of previous works [Rabbitt *et al.*, 2015]. However, most of the available therapeutic robotic platforms target supporting either children with special conditions such as autism or assisting elderly people in their daily lives. In particular, robotic platforms for patient education and self-management interventions are still scarce. There are a few lines of work focusing on self-management and awareness of type 1 diabetes in children [Kruijff-Korbayová *et al.*,

2014; van der Drift *et al.*, 2014] and motivation coaching for healthy living, weight loss and exercise [Leme *et al.*, 2019]. To the best of our knowledge, the application of social robots in supporting individuals suffering multiple conditions has remained an unexplored area.

Computer vision can be used to detect, track, and re-identify patients without the need for any specific sensors or markers to be worn or carried. In their place, technical requirements include the need for the computer to infer the location, pose and movement of the trainee (and other people in its vicinity). Person following [Honig *et al.*, 2018] is a key capability for the machine to be able to observe and guide a patient. For this purpose, unmanned aerial vehicles (commonly referred to as drones) may present a flexible solution; rather than fast-flying quad-copters, blimps may offer a more stable and safer platform [Yao *et al.*, 2019]. Their use in related applications such as monitoring older people in care homes has been suggested (but not yet developed) [Srisamosorn *et al.*, 2016]. However, the computer vision systems used need to be made more robust and reliable before a flying socially-aware robot for monitoring patients could be deployed and trialled, especially in less controlled environments outside the clinic.

A further potential solution in this field is Ambient Assisted Living (AAL), which targets the use of multiple devices and sensors around the home to support personal healthcare monitoring. AAL offers an opportunity for non-intrusive tracking of patient condition through smart home technology [Forbes *et al.*, 2019]. Though traditionally difficult to apply this for accurate patient monitoring (particularly in open areas), recent advancements demonstrate detailed activity profiling can be gained from non-intrusive RFID chips situated in locations around an individual's home [Oguntala *et al.*, 2019], even in large rooms [Obeidat *et al.*, 2019]. Though work has targeted AAL to support assisted living and fall prediction for the elderly [Massie *et al.*, 2018], we are unaware of any current AAL approaches which comprehensively support multimorbidities in every age group.

3 Generic self-management framework

Generic self-management frameworks need to be capable of covering a wide range of devices and conditions in order to support personalisation, if trained models are to move away from the current one-size-fits all systems driven by centralised datasets. Further they must ensure patient privacy and data security despite the large volumes of data needed for training machine learning models that will provision edge computing devices and intelligent decision support systems. Frameworks must also be flexible, able to react to changes in the environment and adapt reasoning appropriately. We argue the key to addressing these requirements is treating self-management as an AI planning problem; where AI methods support and recommend interventions centred around the likely achievement of goals and actions recorded in plans. To achieve this we need constructs that can be configured to a given self-management plan; and a generic architecture that can support the reuse of these constructs to enable evidence-based community care.

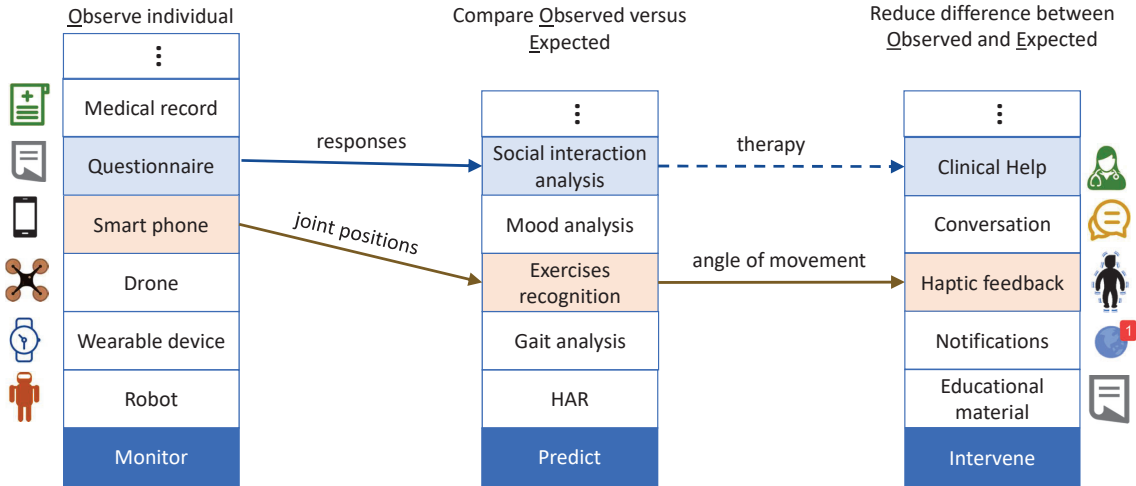


Figure 1: Instance of a blueprint for stroke rehabilitation configured with two Monitor-Predict-Intervene (MPI) component paths.

3.1 Reasoning with the MPI Cycle

We propose inclusion of three components: digital Monitoring (M) to track patient condition(s) and underpin anticipatory Predictions (P) to trigger real-time Interventions (I) for additional support when it is required. We refer to the combination of these components as an MPI. Reasoning is facilitated by a self-management plan consisting of guidelines, health recommendations, patient goals, decisions and a trace of previous lessons learned. These lessons learned data supports tailoring of plans to an individual. Given a self-management plan (e.g. level and frequency of exercise, targeted levels of anxiety), the MPI cycle monitors adherence by a combination of patient self-reported outcome measures (e.g. pain) and automated real-time monitoring (e.g. of physiological response) by a variety of existing and yet-to-be-created devices (e.g. wearables, implants, in-home sensors, drones, robots). Differences in observed and expected adherence, combined with environmental and personal factors can be used to predict likely trends and outcomes. Thereafter, autonomous and remotely supported proactive interventions by health professionals are initiated and plans are adapted collaboratively, replacing the current reactive patient-clinician interaction.

An evidence-based approach to self-management is facilitated by having access to previous plans and lessons learned as well as reusing care-plans from similar patients that have been successful in relation to outcome measures. The reasoning capability of the MPI increases with increasing adoption of the framework. As the richness of the evidence base grows (from initial guidelines to personalised plans) the impact on the community and their common self-management conditions will be improved.

Thus MPIs form the building blocks for a multi-faceted self-management plan - a plan that can cater for multimorbidities. For example, consider Figure 1, which features two MPIs created to support stroke rehabilitation. Stroke survivors commonly struggle with freedom of movement in their joints and have a tendency to develop depression dur-

ing treatment. In this image, a patient is managing their joint strengthening exercises using an exercise-MPI (consisting of the components shaded orange), while managing their mood through an emotion-MPI (shaded blue). The exercise-MPI monitors bending of the joint via machine vision technology and a smart phone camera. Reasoning on the monitored data allows the system to predict whether the exercises are being performed correctly (e.g. the angle of movement is satisfactory). The system can then intervene by actuating haptic feedback through a patient's wearable sensor, guiding the patient to perform the exercise correctly. Similarly, the emotion-MPI monitors mood by reasoning on the patient's responses to questionnaires. If mood is predicted to be below the threshold determined by a clinician, an intervention can organise clinical help before this evolves into depression. However if mood is above this threshold, no intervention is necessary. The goal of each MPI is to Observe (O) the patient through monitoring to predict a comparison with the Expected (E) outcome as established by their clinician. If the observed actions deviate sufficiently from the expected, then an intervention is necessary to minimise the difference ($\min(\delta(O, E))$).

3.2 Reusing self-management plans

We propose the idea of a *blueprint* through which an individual's self-management plan for multimorbidities can be formulated. Essentially a blueprint is a combination of one or more MPIs and their relationships (e.g. contraindications) which has been configured jointly with a clinician. As an example, the aforementioned Figure 1 shows a patient's blueprint for the self-management of stroke rehabilitation, as it combines the exercise-MPI and the emotion-MPI.

The reasoning necessary when combining multiple different MPIs is an important area of research; it must ensure adherence to a patient's collective self-management plans and suggest interventions that are suitable for all the multimorbidities, whilst maintaining knowledge of any contraindications between those conditions. Central to this is a knowledge structure, an MPINet ontology, where known MPI re-

relationships can be recorded and used to form generalised blueprints. These can be refined by co-occurrences inferred from collected data.

A shared community of blueprints is formed from individuals who share the same or a similar set of co-occurring conditions. Configurations embedded in blueprints can then be reused and adapted to suit new individuals joining that community (see Figure 2). In this image, blueprints have been extracted from members of the community who are similar to the new individual, where φ is a reuse function containing adaptation knowledge. The output is a blueprint configured to the individual’s needs founded on evidence-based practice. This can be formalised as:

$$O', E' = \varphi[(O, E)^1, (O, E)^2, \dots] \quad (1)$$

where O' the configuration for what is to be observed and E' represents the modified expectations in this new blueprint.




Figure 2: Reuse of blueprints from similar individuals.

Personalisation of blueprints is achieved through the recommendation of contextually-relevant MPIs and customising general community blueprints. Similarities in self-management goals within a community can be used for personalisation and to anticipate known common complications of a condition, which can be extended to enable forecasting for healthcare service demand. Metric learning algorithms that are suited for evidence-based reasoning lend well to learning personalised models on the basis of similarity computations, and also lend themselves well to few-shot learning approaches. Recent advances in similarity driven attention mechanisms currently use simplistic metrics, hence in future research we must establish how similarity on the basis of the user contexts can be computed in a differentiating manner.

Configuration of blueprints shares many similarities with On-Line Case-Based Planning (OLCBP). Both require continuous monitoring of expected versus observed and acting upon deviations in a timely manner. Much like the snippets that exist within a plan [Ontonón *et al.*, 2010], MPIs act as the basic constituent pieces of a blueprint and are configured with individual goals. Though failure to achieve individual MPI goals does not necessarily define the failure of a self-management plan, we suggest this is an early warning of the need for additional care. Importantly recording such failures and the follow-on adaptation of both the self-management goal and agreed activities are useful experiential content that are crucial when developing evidence-based practices for the community.

3.3 Role of federated machine learning

Ideally, MPI models will learn from the data generated by all users of the system, i.e. in a distributed manner. This requires us to adopt ideas from Federated Learning (FL) [McMahan *et al.*, 2017] and meta-learning. Specifically maintaining control of device data, where training involves a shared global

model under the coordination of a central server acting as a curator, selecting which participating devices (i.e. the federation) to incorporate when training models. This form of learning is ideal for community health care ensuring privacy by default, respecting data ownership, and maintaining locality of data (without centralising data) for application deployment at scale. An interesting direction of research here is to the use of data provenance records combined with metrics evaluating quality and trust of individuals to influence the global computational curator’s decisions on sampling of MPIs on devices. With complex model architectures there is also a need to share and describe the architectural properties so as to inform the curator about compatibility. This perhaps calls for a meta-language for architectural descriptors. Learning from few labelled data is important for technology to operate at scale. It will be useful to extend the FL paradigm to evidence-based reasoning methods such as matching networks [Vinyals *et al.*, 2016] to enable few-shot learning while using a federated strategy.

This type of environment is potentially suitable for implementing a multi-agent framework [Moreno A, 2003], whereby autonomous, adaptive and interactive software components, provide notions that specifically meet the MPI challenges to form the base of a robust and scalable self-management infrastructure. A multi-agent framework will also enable the collection of patient-reported well-being information (as done in [Ibrahim *et al.*, 2015]), integrating them with sensor and device-generated data via API-based real-time data collection and streaming platforms such as the in-house build RADAR-base platform [Ranjan *et al.*, forthcoming 2019]. The combination will yield an intelligent and real-time framework for schematised, secure and role-based data collection, harmonisation and integration based on unified temporal intervals from all devices, sensors and individuals.

4 The Future of Digital Self-Management

Tackling self-management of multimorbidities requires a model-aware and data-agnostic machine learning platform - a modular digital health ecosystem, where multiple monitoring technologies enable comprehensive predictions that trigger the most appropriate intervention strategy from a broad range of possibilities. With that in mind, we suggest the proposed MPI architecture could answer many of the research challenges discussed for individual technologies in Section 2 and greatly improve the self-management experience from a user perspective. However, there is also a need to consider how this architecture could be practically utilised at scale while preserving patient privacy.

4.1 Wearables and Haptic Feedback

We aim to develop a framework to allow wearable technologies to provide real-time haptic feedback to support users effectively perform rehabilitation exercises and physical activities. Haptic feedback uses digital cues to stimulate the feeling of touch in users. Whereas current wearable technologies are reliant on the user reacting to textual notifications or messages, haptic feedback provides an opportunity to correct execution of exercises on-the-fly. The result would be

a decreased reliance on the external technology of a smart phone and less expectation for users to continually refer to their phone throughout an exercise session.

There are two primary challenges to provisioning haptic feedback remotely; (1) currently, haptic feedback for therapy is performed one-to-one, either with a healthcare professional activating the necessary feedback while locally observing the patient or the skeleton following a pre-programmed set of instructions; and (2) current telecommunication networks do not have the throughput required to enable such feedback remotely at scale. We believe that the former research challenge could be answered through development of a learned model which can provide activation of the necessary components to actuate haptic feedback. The latter challenge may be answered with effective advanced networking infrastructures, such as 5G.

4.2 Robots and Conversational Dialogue

The use of therapeutic robots to support community-dwelling individuals participate in rehabilitation, fitness exercises, social engagement and outdoor physical activities are not only a viable solution to the problem of shortage in care resources, but also potentially address loneliness in the elderly. However, current SAR systems are limited in their capability to offer truly personalised support. They do not take the user state into consideration, beyond verbal and interaction logs, suffer from low engagement, and uni-directional information flow. Expanding this to exchange more information between the user and the system through multiple modalities will enable a more versatile and natural interaction, ensuring that the user maximally benefits from the system. In addition, it provides a means to build a richer user-context that can help to improve, structure and personalise a robot's behaviours and conversational dialogue. To address these issues, novel methods are needed for semantically integrating various data sources relating to (1) mood evaluations based on visual cues (i.e., facial, bodily cues); (2) physical performance, psychological variables and environmental variables collected from wearables; and (3) adaptive activity programmes and behaviour change strategies in real-time. Integrated information can be then used to build effective reasoning and intervention mechanisms for the interactive robot companion that can support citizens to lead active and social lives.

4.3 Autonomous Computer Vision

Computer-vision models can provide visual feedback to support patients with rehabilitation activities. This involves the analysis of large volumes of imagery data, recognition of fine-grained human movement details and generation of visual feedback interventions. Drones could offer a rich source of information for on-going monitoring and observation of patients in different contexts, including both indoor self-management activities and outdoor exercises. Utilising a visual source that captures patient's movement and behaviours from different angles provides a unique opportunity for learning and improving practices, but at the same time poses some key challenges. These are due to the inherent vision problem, in particular, performing accurate detection, tracking and classification under different poses and light conditions,

which can be exacerbated by the freedom of movement afforded to a drone. Additionally, the anticipated skewed distribution of the data will pose another challenge that needs to be addressed. In particular, having enough representative training examples that captures the different scenarios of a specific case study or patient can potentially be very difficult, meaning that few-shot learning strategies are likely to be very relevant here. The benefits for patients go beyond support with performing exercises and documenting their self-management progress; for example, people with movement difficulties are often anxious about walking in unfamiliar environments as they are unaware of any obstacles they may face - this can lead them to withdrawing from social activities and an increased feeling of loneliness. Being accompanied by a drone that advises them of upcoming obstacles and how they can be avoided, may reduce anxiety associated with such situations.

4.4 Ambient Assisted Living

We aim to use Ambient Assisted Living (AAL) technology to support non-intrusive monitoring of patients within their own home. Several interesting research challenges arise from this field. In particular, it will be interesting to see how the constant monitoring of AAL devices allows the identification of self-management activities as they take place within aspects of daily living. For example, if a patient has climbed the stairs 5 times in an afternoon as part of cleaning their home, this action may cover some of the ambulatory physical activity scheduled as part of managing their lower-back pain condition. Understanding the situations in which daily living activities can replace self-management exercises (or alternatively, the situations in which they cannot) is an intriguing avenue for exploration. Furthermore, there is the technical challenge of being able to seamlessly change monitoring devices that feed into a single MPI e.g. as the patient transitions from one room to the next. A comprehensive self-management solution involving AAL devices should maximise comfort of the patient and ensure that they are able to relax within their home without feeling like they are being continually observed. Developing an individual's trust requires further understanding of human machine interfaces by bringing together behavioural psychologists and computer scientists.

4.5 A Comprehensive Self-Management Solution

The distributed MPI model comprises a number of devices collaborating autonomously, and a substantial amount of real-time data generated passively by sensors and devices or actively reported by the users (e.g. well-being reports) for a single non-fragmented self-management solution. As such, the MPI framework constitutes a decentralised, heterogeneous and open environment that operates on multiple computing systems to manipulate, share and analyse strictly-confidential patient records. Moreover, to build a platform to continuously inform patients using their own records, requires mechanisms to mine the 'big' and heterogeneous data for relevant knowledge, and learn the adapted and personalised recommendations and possible interventions in real-time to aid in self-management. These methods are underpinned by the growing availability of advanced networking technology (e.g.

5G), where the faster connection speeds promise huge benefits to future healthcare systems. Such advances allow high-intensity data processing to be supported by near real-time decision-making. We believe these will play a crucial role in supporting at-home care (e.g. wearable sensors applications and online consultations) and remote treatments (e.g. digital diagnosis). In particular, 5G will facilitate the transfer of expertise over a great distance in real-time using technologies such as robotics and haptic feedback.

To give an example, a stroke survivor patient may use AAL sensors for non-intrusive tracking of physical activity and interact with a SAR to self-report on their pain and mental condition while in the home. When journeying outside the house, the patient initially uses a drone with computer vision capabilities to monitor their journey, before transitioning to use of a wearable technology as their walking ability improves. The combination of monitoring technologies to empower interventions is a comprehensive solution to their self-management. We believe that not only will this greatly improve adherence to the various aspects of a self-management plan, it will also have an irreplicable effect on patient quality of life. Furthermore, integration with the MPI framework will bring AI power to off-the-shelf hardware rather than building expensive devices, thus encouraging socially-inclusive care.

5 Conclusion

In conclusion, in this position paper we have discussed the need for an affordable comprehensive self-management solution to improve quality of life for patients living with multiple morbidities. We have presented our ideas towards the creation of MPI components - constituent pieces of wider-scale framework which allows a configurable and personalised self-management plan for individual patients, while encouraging reuse plans from within a community of similar patients and morbidities. The platform and MPIs potentially offer several interesting avenues of research, not only around how new and existing technologies can be integrated, but also around areas such as comparing the similarity of MPIs and understanding the ramifications of using them for reuse and adaptation. Most importantly, we believe that a wide-scale integrated framework of devices is necessary for supporting patients with confidently managing their condition(s), improving their quality of life in the future.

References

- [Bach *et al.*, 2016] Kerstin Bach, Tomasz Szczepanski, Agnar Aamodt, Odd Erik Gundersen, and Paul Jarle Mork. Case representation and similarity assessment in the self-back decision support system. In Ashok Goel, M Belén Díaz-Agudo, and Thomas Roth-Berghofer, editors, *Case-Based Reasoning Research and Development*, pages 32–46, Cham, 2016. Springer International Publishing.
- [Barnett *et al.*, 2012] Karen Barnett, Stewart W Mercer, Michael Norbury, Graham Watt, Sally Wyke, and Bruce Guthrie. Epidemiology of multimorbidity and implications for health care, research, and medical education: a cross-sectional study. *The Lancet*, 380(9836):37–43, 2012.
- [Chavarriaga *et al.*, 2013] Ricardo Chavarriaga, Hesam Sagma, Alberto Calatroni, Sundara Tejaswi Digumarti, Gerhard Tröster, José del R Millán, and Daniel Roggen. The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters*, 34(15):2033–2042, 2013.
- [Chen *et al.*, 2017] Yoke Yie Chen, Nirmalie Wiratunga, Stewart Massie, Jenny M Hall, Kate Stephen, Amanda Croall, Jacky MacMillan, Lesley Murray, Geoff Wilcock, and Sandra M MacRury. Designing a personalised case-based recommender system for mobile self-management of diabetes during exercise. In *In Workshop proceedings of the 2nd Knowledge Discovery from Health Workshop at the International Joint Conference on AI*, 2017.
- [Feehan *et al.*, 2018] Lynne M Feehan, Jasmina Geldman, Eric C Sayre, Chance Park, Allison M Ezzat, Ju Young Yoo, Clayon B Hamilton, and Linda C Li. Accuracy of fitbit devices: Systematic review and narrative syntheses of quantitative data. *JMIR mHealth and uHealth*, 6(8):e10527, 2018.
- [Forbes *et al.*, 2019] Glenn Forbes, Stewart Massie, and Susan Craw. Fall prediction using behavioural modelling from sensor data in smart homes. *Artificial Intelligence Review*, pages 1–21, 2019.
- [Honig *et al.*, 2018] S. S. Honig, T. Oron-Gilad, H. Zaichyk, V. Sarne-Fleischmann, S. Olatunji, and Y. Edan. Toward socially aware person-following robots. *IEEE Transactions on Cognitive and Developmental Systems*, 10(4):936–954, Dec 2018.
- [Ibrahim *et al.*, 2015] Zina Ibrahim, Lorena Fernández de la Cruz, Argyris Stringaris, Robert Goodman, Michael Luck, and Richard Dobson. A multi-agent platform for automating the collection of patient-provided clinical feedback. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, page 831–839, 2015.
- [Kruijff-Korbayová *et al.*, 2014] I. Kruijff-Korbayová, E. Oleari, I. Baroni, B. Kiefer, M. C. Zelati, C. Pozzi, and A. Sanna. Effects of off-activity talk in human-robot interaction with diabetic children. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 649–654, Aug 2014.
- [Leme *et al.*, 2019] Bruno Leme, Masakazu Hirokawa, Hideki Kadone, and Kenji Suzuki. A socially assistive mobile platform for weight-support in gait training. *International Journal of Social Robotics*, Apr 2019.
- [Lynch *et al.*, 2018] Elizabeth A Lynch, Taryn M Jones, Dawn B Simpson, Natalie A Fini, Suzanne S Kuys, Karen Borschmann, Sharon Kramer, Liam Johnson, Michele L Callisaya, Niruthikha Mahendran, et al. Activity monitors for increasing physical activity in adult stroke survivors. *Cochrane Database of Systematic Reviews*, (7), 2018.
- [Marling *et al.*, 2012] Cindy Marling, Matthew Wiley, Razvan Bunescu, Jay Shubrook, and Frank Schwartz. Emerging applications for intelligent diabetes management. *AI Magazine*, 33(2):67–67, 2012.

- [Massie *et al.*, 2018] Stewart Massie, Glenn Forbes, Susan Craw, Lucy Fraser, and Graeme Hamilton. Fitsense: Employing multi-modal sensors in smart homes to predict falls. In *International Conference on Case-Based Reasoning*, pages 249–263. Springer, 2018.
- [McMahan *et al.*, 2017] H. B. McMahan, E. Moore, Ramage, S. D., Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1273 – 1282, 2017.
- [Montani *et al.*, 2000] Stefania Montani, Riccardo Bellazzi, Luigi Portinale, Giuseppe d’Annunzio, Stefano Fiocchi, and Mario Stefanelli. Diabetic patients management exploiting case-based reasoning techniques. *Computer methods and programs in biomedicine*, 62(3):205–218, 2000.
- [Moreno A, 2003] Garbay C. Moreno A. Software agents in health care. *Artificial intelligence in medicine*, 27(3):229–232, 2003.
- [Morris *et al.*, 2012] Jacqui Morris, Tracey Oliver, Thilo Kroll, and Steve MacGillivray. The importance of psychological and social factors in influencing the uptake and maintenance of physical activity after stroke: a structured review of the empirical literature. *Stroke research and treatment*, 2012, 2012.
- [Morris *et al.*, 2014] Jacqui H Morris, Stephen MacGillivray, and Sarah Mcfarlane. Interventions to promote long-term participation in physical activity after stroke: a systematic review of the literature. *Archives of physical medicine and rehabilitation*, 95(5):956–967, 2014.
- [Obeidat *et al.*, 2019] Huthaifa Obeidat, Ali AlAbdullah, Nazar Ali, Rameez Asif, Omar Obeidat, Mohammed Bin-Melha, Wafa Shauib, Raed Abd-Alhameed, and Peter Excell. Local average signal strength estimation for indoor multipath propagation. *Accepted to IEEE Access*, April 2019.
- [Oguntala *et al.*, 2019] George A. Oguntala, Raed A. Abd-Alhameed, James M. Noras, Yim-Fun Hu, Eya N. Nnabuike, Nazar Ali, Issa T. Elfergani, and Jonathan Rodriguez. Smartwall: Novel rfid-enabled ambient human activity recognition using machine learning for unobtrusive health monitoring. *Accepted to IEEE Access*, May 2019.
- [Ontonón *et al.*, 2010] Santi Ontanón, Kinshuk Mishra, Neha Sugandh, and Ashwin Ram. On-line case-based planning. *Computational Intelligence*, 26(1):84–119, 2010.
- [Rabbitt *et al.*, 2015] Sarah M. Rabbitt, Alan E. Kazdin, and Brian Scassellati. Integrating socially assistive robotics into mental healthcare interventions: Applications and recommendations for expanded use. *Clinical Psychology Review*, 35:35 – 46, 2015.
- [Ranjan *et al.*, forthcoming 2019] Yatharth Ranjan, Zulqarnain Rashid, Callum Stewart, Mark Begale, Denny Verbeeck, Sebastian Boettcher, Pauline Conde, Richard Dobson, and Amos Folarin. RADAR-base: An open source mhealth platform for collecting, monitoring and analyzing data using sensors, wearables, and mobile devices. *Journal of Medical Informatics Research mHealth and uHealth*, forthcoming, 2019.
- [Reiss and Stricker, 2012] Attila Reiss and Didier Stricker. Introducing a new benchmarked dataset for activity monitoring. In *Wearable Computers (ISWC), 2012 16th International Symposium on*, pages 108–109. IEEE, 2012.
- [Sani *et al.*, 2018] Sadiq Sani, Nirmalie Wiratunga, Stewart Massie, and Kay Cooper. Personalised human activity recognition using matching networks. In Michael T. Cox, Peter Funk, and Shahina Begum, editors, *Case-Based Reasoning Research and Development*, pages 339–353. Springer International Publishing, 2018.
- [Smith *et al.*, 2012] Susan M Smith, Hassan Soubhi, Martin Fortin, Catherine Hudon, and Tom O’Dowd. Managing patients with multimorbidity: systematic review of interventions in primary care and community settings. *BMJ*, 345, 2012.
- [Srisamosorn *et al.*, 2016] Veerachart Srisamosorn, Noriaki Kuwahara, Atsushi Yamashita, Taiki Ogata, and Jun Ota. Design of face tracking system using environmental cameras and flying robot for evaluation of health care. In *Digital Human Modeling: Applications in Health, Safety, Ergonomics and Risk Management*, LNCS. Springer, 2016.
- [van der Drift *et al.*, 2014] Esther J.G. van der Drift, Robbert-Jan Beun, Rosemarijn Looije, Olivier A. Blanson Henkemans, and Mark A. Neerincx. A remote social robot to motivate and support diabetic children in keeping a diary. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction, HRI ’14*, pages 463–470, New York, NY, USA, 2014. ACM.
- [Vinyals *et al.*, 2016] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, and Daan Wierstra. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [Wolff *et al.*, 2002] Jennifer L. Wolff, Barbara Starfield, and Gerard Anderson. Prevalence, Expenditures, and Complications of Multiple Chronic Conditions in the Elderly. *JAMA Internal Medicine*, 162(20):2269–2276, 11 2002.
- [Yao *et al.*, 2019] Ning-shi Yao, Qiu-yang Tao, Wei-yu Liu, Zhen Liu, Ye Tian, Pei-yu Wang, Timothy Li, and Fumin Zhang. Autonomous flying blimp interaction with human in an indoor space. *Frontiers of Information Technology & Electronic Engineering*, 20(1):45–59, Jan 2019.