

Probabilistic Models versus Discriminate Classifiers for Human Activity Recognition with an Instrumented Mobility-Assistance Aid

Mitesh Patel, Rami Khushaba, Jaime Valls Miro, Gamini Dissanayake

ARC Centre of Excellence for Autonomous Systems

Faculty of Engineering and IT, University of Technology Sydney

NSW 2007, Australia

{miteshkumar.n.patel, rami.khushaba, javalls, gdissa}@eng.uts.edu.au

Abstract

Detection of individuals' intentions and actions from a stream of human behaviour is an open and complex problem. There is however an intrinsic need to automatically recognise the activities performed by users of mobility assistive aids to better understand their behavioural patterns, with the ultimate objective of improving the utility of these devices. While discriminative algorithms such as Support Vector Machines (SVM) are well understood, generative probabilistic approaches to machine learning such as Dynamic Bayesian Networks (DBN) have only recently started gaining increasing interest within the robotics community. In this paper, a comprehensive evaluation of these techniques is carried out for human activity recognition in the context of their applicability to assistive robotics. Results show comparable recognition rates, offering valuable insights into the advantageous characteristics of DBN in relation to their dynamic and unsupervised nature for realistic human-robot interaction modelling.

1 Introduction

There is considerable interest within the artificial intelligence (AI) and robotics communities in how a single agent can autonomously make decisions in large, uncertain domains and even more so when human users are to work in close collaboration with these robotic agents. In this ever growing field of human-robot interaction (HRI), the use of modern AI techniques has become widespread as a mechanism to understand human intention. They have the potential to allow in designing more flexible robotic systems, as they are seen as capable of synthesising complex behaviours normally associated with human intelligence. One of the most relevant goals of an "intelligent" HRI system is to successfully execute the

explicit commands given by the user, while at the same time account for the implicit cues that are not so easily observed [Schrempf and Hanebeck, 2005].

Intention recognition can be described as the process of becoming aware of the intention of an agent (the user in our case), or as the problem of inferring an agent's intention through its actions and the associated effects on the environment [Heinze, 2003]. Intention recognition becomes an even more complex phenomenon when multiple agents try to interact with each other to achieve a common goal. Due to noisy and partial observations the communication of intentions between the agents becomes an important issue. Intention recognition has found its application in many research areas such as user assistance in aviation monitoring [Heinze, 2003], traffic monitoring [Pynadath, 1999], human robot co-operation [Schrempf and Hanebeck, 2005] and many other applications.

2 Motivation

Improving the quality of life of people suffering from impaired mobility is one of the key concerns for health care provision. Research indicates that, by 2040, the number of elderly people in the industrialised world will increase by 50%, and those reaching the age of 85 and up will rise by 100% [Begovich, 2005]. Assistive technology is increasingly finding its usage in this area to offset the impact of impairments resulting from aging process, injury and related disorders. Typically, these technologies have focused on assisting users with mobility impairments (mainly with wheelchairs, walker and canes). Increasing growth in the number of people with motor disabilities has resulted in considerable research being conducted into developing robotic assistive technologies to address the difficulties that this population faces, and several intelligent systems have been developed with the older adult population in mind. These include for instance the Nursebot project [Pineau *et al.*, 2003], the PAMM project [Dubowsky *et al.*, 2000], the COACH project [Hoey *et al.*, 2007], and the UTS Assis-

tive Walker project [Miró *et al.*, 2009]. Researchers in these projects have explored different avenues to develop robotic assistive technologies, using a variety of probabilistic and stochastic models amongst others.

3 Previous Work & Proposition

A significant amount of work has been done on intention recognition in the context of assistive systems. Hirata *et al.* [Hirata *et al.*, 2006] instrumented a passive walker with sensors and actuators. Their work was limited to recognising three states: walking, stopping and emergency (including falling), which were inferred based on the distance between the user and the walker (measured by a laser range finder), and the velocity of the walker. Morris *et al.* [Morris *et al.*, 2003] developed a mobility assistant device at CMU incorporating modules for obstacle avoidance, localisation, mapping, path planning and people tracking. Graf and Hgele [Graf and Hgele, 2001] designed a mobility assistant, "Care-O-Bot", as a part of a larger home care project for the elderly at the Fraunhofer Institute. The mobility assistant had two modes of operation for navigation: a direct user control where the robot took readings from a user intent sensor to determine the direction and speed of travel, and a target mode which allowed users to input a destination based on a map, with the robot guiding them to the destination in a reactive manner along the calculated route. Omar *et al.* [Omar *et al.*, 2010] proposed an activity recognition technique based on Hidden Markov Models (HMMs) and Conditional Random Fields (CRFs) for an instrumented passive rollator walker. The model recognised a number of user states: not touching the walker, stop/standing, walking forward, turn left, turn right, walking backwards and transfers (sit to stand/stand to sit). Miro *et al.* [Miró *et al.*, 2009] used Partially Observable Markov Decision Process Models (POMDP) as the probabilistic framework to recognise a similar set of short term user's intended behaviours. While successfully implemented, POMDP were shown to be computationally expensive, and a similar problem definition was later realised using Dynamic Bayesian Networks (DBN) [Patel *et al.*, 2010].

While various techniques have been proposed in these works for activity recognition in the context of HRI mobility assistive tasks, they are often limited in scope and tailored to some specific (learning) technique. The emphasis in this paper, on the other hand, is in presenting a critical appraisal of the two most distinctive AI approaches for learning and clustering, i.e. probabilistic Dynamic Bayesian Network (DBN) models and statistical Support Vector Machine (SVM) classifiers, for human-driven activity recognition. This is in the context of realistic assistive robotic settings, and based on an expectation of minimum indicative inputs from the user, with no explicit control interfaces, like voice, switches or



Figure 1: Rear view of the instrumented power rollator walker.

other forms of shared-control.

DBNs are powerful yet convenient tools for modelling dynamic systems. They allow modelling a system using graph-theoretic representations, and to integrate the various noisy observations together into a single consistent probabilistic framework. They have proved powerful tools in robotic areas such as perception, SLAM and scene analysis. SVMs, on the other hand, are a kernel-based approach with a strong mathematical background, which have become an increasingly popular de-facto tool for supervised machine learning tasks involving classification and regression. SVMs map the original training data into a high dimensional feature space using non-linear kernel functions, which allow them to construct optimal separating hyper-planes that greatly increases the power of learning and generalisation in higher dimensional spaces. The flexibility of the kernel-induced feature space is controlled by setting an upper bound for generalisation risk [Scholkopf and Smola, 2002], which has proved to improve the robustness against noisy sensor observations.

This paper is organised as follows. Section 4 describes the walker robotic platform and the mechanism employed to obtain the required activity data sets. Sections 5 and 6 respectively describe the proposed probabilistic and the discriminative models designed for the intention recognition. Section 7 presents the results of the experiments carried out and an analysis of the outcomes. Finally, Section 8 concludes and discusses future work.

4 Experimental Set-up

The power walker employed in this work is depicted in Figure 1. It is a modified commercial rollator walking frame with four wheels. The base frame has been instrumented with actuators and incremental encoders to the two rear wheels (front wheels are passive), two infra-

User Activities
General Assistive Navigation (GAN)
to Living Room (LRO)
to Kitchen (KIT)
to Bathroom (BAT)
to Medical Facility (MED)
to Bedroom (BED)
to Laundry (LAU)
Walker Go Away (WGA)
Recall Walker (RW)
Stand Up (SU)
Sit Down (SD)

Table 1: User activities.

red (IRs) proximity sensors to detect the presence of the user, four strain gauges (SGs) to detect indicative pressure, two on each of the walker’s handle-bars, a low-level micro-controller and a high-level control computer.

The strain gauges are two Micro Measurements 120UR. The differential force between the vertical axis of each handle-bar is used to establish how the user is holding on the handle-bar in readiness to start a task such as sitting down, standing up or ambulation.

The IR subsystem is made up of two Sharp GP2Y0A02YK, which are used to estimate whether the driving user is standing behind the walker and how far they are from it. Sensing range after calibration is [20,150]cm. The motorised actuation subsystem is based around two powerful Matsushita Electric GMX-8MC045A 24VDC reversible gear-head motor with optical encoder and rotary mechanical couplings. The motors are PWM driven using a national semiconductor LMD18200 3A, 55V H-Bridge motor driver.

The hardware also includes two push button switches installed under the handle-bars. At this stage these are used to simulate an RF switch (replacement is in progress), which indicates for the walker to come back from a parking position to where it last left the user. In the current framework, this feature has been added so that in certain locations of the house such as bedroom, living room or kitchen where the user tends to spend more time unaided, the walker can stand by in a safe place and is not a hindrance to other people using the same area.

4.1 Data Generation Set-Up

Large training and testing data sets are required for a meaningful analysis of the proposed strategies. In the scope of this project it was impractical to collect all the real-time data needed from a large pool of representative users, moreover given the lengthy ethical approval process required. Hence, sensorial observations were gener-




Figure 2: Division of office area into different point of interest used to simulate the daily activities.

ated for the chosen activities listed in Table 1 as a realistic representation of the daily tasks a typical walker user would encounter.

Sensor ranges and the behaviours expected for different users and actions have been estimated from previous smaller scale experiments performed on the walker platform with a number of able users [Patel *et al.*, 2010]. Hence, observed continuous data was generated for the two infra-red sensors (IRT, IRW) and left and right strain gauges (LSG, RSG), whereas the RF switch was model as discrete on/off. Along with data from these physical sensors, time of day (TOD) we also added as another discrete observation to emphasise the time-dependency of the different activities on the time of the day at which they take place. The day was assumed to span into Morning, Late-Morning, Afternoon, Late-Afternoon, Evening and Late-Evening. In order to simulate discrete localisation (LOC) observations (attained with a Monte Carlo localiser particle filter from laser and wheel encoder readings [Patel *et al.*, 2010]), our mapped office environment was considered as a representation of a typical home environment, and the geometrical space was divided into the relevant points of interest the user will normally visit during the day shown in Figure 2.

5 DBNs as Probabilistic Generative Models

Dynamic Bayesian Networks are a branch of Bayesian Networks (BN) for modelling sequential data. BNs are probabilistic graphical models represented by directed acyclic graphs (DAG) in which nodes are variables and arcs show the conditional independencies among the variables [Castillo *et al.*, 1997]. Each node on the net has a probability table associated containing the conditional probabilities of each of the values that node can take with respect to each of the possible combination of

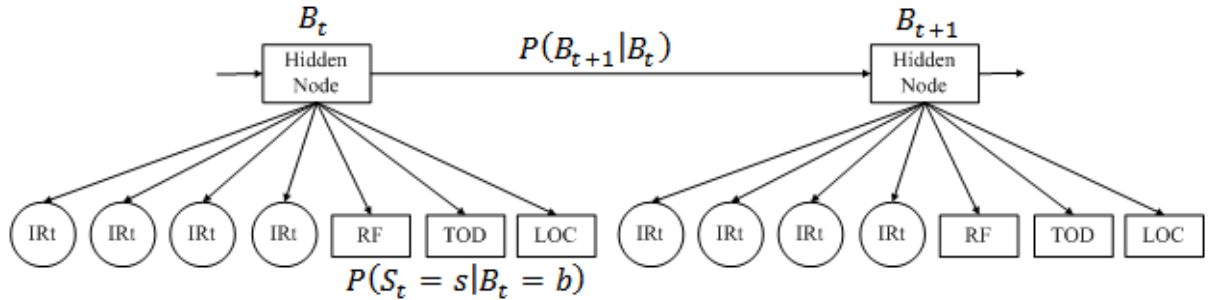


Figure 3: Two-slice DBN, with observed nodes IRt (Infra-red Torso), IRw (Infra-red Waist), LSG (Left Strain Gauge), RSG (Right Strain Gauge), RF (wireless Radio Frequency switch), LOC (Localisation) and TOD (Time-of-Day).

its parent nodes.

DBN provides probability distributions over semi-infinite collections of random variables. The variables are partitioned into $Z_t = (S_t, B_t, B'_t)$ to represent the input, hidden and output variables of a state-space model. The network only considers the discrete-time stochastic process, and it increases the index time t by one every time a new observation is recorded [Murphy, 2002].

A DBN is defined to be a pair, $(B_1, B_{->})$, where B_1 is a Bayesian Network which defines the prior $P(Z_1)$, and $B_{->}$ is a two slice temporal Bayes net which defines $P(Z_t/Z_{t-1})$ by means of a directed acyclic graph as given by:

$$P(Z_t/Z_{t-1}) = \prod_{i=1}^N P(Z_t^i | Pa(Z_t^i)) \quad (1)$$

Here Z_t^i is the i_{th} node at time t , which could be a component of S_t , B_t or B'_t , and $Pa(Z_t^i)$ are the parents of Z_t^i . Except for the first nodes in the first time slice, all subsequent nodes have an associated conditional probability distribution (CPD). Unlike the Hidden Markov Model (HMM), the DBNs represent the hidden state in terms of a set of random variables, whereas in HMMs the state space consists of a single random variable. Hence the inference of HMMs becomes computationally very expensive due to huge transition matrix as compared to DBNs [Russell and Norvig, 2003].

To construct a DBN three clusters of information needs to be defined: the prior distributions (initial probabilities of the state variables), the transition model and the conditional probability distribution between states, and the sensor models. To specify the transition probability model we must also specify the topology of the connections between successive slices, and between the states and the evidence variables within each time slice.

5.1 DBN Model for Intention Recognition on Walker Platform

The DBN model structure is designed with the user intention as the hidden state, and the sensor readings as

the observations, as graphically depicted in Figure 3. As can be seen, the DBN model structure has got one discrete hidden node and seven observed nodes. As described in Section 4.1, four of the seven observed sensor nodes are described by a continuous Gaussian distribution, while the other three are represented with discrete variables. The two slice DBN model is unrolled infinitely. The only connection between each time slice is via the hidden states.

We assume that the set of all possible intentions of the user is B . The intention at time t is represented by the random variable B_t . The reading of sensors at time t is given by random variable S_t and the user intention is denoted b' at $t+1$ and b at time t . The parameters also include the transition probabilities given by $\theta_{b',b} = Pr(B_{t+1} = b' | B_t = b)$ which specifies the probability of the intention at time $t+1$ is b' given that the behaviour at time t is b and the conditional probabilities of the observed node $\phi_b = Pr(S_t = s | B_t = b)$ which specified the probability of the value measured by sensor S at time t is s given the intention is b .

The hidden node of the DBN model has 11 hidden states, those described by the intended user activities collected in Table 1. The inter-state dependency of this hidden node is shown in Figure 4. The connectivity between each node indicates the probability of ending in state B at time $t+1$ provided the user was in a particular state at time t . No connectivity between states reflects the impossibility to end in state B at time $t+1$ provided the user was in some particular state at time t .

The conditional probabilities ϕ of the discrete observation nodes (RF, LOC, TOD in our case) have been derived by assuming standard behaviours of the user at particular instances (e.g. the possibility of a user going to the laundry in evening or morning is perceived to be relatively low as it can be termed as a scheduled phenomenon, whereas the possibility of a user going to bathroom during any time of the day is highly possible as its not a scheduled phenomenon for the user). We

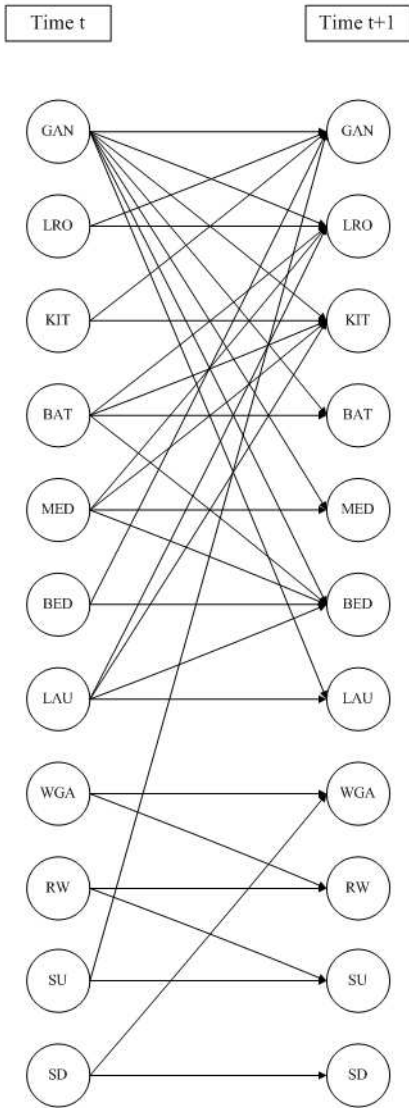


Figure 4: State transition matrix

also used Differential Evolution (DE) [Price *et al.*, 2005] optimisation technique to further optimise the assumed conditional probabilities based on the collected evidence, yet as will be discussed in Section 7, the results obtained after manually modelling the parametric distributions based on the collected data, and those obtained from DE optimisation were almost identical.

The transition probabilities for interstate connectivity specified in Figure 4 are manually defined based on prior knowledge [Patel *et al.*, 2010], and common laws of operation at what is perceived as accepted behaviour from the intended user pool (e.g. the user is unlikely to sit down immediately after standing up or it is impossible for the walker end up in state walker go away state after the user has just stood up and was in stand up state).

6 SVMs as Discriminative Classifiers

Support vector machines are a system for efficiently constructing and training the optimal separating hyper-planes in the kernel-induced feature space while enforcing the learning biases suggested by the generalisation theory. SVMs also produce sparse dual representations of the hypothesis, resulting in an extremely efficient algorithm that play a crucial role in the practical implementation and analysis of these machines [Christianini and Shawe-Taylor, 2000]. Additionally, due to the Mercer’s condition on the kernels functions [Cortes and Vapnik, 1995], the corresponding optimisation problems are convex and hence have no local minima. This in turn makes a clear distinction between SVM and other pattern recognition algorithms like neural networks.

Given a training set of instance-label pairs (\mathbf{x}_i, y_i) , $\mathbf{x}_i \in \mathcal{R}^n$ and $y \in \{1, -1\}^l$, $i = 1 \dots l$, with l being the number of samples, the instances x_i are mapped into a higher dimensional space $\phi : \mathcal{R}^n \rightarrow F$. SVM then constructs an optimal hyper-plane with maximum-margin and bounded error to divide these instances into two classes by solving the following optimisation problem:

$$\min_{\mathbf{w}, b, \zeta} \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{i=1}^l \zeta_i \quad (2)$$

$$y_i (\mathbf{w}^T \phi(\mathbf{x}_i) + b) \geq 1 - \zeta_i,$$

The first term in the cost function represented by this equation maximises the margin of separation between classes in the higher dimensional space. In simple words, it is responsible for maximising the distances between the nearest training points that belong to different classes along the sides of the constructed hyper-plane. This can be graphically seen on the example depicted in Figure 5. It can be observed how the dashed separation line on the right figure provides better separation between the data points that belong to the two different classes than the dashed line on the left side. This can also be expressed in terms of a classifier with a smaller margin having a higher expected risk of miss-classification, while that with a largest margin promises better classification rates when faced with unseen data. The second term in Equation 2 provides an upper bound for the error in the training data. Due to the inherent noises which regularly accompany the sensor observations, errors in the training data (i.e., samples that are placed on the wrong side of the hyper-plane shown by ζ_i) is practically inevitable. At the same time, the bounded error in the training data can prevent over-fitting. The constant $C > 0$ is the penalty parameter of the error term. $K(x_i, x_j) \equiv \phi(x_i)^T \phi(x_j)$ is the kernel function, for which the most frequent variations include

- Linear: $K(x_i, x_j) = x_i^T x_j$
- Polynomial: $K(x_i, x_j) = (\gamma x_i^T x_j + r)^d \gamma > 0$.




Figure 5: Examples of SVM hyper-planes separation margins.

- Radial Basis Function: $K(x_i; x_j) = \exp(-\gamma \|x_i - x_j\|^2), \gamma > 0$
- Sigmoid: $K(x_i; x_j) = \tanh(\gamma x_i^T x_j + r)$.

Here, γ , r , and d are kernel parameters. More details about the SVM are available in [Chen *et al.*, 2005] and [Cortes and Vapnik, 1995]. Given that the SVM is a binary classifier, then for implementing a multi-class problem various strategies exist. The one utilised in this paper is the "one-against-one" multi-class scheme included in libSVM <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>, a well known SVM library implementation. An optimisation over the various libSVM parameters described above against the labelled data was used to achieve the best possible performance.

7 Results

The DBN and SVM algorithms were tested off-line using the data generated as described in Section 4.1. The model was designed to recognise one of the eleven user activities. 1008 data samples were generated for each user intentions. Out of the total 11088 data samples simulated, 50% were used for training and 50% for testing. The data samples were generated keeping in mind daily routine activities performed by frail and elderly people at home or at a retirement nursing home. In this context, the average walker usage in the morning would be a timely sequence such as calling the platform -> stand up -> proceed to the bathroom -> go to the kitchen -> go to the bedroom -> sit down -> instruct the walker to go away, for instance. Unsupervised Expectation Maximisation (EM) was used by the DBN to learn the model parameters. Based on past observations, EM alternates between computing the expectations and updating the transition and conditional probability parameters that maximise the a-posteriori probability of the inferred intention states. Table 2 shows the confusion matrix for each of the user intended activities, while Table 3 collects the confusion matrix of the results achieved

with SVM static classification on the same data. Ground truth was established on data generation. In order to provide a more automatic process of knowledge fitting to observed data, conditional probability distributions were also optimised with DE, hence imposing a more strict verification of the results. Given sufficient time, the optimisation process from seeding random distributions attained similar accuracies when compared with manually defined conditional probabilities.

An analysis of the results shows how both techniques achieved near perfect classification accuracies for some classes (GAN, WGA, RW, SU, DU), while the rest of the classes were equally divided between the SVM and DBN. DBN managed to outperform SVM on three classes (KIT, BAT, and LAU), while SVM outperformed DBN on the remaining ones (LRO, MED, and BED). This can be attributed to the nature of the data itself facilitating certain classes to be better represented by one technique or the other, but the important factor is that overall recognition rates remain on parallel terms. Certain states (LRO, MED, KIT, BED, LAU and BAT) can be seen to be lower in both instances, which is justifiable as predictions for these states rely significantly on changes to LOC and TOD observations, with the remaining observations being less discriminative on those states. For instance, simulated data to infer KIT in TOD = Morning or Afternoon, is equally plausible in the simulated data set from various locations LOC. Since the remaining sensorial observations remain pretty much the same for a given user, it is difficult for any of the two models considered to be highly discriminative in those instances. Additional sensors or more training data would be the natural solution to overcome these deficiencies and help improve accuracy rates.

8 Discussion & Future Work

This paper has discussed human intention recognition in the context of assistive robotic walkers with AI models. The proposed algorithms have been designed based

Conf. Matrix	GAN	LRO	KIT	BAT	MED	BED	LAU	WGA	RW	SU	SD
GAN	96.83	0	0	0	0	0	0	0	0	0	3.17
LRO	0	57.14	9.12	0	0	0	33.73	0	0	0	0
KIT	0	11.90	63.28	2.38	0	0	21.42	0	0	0	0
BAT	0	0	21.82	67.86	0	4.71	0	0	0	0	0
MED	0	9.92	0	15.47	50.40	9.12	15.07	0	0	0	0
BED	0	0	31.74	19.84	0	48.42	0	0	0	0	0
LAU	0	13.49	0	0	0	0	86.52	0	0	0	0
WGA	0	0	0	0	0	0	0	100	0	0	0
RW	0	0	0	0	0	0	0	0	100	0	0
SU	0	0	0	0	0	0	0	0	0	100	0
SD	0	0	0	0	0	0	0	0	0	0	100

Table 2: DBN confusion matrix results for the activity recognition on the walker platform, with an overall accuracy of 79.22%.

Conf. Matrix	GAN	LRO	KIT	BAT	MED	BED	LAU	WGA	RW	SU	SD
GAN	99.00	0	0	0	0	0	0	0	0	0	0.99
LRO	0	85.51	9.12	0	0	0	5.35	0	0	0	0
KIT	0	28.57	58.73	0	0	7.93	4.76	0	0	0	0
BAT	0	0	18.65	51.38	2.57	27.38	0	0	0	0	0
MED	0	21.42	0	0	51.58	23.01	3.57	0	0	0	0.39
BED	0	0	24.80	0.39	3.76	71.03	0	0	0	0	0
LAU	0	40.47	0	0	0	0	59.52	0	0	0	0
WGA	0	0	0	0	0	0	0	100	0	0	0
RW	0	0	0	0	0	0	0	0	100	0	0
SU	0	0	0	0	0	0	0	0	0	100	0
SD	0.59	0	0	0.19	0	0	0	0	0	0	99.20

Table 3: SVM Confusion matrix results for the activity recognition on the walker platform, with an overall accuracy of 79.63%.

on probabilistic models (DBN) and discriminative statistical classifiers (SVM), and have been tested using a large data set generated from realistic sensor models and activity scenarios representative of the intended user population. Results have shown that unsupervised learning DBN algorithms can equally perform the results obtained with more established machine learning techniques such as widely used supervised learning with SVMs. DBNs being dynamic in nature are able to better account for conditional dependencies between states, hence can infer more likely behaviours where the transitional dependency on previous states is more apparent. While this is hard to encapsulate in the simulated data, it is expected this will surface more clearly when real data is used, particularly for unseen cases. The results also show that higher-level user behaviours associated with walker usage can be inferred using a minimalist set of sensors inputs with reasonable accuracy. In the future, we would like to further improve the recognition accuracy of user activities by better exploiting the dynamic

properties inherent in the structure of DBNs, and also to compare with other optimisation algorithms able to derive the best full posterior distributions over the model parameters for the given evidence and possible priors, e.g. Bayesian Learning. Moreover, based on the findings which indicate how SVM is able to provides better recognition accuracy for some states and DBN for others, it seems natural to exploit the combined strengths of DBN and SVM with a hybrid model as a possible vehicle to enhance the accuracy of the overall activity recognition model. Proceedings are also underway to be able to test the algorithms in real time on our in-house walker platform, first with able users but also with the representative user population once ethical approval is attained.

Acknowledgments

This work is supported by the Australian Research Council (ARC) through its Centre of Excellence programme, and by the New South Wales State Govern-

ment. The ARC Centre of Excellence for Autonomous Systems (CAS) is a partnership between the University of Technology Sydney, the University of Sydney and the University of New South Wales.

References

- [Begovich, 2005] M. Begovich. Arma - assisting robotic mobility aid, 2005.
- [Castillo *et al.*, 1997] Enrique Castillo, Jose Manuel Gutierrez, and Ali S. Hadi. *Expert Systems and Probabilistic Network Models*. Springer Verlag, 1997.
- [Chen *et al.*, 2005] P. H. Chen, C. J. Lin, and B. Scholkopf. A tutorial on nu-support vector machines. *Applied Stochastic Models in Business and Industry*, 2005.
- [Christianini and Shawe-Taylor, 2000] N. Christianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines and other kernel-based learning methods*. Cambridge University Press, 2000.
- [Cortes and Vapnik, 1995] C. Cortes and V. Vapnik. Support vector networks. *Machine Learning*, 1995.
- [Dubowsky *et al.*, 2000] Steven Dubowsky, Frank Génot, Sara Godding, Hisamitsu Kozono, Adam Skwersky, Haoyong Yu, and Long Shen Yu. PAMM - A robotic aid to the elderly for mobility assistance and monitoring: A helping-hand for the elderly. In *IEEE International Conference on Robotics and Automation, (ICRA 2000)*, pages 570–576, 2000.
- [Graf and Hgele, 2001] Birgit Graf and Martin Hgele. Dependable interaction with an intelligent home care robot. In *In Proceedings of ICRA Workshop on Technical Challenge for Dependable Robots in Human Environments*, page 2, 2001.
- [Heinze, 2003] C. Heinze. *Modeling Intention Recognition for Intelligent Agent Systems*. PhD thesis, The University of Melbourne, 2003.
- [Hirata *et al.*, 2006] Y. Hirata, A. Muraki, and K. Kotsuge. Motion control of intelligent passive-type walker for fall-prevention function based on estimation of user state. pages 3498–3503, may. 2006.
- [Hoey *et al.*, 2007] J. Hoey, A. von Bertoldi, P. Poupart, and A. Mihailidis. Assisting persons with dementia during handwashing using a partially observable markov decision process. In *Proceedings of the International Conference on Vision Systems*, March 2007.
- [Miró *et al.*, 2009] J. V. Miró, V. Osswald, M. Patel, and G. Dissanayake. Robotic assistance with attitude: A mobility agent for motor function rehabilitation and ambulation support. In *IEEE International Conference on Rehabilitation Robotics, 2009. (ICORR 2009)*, pages 529–534, June 2009.
- [Morris *et al.*, 2003] Aaron Christopher Morris, Raghavendra Rao Donamukkala, Anuj Kapuria, Aaron Steinfeld, J. Matthews, J. Dunbar-Jacobs, and Sebastian Thrun. Robotic walker that provides guidance. In *Proceedings of the 2003 IEEE Conference on Robotics and Automation (ICRA 2003)*, pages 25–30, May 2003.
- [Murphy, 2002] Kevin Patrick Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, Berkeley, 2002.
- [Omar *et al.*, 2010] Farheen Omar, Mathieu Sinn, Jakub Trzuskowski, pascal Poupart, James Tung, and Allan Caine. Comparative analysis of probabilistic models for activity recognition with an instrumented walker. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI 2010)*, July. 2010.
- [Patel *et al.*, 2010] M. Patel, J. V. Miró, and G. Dissanayake. Dynamic bayesian networks for learning interactions between assistive robotic walker and human users. In *The 33rd Annual German Conference on Artificial Intelligence (KI)*, pages 333–340, Sept. 2010.
- [Pineau *et al.*, 2003] Joelle Pineau, Michael Montemerlo, Martha E. Pollack, Nicholas Roy, and Sebastian Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42(3-4):271–281, 2003.
- [Price *et al.*, 2005] K. V. Price, R. M. Storn, and J. A. Lampinen. *Differential Evolution: A Practical Approach to Global Optimization*. Springer-Verlag Berlin Heidelberg, 2005.
- [Pynadath, 1999] D. Pynadath. *Probabilistic Grammars for Plan Recognition*. PhD thesis, The University of Michigan MI, 1999.
- [Russell and Norvig, 2003] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, N.J.:Prentice Hall/Pearson Education, 2003.
- [Scholkopf and Smola, 2002] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, Massachusetts, 2002.
- [Schrempf and Hanebeck, 2005] Oliver C. Schrempf and Uwe D. Hanebeck. A generic model for estimating user intentions in human-robot cooperation. In *Proceedings of the 2nd International Conference on Informatics in Control, Automation and Robotics (ICINCO 2005)*, volume 3, pages 251–256, Barcelona, Spain, sep 2005.