# Caring Analytics for Adults With Special Needs

**Marilyn Wolf and Honggab Kim**
Georgia Institute of Technology

**Jie Xu**
University of Miami

**Mihaela van der Schaar**
University of California, Los Angeles

*Editor's notes:*
Medical cyber–physical systems are expected not only to provide a more appropriate control of human physiological processes, but also to enable a significant improvement in care for adults with special needs. This paper describes a caring analytics system that collects and mines heterogeneous sensor network data for identifying the actions that should be taken for improving the care and quality of life of adults with special needs.
—*Paul Bogdan, University of Southern California*

■ **We propose a** novel caring analytics system for assisting with the long-term care of adults with special needs. Our proposed system combines sensor network-driven activity analysis and online learning algorithms to analyze each resident's care. The analysis should result in a variety of reports and alerts on activities of interest (is the resident eating regularly?) as well as recommendations (try a different type of food). We do so in a complex environment: each home contains several residents, one or more caregivers, and visitors. Our system must extract the activity of each resident from this noisy environment. Moreover, the conditions of the residents vary widely, and the recommendation system must be robust even though the available information may be limited.

The special needs community represents a wide range of cognitive and physical disabilities, sometimes in combination known as multiple involvements. People with special needs may need different amounts and types of attention and care in daily life depending on their diagnosis. Over the past 30 years, the United States has moved away from institutional care toward community-based care for people with special needs. While community-based care provides a more inclusive environment, assisted living or group homes do not have the staff and facilities to monitor and analyze the status and progress of the residents to nearly the same level as was possible for institutions. Much of the care provided to people with physical and cognitive special needs is not medical; instead, daily care concentrates on basic needs, activities, and a fulfilling life.

Service providers who provide care for adults with special needs have started to explore the use of smart homes to improve the quality and cost-effectiveness of special needs care. Imagine! Colorado, a Colorado Medicaid service provider, is one example of an organization that has run two smart homes for special needs adults for several years (http://www.imaginecolorado.org/smarthomes). However, making best use of the rich data provided by these sensor networks remains a challenge.

Our system builds on previous work in sensor networks and machine learning. This problem presents several technical challenges, including model-free analysis, analyzing data at multiple time scales,

*Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.*

*Digital Object Identifier 10.1109/MDAT.2015.2441717*
*Date of publication: 04 June 2015; date of current version: 01 September 2015.*

  

and accurate classification of noisy data streams. Our long-term goal is to analyze daily activity and make customized recommendations for each resident based on their individual needs.

The next section describes the technical challenges of our problem and briefly discusses previous work. We then go on to discuss our experimental setup, the architecture of our caring analytics system and our algorithms for activity analysis and learning-based recommendation.

## Challenges and previous work

The very practical problem of long-term care introduces several technical challenges in the design of algorithms for cyber–physical systems. We need powerful activity analysis algorithms (tracking, in the parlance of computer vision) that can infer useful activities from limited sensor data. We need to be able to identify the identity of the residents from the tracks without relying on explicit tagging and in the presence of multiple people in the home. Our partners at Imagine! have found that residents and staff rarely use the RFID tags that have been provided for them; as a result, our algorithms must infer identity from appearance and activity. Limitations exist on where sensors can be placed due to both privacy concerns and the practicalities of mounting sensors. While some of our data will come from cameras that can use appearance to help identify a subject, we must also rely on nonspecific sensors such as electric eyes at the doors of rooms. Hence, we need to be able to use sparse data from sensor networks to generate continuous, reliable tracks over extremely long periods—ideally years—to adequately understand the daily life of the residents.

Once we have extracted a set of activities from the sensor data, we need to make useful recommendations without relying on expert training. While the caregivers can provide observations and suggestions to the recommendation system, neither they nor the medical and therapeutic staff will be experts in computer science, computer engineering, or signal processing. The recommendation system must combine personalized and group learning. Because people with special needs exhibit a wide variety of conditions, the recommendation system must be able to adapt its recommendations to the observed behavior of each person—what works for one person may not work for another. However, some general strategies may also work for over a broad population. Group learning can help to codify strategies for dealing with the residents, which can in turn help to train new caregivers as they come on board. The wide range of conditions also means that we cannot rely on a predetermined set of candidate behaviors and recommendations. Instead, we must infer appropriate recommendations from experience with the residents. The wide range of conditions that we must be able to handle distinguishes our problem from elder care; although the elderly exhibit natural variation, the aging process generally follows a well-understood trajectory.

Several groups have developed activity analysis systems for healthcare and other applications. Much of this work has concentrated on well-defined tasks of relatively short duration. Acampora et al. surveyed research in ambient intelligence in healthcare [1]. They identified several characteristics of such systems: context-aware, personalized, anticipatory, adaptive, ubiquitous, and transparent. Fine and Singer [2] introduced hierarchical hidden Markov models (HHMMs); they used a generalization of the Baum–Welch algorithm to estimate model data. Patel et al. [3] used HHMMs to model home use of a walker. Pentney et al. [4] used a Markov logic to model learning of activities. Rashidi et al. [5] used a context-driven algorithm to identify sequences and a clustering algorithm to identify common contexts. Robben et al. [6] used ambient sensors to measure the activity of two persons over a long period; however, the individuals they measured were isolated and thus they did not need to separate observations among possible persons, nor did they make recommendation based on their analysis. Dawadi et al. [7] used sensors to evaluate the performance of older adults on predefined tasks. Wan et al. [8] segmented sensor data to identify a set of discrete activities.

Diagnosis and recommendation algorithms have also been developed for medical applications. Bennett et al. [9] use a finite horizon Markov decision process (MDP)—effectively, a finite decision tree—to create a framework to reduce care costs. In comparison with previous work, Bennett et al. achieve substantial improvements (in the metric they use) because their method looks ahead: at each decision point, it optimizes with respect to all possible future paths of outcomes and decisions (weighted according to probabilities) rather than simply optimizing with respect to the decision to be taken at the moment. Our analysis and recommendation framework can also be compared with well-known machine learning techniques including decision trees [10] and Bayesian networks [11]. Given an initial context and a sequence of actions and observations, those methods can determine the likelihood of the next state when the next action is applied. In the decision tree [10] or Bayesian network [11] frameworks, learning and reward maximization are uncoupled: learning is performed on the training data and reward is evaluated on the test data.

## Approach

### The long-term care environment

Adults with special needs require assistance for daily life. Their conditions range widely, including both cognitive (intellectual disability, autism, bipolar disorder, etc.) and physical (cerebral palsy, paralysis, epilepsy, etc.) disabilities. Adults with special needs require a wide range of care: daily activities (dressing, eating, and toileting); monitoring to ensure their safety; monitoring of medical issues; etc. Beginning in the 1970s, the United States started to move from an institutional care model to a community-based care model for special needs adults. Although institutions had permanent, trained staffs, much long-term care is provided either by family members or by caregivers who are not trained as nurses; a recent article by Dukakis [12] describes the importance of caregivers and the difficulties in finding them. Long-term care workers are often highly motivated by the opportunity to work with special needs adults, but financial compensation for such work is usually poor and the work itself can be frustrating—progress can be difficult to see and the paperwork required to document each resident's care is burdensome. As a result there is high turnover among caregivers. We hope that our cyber–physical systems will not only provide assistance to the special needs adults directly but also support caregivers.

In the United States, Medicaid provides funding for the long-term care of special needs adults. Each resident is required to have an individualized service plan (ISP) that is updated yearly. The plan is created and implemented by the care team. The plan lists a set of goals; depending on the individual, these can range from learning to eat with a fork to acquiring job-related skills. Progress on the goal is tracked throughout the year; progress from one year's plan is used to help formulate the next ISP. Systems such as ours can help to track progress on the ISP and make recommendations to help fulfill the plan's goals. Monitoring allows caregivers to monitor and assist the resident as much as possible: alerts tell them when the resident needs attention; long-term records give them a better understanding of overall health and behavior. Monitoring is a much broader term than surveillance: it encompasses progress on the ISP, reporting on trends in resident behavior to identify problems early, and evaluations that help to develop new goals.

### Architecture

Figure 1 shows the organization of our proposed system for long-term care analytics. In this person-in-the-loop scenario, our online tools provide alerts and information for use by staff. Some activities to be monitored for each resident are identified by his/her ISP; other daily activities, such as meals, recreation, and sleeping, can also be analyzed.

Sensor data can be used to create three types of outputs that are important to the care of special needs adults:

- reports on the condition of the resident, including both daily reports for use by staff and loved ones as well as reports on the resident's progress on

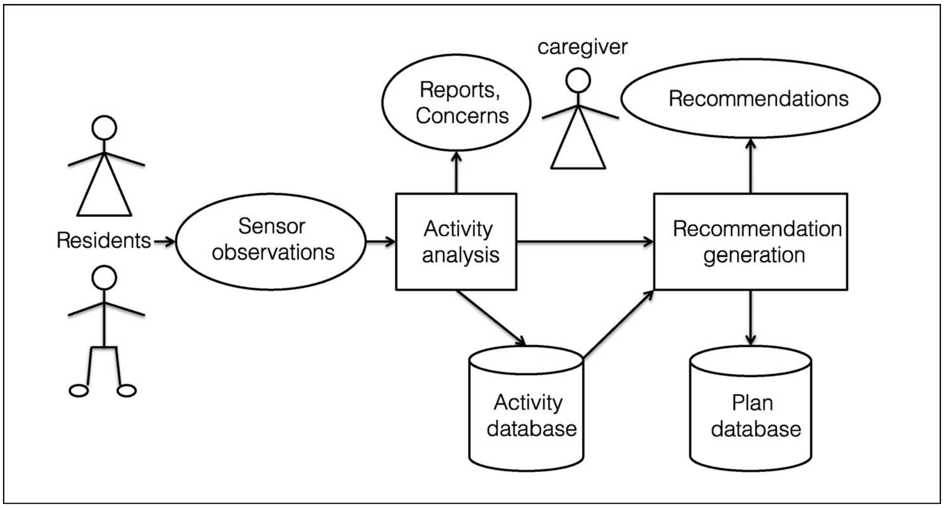


**Figure 1. Block diagram for the caring analytics architecture.**

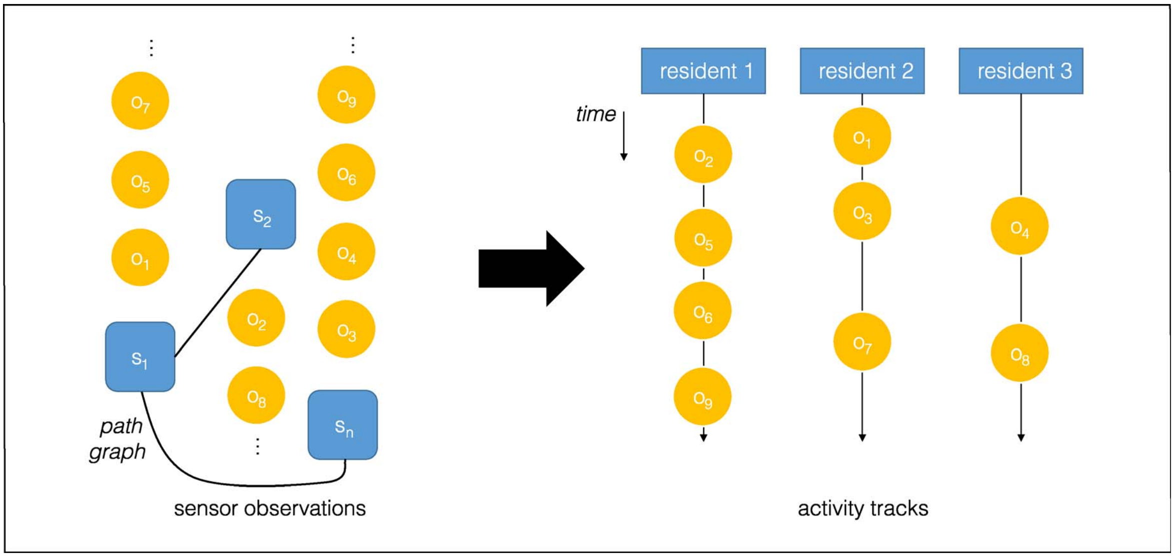


**Figure 2. Activity analysis from sensor observations.**

their ISP; these reports will be generated based on the identification of patterns in the resident's activity;

- concerns that alert the staff about changes to the resident's condition; these concerns will be generated by unusual patterns or changes to patterns, with learning to reject unnecessary concerns;
- recommendations on the progress toward goals in the ISP and what actions can be taken to achieve the goals.

This is a soft real-time system that must perform online analysis of event-oriented signals. The algorithms we use for both activity analysis and recommendations are designed to support the incremental updates and analysis required for this type of online operation. Activity analysis classifies events and observations into tracks for each resident that are fed to the recommendation system as well as storing the tracks in case staff wants to further analyze the activity of a resident. The recommendation system keeps a history of activities, the interventions posed by staff, and the results of those interventions. That history is used by its learning algorithms to generate new recommendations.

Given the computational requirements of activity analysis and recommendation generation, we believe that a distributed architecture is best suited to the application: perform initial sensor analysis (computer vision, etc.) locally, then ship observations to the cloud for higher level processing.

## Activity analysis

As shown in Figure 2, activity analysis maps sensor observations onto resident activity. The residence has, in general, several sensors, each of which generates a stream of observations. For example, a camera may generate an observation when a person enters the camera's field of view; the camera could provide a time stamp as well as an appearance model for each person. An electric eye at a door would fire when someone crossed the door's threshold. That type of sensor can provide a time stamp but no other information. We are also given a graph model, known as a path graph, of the spatial relationships between sensors; an edge connects two nodes for which there is a direct path between the corresponding sensors. The activity analysis system maps those observations onto a set of tracks, one for each resident. It uses Bayesian models to identify the most likely assignment of observations to tracks. Appearance is one factor—all observations in a track should have a similar appearance. Spatiotemporal relationships also help narrow the choices. For example, if a person is unlikely to move a large distance in a very short amount of time, then two observations close in time but far away in distance are unlikely to be of the same resident. Markov chain Monte Carlo (MCMC) algorithms have been used by several authors [13]–[15] to solve this tracking problem.

We have developed an algorithm for activity analysis from sparse sensor networks [16], [17]. Our original algorithm was designed to work with camera networks that could provide appearance models to help identify subjects. We have now explored the use of this algorithm in mixed sensor networks that include some sensors that do not provide appearance models. The sensors generate a set of observations by observing an unknown number $K$ of objects. We use a new form of path graph whose structure helps eliminate certain types of unlikely tracks. An observation includes an appearance measurement $(o_i)$, a timestamp $(t_i)$, and a sensor index. The goal of activity analysis is to partition the observations into subsets $\omega_K = \{Y_1, Y_2, \ldots, Y_K\}$ where each $Y_i$ represents the activity of a single resident. An observation can belong to only one activity. We wish to find a partition instance that maximizes the conditional probability $P(\omega_K|Y)$, which can be formulated as

$$P(\omega_K|Y) \propto \frac{1}{c}\prod_{k=1}^{K}\Bigg[P_b(x_{k,1})\prod_{i=2}^{N_k}P(x_{k,i}|x_{k,i-1}) \times p(o_{k,i}-o_{k,i-1}|x_{k,i},x_{k,i-1}) \times p(t_{k,i}-t_{k,i-1}|x_{k,i},x_{k,i-1})\Bigg],$$

where $P_b(x_{k,1})$ is the probability of a new person appearing in the observations, $P(x_{k,i}|x_{k,i-1})$ is a transition probability, $p(o_{k,i}-o_{k,i-1}|x_{k,i},x_{k,i-1})$ is an appearance similarity density, and $p(t_{k,i}-t_{k,i-1}|x_{k,i}, x_{k,i-1})$ is a travel time density.

This form of tracking problem presents a huge search space—the number of possible partitions of the observations into tracks is given by Bell's number $B(n)$. This maximum posterior estimation problem is reduced to a maximum weight matching problem, which can be solved by the Hungarian method or MCMC optimization. The Hungarian method can provide an optimal solution. However, it is less scalable to large amounts of data and parallel operation than are MCMC-based algorithms. MCMC-based algorithms seek an optimal solution by recursively generating plausible partition instances based on a current partition instance. For the Monte Carlo operation, Kim et al. [16] propose three moves: "update" moves an observation from one subset to another; "split" divides a subset into two subsets; and "merge" combines two subsets into one. One of the challenges in using MCMC for activity analysis is generating an initial estimate of the posterior and of the number of tracks. Our algorithm places each event in a separate subset and shows that both the number of subsets and the posterior quickly converge to the actual number of activities and the maximum *a posteriori*. We have developed a distributed version of our algorithm that parallelizes well and exhibits good memory behavior.

Different observations come with different amounts of appearance and identification, depending on the type of sensor used. For example, a camera can report approximate body size, color of clothing and hair, etc. However, many other types of useful sensors do not give such specifics. For example, electric eyes at the doors do not report who walked through the door or even the direction of travel. Similarly, monitors on water faucets and electric outlets provide useful information but without personal identification. Our inference approach helps us propagate identification information from observations that provide it to observations that do not.

## Personalized recommendations

Once observations of the resident have been organized into activity tracks, the system uses learning algorithms to make use of those activities. Given the outcomes of those activities (did the resident complete his or her exercise?), the recommendation will classify activities and outcomes and make recommendations on what future actions to recommend to best meet goals.

The resident's state evolves stochastically as care actions are taken on the resident. For each care action taken in a given state, a reward is obtained. The reward can be either positive or negative: for the intermediate care actions, it represents the care cost incurred; for the final action that ends the care, it represents the final care benefit after the goal of care is completed. The optimal planning problem can be solved as a Markov decision process (MDP) using the estimated transition probabilities and reward functions. However, one of the key challenges for many long-term care systems is that access to relevant data is limited. Hence, transfer of knowledge from past scenarios is key. Existing data sets often have distributions that do not necessarily capture the information needed for performing the best care for a new resident or one whose condition is changing. To improve the recommendation for each resident, we will learn not only from that resident's activity but also from the activity from other residents.

A care plan is a set of care actions that are recommended to the care team in the various states.

The goals of the care plan are set by the resident's ISP. The overall care is evaluated on the basis of the sequence of costs and benefits. We define the long-term care reward for resident $t$

$$R = \sum_{t=1,\ldots,L(t)} \gamma^l r_l,$$

where $r_i$ is the intermediate reward and $\gamma$ is a discount factor.

The state transitions are probabilistic, representing possible future actions of the resident. The transition probabilities and the reward function are Markovian; they depend only on the current state of the special needs adult. This is a reasonable approximation in practice. The optimal plan maximizes the long-term expected reward in each state, which is defined using the Bellman equation

$$V(a|s) = r(a|s) + \gamma \sum_{s'} p(s'|s,a) V(s').$$

Due to the lack of a training data set, it is initially impossible to construct a good plan for a new resident. Instead, we have a set of $K$ existing plans $\{p_1, \ldots, p_K\}$ constructed using existing data sets. However, the relationship and the effectiveness of these plans on the new special needs adults are unknown *a priori*. Our algorithm begins by exploring the existing plans for new special needs adults. After accumulating sufficient data, it uses the learned similarity between the special needs adults to build new individualized plans that maximize the long-term care reward. There are two major questions that remain to be addressed: 1) which existing plan to apply when we are exploring the existing plans before accumulating sufficient data; and 2) how to build new individualized plans after data are accumulated.

Ideally, we would like to always apply the most effective existing plan. However, this is impossible since we cannot evaluate the effectiveness of existing plans before the system starts. To do this, we make the optimal tradeoff between exploring different existing plans in order to learn their effectiveness online and exploiting the most effective plans that are learned so far in order to maximize the long-term care reward. Moreover, since the existing policies may have different performance on different special needs adults, we also adaptively cluster patterns of

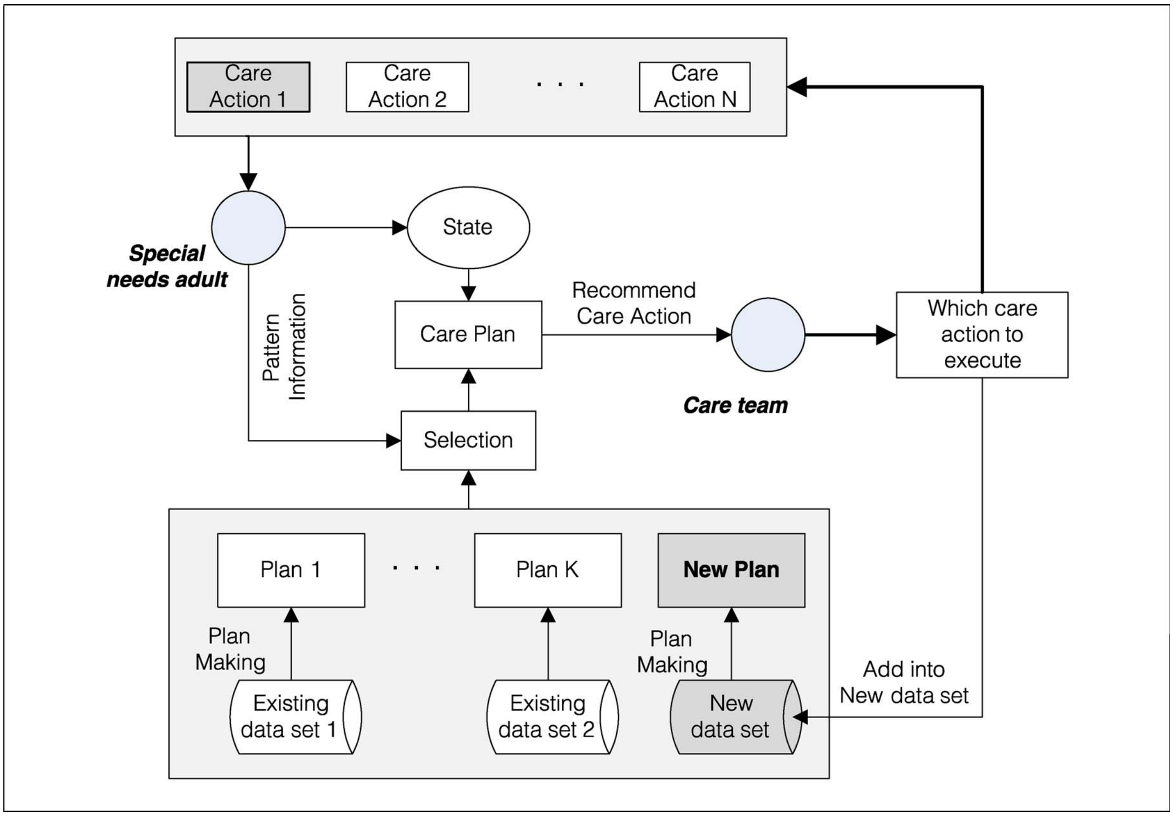


**Figure 3. The recommendation generation process.**

the residents according to the effectiveness of different existing policies.

As shown in Figure 3, when a new resident enters the system, we use the resident's diagnosis to place him or her into an initial cluster. We then determine if there exist any underexplored existing plans according to a carefully designed threshold. If one exists, the system recommends care actions to the care team according to this plan for this resident; otherwise, the system recommends care actions to the care team according to the plan with the highest estimated long-term care reward. We also check if there have been sufficiently many cases in the current cluster according to another carefully designed threshold. If yes, the current cluster is further partitioned into smaller clusters so that learning can be refined.

Special needs adults whose most effective plans are the same are considered to be similar. The data about these adults are grouped together to construct new individual plans when sufficient data/experiences are collected. Since there are $K$ existing plans, we create $K$ groups of residents that have been received so far: plan $p_k$ is the most effective for special needs adults in group $k$. For each group, the transition probabilities and reward functions are estimated using the data about adults in this group based on which an individualized care plan can be constructed by solving the MDP described earlier.

Using our algorithm, we can derive a following confidence bound on the learned effectiveness of existing plans. Having such a confidence bound is essential, since it provides the care givers with information about the effectiveness of each plan, thereby being able to reduce ambiguity about plans and increased confidence in these plans by caregivers. We can also prove that the algorithm converges in sublinear time to the optimal performance.

## Results

We performed experiments to determine the effectiveness of both our activity analysis and recommendation algorithms in the long-term care environment.

We evaluated our activity analysis algorithm under different sensor configurations. Our original algorithm was designed to work with camera data that included appearance models as well as timestamps and camera location. In many cases, the number of cameras will be limited, and we must rely on electric eyes or other nonidentifying sensors to sense activity. We conducted a new simulation experiment to determine the effectiveness of our algorithm with a mixture of appearance-bearing and non-appearance-bearing observations.

We compared three scenarios using a set of simulated data: a set of 121 simulated cameras on a grid (EXP-I); and a combination of 96 cameras and 25 pressure sensors on a grid (EXP-II); and a partial grid of 96 cameras (EXP-III). Figure 4 shows the

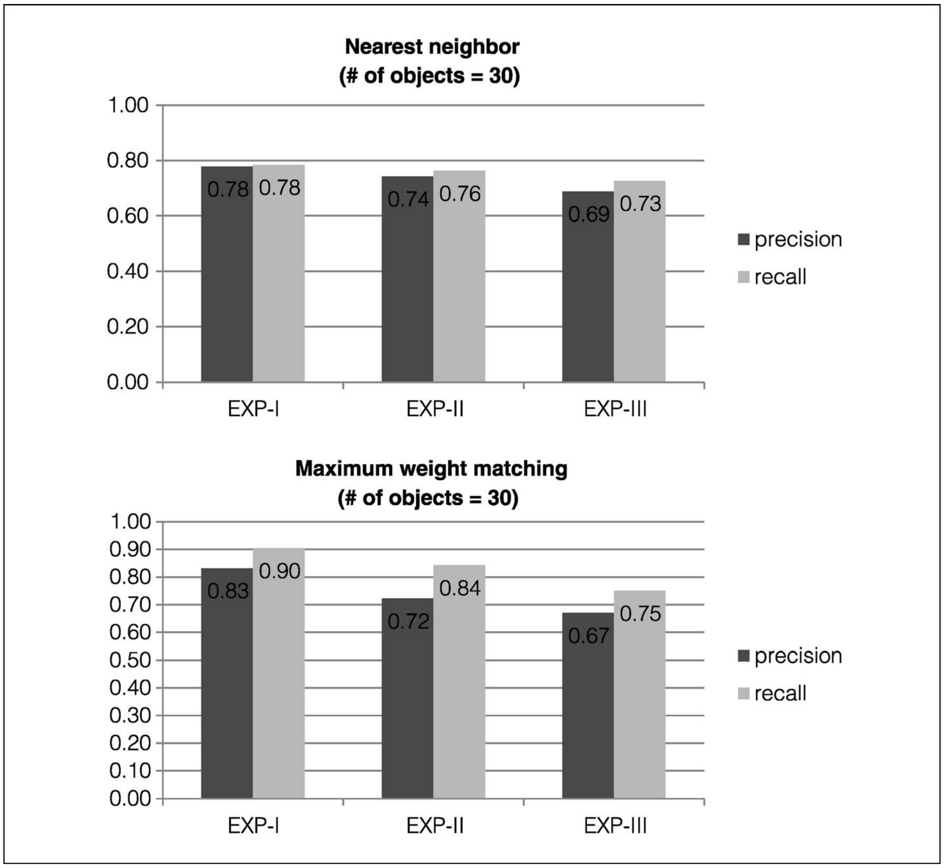


**Figure 4. Results of two different tracking algorithms in three different scenarios.**

**Table 1 Normalized reward for individual activities.**

| | Running | Walking Around | Walking Normal |
|---|---|---|---|
| Our method | 0.98 | 0.98 | 0.97 |
| Weighted Majority | 0.84 | 0.77 | 0.69 |
| AdaBoost | 0.85 | 0.80 | 0.70 |

results of the camera/pressure sensor experiment comparing our algorithm using two different algorithms to solve the tracking problem: nearest neighbor matching and maximum weighted matching. (Precision is the fraction of recognized activities that are relevant; recall is the fraction of ground truth activities that were recognized.) Using pressure sensors somewhat reduced accuracy and completeness (EXP-II) but using them was better than relying on a reduced set of cameras (EXP-III).

We evaluated our recommendation algorithm using real-world sensor data collected from various users. Three major components of the system need to be deployed for data collect and system evaluation. The server responsible for user and scenario management, context, and activity classification was deployed to the medical network servers of the University of California, Los Angeles (UCLA). The domain expert client was given the collaborators at the UCLA's Department of Neurology. The end-user component is a physical package containing four body worn inertial measurement units (IMUs) with Velcro attachments, a Nexus 7 table, and the associated

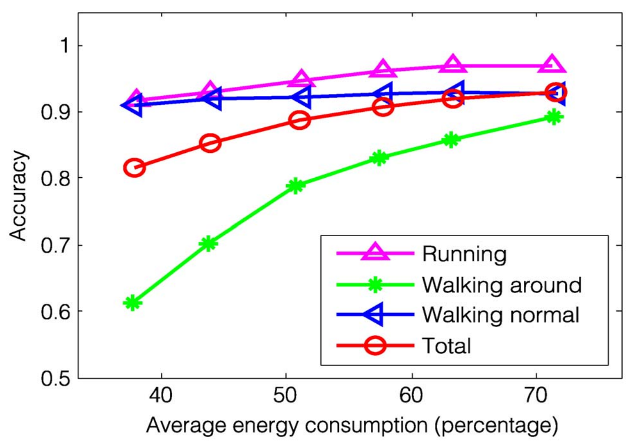


**Figure 5. Tradeoff between accuracy and energy consumption (normalized to the maximum energy consumption).**

mobile application. Note, however, that the proposed method can also be applied a standalone system which does not feature a client–server architecture. The server continuously receives classification requests from end-users through wireless transmissions (Wifi, Bluethooth, or cellular).

We considered three possible activities, "running," "walking around," and "walking normal," as possible plans. The activity "walking around" refers to nonsustained walking segments that are typical of walking in confined spaces, while "walking normal" refers to sustained long-distance walk typical of open spaces. If the resident performs well in one plan, the recommendation can change for him/her to move to the next plan. A track identified by activity analysis can be classified into these three categories based on the pace of observations (running versus walking) and the physical length of the tracks (walking around versus walking normal).

The results in Table 1, which breaks out the reward values for different types of activities, show that our algorithm can correctly track the various activities and that it outperforms the considered benchmarks. As a benchmark we compare our methods reward against the well-known weighted majority and Adaboost methods.

Finally, we investigate the tradeoff between energy consumption (i.e., classification cost) and classification accuracy. Figure 5 illustrates the accuracy and energy consumption tradeoff curve of the proposed algorithm. The energy consumption is normalized to the maximum power consumption when all sensors are activated. Note that WM and AdaBoost use all sensors for all requests and, hence, they are not able to make tradeoff between energy consumption and accuracy. As can be seen from the figure, a higher accuracy can be obtained at a cost of higher energy consumption for all three activities as well the overall performance.

## Future Steps

Based on our experiments, we are working with Imagine! Colorado to deploy a prototype system in a smart home with real residents. As we look forward to that ambitious project, we plan to work on several important extensions. We plan to extend our activity analysis algorithm to also identify patterns of activity based on metrics for similarity between tracks. By recursively defining the similarity measure, we can identify patterns made of subpatterns, which will allow us to identify patterns of activity over multiple time scales, such as daily versus weekly patterns. We also need to analyze the robustness of our recommendation algorithms. We will see wide variations in the capabilities and behaviors of residents, and group-oriented learning should not be confused by these variations. As a practical matter, we need to find a recipe for the right mixture of sensors that allows us to capture the activities that are of most interest to the residents and staff. Both privacy and the practicality of mounting sensors in a house limit both the number and types of sensors that we can use.

**Long-term care** of special needs adults presents new challenges for data analysis, recommendation systems, and sensor networks. We also believe that a Big-Data-centric, broad spectrum approach is the only way to bridge the gap in care left by the closure of large, highly staffed institutions. We think that the combination of new sensors, transfer learning, and the ability to provide confidence bounds to professionals by such systems will help the healthcare community to fulfill its potential to provide adults with special needs with fulfilling lives. ■

## Acknowledgment

We greatly appreciate the help of Kevin Harding and Greg Wellems of Imagine! Colorado who explained to us the importance of smart homes and their needs for data analytics. The work of M. van der Schaar and J. Xu was supported by the U.S. Air Force Research Laboratory (AFRL) under the DDAS program. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the authors and do not reflect the views of AFRL.

## ■ References

[1] G. Acampora, D. J. Cook, P. Rashidi, and A. V. Vasilakos, "A survey on ambient intelligence in healthcare," *Proc. IEEE*, vol. 101, no. 12, pp. 2470–2494, Dec. 2013.

[2] S. Fine and Y. Singer, "The hierarchical hidden Markov model: Analysis and applications," *Mach. Learn.*, vol. 32, no. 1, pp. 41–62, 1998.

[3] M. Patel, J. V. Miro, and G. Dissanayake, "A Hierarchical hidden Markov model to support activities of daily living with an assistive robotic walker," in *Proc. 4th IEEE RAS & EMBS Int. Conf. Biomed. Robot. Biomechatron.,* Jun. 24–27, 2012, pp. 1071, 1076.

[4] W. Pentney, M. Philipose, J. Blimes, and H. Kautz, "Learning large scale common sense models of everyday life," in *Proc. 22nd Nat. Conf. Artif. Intell.*, 2007, pp. 465–470.

[5] P. Rashidi, D. J. Cook, L. B. Holder, and M. Schmitter-Edgecombe, "Discovering activities to recognize and track in a smart environment," *IEEE Trans. Knowl. Data Eng.*, vol. 23, no. 4, pp. 527–539, Apr. 2011.

[6] S. Robben, M. Pol, and B. Kröse, "Longitudinal ambient sensor monitoring for functional health assessments: A case study," in *Proc. ACM Int. Joint Conf. Perv. Ubiquitous Comput., Adjunct Publ.*, 2014, pp. 1209–1216, DOI: 10.1145/2638728.2638812.

[7] P. Dawadi, D. Cook, C. Parsey, M. Schmitter-Edgecombe, and M. Schneider, "An approach to cognitive assessment in smart home," in *Proc. Workshop Data Mining Med. Healthcare*, 2011, pp. 56–59.

[8] J. Wan, M. J. O'Grady, and G. M. P. O'Hare, "Dynamic sensor event segmentation for real-time activity recognition in a smart home context," *Pers. Ubiquitous Comput.*, vol. 19, no. 2, pp. 287–301, 2015.

[9] C. C. Bennett and K. Hauser, "Artificial intelligence framework for simulating clinical decision-making: A Markov decision process approach," *Artif. Intell. Med.*, vol. 57, no. 1, pp. 9–19, Jan. 2013.

[10] J. R. Quinlan, "Probabilistic decision trees," in *Machine Learning: An Artificial Intelligence Approach.* San Francisco, CA, USA: Morgan Kaufmann, 1990, vol. 3, pp. 140–152.

[11] D. Heckerman, *A Tutorial on Learning With Bayesian Networks*. Berlin, Germany: Springer-Verlag, 2008.

[12] O. Dukakis, "Who cares for the home-care aides?" *CNN Opinion*, Mar. 6, 2014. [Online]. Available: http://www.cnn.com/2014/03/06/opinion/dukakis-caregivers/index.html?iref=allsearch

[13] D. Marinakis, G. Dudek, and D. Fleet, "Learning sensor network topology through Monte Carlo expectation maximization," in *Proc. IEEE Int. Conf. Robot. Autom.,* Apr. 2005, pp. 4581–4587.

[14] S. Oh, S. Russell, and S. Sastry, "Markov chain Monte Carlo data association for general multiple-target tracking problems," in *Proc. 43rd IEEE Conf. Decision Control,* Dec. 2004, vol. 1, pp. 735–742.

[15] H. Pasula, S. Russell, M. Ostland, and Y. Ritov, "Tracking many objects with many sensors," in *Proc. 16th Int. Joint Conf. Artif. Intell.*, San Francisco, CA, USA, 1999, vol. 2, pp. 1160–1167.

[16] H. Kim, J. Romberg, and W. Wolf, "Multi-camera tracking on a graph using Markov chain Monte Carlo," in *Proc. 3rd ACM/IEEE Int. Conf. Distrib. Smart Cameras*, 2009, DOI: 10.1109/ICDSC.2009.5289352.

[17] H. Kim and M. Wolf, "Distributed tracking in a large-scale network of smart cameras," in *Proc. 4th ACM/IEEE Int. Conf. Distrib. Smart Cameras*, 2010, pp. 8–16.

**Marilyn Wolf** is a Rhesa "Ray" S. Farmer, Jr., Distinguished Chair and Georgia Research Alliance Eminent Scholar at the Georgia Institute of Technology, Atlanta, GA, USA. She was with AT&T Bell Laboratories from 1984 to 1989. She was on the faculty of Princeton University from 1989 to 2007. Her research interests include cyber–physical systems, embedded computing, embedded video and computer vision, and VLSI systems.Wolf has a BS, an MS, and a PhD in electrical engineering from Stanford University, Stanford, CA, USA (1980, 1981, and 1984, respectively). She received the ASEE Terman Award and the IEEE Circuits and Systems Society Education Award. She is a Fellow of the IEEE and the Association for Computing Machinery (ACM) and an IEEE Computer Society Golden Core member.

**Honggab Kim** is currently working toward a PhD at the Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. After completing his MS studies, he joined the Electronics and Telecommunications Research Institute, where he designed and developed component-based frameworks to simultaneously process heterogeneous geospatial data in real time. His major research area is distributed tracking algorithms for a large-scale network of smart cameras. His research interests include image processing, computer vision, pattern recognition, and visual senor networks. Kim has a BS in control engineering and an MS in electrical engineering from Chonbuk National University, Korea (1997 and 1999, respectively).

**Mihaela van der Schaar** is Chancellor's Professor of Electrical Engineering at the University of California Los Angeles, Los Angeles, CA, USA. She is also the Founding and Managing Director of the UCLA Center for Engineering Economics, Learning, and Networks. She was a Distinguished Lecturer of the Communications Society for 2011–2012 and the Editor-in-Chief of the IEEE Transactions on Multimedia (2011–2013). She received the National Science Foundation (NSF) CAREER Award (2004), the Best Paper Award from the IEEE Transactions on Circuits and Systems for Video Technology (2005), the Okawa Foundation Award (2006), the IBM Faculty Award (2005, 2007, 2008), the Most Cited Paper Award from *EURASIP: Image Communications Journal* (2006), the Gamenets Conference Best Paper Award (2011), and the 2011 IEEE Circuits and Systems Society Darlington Award Best Paper Award. She holds 33 granted U.S. patents. She is a Fellow of the IEEE.

**Jie Xu** is currently an Assistant Professor in the Department of Electrical and Computer Engineering, University of Miami, Miami, FL, USA. His research spans the area of machine learning, data mining, and game theory. His interests lie in both developing the theory in these areas and applying it in real-world engineering systems, such as communication networks, cybersecurity, online social platforms, and healthcare informatics. Xu has a BS and an MS in electronic engineering from Tsinghua University, Beijing, China (2008 and 2010, respectively) and a PhD in electrical engineering from the University of California, Los Angeles (UCLA), Los Angeles, CA, USA (2015).

■ Direct questions and comments about this article to Marilyn Wolf, School of Electrical and Computer Engineering, Georgia Institute of Technology (Georgia Tech), Atlanta, GA 30332 USA; wolf@ece.gatech.edu.