

# Home Occupancy Agent: Occupancy and Sleep Detection

Stephen Makonin and Fred Popowich

**Abstract**—Smart homes of the future will have a number of different types of sensors. What types of sensors and how they will be used depends on the behaviour needed from the smart home. Using the sensors to automatically determine if a home is occupied can lead to a wide range of benefits. For example, it could trigger a change in the thermostat setting to save money, or even a change in security monitoring systems. Our prototype *Home Occupancy Agent* (HOA), which we present in this paper, uses a rule based system that monitors power consumption from meters and ambient light sensor readings in order to determine occupancy. The agent is also able to determine when the occupants are asleep, and thus provide the potential for further energy saving opportunities.

**Index Terms**—Smart Home, Intelligent Agent, Occupancy Detection, Power Consumption, Ambient Light Sensors, Sleep Detection, Energy Conservation, Sustainability.

## I. INTRODUCTION

Research into *intelligent homes* has focused on a range of different issues, including the *human-home* interaction [1], and occupant safety [2]. There remain some very basic issues, though, that play a key role to a wide range of smart home investigations. One of these basic issues is how to detect if someone is at home, and if so, whether the occupant(s) are active or alternatively in a sleep period. Investigation into this question leads to other questions. What kind of architecture can facilitate sensor configuration and installation? Exactly what information can be determined for a given type of sensor? Can this be done using very simple, low cost, energy efficient sensors? Indeed there are a number of research projects focusing on various aspects of these questions [3], [4], [5], [6].

The automatic determination of whether someone is home, or whether the occupants are asleep, can be difficult, but has a number of exploitable benefits in terms of energy conservation, sustainability, and cost savings for the homeowner. Consider an intelligent agent (IA) that could turn down the heat by 2°C when no one is home. Such automatic behaviour could save the home owner money by lowering heating costs and could also save the home owner time, since (s)he would not need to (remember to) override the thermostat setting.

In this paper, we explore how different types of sensors can be used to get information in the context of determining whether a home is occupied or unoccupied, and specifically whether it is active or inactive. We adopt a *bottom-up* approach; looking at each type of sensor in detail, first. *Home*

*Occupancy Agent* (HOA) is our first single agent system prototype for a smart home. With HOA, we start to put forth some answers to the questions posed above. In section 2 we introduce a generalized system design of a sensor network with a centralized intelligent agent. We then discuss the ideas behind monitoring power consumption in section 3, and formulate an algorithm that can be used to determine household occupancy and activity. In section 4, we investigate ambient light sensors and their use, and introduce a *profile matrix*, discussing how to use it in conjunction with ambient light sensor readings to better determine household and room activity. We then take some of our observations with ambient light sensors and derive an algorithm to determine nightly sleeping periods of occupants (section 5). In section 6, we describe our implementation of HOA in a *real home* using *off the shelf products* and report on our test results. We relate our research to similar studies in section 7, and provide some conclusions and directions for future research in section 8.

## II. SYSTEM DESIGN

Our approach to designing a smart home is to have sensors communicate with IA software running on appliances and devices within the home. Some of these agents can then interact with the user(s). Types of agents that will run in a smart home can be categorized as either centralized or localized agents. Moslehi notes that there is a *virtual hierarchy* and a *temporal dimension* to any agent system that needs to be developed for the smart grid [7], [8]. These ideas that were developed for the smart grid can be scaled down and used within a smart home. Both systems will have a centralized server system, and localized controls and sensors.

*Centralized agents* will have a responsibility to collect and store localized sensor data on a continual, periodic basis. These agents can then report aggregations of this data and present them to the user either via the web or an in-home display (IHD). Whereas, *localized agents* monitor the sensory equipment directly connected to them and make localized intelligent decisions based on the data from the sensors they monitor. Agents will need to communicate with other agents within the home through a series of events, triggers, parameters, and alarms [7].

Figure 1 shows a generalized smart house architecture derived from Moslehi's investigations. The system architecture needs to be general enough to be implementable in a wide variety of different environments. The architecture also needs to be extensible in terms of being able to add components, such as sensors and controls without major rework of the *plug and*

S. Makonin and F. Popowich are with the School of Computing Science at Simon Fraser University.

Website: <http://www.cs.sfu.ca/>

Email: {smakonin, popowich}@sfu.ca

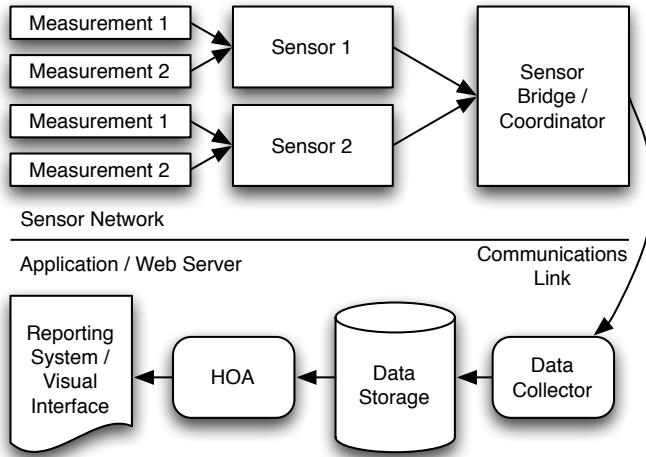


Fig. 1. A generalized architecture for a Home Occupancy Agent.

play system architecture. Note that we are proposing an agent that focuses on a more centralized approach. That is to say that the agent will be running on a server and accessing data collected from local sensors via a database to make decisions. We do this to simplify our initial investigation. Another simplification for our initial implementation is that HOA will be a *rule-based* intelligent agent. If we were to deal with a large number of sensors and a wide range of intelligent behaviours, then a distributed HOA system should be investigated. The general architecture from Figure 1 will form the basis for introducing and evaluating a specific implementation of HOA in section 6.

In general terms, the different factors we can monitor are: *consumption* of power from appliances or devices; *environmental* sources such as temperature, light levels, time/date, noise; and the *movement* of individuals, air. All three of these factors are significant to determine home occupancy or activity patterns, but for our current study, we are only looking at the first two in any detail.

### III. POWER CONSUMPTION

When monitoring power consumption for detecting occupant activity we look for spikes in kilowatt (kW) readings that correlate directly with occupants turning appliances on and off. Figure 2 shows a graph of kW readings recorded from the two meters in our test home. Our test home also has a heat pump, which is not directly controlled by the occupants. Notice that it is easy to tell when the heat pump is operating and when it is not. Monitoring consumption will be an ideal choice in the future due to high priority initiatives from electric utility companies who will be installing *smart meters* on every home. This means that power consumption data will be available to use and study.

In Figure 2, we show electric power consumption ( $C^E$ ) for a household where there is a meter attached to the heat pump. With information from this meter, we can filter out HVAC system consumption because it: (a) creates large spikes, and (b) will often run when no one is home creating false-positives in occupancy determination. Once filtered out, we can get

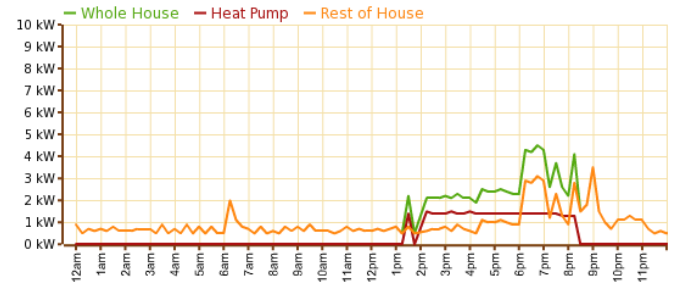


Fig. 2. Example of power consumption of a household over a 24-hour period (measured in kW).

a more accurate activity determination and find a baseline kW reading for the home, called *rest of house* ( $C_R^E$ ). Any spike above some *safe margin* ( $C_M^E$ ) would be then considered activity initiated by the occupant, for example, boiling a pot of water to make tea.

We must set  $C_M^E$  to a reasonable value. In Figure 2,  $C_M^E = 0.9\text{kW}$  would be a reasonable margin, because our test home has an ambient consumption ( $C_R^E$ ) that oscillates anywhere from  $0.2\text{kW}$  to  $0.7\text{kW}$ . All homes will have a certain amount of ambient consumption (fridge, freezer, electronics on stand-by mode). Only when meters are installed at every wall outlet would be possible to have no ambient consumption and  $C_M^E = 0\text{kW}$ . From the above observations we can derive the following formulae:

$$C_A^E = C_{A1}^E + C_{A2}^E + \dots + C_{Aa}^E \quad (1)$$

$$C_R^E = C_T^E - C_A^E \quad (2)$$

$$activity? = \begin{cases} true, & \text{if } C_R^E \geq C_M^E \\ false, & \text{otherwise} \end{cases} \quad (3)$$

In (1)-(3),  $C_T^E$  is the kW reading of the *whole house*,  $C_{A1}^E$  to  $C_{Aa}^E$  is a set of kW readings from occupant independent consumption activities (or appliances) that we want to filter out, and  $a$  is the number of filtered out activities (or appliances).

### IV. AMBIENT LIGHT SENSORS

When monitoring ambient light sensors for detecting occupancy we look for illuminance spikes measured in lux (lx). These spikes correlate directly with occupants turning lights on and off in different rooms. By placing ambient light sensors in various rooms we can thus determine occupant activity.

Figure 3 shows a graph of ambient light level readings recorded from each room in our test house. Immediately it is obvious that the readings from room to room are dramatically different and that the data in most rooms are not smooth. This lack of *smoothness* can be attributed to sunlight levels changing due to: different atmospheric events, such as the sun being blocked by a cloud (not a sign of occupant activity); as well as curtains or blinds being open or shut (which is a sign of occupant activity). This means that using ambient light sensors is not as straight forward as using power consumption.

Some rooms will have more sunlight than others. For example, a north facing room will have small amounts of light

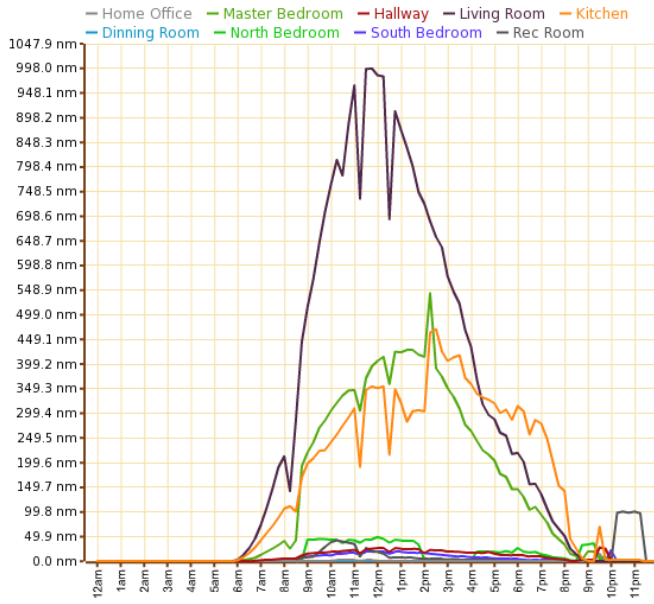


Fig. 3. Example of ambient light levels of a household for each room over a 24-hour period (measured in lux).

when the curtains are open and practically no light when the curtains are closed. Rooms that are south facing have so much sunlight that we cannot tell if the lights are turned on or off during the day. From this we can safely assume that there would be no need for an occupant to turn on a light. Looking at Figure 3 we can observe that light turned on at night reach a light level of  $\sim 100$ lx whereas light levels from the sun on a bright day easily exceeds 200lx.

Our test data from Figure 3 has provided some interesting anomalies. For example, the basement recreation room and basement home office are rooms that have east facing windows and even with the curtains closed there is a gradual spike of light in those rooms. North facing rooms, like the north bedroom, do not display this anomaly. Another anomaly was the widening of light curve data. For instance, as winter turns to spring, and spring to summer there is a gradual lengthening of daylight time. Collection of this *day period* information [9] will thus also have value. A day period is a point in time that states when *sunrise*, *noon*, and *sunset* start. Using day periods helps us normalize the data by removing these seasonal changes.

We must take into account these anomalies when trying to accurately determine environmental occupant activity via ambient light ( $E^L$ ). A *profile matrix* becomes part of the solution; at each sensor location and during each day period, lux spikes may or may not be considered occupant activity. For example, a lux spike in a dark north facing room during the day would be activity but that same spike in a bright south facing room would not. If we let  $PM$  be the 2D *profile matrix* indexed by *sensor position* ( $n$ ), and the *day period* ( $p$ ),  $E_n^L$  be the lux reading from the ambient light sensor at  $n$ , then we can use the following formulation to determine occupant activity:

$$activity(n)? = \begin{cases} true, & \text{if } E_n^L \geq PM(n, p) . \\ false, & \text{otherwise} . \end{cases} \quad (4)$$

Now that we have introduced measurements associated with consumption and environment factors, we are now in a position to see if we can use these factors to allow a *rule-based* intelligent agent (IA) to come to a conclusion with respect to occupant (in)activity. The next section examines sleep detection, a challenging activity since it corresponds to a situation where occupants are in the house, but are inactive.

## V. SLEEP DETECTION

Our initial examination of sleep detection will use only ambient light sensors to determine the occupants' nightly sleeping period. Any implementation of such an algorithm can be seen as brittle because; (a) sleeping patterns from one home to another can be extremely different, and (b) even within a given household sleeping patterns change over time. One idea would be to look for a lux spike in certain ambient light sensors and create a general rule for our IA to follow. We use  $E_{SS}^L$  to denote the sensor used when occupants go to sleep (Sleep Start), and we use  $E_{SE}^L$  to denote the sensor used when occupants wake up (Sleep End). The following simple algorithm can be used to determine the sleeping period:

---

### Algorithm 1 Detect Sleep

---

**Require:**  $date, lux_{min}$

```

 $sleep_{start} \leftarrow 0$ 
 $sleep_{end} \leftarrow 0$ 
for  $t = 24 : 00 \rightarrow sunset(date), decrement$  do
  if  $E_{SS}^L \geq lux_{min}$  then
     $sleep_{start} \leftarrow t$ 
  end if
end for
for  $t = 00 : 00 \rightarrow noon(date), increment$  do
  if  $E_{SS}^L \geq lux_{min}$  then
     $sleep_{end} \leftarrow t$ 
  end if
end for
return  $sleep_{start}, sleep_{end}$ 

```

---

In the above algorithm we make two assumptions for simplicity that the occupants: (1) go to sleep before midnight, and (2) wake up before noon (12:00pm).

## VI. EXPERIMENTATION SETUP AND EVALUATION

We implemented HOA in a *real home* using *off the shelf products*. The advantages of choosing this solution was that there were no extensive renovations done to install smart light switches. The installation of the meters did require a non-temporary solution but that was the choice of the home owner. Figure 4 illustrates the different components of the HOA system.

Data is retrieved from a number of different sources at 15 minute intervals. Weather [10] and day period [9] data sets

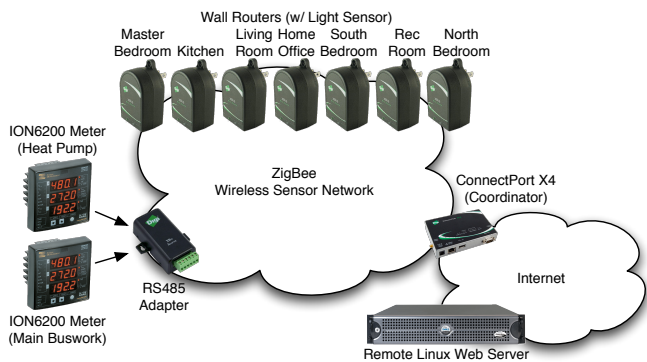


Fig. 4. The specific implementation of the generalized architecture. Data is collected by Python scripts executed by Cron on the remote server. Data is stored in a PostgreSQL database with graphical reports being generated in HTML5.

TABLE I  
ACCURACY RESULTS FOR DETERMINING OCCUPANCY

Occupancy By	$F_1$ Score	Precision	Recall
Power Consumption	1.000	1.000	1.000
Ambient Light Sensors	0.943	0.926	0.962
Nightly Sleeping Detection	0.892	1.000	0.806

are retrieved from external third party websites. Ambient light sensor data is collected over a ZigBee home area network (HAN). There are 7 sensors installed throughout our test home in various rooms. Power consumption data is collected from 2 meters over the same ZigBee HAN. All data is stored in a central database. HOA accesses the database and analyzes the historical data looking to see if anyone is home or not.

We collected data from January 22<sup>nd</sup> to August 28<sup>th</sup> 2010 for a home in Burnaby, Canada. This data was used as the input to algorithms to determine occupancy by 1) power consumption and by 2) light levels. We then used the data to determine if the occupants were asleep. The HOA results are then compared with correct results using the standard precision-recall methods. Table I shows how HOA performed.

It is not surprising that determining occupancy by power consumption is most accurate, particularly if you are able to cancel out background consumption from known non-occupant triggered events. Although we were able to obtain an f-score of 100%, this was only because of very detailed information being used to determine the threshold value in equation (3).

Determining occupancy from light sensor data did work very well. From a more detailed examination of our results, we found there were 19 out of 219 days where the day's precision or recall scores were below 80% due to unusual light levels. Unusual levels included extremely low sunlight levels, where the algorithm could not tell if it was daylight or light from indoor lamps. The intensity and brightness of natural light changes from season to season presented challenges for our algorithm. Like normalizing the changing lengths of days we would also need to normalize the day light intensity to get a better accuracy for both the ambient light sensor scores and the nightly sleeping detection scores.

The sleep detection algorithm performed quite well, with an f-score of 89%. However, it could easily be improved. Firstly,

polling every 15 minutes was too infrequent, if the occupants took less than 15 minutes to go to bed then the trigger for  $E_{SE}^L$  never fired causing inaccuracies. Secondly, some sensors were not optimally placed. For example, there was no clear line of sight from the sensor to the lighting source. We were limited in sensor placement because the sensor was a small unit that plugged directly into a wall outlet. Thirdly, there were anomalies of 1.2lx spikes during summer that caused the algorithm to infer that the sleep period ended prematurely.

## VII. RELATED WORK

Current research revolves around identifying what the occupant is doing. PlaceLab [1], [4] and MavHome [5], [6] both use sensors as a commodity. We are interested in a slightly different approach. We want to use the least amount of sensors possible to determine occupancy so that a home can run in a more energy efficient manner. Since each sensor consumes energy, it would not be beneficial if the cost of running a smart home system was more than the amount of savings the system could realize. We fully agree with MavHome's objective to *minimize operation cost* [6].

We feel the use of real homes is important in evaluating the performance IA systems. PlaceLab papers [1], [4] also agree with this point of view. PlaceLab and MavHome also require the occupants to record their activities as a way to test the performance of the agent to the actual events by using a Bayes learner [3], [4]. This is very time consuming and onerous to the occupant. Our HOA work is focused on identifying more general events. With more specific event determination we feel that unsupervised learning with a rule-based system might be a more appropriate avenue to pursue.

We want all sensors to be ubiquitous. We do not see that wearable sensors like RFID gloves as part of the solution in an average household. PlaceLab and others [2] have experimented with these. At this point the technology used in wearable sensors is cumbersome and not practical. Wilson *et al* also discuss why wearable sensors are not appropriate for people with disabilities [3]. PlaceLab [1] has used video cameras and microphones as sensors in their systems. We feel that these types of sensors are too costly and too intrusive as Abascal has pointed out [11].

Ambient light sensors in a home are susceptible to small changes in lighting levels. Mohamaddoust *et al*'s system only acted on changes of 50lx [12]. This is far greater than the level we have observed in our test home where light sources can change by as little as 3lx. Also their system was a simulated test in a wide open room with multiple light sensors. In real homes, there can be several small rooms in which multiple sensors would be too costly. Light from the sun and atmospheric changes had a dramatic effect on our ability to evaluate light levels, and is an important factor to consider. Mohamaddoust *et al* state the same for their LACS specifications but later on when testing do not consider these scenarios [12]. This again stresses the importance of testing research in a real home.

### VIII. CONCLUSIONS AND FUTURE WORK

With our key constraint of using a minimal number of inexpensive sensors, we were able to use power consumption and ambient light sensors to accurately detect occupancy and sleep periods. During our testing period, HOA found that on average a home was occupied 67.87% of the day; 36.72% of that was active occupancy and 31.15% was inactive (sleep) occupancy.

Seasonal effects (e.g. summer, winter), habitual effects (e.g. change in sleep pattern) and environment effects (e.g. light bulbs burning out changing light levels) can have performance impacts on the correctness of determining occupant activity. Identifying the *weak spots* of each the type of sensor is critical in identifying what different complementary sensors are needed to be deployed in future research. Careful and exhaustive *testing* of all algorithms we choose to implement is equally as important. Developing HOA with a *bottom-up* approach will allow us to incrementally improve the detection of occupancy. Some of these incremental improvements involve looking at other common, low cost sensors that we can deploy as part of our research.

There are many directions that we can take HOA. A high priority item for us is to incorporate a *learning-based* algorithm to complement our *rule-based* ones. Adding more intelligence presents opportunities to explore *Reinforcement Learning*, and *multi-agent systems* (MAS) or distributed intelligence. We feel that ultimately it would be better if more than one sensor could be used to determine occupancy activity as a way to confirm that the activity is truly from the occupant. For example, a power spike in the kitchen could be confirmed as activity because the PIR motion sensor was triggered as the occupant moved.

Lastly, we would like to evaluate and consider sensor performance due to sensor placement. Optimality was not considered in the current sensor layout as discussed above. We also need to look at determining the optimal frequency of collecting data. Using 15 minute intervals, as discussed above, is too infrequent and short events can be missed. We are also considering moving to more of a *push* system, where data is gathered more locally and the push out the the database server.

### ACKNOWLEDGMENTS

This work was supported in part by a grant from the National Sciences and Engineering Research Council of Canada (NSERC), along with the GRAND (Graphics, Animation and New Media) Network of Centres of Excellence of Canada.

### REFERENCES

- [1] S. Intille, K. Larson, E. Tapia, J. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson, "Using a live-in laboratory for ubiquitous computing research," *Pervasive Computing*, vol. 2006, no. 4, pp. 349–365, 2006.
- [2] M. Raad and L. Yang, "A ubiquitous smart home for elderly," *Information Systems Frontiers*, vol. 11, no. 5, pp. 529–536, 2009.
- [3] D. Wilson, B. Wolfe, and C. Atkeson, "Using many simple sensors for automatic monitoring in the home," *Robotics Institute, Carnegie Mellon, PA*, 2008.
- [4] E. Tapia, S. Intille, and K. Larson, "Activity recognition in the home using simple and ubiquitous sensors," *Pervasive Computing*, pp. 158–175, 2004.

- [5] D. Cook, M. Youngblood, E. Heierman III, K. Gopalratnam, S. Rao, A. Litvin, and F. Khawaja, "Mavhome: An agent-based smart home," in *Pervasive Computing and Communications, 2003.(PerCom 2003). Proceedings of the First IEEE International Conference on*. IEEE, 2003, pp. 521–524.
- [6] S. Das, D. Cook, A. Battacharya, E. Heierman III, and T. Lin, "The role of prediction algorithms in the mavhome smart home architecture," *Wireless Communications, IEEE*, vol. 9, no. 6, pp. 77–84, 2003.
- [7] K. Moslehi and R. Kumar, "Vision for a self-healing power grid," *ABB Review*, pp. 21–25, 2006.
- [8] —, "Smart grid-a reliability perspective," in *Innovative Smart Grid Technologies (ISGT), 2010*. IEEE, 2010, pp. 1–8.
- [9] "Sunrise and sunset calculator," 2011. [Online]. Available: <http://www.timeanddate.com/worldclock/sunrise.html>
- [10] "Environment canada weather office," 2011. [Online]. Available: [http://www.weatheroffice.gc.ca/canada\\_e.html](http://www.weatheroffice.gc.ca/canada_e.html)
- [11] J. Abascal, "Ambient intelligence for people with disabilities and elderly people," in *ACM's Special Interest Group on Computer-Human Interaction (SIGCHI), Ambient Intelligence for Scientific Discovery (AISD) Workshop, Vienna*. Citeseer, 2004.
- [12] R. Mohamaddoust, A. Haghight, M. Kafash, and M. Sharif, "Designing the lighting control system based on wsn with optimization of decision making algorithm," in *The International Conference on Computational Intelligence and Communication Networks (CICN2010)*, 2010.



**Stephen Makonin** is a PhD student in the School of Computing Science at Simon Fraser University. He is also an applied researcher in the Technology Centre at the British Columbia Institute of Technology. His research interests focus on intelligent systems in smart homes for energy conservation and sustainability. He often blogs about his research at <http://eco-sustain.org/>.



**Fred Popowich** is a Professor of Computing Science at Simon Fraser University, Director of the Vancouver Institute for Visual Analytics, and member of the Institute for Computing, Information and Cognitive Systems at the University of British Columbia. He received his PhD in Cognitive Science from the University of Edinburgh in 1989. His research involves several aspects of Artificial Intelligence with an overall objective to make it easier for human beings to interact with computers, or to make it easier for human beings to interact with each other. He is also active in technology commercialization.