



30th Annual Workshop of  
the Swedish Artificial Intelligence Society  
SAIS 2017

May 15–16, 2017  
Karlskrona, Sweden

Editor: Niklas Lavesson  
Blekinge Institute of Technology, Sweden

Linköping Electronic Conference Proceedings No. 137  
ISSN: 1650-3686  
eISSN: 1650-3740  
ISBN: 978-91-7685-496-9

## Preface

The Swedish Artificial Intelligence Society (SAIS) was established on June 11, 1982. However, the first meeting was held on Jan 30, 1981 in Uppsala. Since then, SAIS has strived to arrange an annual national workshop, primarily to serve as a platform for junior researchers and doctoral students to present their work and for senior researchers and professionals to discuss topics, challenges, and opportunities related to artificial intelligence research and innovation in Sweden.

The SAIS workshop now returns to Blekinge for the first time in 15 years. SAIS 2002 was organized by then local chair Professor Paul Davidsson at Blekinge Institute of Technology in Ronneby. This year, the 30th Annual Workshop, SAIS 2017, is held on May 15–16 at Blekinge Institute of Technology in Karlskrona.

The technical program of the workshop contains 12 presentations of peer-reviewed papers submitted, with topics covering a variety of basic research directions in artificial intelligence and related subject areas as well as applications.

The workshop program features keynote presentations by distinguished invited speakers Professor Johannes Fürnkranz (Technische Universität Darmstadt), Associate Professor Tina Eliassi-Rad (Northeastern University), and Tomas Akenine-Möller (Lund University, NVIDIA).

# Contents

Preface . . . . .	3
Organization . . . . .	5
Local Organization . . . . .	5
Program Committee . . . . .	5
Program . . . . .	6
Monday, May 15 . . . . .	6
Tuesday, May 16 . . . . .	7
Papers . . . . .	9
Long-Term Accuracy in Sea Navigation without using GNSS Systems <i>Mårten Lager, Elin Anna Topp, and Jacek Malec</i> . . . . .	10
Modelling regimes with Bayesian network mixtures <i>Marcus Bendtsen and Jose M. Peña</i> . . . . .	20
PastVision: Exploring "Seeing" into the Near Past with a Thermal Camera and Object Detection-For Robot Monitoring of Medicine Intake by Dementia Patients <i>Martin Cooney and Josef Bigun</i> . . . . .	30
Multi-expert estimations of burglars' risk exposure and level of pre-crime preparation based on crime scene data <i>Martin Boldt, Veselka Boeva, and Anton Borg</i> . . . . .	39
Advanced Data-driven Techniques for Mining Expertise <i>Milena Angelova, Veselka Boeva, and Elena Tsiporkova</i> . . . . .	45
Multi-Task Representation Learning <i>Mohamed-Rafik Bouguelia, Sepideh Pashami, and Slawomir Nowaczyk</i> . . . . .	53
Improved Inter Terminal Transportation using Agent Technology <i>Lawrence Henesey</i> . . . . .	60
Energy Efficiency in Machine Learning: A position paper <i>Eva Garcia-Martin, Niklas Lavesson, Håkan Grahn, and Veselka Boeva</i> . . . . .	68

## Organization

### Local Organization

**Chair:** Niklas Lavesson

**Coordinator:** Eva-Lotta Runesson

**Supporting staff:** Eva García Martín

### Program Committee

**Marcus Bjärelund**, AstraZeneca

**Eva Blomqvist**, Linköping University

**Anton Borg**, Blekinge Institute of Technology

**Paul Davidsson**, Malmö University

**Göran Falkman**, University of Skövde

**John Folkesson**, Royal Institute of Technology

**Johan Hagelbäck**, Linnaeus University

**Fredrik Heintz**, Linköping University

**Anders Holst**, Swedish Institute of Computer Science

**Magnus Johnsson**, Lund University

**Lars Karlsson**, Örebro University

**Jonas Kvarnström**, Linköping University

**Jacek Malec**, Lund University

**Slawomir Nowaczyk**, Halmstad University

**Panagiotis Papapetrou**, Stockholm University

**Cecilia Sönströd**, University of Borås

## Program

### Monday, May 15

**09:00 - 09:30 Coffee and registration - outside of Room J1610**

**09:30 - 09:50 Welcome and general information - Room J1610**

**09:50 - 10:50 Invited talk: Tomas Akenine-Möller (Lund University, NVIDIA) - Room J1610**

Chair: Niklas Lavesson, Blekinge Institute of Technology

**10:50 - 11:10 Fredrik Heintz. AI in Sweden: challenges and opportunities - Room J1610**

Chair: Håkan Grahn, Blekinge Institute of Technology

**11:10 - 11:30 Maj Stenmark, Mathias Haage, and Elin Anna Topp. Simplified Programming of Re-usable Skills on a Safe Industrial Robot — Prototype and Evaluation – Room J1610**

Chair: Håkan Grahn, Blekinge Institute of Technology

**11:30 - 13:00 Lunch - Room J1504**

**13:00 - 14:00 Invited talk: Tina-Eliassi Rad (Northeastern University) - Room J1504**

Chair: Paul Davidsson, Malmö University

**14:00 - 15:00 Session I - Room J1610**

Chair: Fredrik Heintz, Linköping University

- Martin Cooney and Josef Bigun. PastVision: Exploring "Seeing" into the Near Past with a Thermal Camera and Object Detection – For Robot Monitoring of Medicine Intake by Dementia Patients
- Awais Ashfaq, Markus Lingman, Miltiadis Triantafyllou, Slawomir Nowaczyk, and Anita Sant'Anna. Unravelling healthcare complexity - a machine learning perspective
- Milena Angelova, Veselka Boeva, and Elena Tsiporkova. Advanced Data-driven Techniques for Mining Expertise

**15:00 - 15:30 Coffee break**

**15:30 - 16:00 SAIS Master's thesis award - Room J1610**

The SAIS board is happy to award Thomas Rosenstatter, Halmstad University, the 2017 SAIS Master's Thesis Award, with the following motivation:

*The award goes to Thomas Rosenstatter for his thesis entitled "Modelling the Level of Trust in a Cooperative Automated Vehicle Control System". The thesis introduces a trust system that allows an autonomous vehicle, in this case a car, to make more reliable and robust decisions by taking into account current information about its context. The system evaluates the current situation and generates a trust index, indicating the level of trust in the environment, the vehicle itself, and other vehicles surrounding it. This work was partly evaluated as part of the winning team in the Grand Cooperative Driving Challenge 2016. The thesis addresses a topic that is timely and of high practical relevance in today's AI community. The thesis is well written and has the potential of both impacting future research in the field, and practical applications.*

**16:00 - 17:00 SAIS annual meeting - Room J1610**

**18:30 - 21:00 Dinner**

Restaurant Landbron, Landbron 1, 371 33 Karlskrona

**Tuesday, May 16**

**10:30 - 11:30 Session II - J1610**

Chair: Anders Holst, SICS

- Marcus Bendtsen and Jose M. Peña. Modelling regimes with Bayesian network mixtures
- Mohamed-Rafik Bouguelia, Sepideh Pashami and Slawomir Nowaczyk. Multi-Task Representation Learning
- Eva García-Martín, Niklas Lavesson, Håkan Grahn, and Veselka Boeva. Energy Efficiency in Machine Learning

**11:30 - 13:00 Lunch at BTH - Room H402B**

**13:00 - 14:00 Invited talk: Johannes Fürnkranz (Technische Universität Darmstadt) - Room J1610**

Chair: Jacek Malec, Lund University

**14:00 - 15:00 Session III - Room J1610**

Chair: Johan Hagelbäck, Linnaeus University

- Yuantao Fan, Pablo Del Moral, and Slawomir Nowaczyk. Evaluation of Two Level Classifier for Predicting Compressor Failures in Heavy Duty Vehicles
- Bin Sun, Wei Cheng, Prashant Goswami, and Guohua Bai. An Overview of Parameter and Data Strategies for kNN-Based Short-Term Traffic Prediction
- Lawrence Henesey. Improved Inter Terminal Transportation using Agent Technology

**15:00 - 15:30 Coffee break**

**15:30 - 16:30 Session IV - Room J1610**

Chair: Slawomir Nowaczyk, Halmstad University

- Martin Boldt, Veselka Boeva, and Anton Borg. Multi-expert estimations of burglars' risk exposure and level of pre-crime preparation based on crime scene data
- Mårten Lager, Elin Anna Topp, and Jacek Malec. Long-Term Accuracy in Sea Navigation without using GNSS Systems
- Radu-Casian Mihailescu and Paul Davidsson. Integration of Smart Home Technologies for District Heating Control in Pervasive Smart Grids

**16:30 Workshop closing**



## Papers

# Long-Term Accuracy in Sea Navigation without using GNSS Systems

Mårten Lager\*, Elin A. Topp, Jacek Malec†

May 8, 2017

## Abstract

Many ships today rely on Global Navigation Satellite System (GNSS), for their navigation, where GPS (Global Positioning System) is the most well known. Unfortunately, the GNSS systems make the ships dependent on external systems, which can be malfunctioning, be jammed or be spoofed.

There are today some proposed techniques where, e.g. bottom depth measurements are compared with known maps using Bayesian calculations, which results in a position estimation. Both maps and navigational sensor equipment are used in these techniques, most often relying on high accuracy maps, with the accuracy of the navigational sensors being less important.

Instead of relying on high accuracy maps and low accuracy navigation sensors, this paper presents an idea of the opposite, namely using low accuracy maps, but compensating this by using high accuracy navigational sensors and fusing data from both bottom depth measurements and magnetic field measurements.

## 1 Introduction

Finding the way over great seas has been important for thousands of years. The compass was invented almost a thousand years ago, and the first nautical sea charts were used in Italy in the 13th century. With a compass and a nautical sea chart, it is possible to perform *dead reckoning* to estimate the current location based on the previous location and

the compass direction. In this technique, it is also possible to compensate for the vessel's drift and sea current. But if one is not able to estimate the drift and sea current accurately, the position error starts increasing, as each estimation of the position is relative to the previous one, which means that the position error is accumulated over time. This deficiency can be overcome in different ways. By regularly determining the position compared to known landmarks, the accumulation of error is reset. But if landmarks cannot be found because one is on open water, either there is a need for increasing the accuracy of the dead reckoning by using better equipment (e.g. compass, logs (for speed), gyro, accelerometers, inertial sensors), or there is a need to use information about the environment that can be seen out on open waters. During the 18th century the celestial navigation was invented, which uses angle measurements to the sun, moon and stars to greatly improve the long-term accuracy of navigation. Nowadays the celestial navigation has almost completely been abandoned, because GNSS can determine the position accurately and efficiently. The most common and oldest GNSS system is GPS, but there are also other systems, e.g. Galileo and Glonass.

The GNSS systems have made it very simple to determine a vessel's position with good accuracy, but there are still some disadvantages. One important disadvantage is that the ship needs to rely on external information from the GNSS satellites which is sent to the GNSS receiver onboard. It is quite simple to jam the radio reception from the GNSS satellites, which results in that it is not possible to determine the position any more. Even worse, it is possible to spoof the GNSS transmission information with advanced equipment, resulting in that an incorrect position is provided [8].

---

\*M. Lager is an industrial Ph.D student with Saab Kockums in Malmö and with the Department of Computer Science, Lund University, Sweden.

†E.A. Topp and J. Malec are with the Department of Computer Science, Lund University, Sweden.

If a dependency on the GNSS system is not desired, and a more modern and a less time consuming technique than the celestial technology from the 18th century is wanted, there might be some alternatives. Some alternatives can be found in the following research papers.

Reference [2] describes how systems on an airplane measure the elevation and compare it to a known digital elevation map with a *Particle Filter* (PF) algorithm. By doing so, the algorithm's estimation of the position eventually converges to the correct position. A similar technique is used in [3] for surface and underwater navigation, where the bottom depth is compared to a high accuracy bottom map with a PF algorithm. The same paper also describes how almost the same PF algorithm can be used to estimate the position by comparing measured distances to the surrounding shore line with a map of the same area. There is other information which can be used by particle filters for positioning. The earth magnetic field surrounds the earth, and is disturbed by ferromagnetic elements. In an indoor environment, these disturbances are normally bigger than the earth magnetic field itself [5], and both [4] and [5] suggest how to estimate a position in an indoor environment with a PF comparing magnetometer measurements to a known magnetic map.

This paper presents an idea of how to perform sensor fusion based on various types of PF calculations in order to estimate a ship's position. By using various types of measurements for the PF calculations, and by relying on high accuracy navigation sensors, the probability of being able to obtain the current position without having to rely on high accuracy maps or GNSS data is high.

This paper is organized as follows: In Section 2, first a brief discussion is given about Bayesian calculations and how these can be used to estimate the position of a ship. Then Particle Filters (PF) are explained more in detail, and it is described how e.g. bottom depth and magnetic fields can be used for the Correction Step of the PF. Based on the current available research, limitations and opportunities of this research is described in Section 3. In Section 4 an idea of how to implement an algorithm that correct the limitations is given. Further, what has been implemented so far and what is left of the implementation is presented. In Section 5, concluding remarks are given.

## 2 Probabilistic position estimation

### 2.1 Bayesian position estimation

The key problem we have is that we would like to estimate the position, but are not able to measure the position directly. Instead we can measure other information, such as how the ship is moving, bottom depth and magnetic field data. This can be modeled as a *Hidden Markov Model* (HMM), where the state (i.e. the position, attitude, velocity and acceleration) influences the data which can be measured. Figure 1 illustrates how the state  $x(t)$  (i.e. position) influences the data which can be measured, denoted by  $y(t)$ .




Figure 1: Each state ( $x(t)$ ) in an HMM influences the data which can be measured ( $y(t)$ ). The position is not possible to measure directly, but the position will influence which bottom depth and magnetic field vector that is measured. The state in time  $t$  ( $x(t)$ ), contains the state position, attitude, and acceleration. The measured quantities ( $y(t)$ ) which are influenced by the state, are attitude measurements, acceleration measurements, bottom depth measurements, magnetic field measurements, etc.

To get the best possible estimation of the state, not only the present measurement is to be analyzed, but all previous data. The equations for calculating the probability of being in one state and going to another state given the measurements at time  $t$ , are given as follows:

$$p(x_{t+1} | \mathbb{Y}_t) = \int_{\mathbb{R}^n} p(x_{t+1} | x_t) p(x_t | \mathbb{Y}_t) dx_t \quad (1)$$

$$p(x_t | \mathbb{Y}_t) = \frac{p(y_t | x_t) p(x_t | \mathbb{Y}_{t-1}) dx_t}{p(y_t | \mathbb{Y}_{t-1})} \quad (2)$$

These equations are not analytically solvable, and therefore a filter will instead be used for estimation of the position. If the measurements and transition functions would have been linear and the measurement and process noise Gaussian, a Kalman Filter would have been the optimal choice to compute the position [6]. In our case the transition functions are non-linear and the measurements have no Gaussian distribution, but instead a highly multi-modal distribution. There are some non-optimal extensions to Kalman Filters to handle the issues with non-linearity and not having a Gaussian distribution [1]. However, PF are more flexible and have a built-in capability to handle multi-modal distributions. Therefore, the PF algorithm will be used in this paper.

## 2.2 Particle Filters for estimation of the position

A *Particle Filter* (PF) is a Bayesian sequential Monte Carlo method (SMC). It keeps track of an object through a *Probability Density Function* (PDF), which may be non-Gaussian and even multi-modal [1, 7].

The objective of the PF is to evaluate  $p(\mathbb{X}_t | \mathbb{Y}_{0:t})$ , where  $\mathbb{X}_t$  is the vector of all available states in the time  $t$ , and  $\mathbb{Y}_{0:t}$  are all measurements up to the time  $t$ . Instead of directly calculating  $p(\mathbb{X}_t | \mathbb{Y}_{0:t})$ , what has happened before time  $t$  is modeled into a large set of particles, where the number of particles in each location and their weights estimate how likely the position is. For each new time step, the  $p(x_t | \mathbb{Y}_t)$  is calculated for each particle.

Initially, the PF algorithm starts with a large number of random samples (particles), where each particle is given a weight that characterizes its quality. At the beginning, each particle has the same weight. The estimated state is given by calculating the weighted sum of all particles. There are three important steps in the PF algorithm:

- Prediction
- Correction (Filtering)
- Re-sampling

During the *Prediction Step*, each particle is moved according to a random value of the state model including the modeled noise. In our case,

we have a good idea of how the ship is moving because of the navigational sensor equipment, which measure the ship's Reference Data (RD), estimating both the position via dead reckoning calculation and the orientation. We also know the noise that these sensors have. This will give us a PDF of to where the state has moved, and for each particle we then pick a random value from that distribution.

During the *Correction step* the weight of each particle is regenerated according to the sensor readings that can validate the probability of each state. In our case, we use the bottom depth and/or magnetic field vector compared to maps, to estimate how likely it is that each particle is in the correct position. This step is also known as the filtering step.

During the *Re-sampling step* new particles are re-sampled randomly according to the PDF (including weights) of the old particles. By this step, there will be many new particles in states where the *Correction step* has judged the probabilities for the particles to be high, and few where the probabilities were low. The old particles from the previous step will not be used any more, and can now be discarded.

One cycle with the *Prediction step*, *Correction step* and *Re-sampling step* is now complete, and the algorithm continues with iterations for the newly sampled particles, see figure 2.

The complete algorithm looks as follows:

### 1. Initialization

- $t = 0$
- Generate N initial samples with an initial distribution of the position.

### 2. Prediction:

Predict how the particles are moving to the next position regards to the RD.

### 3. Correction (Filtering):

Compute the weights for each particle and normalize the weights, i.e. compute how likely it is that the particle is positioned where it is, regards to measured bottom-depth.

### 4. Re-sampling:

Generate a new set of N particles according to how the previous particles are distributed including their weights.

### 5. Increase t, and **iterate** to step 2.

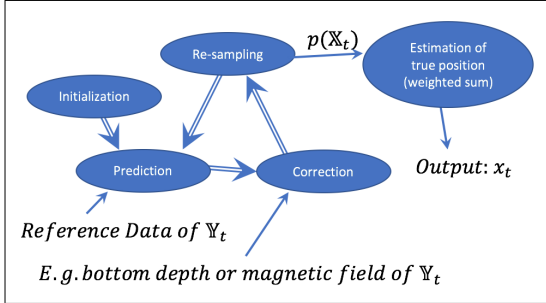


Figure 2: A block diagram of a Particle Filter. After initialization, the particle filter iterates between *Prediction*, *Correction* and *Re-sampling*. After the *Re-sampling Step*, a probability of each state  $p(\mathbb{X}_t)$  is estimated from the particles. From this, a weighted sum of all particles can be used as an estimation of the true position.

### 2.3 Depth data in the particle filter

Reference [2] describes how elevation can be used in the PF for navigation of an airplane. In [3] almost the same technique is used for the domain of naval ships, where a PF is used for estimating the position based on how the bottom depth measurements varies when moving in a trajectory.

The state of the system is denoted as  $x_t$ . The state can contain different variables depending on what sensors are available and how complex we want the algorithm to be. [3] suggests using the Cartesian position  $(X, Y)$  and the crab angle  $\delta$  for  $x_t$  (the crab angle  $\delta$  is the angle between the direction the ship is pointing towards and the direction of the velocity of the ship).

$$x_t = (X_t \quad Y_t \quad \delta_t)^T \quad (3)$$

The following equation with discrete time with the sample time  $\Delta$ , model the state:

$$x_{t+1} = f(x_t, u_t, w_t) = \begin{pmatrix} X_t + v_t \Delta \sin(\varphi_t - \delta_t) \\ Y_t + v_t \Delta \cos(\varphi_t - \delta_t) \\ \delta_t \end{pmatrix} + w_t \quad (4)$$

In this equation  $u_t = (v_t \quad \varphi_t \quad \theta_t \quad \phi_t)^T$  is the input signal, which consists of the speed  $v_t$ , compass angle  $\varphi_t$ , the sensor azimuth angle  $\phi_t$  relative to the stern of the ship and the sensor elevation relative to vessel  $\theta_t$ . The  $w_t$  is the process noise. The

range to the sea floor is measured in the direction from the sonar sensor. The measurement relation is given by the following equation [3]:

$$y_t = h(x_t, u_t) + e_t = r(x_t, \phi_t + \varphi_t + \theta_t) + e_t \quad (5)$$

where  $r(x_t, \phi_t + \varphi_t + \theta_t)$  is the range measured from the position of  $x_t$  with the azimuth angle  $\phi_t + \varphi_t$  and elevation angle  $\theta_t$ . A sonar sensor that measures the range to the bottom typically has a fixed elevation and azimuth, where the sensor normally is pointing straight downwards.

We now have the model for how to go from one state to the next state in function (4) and we have the model for how the measured value  $y_t$  depends on the state in function (5).

If they would have been analytically solvable, the Bayesian calculations in function (1) and function (2) would have been used for calculating the position given all history. Now that they are not, PF instead is used, according to the algorithm in section 2.2.

### 2.4 Magnetic data in the particle filter

The earth is surrounded by a magnetic field, a phenomenon, which has been used by compasses for many decades. The compass needle points towards the magnetic north, which might give the user the idea that the magnetic field is horizontal to the surface of the earth. The magnetic field is in fact more accurately represented by its *declination* and *inclination* [4]. The declination describes the horizontal deviation of the magnetic field, and it is this field which is measured by the compass. The inclination describes the vertical deviation of the field, and this is more or less neglected in the compass by, e.g., arranging the compass needle on a floating device on a water bed. The magnetic field also varies depending on the time of the day, but the fluctuations are relatively small with fluctuations between 10 nT and 30 nT, which is less than 0.1% of the average magnitude of 48.19  $\mu$ T [4].

Each ferromagnetic element disturbs this magnetic field, and these disturbances can for indoor environments be even greater than the natural magnetic field of the earth [5]. For ships, the local ferromagnetic elements onboard the ship disturb the compass. When navigating, bigger ships

often have had a binnacle, which has two movable compensating magnets trying to compensate for the magnetic disturbances of the ship. There are also techniques for manually correcting the compass course depending on which direction the compass points towards. For a long time, the compass direction has been the desired sensor measurement, and the disturbances the thing that is to be minimized.

In [5] and [4], the disturbances are instead considered as a signal rather than as noise. For indoor environments, the many ferromagnetic elements create a complex magnetic environment where the magnetic vector varies greatly depending on where the sensor is located. The magnetic field is also quite stable if no major furniture or iron walls are moved. In [5] and [4] all three dimensions of the magnetic field vector are considered, i.e. not only the magnetic intensity. This information is compared to a magnetic map with a PF, and in conjunction with some sort of odometry, such as wheel encoders or inertial sensors, it has in [4] been possible to precisely localize a human or robot. In [5] only cheap smartphone sensors are used, where the 3-axis magnetic field and acceleration are used for determining the position of the user.

In [4], the following PF algorithm is presented.

#### 1. Initialization

- Generate N particles and give them a random starting position, heading and drift rate.

#### 2. Prediction - For each particle:

- Increase/decrease the drift rate, and update the heading according to the drift.
- Update the heading according to measured heading changes.
- Update position according to traveled distance and the heading.

#### 3. Correction

- Check if each particle is within the mapped area, and if so, calculate the weight of each particle. The weight is calculated by a likelihood function that compares the difference between the magnetic field in the map, and the measured magnetic field.

- Normalize the particle weights to sum 1.

#### 4. Re-sampling

- Resample the particles

#### 5. Iterate to step 2.

In [4], there are three different alternatives for the likelihood functions calculating the weights. The simplest function only measures and compares the magnetic intensity, the second measures and compares the horizontal and vertical magnetic field component, and the third measures and compares the full 3-dimensional magnetic vector. As can be expected, [4] shows that the third algorithm performs better than the second one, and the second algorithm performs better than the first one. Especially the robustness and the time for filter convergence have improved when going for higher dimensions.

Although [4] and [5] have explored indoor environments, the same algorithms are applicable for outdoor environments. The magnetic field does not fluctuate as fast as in indoor environments, but on the other hand it is more stable, because no furniture or building parts are moved around as in the indoor environments. There are satellite maps available covering the entire magnetic field of the earth, and in some areas of the world, accurate magnetic field maps have been created. Therefore, the magnetic field is a good candidate to be used for the PF algorithm when estimating a ship position, at least as a complement to the bottom depth.

### 2.5 Using other data in the particle filter

The bottom depth and the magnetic field are good candidates to use for the PF algorithm when estimating the position, but there are other alternatives. In addition to the bottom depth, [3] also uses range measurements to land objects in another PF algorithm. This range is measured by a radar, and is compared to a sea chart database.

It is also possible to not only use the depth directly to the bottom. If the ship is equipped with a sonar system, it is also possible to use multiple bottom depth measurements covering a large area at once. This will increase the performance of the PF even further, as it will be possible to evaluate if

the bottom readings matches the map with better precision.

The strength of the PF algorithm is that it is very flexible when it comes to which measurements to use. The important thing is that the measurements shall vary enough when changing position, and that it shall have varied in the same way (or in a predicted way) when the map was created, and when doing the PF measurement. Other candidates which could be used for the PF algorithm are:

- Celestial navigation items such as star positions, where a star either is present in a proposed direction, or is not.
- Gravitation, which vary depending on where the ship is located on the earth.
- Various types of available bearing measurements, depending on which sensors the ship is equipped with. For instance, radio and radar sources with known map locations can be used, if the ship's sensors are able to estimate the bearing to that kind of sources.

### 3 Limitations with current research

The papers referenced from this paper show that it is possible to do accurate position estimations if having sensors measuring data which has previously been mapped into accurate maps. Many of the references also evaluate how accurate the position estimation can become, when already having accurate maps available. However, there are some limitations with the current research.

The algorithms proposed in the studied research papers require that there are highly accurate maps. This is not the case out on open water, and not even in most coastal areas. The reality is that different areas have been mapped with various accuracy, where highly trafficked areas more often have better accuracy than less trafficked areas. The algorithm for positioning in e.g. [4] assumes that it can get the true bottom depth in any position of the map, but from a normal sea chart it is more likely that it is possible to compute some sort of likelihood distribution of the bottom depth for each position. In figure 3, a 1000 m wide part of a sea chart is

presented as an example with the position of interest marked with an X. In figure 4 an example of the bottom depth likelihood distribution is presented, which gives the algorithm an estimation of the bottom depth, when no accurate bottom depth is available. It should also be noted, that when creating a sea chart, the most important thing is that there is no shallower area in the map than what is presented. If there are indications that there are bottom depths of 15 m, 18 m and 22 m in an area, it is quite unlikely that there are any depths of 10 m in the middle of these indications, but there could be bottom depths of 30 m.

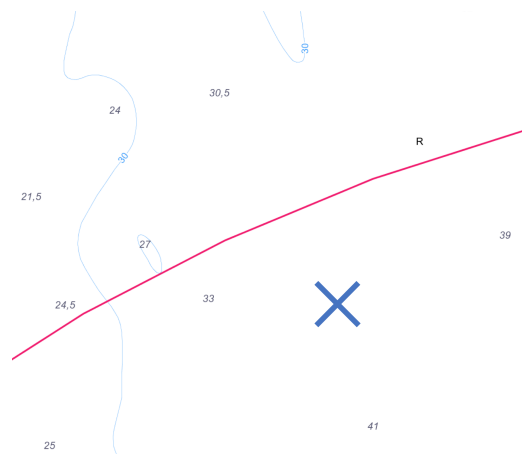


Figure 3: An example of a sea chart and the current position of interest marked with a blue X. (The red line marks the area surrounding a lighthouse in the sea chart.)




Figure 4: An example of how a bottom depth likelihood distribution could look like for the position of interest in figure 3.

The user platforms that are most likely to have a need for a system for accurate position estimation techniques which eliminates the need for GNSS

systems, are not a cheap platform with moderate navigation systems. The most probable platform is instead an advanced platform with accurate and expensive navigation sensors, where the RD, bottom depth and speed can be measured with high accuracy. The extra expense for buying and integrating a magnetometer sensor for measuring also the magnetic field vectors are foreseen to be a minor investment on platforms like this.

The current research in this field has mainly focused on achieving good performance of the positioning systems when having limited performance of the sensor suite, but nearly unlimited accuracy of the map. In this paper, the goal is to investigate if it is possible to do it the other way around. The main research question is therefore: **Is it possible to navigate accurately enough without GNSS systems, only relying on high performance navigation sensors and normal sea chart and magnetic charts?** In some areas, the sea charts are accurate, and in other areas the magnetic charts can be more accurate. It would therefore be a good feature to be able to estimate the position both with magnetic measurements and bottom depth measurements at the same time, and weighting the fusion between the two, depending on which data that has the highest confidence level. When both maps are not accurate enough, dead reckoning can instead be used for a while until going into areas with enough accuracy in the maps.

## 4 Proposal of a new implementation

### 4.1 Proposed Implementation

The future goal is to implement something in line with the following algorithm:

1. **Initialization** - Generate N particles and give them a random starting pose around a manual estimation of the starting position.
2. **Prediction Step** - Move each particle according to the ship's total navigation sensor suite measurements including their probability distributions.
3. **Correction Step** - Calculate the weight for each particle given all available sensors and

maps.

4. **Re-sampling Step** - Re-sample the particles.
5. Iterate to step 2).

The most interesting step in the algorithm is the *Correction Step*, which will need the following parts to work:

#### 4.1.1 Map data

To make the PF algorithm operable, there must be data supporting the likelihood calculations, which estimates how likely it is that the current measurement has been performed at each location. If the goal is that the PF should be possible to use on most places, it is important not to require high accuracy maps. The best is if it is possible to use the best available information in each area, i.e. high accuracy maps where available and normal sea charts for bottom depth information if only those are available. To support the PF algorithm, a function is needed to create bottom depth values from a sea chart database, including confidence estimations. If the algorithm e.g. can read surrounding bottom depth coordinates and surrounding bottom depth curves (see figure 3), an estimation of the bottom depth and confidence estimations can be calculated. In the simplest form, only a single bottom depth estimation is given from the function, and almost as simple as this would be to let the sea chart give an interval of valid bottom depths, e.g. 20-30 meters. The best support for the PF algorithm would be to give an accurate PDF estimation (see figure 4) based on knowledge of how sea charts are created. The same type of algorithm would be needed for magnetic data.

#### 4.1.2 Fusion of sensor data

On an advanced ship with high precision navigation sensors, it can be acceptable to use the navigation sensors for dead reckoning without using global positioning techniques for some time. It will take a long time before the drift has become large enough for resulting in a completely inaccurate position. When using the PF algorithm to correct the position, it is therefore important to not spoil the advantages of the already well working navigation system, by lean too much against the estimation of



the position from the PF compared to the dead reckoning algorithm. If e.g. there is uncertain evidence that estimates a particle is in the wrong position, it is better to let the particle remain, than removing it. The worst thing that could happen is if all particles at the correct position eventually is discarded, which could happen if the local measurements have not been accurate enough, the maps are not accurate enough, or the map measurements have changed, e.g. due to some external effect. To meet this challenge, we propose dividing the particles into subsets in the beginning of each Correction Step. Then the particles in each subsets are corrected according to the correction rule in the particular subsets. There are different alternatives of how to divide the particles into sub-groups, when using magnetic and bottom depth data for the PF algorithm. We propose the following alternatives:

1. Divide the particles into three subsets, where one subset of particles will be weighted according to a PF algorithm working with bottom depth, one subset of particles will be weighted according to a PF algorithm working with magnetic fields, and the last subset will have equal weights, where only the dead reckoned position from the RD matters. The size of each subset can then be determined by the quality of the bottom depth and magnetic maps/measurements compared to RD accuracy. The advantage is that e.g. bad magnetic measurements or maps not will damage the subsets where magnetism is not taken into consideration. The drawback is that it will take longer time before the PF converges to the correct position.
2. Divide the particles into two subsets, where one subset of particles will be weighted according to a PF algorithm working with both bottom depth and magnetic fields at the same time, and the other subset will have equal weights, where only the RD matters. The size of each subset including its internal subsets can be determined by the quality of the bottom depth and magnetic maps/measurements. By combining both magnetic fields and bottom depth into  $\mathbb{Y}$ , the PF will be able to calculate the probability density  $p(x_t | \mathbb{Y}_t)$  very efficiently. The drawback is that particles can

be discarded incorrectly if any measurements or the maps are inaccurate.

3. By combining 1) and 2) and having four sets of particles, the advantages can be taken from each solution. If bottom depth is better than magnetic fields in one area, the particles can e.g. be divided according to table 1.

Table 1: Example distribution of particles

Subset	Nbr of particles
Depth and magn. field	50%
Only bottom depth	20%
Only magnetic field	10%
No PF (only RD)	20%

In this way, the strength in combining data to support the PF are used by half of the particles. The other half of the particles are more carefully used, so that some particles will survive even if local measurement errors occur or maps are inaccurate.

## 4.2 Present situation of the implementation

To start the investigation of the possible solutions, a Python program has been created to explore the possibilities with using depth measurements for the PF, see figure 5 (left). In this program, a sea chart is digitized into 10x6 squares by manually setting a lower and upper boundary of the bottom depth in each square. In the initialization of the program, the ship is placed in an (for the algorithm) unknown position, marked with a green dot. The program then iterates through the following algorithm:

1. **Initialization** - Generate 1000 particles and put them into a random square of the 60 available squares. This can be seen in figure 5 (left), where some small blue dots can be seen in the middle of each square, where the size of the dots indicate how many particles are located in each square.
2. **Prediction Step** - The user then moves the ship by pressing some direction and speed in

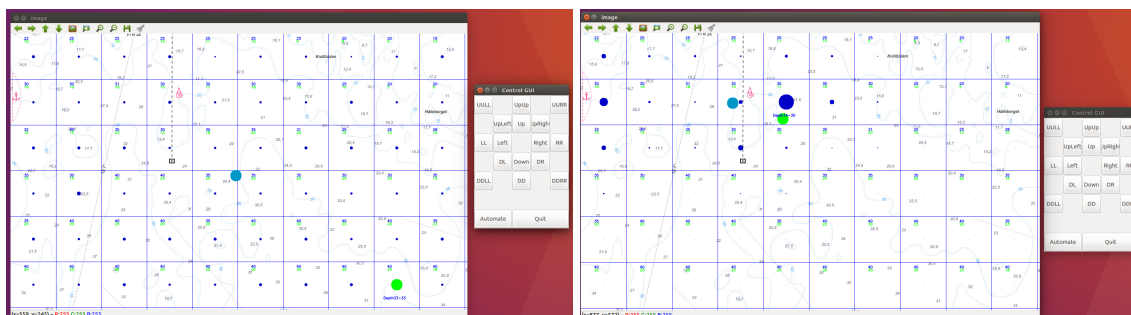


Figure 5: (left) In the initial state, the particles (indicated by the size of the blue dots) are randomly distributed between all 60 squares. The light blue dot in the middle of the sea chart indicates the weighted sum of all particles, and is the current estimation of the true position of the ship. The green dot to the lower right, indicates where the ship is located. The ship can then be moved by clicking the GUI. (right) After 8 random steps, the square with the highest number of particles is the same square as the ship. The weighted sum of all particles (the estimation of the position), is still in the wrong square.

the GUI (or alternatively the Automate button for a random movement). A random error is added to both the direction and speed, and then the ship is moved according to those values. Each particle is also moved according to an estimation of the random movement. (If the ship or particles hit the boarder of the map, they stay on the border.)

3. **Correction Step** - When the ship has moved, it will measure 10 random bottom depths in the square, and removing the biggest and smallest values in order to increase the resistance against error measurements. Then it starts comparing each particle's square's minimum and maximum depth to the measured values by the ship. If the measurements are within the interval, 10 is given as a weight for the particle. If not in the interval, 2 is given. Misplaced particles will then for each iteration decline in number, in favor for well-placed particles.

4. **Re-sampling Step** - Next, 1000 new particles are re-sampled according to the weighted sums in each square.

5. Iterate to step 2).

After an example run of the program, about eight moves from an initial location in south-east, mainly in the direction north-west, the square with the

most particles is the same square as the square where the ship is located, which can be seen in figure 5 (right). After 10 moves, the weighted sum of all particles is located in the same square as the ship, which is shown in figure 6.

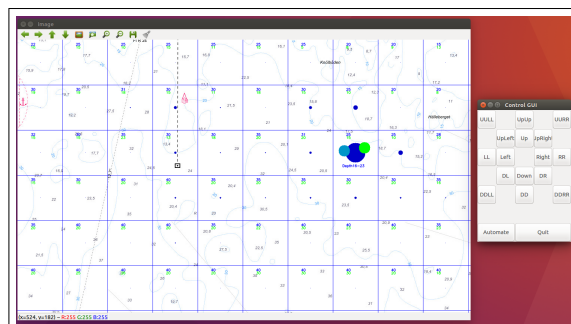


Figure 6: After 10 random steps, both the square with the highest number of particles and the weighted sum of all particles are in the square where the ship is located.

### 4.3 Further development of the implementation

There are several ways of refining the program. Some of the most important features are as follows:

- Adding support also for magnetic field maps.

- Using data from a real sea chart database, and creating a function for estimating bottom depths from that data set.
- Evolving the Correction Step by implementing the features in table 1 in section 3.
- Evolve the principles of how to be weighting the different subsets to increase the stability and the robustness. This can possibly be done on-line using machine learning, by feeding the "true position" from the GPS.

After finishing the implementation, a comparison needs to be performed on data originating from a real ship. The navigation performance can then be compared to the GPS location of the same ship.

## 5 Conclusion

It has already been shown that PF algorithms can be used for estimating positions [1, 2, 4, 6], at least for other domains than for naval ships. In this paper, an idea of how to use this knowledge in an algorithm more suitable for real world scenarios has been presented. A brief explanation of how to do the implementation has been discussed, and the present implementation has been presented along with proposed future upgrades.

If it is possible to navigate accurately enough without GNSS systems, only relying on high performance navigation sensors and normal sea chart and magnetic charts remains unclear. Further implementation and testing with real ship data is first needed.

## ACKNOWLEDGMENT

This work was partially supported by the Wallenberg Autonomous Systems and Software Program (WASP).

## References

- [1] Frank Daellaert, Dieter Fox, Wolfram Burgard, Sebastian Thrun "Monte Carlo Localization for Mobile Robots", Proc. IEEE International Conference on Robotics and Automation (ICRA-99), 1999
- [2] Fredrik Gustafsson, Fredrik Gunnarsson, Niclas Bergman, Urban Forssell, Jonas Jansson, Rickard Karlsson, and Per-Johan Nordlund "Particle Filters for Positioning, Navigation, and Tracking", IEEE Transactions on signal processing, vol. 50, No. 2, 2002
- [3] Rickard Karlsson and Fredrik Gustafsson "Bayesian Surface and Underwater Navigation", IEEE Transactions on Signal Processing, vol. 54, No. 11, 2006
- [4] Martin Frassl, Michael Angermann, Michael Lichstenstern, Patrick Robertson, Brian J. Julian, Marek Doniec "Magnetic Maps of Indoor Environments for Precise Localization of Legged and Non-legged Locomotion", IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2013
- [5] Etienne Le Grand and Sebastian Thrun "3-Axis Magnetic Field Mapping and Fusion for Indoor Localization", 2012 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), 2012
- [6] Michael Angermann, Patrick Robertson "Inertial-Based Joint Mapping and Positioning for Pedestrian Navigation", Proc. ION GNSS, 2009
- [7] Lawrence A. Klein "Sensor and Data Fusion - A Tool for Information Assessment and Decision Making", Publisher: SPIE Press, 2012
- [8] Todd E. Humphreys, Brent M. Ledvina, Mark L. Psiaki, Brady W. O'Hanlon, and Paul M. Kintner, Jr., Cornell "Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer", ION GNSS Conference, 2008

# Modelling regimes with Bayesian network mixtures

Marcus Bendtsen and Jose M. Peña

marcus.bendtsen@liu.se | jose.m.pena@liu.se

Linköping University, Department of Computer and Information Science, Sweden

## Abstract

Bayesian networks (BNs) are advantageous when representing single independence models, however they do not allow us to model changes among the relationships of the random variables over time. Due to such *regime changes*, it may be necessary to use different BNs at different times in order to have an appropriate model over the random variables. In this paper we propose two extensions to the traditional hidden Markov model, allowing us to represent both the different regimes using different BNs, and potential driving forces behind the regime changes, by modelling potential dependence between state transitions and some observable variables. We show how expectation maximisation can be used to learn the parameters of the proposed model, and run both synthetic and real-world experiments to show the model's potential.

## Keywords

Bayesian networks, hidden Markov models, regimes, algorithmic trading.

## 1 INTRODUCTION

Introduced by Judea Pearl [1], Bayesian networks (BNs) consist of two components: a qualitative representation of independencies amongst random variables through a directed acyclic graph (DAG), and a quantification of certain marginal and conditional probability distributions, so as to define a full joint probability distribution over the random variables. A feature of BNs, known as the local Markov property, implies that a variable is independent of all other non-descendant variables given its parent variables, where the relationships parent and descendant are defined with respect to the DAG of the BN. Let  $\mathbf{X}$  be a set of random variables in a BN, and let  $pa(X_i)$  represent the set of variables that consists of the parents of variable  $X_i \in \mathbf{X}$ , then the local Markov property allows us to factorise the joint probability distribution according to Equation 1.

$$p(\mathbf{X}) = \prod_{X_i \in \mathbf{X}} p(X_i | pa(X_i)) \quad (1)$$

From Equation 1 it is evident that the independencies represented by the DAG allow for a representation of the full joint distribution via smaller marginal and conditional probability distributions, thus making it easier to elicit the necessary parameters, and allowing for efficient computation of posterior probabilities. For a full treatment of BNs, please see [2, 3, 1].

While a BN has advantages when representing a single independence model, it does not allow us to model changes of the independencies amongst the modelled variables over time. One reason why we would take such changes into consideration is that we may wish to use different models for different sequential tasks, such as buying and selling shares in a stock market. This was the main reason for introducing gated Bayesian networks (GBNs) [4, 5], allowing the investor to create different BNs for the different phases of trading.<sup>1</sup> Another reason may be that the system that the modelled variables represent undergoes regime changes, i.e. there may be

---

<sup>1</sup>The GBN model also allows completely different random variables within each BN, something that we shall not explore further in this paper.

*states of the world* among which the independencies and distributions over the variables are different [6].

From the view of graphical models, the archetype approach for modelling regimes is to use a hidden Markov model (HMM), where the regimes are modelled using hidden random variables, and we observe the random variables that we are modelling under different states of these hidden variables. When using standard HMMs, it is common to assume that the observable variables are independent of each other given the hidden regime variable, and not to model any potential dependencies among the observed variables directly.

In this paper we are proposing an extension of the HMM, which we shall call GBN-HMM, where we bring in two of the fundamental ideas behind the GBN model. First, we shall allow for different BNs over the observable variables under the different states of the hidden variables, to have a regime-dependent model over the observable variables. Second, we shall model a potential dependence between one of the observable variables and the next hidden state. The second extension stems from one of the building blocks of GBNs, where the change of state is dependent on the posterior probability of a specific variable. The main difference between the GBN-HMM and the GBN is that GBNs identify one distinct BN as the model that represents the current regime, whereas the GBN-HMM defines a mixture of independence models, thus being a generative model of the data.

The rest of the paper is disposed as follows. In Section 2 we shall consider other existing extensions of the HMM, found in the literature, that are related to the extension that we shall propose. In Section 3 we will introduce and define the model that we are proposing, describing some of its underlying properties. Since there are hidden variables in the proposed model, parameter estimation is not immediately straightforward, and we shall therefore explore how we can use expectation maximisation (EM) in Section 4 to estimate the parameters of our model. In Section 5 we wish to demonstrate the appropriateness of the GBN-HMM using synthetic data, and compare it with the HMM as well as two other HMM variants. We then turn our attention to using the GBN-HMM in a real-world situation, namely trading shares in a stock market, in Section 6. Finally, we shall end this paper with our conclusions and a summary in Section 7.

## 2 RELATED WORK

HMMs have been applied and extended extensively throughout the literature, and we shall here not attempt an overview of all that has been explored. The interested reader may instead wish to consider the summary provided by Murphy [7]. Instead, we shall pay brief attention to a few existing variations that have a connection with the ideas that we are putting forward in this paper.

In [8] a HMM is described where some *control signal* is given as input to the hidden state and the observable variables, and offer an EM algorithm to update the parameters of the observational and transition distributions conditional on a sequence of input. As a variation on this theme, [9] proposes that transitions between hidden states in a HMM may not only depend on the immediately previous state, but also on the immediately preceding observation. This potential dependence between the observed variables at time  $t$  and the hidden state at time  $t + 1$  is also present in the GBN-HMM that we are proposing. We shall use the model proposed in [9] as a comparison model in our experiments.

The auto-regressive HMM (AR-HMM), also known as the regime switching Markov model [10], incorporates potential dependence directly between an observable variable at time  $t$  and its counterpart at  $t + 1$ . While the AR-HMM may be extended to higher orders, i.e. allowing for even longer dependence than only between  $t$  and  $t + 1$ , the dependence is between counterparts in each time slice. However, dynamic Bayesian multinets (DBMs) proposed in [11] allow not only for dependence across time slices among observational counterparts, but arbitrarily among the observed variables. Furthermore, DBMs allow these potential dependencies to change depending on the hidden states, thus allowing for a more complex dependence structure across time. The model that we are proposing does not include potential direct dependence among observable variables across time, but rather within each time slice.

In the next section we shall formally introduce the GBN-HMM that we are proposing, and then subsequently discuss parameter estimation and experiments comparing the GBN-HMM with other HMM variants.

### 3 MODEL DEFINITION

The GBN-HMM that we are proposing consists of a set of discrete random variables  $H_{1:T} = \{H_1, H_2, \dots, H_T\}$  that represent the hidden state at each time  $t \in [1, T]$ . We call these the hidden state variables, and they each have the same number of possible states  $N$ . We use  $h_t$  to denote a specific instantiation of the variable  $H_t$ , and use  $h_{j:k}$  to denote a sequence of states from time  $j$  to  $k$ . For each  $t$ , we will also model a set of discrete random variables  $\mathbf{O}_t = \{O_t^1, O_t^2, \dots, O_t^M\}$  for which we can observe their values. We will refer to these variables as the observable variables. We let  $\mathbf{o}_t = \{o_t^1, o_t^2, \dots, o_t^M\}$  be a particular instantiation of the observable variables at time  $t$ , and use  $\mathbf{o}_{j:k}$  and  $\mathbf{o}_{j:k}$  when considering all observable random variables and their respective values from time  $j$  to  $k$ .

Since we wish to model the observable variables depending on the current state, we will have one BN for each state of the hidden state variable  $H_t$ , that is, there are  $N$  BNs over the variables  $\mathbf{O}_t$ , and the value of  $H_t$  selects one of these. One of the variables in  $\mathbf{O}_t$  is of particular interest, as we will model a potential dependence between this variable and the state of  $H_{t+1}$ . We will refer to this variable as the  $Z$  variable when we need to distinguish it from the other observable variables. Notation wise we let  $Z_t$  represent the  $Z$  variable at time  $t$ , and  $z_t$  an instantiation of the  $Z$  variable at time  $t$ .

Note that, although not made explicit, we have made use of certain independence assumptions among the variables  $H_{1:T}$  and  $\mathbf{O}_{1:T}$ . First, we assume that  $\mathbf{O}_t$  are conditionally independent of all previous random variables  $\mathbf{O}_{1:t-1}$  and  $H_{1:t-1}$ , given the current hidden state variable  $H_t$  (thus knowing the current state renders the past irrelevant). Second, the current hidden state variable  $H_t$  is conditionally independent of  $\mathbf{O}_{1:t-1} \setminus Z_{t-1}$  and  $H_{1:t-2}$  given  $H_{t-1}$  and  $Z_{t-1}$  (thus knowing the value of the previous state and  $Z$  variable renders the rest of the past irrelevant). We can represent these assumptions using a graph, an example of which is depicted in Figure 1. In the figure we can see that it is  $O_t^3$  that is the  $Z$  variable, as we are modelling a potential dependence between it and the next hidden state.

The final assumption that we will make is that of stationarity of the model. That is, the distribu-

tions and independencies that govern the model are independent of  $t$ . This implies that the probability of moving from one hidden state to another is the same regardless of  $t$ , and that the BNs selected by  $H_i$  are the same as for  $H_j$  for all  $i, j \in [1, T]$ . Furthermore, the  $Z$  variable is always the same observable variable, regardless of  $t$  or the state of  $H_t$ .

#### 3.1 Factorisation

Using the independence assumptions implied by the model, and the chain rule of probability, we can factorise the joint distribution over  $H_{1:T}$  and  $\mathbf{O}_{1:T}$  into marginal and conditional distributions that together require fewer parameters than the full joint. To illustrate this factorisation in a succinct manner, we shall factorise the GBN-HMM given in Figure 1. We assume that the hidden state variables have two states, i.e.  $N = 2$ , however expanding this example to any number of observable variables, hidden states and time steps is straightforward. We begin the example by observing that we can isolate the variables  $O_3^1, O_3^2$  and  $O_3^3$  by conditioning on  $H_3$  alone, which follows from Equation 2.

$$\begin{aligned} & p(O_1^1, O_1^2, O_1^3, \dots, O_3^1, O_3^2, O_3^3, H_1, H_2, H_3) = \\ & p(O_3^1, O_3^2, O_3^3 | \cancel{O_1^1}, \cancel{O_1^2}, \cancel{O_1^3}, \dots, \cancel{H_1}, \cancel{H_2}, H_3) \times \\ & p(O_1^1, O_1^2, O_1^3, \dots, H_1, H_2, H_3) = \tag{2} \\ & p(O_3^1, O_3^2, O_3^3 | H_3) \times \\ & p(O_1^1, O_1^2, O_1^3, \dots, H_1, H_2, H_3) \end{aligned}$$

Since the hidden variables in a GBN-HMM select among several BNs over the observable variables, the two states of  $H_3$  select between two joint distribution specifications over  $O_3^1, O_3^2$  and  $O_3^3$ . If we let  $pa_j(O_3^j)$  represent the parents of the variable  $O_3^j$  with respect to the DAG of the BN selected by  $H_3 = j$ , then using the local Markov property of BNs we can continue the factorisation according to Equation 3.




Figure 1: Graph representation of the GBN-HMM with three time steps.

$$\begin{aligned}
 & p(O_3^1, O_3^2, O_3^3 | H_3) \times \\
 & p(O_1^1, O_1^2, O_1^3, \dots, H_1, H_2, H_3) = \\
 & p(O_3^1, O_3^2, O_3^3)^{\delta(H_3=1)} p(O_3^1, O_3^2, O_3^3)^{\delta(H_3=2)} \times \\
 & p(O_1^1, O_1^2, O_1^3, \dots, H_1, H_2, H_3) = \quad (3) \\
 & \prod_{j=1}^2 \prod_{i=1}^3 p(O_3^i | pa_j(O_3^i))^{\delta(H_3=j)} \times \\
 & p(O_1^1, O_1^2, O_1^3, \dots, H_1, H_2, H_3)
 \end{aligned}$$

In Equation 3 we let  $\delta(H_3 = j)$  represent the Kronecker delta, i.e. when  $H_3$  takes the value  $j$  it equates to unity, otherwise zero.

The next step of the factorisation is to break out  $H_3$  from the remaining variables, which follows from Equation 4. It should then be clear that we can continue the same operations for the remainder of the variables, ending the factorisation with a marginal distribution over  $H_1$ .

$$\begin{aligned}
 & \prod_{j=1}^2 \prod_{i=1}^3 p(O_3^i | pa_j(O_3^i))^{\delta(H_3=j)} \times \\
 & p(H_3 | \mathcal{O}_1^1, \mathcal{O}_1^2, \mathcal{O}_1^3, \mathcal{O}_2^1, \mathcal{O}_2^2, \mathcal{O}_2^3, \mathcal{H}_1, H_2) \times \\
 & p(O_1^1, O_1^2, O_1^3, O_2^1, O_2^2, O_2^3, H_1, H_2) = \quad (4) \\
 & \prod_{j=1}^2 \prod_{i=1}^3 p(O_3^i | pa_j(O_3^i))^{\delta(H_3=j)} p(H_3 | H_2, O_2^3) \times \\
 & p(O_1^1, O_1^2, O_1^3, O_2^1, O_2^2, O_2^3, H_1, H_2)
 \end{aligned}$$

The GBN-HMM factorisation for  $T$  time steps, with  $N$  hidden states and  $M$  observable variables, is given in Equation 5.

$$\begin{aligned}
 & p(H_1) \prod_{t=2}^T p(H_t | H_{t-1}, Z_{t-1}) \times \\
 & \prod_{t=1}^T \prod_{j=1}^N \prod_{i=1}^M p(O_t^i | pa_j(O_t^i))^{\delta(H_t=j)} \quad (5)
 \end{aligned}$$

### 3.2 Likelihood

Considering a specific sequence of observations  $\mathbf{o}_{1:T}$  and hidden states  $h_{1:T}$ , we can use the factorisation to compute the likelihood of this data under a set of parameters  $\Theta$ . We let  $\pi_i$  represent the probability  $p(H_1 = i | \Theta)$ ,  $a_{ijk}$  the probability  $p(H_t = j | H_{t-1} = i, Z_{t-1} = k, \Theta)$ , and  $b_j^i(\mathbf{o}_t)$  represent the probability  $p(O_t^i = o_t^i | pa_j(O_t^i) = \mathbf{o}_t^{pa_j(O_t^i)}, \Theta)^{\delta(H_t=j)}$ , where we let  $\mathbf{o}_t^{pa_j(O_t^i)}$  represent the values that the parent set takes in  $\mathbf{o}_t$ . Then the likelihood  $p(\mathbf{o}_{1:T}, h_{1:T} | \Theta)$  can be expressed by Equation 6.

$$\begin{aligned}
 & p(\mathbf{o}_{1:T}, h_{1:T} | \Theta) = \\
 & \pi_{h_1} \prod_{t=2}^T a_{h_{t-1}, h_t, z_{t-1}} \prod_{t=1}^T \prod_{i=1}^M b_{h_t}^i(\mathbf{o}_t) \quad (6)
 \end{aligned}$$

If we could observe both  $\mathbf{o}_{1:T}$  and  $h_{1:T}$  then estimating the parameters  $\Theta$  that maximised the likelihood would be straightforward. However, since  $H_{1:T}$  are hidden variables we cannot observe their values, and must therefore apply a more involved technique for estimating  $\Theta$ .

## 4 PARAMETER ESTIMATION

The canonical way of solving the parameter estimation problem in regular HMMs (and in their ex-

tensions) is to employ EM. We shall also adopt this approach, and in this section describe the computations necessary for iteratively updating the parameters  $\Theta$  for the GBN-HMM that we are currently proposing.

As before, let  $\mathbf{o}_{1:T}$  represent a sequence of observations over the variables  $\mathbf{O}_{1:T}$  and let  $h_{1:T} = \{h_1, h_2, \dots, h_T\}$  represent a sequence of states. Let  $\mathcal{H}$  represent the set of all state sequences  $h_{1:T}$ . The current parameters for our model are denoted  $\Theta'$ , and we seek parameters  $\Theta$  such that  $p(\mathbf{o}_{1:T}|\Theta) \geq p(\mathbf{o}_{1:T}|\Theta')$ . It can be shown [12] that this task can be converted into a maximisation problem of  $Q(\Theta, \Theta') = \sum_{h_{1:T} \in \mathcal{H}} p(\mathbf{o}_{1:T}, h_{1:T}|\Theta') \log p(\mathbf{o}_{1:T}, h_{1:T}|\Theta)$ .

Substituting  $p(\mathbf{o}_{1:T}, h_{1:T}|\Theta)$  in the  $Q$  function with the likelihood expression in Equation 6, gives us the expanded  $Q$  function in Equation 7. From this expansion we can conclude that the individual terms do not interact, thus they can be maximised separately.

$$\begin{aligned}
 Q(\Theta, \Theta') = & \\
 & \sum_{h_{1:T} \in \mathcal{H}} p(\mathbf{o}_{1:T}, h_{1:T}|\Theta') \log p(\mathbf{o}_{1:T}, h_{1:T}|\Theta) = \\
 & \sum_{h_{1:T} \in \mathcal{H}} p(\mathbf{o}_{1:T}, h_{1:T}|\Theta') \log \pi_{h_1} + \\
 & \sum_{h_{1:T} \in \mathcal{H}} p(\mathbf{o}_{1:T}, h_{1:T}|\Theta') \sum_{t=2}^T \log a_{h_{t-1}, h_t, z_{t-1}} + \\
 & \sum_{h_{1:T} \in \mathcal{H}} p(\mathbf{o}_{1:T}, h_{1:T}|\Theta') \sum_{t=1}^T \sum_{i=1}^M \log b_{h_t}^i(\mathbf{o}_t)
 \end{aligned} \tag{7}$$

The derivation of which values for the individual terms that maximise the  $Q$  function is relatively lengthy. We therefore defer all details to the supplementary material<sup>2</sup>, and here only account for the results of the derivation and show how to compute the necessary quantities.

#### 4.1 Estimating new parameters

Computing new parameters  $\pi_i$  for the initial hidden state distribution that maximise the  $Q$  function is done according to Equation 8. Here we are taking the conditional probability of each possible state  $N$

<sup>2</sup>Please find the supplementary material here: [https://www.ida.liu.se/~marbe92/pdf/gbn-hmm\\_supp.pdf](https://www.ida.liu.se/~marbe92/pdf/gbn-hmm_supp.pdf)

given the observed data and the current parameters  $\Theta'$ .

$$\pi_i = \frac{p(\mathbf{o}_{1:T}, h_1 = i|\Theta')}{p(\mathbf{o}_{1:T}|\Theta')} \tag{8}$$

The new parameters  $a_{ijk}$  can be computed using Equation 9, where we use  $\delta(z_{t-1} = k)$  to represent the Kronecker delta which is unity when  $z_{t-1}$  takes on value  $k$ , and zero otherwise. Essentially, we are taking into consideration the expected number of times that we have observed a transition from state  $i$  to  $j$  when  $z$  took value  $k$ , divided by the expected number of times we have seen transitions away from  $i$  when  $z$  took value  $k$ .

$$\begin{aligned}
 a_{ijk} = & \\
 & \frac{\sum_{t=2}^T p(\mathbf{o}_{1:T}, h_{t-1} = i, h_t = j|\Theta') \delta(z_{t-1} = k)}{\sum_{t=2}^T p(\mathbf{o}_{1:T}, h_{t-1} = i|\Theta') \delta(z_{t-1} = k)}
 \end{aligned} \tag{9}$$

The final set of parameters that we shall compute to maximise  $Q$  are the parameters of the distributions over the observed variables. We let  $b_{jkl}^i$  denote the parameter of the distribution for observable variable  $i$  when it takes on value  $l$ , given the hidden state  $j$  and its  $k$ :th parent configuration. An observation  $\mathbf{o}_t$  will identify one such parameter for each observable variable under a specific hidden state. We let  $\delta(\mathbf{o}_t, b_{jkl}^i)$  represent the Kronecker delta such that it is unity when the parameter identified by  $\mathbf{o}_t$  given  $h_t = j$  is  $b_{jkl}^i$ , and zero otherwise, and likewise let  $\delta(\mathbf{o}_t, b_{jk}^i)$  be unity when the  $k$ :th parent set is identified given hidden state  $j$  (regardless of the value of  $l$ ). We can then compute each  $b_{jkl}^i$  such that  $Q$  is maximised using Equation 10. This can again be seen as dividing the number of times that we expect to encounter a certain event  $(j, k, l)$  with the number of times we expect to encounter a superset of these events  $(j, k)$ .

$$b_{jkl}^i = \frac{\sum_{t=1}^T p(\mathbf{o}_{1:T}, h_t = j|\Theta') \delta(\mathbf{o}_t, b_{jkl}^i)}{\sum_{t=1}^T p(\mathbf{o}_{1:T}, h_t = j|\Theta') \delta(\mathbf{o}_t, b_{jk}^i)} \tag{10}$$

#### 4.2 Computing necessary quantities

While Equation 8, 9 and 10 describe which quantities are needed to compute the values necessary to maximise  $Q$ , the calculation of these quantities are not immediately available. In this section we turn



our attention to the computation of these necessary quantities. As before, we defer some of the details to the supplementary material, and here offer the results from the derivation.

The two quantities that we require, which we shall call  $\gamma$  and  $\xi$ , are presented and expanded in Equation 11 and 12. Apart from the quantities  $\alpha$  and  $\beta$ , the expansions consists of known quantities (readily available from the model under parameters  $\Theta'$ ).

$$\begin{aligned} \gamma_j(t) &= p(\mathbf{o}_{1:T}, h_t = j | \Theta') = \\ & p(\mathbf{o}_{t+1:T} | \mathbf{o}_t, h_t = j, \Theta') p(\mathbf{o}_{1:t}, h_t = j | \Theta') = \\ & \beta_j(t) \alpha_j(t) \end{aligned} \quad (11)$$

$$\begin{aligned} \xi_{ij}(t) &= p(\mathbf{o}_{1:T}, h_{t-1} = i, h_t = j | \Theta') = \\ & p(\mathbf{o}_{t+1:T} | \mathbf{o}_t, h_t = j, \Theta') p(\mathbf{o}_t | h_t = j, \Theta') \times \\ & p(h_t = j | \mathbf{o}_{t-1}, h_{t-1} = i, \Theta') \times \\ & p(\mathbf{o}_{1:t-1}, h_{t-1} = i | \Theta') = \\ & \beta_j(t) \prod_{k=1}^M b_j^k(\mathbf{o}_t) a_{ijz_{t-1}} \alpha_i(t-1) \end{aligned} \quad (12)$$

What is left to do is to define recursively  $\alpha$  and  $\beta$ , and then all required quantities are either already available or computable. We finish this section by defining these two quantities in Equation 13 and 14.

$$\begin{aligned} \alpha_j(t) &= p(\mathbf{o}_{1:t}, h_t = j | \Theta') = \\ & \prod_{k=1}^M b_j^k(\mathbf{o}_t) \sum_{i=1}^N a_{ijz_{t-1}} \alpha_i(t-1) \end{aligned} \quad (13)$$

$$\begin{aligned} \beta_j(t) &= p(\mathbf{o}_{t+1:T} | \mathbf{o}_t, h_t = j | \Theta') = \\ & \sum_{i=1}^N \beta_i(t+1) \prod_{k=1}^M b_i^k(\mathbf{o}_{t+1}) a_{jiz_t} \end{aligned} \quad (14)$$

Note that the equations given here are slightly different from those used when estimating the parameters of a traditional HMM. In Equation 9 we are only considering cases under different values of the  $Z$  variable, and in Equation 10 we are considering different parent configurations rather than just the hidden state. Also, the definition of  $\beta$  in Equation 14 includes conditioning on  $\mathbf{o}_t$ , since the  $Z$  variable at time  $t$  may influence the hidden state at  $t+1$ .

The only part that is left to take into consideration is how we find the parent sets of each observable variable within each hidden state, i.e. how do we learn the structure of the BNs. We shall take this into consideration in the next section, and then move on to synthetic and real-world experiments.

### 4.3 Structure learning

Taking the approach of [13], we wish to identify the model over the observable variables that, together with the parameters, maximises the last term of Equation 7. While advances in exact learning of graphical model structures have been made [14, 15], we shall here rely on a heuristic approach. Therefore, we use a greedy thick thinning algorithm [16] to identify the structure over the observed variables, such that the term over the observable variables is maximised in Equation 7. Thus within each iteration of the EM algorithm, we also heuristically identify the best structure over the observed variables within each regime.

## 5 EXPERIMENTS USING SYNTHETIC DATA

We shall in this section account for our experiments using synthetic data to compare the GBN-HMM with three other models. The comparison models are: the standard HMM with observation variables that are independent of each other given the hidden state, the SDO-HMM proposed in [9], where observations are again independent given the hidden state, but where we have (using our term) a  $Z$  variable, and finally a version of our GBN-HMM but without the  $Z$  variable, which we shall call MULTI-HMM (due to their relationship to Bayesian multi-nets).

### 5.1 Methodology and data generation

A single sample was generated as follows (with input to the procedure the predictive power of the  $Z$  variable):

Four BNs were created by randomly generating four DAG structures<sup>3</sup> over four variables, and then

<sup>3</sup>We used the R package bnlearn which uses the method

uniformly at random generating parameters for the resulting conditional distributions.<sup>4</sup> The number of states for each variable was determined uniformly between two and five, except for the  $Z$  variable which was given four states.

The first data point in the sample was generated from the first BN. The value of the  $Z$  variable then determined which BN to take the second data point from, with a certain level of predictiveness (the supplied predictive power). For instance, if the  $Z$  variable took value two, and the predictive power was 0.6, then there was a 60% chance that the next data point would come from the second BN, and a 40% that the next data point would come from the same BN as the previous data point. We repeated this until there were 1000 data points in the sample.

Following this procedure we generated 50 samples for each of the predictive powers 0.6, 0.7, 0.8 and 0.9.

For the synthetic experiments we were interested in how well the models fit held out test data. Therefore, for each sample, we employed a 5-fold cross-validation procedure using two thirds of the data to determine the number of hidden states, estimate the parameters of the models, and to learn the BN structures for GBN-HMM and MULTI-HMM. For SDO-HMM and GBN-HMM the models were told which  $Z$  variable to use. The remaining third was treated as held out test data, the likelihood of which will be reported.

## 5.2 Results and discussion

In Table 1 the results from the synthetic experiments are reported. Each row represents a certain predictive power. The values in the table are the means of the log-likelihoods of the held out test data, over the 50 samples, given each model.

Already when the  $Z$  variable has a predictive power of 0.6, the GBN-HMM had a considerably better fit to the data than both HMM and MULTI-HMM (note that this is log-scale). However, the SDO-HMM was also able to utilise this predictive power to get a similar fit as the GBN-HMM. As the predictive power of the  $Z$  variable increased to 0.7, the difference between the GBN-HMMs' fit of the data and the other models increased, suggesting that taking this predictiveness into account,

proposed in [17] to generate DAGs uniformly at random.

<sup>4</sup>Using the method described in [18].

Table 1: Means of log-likelihoods of held out data, using different predictive powers of the  $Z$  variable.

HMM	SDO-HMM	MULTI-HMM	GBN-HMM
Predictive power = 0.6			
-1546.438	-1537.453	-1558.732	<b>-1536.529</b>
Predictive power = 0.7			
-1538.541	-1526.940	-1550.820	<b>-1518.590</b>
Predictive power = 0.8			
-1535.269	-1509.830	-1546.112	<b>-1506.726</b>
Predictive power = 0.9			
-1513.058	-1476.529	-1526.843	<b>-1475.436</b>

and allowing for multiple BNs, can improve the appropriateness. When we look at the outcomes when the predictive power was increased to 0.8 and 0.9, the two models that do not utilise a  $Z$  variable (HMM and MULTI-HMM) drift further from the GBN-HMM, while the HMM-SDO reversed and came closer again. Although the GBN-HMM outperforms the other models throughout all experiments, it is interesting to see that the SDO-HMM can outperform HMM and MULTI-HMM by utilising the  $Z$  variables predictive power.

While the experiments that we have reported in this section work well as a confirmation of the proposed model's appropriateness, we shall now turn our attention to experiments where we wish to employ the model for a specific task. In Section 6 we shall explore the performance of the four models when they are used for systematic stock market trading.

## 6 TRADING THE STOCK MARKET

In this section we shall employ the models under comparison for trading stock shares, with the goal of balancing the risk and reward of such trading.

We shall first offer a brief introduction to some of the ideas and concepts surrounding systematic stock trading, and then employ the GBN-HMM in such trading, using the same models as in Section 5 as comparison (HMM, SDO-HMM and MULTI-HMM).

## 6.1 Systematic stock trading concepts

The general idea of systematic stock trading is to use some collected data to create rules that identify opportune times to own certain stock shares, and times when it is less beneficial to own them. Usually this is referred to as generating buy and sell *signals*. For the purpose of the experiments that we shall undertake, this type of all-or-nothing approach will suffice. However, in a more mature systematic trading system one may very well wish to trade several different shares at different quantities, utilising diversification in one's favour.

If signals from a system are executed, then this will generate a certain risk and reward in terms of the initial investment. For instance, if we execute a buy signal then any change in the price of the bought shares will also give us a proportional (positive or negative) return on our investment. Naturally, one seeks a positive return on one's investment, however simply using the raw return as the only goal of investment is not necessarily the best approach. Instead it is common to take into consideration the variation of the returns an investment yields. Therefore we shall seek a high Sharpe ratio (named after Nobel Laureate William F. Sharpe), where we take the mean of our returns, less the *risk free rate*, divided by the standard deviation of our returns. Here, the risk free rate is the return that we can expect from interest, or some other "safe" asset such as government bonds. As our comparison will be among models, rather than investment strategies, we shall remove the risk free rate from the Sharpe ratio and simply consider the mean return divided by the standard deviation of the returns.

The type of data that is used in stock trading systems vary greatly, however a common approach is to take the historical price and apply so called *technical analysis indicators* to gauge whether prices are trending, shares are overpriced, etc. For our purposes we shall consider two such indicators: the relative difference between two *moving averages*, often referred to as MACD [19], and the relative strength index (RSI) [20], which compares recent price increases with recent price decreases. The MACD is computed by first calculating two moving averages with different length windows, one using the most recent five days of prices, and one

using the most recent ten days of prices. The difference between the two then becomes a gauge for the trend in the market, if it is positive it means that the five day moving average is above the ten day moving average, indicating an upswing in price (and vice versa). The RSI computes the average of all price increases the past 14 days and divides by the average of all price decreases the past 14 days, a high RSI indicates that prices have been increasing strongly and may therefore be overpriced (and vice versa). For sake of brevity we shall leave out the exact calculations of these indicators, and refer the interest reader to the referred literature.

## 6.2 Methodology

The MACD and RSI gave us two observable variables in our models, and we additionally considered the first order backward difference of these variables (i.e. we approximated the indicators' first order derivatives), giving us a total of four observed variables. The MACD was discretised into two states, positive and negative, and used as the  $Z$  variable. The rest of the indicators were discretised into four states, using their respective mean and one standard deviation below and above their mean as cut points.

We used daily data between 2003-01-01 and 2012-12-28 for seven actively traded stocks: Apple (AAPL), Amazon (AMZN), IBM (IBM), Microsoft (MSFT), Red Hat (RHT), Nvidia (NVDA) and General Electric (GE). To create multiple simulations from this data we divided the data into ten blocks (one year per block), and created seven simulations by first using block one, two and three as training data and block four as testing data, and then block two, three and four as training data and block five as testing data, and so on.

As in the experiments in Section 5, we employed a 5-fold cross-validation procedure using the training data to decide upon the number of hidden states, the parameters of the models, and the BN structures within the GBN-HMMs and MULTI-HMMs. For SDO-HMM and GBN-HMM the models were told to use the MACD variable as the  $Z$  variable.

While ones first intuition may be to attempt to label the hidden states of our models as "buy", "sell", etc. and thereby generate signals that can be executed, this is not the approach we will take

in this application. We do not know how many hidden states will be identified in each simulation, thus it would require some automatic labelling based on the number of states and historical advantage of different types of labelling. Instead, we shall build our rules as follows:

- On day  $t$ , when we know the values of  $\mathbf{O}_{1:t}$ , we shall make a prediction of the MACD variable at time  $t + 1$ .
- If  $p(\text{MACD}_{t+1} = \text{positive} \mid \mathbf{O}_{1:t}) > \theta$ , then generate a buy signal.
- If  $p(\text{MACD}_{t+1} = \text{negative} \mid \mathbf{O}_{1:t}) > \theta$ , then generate a sell signal.

We are thus generating buy and sell signals when enough of the probability mass indicates that the MACD is positive/negative. The particular  $\theta$  used was determined for each model by generating trade signals using the training data. For each simulation we generated signals for each block reserved for training (three blocks per simulation) and calculated the Sharpe ratio per block using different  $\theta$  (0.50, 0.55, ..., 0.90, 0.95). The  $\theta$  used on the test data was then the  $\theta$  with the highest average Sharpe ratio over the training blocks.

### 6.3 Results and discussion

Signals were generated for each held out test block, and the annual return and standard deviation was calculated for each block and model, giving rise to an annual Sharpe ratio for each model and traded stock. The annual Sharpe ratios are given in Table 2.

From the table we can see that the use of multiple BNs (i.e. MULTI-HMM and GBN-HMM) yields a higher annual Sharpe ratio for five of the seven stocks, losing out to SDO-HMM for RHT and HMM for IBM. In four out of the five cases where using multiple BNs was better, the GBN-HMM outperformed the MULTI-HMM. Thus in general, allowing for multiple BNs over the observable variables does increase the performance of the trading systems. Similarly, when considering the models that include a  $Z$  variable against those which did not, we see that the  $Z$  variable models won five against two. When comparing the use of both multiple BNs and a  $Z$  variable, i.e. the GBN-HMM, the outcome is four against three in favour of the

Table 2: Annual Sharpe ratio comparison.

HMM	SDO-HMM	MULTI-HMM	GBN-HMM
Apple (AAPL)			
0.844	0.708	<b>0.849</b>	0.718
Amazon (AMZN)			
0.466	0.580	0.449	<b>0.592</b>
IBM (IBM)			
<b>0.713</b>	0.521	0.699	0.616
Microsoft (MSFT)			
0.091	-0.189	-0.307	<b>0.219</b>
Red Hat (RHT)			
-0.198	<b>0.111</b>	-0.780	-0.085
Nvidia (NVDA)			
0.113	0.211	0.262	<b>0.308</b>
General Electric (GE)			
0.0621	0.362	-0.378	<b>0.419</b>

GBN-HMM. So even when all the other models are counted as one, the GBN-HMM wins. It should be noted that the models are all generative, thus they have been learnt with the goal of explaining the data generating process, and not to the specific task of stock trading. The case of SDO-HMM outperforming GBN-HMM on RHT is evidence of this difference between goals, as the GBN-HMM should always explain the data better, or the same, as the SDO-HMM, as the former is capable of mimicking the same structure as the latter.

It seems that the different models are advantageous under different circumstance, although the GBN-HMM seems to have an advantage in general. However, since GBN-HMMs embrace the other three models, we could take the structure learning further than only for the individual BNs, and learn which one of the four models considered is the most appropriate for the current task. We however leave such exploration to future work.

## 7 CONCLUSIONS & SUMMARY

Many real-world systems undergo changes over time, perhaps due to human intervention or natural

causes, and we do not expect probabilistic relationships among the random variables that we observe to stay static throughout these changes. We therefore find the use of multiple BNs for the different resulting regimes intriguing. In this paper we have proposed a model, which we call GBN-HMM, that incorporates these regime changes by using different BNs for the different regimes. Furthermore, the GBN-HMM allows us to model potential driving forces behind the regime changes by utilising some observational variables. We have shown the benefits of using the GBN-HMM in comparison with three related models, both by comparing fitness to data using synthetic data, and in a real-world systematic trading task.

## References

- [1] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers, 1988.
- [2] F. V. Jensen and T. D. Nielsen, *Bayesian networks and decision graphs*. Springer, 2007.
- [3] K. B. Korb and A. E. Nicholson, *Bayesian artificial intelligence*. Taylor and Francis Group, 2011.
- [4] M. Bendtsen and J. M. Peña, “Gated Bayesian networks for algorithmic trading,” *International Journal of Approximate Reasoning*, vol. 69, pp. 58–80, 2016.
- [5] M. Bendtsen, “Bayesian optimisation of gated Bayesian networks for algorithmic trading,” in *Proceedings of the Twelfth Annual Bayesian Modeling Applications Workshop*, pp. 2–11, 2015.
- [6] M. Bendtsen, “Regimes in baseball player’s career data,” *Data Mining and Knowledge Discovery*, 2017, accepted.
- [7] K. P. Murphy, *Machine learning: a probabilistic perspective*. The MIT press, 2012.
- [8] S. Bengio and Y. Bengio, “An EM algorithm for asynchronous input/output hidden Markov models,” in *Proceedings of the International Conference On Neural Information Processing*, pp. 328–334, 1996.
- [9] Y. Li, “Hidden Markov models with states depending on observations,” *Pattern Recognition Letters*, vol. 26, no. 7, pp. 977–984, 2005.
- [10] J. D. Hamilton, “A new approach to the economic analysis of nonstationary time series and the business cycle,” *Econometrica*, vol. 57, no. 2, pp. 357–384, 1989.
- [11] J. A. Bilmes, “Dynamic Bayesian multinets,” in *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pp. 38–45, 2000.
- [12] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2013.
- [13] N. Friedman, “The Bayesian structural EM algorithm,” in *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pp. 129–138, 1998.
- [14] D. Sonntag, J. M. Peña, A. Hyttinen, and M. Järvisalo, “Learning optimal chain graphs with answer set programming,” in *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pp. 822–831, 2015.
- [15] C. Yuan and B. Malone, “Learning optimal Bayesian networks: a shortest path perspective,” *Journal of Artificial Intelligence Research*, vol. 48, no. 1, pp. 23–65, 2013.
- [16] D. Heckerman, “A tutorial on learning with Bayesian networks,” Tech. Rep. MSR-TR-95-06, Microsoft Research, March 1995.
- [17] G. Melançon, I. Dutour, and M. Bousquet-Mélou, “Random generation of directed acyclic graphs,” *Electronic Notes in Discrete Mathematics*, vol. 10, pp. 202–207, 2001.
- [18] J. S. Ide and F. G. Cozman, “Random generation of Bayesian networks,” in *Brazilian Symposium on Artificial Intelligence*, pp. 366–376, 2002.
- [19] J. J. Murphy, *Technical analysis of the financial markets*. New York Institute of Finance, 1999.
- [20] W. J. Wilder, *New concepts in technical trading systems*. Trend Research, 1978.

# PastVision: Exploring “Seeing” into the Near Past with Thermal Touch Sensing and Object Detection For Robot Monitoring of Medicine Intake by Dementia Patients

Martin Cooney, Josef Bigun

School of Information Technology, Halmstad University, Halmstad, Halland, Sweden, martin.daniel.cooney@gmail.com

## Abstract

We present PastVision, a proof-of-concept approach that explores combining thermal touch sensing and object detection to infer recent actions by a person which have not been directly observed by a system. Inferring such past actions has received little attention yet in the literature, but would be highly useful in scenarios in which sensing can fail (e.g., due to occlusions) and the cost of not recognizing an action is high. In particular, we focus on one such application, involving a robot which should monitor if an elderly person with dementia has taken medicine. For this application, we explore how to combine detection of touches and objects, as well as how heat traces vary based on materials and a person's grip, and how robot motions and activity models can be leveraged. The observed results indicate promise for the proposed approach.

## Keywords

Thermal Sensing, Action Recognition, Home Robots, Monitoring, Medicine Adherence.

## 1 INTRODUCTION

This paper explores a concept for how to infer a person's recent past actions using thermal touch sensing and object detection, as shown in Figure 1. We focus on the case of a home robot which should monitor medicine intake, toward potentially supporting health and well-being in elderly persons with dementia.

An “action”, sometimes called a sub-activity or actionlet, refers to a basic human behavior; actions can be combined to form more complex “activities”. A “past” action here refers to an action which was not sensed by the system at the time the action was performed. A “thermal camera” outputs grids of pixels whose intensities depend on the temperatures of remote objects passively sensed through radiated long-wavelength infrared light; we refer to spots where heat has been transferred through touch as “heat traces” or “thermal touches”. “Object detection” involves both finding and recognizing objects. “Well-being” refers to a subjective feeling of being well, related to happiness and good quality of life.

Inferring past actions will be useful in scenarios when sensing can fail and the cost of not recognizing an action is high (e.g., when occlusions could interfere with monitoring for terrorist attacks or acute health problems). The scenario we focus on in the current paper, home robots intended to care for elderly persons with dementia, is characterized by both of these properties. A robot might not be able to sense a person if




Figure 1. PastVision: basic concept

it is in a different part of a home to conduct a task or to provide a person with some time alone, or if there is a temporary sensor failure. Yet monitoring could help to save lives by allowing a robot to intervene in emergencies; e.g., by sending out an alarm and providing health care professionals with information. Emergencies can arise from various causes such as dehydration, wandering at night, burns from fires or hot water, or non-adherence to medicine regimens. The latter problem has been reported as particularly serious and widespread. Taking too little, too much, or the wrong medication can have disastrous consequences, including death;

forgetting is a problem for dementia patients; and in some situations, such as when doses are frequent, compliance by only half of patients with a regiment has been reported [22]. Thus, we focus on monitoring of medicine intake.

Various approaches could be used by a machine to infer if a person has recently taken medicine. For example, cameras, RFID tags, and weighing scales could be used to determine, e.g., if a medicine package has been moved or weighs less, sounds of opening and swallowing could be detected, or a robot could directly ask a person what they have done. Some shortcomings of such approaches include that medicine might have been moved absentmindedly or replaced in a similar position after intake, frequent weighing and scrutiny could make a person feel little peace of mind, sounds can be easily muffled by noise (e.g., from a television, phone call, or alarm), constantly being asked could be irritating, and a person, especially with dementia, might not remember everything they have done.

Here we turn our attention to one possibility for inferring past actions, based on thermal touch sensing and object detection. Our proposal rests on an assumption that activities of daily living often involve touches and objects: we touch appliances when cooking, eat with cutlery and dishes, dress ourselves with clothes taken from cupboards, clean ourselves and our environments with brushes and cloths, and touch packages to take medicine. Furthermore we guessed that the result in a typical scenario of a home heated to around 20 degrees [36] would be some heat traces on objects which could be sensed after actions had occurred.

The challenge was that how or if thermal sensing can be combined with object detection was unclear. Furthermore, we did not know what can be sensed thermally in a typical scenario and how a robot could move to facilitate inference.

To address this challenge we propose an approach which we call PastVision, based on building a prototype, while also exploring some typical thermal properties of materials and grips, and strategies for a robot to move to reduce uncertainty via locomotion and manipulation.

To assist others interested in exploring this promising area, source code for an implementation of PastVision, as well as a video, will also be made available online (from martin-cooney.com).

## 2 RELATED WORK

The current work relates to recognition and inference of (past) actions, thermal touch sensing, and motion planning for healthcare robots dealing with medications.

### 2.1 Past action inference

Various innovative work has been conducted on autonomous inference, also of past events, actions, touches, and activities. Timing was used by a robot to infer who was doing an action, by correlating motion commands and visually sensed movements [13].

Adaptation of another robot's model of itself based on observed motions was used to infer a person's goal of pressing buttons [14]. And, a robot was able to infer the rules of games such as hide-and-go-seek and tag, by correlating its own motor commands and the actions of salient detected objects [10].

Inference of the past is also common in various fields. Looking up, evidence can be seen in the night sky of events such as star formations and deaths which have occurred millions of years ago. Looking down, fossils and rocks describe various events in the evolutionary history of species and the history of our world which we have not directly witnessed.

As well, many approaches have been used toward ascribing meaning to actions (e.g. [7]). For example, sounds were used to recognize some actions in homes [31]. And, similar to our approach, information from objects has been used to recognize actions like reading a book [27].

Also, touch-based actions have been recognized in various ways [30, 16]. For example, Lee demonstrated a way to separate spatiotemporally co-occurring touches using spatial Independent Component Analysis (ICA) and time series clustering [19]. Typical grips people use for opening packages have also been described (e.g., "spherical", "cylindrical" or "lateral" grips for bottles), which could be recognized [3].

Furthermore, some recent work has tackled "predicting" actions, which is typically framed as a problem of early recognition. For example, a human's next actions have been predicted by modelling object interactions using a conditional random field (CRF) on RGB-D data (for some activities such as taking medicine) and generating likely next moves [17]. Another system leveraged action sequences and objects [20]. And, a "memory" model was used to associate observed actions with previously learned ones [35].

What has received little attention is the combination of past inference and action recognition (possibly because actions cannot be sensed after completion by typical sensors such as cameras and microphones); furthermore, we are not aware of a previous work which proposed a way to ascribe meaning to touches by combining object detection and thermal touch sensing. (On the side, we also suggest how the basic approach of early recognition for action prediction can be inverted to see deeper into the past).

### 2.2 Thermal touch sensing

Thermal sensing has been used for many applications, some also related to healthcare and detecting touches.

In healthcare, use of infrared thermography has been described for detection of Raynaud's, fever, injury, breast cancer, diabetes neuropathy, dental and brain problems [25], apnea/hypopnea [21] and potentially skin cancer [5]. Moreover, thermal cameras have been used, along with an RGB camera and GPS, to find reclining people in emergencies [26], and to detect falls in a home [37].

Thermal sensing has also been studied in the area of human interfaces as a way to detect touches. Described as a concept by Iwai and Sato for a table-shaped interface [15], Benko et al. reported how touches could also be detected on a spherical display by normalizing, binarizing, and detecting and tracking connected components [4]. Larson and colleagues, similar to the current paper, sought to explore various facets of thermal touch sensing for their application [18]. The authors described a video-based approach for detecting heat traces which searches within a region in which a person's hands have moved recently; Bayesian estimation is conducted per pixel based on spatiotemporally smoothed temperatures, changes over time, and background subtraction, and belief is compared to a threshold for binary classification. Furthermore, the authors distinguished touches from hovering, recognized shapes of gestures and touch pressures, also for touches with multiple fingers, and suggested that interfaces could also take advantage of thermal reflections. This latter suggestion was investigated by Shirazi and colleagues [29], who also examined properties of some typical materials: glass, tile, MDF, and aluminum [1].

For the current study, it was considered that it would be useful to have an indication for how feasible it is to sense thermal touches to monitor medicine intake: e.g., how long touches remain visible on typical packaging materials, and how people typically touch medication packages.

### 2.3 Motion planning for robotic medication reminder systems

Various robotic systems have been proposed to remind people about medicine, bring medication, or check if a person has taken medication; and usability studies have confirmed that elderly see the usefulness in such systems [34]. One early work reported an aim for the nurse robot Pearl to remind in such a way that people comply without feeling annoyed or becoming overly reliant [23]. Also, some other robots capable of providing medication reminders have been designed to care for persons with mild dementia or Chronic Obstructive Pulmonary Disease (COPD) in smart environments; moreover some robots have been designed to fetch medication when users are indisposed [11]. Similarly, a small robot was proposed for keeping track of medicine and navigating to an elderly person to provide medicine [6]. For checking medicine intake, RFID tags were proposed for detecting if a person extracted and replaced a sheet of tablets from a tray on top of a small mobile robot [33]; and, some work outside of robotics has examined detecting swallowing by sound or muscle activation [2].

Additionally, much work has shown how robots can move to achieve various goals. For example, social force models have been used to describe how robots can locomote in a human-like, safe way (e.g., [28]). Another strategy was proposed for how a robot can try to move to detect people using a thermal and RGB camera [9].

Some other robots have manipulated objects to better segment them visually for grasping [12] and learn affordances [32]. The current work proposes how locomotion and manipulation can be used by a robot to facilitate past action inference for medication intake monitoring.

### 2.4 Contribution

The contribution of the current paper is exploring the concept and feasibility of combining thermal touch sensing and object detection to infer past actions, which we call PastVision, in the context of robotic monitoring of medicine intake; along the way, some typical thermal touch parameters and strategies for robot motion were also explored for this context.

## 3 METHODS

### 3.1 PastVision: making basic sense of thermal touches via object detection

This section describes step by step our approach for inferring past actions based on thermal touch sensing and object detection, illustrated in Figures 2, 3, and 4.

In general, the precondition for our algorithm is that a system has inferred that there might have been some missed data. The input to the algorithm is one thermal and one RGB image. The algorithm then outputs labels of touched objects. The postcondition is that the system has inferred what it has missed (a person's recent interactions with objects).

In particular, we focus here on a simplified basic scenario in which one robot is monitoring one person's medication intake; the scenario can be further simplified if the robot has been constrained to surveil a designated area in which only objects of interest are placed. During monitoring, some problem occurs with the robot's sensing, which is then fixed several seconds later. For example, this could be due to a person temporarily blocking the robot's view. The robot in front of some medicine packages then seeks to infer what the person has done. Here we consider some typical objects related to medicine intake: pill bottles, medicine boxes, flat sheets of medicine, water bottles, cups, and glasses (for oral medicine), as well as creams and syringes (for topical medicines and injections).

The process for inferring past actions targeting such objects is shown in Figure 2:

- (a-b) a thermal and RGB image are recorded
- (c-d) registration is conducted using a simple mapping found ahead of time
- (e) the thermal image is thresholded to detect warm regions
- (f) objects are detected within the RGB image, yielding bounding boxes and predicted classes



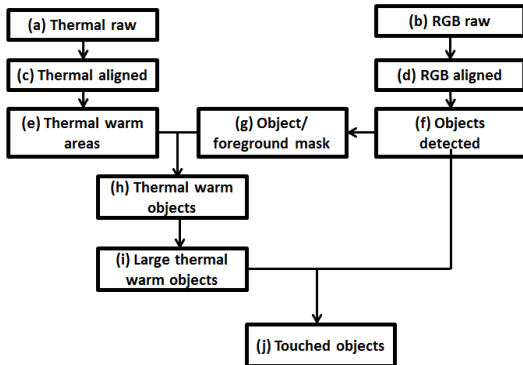


Figure 2. PastVision: process flow

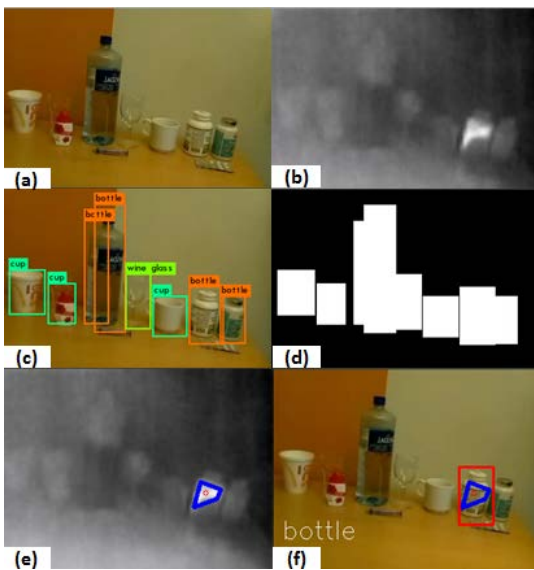


Figure 3. PastVision: (a) RGB image, (b) thermal image, (c) objects detected, with a few false positives and negatives, (d) mask image to reduce noise, (e) thermal image with touched region and contour centroid drawn, (f) touched object identified

- (g) The bounding boxes are used to prepare a mask image to extract foreground regions in the thermal image.
- (h) The intersection (the bitwise and) of the object mask and thresholded thermal image is computed to remove noise, which can arise due to heat from light sources, thermal reflections or unintended touches.
- (i) Connected components (contours) are detected. A threshold is used on contour size to further remove small noise.
- (j) For each contour, the centroid of the contour is calculated, and distances are found to centers of detected objects. The algorithm outputs the labels of detected objects with bounding box centers closest to the centroids of detected heat traces.



Figure 4. Example of a simplified approach for extracting heat traces and not humans via thresholds, morphology, and a basic shape model for touching: (a) thermal image, (b) RGB image with heat traces drawn by the algorithm

We note some considerations in using this approach:

- Action inference. There is no guarantee that inference is correct. For example, if a person touches a water bottle their intention could be to drink but could also be to dilute paint, water a plant, or cool their forehead. For this reason, our algorithm also approximates detected contours with the Ramer-Douglas-Peucker algorithm, in the expectation that shape information will provide useful information for inferring what a person has done. More information about context (e.g., current actions, timing models) will also be useful.
- Videos containing humans. Our approach is not limited to single images of objects; we also offer an example of the algorithm used in video, shown in Figure 4; we note that due to the simplified scenario hot objects such as humans can be ignored by the algorithm using some simple processing. More complex approaches can be used to enhance robustness of sensing for more challenging scenarios (possibly at the cost of incurring some waiting time). For example, motion estimation can be used to detect pixels with continuously changing heat, with a precalculated model of heat decay fit to every pixel to eliminate false pixels and confirm correct pixels. The approach used by Larson and colleagues of detecting heat traces near where a person's hands move could also be applied, although this approach was designed for a user interface with assumptions which might not hold for monitoring medicine intake (a robot might not have a background model, and a person's hands might not be visible). Also, cues like skin color could be used for detecting people.
- Feasibility for the application. For monitoring medicine intake, one problem could be if current object detection methods cannot recognize required categories. To gain insight, we checked the degree to which typical medication dosage forms are found within a common object recognition dataset, used for the ImageNet Large Scale Recognition Challenge (ILSVRC). We found pill bottle (ImageNet category number n03937543), water bottle (n04557648), cup (n07930864), goblet (n03443371), beer glass (n02823750), sunscreen (n04357314), lotion (n03690938), Band Aid (n02786058), face powder (n03314780), and syringe (n04376876). Two categories which appeared not to be represented

were paper boxes and flat sheets of lozenges. (As a check, an image of a paper box submitted to an ImageNet classifier was perceived to be of an “eraser” and lozenges were also not detected in the example image in Figure 3). However, it seemed many categories of interest for monitoring medicine intake can already be recognized, suggesting the promise of this approach for the application.

- Proof-of-concept. We note that the examples in Figures 3 and 4 show a highly simplified case chosen for initial proof-of-concept investigation: images were recorded in a lab environment at regular temperature with few objects and a simple background.
- Seeing further into the past via activity recognition models. We propose that, in the same way that early recognition of activities allows action prediction, recognizing the last action in an activity could allow inference into the more distant past. Other potential benefits of considering activities include increased certainty of inference if multiple related actions are observed, and possible ability to infer difficult-to-sense actions involving metals, for which heat traces disappear quickly. We provide two examples below of potential action sequences for medicine intake, while noting that many variations are possible: oral intake might include fetching objects, pouring water into a glass, opening a medicine bottle, taking out a pill, swallowing it, and drinking water; injections or topical applications might include fetching objects, baring a body part, washing and/or disinfecting, applying or injecting, and covering. Automatic modeling of such activities is possible, e.g., through “fluent learning” combining interval calculus and co-occurrence frequencies [8].

Although our approach and scenario are highly simplified, we verified that our system was able to detect unobserved object interactions which would be highly difficult for previous systems using only thermal touch detection or object detection.

### 3.2 Typical properties of heat traces

We proposed an approach for how to infer which objects have been recently touched by correlating extracted locations of heat traces and objects, but some questions remained with regard to what can be recognized. In regard to timing, if only a few milliseconds into the past can be recognized our approach might not be very useful. As well, we did not know to what degree it might be possible to recognize different kinds of touches, like typical ways of opening a medicine package or just picking up a package and putting it down.

To answer the questions, two simplified tests were conducted. First we touched some typical materials for medicine packages for a set time (2 seconds) lightly with the pad of one finger and checked how visible the heat trace was after 30 seconds. Materials tested were polyethylene terephthalate (PET, for a pill bottle), high-density polyethylene plastic (HD-PE, for a sunscreen

lotion bottle), paper (for a medicine box), glass (for a drinking glass), and ceramic (for a cup). Ambient room temperature was measured to be 24.7 degrees, and hand temperature to be 36.8 degrees.

For the second test we recorded some data of opening a pill bottle with spherical, cylindrical or lateral grips.

As a result of the first test, we were surprised by the extent that short touches to all materials tested except ceramic resulted in heat traces which could be seen for over thirty seconds, as shown in Figures 5 and 6. We think this indicates the feasibility for our proposed approach because we expect that people will often touch longer than two seconds to open a medicine package, especially elderly persons with dementia.

For the second test, patterns for spherical and lateral grips, shown in Figure 7, appeared to be similar. This is because, although the fingers on the top are positioned differently, in unscrewing the bottle, fingers are removed and replaced in different positions several times, leading to unclear smudging at the top for both. Heat traces for the cylindrical grip had a different appearance, due to some smudging from the lower half of a hand brushing against the bottle. The possibility was also suggested that these grips could be distinguished from some absent-minded touching if the latter only affects one region of the bottle.


### 3.3 Maximizing information gain via robot motion planning

In the last section, our investigation suggested that being able to see a package from various angles could yield useful information. Based on this, we propose that robots can enhance past action recognition by (1) locomoting and (2) manipulating objects.


(1) Locomotion: A robot could seek to find a position which allows better observation if its view becomes occluded and is not expected to improve immediately. Here we consider an example of a person moving in front of the robot to interact with some objects, who then becomes stationary. We propose that the robot should move to:

- minimize occlusions by the person on objects in its view
- be close enough to see objects clearly and appear socially positive to a person
- be far enough not to prevent object interactions or bother a person
- minimize work

To achieve such requirements one way to calculate how a robot should move is to use a social force model [28]. Social force models seek to model natural locomotion via forces pulling a robot toward a goal state and pushing it away from obstacles. Individual forces, coinciding or conflicting, are summed to calculate a net force acting on




**Figure 5.** Heat trace decay over time for some typical materials for medicine intake: for (a) PET, (b) HD-PE, (c) paper, (d) glass, and (e) ceramic; (left) RGB data, (center) thermal data soon after touching, (right) thermal data after 30s




**Figure 6.** Heat trace decay over time for PET. After two minutes the heat trace was difficult to see

the robot. All of the proposed requirements above except the first are common in robot social force models. We propose to model such a requirement via torque (a force applied to the robot position with the object position as an axis of rotation). Similar to the Collision Prediction (CP) specification, which computes a hypothetical time when pedestrians would become closest, the direction of force can be calculated by translating the robot position along the lever arm connecting it to the objects until the distance between robot and human is minimal. Intuitively, this means a robot could try to move circularly left if a human is in front and to its right, and circularly to the right if a human is standing in front and to the left. The proposed model is shown in Figure 8.

A more complex approach might involve predicting when and where a human will move next to determine how



**Figure 7.** Typical grips for opening packages: (a) spherical, lateral, and cylindrical; (b) an example of a heat trace for a spherical grip



**Figure 8.** A social force model can be used to model how a robot can move to monitor medicine intake when occluded. Symbols:  $o$ : object position,  $h_k$  and  $r_k$  human and robot positions at time  $k$ ,  $h'_k$  hallucinated human position to calculate torque on the robot (and force  $f_{oh}$ );  $f_{h+}$  and  $f_{h-}$  attractive and repulsive forces toward a human,  $f_{o+}$  and  $f_{o-}$  attractive and repulsive forces toward objects; the dashed line represents the robot's orientation

a robot should move. For example, people's bodies might start to turn as they shift their attention to some new object. A robot could also use knowledge of typical activities; e.g., a person picking up some bread might next move to a toaster.

(2) Manipulation: We propose that a robot can pick up objects to better determine if they have been touched recently and how. This addresses the problem that touches might not be fully visible from the robot's initial perspective.

To gain some insight, we acquired some data by conducting some actions (either pretending to take medicine, or just picking up and putting down packages) and then commanding our robot to pick up packages and observe them from different angles. (A more complex approach could involve checking that the robot's gripper is not covering a heat trace, and use image stitching for

cylindrical or spherical objects and segmentation.) Figure 9 demonstrates how inference can be facilitated by gaining extra information from viewing objects close-up and from various perspectives.


### 3.4 Implementation: Software and Hardware

To explore the PastVision concept, some tools were used, which are described below for reproducibility. For software, we used Robot Operating System (ROS) to send commands and data, OpenCV 3.2.0 for image processing, and Darknet code for YOLO (You only look once) object detection. For the latter, we used YOLO9000 version 2 because it was easy to use (predicting bounding boxes and class probabilities at once), fast, and provided excellent performance [24]. YOLO uses a single convolutional neural network with many layers (deep learning) trained on a mixture of detection and recognition data, and non-maximum suppression with a confidence threshold (0.25).

For hardware, we used a thermal camera and small computer attached to a robot, and a remote desktop for control and processing. An inexpensive 80 x 60 forward looking infrared (FLIR) camera was used, which is capable of detecting long wave infrared wavelengths from 8 to 14 microns with a 51-degree horizontal field of view (63.5 degree diagonal) and thermal sensitivity less than 0.05K (we consider this to be sufficient for the context of detecting touches to medicine packages, as touches to typical materials can result in much larger changes in surface temperature [1]). Thermal images were read over SPI by a Raspberry Pi 3, with RGB images obtained over CSI. Unoptimized code, showing thermal and RGB streams both independently and overlaid while recording data, ran at approximately 8.6fps. Processing was conducted on a desktop (i5 2400 CPU @ 3.1GHZ).

For a robot we used a Baxter humanoid upper body on a Ridgeback mobile base (approx. 100 x 80 x 180 cm (l x w x h), weight approx. 210 kg). The robot was equipped with various actuators: two seven degree of freedom (DOF) arms incorporating springs for safety with a reach of 1.2m capable of lifting approx. 2kg (25kg with safety disabled), four independent omnidirectional wheels with maximum speed 1.1 m/s, and a touch display screen showing a face. For sensors, the robot had a LIDAR, a 360 degree ring of 12 ultrasonic range sensors, three additional RGB cameras, two infrared range sensors, force sensors and accelerometers in each arm joint, base encoders, and an inertial measurement unit (IMU). Processing on the robot took place on a i7-3770 Processor (8MB, 3.4GHz) with HD4000 Graphics driver.

We note that, although we found these tools to be useful for our exploration, the proposed approach is not dependent on this implementation and can be applied with different sensors and robots: the thermal and RGB cameras used are standard off-the-shelf components, and various robots are capable of locomotion and manipulation.



**Figure 9.** Robotic manipulation of touched objects can yield additional information to aid inference: (a-b) depict the front and backside of a bottle which has been opened, and (c-d) the front and backside of a bottle which has only been picked up and replaced; the thermal images in (a) and (c) appear similar, with minor signs of touching, but (b) shows heavy touching on the backside compared to (d).

## 4 DISCUSSION

In the current paper, we presented PastVision, an approach to infer past actions by correlating thermally sensed touches and object detection, with a focus on facilitating robust robotic monitoring of medicine intake in dementia patients. We furthermore explored the concept, checking how long touches on typical materials used for medication packaging persisted and how heat traces from typical grips related, and suggesting how robot motion planning can be used to enhance inference results.

### 4.1 Ethics

We believe that privacy is an important consideration and we do not suggest that monitoring is for everyone; it should only be provided for those who want it and foresee benefits which outweigh the downsides. Also, we advocate that monitoring should not lead to control being taken away from people, but rather that it should enhance people's confidence and increase their control over conditions such as dementia. For example, a person with dementia who loves cooking might feel better about using a hot stove if it is believed that someone is there who can remind them if they forget to turn off the stove at the right time.

### 4.2 Limitations

There are many challenges affecting thermal and RGB cameras. Thermal cameras must deal, inter alia, with high variances in ambient temperatures (e.g., in winter or summer), heat contamination when objects are touched multiple times, and thermal reactions from remote heat sources. RGB cameras can be affected by illumination and shadows, and we consider general object detection against complex backgrounds (e.g. with wallpaper or art) to still be an open problem in computer vision. Thus, we emphasize that the current work is a feasibility study,

whose results are limited to a highly simplified scenario chosen for initial investigation: images were recorded in a lab environment with controlled temperature, few objects and a simple background. We intend to extend these results in future work.

### 4.3 Future Work

Our next step will be to conduct further tests and obtain some quantitative results, e.g., for system accuracy in inferring touched objects and discriminating touch types (medicine intake vs. just touching), as well as durations for which traces can be detected. Also, the surface has only been scraped in terms of what kinds of past action inference can be conducted. For example, for thermal inference, temperature changes not only in objects but in people will facilitate inference (e.g., cooling of a person's mouth and hands might provide extra confirmation that the person has taken some oral medication). A variation of the intersection over union (IOU) metric could also be used for detecting which object has been touched, since centroid distances do not take into account bounding box locations. And, how to make sense of spatiotemporally co-occurring thermal touches will be an interesting problem.

Other modalities such as sound and olfaction will also be useful. For example, sounds such as nose-blowing might indicate a person has caught a cold, and hair-drying might indicate a person has taken a shower recently or been out in the rain. Also, by touching objects people can leave scents; the promise of olfaction as an inference modality is suggested by the amazing abilities of some animals to track and locate objects people have touched via smell. We believe such work will contribute to the ability of robots to monitor and care for humans, toward supporting well-being.

## 5 REFERENCES

- [1] Abdelrahman Y, Shirazi AS, Henze N, Schmidt A. 2015. Investigation of Material Properties for Thermal Imaging-Based Interaction. CHI 2015. <http://dx.doi.org/10.1145/2702123.2702290>
- [2] Amft O, Troster G. 2006. Methods for Detection and Classification of Normal Swallowing from Muscle Activation and Sound. Pervasive Health Conference and Workshops.
- [3] Bell AF, Walton K, Chevis JS, Davies K, Manson C, Wypych A, Yoxall A, Kirkby J, Alexander N. 2013. Accessing packaged food and beverages in hospital. Exploring experiences of patients and staff. *Appetite* 60, 231-238.
- [4] Benko H, Wilson AD, Balakrishnan R. 2008. Sphere: Multi-Touch Interactions on a Spherical Display. Proceedings of the 21st annual ACM symposium on User interface software and technology (UIST 2008), 77-86.
- [5] Bhowmik A, Repaka R, Mulaveesala R, Mishra S. 2015. Suitability of frequency modulated thermal wave imaging for skin cancer detection: A theoretical prediction. *Journal of Thermal Biology* 51(2015) 6582.
- [6] Chelvama YK, Zamina N, Steeleb GS. 2014. A Preliminary Investigation of M3DITRACK3R: A Medicine Dispensing Mobile Robot for Senior Citizens. *Procedia Computer Science* 42, 240-246.
- [7] Chen L, Hoey J, Nugent CD, Cook DJ, Yu Z. 2012. Sensor-Based Activity Recognition. *Transactions on Systems, Man, and Cybernetics. Part C: Applications and Reviews*, Vol. 42, No. 6.
- [8] Cohen PR, Sutton C, Burns B. 2002. Learning Effects of Robot Actions using Temporal Associations. Proceedings of the 2nd International Conference on development and Learning. DOI: 10.1109/DEVLRN.2002.1011807
- [9] Correa M, Hermosilla G, Verschae R, Ruiz-del-Solar J. 2012. Human Detection and Identification by Robots Using Thermal and Visual Information in Domestic Environments. *J Intell Robot Syst*, 66:223243. DOI 10.1007/s10846-011-9612-2
- [10] Crick C, Scassellati B. 2008. Inferring Narrative and Intention from Playground Games. Proceedings of the 12<sup>th</sup> IEEE Conference on Development and Learning.
- [11] Dragone M, Saunders J, Dautenhahn K. 2015. On the integration of adaptive and interactive robotic smart spaces. *Paladyn, Journal of Behavioral Robotics* 6(1):165179
- [12] Fitzpatrick P. 2003. First Contact: an Active Vision Approach to Segmentation. Proceedings of 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003). DOI: 10.1109/IROS.2003.1249191
- [13] Gold K, Scassellati B. 2007. A Bayesian Robot That Distinguishes Self from Other. 29th Annual Meeting of the Cognitive Science Society.
- [14] Gray J, Breazeal C, Berlin M, Brooks A, Lieberman J. 2005. Action parsing and goal inference using self as simulator. *RO-MAN 2005*: 202-209
- [15] Iwai D, Sato K. 2005. Heat sensation in image creation with thermal vision. In Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology (ACE '05 ), 213216.
- [16] Jung MM, Poel M, Poppe R, Heylen DKJ. 2017. Automatic recognition of touch gestures in the corpus of social touch. *J Multimodal User Interfaces*, 11:8196. DOI 10.1007/s12193-016-0232-9
- [17] Koppula HS, Saxena A. 2013. Learning Spatio-Temporal Structure from RGB-D Videos for Human Activity Detection and Anticipation. Proceedings of the 30th International Conference on Machine Learning.

- [18] Larson E, Cohn G, Gupta S, Ren X, Harrison B, Fox D, Patel SN. 2011. HeatWave: Thermal Imaging for Surface User Interaction. CHI 2011, Session: Touch 3: Sensing, 2011.
- [19] Lee K, Ikeda T, Miyashita T, Ishiguro H, Hagita N. 2011. Separation of Tactile Information from Multiple Sources Based on Spatial ICA and Time Series Clustering. IEEE/SICE Int. Symposium on System Integration (SII), pp. 791-796
- [20] Li K, Fu Y. 2014. Prediction of Human Activity by Discovering Temporal Sequence Patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 36, No. 8.
- [21] Murthy JN, van Jaarsveld J, Fei J, Pavlidis I, Harrykissoon RI, Lucke JF, Faiz S, Castriotta RJ. 2009. Thermal Infrared Imaging: A Novel Method to Monitor Airflow During Polysomnography. SLEEP, Vol. 32, No. 11.
- [22] Osterberg L, Blaschke T. 2005. Adherence to Medication. N Engl J Med, 353:487-97.
- [23] Pollack M, Engberg S, Thrun S, Brown L, Colbry D, Orosz C, Peintner B, Ramakrishnan S, Dunbar-Jacob J, McCarthy C. 2002. Pearl: A Mobile Robotic Assistant for the Elderly, in AAAI Workshop on Automating as Caregiver.
- [24] Redmon J, Divvala S, Girshick R, Farhadi A. 2016. You only look once: Unified, real-time object detection. In:CVPR.
- [25] Ring EFJ, Ammer K. 2012. Infrared thermal imaging in medicine. Physiol. Meas. 33(3), R33.
- [26] Rudol P., Doherty P. 2008. Human Body Detection and Geolocalization for UAV Search and Rescue Missions Using Color and Thermal Imagery. IEEEAC Paper 1274.
- [27] Shiga Y, Dengel A, Toyama T, Kise K, Utsumi Y. 2014. Daily activity recognition combining gaze motion and visual features. UbiComp Adjunct: 1103-1111
- [28] Shiomi M, Zanlungo F, Hayashi K. et al. 2014. Towards a Socially Acceptable Collision Avoidance for a Mobile Robot Navigating Among Pedestrians Using a Pedestrian Model. Int J of Soc Robotics 6(3): 443-455. doi:10.1007/s12369-014-0238-y
- [29] Shirazi AS, Abdelrahman Y, Henze N, Schneegass S, Khalilbeigy M, Schmidt A. 2014. Exploiting Thermal Reflection for Interactive Systems. Session: Novel Mobile Displays and Devices, CHI 2014.
- [30] Silvera-Tawil D, Rye D, Velonaki M. 2015. Artificial skin and tactile sensing for socially interactive robots: A review. Robotics and Autonomous Systems 63 230-43.
- [31] Stork JA, Spinello L, Silva J, Arras KO. 2012. Audio-Based Human Activity Recognition Using Non-Markovian Ensemble Voting. IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication.
- [32] Stoytchev A. 2005. Behavior-Grounded Representation of Tool Affordances, ICRA.
- [33] Takacs B, Hanak D. 2008. A prototype home robot with an ambient facial interface to improve drug compliance. Journal of Telemedicine and Telecare 2008; 14: 393395.
- [34] Tiwari P, Warren J, Day K, MacDonald B, Jayawardena C, Kuo IH, Igic A, Datta C. 2011. Feasibility study of a robotic medication assistant for the elderly (2011). Conference: Proceedings of the Twelfth Australasian User Interface Conference - Volume 117.
- [35] Wang L, Zhao X, Si Y, Cao L, Liu Y. 2017. Context-Associative Hierarchical Memory Model for Human Activity Recognition and Prediction. IEEE Transactions on Multimedia, Vol. 19, No. 3.
- [36] Wang Z. 2006. A field study of the thermal comfort in residential buildings in Harbin. 2005. Building and Environment 41, 10341039. doi:10.1016/j.buildenv.2005.04.020
- [37] Wong WK, Chew ZY, Lim HL, Loo CK, Lim WS. 2011. Omnidirectional Thermal Imaging Surveillance System Featuring Trespasser and Faint Detection. International Journal of Image Processing (IJIP), Volume 4, Issue 6. version 32.

## 6 ACKNOWLEDGEMENT

We thank those who helped us. We received support from the Swedish Knowledge Foundation for the SIDUS AIR project, and the first author is part of the REMIND project on medicine reminder systems.

# Multi-expert estimations of burglars' risk exposure and level of pre-crime preparation based on crime scene data (position paper)

Martin Boldt, Veselka Boeva and Anton Borg

Department of Computer Science and Engineering  
Blekinge Institute of Technology, Sweden  
{*martin.boldt, veselka.boeva, anton.borg*}@bth.se

May 3, 2017

## Abstract

Law enforcement agencies strive to link crimes perpetrated by the same offenders into crime series in order to improve investigation efficiency. Such crime linkage can be done using both physical traces (e.g., DNA or fingerprints) or “soft evidence” in the form of offenders’ modus operandi (MO), i.e. their behaviors during crimes. However, physical traces are only present for a fraction of crimes, unlike behavioral evidence. This position paper presents a method for aggregating multiple criminal profilers’ ratings of offenders’ behavioral characteristics based on feature-rich crime scene descriptions. The method calculates consensus ratings from individual experts’ ratings, which then are used as a basis for classification algorithms. The classification algorithms can automatically generalize offenders’ behavioral characteristics from cues in the crime scene data. Models trained on the consensus rating are evaluated against models trained on individual profiler’s ratings. Thus, whether the consensus model shows improved performance over individual models.

**Keywords:** Multi-expert decision making, classification, crime linkage, offender profiling.

## 1 Introduction

For crime categories that involve serial offenders, i.e. where the same offender commits two or more

crimes in the same crime category (e.g. burglaries or robberies), law enforcement agencies strive to link those crimes into crime series [9, 14]. The linking of crimes enables investigators to get a more complete understanding, based on the combined knowledge and evidence collected from the different crime scenes, compared to investigating each crime in isolation [16]. In addition, such linking also enable more efficient use of the police force’s scarce resources compared to investigating each crime individually [14].

The linking of crimes into series can be done based on physical evidence, e.g., DNA or fingerprints. However, such evidence is only available in a fraction of all crime scenes and the processing of physical evidence is also costly and time-consuming [8]. Thus, it is many-times difficult for law enforcement agencies to handle large amounts of physical evidence from high-volume crime categories, e.g., burglaries [15]. However, “soft evidence” in the form of offender’s modus operandi (MO), i.e. habits, techniques and peculiarities of behavior when committing an offense [12], could be used for linking crimes committed in a similar fashion [1, 10]. This is done by 1) registering behavioral patterns at the crime scenes, 2) interpreting the patterns using criminal and geographical profiling, 3) comparing the behaviors of (unknown) offenders at different crimes, and 4) consider linking similar crimes into series [2, 19]. The process of linking crimes based on soft evidence is known as *crime*

*linkage* and it rests on two key assumptions. First, the offender consistency hypothesis [6] stating that offenders display similar behaviors across time and place. Second, the offender specificity hypothesis [13] describing that offenders have an approach that deviate or is distinct from other offenders' approaches. These assumptions apply for offender behavior in residential burglaries [25]. The accumulated evidence published during the past 20 some years provide a base for conducting crime linkage based on behavioral consistency and distinctiveness for some offenders [20], some of the time, in various crime categories and types, e.g. commercial and residential burglary [2, 5, 11, 24].

A drawback with crime linkage is that crime analysts and offender profilers are required to manage a substantial amount of information that brings along a heavy cognitive load [17]. Law enforcement agencies have created computerized databases enclosing a large number of reported crimes that analysts can search for similarities [26]. Various data science methods and algorithms can be used in combination with such databases, e.g. in order to calculate similarity scores based on pair-wise comparisons of crimes. Further, learning algorithms from the area of computer science and the field of machine learning can be used for grouping crimes with similar characteristics together, e.g., by using cluster algorithms [4]. Such intelligent models can be packaged in decision-support systems that assist crime analysts in the crime linkage process.

In a previous initial study, we investigated the possibility of using intelligent models to automatically estimate both offenders' risk exposure and their level of pre-crime preparation for residential burglaries [3]. We did this using soft evidence in the form of MO characteristics that were recorded from crime scenes. The main motivation for an automated approach is because volume crimes (e.g. various forms of thefts, including burglaries) occur with such a high frequency that criminal profilers can not analyze and rate each individual crime instance manually. Therefore, models that automatically estimate behavioral characteristics for unknown offenders using crime scene data, would be valuable for law enforcement agencies. Mainly since they would allow law enforcement agencies to use those behavioral characteristics in the crime linkage process, i.e. when linking sets of crimes to common offenders. Another benefit of such models would be

that the behavioral characteristics that they calculate could be compared for crimes *between* different crime categories (e.g. burglaries and Diesel thefts) that otherwise more or less lack comparable MO behaviors. The previous initial study included an experimental evaluation of 16 learning algorithms and the models trained by the Naïve Bayes Multinomial algorithm both showed interesting results and was the most suitable candidates for the problem at hand. However, further research is required before such models could be used in an operational setting by law enforcement agencies.

In this work a planned extended study is described, which investigates whether multi-expert decision making concepts can improve the performance of intelligent models when estimating offender characteristics. The remainder of this position paper is structured as follows, in Section 2 we summarize the previous study. In Section 3 we elaborate on how we intend to extend the previous study. Finally, in Section 4 we present some conclusions and avenues for future work.

## 2 Estimating offender characteristics

In the previous study, two criminal profilers from the national offender profiling group within the Swedish police manually rated offenders' risk exposure and level of pre-crime preparation for 50 burglaries each. However, there was an overlap of 25 burglaries to allow the analysis of inter-rater agreement, which turned out to be moderate. For each burglary a feature-rich and structured representation of the crime scene was available, containing some 137 features, e.g. type of residence and which entrance method the offender used. Based on the manual ratings of offender characteristics provided by the profilers, a labeled dataset of 75 instances was created. Each instance has an offender risk exposure score as well as a pre-crime preparation score together with the feature-rich crime scene data. Both risk and plan scores were rated using the following scale: *low, decreased, increased, high*.

An experiment was executed in which 16 machine-learning algorithms were evaluated using a supervised-learning approach on the labeled dataset. Performance evaluation was done us-



ing stratified 10-times 10-fold cross-validation tests. Five performance metrics were used, with AUC as the primary metric. Statistical analysis of the AUC performance was done using the non-parametric Kruskal-Wallis test in combination with the Nemenyi post-hoc test. Models trained by the Naïve Bayes Multinomial algorithm outperformed more competing models than any of the other algorithms, and was therefore selected as the most suitable candidate for the problem at hand. The AUC measures were 0.79 (sd=0.15) and 0.77 (sd=0.16) for estimating offenders' risk and preparation scores respectively. The classification performance of the models were not excellent, but given that this was an initial study (with a quite limited dataset) the results indicate that models can pick up cues in the feature-rich crime scene data that are useful when generalizing offenders' risk and preparation scores.

Next, the models trained by Naïve Bayes Multinomial were then used for calculating both offenders' risk exposure scores as well as pre-crime preparation scores for 15,598 residential burglaries that all contained feature-rich and structured crime descriptions. For a subset of 153 burglaries (out of the 15,598) the police provided anonymized identifiers of the offenders, which allowed us to construct 41 linked crime series of linked burglaries. For each of the 41 crime series we calculated the variation of both the risk exposure and pre-crime preparation scores. We then compared those variations with randomly constructed "crime series" that contained equally many burglaries. Differences in score consistency between linked series and random series were studied using a Wilcoxon signed rank test, which showed that scores were significantly more consistent in linked series compared to random ones. Further analysis revealed that the scores also showed promising distinctiveness between linked series, as well as consistency for crimes within series compared to randomly sampled crimes. This indicates the usefulness of automatic models for estimating offenders' risk exposure and degree of pre-crime preparation. Further, that such behavioral scores could be used as a complement to traditional crime scene data in the crime linkage process when law enforcement officers try to link crimes together that most probably are committed by the same offender.

### 3 Multi-expert decision making

As crime profilers have different education, experience and domain knowledge it is interesting to consider the manual rating of offender's crime characteristics using a multi-expert decision making approach. Such an approach enables both inter-rater agreement, as well as weighting of each individual rater's relative importance in order to reach a consensus decisions regarding the scores. When linking crimes it is important that the decisions are as correct as possible, since the decisions made in the crime linkage process regarding which crimes to link into, or exclude from, series can have impact on the daily lives of several persons. It is therefore important that manual individual ratings of offender characteristics by profilers are aggregated in a systematic way into sound team evaluation ratings.

Although the work discussed in this paper investigates offender's risk exposure and planning for residential burglaries, the method is applicable on other crime types as well. Profilers, when manually rating offender's characteristics, use the same working method in more violent crimes as well, e.g., in rape, assault, and murders. As such, the possibility to have an initial automated estimate to compare the law enforcement officers' decision against is useful. By training models on both aggregated decisions (based on all profilers ratings), as well as individual raters decisions allow for detection of outlier ratings. The detection of outlier ratings is interesting as an outlier may be more important than the aggregated decision. This might be in case the individual profiler possesses unusual and valuable domain knowledge otherwise lacking in the group, i.e. when an outlier is the correct decision, and which is not represented in the profiler weights. Such outlier ratings need to be manually investigated in more detail.

#### 3.1 Proposed approach

In this work, we propose a method that mimics a multi-expert estimation process of a team of criminal profilers who are involved in the rating of a given set of crimes. The defined process can be divided in the following two phases: (i) an individual estimation phase engaging individual profilers,

and (ii) a consensus and decision phase involving the interaction between the team members. During the estimation phase each profiler is asked to rate/score, independently from the other profilers, both offender's risk exposure and level of pre-crime preparation for a set of crimes. In addition, each profiler gives his/her opinion about the expertise of the other profilers, by assigning weights to each one, including himself/herself. During the consensus and decision phases, the crime scores and profiler weights provided by each individual profiler are aggregated into team-based consensus scores. Then, learning algorithms are applied to generalize from these team-based consensus scores, using the feature-rich structured crime scene descriptions.

The proposed multi-expert decision making method is based on an aggregation approach developed by Tsiporkova and Boeva [22]. It resolves potential conflicts between the team members by mimicking a multi-step decision making process during which the decision makers have an opportunity to discuss and exchange views, ideas, information, etc. At each decision step each expert (profiler) aggregates the outcome of the previous step according to the set of weights he/she has assigned to himself/herself and the rest of the profilers. These weights express the relative degree of influence each profiler is inclined to accept from the rest of the team when forming a judgment for the different crimes.

The proposed multi-expert approach has a few advantages. First, that profilers are not allowed to completely ignore colleagues by using zero weights and in this way putting the decision process in a deadlock. This requirement is imposed in order to ensure that the applied recursive aggregation process is convergent [21]. Second, that the approach reflects different team interaction styles by incorporating mechanisms to deal with the reputation of the profilers [22, 23]. The goal is to minimize the total dissatisfaction of each crime rating by considering the profilers' particular views and implementing reputation-enhanced collaboration, but avoiding any kind of explicit public profilers rating. The latter is crucial since it contributes to a positive collaboration atmosphere. Finally, the approach translates the chosen six level interval scale used by individual profilers, which spans *low* to *high* represented as [1–6], into a continuous scale, i.e. all values between 1.0 and 6.0 (inclusive). As it was dis-

cussed in the work by Sullivan et al. [18] responses can be rated or ranked in an ordinal scale, but the distance between responses is not measurable. In other words, one cannot assume that the difference between responses is equidistant even though the numbers assigned to those responses are. In contrast to this when we have interval data, the difference between responses can be calculated and the numbers can be used for further analysis and discrimination of the rating crimes.

We also plan to use the proposed multi-expert decision making model to simulate different team interaction styles and study how these affect the estimation performance of the applied learning models. The three team interaction styles described by Cooke and Szumal [7] will be investigated, that is the *constructive*, *passive*, and *aggressive* interaction styles. The constructive style is characterized by a balanced concern for personal and team outcomes. The passive style places greater emphasis on fulfillment of affiliation goals only, i.e. maintaining harmony in the team. While the aggressive style is characterized with that personal ambition is placed above concern for team outcome. Cooke and Szumal demonstrate that groups that predominantly use a constructive interaction style produce solutions that are: 1) superior in quality to those produced by passive groups, and 2) superior in acceptance (satisfaction) to those produced by either passive or aggressive groups.

### 3.2 Study design

The planned extended study consists of the following three phases. In the first phase the 4-6 profilers included in the study will participate in a workshop that presents the study and the tasks that should be carried out. During the workshop the profilers should individually rate the risk and plan scores for 25 burglaries. Next, they should come to joint decisions (both for risk and plan scores) for each of the 25 burglaries. This workshop will apart from describing the problem domain, allow participants to discuss the crime scene data available and how to interpret and judge it by relating to concrete examples represented by the 25 burglaries. These discussions will also give the profilers a more detailed opinion about the other profilers' competence, which will be valuable in the next phase.

The second phase of the study includes two tasks:

(i) that each rater individually rates the relative degree of influence he/she puts on each of the other profilers and him/herself, and (ii) that each profiler independently rates the risk and plan scores for another 125 burglaries. This work could be done little-by-little over the course of a couple of weeks when the profilers have some time to spare. Once all profilers have finished their tasks, the data is collected. Then it is possible to train individual models for estimation of both risk and plan scores based on the ratings from each individual profiler, i.e. 4-6 models (one per profiler) is created. The learning algorithm used for training the models will be Naïve Bayes Multinomial as it showed best performance in the previous study. Next, the consensus scores are calculated using the ratings from all profilers and the matrix of profiler weights according to the iterative method previously described in Section 3.1. An experiment is used to compare the estimation performance of the individual models against the consensus model over the full labeled dataset consisting of 150 burglaries. Evaluation is handled using stratified 10-times 10-fold cross-validation and suitable statistical tests. This is done twice, one time for the models that estimate offender risk exposure scores, and another time for the models that estimate offender pre-crime preparation scores.

The third phase of the study involves a simulation of the three team interaction styles according to the work by Cooke and Szumal, i.e. the constructive, passive, and aggressive styles. The matrix of profiler weights will be changed according to each of the three interaction styles. For each new weight matrix alternative consensus decisions for both the risk and plan scores will be calculated, and new models will be trained on these new scores. Finally, the estimation performance of each model will be compared to each other using a similar experimental setup as in phase two above.

## 4 Conclusion and future work

This position paper presents a method for aggregating multiple experts ratings of offender behavioral characteristics using a feature-rich crime scene dataset. The method calculates consensus ratings from individual experts' ratings, which then are used as a basis for machine learning algorithms.

The learning algorithms can be used to automatically generalize offenders' behavioral characteristics from cues in the crime scene data. Models trained on the consensus rating are evaluated against models trained on individual profilers ratings. Thus, investigating if there is any improvement in the performance of the consensus model compared to the individual models.

As future work it would be interesting to investigate if the use of consensus ratings as seeds for clustering algorithm could generate more informative clusters. This could further be extended as an interactive clustering approach.

## References

- [1] Bennell, C. and Canter, D.V., Linking commercial burglaries by modus operandi: tests using regression and ROC analysis, *Science & Justice*, 42, 2002, pp. 153-164.
- [2] Bennell, C. and Jones, N.J., Between a ROC and a hard place: A method for linking serial burglaries by modus operandi, *Journal of Investigative Psychology and Offender Profiling*, 2, 2005, pp. 23-41.
- [3] Boldt, M., Borg, A., Svensson, M., Hildeby, J., Predicting burglars' risk exposure and level of pre-crime preparation using crime scene data, (to appear in) *Journal of Intelligent Data Analysis*, 2018, 22(1).
- [4] Borg, A. and Boldt, M., Clustering Residential Burglaries Using Modus Operandi and Spatiotemporal Information, *International Journal of Information Technology & Decision Making*, 2016, 15(1), pp. 23-41.
- [5] Bouhana, N., Johnson, S.D. and Porter, M., Consistency and specificity in burglars who commit prolific residential burglary, *Legal and Criminological Psychology*, 2014, 21(1), pp. 77-94.
- [6] Canter, D., Psychology of offender profiling, Handbook of psychology in legal contexts, Chichester, UK, John Wiley and Sons, 2000.
- [7] Cooke, R.A. and Szumal, J.L. The impact of group interaction styles on problem-solving ef-

- fectiveness, *Journal of Applied Behavioral Science*, 1994, 30, pp. 415-437.
- [8] Davies, A., The use of DNA profiling and behavioural science in the investigation of sexual offences, *Medicine, Science and the Law*, 1991, 31, pp. 95-101.
- [9] Grubin, D., Kelly, P. and Brunsdon, C., Linking serious sexual assaults through behaviour [Internet], 2001 [cited 2017-03-30], Available from: <http://webarchive.nationalarchives.gov.uk/20110314171826>
- [10] Hazelwood, R.R. and Warren, J.I., Linkage analysis: Modus operandi, ritual, and signature in serial sexual crime, *Aggression and Violent Behavior*, 8, 2003, pp. 587-598.
- [11] Markson, L., Woodhams, J., and Bond, J.W., Linking serial residential burglary: comparing the utility of modus operandi behaviours, geographical proximity, and temporal proximity, *Journal of Investigative Psychology and Offender Profiling*, 2010, 7(2), pp. 91-107.
- [12] O'Hara, C.E. and O'Hara, G.L., *Fundamentals of Criminal Investigation* (7th ed.), Springfield, IL, Charles C Thomas Publisher Ltd., 1956.
- [13] Pervin, L.A., *Current Controversies and Issues in Personality*, New York, John Wiley and Sons, 2002.
- [14] Reich, B.J. and Porter, M.D., Partially supervised spatiotemporal clustering for burglary crime series identification, *Journal of Royal Statistical Society, Series A (Statistical Society)*, 178, 2015, pp. 465-480.
- [15] Roman J., Reid, S., Reid, J., Chalfin, A., Adams, W. and Knight, C., The DNA Field Experiment: Cost-effectiveness Analysis of the Use of DNA in the Investigation of High-volume Crimes [Internet], 2008 [cited 2017-03-30], Available from: [http://www.urban.org/UploadedPDF/411697\\_dna\\_field\\_experiment.pdf](http://www.urban.org/UploadedPDF/411697_dna_field_experiment.pdf)
- [16] Rossmo, K., *Geographic Profiling*, Boca Raton, CRC Press, 1999.
- [17] Santtila, P., Korpela, S. and Häkkinen, H., Expertise and decision making in the linking of car crime series, *Psychology, Crime & Law*, 2004, 10(2), pp. 97-112.
- [18] Sullivan, G.M. and Artino, A.R., Analyzing and Interpreting Data from Likert-Type Scales, *Journal of Graduate Medical Education*, 2013, pp. 541-542.
- [19] Tonkin, M., Woodhams, J., Bull, R. and Bond, J.W., Behavioural case linkage with solved and unsolved crimes, *Forensic Science International*, 212, 2012, pp. 146-153.
- [20] Tonkin, M. and Woodhams, J., The feasibility of using crime scene behaviour to detect versatile serial offenders: An empirical test of behavioral consistency, distinctiveness, and discrimination accuracy, *Legal and Criminological Psychology*, 2015, 22(1), pp. 99-115.
- [21] Tsiporkova, E. and V. Boeva, V. Nonparametric recursive aggregation process, *Kybernetika, The Journal of the Czech Society for Cybernetics and Inf. Sciences* 2004, 40(1), pp. 51 - 70.
- [22] Tsiporkova, E. and Boeva, V., Multi-step ranking of alternatives in a multi-criteria and multi-expert decision making environment, *Information Sciences*, 2006, 176(18), pp. 2673-2697.
- [23] Tsiporkova, E., Tourwè, T. and Boeva, V. A Collaborative Decision Support Platform for Product Release Definition. *The Fifth International Conference on Internet and Web Applications and Services ICIW 2010*, 2010, pp. 351-356.
- [24] Woodhams, J., Hollin, C.R. and Bull, R., The psychology of linking crimes: A review of the evidence, *Legal and Criminological Psychology*, 2010, 12(2), pp. 233-249.
- [25] Wright, R.T. and Decker, S.H., *Burglars on the job*, Boston, MA, Northeastern University Press, 1994.
- [26] Yokota, K. and Watanabe, S., Computer-Based Retrieval of Suspects Using Similarity of Modus Operandi, *International Journal of Police Science & Management*, 2002, 4(1), pp. 5-15.

## Advanced Data-driven Techniques for Mining Expertise

Milena Angelova<sup>1</sup>, Veselka Boeva<sup>2</sup> and Elena Tsiporkova<sup>3</sup>

<sup>1</sup>Technical University of Sofia-branch Plovdiv, Bulgaria [mangelova@tu-plovdiv.bg](mailto:mangelova@tu-plovdiv.bg)

<sup>2</sup>Blekinge Institute of Technology, Karlskrona, Sweden

<sup>3</sup>Sirris, The Collective Center for the Belgian technological industry, Brussels, Belgium

### Abstract

In this work, we discuss enhanced techniques that optimize expert representation and identify subject experts via clustering analysis of the available online information. We use a weighting method to assess the levels of expertise of an expert to the domain-specific topics. In this context, we define a way to estimate the expertise similarity between experts. Then the experts finding task is viewed as a list completion task and techniques that return similar experts to ones provided by the user are considered. In addition, we discuss a formal concept analysis approach for clustering of a group of experts with respect to given subject areas. The produced grouping of experts can further be used to identify individuals with the required competence.

### Keywords

Data mining, expert finding, health science, knowledge management.

## 1 INTRODUCTION

Nowadays, organizations search for new employees not only relying on their internal information sources, but they also use data available on the Internet to locate required experts. As the data available is very dispersed and of distributed nature, a need appears to support this process using IT-based solutions, *e.g.*, information extraction and retrieval systems, especially expert finding systems. Expert finding systems however, need a lot of information support. On one hand, the specification of required "expertise need" is replete with qualitative and quantitative parameters. On the other hand, the expert finders need to know whether a person who meets the specified criteria exists, how extensive his/her knowledge or experience is, whether there are other persons who have the similar competence, how he/she compares with others in the field, *etc.* Consequently, techniques that gathers and makes such information accessible are needed.

Many practical scenarios of organizational situations that lead to expert seeking have been extensively presented in the literature. For instance, Autonomy [1] analyses users' search and publication histories to determine concepts that are indicative of their expertise. Yenta [2] and Tacit KnowledgeMail [3] determine user expertise from email message traffic. Referral Web from AT&T [4] provides access to experts across an expertise or community, aiming to make the basis for referral transparent to the user. In recruitment industry, the problem of finding expertise is an one of seeking for job candidates given the required skills as well as some additional information, such as, location and/or company names [5]. Several Web-based expert seeking tools that support both type players at the job market have recently appeared [6][7]. For instance, in [6],

a personalized job seeking for an applicant is proposed by given him/her benchmark against the current job market.

Expert finders are usually integrated into organizational information systems, such as knowledge management systems, recommender systems, and computer supported collaborative tasks. Initial approaches propose tools that rely on people to self-assess their skills against a predefined set of keywords, and often employ heuristics generated manually based on current working practice [8]. Later approached try to find expertise in specific types of documents, such as e-mails [9][10] or source code [11]. Instead of focusing only on specific document type systems that index and mine published intranet documents as sources of expertise evidence are discussed in [12]. In the recent years, research on identifying experts from online data sources has been gradually gaining interest [13][14][15][16][17].

In this work, we discuss enhanced techniques that optimize expert representation and identify subject experts via clustering analysis of the available online information. In [23], we have proposed a weighting method to assess the levels of expertise of an expert to the domain-specific topics. In this context, we have further defined a way to estimate the expertise similarity between experts. Then the experts finding task is viewed as a list completion task and techniques that return similar experts to ones provided by the user are considered. In addition, we have proposed a formal concept analysis approach for clustering of a group of experts with respect to given subject areas [33]. The produced grouping of experts can further be used to identify individuals with the required competence. The proposed expert finding techniques have been evaluated on data extracted from PubMed repository.

## 2 PROPOSED SOLUTIONS

Many scientists who work on the expertise retrieval problem distinguish two information retrieval tasks: *expert finding* and *expert profiling*, where *expert finding* is the task of finding experts given a topic describing the required expertise [18], and *expert profiling* is the task of returning a list of topics that a person is knowledgeable about [19].

In this work, we consider data-driven techniques that deal with both expertise retrieval tasks. Initially, we need to describe the expertise of each involved person by creating his/her expert profile, *i.e.* each person is associated by a list of subjects he/she is an expert in.

### 2.1 Expert profiling

The data needed for constructing the expert profiles could be extracted from various Web sources, *e.g.*, LinkedIn, the DBLP library, Microsoft Academic Search, Google Scholar Citation, PubMed etc. There exist several open tools for extracting data from public online sources. For instance, Python LinkedIn is a tool which can be used in order to execute the data extraction from LinkedIn. In addition, the Stanford part-of-speech tagger [20] can be used to annotate the different words in the text collected for each expert with their specific part of speech. Next to recognizing the part of speech, the tagger also defines whether a noun is plural, whether a verb is conjugated, etc. Further the annotated text can be reduced to a set of keywords (tags) by removing all the words tagged as articles, prepositions, verbs, and adverbs. Practically, only the nouns and the adjectives are retained.

However, an expert profile may be quite complex and can, for example, be associated with information that includes: e-mail address, affiliation, a list of publications, co-authors, but it may also include or be associated with: educational and (or) employment history, the list of LinkedIn contacts etc. All this information could be separated into two parts: the expert's personal data and information that describes the competence area of expert.

The expert's personal data can be used to resolve the problem with ambiguity. This problem refers to the fact that multiple profiles may represent one and the same person and therefore must be merged into a single generalized expert profile, *e.g.*, the clustering algorithm discussed in [21] can be applied for this purpose. A different approach to the ambiguity problem has been proposed in [22]. Namely, the similarity between the personal data (profiles) of experts is used to resolve the problem with ambiguity. The split and merge of expert profiles is driven by the calculation of similarity measure between the different entities composing the profile, *e.g.* expert name, email address, affiliations, co-authors names etc. Thus the similarity between the personal data of two expert profiles is defined by the weighted mean of similarities between the corresponding fields of their profiles [22].

In [23], we use a Dynamic Time Warping (DTW) based approach to deal with the ambiguity issue. In general, the

DTW alignment algorithm finds an optimal match between two given sequences (*e.g.*, time series) by warping the time axis iteratively until an optimal matching (according to a suitable metric) between the two sequences is found [24]. Due to its flexibility, DTW is widely used in many scientific disciplines and business applications as *e.g.*, speech processing, bioinformatics, matching of one-dimensional signals in the online hand writing communities etc. A detail explanation of DTW algorithm can be found in [24][25].

In view of the above, an expert profile can be defined as a list of keywords (domain-specific topics), extracted from the available information about the expert in question, describing her/his subjects of expertise. Assume that  $n$  different expert profiles are created in total and each expert profile  $i$  ( $i = 1, 2, \dots, n$ ) is represented by a list of  $p_i$  keywords.

### 2.2 Assessing of expertise

An expert may have more extensive knowledge or experience in some topics that in others and this should be taken into account in the constructions of expert profiles. Thus the gathered information about each individual expert can further be analyzed and used to access her/his levels of expertise to the different topics that compose her/his expert profile.

There is no standard or absolute definition for accessing expertise. This usually depends not only on the application area but also on the subject field. For instance, in the peer-review setting, appropriate experts (reviewers, committee members, editors) are discovered by computing their profiles, usually based on the overall collection of their publications [26]. However, the publication quantity alone is insufficient to get an overall assessment of expertise. To incorporate the publication quality in the expertise profile, Cameron used the impact factor of publications' journals [26]. However, the impact factor in itself is arguable [27][28]. Therefore, Hirsch proposed another metric, the "HIndex", to rank individuals [29]. However, this index works fine only for comparing scientists working in the same field, because citation conventions differ widely among different fields [29]. Afzal et al. proposed an automated technique which incorporates multiple facets in providing a more representative assessment of expertise [30]. The developed system mines multiple facets for an electronic journal and then calculates expertise' weights.

In [23], we have proposed a weighting method to assess the levels of expertise of an expert to the domain-specific topics. Namely, weights are used to access the relative levels of knowledge or experience an individual has in the topics he/she has shown to have an expertise. Let us suppose that a weighting method appropriate to the respective area is used and as a result each keyword (domain-specific topic)  $k_{ij}$  of expert profile  $i$  ( $i = 1, \dots, n$ ) is associated with a weight  $w_{ij}$ , expressing the relative level (intensity) of expertise the expert in question has in the topic  $k_{ij}$ , *i.e.*  $\sum_{j=1}^{p_i} w_{ij} = 1$  and  $w_{ij} \in (0, 1]$  for  $i = 1, \dots, n$ .

In this way, each expert can be presented by a richer expert profile describing the topics (keywords) in which he/she is an expert plus the levels (weights) of knowledge or

experience he/she has in the different topics. Namely, each expert is represented by two components: a list of keywords (topics) and a vector of weights expressing the relative levels of expertise the expert has in the different topics.

### 2.3 Expertise similarity

The calculation of expertise similarity is a complicated task, since the expert expertise profiles usually consist of domain-specific keywords that describe their area of competence without any information for the best correspondence between the different keywords of two compared profiles. Therefore, it is proposed in [22] to measure the similarity between two expertise profiles as the strength of the relations between the semantic concepts associated with the keywords of the two compared profiles. Another possibility to measure the expertise similarity between two expert profiles is by taking into account the semantic similarities between any pair of keywords that are contained in the two profiles.

Accurate measurement of semantic similarity between words is essential for various tasks such as, document (or expert) clustering, information retrieval, and synonym extraction. Semantically related words of a particular word are listed in manually created general-purpose lexical ontologies such as WordNet. WordNet is a large lexical database of English [31][32]. Initially, the WordNet networks for the four different parts of speech were not linked to one another and the noun network was the first to be richly developed. This imposes some constraints on the use of WordNet ontology. Namely, most of the researchers who use it limit themselves to the noun network. However, not all keywords representing the expert profiles are nouns. In addition, the algorithms that can measure similarity between adjectives do not yield results for nouns hence the need for combined measure. Therefore, a normalized measure combined from a set of different similarity measures is defined and used in [33] to calculate the semantic relatedness between any two keywords. In the considered context the expertise similarity task is additionally complicated by the fact that the competence of each expert is represented by two components: a list of keywords describing her/his expertise and a vector of weights expressing the relative levels of knowledge/expertise the expert has in the different topics.

Let  $s$  be a similarity measure that is suitable to estimate the semantic relatedness between any two keywords used to describe the expert profiles in the considered domain. Then the expertise similarity  $S_{ij}$  between two expert profiles  $i$  and  $j$  ( $i \neq j$ ), can be defined by using the weighted mean of semantic similarities between the corresponding keywords

$$S_{ij} = \sum_{l=1}^{p_i} \sum_{m=1}^{p_j} W_{lm} \cdot s(k_{il}, k_{jm}), \quad (1)$$

where  $W_{lm} = w_{il} \cdot w_{jm}$  is a weight associated with the semantic similarity  $s(k_{il}, k_{jm})$  between keywords  $k_{il}$  and  $k_{jm}$ , and  $W_{lm} \in (0, 1]$  for  $i = 1, \dots, p_i$  and  $m = 1, \dots, p_j$ . It can easily be shown that  $\sum_{l=1}^{p_i} \sum_{m=1}^{p_j} W_{lm} = 1$ .

### 2.4 Identifying experts through clustering

In [33], we have proposed a formal concept analysis approach for grouping a given set of experts with respect to pre-defined subject areas. Initially, the domain of interest is described at some level of abstraction by partitioning the domain to a number of subject areas. Next each expert is represented by a vector of contributions (membership degrees) of the expert to the different areas. This defines an overlapping partition, which is further analysed and refined into a disjoint one by applying Formal Concept Analysis (FCA).

A conceptual model of the domain of interest, such as a thesaurus, a taxonomy etc., can be available. In this case, usually a set of subject terms (topics) arranged in hierarchical manner (tree structures) is used to represent concepts in the considered domain. Further it can be supposed that the tree structure describing the considered domain has  $k$  main branches (broad subject categories). Another possibility to represent the domain of interest at a higher level of abstraction is to partition the set of all different keywords used to define the expert profiles into  $k$  groups (main subject areas). The latter idea has been proposed and applied in [34]. Initially, a common set of all different keywords is formed by pooling the keywords of all the expert profiles. Then the semantic distance between each pair of keywords is calculated and the keywords are partitioned by applying a selected clustering algorithm.

As discussed above, the domain of interest can be presented by  $k$  main subject categories  $C_1, C_2, \dots, C_k$ . Let us denote by  $b_{ij}$  the number of keywords from the expert profile of expert  $i$  that belong to category (subject area)  $C_j$ . In [33], we have assumed that each expert  $i$  ( $i = 1, 2, \dots, n$ ) is described by only a list of the domain-specific topics (keywords) in which he/she is an expert. Then this representation can be converted into a vector  $e_i = (e_{i1}, e_{i2}, \dots, e_{ik})$ , where  $e_{ij} = b_{ij}/p_i$  ( $j = 1, 2, \dots, k$ ) and  $p_i$  is the total number of keywords in the expert profile.

In this way, each expert  $i$  is represented by a  $k$ -length vector of membership degrees of the expert to  $k$  different subject categories, *i.e.* the above procedure generates a *fuzzy* clustering. Thus an expert will have a membership degree of 1 to a certain subject area only in case all the keywords of her/his expert profile are assigned to the category in question. The resulting fuzzy partition can easily be turned into a *crisp* one by assigning to each pair (expert, area) a binary value (0 or 1), *i.e.* for each subject area we can associate those experts who have membership degrees greater than a preliminary given threshold (*e.g.* 0.5). Notice, this partition is not guaranteed to be disjoint in terms of the different subject area, since there will be experts who will belong to more than one subject category.

The above overlapping partition is further analyzed and refined into a disjoint one by applying Formal Concept Analysis. FCA is a principled way of automatically deriving a hierarchical conceptual structure from a collection of objects and their properties [35]. The approach takes as input a matrix (referred to formal context) specifying a set of objects and the properties thereof, called attributes.

In our case, a (formal) **context** consists of the set of the  $n$  experts, the set of main categories  $\{C_1, C_2, \dots, C_k\}$  and an indication of which experts are associated with which subject category. Thus the context is described as a matrix, with the experts corresponding to the rows and the categories corresponding to the columns of the matrix, and a value 1 in cell  $(i, j)$  whenever expert  $i$  is associated with (has expertise in) subject area  $C_j$ . Subsequently, a (formal) **concept** for this context is defined to be a pair  $(X, Y)$  such that

- $X$  is a subset of experts and  $Y$  is a subsets of subject areas, and every expert in  $X$  belongs to every area in  $Y$
- for every expert that is not in  $X$ , there is a subject area in  $Y$  that does not contain that expert
- for every subject area that is not in  $Y$ , there is an expert in  $X$  who is not associated with that area.

The family of these concepts obeys the mathematical axioms defining a **concept lattice** [35]. The built lattice consists of concepts where each one represents a subset of experts belonging to a number of subject areas. The set of all concepts partitions the experts into a set of disjoint expert areas.

Evidently, the produced grouping of experts facilitate the identification of individuals with the required competence. For instance, if we need to recruit experts who have expertise simultaneously in two subject categories, we can directly locate those who belong to the concept that unites the corresponding categories. In addition, such a grouping of experts can be performed with respect to any set of subject areas describing the domain of interest, *e.g.*, the experts could be clustered on a lower level of abstraction by using more specific topics. It is even possible to further produce a grouping of experts belonging to a particular concept around topics specifying the subject areas associated with this concept.

## 2.5 Finding similar experts

The experts finding task can also be viewed as a list completion task, *i.e.* the user is supposed to provide a small number of example experts who have been used to work on similar problems in the past, and the system has to return similar experts. In [23], we have proposed techniques that return similar experts to ones provided by the user.

The concept of expertise spheres has been introduced in [22]. Conceptually, these expertise spheres are interpreted as groups of experts who have strongly overlapping competences. In other words, the expertise sphere can be considered as a combination of pieces of knowledge, skills,

proficiency etc. that collectively describe a group of experts with similar area of competence. Consequently, the user may find experts with the required expertise by entering the name(s) of example expert(s) and the system will return a list of experts with close (similar) expertise by constructing the expertise sphere of the given expert(s).

In order to build an expertise sphere of an expert it is necessary to identify experts with similar area of competence, *i.e.* for each example expert  $i$  a list of expert profiles which exhibit at least minimum (preliminary defined) expertise similarity with her/his expert profile needs to be generated. An expert profile  $j$  will be included in the expertise sphere of  $i$  if the following inequality holds  $S_{ij} \geq T$ , where  $T \in (0, 1)$  is a preliminary defined threshold. The experts identified can be ranked with respect to their expertise similarities to the example expert.

Another possibility is to present the domain of interest by several preliminary specified subject categories and then the available experts can be grouped with respect to these categories into a number of disjoint expert areas (clusters) by using some clustering algorithm, as *e.g.* [33][34]. In the considered context each cluster of experts can itself be interpreted as an expertise sphere. Namely, it can be thought as the expertise area of any expert assigned to the cluster and evidently, the all assigned experts are included in this sphere. In this case, in order to select the right individuals for a specified task the user may restrict her/his considerations only to those experts who are within the expert area (cluster) that is identical with (or at least most similar to) the task's subject. The specified subject and the expert area can themselves be described by lists of keywords (subject profiles), *i.e.* they can be compared by way of similarity measurement. In this scenario, weights can also be introduced by allowing the user to express her/his preferences about the relative levels of expertise the experts in query should have in the specified topics. In addition, the subject profiles that are used to present the different clusters of experts can also be supplied with weights. The experts in the selected cluster can be ranked with respect to the similarity of their expert profiles to the specified subject profile.

In case of a newly extracted (registered, discovered) expert we can classify him/her into one of the existing clusters of experts by determining his/her expertise sphere. Namely we initially calculate the expert's expertise spheres with respect to any of the considered expert areas. Then the expert in question is assigned to that cluster of experts for which the corresponding expertise sphere has the largest cardinality, *i.e.* the overlap between the two sets of experts is the highest.

## 3 EVALUATION AND DISCUSSION

### 3.1 PubMed data

The data needed for constructing the expert profiles are extracted from PubMed, which is one of the largest repositories of peer-reviewed biomedical articles published worldwide. Medical Subject Headings (MeSH) is a controlled vocabulary developed by the US National



Library of Medicine for indexing research publications, articles and books. Using the MeSH terms associated with peer-reviewed articles published by Bulgarian authors and indexed in the PubMed, we extract all such authors and construct their expert profiles. An expert profile is defined by a list of MeSH terms used in the PubMed articles of the author in question to describe her/his expertise areas.

### 3.2 Metrics

Unfortunately, large data collections such as *e.g.* LinkedIn, the DBLP library, PubMed etc. contain a substantial proportion of noisy data and the achieved degree of accuracy cannot be estimated in a reliable way. Accuracy is most commonly measured by precision and recall. Precision is the ratio of true positives, *i.e.* true experts in the total number of true experts in a given domain. However, determining the total number of true experts in a given domain is not feasible.

In the current work, we use *resemblance*  $r$  and *containment*  $c$  to compare the expertise retrieval solutions generated on a given set of experts by using the expertise retrieval techniques discussed in Section 2.5.

Let us consider two expertise retrieval solutions  $S = \{S_1, S_2, \dots, S_k\}$  and  $S' = \{S'_1, S'_2, \dots, S'_k\}$  of the same set of experts. Then the similarity between two expertise retrieval results  $S'_i$  and  $S_i$ , which are constructed for the same example expert, can be assessed by *resemblance*  $r$ :

$$r(S'_i, S_i) = |S'_i \cap S_i| / |S'_i \cup S_i|, \quad (2)$$

where  $S_i$  and  $S'_i$ ,  $i = 1, 2, \dots, k$ , are corresponding expertise retrieval results. The first solution  $S$  is generated on the considered data set without taking into account the expert levels of expertise in different topics while the second one  $S'$  is a solution built by applying the proposed weighting method.

We also use *containment*  $c$  that assesses how  $S'_i$  is a subset of  $S_i$ :

$$c(S'_i) = |S'_i \cap S_i| / |S'_i| \quad (3)$$

The values of  $r$  and  $c$  are in the interval  $[0, 1]$ .

We also use *Silhouette Index* (SI) to evaluate the quality of the cluster solution generated by the FCA-based approach considered in Section 2.4. Silhouette Index is a cluster validity index that is used to judge the quality of any clustering solution  $C = \{C_1, C_2, \dots, C_k\}$  of the considered data set, which contains the attribute vectors of  $m$  objects. Then the SI is defined as

$$s(C) = 1/m \sum_{i=1}^m (b_i - a_i) / \max\{a_i, b_i\}, \quad (4)$$

where  $a_i$  represents the average distance of objects  $i$  to the other objects of the cluster to which the object is assigned,

and  $b_i$  represents the minimum of the average distances of object  $i$  to object of the other clusters. The value of SI from -1 to 1 and higher value indicates better clustering results.

### 3.3 Implementation and availability

Publications originating from Bulgaria have been downloaded in XML format from the Entrez Programming Utilities (E-utilities). The E-utilities are the public API to the NCBI Entrez system and allow access to all Entrez databases including PubMed, PMC, Gene Nucleotide and Protein. The E-utilities use a fixed URL syntax that translates a standard set of input parameters into the values necessary for various NCBI software components to search for therefore the structured interface to the Entrez system, which currently includes 38 databases covering a variety of biomedical data, including biomedical literature. To access these data, a piece of software first makes an API call to E-Utilities server, then retrieves the results of this posting, after which it processes the data as required. Thus the software can use any computer language that can send a URL to the E-utilities server and interpret the XML response.

For calculation of semantic similarities between MeSH headings, we use MeSHSim which is an R package. It also supports querying the hierarchy information of a MeSH heading and information of a given document including title, abstraction and MeSH headings [36].

In our experiments, we have applied the DTW-based algorithm to resolve the problem with ambiguity. For this purpose, we have used a Python library *cdtw*. It proposes a DTW algorithm for spoken word recognition which is experimentally shown to be superior over other algorithms [37].

### 3.4 Results and discussion

We have extracted a set of 4343 Bulgarian authors from the PubMed repository. After resolving the problem with ambiguity the set is reduced to one containing only 3753 different researchers. Then each author is represented by two components: a list of all different MeSH headings used to describe the major topics of her/his PubMed articles and a vector of weights expressing the relative levels of expertise the author has in the different MeSH terms composing her/his profile. The weight of a MeSH term that is presented in a particular author profile is the ratio of repetitions, *i.e.* the repetitions of the MeSH term in the total number of MeSH terms collected for the author. This weighting technique could additionally be refined by considering the MeSH terms annotating the recent publications of the authors as more important (*i.e.* assigning higher weights) than those met in the old ones. This idea is not implemented in the current experiments.

Experts	MeSH headings
1	Kidney Transplantation; Liver Transplantation
2	Health Behavior
3	Drinking; Health Behavior; Health Knowledge, Attitudes, Practice; Program Evaluation
4	Models, Biological; Temperature; Models, Neurological; Water
5	Computer Simulation; Models, Molecular; Protons; Thermodynamics; Molecular Conformation
6	Vibration; Models, Molecular; Infrared Rays; Hydrogen Bonding
7	Monte Carlo Method; Models, Theoretical; Phase Transition; Thermodynamics
8	Photosynthesis; Quantum Theory
9	Health Behavior; Decision Support Techniques; ... (more than 20 MeSH terms)
10	Polymorphism, Genetic

**Table 1** Expert MeSH heading profiles.

Experts	MeSH heading weights
1	0.5; 0.5
2	1
3	0.25; 0.25; 0.25; 0.25
4	0.166; 0.333; 0.166; 0.333
5	0.285; 0.285; 0.142; 0.142; 0.142
6	0.5; 0.166; 0.166; 0.166
7	0.428; 0.285; 0.142; 0.142
8	0.75; 0.25
9	0.022; ...; 0.045; ...; 0.068; ...; 0.25
10	1


**Table 2** Expert MeSH heading weights.

Examples of 10 expert MeSH heading profiles can be seen in Table 1. The corresponding weight vectors calculated as it was explained above can be found in Table 2.

We build expertise spheres of the ten example experts whose profiles are given in Table 1. Initially, we construct the expertise spheres of these authors by applying the weighting method discussed in Section 2.2. Respectively, the expertise spheres of the same authors without taking into account the intensity of their expertise in the different MeSH topics containing in their profiles are also produced. Next the resemblance  $r$  and the containment  $c$  are used to compare the two expertise retrieval solutions generated on the set of extracted Bulgarian PubMed authors for the example expert profiles.

Figure 1 depicts  $r$  and  $c$  scores which have been calculated on the expertise retrieval results produced for the example experts by identifying for each expert profile a fixed number (50) of expert profiles that are most similar to the

given one. As one can notice the obtained results are quite logical. Namely, the returned expertise retrieval results are identical ( $r = 1$  and  $c = 1$ ) when the experts have equally distributed expertise in the different MeSH headings presented in their profiles (e.g., see experts: 1, 2, 3 and 10). However, in the other cases (see experts: 4, 5, 6, 7 and 8) the resemblance between the corresponding expertise retrieval results is not very high (maximum 0.4). Evidently, the produced expertise retrieval results can be significantly changed by using a weighting method for assessing expert expertise. The latter is also supported by the results generated for the containment  $c$ .



**Figure 1**  $r$  and  $c$  scores calculated on the expertise retrieval results that are generated for the example experts given in Table 1 by selecting for each expert profile a fixed number of the most similar expert profiles.

The MeSH headings are grouped into 16 broad categories. We have produced a grouping of all the extracted authors with respect to these subject categories by applying the formal concept analysis clustering approach explained in Section 2.4. In this experiment, we have assumed that each author is described by only a list of MeSH terms. Next each author is further represented by a 16-length vector of membership degrees (contributions) of the expert to the different categories. The membership degree to the category is calculated as a ratio between the number of MeSH headings from the author profile that belong to the category and the total number of headings in her/his profile. The calculated membership degrees are turned into binary values by using a preliminary determined threshold. In the considered context the threshold is set to be equal to the median of all different membership degrees. Thus for each subject category we have associated those authors who have membership degrees greater than the determined threshold. Ten of the authors have not been assigned to any category, since there are no membership degrees above the calculated threshold in their profiles. Then a formal context presented by a  $3753 \times 16$  matrix, with the authors corresponding to the rows and the subject categories corresponding to the columns is built. Finally, a formal concept lattice for the built context is generated by using a data mining prototype Lattice Miner. It produces a lattice of 234 concepts. The non-empty concepts are 198, where each one represents a

subset of authors who belong to a number of subject categories. Thus the extracted Bulgarian health science experts are partitioned into 198 disjoint expert areas with respect to the main MeSH categories. 2166 researchers have been partitioned among 14 singleton concepts, 10 authors belong to the empty concept and the other 1587 researchers demonstrate multiple expertise. The number of authors partitioned into the main MeSH categories (singleton concepts) are given in Table 3.

Category label	Category name	Number of authors
A	Anatomy	45
B	Organisms	101
C	Diseases	68
D	Chemicals and Drugs	158
E	Analytical, Diagnostic and Therapeutic Techniques and Equipment	663
F	Psychiatry and Psychology	97
G	Phenomena and Processes	797
H	Disciplines and Occupations	38
I	Antropology, Education, Socialogy and Social Phenomena	14
J	Tehnology, Industry, Arguculture	20
K	Humanities	2
L	Information Science	37
M	Named Groups	1
N	Health Care	125

**Table 3** Number of authors partitioned into the main MeSH categories (singleton concepts).

Evidently, the produced grouping of experts well capture the expertise distribution in the considered domain with respect to the main subjects. In addition, it facilitates the identification of individuals with the required competence. For instance, if we need to recruit researchers who have expertise simultaneously in 'Phenomena and Processes' and 'Health care' categories, we can directly locate those who belong to the concept that unites the corresponding categories ({G, N}). Selected non-singleton concepts are given in Table 4. Most of these concepts unite two categories, *i.e.* the corresponding authors are active in two scientific areas. Logically the number of authors who have expertise in more than two subject categories is not very high.

It is difficult to evaluate the obtained expert partitioning as there are no benchmark ones available. Therefore, we have conducted another experiment in [33]. It performs the semantic-aware expert clustering algorithm, proposed in [34], with our test data. Initially, the constructed expert profiles represented by the 16-length vectors of membership degrees are used to calculate the Euclidean distance between each pair of vectors. Then the authors are clustered by using  $k$ -means clustering algorithm.

However, in order to determine the optimal number of clusters for the considered set of experts we have initially applied  $k$ -means clustering algorithm for different values of  $k$  and then we have evaluated the obtained clustering solutions by SI. In comparison to the partitioning algorithms as  $k$ -means the FCA-based approach does not need prior knowledge about the optimal number of clusters in order to produce a good clustering solution. Notice that the SI score generated on the expert clustering produced by the proposed approach is 0.698.

United categories	Number of authors
{G, N}	106
{E, N}	55
{C, G}	59
{E, L}	36
{F, N}	23
{F, I}	12
{E, G, N}	56
{E, H, J, L}	8
{G, H, L, N}	6
{E, G, I, L, N}	11
{F, G, H, I, N}	7

**Table 4** Number of authors partitioned into united MeSH categories (non-singleton concepts)

#### 4 SUMMARY

In this paper, we have discussed enhanced data-driven techniques for expert representation and identification. The proposed techniques have been tested and evaluated on data extracted from PubMed repository.

For future work, we aim to pursue further evaluation, validation and refinement of the discussed expert mining techniques on richer data coming from different application areas, subject fields and online sources, *e.g.* such as LinkedIn, Google Scholar, the DBLP library, Microsoft Academic Search, etc.

#### 5 REFERENCES

- [1] Autonomy Technology White Paper (<http://www.autonomy.com>)
- [2] Foner, L. "Yenta: a multi-agent referral system for matchmaking system", Proceedings of the First International Conference on Autonomous Agents, Marina Del Ray, CA, 1997.
- [3] Tacit Knowledge Systems' KnowledgeMail (<http://www.tacit.com>)
- [4] Kautz, H., Selman, B., Shah., M., "Referral Web: combining social networks and collaborative filtering" in *Communications of the ACM*, Vol. 40, Issue 3, pp. 63-65, 1997.

- [5] Ha-Thuc, V., Venkataraman, G., Rodriguez, M., Sinha, S., Sundaram, S., Guo, L. "Personalized Expertise Search at LinkedIn", 2016.
- [6] <https://maj.io/#/>
- [7] <http://yagajobs.co.uk>
- [8] Seid, D., Kobsa, A. "Demail: A hybrid architecture for expertise modelling and recommender systems". 2000.
- [9] Campbell, C.S., "Expertise identification using Bibliography 189 email communications", 12th Int. Conf. on Inform. and Knowl. Manag. ACM Press. 2003.
- [10] D'Amore, R. "Expertise community detection", 27th Annual Int. ACM SIGIR Conf. on Research and Development in Information Retrieval. ACM Press. 2004.
- [11] Mockus, A., Herbsleb, J.D. "Expertise browser: a quantitative approach to identifying expertise", 24th Int. Conf. on Software Engineering. ACM Press. 2002.
- [12] Hawking, D. "Challenges in enterprise search", 15th Australasian Database Conference. Australian Computer Society, Inc. 2004.
- [13] Tshiporkova, E., Tourwé, T. "Tool support for technology scouting using online sources" in *Springer* pp. 371–376. 2011.
- [14] Singh, H. "Developing a Biomedical Expert Finding System Using Medical Subject Headings" in *Healthcare Informatics Research* Vol. 19, Issue 4, pp. 243-249. 2013.
- [15] Hristoskova, A. "A Graph-based Disambiguation Approach for Construction of an Expert Repository from Public Online Sources", 5th IEEE Int. Conf. on Agents and Artificial Intelligence. 2013.
- [16] Abramowicz, W. "Semantically Enabled Experts Finding System - Ontologies, Reasoning Approach and Web Interface Design" in *ADBI* Vol. 2, pp. 157-166. 2011.
- [17] Bozzon, A. "Choosing the Right Crowd: Expert Finding in Social Networks", *EDBT/ICDT'13*. Genoa, Italy. 2013.
- [18] Craswell, N. "Overview of the TREC-2005 Enterprise Track", 14th Text Retrieval Conference. 2006.
- [19] Balog, K. "People search in the enterprise". PhD thesis, Amsterdam University. 2008.
- [20] Toutanova, K. "Enriching the knowledge sources used in a maximum entropy part of speech tagger", the Joint SIGDAT Conference on Empirical Methods in NLP and Very Large Corpora. EMNLP/VLC-2000. 2000.
- [21] Buelens, S., Putman, M., "Identifying experts through a framework for knowledge extraction from public online sources". Master thesis, Gent University, Belgium, 2011.
- [22] Boeva, V., Krusheva, M., Tshiporkova, E. "Measuring Expertise Similarity in Expert Networks", Proceedings of the 6<sup>th</sup> IEEE Int. Conf. on Intelligent Systems, pp. 53-57, 2012.
- [23] Boeva, V., et al. "Data-driven Techniques for Expert Finding", ICAART 9<sup>th</sup> International Conference on Agents and Artificial Intelligence, pp. 535-542, Porto, 2017.
- [24] Stankoff, D., Kruskal, J. "Time warps, string edits, and macromolecules: the theory and practice of sequence comparison", Addison Wesley Reading Mass. 1983.
- [25] Sakoe, H. and Chiba, S. "Dynamic programming algorithm optimization for spoken word recognition". In *IEEE Trans. On Acoust, Speech, and Signal Proc.*, ASSP-26, pp. 43-49, 1978.
- [26] Cameron, D, L. "SEMEF: A Taxonomy-based Discovery of Experts, Expertise and Collaboration Networks". MS thesis, The University of Georgia. 2007.
- [27] Hecht, F. "The Journal Impact Factor: A Misnamed, Misleading, Misused Measure" in *Cancer GenetCytogenet*, Vol. 4, pp. 77-81, Elsevier Science Inc. 1998.
- [28] Seglen, P. O. "Why the impact factor of journals should not be used for evaluating research" in *BMJ*. Vol. 314, Issue 7079, pp. 497. 1997.
- [29] Hirsch, J. E. "An index to quantify an individual's scientific research output" in *PNAS* Vol. 102, Issue 46, pp. 16569-16572. 2005.
- [30] Afzal, M.T., Maurer, H. "Expertise Recommender System for Scientific Community" in *Journal of Universal Computer Science* Vol. 17, Issue 11, pp. 1529-1549. 2011.
- [31] Fellbaum, C., "WordNet: An Electronic Lexical Database". MIT Press, Cambridge. 2001.
- [32] Miller, G. A. "WordNet: A lexical Database for English" in *Communications of the ACM* Vol. 38, Issue 11, pp. 39-41. 1995.
- [33] Boeva, V. et al. "Measuring Expertise Similarity in Expert Networks", In 6<sup>th</sup> IEEE Int. Conf. on Intelligent Systems, IS 2012 IEEE Sofia Bulgaria, pp. 53-57. 2012.
- [34] Boeva, V. et al. "Semantic-aware Expert Partitioning" Artificial Intelligence: Methodology, Systems, and Applications in *LNAI. Springer* Int. Pub. Switzerland. 2014.
- [35] B. Ganter, B., Stumme, G. and Wille, R. Formal Concept Analysis: Foundations and Applications, LNAI, no. 3626, Springer-Verlag, 2005.
- [36] Zhou, J., Shui, Y. "The MeSHSim package".
- [37] Paliwal, K.K. et al. "A modification over Sakoe and Chiba's Dynamic Time Warping Algorithm for Isolated Word Recognition" in *Signal Proc* Vol. 4, pp. 329-333. 1983.

# Multi-Task Representation Learning

Mohamed-Rafik Bouguelia      Sepideh Pashami      Sławomir Nowaczyk\*

## Abstract

The majority of existing machine learning algorithms assume that training examples are already represented with sufficiently good features, in practice ones that are designed manually. This traditional way of preprocessing the data is not only tedious and time consuming, but also not sufficient to capture all the different aspects of the available information. With big data phenomenon, this issue is only going to grow, as the data is rarely collected and analyzed with a specific purpose in mind, and more often re-used for solving different problems. Moreover, the expert knowledge about the problem which allows them to come up with good representations does not necessarily generalize to other tasks. Therefore, much focus has been put on designing methods that can automatically learn features or representations of the data instead of learning from handcrafted features. However, a lot of this work used ad hoc methods and the theoretical understanding in this area is lacking.

## 1 Motivation

Representation learning is concerned with automatically transforming raw input data into representations or features that can be effectively exploited in machine learning tasks. Existing unsupervised approaches to representation learning such as [1, 2, 3, 4, 5] yield general features capturing dimensions of variation that may or may not be essential to a given task. On the other hand, supervised approaches to representation learning such as [6, 7, 8, 9, 10] can be overly specific as they allow to exclusively learn representations that help to discriminate among class labels related to a specific task. Nevertheless, such approaches have, especially recently, been extensively studied in the deep learning community [7, 11]. In this case, however, the learned

representations cannot be directly applied to another task as they are explicitly tailored for a specific task.

The motivation for this research comes from our previous work on mapping of raw sensor data from Volvo trucks into low-dimensional representation, both in a supervised and unsupervised manner. Such a representation is needed for predictive maintenance solutions, as using the original raw data is not feasible. The overall goal is to extract general features which are suitable for more than one task, for example, estimating remaining useful life of several different components. Since those components can be related to different aspects of the truck operation, the representations that allow accurate predictions are related, but not necessarily the same. Achieving sufficient generality of the resulting features is not possible given current state of the knowledge in the field; more in-depth study of the underlying problem is needed before practical solutions can be developed.

We will contribute with extending the current representation learning methodology along two separate but interdependent directions. The first direction is considering training setup in which not one, but rather multiple related tasks are provided. This idea allows for a well-defined formalization of concepts such as complexity, diversity or incongruity among tasks to which the learned representation is expected to be applied to in the future. The second direction is aiming for a diverse set of representations, with clear and well-defined purposes and motives, instead of a single, all-encompassing one. Imposing such a meaningful structure onto the result allows for incremental generation and evaluation, as well as for explicit tradeoff between accuracy, generality and compression provided by the learned representation.

On the one hand, semi-supervised representation learning methods, like the one proposed in [12], aim to learn a representation based on few labeled data. However, the labels are still related to a single task and the result does not necessarily generalize well to multiple related ones. In parallel to representation

---

\*Authors are with Center for Applied Intelligent Systems Research (CAISR), Halmstad University, Sweden. Emails: mohbou@hh.se, seppas@hh.se and slanow@hh.se

learning, there has recently been considerable progress on the important problem of multi-task learning, which exploits similarities across several learning problems [13, 14]. The results show often significantly improve performance compared to learning each task independently. In the past few years different studies, such as [15, 16, 17, 18, 19], have advocated that a representation which captures properties that are invariant across tasks can significantly improve the performance. This generated an increasing interest in learning representations from multiple tasks, perhaps most noticeable in the computer vision domain. Nonetheless, despite the empirical success, formal justifications of why this works remain largely unexplored.

We formulate such multi-task representation learning problem as an extension of the classical supervised machine learning problem. The latter states that, given a set of training examples sampled according to an unknown underlying distribution  $\mathcal{D}$ , labelled by an unknown function  $\mathcal{T}$ , the goal is to find a hypothesis  $\mathcal{H}$  that minimizes the probability of  $\mathcal{H}$  and  $\mathcal{T}$  disagreeing. The proposed extension states that, in addition to underlying distribution of data examples, now  $\mathcal{D}_E$ , there also exists an underlying unknown distribution over tasks,  $\mathcal{D}_T$ . A number of training tasks, sampled from this distribution, provide the (set of) labelling for the training data. The goal is to produce a group of representations with the lowest error on unseen data across all the expected tasks. A measure of the capacity (or expressive power) of a function family, such as the VC dimension [20], can be generalized to capture the essential complexity across multiple tasks. This way measures for approximating the true error, on unseen data across all the expected tasks, can be modeled depending on the a priori assumptions about the inherent difficulty of the particular problem instance. Transformations of raw input into features can be based on capturing different dimensions of variation in data, and either being essential for a particular class of tasks, or providing broad benefits by being generally useful for many different tasks. Starting from the establishment of theoretical foundations for this problem, the end-goal is to create an algorithm that efficiently produces such representations.

Representation learning has clearly demonstrated early success with deep learning in application areas such as computer vision, natural language processing and speech recognition. However, replicating those results in other domains has proven difficult, in part due to lack of sufficient theoretical foundations. It

is well understood that the performance of machine learning methods is heavily dependent on the choice of the data representation (or features) on which they are applied. Unlike traditional feature engineering which requires labor-intensive effort, representation learning allows computers to autonomously create specific features which are appropriate for a particular problem. A good representation of data can also provide a substitute for storing raw data when dealing with big data in real-world applications. This research direction will have major impacts in various domains, as it enables building systems and algorithms that learn to perform new tasks based on experience gained from previous tasks. It will have major impact in a large number of application domains where machine learning is a key aspect; these includes text mining, patient healthcare data analysis, social network analysis, multiple object classification in computer vision, and predictive maintenance in the automotive industry, to mention but a few. We believe that it is a promising line of research, making progress towards real Artificial Intelligence.

## 2 Survey of the field and open challenges

Representation learning is challenging primarily for three general reasons: the immense space of possible solutions that should be considered; the difficulty in establishing a clear measurable objective for the learning process; and the insufficient understanding of how the properties of the problem instance match against the parameters of the representation learning process. This combination makes it hard to design efficient algorithms for determining which representation will ultimately be relevant for the expected distribution of tasks, as well as to propose a compelling theoretical foundation for such work.

The first of those main challenges is related to the generation of representations. Necessarily, for a given problem instance, a set of representations needs to be considered; such family of representations is usually generated by the same algorithm and corresponds to a certain family of functions. An important question for generating representations is that it is not clear whether there exists a single family of representations that is sufficient for all problems, or should multiple families be used, depending on the problem instance. Some comparative studies such as [21, 22, 23, 24]

have been carried out for various application domains. However, so far, the properties that can influence the choice of one family of representations over another are still unknown [25]. Another aspect closely related to the generation of representations is about encouraging diversity among the set of representations. For example, it took quite some time after Breiman's 2001 paper [26] before modern ways to measure the diversity of trees within random forest were suggested. A similar development is needed for representation learning. A challenging aspect in this context is to balance specificity and diversity in an optimal way that leads to an improved accuracy. Promoting diversity among representations is very important, yet, it is not a well-studied aspect and there are no explicit metrics in the literature that allow capturing the diversity among a set of representations. To the best of our knowledge, the only directly relevant paper in this context is [27], which proposes a strategy to produce an ensemble of diverse representations specifically for the unsupervised case. This is done by controlling the trade-off between minimizing reconstruction error and maximizing diversity between reconstructions. A similar method has been applied in [28] for the problem of fall detection. However, both of those results are specific for autoencoders only and not directly applicable in a multi-task representation learning context.

The second main challenge is related to the evaluation of representations. One aspect that makes it different from most machine learning problems (such as classification), is the difficulty in establishing a clearly defined objective. The standard way of performing this evaluation is to measure the representation or feature learning algorithm in terms of its usefulness with respect to a particular task [23, 24]. This is typically done at regular intervals (to enable early stopping), by evaluating the performance of a cheap classifier trained using the learned features. However, a first issue is that regularly alternating between learning features and training a classifier produces a substantial computational overhead. That raises the important question of how to balance between extraction and classification stages. One principled solution to this problem is to use a tree of classifiers as proposed in [29]. In this solution, test inputs traverse along individual paths, where each path extracts different features for inputs that benefit from them the most. Similar methods include those proposed in [30, 31]. However, not only is this problem NP-hard, but most importantly, such methods give an incomplete evaluation of the features. In particular,

the extension to the case of multi-task representation learning is challenging. As indicated in [25], these issues strongly motivate the use of unsupervised evaluation measures. For example, for auto-encoders [32, 33, 34], the reconstruction error on the test data can readily be used as an evaluation measure. However, such a measure can be unreliable because systems that learn more features as the time goes tend to overfit and systematically produce a lower test reconstruction error. Besides the accuracy, either with respect to a particular task or to a group of tasks, however, one can imagine a number of other criteria for assessing the quality and usefulness of a representation. For example, data compression has been studied in [35, 36], with the aim to transform the data into a compact but expressive form. Finally, the generality of representations across tasks is an important aspect that needs to be taken into account in multi-task representation learning; however, there is not yet any formal definition of this property. In particular, the aspect of balancing all those important metrics, both when generating and evaluating a multi-task representation, is usually not considered.

Finally, the third main challenge is understanding properties of the problem instance from the perspective of representation learning. There exist very few theoretical studies which expose problem properties that are relevant in the context of a multi-task representation learning. For example, the recent work [15] establishes theoretical results about the benefit of learning a representation from multiple tasks compared to learning from each task separately, based on basic properties such as the sample size, the data dimensionality, and the number of tasks. However, this work is only demonstrated for the specific case of subspace learning (i.e., linear feature learning), and does not take into consideration other relevant properties, such as the complexity of tasks or the similarity between tasks. Most existing multi-task representation learning methods such as [16, 17, 18, 19] assume that an expert can determine which tasks are related, or they implicitly assume that the available tasks are related and can be readily used to perform a joint training. However, in real-world, this assumption may not be always satisfied. Learning common representations across many unrelated or dissimilar tasks can lead to poor representations that decrease the performance compared to learning representations from each task separately, as discussed in [37, 38]. Therefore, similarity among tasks is an important problem property that needs to be taken into consideration.

### 3 Research questions

The challenges discussed above are very broad, therefore, we propose the following five main research questions as a good starting point.

1. *What problem instance properties require the generation of multiple diverse representations over generating a single multi-dimensional representation?*

Usual representation learning methods assume that the goal is learning a single low-dimensional representation of the data. However, in many cases it can be more beneficial to create a set of several independent representations. A set of multiple representations can always be seen as one higher-dimensional representation; yet, having multiple representations provides an explicit structure that can be exploited in various ways. In particular, if a set of representations is expected to be useful across a wide spectrum of tasks, such a structure offers a convenient way of expressing the trade-off between accuracy on each task and the diversity of representations within this set. In this context, the questions that need answering are when to learn multiple representations, how many of them should be created, and how to promote the diversity within the set of representations on several different levels (e.g., on the algorithm, the mapping and the data levels).

2. *How to generate a family of representations which ensures a good coverage of the space of mapping functions that are appropriate for a given problem instance?*

In practice, the representations are necessarily generated according to some algorithm, which induces a particular family of representations. In order to establish the theoretical foundations for the multi-task representation learning problem, it is crucial that one can measure whether this family is expressive enough to provide sufficient coverage, appropriate for the given problem instance. The way of connecting the expressiveness of the family with the complexity of the expected tasks it should be used on is going to be an important contribution. An essential research question is: can existing concepts such as VC dimension be extended to capture this match? One idea is to use different subsets of training tasks to achieve high diversity of representations and to get a reliable estimation of the expected coverage that can be achieved. However, the actual concrete method

for doing that needs to be developed. Further, how to generate a large set of non-redundant representations, from simple to more complex ones, which are able to approximate any and all of the expected future task?

3. *Does identifying and grouping related tasks, followed by learning multiple representations separately for each group, lead to improved outcome?*

Existing results [15] establish the benefits of learning a representation from multiple related tasks compared to learning from each task separately. At the same time, it has been shown that trying to learn a common representation across unrelated or dissimilar tasks can decrease the performance [37, 38]. Therefore, a natural question is whether it is possible to automatically find an appropriate partitioning of tasks that leads to learning better representations from tasks within each group? Under what conditions are such a step necessary? Establishing that will require new advancements in determining how the similarity among tasks should be measured, and which properties of the problem instance affect it to a different degree.

4. *How to evaluate representations in a multi-task setting based on several aspects such as the compactness of representations, the accuracy relative to each individual task, and the generality of across all tasks?*

Being able to evaluate representations is absolutely crucial, however, today in many cases it is done in an ad hoc manner. There is a need to define measures based on a well justified theoretical description of the problem at hand. In this context, an important question is, how can one establish conditions of whether a representation or set of representations is “good enough” for any given set of tasks? In particular, a good representation can be defined as a one which leads to a low prediction error (i.e., high accuracy) over the whole population of data and expected tasks (according to the unknown underlying distribution). Hence, given reasonable assumptions, which evaluation measures can be proven to be good approximations of the true error in multi-task representation learning? Moreover, as the data compression plays a non-negligible role in the context of representation learning, it is important to select a reasonably small number of representations that are common across tasks, while improving the accuracy relative to learning each task



independently. Therefore, another question relates to the number (or ratio) of representations that are “sufficient” to improve the accuracy, across a set of tasks. More generally, how can we quantify the expected benefit of representation learning from multiple tasks? Finally, how to balance all these different aspects of the evaluation: the data compression (or compactness of representations), the accuracy relative to each task, and the generality of representations across tasks?

*5. How to define the complexity of tasks (in addition to other problem instance properties) to address all the above questions in a principled way?*

In order to answer all the above research questions in a principled way, problem instance properties that potentially influence the choice of algorithms, representation families and quality measures need to be defined and formalized. The success of a set of representations for a given problem instance is related to the difficulty or complexity of tasks that one needs to deal with. In this context, the complexity of tasks is the most important property, which leads to an important question: what would be a good measure for the complexity of tasks? What is the equivalent of VC dimension for a family of tasks? Necessarily, such a measure needs to consider the similarity between tasks. Even though each task within a training set can be simple, if these tasks are very different, one may need a quite diverse and expressive family of representations. On the contrary, a much simpler family of representations can be sufficient for a set of very individually difficult but overall similar tasks. Task complexity is of course only one, even if arguably the most important, property of the problem that needs to be studied. Other examples include the amount of noise, the size (and overlap) of data available for each task as well as the dimensionality and the heterogeneity of the data.

## 4 Methods, approaches and ideas

### Defining relevant properties of the problem.

The starting point is to define the most relevant properties that can be used to describe a problem instance, for example to investigate a measure for estimating the complexity of expected tasks based on the available training tasks. One possibility is generalizing the

VC dimension [20], which is related to the inherent complexity of a space, for a set of tasks. The goal is to capture how difficult are the tasks we are expecting to have to deal with in the future. Another is a measure for modeling the similarity between tasks, possibly modeled either based on a direct comparison between parameters learned from the different tasks, or based on how well the parameters learned from one task, perform other tasks. Based on such properties one can describe the problem instance, together with additional basic features such as the number of tasks, the dimensionality of the original data representation, the level of noise, and the data size per task. Those properties can lead to an upper bound on the performance of representations across tasks. For example, the PAC learning framework [39] (Probably Approximately Correct learning) enables mathematical analysis of machine learning which stipulates that with high probability, a learned hypothesis (e.g., a classification model) will have low generalization error for a given classification task. PAC learning does not take into consideration the possible existence of multiple related tasks, nor does it concern itself with data representations.

**Generating representations.** A strategy for generating adequate families of representations can be based on efficient methods for generating large sets of representations that are able to approximate or represent any function of certain properties. For example, the field of Functional Data Analysis primarily focuses on smooth functions, which is probably too broad for our needs. The coverage of the space of functions by the generated representations will be measured in order to ensure a trade-off between the “exploration of all possible representations” and the “exploitation of the best generated representations”. In order to produce more useful representations, new methods for determining which tasks are related and therefore can be automatically grouped together based on the similarity, are needed.

**Encouraging diversity among representations.** The generation of representations should be directed by an evaluation process which allows selecting, among all the possible representations, the ones that fulfill a range of assessment criteria, in particular, preservation of diversity among the generated representations and generality of the representations across tasks. First, on the algorithm level: if representations are created by sufficiently different algorithms, they are likely to be different. This can be done, for exam-

ple, by explicitly controlling the bias and variance for a family of algorithms. Second, on the mapping level: representations are functions from one feature space to another, and those functions can be compared based on their mathematical properties, as defined in either Hilbert or Banach spaces. Third, on the example level: one can measure how do the relative positions change for the data points in the training sets. All of these can be done in either supervised and unsupervised manner, or as a combination of both approaches.

**Algorithm for multi-task representation learning.** The final goal, clearly, is an efficient algorithm which benefits from the results of the other work packages. The algorithm takes as input a dataset and a set of tasks, and produces as output a set of representations that are expected to generalize well across unseen tasks.

## 5 Preliminary results

This idea builds on our previous work of evaluating several approaches for both supervised and unsupervised mapping of raw sensor data from Volvo trucks into low-dimensional representation. Such a representation is needed for predictive maintenance solution, as using the original raw data is not feasible. The overall goal is not to find the best low-dimensional representation tailored to a very specific task, but rather to identify the method for learning a widely applicable representation.

For example, general low-dimensional representations of the data are calculated to find various truck configuration. Data originates from 79974 unique Volvo trucks and is recorded during a full year. The data of a single truck is represented with a bivariate histogram, where the axes correspond to a pair of sensors: turbocharger speed vs boost pressure. Each task describes various truck configurations, e.g., engine, gearbox, country of operation or brand, while the bivariate histograms describe the usage of the truck. We have performed a comparison of techniques based on t-distributed stochastic neighbor embedding (t-SNE) and convolutional autoencoders (CAE) in a supervised fashion over 74 different 1-vs-Rest tasks using random forest. The results show that t-SNE is most effective for 2D and 3D, while CAE could be recommended for 10D representations. Fine-tuning of the results shows slight improvement using low-dimensional representation in comparing to the original data representation.

## 6 References

- [1] Coates, A., Lee, H., & Ng, A. Y. (2010). An analysis of single-layer networks in unsupervised feature learning. *Ann Arbor*, 1001(48109), 2.
- [2] Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- [3] Bengio, Y. (2012). Deep learning of representations for unsupervised and transfer learning. *ICML Unsupervised and Transfer Learning*, 27, 17-36.
- [4] Bengio, Y., Yao, L., Alain, G., & Vincent, P. (2013). Generalized denoising auto-encoders as generative models. In *Advances in Neural Information Processing Systems*.
- [5] Dosovitskiy, A., Springenberg, J. T., Riedmiller, M., & Brox, T. (2014). Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems*.
- [6] Zhuang, F., Cheng, X., Luo, P., Pan, S. J., & He, Q. (2015, July). Supervised Representation Learning: Transfer Learning with Deep Autoencoders. In *IJCAI*.
- [7] Bengio, Y. (2013, July). Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing* (pp. 1-37). Springer Berlin Heidelberg.
- [8] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097-1105).
- [9] Zou, F., Wang, Y., Yang, Y., Zhou, K., Chen, Y., & Song, J. (2015). Supervised feature learning via l2-norm regularized logistic regression for 3d object recognition. *Neurocomputing*, 151, 603-611.
- [10] Fan, H., Cao, Z., Jiang, Y., Yin, Q., & Doudou, C. (2014). Learning deep face representation. *arXiv preprint arXiv:1403.2802*.
- [11] Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural networks*, 61, 85-117.
- [12] Banijamali, E., & Ghodsi, A. (2016). Semi-Supervised Representation Learning based on Probabilistic Labeling. *arXiv preprint arXiv:1605.03072*.
- [13] Evgeniou, T., Micchelli, C. A., & Pontil, M. (2005). Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6(Apr), 615-637.
- [14] Evgeniou, T., & Pontil, M. (2004, August). Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 109-117). ACM.
- [15] Maurer, A., Pontil, M., & Romera-Paredes, B. (2016). The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81), 1-32.
- [16] Gong, P., Zhou, J., Fan, W., & Ye, J. (2014, August). Efficient multi-task feature learning with calibration. In *Proceedings of the 20th ACM SIGKDD* (pp. 761-770).
- [17] Zhao, H., Stretcu, O., Negrinho, R., Smola, A., & Gordon, G. (2017). Efficient Multi-task Feature and Relationship Learning. *arXiv preprint arXiv:1702.04423*.
- [18] Argyriou, A., Evgeniou, T., & Pontil, M. (2008). Convex multi-task feature learning. *Machine Learning*, 73(3).
- [19] Argyriou, A., Evgeniou, T., & Pontil, M. (2007). Multi-task feature learning. *Advances in neural information processing systems*, 19, 41.

- [20] Vapnik, V. N., & Chervonenkis, A. Y. (2015). On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of Complexity* (pp. 11-30). Springer International Publishing.
- [21] Renshaw, D., Kamper, H., Jansen, A., & Goldwater, S. (2015). A comparison of neural network methods for unsupervised representation learning on the zero resource speech challenge. In *INTERSPEECH* (pp. 3199-3203).
- [22] Cruz-Roa, A., Arevalo, J., Basavanthally, A., Madabhushi, A., & Gonzalez, F. (2015, January). A comparative evaluation of supervised and unsupervised representation learning approaches for anaplastic medulloblastoma differentiation. In *Tenth International Symposium on Medical Information Processing and Analysis* (pp. 92870G-92870G). International Society for Optics and Photonics.
- [23] Tokarczyk, P., Montoya, J., & Schindler, K. (2012, July). An evaluation of feature learning methods for high resolution image classification. In *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences, 22nd ISPRS Congress, Melbourne, Australia*.
- [24] Shao, L., Cai, Z., Liu, L., & Lu, K. (2017). Performance evaluation of deep feature learning for RGB-D image/video classification. *Information Sciences*, 385, 266-283.
- [25] Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8), 1798-1828.
- [26] Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32.
- [27] Reeve, H. W., & Brown, G. (2015). Modular Autoencoders for Ensemble Feature Extraction. *NIPS 2015 Workshop on Feature Extraction: Modern Questions and Challenges. JMLR W&CP*, volume 44, 2015.
- [28] Khan, S. S., & Taati, B. (2016). Detecting Unseen Falls from Wearable Devices using Channel-wise Ensemble of Autoencoders. *arXiv preprint arXiv:1610.03761*.
- [29] Xu, Z. E., Kusner, M. J., Weinberger, K. Q., & Chen, M. (2013, June). Cost-Sensitive Tree of Classifiers. In *ICML (1)* (pp. 133-141).
- [30] Khan, S. H., Bennamoun, M., Sohel, F., & Togneri, R. (2015). Cost sensitive learning of deep feature representations from imbalanced data. *arXiv preprint arXiv:1508.03422*.
- [31] Xu, Z. E., Kusner, M. J., Huang, G., & Weinberger, K. Q. (2013). Anytime Representation Learning. In *ICML (3)* (pp. 1076-1084).
- [32] Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P. A. (2010). Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(Dec), 3371-3408.
- [33] Masci, J., Meier, U., Cirean, D., & Schmidhuber, J. (2011, June). Stacked convolutional auto-encoders for hierarchical feature extraction. In *International Conference on Artificial Neural Networks* (pp. 52-59).
- [34] Rifai, S., Vincent, P., Muller, X., Glorot, X., & Bengio, Y. (2011). Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*.
- [35] Gregor, K., & LeCun, Y. (2011). Learning representations by maximizing compression. *arXiv preprint arXiv:1108.1169*.
- [36] Gregor, K., Besse, F., Rezende, D. J., Danihelka, I., & Wierstra, D. (2016). Towards conceptual compression. In *Advances In Neural Information Processing Systems*.
- [37] Kang, Z., Grauman, K., & Sha, F. (2011). Learning with whom to share in multi-task feature learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)* (pp. 521-528).
- [38] Kumar, A., & Daume III, H. (2012). Learning task grouping and overlap in multi-task learning. *arXiv preprint arXiv:1206.6417*.
- [39] Long, P. M. (1995). On the sample complexity of PAC learning half-spaces against the uniform distribution. *IEEE Transactions on Neural Networks*, 6(6), 1556-1559.

## Improved Inter Terminal Transportation using Agent Technology

Lawrence Henesey<sup>1</sup>

<sup>1</sup>Blekinge Tekniska Högskola, Blekinge Institute of Technology, Karlshamn, Sweden lhe@bth.se

### Abstract

Many maritime logistics centres worldwide are experiencing large number of inter-terminal transportation volumes, which raises the complexity of transportation processes between the terminals. Different vehicle systems exist for transporting containers between different terminals, however they often are inefficient due to poor planning or scheduling. We present a solution for dynamic planning of resources by using an agent based simulation tool. The results showed improved resource planning and utilization of different resources in the network of terminals. A cost comparison of different vehicles systems is further analysed in order to identify the best choice of vehicle system for a given scenario.

### Keywords

Inter-terminal transportation, Container terminals, Automated guided vehicle, Agent based simulation.

## 1 INTRODUCTION

Maritime logistics centres, such as large ports, often consist of a number of container terminals (CTs). We investigated different road transport modalities for container transportation between the terminals of Massvlakte area of port of Rotterdam. The Inter-Terminal Transportation (ITT) volume and other processes of transportation have changed since the recent extension of Massvlakte, e.g., Massvlakte II (MV2) will increase ITT volume nearly 5 million TUE/year.

The aim of our work is to provide decision support to ITT planners, who are responsible for scheduling and planning resources in handling containers by providing alternatives on choice of resources employed. In ITT various equipment is used (vehicles, cranes) for improved planning and estimation of ITT, using present truck and MTS systems, and automated vehicle system, which includes automated guided vehicles (AGVs). For this purpose, an agent based simulation model of the ITT system at the Maasvlakte has been designed and developed. According to [1] processes of ITT can be summarized in the following.

- Punctual (neither too early nor late) pickup of containers from source
- Punctual delivery of containers at destination
- Possible bridging of discrepancies in both these tasks by providing buffer areas

According to our best knowledge, limited research has been conducted that focuses on ITT. We found most studies focused on the movement of containers inside container terminals; called intra terminal transportation. On the other hand, none of the studies used multi agents to solve ITT problems except Albert Douma but he only

considers barge handling problems in his studies [2-5] This research space provides us an opportunity to propose an agent based solution for ITT planning, and evaluate it by comparing different transportation vehicles.

## 2 REVIEW OF LITERATURE

An extensive review of literature for container transportation and classification of research published between 2007 and 2014 was conducted. According to Drewry, a maritime research organization, [6] over 90% of the world's general cargo is handled through containers. The main modalities between CTs are road, rail, short-sea, deep-sea, and inland waterway [7]. The importance of improving container terminal (CT) operations and demands for transporting cargo in containers are increasing in parallel. The CTs must be able to effectively and efficiently act as an integral part of transport chain from origin to destination [8].

A comprehensive literature review on terminal operations and their classification was presented by Dirk Steenken et. al. [9] and was further updated in [10]. It has been observed that most research mainly focused on specific part(s) of the terminal system. For instance yard cranes using mathematical modelling [11], yard operations [12], train load planning problem in terminals [13] and berth scheduling [14] are some examples of this type of research. On other hand, the use of simulation is very popular among researchers when previous data is unavailable or when they have to communicate to non-technical staff. In [1], [7] [15], [16], [17], [18] [19], [20], and [21] simulation is adopted to evaluate the container terminals at an aggregate level. Many studies concerning container terminals focus on CT transport using an agent

based approach, for example [22], [16], [18] and [8]. Researchers also have used agent based modelling to model different type of transportation modes. For example, Automated Guided Vehicles (AGVs) and Rubber Tired Gantry (RTGs) are modelled as agents to evaluate their performance in a CT by Hoshino et al [23]. Agent based simulation in which a number of actors are involved in the CT are simulated with set of design theory by Henesey et al in ([24].

In artificial intelligence research, use of distributed control mechanism has increased since the introduction of multi-agent systems. Multi-agent systems serve as platform for distributed planning or distributed control. Multi-agent systems allow us to split a complex problem into multiple sub-problems to simplify the problems and achieve the overall goal. Multi-agent based applications can be found in different fields such as logistics, economics and computer science [2]. Furthermore the efficient management of a container terminal using multi-agent model seems to be an adequate framework for dealing with the design and development of an application [25]. Furthermore the efficiency of CT operations can be improved by considering multi-agent model in a seaport ([26], [20]). Both Albert Douma and Henesey et al. have modelled barge handling problem and overall terminal operations using multi agents respectively [2], [3], [22] and [27].

Multi-terminal environment introduced movement of containers between terminals of the same container port. It introduces new challenges for container transportation planners and researchers. According to our best knowledge, not many researches focused on ITT till 2006. However, after 2006 we can find some studies focusing this area of transportation. Very first study focusing on ITT is published in 1996 [20] in which simulation concepts were introduced. These concepts were further updated in [7]. In later work, simulation of multi-terminal container port, considering port of Rotterdam, was presented by the same research group [7]. A comparison of three transport systems of ITT, including AGVs, Automated Lifting Vehicle (ALV), and MTS, was then published in the same year considering same area of port of Rotterdam as a case [28]. Albert Douma presented his PhD. thesis considering one part of ITT, i.e. barge, in 2008 and introduced new concepts for barge handling problem and used multi agents to model this problem [2]. Other publications, in following years, by Douma et al. [3-5] were also used multi agents focusing on the barge-handling problem at port of Rotterdam. In 2013, Sterzikn, and Kopfer [29] presented a tabu search based heuristic algorithm for inland container transportation focusing on vehicle routing, scheduling, and optimization of number of vehicles. Most recent study focusing on ITT was presented in 2014 by [30]. It presents a mathematical model for analysing inter-terminal transportation. The focus was to minimize container delivery delay

considering key components of ITT, including traffic congestion, multiple vehicle types and loading/unloading times, and arbitrary terminal configurations.

Most of the work published on the ITT domain is written by two research groups, that includes both Ottjes et al. [7] and [20] and Douma et al. [2-5] Both research groups considered port of Rotterdam as a case. However, latest work presented by Ottjes et al., [7] published in 2007, which requires to be updated because of recent changes in port of Rotterdam, while work published by [2-5] focused only on barge handling problem at the same container port. Among these studies research presented by [7] is most relevant to our work. However, this paper differs as we are aiming to apply multi agents to model ITT considering different vehicles system. Our work also differs from research published by Tierney et al. [30] as they use mathematical modelling to address different ITT problems and focus more on a strategic time horizon as we focus more on operational time span, i.e. real-time.

### 3 AGENT TECHNOLOGY

Agent technology has gained a lot of interest by variety of disciplines in computer science since its introduction. According to Wooldridge and Jennings [31], an agent is usually defined as a hardware or software based object with key properties, autonomy, social ability, reactivity, and pro-activeness. Autonomous behaviour of an agent identifies that an agent can operate without direct external intervention. Social ability ensures that an agent can interact with other agents or actors of its environment. It also perceives the information from the environment and reactive property ensures its response to the environment. Its ability of pro-activeness makes it able to take initiatives and exhibit a goal-oriented behaviour. MAS can be defined as a loosely coupled network of actors working together to solve problem or a combination of problems that are beyond the individual capabilities or [32]). Additionally, literature also endorses our point as number of researchers have applied agent technology in transshipment of containers like [2-5], [22], [24], [16], [33] and [34].

To define the agents we follow the suggestions given in [35] and decompose our problem to ensure that agents represent entities in the physical world. We have identified two types of actors involved in ITT operations; container terminal, and transportation vehicles. Transportation vehicles include barge, train, and road vehicle. Road vehicles have three different types of vehicles including truck, AGV, and MTS. Therefore; six (Terminal agent, Barge agent, Truck agent, AGV agent, MTS agent, and Train agent) different agents have been defined in our ITT system. All agents are categorized into two different categories including, terminal agents and transport agents according to their behaviour. All transport agents communicate with terminal agents by

messages passing and event triggering. All agents of the system will have a certain degree of opportunistic behaviour so that they collaborate without hurting each other to get the main goal of the system. Following is the description of all the agents identified in the system.

A terminal agent represents the container terminal of physical world. We only consider those operations that are relevant to ITT. The terminal agent can communicate with all transport agents to accomplish different sub-goals.

Transport Agents under this category are mapping of physical vehicles that are being used for transportation in this ITT model. These vehicles include road vehicles (Truck, AGV, and MTS), trains, and water vehicles (barge). All transport agents can communicate with the terminal agents. They also negotiate with terminal agent for getting orders for transshipment

The manuscript must fit within the required margins. Do not use fonts without serifs like Arial, Helvetica etc. in the text area.

Although we have used barges and trains on our ITT model, we mainly focus on three types of road vehicles during evaluation of our ITT model.

1. Truck: It is a non-automated vehicles used for container transportation with maximum capacity of 2 TEU (Twenty foot Equivalent Unit; box).
2. MTS: A single vehicle that is manned by one driver pulling a number of trailers (with a capacity of 10 TEU), called a multi transport system (MTS).
3. AGV: It is an automated system of vehicle used for container transportation. We consider double stacking AGVs shown with capacity of 4 TEU.

#### 4 SIMULATION METHOD

According to our best knowledge there are three ways to develop a simulation model; code each and every thing by your own, use development toolkit, and use a full simulation tool. In [36] they believe that using Agent Based Simulation (ABS) software for agent simulation modelling is better than using toolkits, due to availability of complete modelling functionalities, professional technical support, user-friendly interfaces, and simplified modelling process in simulation software packages. There are number of free, educational, and/or commercial simulation tools available to facilitate ABS modelling.

We identified a simulation tool called Netlogo that was been used on container terminal problems [37]. In parallel we conducted thin review of some other simulation software packages and selected Anylogic. Both NetLogo and AnyLogic were then compared and analysed to select best suitable tool for our simulation model. In short summary of this comparison in given in Table 1.

**NetLogo.** NetLogo [37] is a multi-agent programmable modelling environment developed by the Centre for Connected Learning & Computer-Based Modelling at the Northwestern University. It is an improved version of StarLogo [37], which was developed by media lab at Massachusetts Institute of Technology (MIT).

**AnyLogic.** AnyLogic (“AnyLogic,” 2014) is a simulation package developed by XJ Technologies Company. It is based on the Eclipse framework, and ensures cross-platform development environment. Programming environment is purely Java oriented in this software package.

	AnyLogic	NetLogo
Programming Language	Java	Scripting
Programming Paradigm	Object Oriented	Procedural
3D animation	Yes	Yes
Drag and Drop Components	Yes	NO
Data Analysis	Yes (Powerful)	Yes (Limited)
Data Import/Export	Easy and several methods	Limited (Text Files Only)
Model Export	Java Applet	Java Applet
Developer Guide	Yes	Yes
Help	Limited Training videos, Paid training sessions	Big online Social media Community

**Table 1.** Comparison of AnyLogic and NetLogo

Table 1 shows that AnyLogic provides Java based Object Oriented programming (OOP) environment opposing NetLogo, which has procedural scripting environment. Having a good background knowledge and experience with Java technologies, and to avoid putting time on learning NetLogo scripting, we chose AnyLogic 6.9 to implement our ABS model. Furthermore, easy data import and export feature was also convincing factor to choose AnyLogic, as we had to import multiple data sets for different scenarios.

#### 4.1 Simulation Method

The simulation model is currently, considering 10 terminals of MV area in this case study. We believe the simulation model is flexible enough and can be used with data of any network consisting several terminals.

We set number of input parameters before running simulations. In addition to loading respective datasets, input parameters determine resources of container terminals, number of vehicles being used to handle

container flow, and capabilities of vehicles. Following we explain three types of parameters shown in input form.

1. Operational Hours. We use three sets (8, 16, and 24 hours) of operational hours in our experiments.
2. Save Results and Enable Validation. These parameters are Boolean (true/false) type of variables. Enable validation is used obtain data on console that is further used for validation and verification. Parameter 'save results' used to save the simulation results in the excel sheet attached with the simulation model.
3. Scenario Selection. Volume distribution changes with respect to selected scenario. This parameter helps the model to use the relevant datasets during simulation experiment. It is important to note that as distances between the terminals do not vary on the basis of selected scenario, so the same datasets for distances are always loaded into simulation for every selected scenario.

Terminal parameters are used to setup capabilities of terminals in terms of tree types of crane resources and service times to handle containers against three transport modalities. All terminals are equipped with gate resources; barges can be served on all terminals except 1, and 7 terminals are capable to handle trains except 3. Following, we define three types of resource capabilities.

1. Gate Capabilities. Includes two parameters; 'Number of vehicles served in parallel', and 'Service time per container'. Prior concerns gate points at terminal which means how many gate cranes are available at terminal to server road vehicles while later defines loading/unloading service time for a container. As all terminals have gate points, these parameters will always be set to ensure the availability of gate resources.
2. Barge Capabilities. Similar to gate resources, barge-handling resources are set by same type of parameters. All terminals serve barges except one terminal, VDCD.
3. Train Handling Capabilities. Trains often leave the container bogies at terminals, those can be unloaded/loaded afterwards in relax hours, and ties themselves with other containers bogies. Therefore, we have change over time for the trains at each terminal that has train terminal. On the other hand "number of trains served" refers to number of trains that can be facilitated for change over at the same time. Three container terminals in MV area, including KDD, VDCD, and DCS, do not train terminals and are connected through train lines.

Three main scenarios that are based on ITT volume varying from 2.1 million to 4.9 million TEU per year are identified. Distribution of volume was obtained from interviews with experts. Simulation time interval is 1 week for scenarios and sub scenarios.

A fix number of barges and trains in the ITT network are simulated. A variable number of road vehicles are employed (AGV, Truck, MTS), presented in Table 2. Only one type of road vehicle can be used in a simulation run. Vehicle settings are presented in Table 3.

Vehicles	Maximum Capacity (TEU)	Average Speed (m/s)	# of Vehicles
AGV	4 (Double Stacking)	6	Variable
MTS	10	6	Variable
Truck	2	6	Variable
Train	70	20	6
Barge	50	3	6

Table 2 Vehicle Setting

Another variable, operations hours, also have three different settings (8, 16, 24 hours) that are used with each set of simulation scenarios. Therefore, 9 sub-scenarios are generated for each three main scenarios. Weekly based simulations will have 56, 112, and 168 maximum acceptable operations hours and three daily operational hour settings. Summary of simulation scenarios is presented in Table 4.

Terminal Name	Gate		Barge		Train Terminal	
	Cranes	Loading/unloading Time per Container (Minutes)	Cranes	Loading/unloading Time per Container (Minutes)	Capacity in terms of Trains	Train Change over Time (Minutes)
RWG	20	2	2	2	3	45
APMTII	8	2	2	2	3	45
ETR	25	2	2	2	3	45
RCT	7	2	2	2	3	45
APMTR	5	2	2	2	3	45
DCS	5	2	2	2	0	NA
ECTDT	5	2	2	2	3	45
ECT-BFT	20	2	2	2	3	45
VDCD	10	2	0	NA	0	NA
KDD	20	2	2	2	0	NA

Table 3 Terminal Setting

Volume/ Week in TUE (Scenario 1 = 40385, Scenario 2 = 65385, Scenario 3 = 94231) # of Trains = 6, # of Barges = 6					
Road Vehicle	Weekly Operations Hours Threshold	Sub-Scenarios for Scenario 1	Sub-Scenarios for Scenario 2	Sub-Scenarios for Scenario 3	Total
C A G V	56	3	3	3	9
	112				
	168				
M T S	56	3	3	3	9
	112				
	168				
T R U C K	56	3	3	3	9
	112				
	168				
Total Scenarios		9	9	9	27

**Table 4** Summary of all scenarios used for experiments

## 5 SIMULATION RESULTS

### 5.1 Scenario1 tested

This scenario has least value for ITT container volume among all three scenarios. Total ITT container volume in this scenario is 2.1 million TEU per year, which gives 40385 TEU per week. Results obtained from simulation experiment for all nine sub-scenarios of this main scenario are grouped with respect to operations hours and analysed in the following sections. We run 20 simulation iterations for several combinations of number of trucks in the scenarios. In Table 5 we show the container volume, average distance covered by every type of vehicles, and average time for handling the given volume of containers. Vehicle combination used in third column "Scenario 1-A (3)" of the table depicts acceptable handling time. We investigate operations hours further in order to find the best-fit value for number of trucks to handle the given scenario combining with barges and trains.

Scenario 1-A (1)			
	Trucks	Trains	Barges
# of Vehicles	360	6	6
Total Volume in TEU	29955	5025	5294
Average Distance Covered by Vehicle (Km):	454.08	148.00	201.86
Average Travel Time (Hr):	48.46	2.04	54.80
Average Service Time (Hr):	2.80	36.18	14.72
Operation Time (Hours):	58.06	58.06	58.06
Scenario 1-A (2)			
# of Vehicles	370	6	6
Total Volume in TEU	30199	4899	5175
Average Distance Covered by Vehicle (Km):	444.83	144.30	203.50
Average Travel Time (Hr):	47.70	1.99	54.76
Average Service Time (Hr):	2.75	35.34	14.39
Operation Time (Hours):	57.35	57.35	57.35
Scenario 1-A (3)			
# of Vehicles	380	6	6
Total Volume in TEU	30238	4874	5162
Average Distance Covered by Vehicle (Km):	436.25	133.20	195.17
Average Travel Time (Hr):	46.48	1.84	52.73
Average Service Time (Hr):	2.69	32.44	14.35
Operation Time (Hours):	58.06	58.06	58.06

**Table 5** Results for Scenario 1-A

### 5.2 Scenario 2 tested

Same steps were followed for "Scenario 1-B" to get best possible solution for the same container volume while using automated guided vehicles.

Results show that overall performance of AGVs is better in terms of number of vehicles as compared to trucks. Results for three combinations that are near most to the solution set are presented in Table 6. Less than half number of AGVs was required to handle the same container volume comparing with number of trucks. Reason for this difference could be two-fold; one double stacking capability of AGVs, second result shows that less average service time was required for loading/unloading container from AGVs as compared to trucks. A best configuration with respect to operations hours is 170 AGVs with 6 trains and 6 barges, where operations time is shown just near to the threshold value.



Scenario 1-B(1)			
	AGV	Trains	Barges
# of Vehicles	150	6	6
Total Volume in TEU	27380	6299	6595
Average Distance Covered by Vehicle (Km):	503.79	178.83	260.49
Average Travel Time (Hr):	53.7	2.47	48.04
Average Service Time (Hr):	1.73	36.08	18.33
Operation Time (Hours):	58.04	58.04	58.04
Scenario 1-B (2)			
# of Vehicles	150	6	6
Total Volume in TEU	27380	6299	6595
Average Distance Covered by Vehicle (Km):	503.79	178.83	260.49
Average Travel Time (Hr):	53.7	2.47	48.04
Average Service Time (Hr):	1.73	36.08	18.33
Operation Time (Hours):	58.04	58.04	58.04
Scenario 1-B(3)			
# of Vehicles	170	6	6
Total Volume in TEU	28896	5435	5943
Average Distance Covered by Vehicle (Km):	466.1	175.75	250.5
Average Travel Time (Hr):	48.87	2.43	44.32
Average Service Time (Hr):	1.61	35.44	16.52
Operation Time (Hours):	55.94	55.94	55.94

Table 6 Results for Scenario 1-B

To better compare in Figure 1 the results from the various equipment types in the transport is presented.




Figure1 Number of Vehicles required

To validate our model we tested it with different scenarios, such as setting 1 vehicle in the network and one crane at each terminal of the network. Preliminary results for the scenario are presented for expert feedback. Suggestions and feedback was noted and model was update accordingly.

Additionally, we did computations on how much time a vehicle should take to cover a given distance with a specific speed of vehicle. The computation results were noted. In the model, logs for time taken to travel from one point to another with specific speed of vehicle were collected and compared with the values computed before. Numbers of test scenario were run with different vehicles and speed of vehicles. Results are cross-verified with computation results for each scenario. We fixed the problems if identified at any point and tested the scenarios again until the satisfactory and expected results are not found.

We have presented and analysed the simulation results in for all possible scenarios defined. We found the most suitable number vehicles for each scenario within the range of maximum valid operations hours. At the end we present the utilization of the gate resource for different terminals in the network.

In conclusion, we found number of required vehicles decreases with increasing operations hours for given parameters of container volume. We also observed that trains and barges perform better by increasing the operations time per day in our model. We can assume it is because of trains have higher change over time and they cannot visit all the terminals within less operation hours. Delayed volume was introduced in the results, which creates more demand for road vehicles to deliver orders with in the given time. It needs more investigation and requires in depth analysis to find the additional reasons, if any.

## 6 SUMMARY

### 6.1 General appearance

Future of containerization indicates more and more transportation between the terminals due to increased demand of containerized transport. Larger terminals are being built resulting in multi-terminals in a single container port. Punctual delivery of consignments is the key of better planning of supply management. The aim of this paper was to provide an effective solution for ITT to meet the future challenges of containerization world.

Though optimization approaches exist, the use of simulation provides decision makers a faster and more intuitive approach to understanding a complex problem. With the loads in terminals varying immensely, the level of uncertainty increases, which motivates the use of simulation for planning. A simulation model was then evaluated through simulation experiments and results are presented and tested against different combination of vehicles in order to locate the best suitable number transport resource for punctual delivery of given container volume. Utilization of terminal resources is then compared while using several settings of gate

cranes. These findings will help ITT planners to do better planning for ITT and related resource in the future to avoid delay in container deliveries.

In conclusion, meeting the goal of this paper we believe our findings and estimations of manned and unmanned transportation vehicles and cranes used in ITT process will help the ITT planners in better planning and estimation for ITT in MV area of port of Rotterdam.

Some pointers for the future we have noted are that research is increasing in the ITT research area. Therefore, there are definitely more opportunities to work in the field of ITT. In future, other transportation modes can be further investigated to evaluate the model and explore its applications. A detailed and through investigation can be performed to improve the performance of barges and trains.

Another improvement in the model can be introduced by integrating it with intra-terminal operations of different terminals of the network to analyse effects of different settings on trough put of individual container terminals as well as whole network of terminals. Automation seems to be a future of container terminals around the world as research paradigm is shifting towards automation. We have evaluated our model using only one type of automated vehicle, i.e. AGV

## 7 REFERENCES

- [1] Duinkerken, M. B., Dekker, R., Kurstjens, S. T. G. L., Ottjes, J. A., & Dellaert, N. P. 2007. Comparing transportation systems for inter-terminal transport at the Maasvlakte container terminals. In P. K. H. Kim & P. D. H.-O. Günther (Eds.), *Container Terminals and Cargo Systems* (pp. 37–61).
- [2] Douma, A. M. 2008. Aligning the operations of barges and terminals through distributed planning. University of Twente.
- [3] Douma, A. M., Schuur, P. C., & Schutten, J. M. J. 2011. Aligning barge and terminal operations using service-time profiles. *Flexible Services and Manufacturing Journal*, 23(4), 385–421.
- [4] Douma, A., Schutten, M., & Schuur, P. 2009. Waiting profiles: An efficient protocol for enabling distributed planning of container barge rotations along terminals in the port of Rotterdam. *Transportation Research Part C: Emerging Technologies*, 17(2), 133–148.
- [5] Douma, A., Schuur, P., & Jagerman, R. 2011. Degrees of terminal cooperativeness and the efficiency of the barge handling process. *Expert Systems with Applications*, 38(4), 3580–3589.
- [6] Drewry. 2012. *Container terminal capacity and performance benchmarks*. London: Drewry Shipping Consultants Ltd ~ The Independent Maritime Adviser.
- [7] Ottjes, J. A., Veeke, H. P. M., Duinkerken, M. B., Rijsenbrij, J. C., & Lodewijks, G. 2007. Simulation of a multiterminal system for container handling. In P. K. H. Kim & P. D. H.-O. Günther (Eds.), *Container Terminals and Cargo Systems* (pp. 15–36). Springer Berlin Heidelberg.
- [8] Peng, Y., & Junqing, S. 2009. Agent Based Container Terminal Optimization. In *IITA International Conference on Control, Automation and Systems Engineering, 2009. CASE 2009* (pp. 607–609).
- [9] Voß, S., Stahlbock, R., & Steenken, D. 2004. Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26(1), 3–49.
- [10] Stahlbock, R., & Voß, S. 2008. Operations research at container terminals: a literature update. *OR Spectrum*, 30(1), 1–52.
- [11] Ng, W. C., & Mak, K. L. 2005. Yard crane scheduling in port container terminals. *Applied Mathematical Modelling*, 29(3), 263–276.
- [12] Kap Hwan Kim. 1997. Evaluation of the number of rehandles in container yards. *Computers & Industrial Engineering*, 32(4), 701–711.
- [13] Ambrosino, D., Bramardi, A., Pucciano, M., Saccone, S., & Siri, S. 2011. Modelling and solving the train load planning problem in seaport container terminals. In *2011 IEEE Conference on Automation Science and Engineering (CASE)* (pp. 208–213). K
- [14] Kim, K. H., & Moon, K. C. 2003. Berth scheduling by simulated annealing. *Transportation Research Part B: Methodological*, 37(6), 541–560. [
- [15] Henesey, L., Davidsson, P., & Persson, J. A. 2009. Agent based simulation architecture for evaluating operational policies in transshipping containers. *Autonomous Agents and Multi-Agent Systems*.
- [16] Huynh, J. M. V. and N. 2010. Building Agent-Based Models of Seaport Container Terminals [text]. Retrieved April 5, 2017 from <http://jmvida.cse.sc.edu/lib/vidal10a.html>
- [17] Kulaka, O., Polata, O., & Guenther, H.-O. 2008. Performance evaluation of container terminal operations. *IT Based Planning and Control of Seaport Container Terminals and Transport Systems*.
- [18] Li, B., & Li, W. 2010. Modelling and simulation of container terminal logistics systems using Harvard architecture and agent-based computing. In *Simulation Conference (WSC), Proceedings of the 2010 Winter* (pp. 3396–3410).
- [19] Liu, C. I., Jula, H., & Ioannou, P. A. 2001. A simulation approach for performance evaluation of proposed automated container terminals. In *2001 IEEE Intelligent Transportation Systems, 2001. Proceedings* (pp. 563–568).

- [20] Ottjes, J. A., Duinkerken, M. B., Evers, J. J. M., & Dekker, R. 1996. Robotised Inter Terminal Transport of Containers. In *Proceedings 8th European Simulation Symposium. Genua [SCS]* (pp. 621–625).
- [21] Yun, W. Y., & Choi, Y. S. 1999. A simulation model for container-terminal operation analysis using an object-oriented approach. *International Journal of Production Economics*, 59(1–3), 221–230.
- [22] Henesey, L. E. 2006. *Multi-Agent Systems for Container Terminal Management*. Blekinge Institute of Technology. PhD Thesis.
- [23] Hoshino, S., Ota, J., Shinozaki, A. ., & Hashimoto, H. . 2005. Highly efficient AGV transportation system management using agent cooperation and container storage planning. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005*.
- [24] Henesey, L. E., Notteboom, T. E., & Davidsson, P. 2003. Agent-based simulation of stakeholders relations: An approach to sustainable port and terminal management. In *International Association of Maritime Economists Annual Conference, 2003* pp. 314–331.
- [25] Rebollo, M., Julian, V., Carrascosa, C., & Botti, V. 2000. *A Multi-Agent System for the Automation of a Port Container Terminal*. Autonomous Agents 2000 workshop on Agents in Industry.
- [26] Van Dam, K. H., Verwater-Lukszo, Z., Ottjes, J. A., & Lodewijks, G. 2006. Distributed intelligence in autonomous multi-vehicle systems. *International Journal of Critical Infrastructures*, 2(2), 261–272.
- [27] Henesey, L., Davidsson, P., & Persson, J. A. 2009. Agent based simulation architecture for evaluating operational policies in transshipping containers. *Autonomous Agents and Multi-Agent Systems*, 18(2), 220–238.
- [28] Duinkerken, I. M. B., Ottjes, J. A., Evers, J. J. M., Kurstjens, S. T. G. L., Dekker, R., Dellaert, N. P., Cpm, D. 1996. Simulation Studies on Inter Terminal Transport at the Maasvlakte. In *Simulation Studies on Inter Terminal Transport at the Maasvlakte*. In Proceeding of 2<sup>nd</sup> Trail PhD Congress 1996 “Defense or attack”. May 1996. Rotterdam (TRAIL). ISBN 90-5584-020-3.
- [29] Sterzik, S., & Kopfer, H. 2013. A Tabu Search Heuristic for the Inland Container Transportation Problem. *Computers & Operations Research*, 40(4), 953–962.
- [30] Tierney, K., Voß, S., & Stahlbock, R. 2014. A mathematical model of inter-terminal transportation. *European Journal of Operational Research*, 235(2), 448–460.
- [31] Wooldridge, M., & Jennings, N. R. 1995. Intelligent Agents: Theory and Practice. *Knowledge Engineering Review*, 10, 115–152.
- [32] Durfee, E. H., & Lesser, V. R. 1989. Negotiating task decomposition and allocation using partial global planning. In M. Huhns (Ed.), *Distributed Artificial Intelligence* (Vol. 2, pp. 229–243). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [33] Sha, M. 2008. A simulation model for intra-terminal transport of Container Terminal operations system. In *IEEE International Conference on Service Operations and Logistics, and Informatics, 2008. IEEE/SOLI 2008*. Vol. 2, pp. 2810–2814.
- [34] Sharif, O., & Huynh, N. 2013. Storage space allocation at marine container terminals using ant-based control. *Expert Systems with Applications*, 40(6), 2323–2330.
- [35] Shen, W., & Norrie, D. H. 1999. Agent-Based Systems for Intelligent Manufacturing: A State-of-the-Art Survey. *Knowledge and Information Systems*, 1(2), 129–156.
- [36] Zhou, Z., Chan, W. K. (Victor), & Chow, J. H. 2009. Agent-based simulation of electricity markets: a survey of tools. *Artificial Intelligence Review*, 28(4), 305–342.
- [37] StarLogo. (2006). Retrieved January 9, 2015, from <http://education.mit.edu/starlogo/>

## 8 ACKNOWLEDGEMENT

We acknowledge the support of Karlshamn Kommun and NetPort for their continued support for research funding.

# Energy Efficiency in Machine Learning: A position paper

Eva García-Martín, Niklas Lavesson, Håkan Grahn and Veselka Boeva  
 Department of Computer Science and Engineering  
 Blekinge Institute of Technology, 371 79, Karlskrona, Sweden  
 Email: {eva.garcia.martin, niklas.lavesson, hakan.grahn, veselka.boeva}@bth.se

## Abstract

Machine learning algorithms are usually evaluated and developed in terms of predictive performance. Since these types of algorithms often run on large-scale data centers, they account for a significant share of the energy consumed in many countries. This position paper argues for the reasons why developing energy efficient machine learning algorithms is of great importance.

## 1 Introduction

Machine learning algorithms are becoming more and more popular due to the availability of large volumes of data and the advancements in hardware that makes it possible to analyze these data. These algorithms are present in large-scale data centers, which account for 3% of the global energy consumption<sup>1</sup>. This energy consumption needs to be reduced due to health issues linked to environmental pollution [1]. There are two ways to address this challenge, either by researching on how to find new, clean, sources of energy that can provide enough energy for the population, or to reduce the actual energy consumption of our devices. We center on the second one: building sustainable and energy-efficient machine learning algorithms.

There is a lot of research conducted in machine learning focusing on improving the predictive performance of algorithms, but recently researchers are becoming more interested on improving energy efficiency as well. This paper argues for the reasons why developing energy efficient algorithms in machine learning is of great importance. We focus on three counterclaims to develop machine learning algorithms considering energy efficiency: i) reducing the energy consumption of machine learning algorithms does not necessarily lead to a reduction of the overall energy consumption, ii) time and energy are strongly correlated, thus being redundant to measure the energy consumption since time is already measured<sup>2</sup>, and iii) it is complicated to measure energy consumption, thus making it time consuming and impractical [2].

The rest of the paper is organized as follows: Section 2 provides the background with the terminology and concepts related to time, energy, and power, together with the related works. Section 3 explains in detail the counterclaims for studying energy efficiency in machine learning. We address those counterclaims in Sections 4, 5 and 6. Section 7 portrays some preliminary results of having measured the energy consumption of a particular algorithm in different datasets. Finally, Section 8 presents the conclusions.

<sup>1</sup><http://www.independent.co.uk/environment/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html>

<sup>2</sup><https://01.org/powertop/overview>

## 2 Background

### 2.1 Terminology

All the equations and notation are based on the work by Dubois et.al [3]. Our scope centers on showing how energy is consumed by an algorithm, presenting the relationship between energy, time, and the instructions of a program. These instructions are to be optimized towards an overall energy reduction.

Energy (joules) is the product of power (watts) and time (seconds),

$$energy = power \cdot time \quad (1)$$

energy being the amount of power consumed during a period of time. Power is a measurement of the rate at which energy is consumed. Since

$$dynamic \quad power = C \cdot V_{dd}^2 \cdot f \quad (2)$$

$V_{dd}$  being the voltage,  $C$  the capacitance and  $f$  the frequency, we observe that the relationship between time and energy is nonlinear. Lowering the frequency of the processor leads to longer executions, but the total energy consumption can be lower, since  $V_{dd}$  can be reduced at a lower frequency, thus lowering the power significantly and lowering the energy consumption.

The total execution time of a program (multicycle computer model) is the multiplication of the number of instructions (IC), the average number of clock cycles per instruction (CPI), and clock cycle time (TPC):

$$T = IC \cdot CPI \cdot TPC \quad (3)$$

The total energy consumed by a program is

$$E = IC \cdot CPI \cdot EPC \quad (4)$$

where EPC is the energy per clock, and it is defined by integrating the power over a period of time, thus,

$$EPC = \frac{1}{2} \cdot C \cdot V_{dd}^2 \quad (5)$$

$EPI = CPI \cdot EPC$ , where EPI is energy per instruction. CPI represents the number of clocks needed to execute an instruction. The key is that CPI is different for the types of instructions, thus, there are certain instructions that need more CPI

than others. For instance, a load instruction has a higher CPI than an ALU instruction. A way to optimize a program to consume less energy and to take less time is to use more instructions with a lower CPI and less instructions with a higher CPI.

## 2.2 Related Work

Energy consumption has been widely studied for many years in the computer engineering community, focusing on designing processors that consume very little energy using Dynamic Voltage Frequency Scaling (DVFS) and on some other power saving techniques [4]. Regarding software, green computing has been introduced as a field that studies ways to address software solutions from a green, sustainable, and energy efficient perspective. Many companies are also starting to be concerned with the energy consumption of their computations<sup>3</sup>. Based on these works, we can observe how there is a trend in building energy efficient software [5, 6, 7].

In relation to machine learning, a concern for building energy efficient algorithms is increasing in the community. For instance, in a panel discussion during the 2016 Knowledge Discovery and Data Mining conference held in San Francisco<sup>4</sup>, there was a very interesting discussion between key researchers in the field addressing which steps should be taken to make algorithms more energy efficient. They centered in the deep learning field, since these algorithms are particularly computationally expensive [8]. Autonomous cars is a field that is directly connected to machine learning and energy efficiency. Autonomous cars often rely on deep learning algorithms, and since most of these cars are powered by batteries, it is important to build accurate models that consume lower amounts of energy than what is currently the case. There is already a research group making deep neural networks energy efficient while maintaining the same levels of accuracy [9].

## 3 Counterclaims

There are several aspects to consider related to developing energy efficient machine learning algorithms. The scope of this position paper centers in three counterclaims: i) reducing the energy consumption of machine learning algorithms does not necessarily lead to a reduction of the overall energy consumption, ii) time and energy are strongly correlated, thus being redundant to measure the energy consumption since time is already measured, and iii) it is complicated to measure energy consumption, thus making it time consuming and impractical.

### Overall energy reduction

The first concern is that even if we spend tremendous efforts on both reducing the energy consumption of algorithms and studying machine learning from an energy efficient perspective, this does not necessarily lead to a significant reduction of energy. For instance, if the algorithms that we are studying are hardly used on large-scale platforms, optimizing their energy consumption will have no impact on the world's energy consumption. So the focus should be to reduce the energy in those algorithms that are being used daily, on large-scale data centers, and by different applications.

<sup>3</sup><https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>

<sup>4</sup><http://www.kdd.org/kdd2016/>

### Correlation between time and energy

The second counterclaim centers on the fact that in many cases energy and time can be so strongly correlated that measuring the energy consumption of algorithms is a waste of resources, since we already have that information from the execution time<sup>5</sup>. Adding the energy consumption variable to the algorithm design process might just create an overhead on the time to publish and release such an algorithm, since it is not so straightforward to measure energy consumption, which leads to the last counterclaim.

### Measuring energy consumption

Measuring energy consumption can be quite troublesome due to the unavailability of tools that can give an accurate approximation of the energy [2]. Some tools are based solely on statistical modeling which are based on the usage of the CPU. Again, giving energy results that are correlated with time. Based on this, it can be unfeasible to try to understand how energy is consumed in an algorithm.

## 4 Impact on overall energy consumption

For the counterclaim: *reducing the energy consumption of machine learning algorithms does not necessarily lead to a reduction of the overall energy consumption* to be true, one of the following two claims need to be satisfied: i) that machine learning algorithms are hardly used in data centers, and that most of the energy consumption of those centers is due to other reasons, ii) that modifying these algorithms does not output a significant difference in terms of the energy consumed.

We can observe how many companies are building artificial intelligence (AI) solutions based on machine learning algorithms in their applications. Facebook has an AI research lab (FAIR)<sup>6</sup> that is applying machine learning to enhance the user experience, with applications that go from face detection to search similar multimedia documents in a database. Google has several research divisions, such as Google Brain and Google Deep Mind, centered in machine learning. This pattern occurs in many other companies, Amazon, Yahoo, etc. Since machine learning solutions are present in these companies, which generate a high percentage of the total Internet traffic, we believe that machine learning algorithms are frequently being used in large-scale data centers.

Addressing the second point, we have already published some work that shows how energy consumption can vary depending on how an algorithm is programmed [10]. In Section 7 we show some results of running an algorithm in different datasets, and the difference in energy consumption.

With the increased interest that machine learning is having in the top companies which govern the Internet traffic, even a small reduction of the energy consumption of an algorithm will have a great impact on the overall energy consumption at a global scale.

<sup>5</sup><https://01.org/powertop/overview>

<sup>6</sup><https://research.fb.com/category/facebook-ai-research-fair/>

## 5 Correlation between energy and time

From Eq. 1 we can see that time and energy are directly connected through power. While time gives an overview of the efficiency of a computation, energy gives a more detailed overview into this matter. Measuring energy consumption can be a complement to measuring only time. If the execution of algorithm  $A$  is longer than the execution of algorithm  $B$ , that does not necessarily lead to a higher energy consumption by  $A$  compared to  $B$ . These algorithms could run in a cellphone, and the battery consumed by  $B$  could be higher. The reasons behind this is Eq. 2, and the main reasons behind DVFS. Reducing the frequency of the processor for a specific process can reduce  $V_{dd}$ , thus significantly reducing the power. While reducing the frequency might make the process run for longer, the total energy consumed could be less.

## 6 Measuring energy consumption

At the moment it is not straightforward how to measure the energy consumption of software. Some researchers have built energy models for specific algorithms [9] and others use statistical tools based on the CPU usage [11].

We proposed a methodology to measure the energy consumption at a fine-grained level of machine learning algorithms in a previous publication [12]. In such publication we used a tool based on statistical models. In this section we focus on describing how to measure the energy consumption with Sniper [13] for each function of the algorithm. We focus on analyzing the energy consumption of each function to understand where exactly is the energy hotspot of the algorithm.

Sniper<sup>7</sup> is an x86 simulator that can run an algorithm and, with the use of McPAT<sup>8</sup> (integrated in Sniper), output the energy consumed by the algorithm giving an overview of the instructions responsible for that. Some examples are given in Section 7.

In order to measure the energy consumption with sniper, we need to inject the code of the algorithm with markers around those functions or regions of interest where we would like a detailed energy measurement<sup>9</sup>. This is done by including `SimMarker()` calls around those functions. The algorithm is then compiled, and inserted in the sniper run, with the `roi` (region of interest) option activated, so that it takes into considerations the markers, saving statistics at every step. Once the execution is finished, we get an overview of the energy and power consumption by calling McPAT, using a script provided by sniper, and passing the different markers of the functions as parameters. In order to get the energy consumption of every function call, we need to have a different marker name for each call, otherwise we will only be able to extract the energy consumption from one function call.

## 7 Experiment

This section shows the energy consumption of three different runs of the same algorithm under three different datasets. The goal is to show how can energy vary depending on the type of execution, and the type of results obtained from measuring energy consumption with sniper.

<sup>7</sup>[http://snipersim.org/w/The\\_Sniper\\_Multi-Core\\_Simulator](http://snipersim.org/w/The_Sniper_Multi-Core_Simulator)

<sup>8</sup><http://www.hpl.hp.com/research/mcpat/>

<sup>9</sup>[http://snipersim.org/w/Multiple\\_regions\\_of\\_interest](http://snipersim.org/w/Multiple_regions_of_interest)




Figure 1: Energy consumption of the VFDT. Dataset = 1 million instances, 5 numeric and 5 nominal attributes.

### 7.1 Algorithm

The algorithm profiled is the Very Fast Decision Tree (VFDT) [14]. VFDT is a decision tree algorithm able to analyze data from a stream, in an online fashion, thus updating the model as the data arrives, not saving any data and reading it only once. It outputs competitive predictive performance results in comparison to standard offline decision tree algorithms, while being able to handle large amount of data. VFDT reads each instance once by once, and then updates the statistics of the attributes and the classes seen at each node. Once enough examples are seen at a specific node, if there is a clear attribute that has a higher information gain in comparison to the others, that attribute becomes a new node on the tree, and a split is made. This process is repeated until there is not more data in the stream.

### 7.2 Datasets




Figure 2: Energy consumption of the VFDT. Dataset = 100k instances, 50 numeric, 50 nominal attributes.

In order to get a general overview of the energy consumed by the VFDT, we run the algorithm under three different datasets. The first one had 1,000,000 instances, 5 numeric attributes and 5 nominal attributes. The second one had 100,000 instances,




Figure 3: Energy consumption of the VFDT. Dataset = 100 thousand instances, 50 numeric attributes.

50 numeric attributes and 50 nominal attributes. The third one had 100,000 instances and 50 numeric attributes.

The goal was to test how energy varied in datasets with different number of attributes and instances, to try to generalize the energy consumption behavior of the algorithm. Keeping track of numeric attributes can be very troublesome, usually needing more energy than for nominal attributes. VFDTc was an update made to VFDT that could handle numerical attributes [15].

### 7.3 Results

Table 1 shows the energy consumed by running VFDT on the three datasets. We observe that there is a significant difference between the energy consumed from the three datasets. Interestingly, although D2 and D3 contain less instances than D1, they consume more energy, since they have more attributes. From the results, we can see that having more attributes consumes more energy than having more instances. D2 has 10 times less instances and 10 times more attributes, and it consumes 1.67 times more energy. D3 shows how expensive it is to analyze numeric attributes. Since basically, the only difference from D2 and D3 are 50 nominal attributes, and the difference is 85.16 J. Adding 50 nominal attributes is only 85.16 J, while the rest is spent on reading the data and analyzing those numeric attributes. One possible way to optimize VFDT in terms of energy consumption is to see which method is best to handle numeric attributes [16].

Figures 1, 2, and 3 show the total energy consumed in those three executions, specifying which instructions are consuming that energy. We can observe that they follow similar patterns, and that most of the energy is consumed due to RAM and cache accesses.

## 8 Conclusion and Future Work

This position paper argues for the importance of developing energy efficient machine learning algorithms. We address three counterclaims: i) reducing the energy consumption of machine learning algorithms does not necessarily lead to a reduction of the overall energy consumption, ii) time and energy are strongly correlated, thus being redundant to measure the energy consumption since time is already measured, and iii) it

Table 1: Results from executing the VFDT on three setups, generated with the random tree synthetic generator [17] Nom=Nominal attributes, Num=Numeric attributes.

Setups	Instances	Nom	Num	Energy (J)
S1	1,000,000	5	5	391.48
S2	100,000	50	50	656.26
S3	100,000	0	50	571.10

is complicated to measure energy consumption, thus making it time consuming and impractical.

We argue that reducing the energy consumption of machine learning algorithms will have an impact on the overall energy consumption, since the top companies responsible for most of the Internet traffic are using these type of algorithms. Examples are: Facebook’s AI research lab, and Google Brain and Google Deep Mind. Moreover, although time and energy are related, measuring energy consumption can offer a unique overview on top of measuring time, since there are cases where the execution time of an algorithm can be longer but the energy consumption lower (for example when using DVFS). Finally, although measuring energy consumption can be complicated, there are some solutions that can model the energy consumption of different algorithms [8]. In Section 6 we presented an approach to measure the energy consumption at the function level of any algorithm using the x86 simulator Sniper. The study concludes with an experiment where we show the energy consumption of an online decision tree under three different datasets.

The planned future work is to investigate the key operations that can reduce the overall energy consumption of different algorithms, together with a generic understanding of the energy complexity of decision trees.

## Acknowledgments

This work is part of the research project Scalable resource-efficient systems for big data analytics funded by the Knowledge Foundation (grant: 20140032) in Sweden.

## References

- [1] M. Naghavi, H. Wang, R. Lozano, A. Davis, X. Liang, M. Zhou, S. E. V. Vollset, A. Abbasoglu Ozgoren, R. E. Norman, T. Vos, *et al.*, “Global, regional, and national agesex specific all-cause and cause-specific mortality for 240 causes of death: a systematic analysis for the global burden of disease study 2013,” *The Lancet*, vol. 385, no. 9963, pp. 117–171, 2015.
- [2] T. Johann, M. Dick, S. Naumann, and E. Kern, “How to measure energy-efficiency of software: Metrics and measurement results,” in *Proceedings of the First International Workshop on Green and Sustainable Software*, pp. 51–54, IEEE Press, 2012.
- [3] M. Dubois, M. Annaram, and P. Stenström, *Parallel computer organization and design*. Cambridge University Press, 2012.

- [4] C. Reams, *Modelling energy efficiency for computation*. PhD thesis, University of Cambridge, 2012.
- [5] A. Freire, C. Macdonald, N. Tonello, I. Ounis, and F. Casheda, "A self-adapting latency/power tradeoff model for replicated search engines," in *7th ACM international conference on Web search and data mining*, pp. 13–22, 2014.
- [6] A. Hooper, "Green computing," *Communication of the ACM*, vol. 51, no. 10, pp. 11–13, 2008.
- [7] S. Murugesan, "Harnessing green it: Principles and practices," *IT professional*, vol. 10, no. 1, pp. 24–33, 2008.
- [8] V. Sze, Y.-H. Chen, T.-J. Yang, and J. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *arXiv preprint arXiv:1703.09039*, 2017.
- [9] T.-J. Yang, Y.-H. Chen, and V. Sze, "Designing energy-efficient convolutional neural networks using energy-aware pruning," *arXiv preprint arXiv:1611.05128*, 2016.
- [10] E. García-Martín, N. Lavesson, and H. Grahn, "Energy efficiency analysis of the Very Fast Decision Tree algorithm," in *Trends in Social Network Analysis - Information Propagation, User Behavior Modelling, Forecasting, and Vulnerability Assessment. (To appear in April 2017)* (R. Missaoui, T. Abdesslem, and M. Latapy, eds.), 2016.
- [11] A. Noureddine, R. Rouvoy, and L. Seinturier, "Monitoring energy hotspots in software," *Automated Software Engineering*, vol. 22, no. 3, pp. 291–332, 2015.
- [12] E. García-Martín, N. Lavesson, and H. Grahn, "Identification of energy hotspots: A case study of the very fast decision tree," in *Green, Pervasive, and Cloud Computing. GPC 2017, LNCS 10232 (To appear in May 2017)* (M. A. et al, ed.), pp. 1–15, 2017.
- [13] T. E. Carlson, W. Heirman, S. Eyerman, I. Hur, and L. Eeckhout, "An evaluation of high-level mechanistic core models," *ACM Transactions on Architecture and Code Optimization (TACO)*, 2014.
- [14] P. Domingos and G. Hulten, "Mining high-speed data streams," in *Proceedings of the 6th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 71–80, 2000.
- [15] J. Gama, R. Rocha, and P. Medas, "Accurate decision trees for mining high-speed data streams," in *Proceedings of the ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 523–528, ACM, 2003.
- [16] R. B. Kirkby, *Improving hoeffding trees*. PhD thesis, The University of Waikato, 2007.
- [17] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer, "MOA: Massive online analysis," *The Journal of Machine Learning Research*, vol. 11, pp. 1601–1604, 2010.