

Personalised Health Monitoring by a Multiagent System

Leo van Moergestel, Brian van der Bijl,
Erik Puik, Daniël Telgen
Department of Computer science
HU Utrecht University of Applied Sciences
Utrecht, the Netherlands
Email: leo.vanmoergestel@hu.nl

John-Jules Meyer
Intelligent systems group
Utrecht University
Utrecht, the Netherlands
Alan Turing Institute Almere, The Netherlands
Email: J.J.C.Meyer@uu.nl

Abstract—By using agent technology, a versatile and modular monitoring system can be built. This paper describes how such a system can be implemented. The roles of agents in this multiagent system are described as well as their interactions. The system can be trained to detect several combinations of conditions and react accordingly. This training can be done with specific patient situations in mind, resulting in a personalised monitoring system. Because of the distributed nature of the system, the concept can be used in many situations, especially when combinations of different sensor inputs are used. Another advantage of the approach presented in this paper is the fact that every monitoring system can be adapted to specific situations by varying the number and types of sensors and the messaging capability. As a case-study, a health monitoring system will be presented.

Keywords—Multiagent-based health monitoring; Multiagent architecture; learning agent.

I. INTRODUCTION

The work in this paper is based on a paper presented at the Intelli 2016 conference [1] and other previous work. Monitoring systems are widely used in many situations. Simple systems collect information that can be inspected by humans or other systems. More advanced systems have the capability to react on the data monitored. Smoke detecting systems with an alarm are examples of these systems. Often a situation arises where more than one monitored condition should be taken into account before an action should be performed. Industrial production systems are examples of complicated situations where many sensors are used to control the process [2]. Another example of a complicated situation is the health condition of the human body [3]. Here, alarm conditions may also depend on individual factors, necessitating the monitoring system to be trained for the specific individual person.

This paper describes a modular agent-based system [4] that can be trained by a medical expert and can monitor the status of a person and react adequately on the conditions encountered. This system has been built using agent technology, resulting in a robust and versatile multiagent-based monitoring system. The concepts presented here can be used in other situations as well [5].

The rest of this paper is organised as follows: Section II will describe the concepts of our approach, the reason for choosing agent technology as well as the architecture of the multiagent system. The agent types introduced in the

architecture description as well as other technical aspects will be explained in more detail in subsequent sections, starting with Section III where the design of the sensor agent is described. The decision agent is the subject of Section IV. Communication and the communication agents are explained in Section V and Section VI. The section is followed by Section VII where the training system will be explained. This training aspect is an important aspect of the system and is treated in detail. The proof of concept and results are presented in Section VIII. Related work will be discussed in Section IX and a conclusion will end the paper.

II. AGENT-BASED MONITORING

The first part of this section will show the requirements and explain the use of agent technology, while the second part will focus on the multiagent architecture.

A. Requirements and technology

A monitoring system should be built to be capable to handle several input values in combination. Depending on the combined values of the inputs a specific action should be executed. The system should be trained to build a knowledge base and utilise known information to decide on a strategy to react to the current situation. This resulted in the following list of requirements:

- the system should monitor different inputs simultaneously;
- it should be easy to add extra monitoring inputs;
- the system should be trained in an effective manner;
- the system should have a set of possibilities to react on certain conditions;
- different types of reaction should be possible depending on the input values.

As a case study, a system in the medical domain has been adopted, but the concepts presented here can easily be used in other domains as well.

For the realisation, agent technology has been used [4]. The reasons for choosing agent technology are:

- Error resistance. By using separate agents, the failure of an agent responsible for sensor input will not bring down the whole monitoring system. There is now a possibility to fall back on a different solution based on the availability of sensor inputs.
- Clear separation of responsibilities and goals. In our design, the sensors will be tied to separate agents that have a clearly defined goal. This is also true for the other agents involved, as will be discussed in the next subsection.
- Modularity. A multiagent system (MAS) is modular by nature and can be easily expanded with new features and possibilities.

B. MAS design

The agents involved have roles and responsibilities. When the different responsibilities are taken into account this will result in the architecture of a multiagent system as depicted in Figure 1.

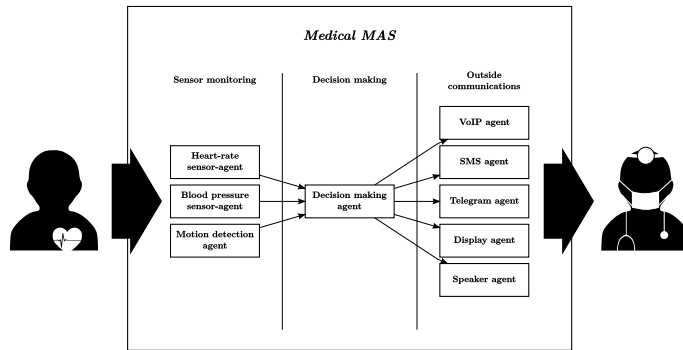


Figure 1. Medical MAS architecture

Three different roles are incorporated in the design resulting in three types of agents.

1) *monitoring agents*: Monitoring agents are responsible for delivering data to the decision agent. The data is coming from sensors. The agents themselves have a rather simple design. It could be possible to tell the agent at what intervals the data should be presented as well as the format expected by the decision agent.

2) *decision making agent*: A central role in the system is played by the decision agent. This agent decides what action should be performed under what conditions given by the monitoring agents. To do so, it has to be trained to build a knowledge base on how to react on certain conditions. This training has to be supervised by an expert, in our case a medical expert. A data acquisition system has been developed to help the expert to efficiently add data to the knowledge base. That system will be explained in the next section.

3) *communication agent*: The system has a set of communication agents that are responsible to communicate with the outside world. These agents are used by the decision agent to send emails, messages for several communication systems, like SMS, and also putting information on a display or generating an audible alarm.

Each MAS could contain any number of agents from the first and third categories (collectively known as utility agents), as well as one central agent having the capability of mapping observations from sensor agents to actions performed by communication agents. This design allows for agents to be added and removed dynamically while keeping core functionality intact. An example medical MAS is shown in Figure 1, including a number of potential utility agents.

C. Internal communication

The categories of agents described above will operate together in a single MAS. Though it would be possible to have agents running outside this MAS and still be able to communicate with the agent within, this aspect falls outside of the scope of this research.

The gold standard for agent development is the JADE developed by Telecom Italia. It provides a set of tools and an extendable framework for creating agents and MASs using the FIPA standard. Within JADE, agents exchange information using the ACL protocol [6] developed by FIPA. The prototype medical MAS has been developed in JADE and as such uses the ACL protocol for communicating information. The ACL protocol allows for a number of languages to encapsulate data, including XML; as the XML format provides a means to represent data in a semantic way that is readable to both humans and computers it appears to be a good choice for inter-agent communication. A typical exchange of messages is shown in Figure 2. Two sensor-agents send a stream of measurements to the decision agent, which periodically calculates its assessment of the situation. When the assessment reaches a certain defined threshold, it starts to send messages to the communication agents. The communication agents can respond with a confirmation or a message indicating either failure to relay the communication or failure to understand the message. If the decision-agent receives a message indicating failure, it tries to alert an operator using the designated fallback.

III. SENSOR AGENTS — BIOLOGICAL FACTORS AND MONITORING

Sensor-monitoring agents are the simplest and most diverse agents in the medical MAS. These agents exist primarily to support modularity: as sensors do not necessarily produce compatible signals to communicate their results, a small, dedicated agent could be programmed to read the sensor-data and communicate it to the decision-making agent in a standard format. This way, the decision-making agent does not need to know the way measurements are performed. A sensor could easily be exchanged by a different kind of sensor, together with its associated agent, without necessitating mayor changes to the decision-making agent. Similarly, new sensors could be added to facilitate patients with diseases requiring certain types of sensors. Though this would require new data to be added to the decision-making agent, the process of acquiring the measurements could easily be added by inserting a new sensor-monitoring agent to the MAS. This part of the system will be discussed in Section III.

In order for any system to be able to assess a patients health, it needs to know about a number of biological factors. For heart failure and related problems — the focus

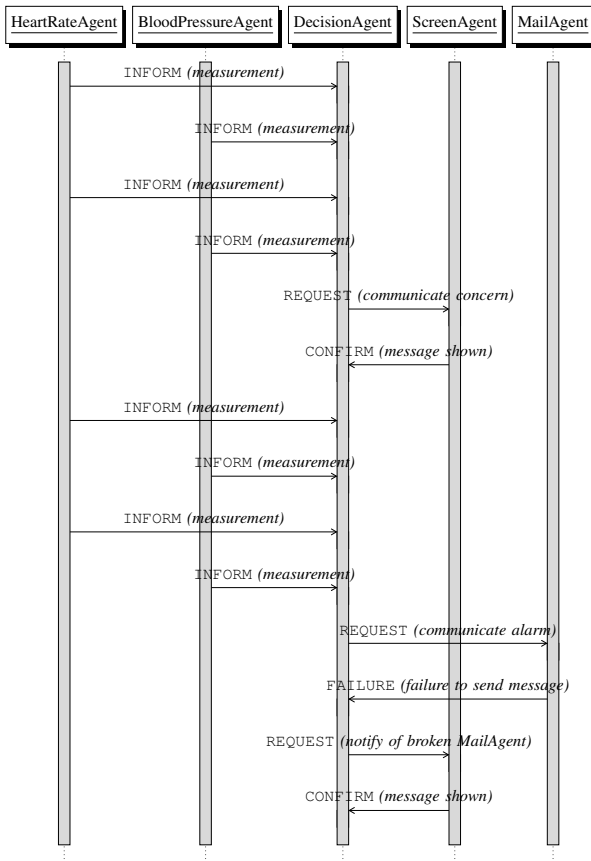


Figure 2. Message exchange within the medical MAS.

of this research, the patients heart rate and blood pressure seem to be obvious choices for factors to be monitored. For other diagnoses another set of variables might contain more meaningful information. However, from an artificial intelligence (AI) point of view, the provision of a decisive list of factors is not in the AI domain but in the medical domain. As such, this resulted in a design of a prototype that is as factor-agnostic as possible. Each factor to be considered is added during initialisation as a feature to a n -feature algorithm, and an appropriate sensor is added to the system to collect the necessary data. By scaling each feature to a value in the range $(-1, 1)$, as described in Section VII-H, the specifics of each feature are abstracted away from the reasoning process: the product does not need to know what a given value represents, only how it affects the output of its prediction-function.

Thus, for the system to work with a given set of variables, three things are needed to add the desired behaviour to the existing product:

- 1) a sensor capable of measuring the new feature,
- 2) a mapping between sensor-output and a value in the range $(-1, 1)$ and
- 3) a dataset for the prediction algorithm featuring the new factor.

To facilitate the development of sensor-agents, an abstract

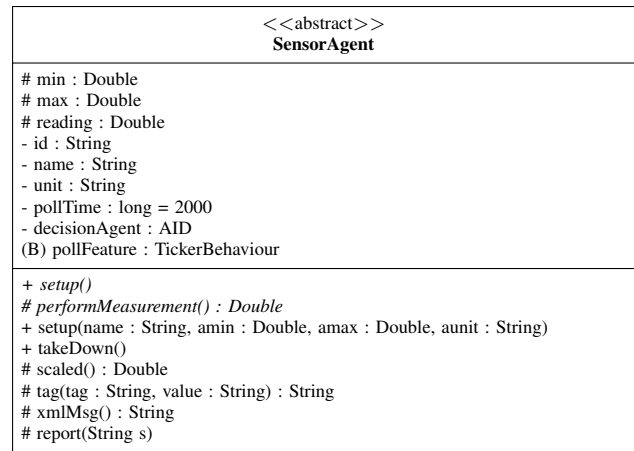


Figure 3. The SensorAgent abstract class.

class has been written to minimise the required amount of boilerplate code. All sensor-agents have a similar structure: a single, repeating behaviour and a standardised method of communicating measurements to the decision-making agent. The `SensorAgent` class, as shown in Figure 3, provides an abstraction of these similarities and only needs a `setup()` method and a `performMeasurement()` method to be implemented to create a concrete class. The `setup()` method should set any relevant variables (of which the variables required by the `SensorAgent` class can be set by calling the `setup(String, Double, Double, String)` methods of the superclass). The `performMeasurement()` method should contain any code needed to perform a measurement, and return its results as a `Double`. The default behaviour of the `SensorAgent` is to execute the measurement-method every 2 seconds and to send the measurement to the decision-making agent using an XML message as shown in Listing 1. The number of 2 seconds can be changed according to the situation and type of measurements to be made. In the example, a subclass called `HeartRateAgent` sends a message containing both the raw values of the measurement (a real number within the range supported by the sensor) and the same measurement scaled to the range of $(-1, 1)$.

Listing 1. A typical message sent from `HeartRateAgent`

```

<measurement>
  <feature>HeartRate</feature>
  <raw>68.950161</raw>
  <scaled>-0.080665</scaled>
</measurement>
                
```

The `SensorAgent` itself contains a number of private and protected attributes to store its specifics such as the name it uses for communication, the specifics of the associated sensor such as minimum / maximum / current values, and the Agent ID of the decision agent. A `SensorAgent` contains a single behaviour, represent here as (B).

As described above, the `SensorAgent` contains two abstract functions needed to implement a subclass. In addition, the class contains a few helper-methods to limit code duplication. The method `scaled()` uses the minimum

and maximum values to apply feature scaling (described in more detail in Section VII-H), `tag(String, String)` and `xmlMsg()` provide a more readable way to generate the XML and `report(String)` is a wrapper around `System.out.println(String)` prepending the output with the agent name to make it easier to distinguish messages when various agents are reporting at the same time.

A. Example: HeartRateAgent

As an example, Listing 2 provides a template for how a `HeartRateAgent` would look as a subclass of `SensorAgent`. The hardware-specific code can be filled in when a sensor has been selected to produce a functional sensor-agent.

Listing 2. A sample `SensorAgent` subclass

```
public class HeartRateAgent extends SensorAgent
{
    public void setup()
    {
        super.setup("HeartRate", 0.0, 150.0, "bpm");
    }

    public Double performMeasurement()
    {
        try
        {
            // read sensor
            // calculate heart rate from reading
            return reading;
        } // feature scaling happens
        catch (Exception e) // when message is sent.
        { // handle exception
            return reading;
        } // return old reading
    }
}
```

IV. THE DECISION AGENT — AGENT REASONING

The decision agent is the central part of our MAS, and is responsible for mapping measurements to assessments of the patient's situation and to request external communication if the situation becomes dire. The agent is trained to perform this task using logistic regression on datapoints provided by a medical expert. The result of this training procedure is a variable θ , which is unique for every patient. Its structure is dependent on the number of factors monitored and the amount of feature-mapping applied, and its contents vary slightly to account for differences between patients, even when the structure is identical. The θ variable and the level of feature-mapping are stored within the agent, and will periodically be used to transform the vector of measurements x into a prediction value representing the certainty the system has of the patients well-being.

In order to make the agent as generalised as possible it will be given a set of general behaviours dependent on variables such as θ , ensuring that most of its behaviour can be changed by sending updated parameters instead of changing the code. The agent will need more information than just the θ variable: it will need a list of "plans" telling it how to react to certain situations. Furthermore, the agent needs to know about the order of the variables within x , the level of feature-mapping, and which agents to contact.

A UML overview of the decision-agent's structure is shown in Figure 4. As shown, the decision agent contains a great amount of variables and operations. The agent contains




Figure 4. The decision-agent class.

matrices (or more precisely, mathematical vectors) for storing θ and x (the latter both scaled and as raw values). A `Map` is used to associate the names of features with their position in x . In addition, it contains a variable `interval` controlling how often a prediction is made, an integer telling the feature mapping function how many polynomials it should generate, a set of `MessageTemplates` allowing it to distinguish various kinds of messages, and a `boolean` indicating whether the agent is operational (*i.e.*, has received a valid set of instructions). The list of `Plans` and the fallback communication method are described below in Section IV-D.

A. Methods

In addition to the `setup()` and `takedown()` methods required by JADE as pseudo-constructors / destructors, the agent contains a number of utility-functions serving to make the code (mostly contained in its behaviours) more readable and to improve maintainability. `binomial(..)`, `combsWithRep(..)`, `numMappedFeatures(..)` and `mapFeatures(..)` are used to perform and verify the feature mapping on the agent.

B. Behaviours

As the decision-agent is the central part of the medical MAS, it contains a large set of behaviours: three cyclic behaviours, which are constantly active, and a ticker behaviour operating on an interval depending on the `interval` variable:

- `handleFailure` listens for messages indicating failure in any of the communication agents. In such an event, it will use a designated fallback-agent to alert an operator that the system might be unable to communicate.

Plan
- below : Double - message : String - recipient : String - agent : AID - limit : Integer - available : boolean = true - timer : Timer
+ Plan(b : Double, m : String, r : String, a : AID, t : Integer) + toString() : String + execute(h : Double) - msg_x() : String - xmlMsg() : String

Figure 5. The plan inner class

- `recvMeasurements` listens for messages from the sensor-agents and saves them in `x` and `x_raw` for use in predictions.
- `recvTheta` listens for messages containing instruction sets. This aspect is explored in Section IV-E.
- `makePrediction` is responsible for periodically multiplying `theta` and `x` to make a prediction regarding the patient's health. This process is described in Section IV-C.

C. Assessing the situation

Every interval milliseconds, the agent uses its then-current knowledge of the features, represented in `x`, to construct a feature-mapped column-vector `x_mapped`. The inner product of `x_mapped` and the row-vector `theta` is passed through the `sigmoid(double)`-method yielding a double between zero and one, representing the probability that the patient is still healthy. As this number decreases, the probability of something being wrong increases. After a prediction has been calculated, the value is compared to the thresholds defined for each plan; if the calculated result is below the threshold for a given plan, the agent will attempt to execute it by messaging a communication agent.

D. Executing plans

Plans are represented by a special `Plan` class, shown in Figure 5. Each time a prediction is made, the agent attempts to invoke the `execute(Double)` method for each plan, passing the predicted probability. Each plan contains a threshold `below`, which is compared to the prediction when `execute(Double)` is called. Execute will send its message to its specified recipient via its specified agent, provided two conditions are met:

- 1) The prediction passed as an argument to `execute(Double)` is lower than or equal to the threshold for the plan and
- 2) The boolean `available` is set to `true`

The value of `available` is initialised as `true`, but is set to `false` when the plan is first executed. At the same time, a timer is started for `limit` seconds, after which `available` is reset to `true`. This prevents the MAS from flooding its recipients with messages as a new prediction is calculated, by default, every five seconds; though it may be meaningful to

provide an occasional update, a realistic poll frequency for the agent to make predictions is likely always higher than a realistic notification frequency. By using a plan specific interval all frequencies can be chosen separately.

E. Receiving instructions

All of the necessary information can be delivered to the agent within a single ACL message; Listing 3 shows a sample XML fragment containing instructions for a decision-agent using two features, including a second-degree feature mapping and two plans.

Listing 3. An initialisation message as sent to the decision-agent.

```

<instructions>
  <features>
    <feature id="SystolicBloodPressure">
      <label>Systolic Blood Pressure</label>
      <min>0</min>
      <max>200</max>
      <unit>mm Hg</unit>
    </feature>
    <feature id="HeartRate">
      <label>Heart Rate</label>
      <min>0</min>
      <max>200</max>
      <unit>bpm</unit>
    </feature>
  </features>
  <mapping>2</mapping>
  <theta>
    <value>2.402548</value>
    <value>2.769392</value>
    <value>3.467782</value>
    <value>-7.500590</value>
    <value>-2.189613</value>
    <value>-11.995721</value>
    <value>-2.301167</value>
    <value>2.064028</value>
    <value>-2.568114</value>
    <value>-2.736256</value>
  </theta>
  <plans>
    <plan>
      <below>0.6</below>
      <message>Watch out!</message>
      <via>ScreenAgent</via>
      <to></to>
      <limit>30</limit>
    </plan>
    <plan>
      <below>0.4</below>
      <message>Panic!</message>
      <via>MailAgent</via>
      <to>brian.vanderbijl@hu.nl</to>
      <limit>3600</limit>
    </plan>
  </plans>
  <fallback>
    <via>ScreenAgent</via>
    <to></to>
  </fallback>
</instructions>

```

The initialisation consists of five parts:

- 1) A `features` node containing information about each feature. The order the features are presented in determines the position of measurements within `x`, and must be identical to the order of features during the learning process.

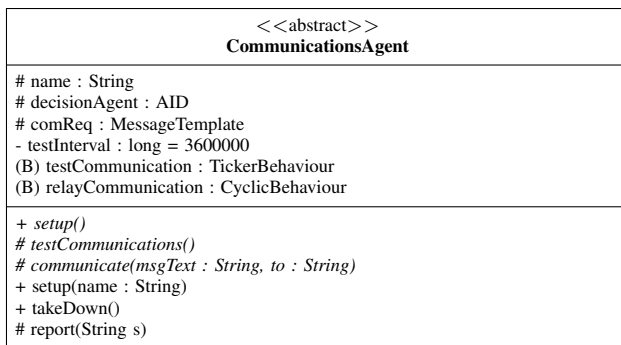


Figure 6. The CommunicationsAgent abstract class.

- 2) A mapping node containing a single integer value determining the amount of feature-mapping.
- 3) A `theta` node containing a series of decimal values representing θ . As with the features, the order is important here.
- 4) A `plans` node containing a set of plans used to respond to predictions. Each plan includes a threshold, below which the plan will be executed, a message to be sent, the name of the agent responsible for relaying the message, an optional recipient (whether this is needed depends on the agent: a mail or SMS agent would require a recipient, whereas a screen agent would not), and a limit in seconds determining how often a plan can be executed in order to avoid flooding messages.
- 5) A `fallback` node containing an agent and a recipient to alert when a communications agent is not functioning properly and cannot be trusted to relay important messages.

When the decision-agent receives a set of instructions, it will confirm whether the length of θ matches the length of x after feature scaling, throwing an error if the two are incompatible. As each level of feature-mapping adds a number of features equal to $\binom{n+d-1}{d}$, the following equality must hold:

$$\text{size}_\theta = \sum_{i=1}^d \binom{\text{size}_x + i - 1}{i}.$$

V. COMMUNICATING RESULTS TO THE OUTSIDE WORLD

The decision-agent as described in Section IV relies on other agents to communicate the results to a medical expert and/or the patients themselves. This choice is deliberate, as it allows new methods of communication to be added “on the fly”, without changing the decision-agent’s behaviour. Each communication agent added to the system represents a new option to communicate the patient’s health and relay concern. As with the sensor-agents, an abstract class has been provided to facilitate the development of additional agents. This class, `CommunicationsAgent`, shown as UML in Figure 6.

Each communication agent has a `String` variable to hold its name and an `AID` representing the decision agent. A variable `testInterval` controls how often the agent performs a

self-diagnostic. Two behaviours are present: one to continually listen for requests for communication, and another to perform the self-tests on an interval dictated by `testInterval`. The class provides the methods for setup, agent destruction and reporting to `stdout`.

Three abstract methods need to be implemented to create a `CommunicationsAgent` subclass:

- 1) `setup()` should set any relevant variables, at the very least including the agent’s name.
- 2) `testCommunications()` should include the code needed to run a self-test, if applicable, and throw an exception if it fails to complete the test. This exception is caught by the `testCommunication` behaviour after which a `FAILURE` message is sent to the decision-agent indicating the communication agent has become unreliable.
- 3) `communicate()` should include all code needed to send a message, such as setting up the necessary objects for IO in Java (provided this needs to be done each time a message is sent; if the method of communication features a persistent object that can be trusted to remain operable, it can be setup in the `setup()` method) and actually sending the message. It will send a `CONFIRM`-message back to the decision agent if the sending process did not encounter any errors; in case of failure it can send either a `NOT_UNDERSTOOD` message to indicate the XML received was illegible, or a `FAILURE` indicating some sort of IO error encountered in trying to relay the message to its recipient.

VI. COMMUNICATIONS AGENTS — REQUESTS FOR OUTSIDE COMMUNICATIONS

Requests for communication from the decision agent are packaged in a small snippet of XML, as shown in Listing 4. The `message`-node contains the body of the email, including the most recent value for each feature.

Listing 4. A typical message sent to MailAgent

```
<request>
  <to>leo.vanmoergestel@hu.nl</to>
  <message>
    Patient health in serious condition!
    - HeartRate = 54.93483905942176
    - SystolicBloodPressure = 86.20808412990199
  </message>
</request>
```

Just as there are various ways to acquire data to facilitate the decision-making agent, there are also many methods to communicate its results. These include telephony, instant messaging, patient-information logs and on-device IO like displays, alarms, etc. The preferred methods of communication may be subject to change over time as the patient’s situation changes, as doctors come and go and as new forms of communication are developed, become widely adopted and are eventually deprecated. Therefore, MAS communication to the outside world should be modular. Just like with sensor-reading agents described in Section III, methods of communications could be implemented by small, trivial single-purpose agents.

By abstracting the way information is delivered, the decision-making agent can communicate its results in a predefined manner indicating the conclusions to be sent and the perceived level of panic. Communications agents can pick up these messages and assume responsibility of relaying the information to the appropriate recipients.

Using the abstract class approach, any method of communication can be added to link the system up to existing medical care, and an existing system can easily be extended to include new ways of communicating. Depending on the availability of usable Java libraries this may be done in relatively small, simple agents.

VII. DATA ACQUISITION FOR AGENT TRAINING

In order to interpret the measurements acquired from the sensors and predict whether the current patient situation constitutes a cause of alarm, the decision agent needs a way to classify potentially high-dimensional data. Each biological factor considered in the model represents an additional dimension for data points. As this information is not guaranteed to be available for various combinations of biological features, it makes sense to explore a way for medical personnel to easily enter such data into the system. Not only does this guarantee the required data can be generated, if not available, it also allows for far greater personalisation, providing the agent with a data-set tailored to its patient. Manual entry, or at least confirmation, also allows an expert intimate knowledge of the agents decision-making process, potentially increasing trust by removing the “black box” aspect of machine learning.

Teaching the system to recognise alarming measurements and differentiate between various levels of threat requires large amounts of information provided by medical personnel, preferably tailored to the patient as thresholds might not be the same for every person. Entering this data can be challenging: as potentially multiple factors need to be taken into account together, it becomes progressively harder for humans to visualise and communicate relevant thresholds. A better way might be to input a set of data-points, together with appropriate assessments of the situation associated with each data-point. These data-points could be used, alone or in conjunction with more general datasets, to train a classification algorithm.

In order to train an agent to make accurate predictions, training data will need to be entered into the system by a medical expert. This should be as easy as possible: the focus should be to quickly train an agent without expending significant time accommodating the system. Unfortunately, entering possibly poly-dimensional data graphically is a difficult task. For one or two dimensional data, clicking points in a scatter plot, as pictured in Figure 7, can be a quick way to enter points; for three dimensional data this becomes harder: a scatter-plot is still possible for data-visualisation, but entry becomes impossible as a mouse or trackpad and a computer screen are both essentially two-dimensional. For even more simultaneous features, only a subset of the features can be plotted at the same time.

An alternative approach would be to require the expert to manually enter all features, as well as the results that the system should predict. Not only is this rather work-intensive, but also prone to omissions: as it is hard for the human mind to visualise all features simultaneously and large gaps are a significant risk.

A better solution would be for the system to dynamically suggest data-points based on the largest knowledge gaps. An expert would then be provided with the parameters for a new datapoint by the algorithm. For this datapoint an assessment of the situation can then be entered. The algorithm continuously updates its collection of datapoints, as well as the model derived from the combination of datapoints and expert assessments, and proceeds to suggest the largest empty areas in its knowledge-continuum as possible locations for new datapoints. This continues until the expert considers the fit of the model to be satisfactory, after which the model is accepted. The expert remains in control of the process of entering datapoints, and can at any time ignore a suggestion or opt to enter the parameters for a new datapoint himself.

This section considers an approach to accomplish this. Each problem will be examined in two dimensions first, as this makes it easier to visualise and demonstrate the applied methods. After the solution has been sufficiently exposed a generalisation can be made in n -dimensions.

To represent gaps in the knowledge-continuum, we create a triangulation of the known datapoints. Each datapoint is considered a vertex in an n -dimensional space, and by triangulating over this set of vertices we can detect sparsely populated areas by the emergence of larger triangles. In contrast, a large amount of datapoints in close proximity will yield a large number of smaller triangles.

Triangles and their higher-dimensional analogues (the tetrahedron in three dimensions, the 5-cell in four, etc.) are collectively referred to as n -simplex or just simplices (singular: simplex). As a triangle (2-simplex) is defined by three vertices of the form (x, y) and a tetrahedron (3-simplex) is defined by four vertices of the form (x, y, z) , an n -simplex is the most basic n -dimensional object defined by $n + 1$ vertices in n -dimensional space.




Figure 7. Scatter plot in two dimensions of a small random dataset.

A. Finding the most valuable points for data-querying

When entering data-points to train an agent, some points are more valuable than others. For example, potential locations completely surrounded by existing data-points all belonging to the same class are unlikely to add any new information to the system. Similarly, points in sparse areas are potentially more valuable, as are points closer to the centre of the point cloud. Figure 8 shows the same scatter plot as Figure 7, but adds a decision boundary and three possible locations for new data points marked by numbers. Location 1 does not appear to be a good addition, as it is very close to existing points and is therefore unlikely to add a great deal of information. Location 2 is not a good suggestion either, as it is very far from the decision boundary — it will likely have the same category as the points surrounding it, especially if a large amount of data has been entered. Location 3 is a better spot for a new data point: it is not a near duplicate of another point, and it lies close to the decision boundary. Depending on the category this point will be assigned to it may significantly change the decision boundary in either direction.

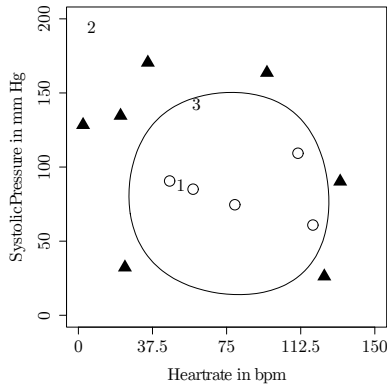


Figure 8. Three possible locations for a new data-point.

B. Data-point-distribution

To find sparsely populated areas to add new data-points, we first create a triangulation containing all data points. For each of these triangles, the circumcentre is calculated, and the collection is ordered based on the area of the triangles. These points can now be evaluated in order to find points close to the current decision-boundary.

C. Triangulating n-dimensional space in simplices

To triangulate a set of points we utilise the Delaunay Triangulation [7]. Most mathematical libraries include a function to quickly get the Delaunay Triangulation of a set of points in n dimensions. Triangulating the example data from Figure 7 yields the triangulation as shown in Figure 9.

D. Calculating the size of each n-simplex

To find the largest simplex we use the determinant of the matrix constructed by adding each vector representing a vertex as a single column, and adding a final row of ones [8]. For a triangle, the absolute value of the result is equal to two factorial times the triangle's area. For a tetrahedron, the absolute value




Figure 9. Triangulation and scatter plot in two dimensions

equals three factorial times the volume. For higher-dimensional shapes, this method continues to yield a scalar multiple of the n -hypervolume of the simplex. As the simplex size is only used for sorting, the scalar multiplication does not influence the ordering and can safely be ignored. As an example, the size of a triangle described by $a = (0, 0)$, $b = (0, 4)$ and $c = (3, 0)$ is given by

$$\text{abs} \left(\begin{vmatrix} 0 & 0 & 3 \\ 0 & 4 & 0 \\ 1 & 1 & 1 \end{vmatrix} \right) = 12 \quad (1)$$

which is twice the area of the triangle.

E. Calculating the circumcentre of each n-simplex

Once the largest data-gap has been found, we want to find its centre to suggest as a new data point. A simplex has multiple definitions of its centre; for this purpose the circumcentre, the point equidistant from all its vertices [9], seems a logical choice. Given a n -simplex defined by vertex $v^{(1)}, v^{(2)}, \dots, v^{(n+1)}$ with a circumcentre c , we know that the distance between any vertex and c must, by definition, be equal. For any two vertices $v^{(a)}$ and $v^{(b)}$, this means:

$$\begin{aligned} \|v^{(a)} - c\| &= \|v^{(b)} - c\| \\ \|v^{(a)} - c\|^2 &= \|v^{(b)} - c\|^2 \\ (v^{(a)} - c) \cdot (v^{(a)} - c) &= (v^{(b)} - c) \cdot (v^{(b)} - c) \end{aligned} \quad (2)$$

We translate each vector by $-v^{(1)}$ so that $v^{(1)}$ becomes the origin (denoted o) and equate the distance to c of each remaining vector with the distance of c to o , yielding the locus for each translated vertex v and the origin o :

$$\begin{aligned} (o - c) \cdot (o - c) &= (v - c) \cdot (v - c) \\ c^2 &= v^2 - 2v \cdot c + c^2 \\ 2v \cdot c &= v^2 \\ v \cdot c &= 0.5v^2 \\ v_1c_1 + v_2c_2 + \dots + v_nc_n &= 0.5\|v\|^2 \end{aligned} \quad (3)$$

Doing this for every vertex $v^{(2)}$ to $v^{(n+1)}$ gives us n equations, allowing us to find the n -dimensional vector c .

We can write these equations in matrix form and solve all equations simultaneously:

Writing

$$S = \begin{pmatrix} v_1^{(2)} - v_1^{(1)} & v_1^{(2)} - v_1^{(1)} & \dots & v_1^{(2)} - v_1^{(1)} \\ v_2^{(3)} - v_2^{(1)} & v_2^{(3)} - v_2^{(1)} & \dots & v_2^{(3)} - v_2^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ v_n^{(n+1)} - v_n^{(1)} & v_n^{(n+1)} - v_n^{(1)} & \dots & v_n^{(n+1)} - v_n^{(1)} \end{pmatrix}$$

$$c = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} \quad r = 0.5 \begin{pmatrix} \|v^{(2)} - v^{(1)}\|^2 \\ \|v^{(3)} - v^{(1)}\|^2 \\ \vdots \\ \|v^{(n+1)} - v^{(1)}\|^2 \end{pmatrix}, \quad (4)$$

we have

$$Sc = r. \quad (5)$$

Given this, we can multiply both sides by S^{-1} to get

$$c = S^{-1}r. \quad (6)$$

As c was translated by $-v^{(1)}$, all that remains is adding $v^{(1)}$ to find the triangle's circumcentre.

F. Avoiding suggesting out-of-bounds points

As shown in Figure 9, Delaunay triangulations are prone to yielding obtuse simplices, in particular around the edges. This can be a problem because an obtuse simplex has a circumcentre outside itself. On the edges, this will result in the algorithm suggesting points outside the sensor's bounds. As these points are meaningless and only serve to distract the user, we would like to avoid generating obtuse simplices.

We solve this problem by introducing a border of false data-points around the edge. These data-points are only used to determine the Delaunay triangulation, and are not present in the actual training-data being generated. The number of data-points is determined by a variable $\beta \in \mathbb{N}_1$: For $\beta = 1$, only the corners of the graph are added. For larger values of β , each axis is subdivided into β parts. As β becomes larger, out-of-bounds points become increasingly unlikely, and suggestions start to gravitate towards existing data-points.

As Figure 10 and Figure 11 show, too large a value for β makes the algorithm increasingly unlikely to suggest points around the edges. Though more central points are preferred, limiting data-points to a central cluster might not be the way to go. A solution for this could be to gradually decrease β over time.

G. Generating the borders

The set of points to be used as a border constitutes of the following:

- a point for each vertex of the n-cube describing the range of data-points




Figure 10. Triangulation for $\beta \in \{1, 2, 3, 8, 12\}$ alongside original triangulation.

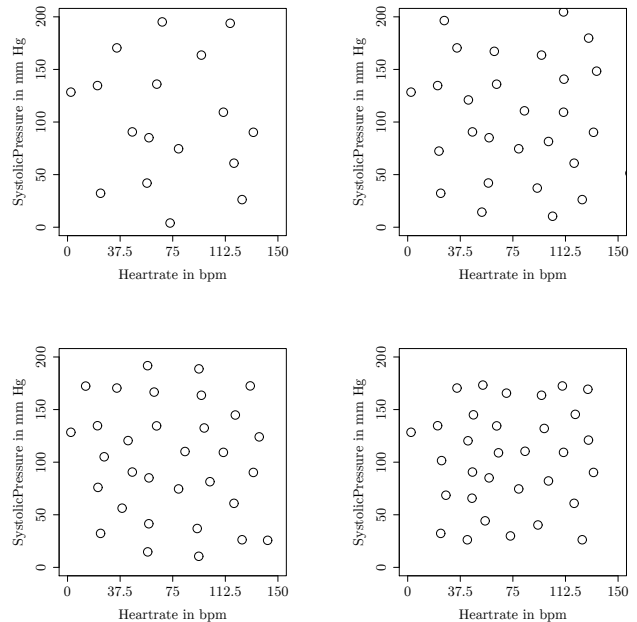


Figure 11. Scatter plot of the first twenty suggestions for $\beta \in \{1, 2, 4, 8\}$. Note that out-of-bounds points are not plotted.

- $(\beta - 1)$ points on each edge (1-face)
- $(\beta - 1)^2$ points on each face (2-face)
- $(\beta - 1)^3$ points on each cell (3-face)
- ...
- $(\beta - 1)^{n-1}$ points on each $(n - 1)$ -face

The number of points denoted by $\#P$ needed given a dimensionality n and a border-saturation β can therefore be

calculated by

$$\#P(n, \beta) = \sum_{i=0}^{n-1} F(n, i)(\beta - 1)^i \quad (7)$$

where $F(n, i)$ is the number of i -faces on a n -cube [10]:

$$F(n, i) = 2^{n-i} \binom{n}{i} \quad (8)$$

The actual value of $P(n, \beta)$ can intuitively be seen as the Cartesian product of n instances of $\text{interval}(\beta)$, also known as its Cartesian Power, of which only those points for which at least one of its members is equal to -1 or 1 are kept. In other words, for which the infinity norm $\|x\|_\infty$ equals 1 .

$$P(n, \beta) = \{x \mid x \in \text{interval}(\beta)^n \wedge \|x\|_\infty = 1\} \quad (9)$$

$$\|x\|_\infty = \max_i |x_i| \quad (10)$$

H. Feature Scaling

The interval-function creates an interval between -1 and 1 in β steps. This is because all features are scaled to lie between -1 and 1 , even though the actual measurements might range from 0 to some arbitrary maximum. This feature scaling is applied to make sure that all features are of the same importance when applying logit later on.

I. Avoiding symmetry

The algorithm presented above tends to favour generating a symmetrical data-set: As the range of values is a perfect n -cube, the first point suggested will be the centre, followed by a group of points equidistant from the first. This is undesirable, as symmetrical data points feature will introduce redundant features when multiplied during the fmap process. It will not help in generating a better hypothesis but will slow down the learning algorithm.

To prevent generating such a duplicate set of data, we will move each suggestion by a small random amount, controlled by a variable δ , that represents the maximal displacement for each point in each dimension. In order to ensure that this displacement will not place points outside the feature boundaries, this displacement will be opposite to the sign of the original location. This results in the data point being moved slightly towards the centre, which generally is the most interesting area to collect data on. We achieve this by replacing each vector element c_i by the weighted mean of $r \cdot 0$ and $(1 - r)c_i$, where $r \sim U([0, \delta])$ is a random variable uniformly distributed on $[0, \delta]$.

VIII. IMPLEMENTATION

For the implementation of the proof of concept, Java agent development framework (Jade) [11] has been used. The Jade runtime environment implements message-based communication between agents running on different platforms connected by a network. The reasons for choosing Jade are:

- the system presented is a multi-agent-based system. Jade provides the requirements for multiagent systems;
- the agent communication standard "Foundation for Intelligent Physical Agents" (FIPA) [12] is included in Jade;
- Jade is Java-based and it has a low learning curve for Java programmers; Java is a versatile and powerful programming language;
- Jade is developed and supported by an active user community.

The prototype has been developed and implemented on a standard Linux-based laptop. It should be possible to operate the system on any small device capable of running Java such as the Raspberry Pi nano [13]. Though the Jade-platform was selected for the prototype, this does not preclude development of a medical MAS in another framework or language. The concepts explored here can be implemented in any language, though support for a solid agent-development framework would be a serious asset. Nevertheless, if better performance is needed, the same principles could be implemented in a lower-level language, such as C, reducing much of the overhead at the cost of lower maintainability.

The prototype has been built and the working has been tested. In summary the following results have been achieved:

- The concept of a medical MAS consisting of three types of agents working together to monitor the patient and communicate the result.
- A method of collecting data from medical experts and utilising this knowledge to teach an agent to evaluate readings provided by sensors.
- The beginnings of a generalised framework upon which to build agents for inclusion in a medical MAS.

The next step will be implementing the system in the real world and testing the usability.

IX. RELATED WORK

Agent-based monitoring for computer networks has been proposed and implemented by Burgess. Burgess [14] [15] describes Cfengine that uses agent technology in monitoring computer systems and ICT network infrastructure. In Cfengine, agents will monitor the status and health of software parts of a complex network infrastructure. In [16], an agent-based monitoring system is proposed. A so-called product agent is responsible to monitor the working of a system in several different phases of its lifecycle. The actions performed by the agent are limited to prevent disasters or misuse. The aforementioned concept of a product agent that supports a product during its lifecycle from production to recycling is described in [17].

A lot of literature is available regarding health monitoring systems. Pantelopoulos and Bourbakis [18] give an overview of wearable sensor-based systems for health monitoring and prognosis. Their work focusses on the hardware implementation of the monitoring systems as well as communication

technologies that might be used by such systems. The work of Milenkovic [19] is dedicated to wireless sensor networks in personal health monitoring. The system they describe collects data that is transferred to a central monitoring system whereas the system described in our paper aims for autonomous operation. Furthermore, monitoring systems that focus on special health related situations exist, such as the work of Marder et al. [20] where a system for monitoring patients with schizophrenia is described. An agent-based health monitoring as a concept for application of agent technology has been proposed by Jennings and Wooldridge in [21].

X. CONCLUSION

In this paper, a complex, expandable and agent-based monitoring system has been proposed and a proof of concept was built. The system turned out to work as expected. The design of the MAS has been described in detail as well as the communication between the different types of agents. Special attention has been given to the way the system builds its knowledge-base, resulting in an efficient system that focusses on the borders of operating space where transitions from one situation to another situation are possible. In the case of the medical monitoring system, this could result in a personal adapted monitoring system that can also be easily changed. Though the system is designed for use in a medical context, the concepts can be used in other domains as well.

REFERENCES

- [1] L. v. Moergestel, B. v. d. Bijl, E. Puik, D. Telgen, and J.-J. Meyer, "A multiagent system for monitoring health," IARIA, Intelli The Fifth International Conference on Intelligent Systems and Applications, Barcelona, Spain, 2016, pp. 57–62.
- [2] L. v. Moergestel, J.-J. Meyer, E. Puik, and D. Telgen, "A versatile agile agent-based infrastructure for hybrid production environments," IFAC Modeling in Manufacturing proceedings, Saint Petersburg, 2013, pp. 210–215.
- [3] J. T. Parer and T. Ikeda, "A framework for standardized management of intrapartum fetal heart rate patterns," American Journal of Obstetrics and Gynecology, vol. 197, no. 1, 2007, pp. 26.e1 – 26.e6.
- [4] M. Wooldridge, An Introduction to MultiAgent Systems, Second Edition. Sussex, UK: Wiley, 2009.
- [5] L. v. Moergestel, J.-J. Meyer, E. Puik, and D. Telgen, "Embedded autonomous agents in products supporting repair and recycling," Proceedings of the International Symposium on Autonomous Distributed Systems (ISADS 2013) Mexico City, 2013, pp. 67–74.
- [6] Foundation for Intelligent Physical Agents. FIPA ACL Message Structure Specification. [Online]. Available: <http://www.fipa.org/specs/fipa00061/SC00061G.html>
- [7] Wolfram MathWorld. Delaunay Triangulation. [Online]. Available: <http://mathworld.wolfram.com/DelaunayTriangulation.html>
- [8] P. Stein, "A note on the volume of a simplex," The American Mathematical Monthly, vol. 73, no. 3, 1966, pp. 299–301. [Online]. Available: <http://www.jstor.org/stable/2315353>
- [9] Wolfram MathWorld. Circumcenter. [Online]. Available: <http://mathworld.wolfram.com/Circumcenter.html>
- [10] R. J. McCann, "Cube face," 2010. [Online]. Available: <http://www.math.toronto.edu/mccann/assignments/199S/cubeface.pdf>
- [11] Telecom Italia. JAVA Agent DEvelopment Framework. [Online]. Available: <http://jade.tilab.com/>
- [12] Foundation for Intelligent Physical Agents. FIPA. [Online]. Available: <http://www.fipa.org/>
- [13] E. Upton. Oracle Java on Raspberry Pi. [Online]. Available: <https://www.raspberrypi.org/blog/oracle-java-on-raspberry-pi/>
- [14] M. Burgess, "Cfengine as a component of computer immune-systems,," Proceedings of the Norwegian Informatics Conference, 1998, pp. 283–298.
- [15] M. Burgess, H. Hagerud, S. Straumnes, and T. Reitan, "Measuring system normality," ACM Transactions on Computer Systems (TOCS) Volume 20 Issue 2, 2002, pp. 125–160.
- [16] L. v. Moergestel, J.-J. Meyer, E. Puik, and D. Telgen, "Monitoring agents in complex products enhancing a discovery robot with an agent for monitoring, maintenance and disaster prevention," ICAART 2013 proceedings, vol. 2, 2013, pp. 5–13.
- [17] L. v. Moergestel, E. Puik, D. Telgen, and J.-J. Meyer, "The role of agents in the lifecycle of a product," CMD 2010 proceedings, 2010, pp. 28–32.
- [18] A. Pantelopoulou and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 40, no. 1, 2010, pp. 1–12.
- [19] A. Milenković, C. Otto, and E. Jovanov, "Wireless sensor networks for personal health monitoring: Issues and an implementation," Computer communications, vol. 29, no. 13, 2006, pp. 2521–2533.
- [20] S. R. Marder, S. M. Essock, A. L. Miller, R. W. Buchanan, D. E. Casey, J. M. Davis, J. M. Kane, J. A. Lieberman, N. R. Schooler, N. Covell et al., "Physical health monitoring of patients with schizophrenia," American Journal of Psychiatry, vol. 161, no. 8, 2004, pp. 1334–1349.
- [21] N. R. Jennings and M. Wooldridge, "Applications of intelligent agents," in Agent technology. Springer, 1998, pp. 3–28.