

Regulated Information Sharing and Pattern Recognition for Smart Cities

John K.C. Kingston

Centre for Secure, Intelligent and Usable Systems, University of Brighton BN2 4GJ
j.k.kingston@brighton.ac.uk

Abstract. This paper describes applications of BOBBIN, a multi-agent system based on a blackboard architecture, to supporting smart cities through regulated information sharing and pattern recognition. The first application uses the knowledge-based processing provided by a blackboard system to ensure that personal information is made available to all and only those who are entitled to see it, thus overcoming the objections raised by the UK Supreme Court to a recently proposed information sharing scheme. The second application extends the first to enable pattern recognition over appropriately regulated information about individuals, with the goal of identifying criminal offences and other patterns of behavior. The final application extends the second to deal with heterogeneous agents, public and private. Each application is illustrated with a scenario.

These three approaches can provide benefits in supporting the growth of intelligent public services and appropriate information sharing within a smart city.

Keywords: Regulated Information; Smart Cities; Pattern Recognition; Blackboard Systems

1 Introduction

This paper describes the use of BOBBIN (the Brighton Blackboard Information System) to support three different information sharing and pattern recognition applications that help to enable “smart cities”. The applications are implemented in proof-of-concept demonstration systems.

The first application is called ‘Named Person Lite’. It is based on the proposed ‘Named Person’ scheme that the Scottish Government planned to set up to allow a single individual with social care responsibilities to collect all relevant information about a child and to act on it if necessary. The scheme was declared illegal by the UK Supreme Court, partly because of concerns that the ‘named person’ was permitted to be selective in the information they shared and the institutions with whom they shared it. This paper proposes that BOBBIN is capable of enabling many of the planned benefits of the scheme whilst avoiding the Supreme Court’s concerns by automating and regulating the sharing of information.

The second application extends the first with a higher ‘blackboard’ layer which draws appropriate information from each individual ‘Named Person Lite’ blackboard agent. The goal is to recognise patterns that are known to recur across different cases in order to detect possible criminal offences or other undesirable behavior. The patterns are based on knowledge of typical precursors or indicators of common offences.

The last application extends the architecture again to draw information from various public and private sources including social media. The goal is to enable the public services to respond in a timely fashion based on near real-time information.

2 Intelligent Sharing of Regulated Information

There has been a long-running public debate between the need for privacy of personal data and information, and the value that can be obtained by sharing information and data. The debate covers sharing of health data for use in medical research (see e.g. [1]); allowing law enforcement access to private data [2]; and online gossip about celebrities (e.g. [3]) among other issues. As the law stands in the UK, most information about individuals can only be shared with the written consent of that individual; there are also restrictions on how that data can be processed [4].

The motivation for the work described in this paper arose from a recent conflict between information sharing and privacy in the context of the Scottish Government's proposed Named Person scheme. The proposal was that each child would have an assigned Named Person who would be available to listen, advise and help a child or young person and their parent(s), directly or by helping them to access other services. The Named Person would also be a clear point of contact if a child, young person or their parents wanted information or advice or to talk about any worries; and a point of contact for other services if they had any concerns about a child's or young person's wellbeing.

The scheme was controversial because it was perceived that the State wanted to appoint a guardian for every child who would replace some of the roles and responsibilities of parents. In the summer of 2016 before the scheme was implemented, the UK Supreme Court declared it to be partially illegal on the grounds of information sharing; the Named Person would have too much information to maintain privacy, and (crucially) there was no way to control how much of that information the Named Person chose to share with other responsible parties.

The Supreme Court did not rule that information sharing was a bad thing; it did, however, rule that unmonitored, unregulated, potentially selective information sharing was a bad thing. What if there was a way of sharing information automatically where the information flow is regulated both in terms of what cannot be shared (due to privacy or confidentiality) and what must be shared?

3 Information Sharing Architecture in BOBBIN

BOBBIN is designed as a 'blackboard' system [5] in which multiple agents communicate with each other via a central 'supervisor agent'. BOBBIN is implemented in JESS, the Java Expert System Shell [6]. This means that APIs with other packages can be programmed using Java.

The information sharing architecture designed within BOBBIN uses a three-tier model in which individual agents occupy the lowest level of the hierarchy; the middle level consists of 'departmental' agents that represent the department or group to which

the individual agents belong; and there is a single ‘supervisor’ agent at the higher level whose role is to mediate information sharing. A diagram of the architecture can be seen in Figure 1.

In a regulated information sharing scenario, an agent (e.g. a social worker) would gather information and send it to its departmental agent in the form of a report. The departmental agent would use knowledge to decide how much of that information could or should be shared with other agents in its department. It might also send the information or a summary of it to the supervisor agent. These decisions would be based on regulatory knowledge; existing information; and any requests for information sent by other departmental agents.

The supervisor agent may then communicate information to other departmental agents. Crucially, the supervisor agent will keep a record of all information that it receives, and that information will be made available to any other departmental agent that has appropriate authority.

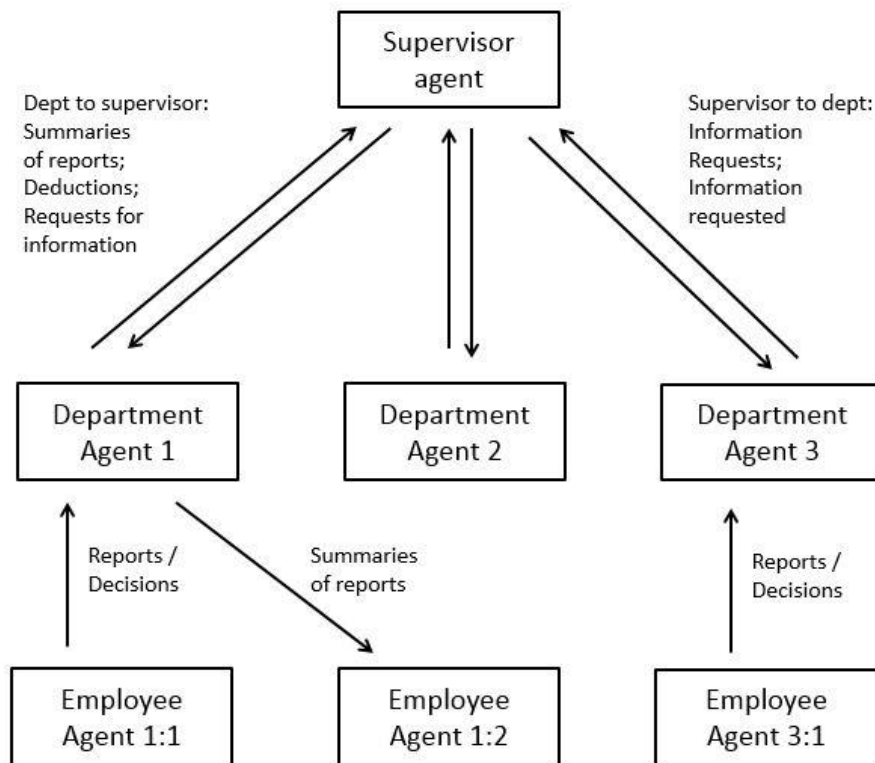


Fig. 1: Information sharing architecture within BOBBIN

Agents never communicate horizontally across the hierarchy; all communication is mediated by the departmental agents and the supervisor agent. This means an agent that

sends an output up the hierarchy cannot control which other agents can see its outputs; that decision is made by the department and supervisor agents.

BOBBIN uses only two templates to represent information. The first template is for Event facts, which describe events that have occurred. The attributes of an Event are:

- Unique ID
- Event type
- Who (was involved)
- When
- Where
- Source (who reported it)
- Importance (0-5 scale; 0 means 'no importance', 5 means 'most important')
- Restriction (on other agents who can see this fact)
- Other

The 'Other' attribute recognises that certain event types may need to record other information apart from the attributes specified. This attribute is likely to be multi-valued; it can contain any type of Java object, including lists or sets. Pattern matching on the 'Other' attribute by rules can be done if necessary because the conditions of JESS rules are allowed to include one or more 'test' functions which can contain any Java predicate.

The second template for data/information is for Tasks. The attributes of Tasks are the same as those for Events (substituting Task Type for Event Type) with one addition: a Status attribute that records whether the task is Required; Optional; Assigned; In_Progress; or Completed. The 'Who' attribute of a Task will record the person/group/agent to whom the task is assigned/who performs the task.

A request for information consists of a Task of type Request_for_Information where the 'Other' field contains a link to an incomplete Event. The supervisor agent is then responsible for supplying the information (if it has it already and the requester – named in the 'Source' attribute -- is authorised) or for forwarding the request to the agent(s) named in the 'Who' attribute. A possible future extension to BOBBIN would be to permit requests for information where 'Who' is unspecified; in such cases, the supervisor agent would monitor incoming information until it found a pattern that matched the request and/or use its knowledge to decide to which department the request should be sent.

4 Regulated Information Sharing: Named Person Lite

A demonstration system has been developed in BOBBIN's information sharing architecture to illustrate how a version of the Named Person scheme might operate within it. This application is called 'Named Person Lite'.

A walkthrough of how the system operates in a scenario is given below. The scenario used was published by the Scottish Government's Get It Right For Every Child initiative to describe how the Named Person scheme would have worked [7]. An edited summary of the scenario is:

A health visitor on a home visit to assess the development of Emma a 2 year old girl notes that the home is very cold and appears unheated. Emma's mum, Ashley says the

heating has been broken in the house for over a week and the private landlord has said that he would come and fix it but has not done so. Emma and her older sister, for whom the health visitor is also the Named Person, are dressed in layers of clothing and are sitting on the couch watching TV under a blanket reluctant to play. Both children have “runny noses” and mum is concerned that they both suffer from night-time coughing.

The health visitor asks mum if she would like her to try to contact the GP for a quick appointment for the children and also get advice and help to address the heating problem. Ashley is enthusiastic to get any help she can as she does not have any local support and was unclear of her rights in relation to the private landlord. The health visitor called the GP surgery and got an urgent appointment for the children the next day. She also called the local authority Crisis Grants Team and informed them of the family’s situation and obtained an appointment with them for later in the day where Ashley is given an emergency grant for a small electric heater. In addition the health visitor gave Ashley the number of her local Citizen’s Advice service to help her get information and advice about how to get the private landlord to accelerate repairs.

In the original scenario, the health visitor was expected to be Emma’s ‘named person’. How could BOBBIN support the same health visitor to achieve the same outcome using a ‘Named Person Lite’ approach?

In the demonstration system, BOBBIN contains a ‘health visitor’ agent for use by an individual health visitor. When the health visitor comes to write a report on her visit, the report is made through BOBBIN. In this scenario, the report includes two Events; it states that there are two children suffering from minor ill-health; there is also a problem with the environment to which greater significance is attached.

BOBBIN asks who is responsible for fixing the environmental problem; a menu of common options is provided. In this case, ‘Private landlord’ is selected. This is added to the Event.

The health visitor agent in BOBBIN includes some rules, one of which is triggered whenever illness and environmental problem(s) are found in the same place. The rule asks the user whether the illness might be caused or triggered by the environmental problem. In this case, this answer is Yes, and that triggers further actions.

Firstly, the health visitor agent offers advice to be passed on to the client. In this case, it prompts the health visitor to suggest to Ashley that she might want to visit her GP and that the local Citizen’s Advice Bureau may be able to help her negotiate with her landlord.

Secondly, it automates the request for a crisis grant by sending a task to a (notional or actual) Crisis Grants agent via the supervisor agent and the Crisis Grants departmental agent. The request provides the Events and asks for a decision on whether an emergency grant can be made available. The Crisis Grants agent is likely to send back a request for information such as how long the problem has persisted for, and how much money is required; this request should appear as a message to the health visitor who can supply the requested information. The Crisis Grants agent can then make a decision in principle based on the information supplied, and the Crisis Grants departmental agent can confirm or reject that decision based on knowledge of the budget and spending priorities. The health visitor does not need to know whether the decision has been made by a human or by an automated agent, or why it was made.

At this stage, the supervisor agent takes action. It has a record of an ongoing request from the Housing department for information on private landlords who are not fulfilling their obligations. While some actions by a landlord might be sufficiently private that the supervisor agent will not disclose them, in this case the supervisor agent determines (using rules) that the Housing department is entitled to see the information that the health visitor has gathered. A possible future enhancement to ‘Named Person Lite’ might be for the supervisor agent to ‘push’ certain information out to agents who might have an interest in it rather than waiting for requests.

In short, the supervisory agent is acting in the place of the Named Person in the role of sharing information between departments. However, it operates strictly according to its rules which are assumed to be regulation-based; any information that is not restricted by regulation is available to all those entitled to see it and only those entitled to see it. It therefore provides a way of sharing information between agents without falling foul of the Supreme Court’s criticisms of the proposed Named Person scheme.

5 Pattern Recognition and Smart Cities

The architecture of BOBBIN is capable of more than just the regulated sharing of information. A second application extends the concept from information sharing to knowledge-based recognition of patterns in information. The goal of this application, as of all the applications described in this paper, is to support a ‘smart city’.

There is no universally accepted definition of what a ‘smart city’ is. Albino et al. [8] review no fewer than twenty-one definitions. The definitions have a similar central conception, but differ in the degree to which the use of information and communication technologies is assumed and in the definitions of what a ‘good’ city is. For current purposes, smart cities will be defined as cities “that use all available technology and resources in an intelligent and co-ordinated manner to develop urban centers that are at once integrated, habitable and sustainable.” [9]

The application proposed here focusses on the concepts of “intelligent” and “co-ordinated”. One of the motivations for the Scottish Government’s proposed Named Person scheme was to protect children who are suffering or at risk of suffering by allowing a single person to co-ordinate all information about that child, rather than having different indicators available to different departments. However, there are (fortunately rare) occasions in towns and cities where numerous children are suffering or at risk from the same source, as happened with the “grooming gangs” in Rochdale [10]. In such cases, a Named Person is needed initially to co-ordinate information about an individual child and then an agent that can co-ordinate information from different Named Persons and intelligently recognize patterns in that information is also required.

The knowledge that drives this pattern matching could be drawn from two sources: similarities observed amongst the experiences of the (actual or potential) victims, or knowledge of typical activities of perpetrators. The former can be handled using data mining techniques – either unsupervised learning to identify any patterns that might exist, or supervised learning to see if any minors or their known contacts fall into the profile of victims or perpetrators. The latter could draw on representations such as the “fraud plans” approach [11] to recognize generic plans and then variations on those plans applied locally.

5.1 Artificial Intelligence and Criminal Profiling

Criminal profiling is a technique for identifying likely suspects in a crime based on similarities in behavior to previous crimes, and (more controversially) to predict likely future crimes and victims. Attempts to apply artificial intelligence to support criminal profiling date back to the 1980s with the FBI's VICAP program for analyzing homicide-related violent crimes [12] which used case based reasoning to compare over 100 features of the modus operandi of a new crime with those already in the database and displays the 10 closest matches. Another program performed a similar function for crimes of arson; in one case, it was used to construct a criminal personality profile describing an individual who could have been responsible for a series of fires at religious homes and houses of worship that summer in a posh New England community, and it not only accurately described the suspect (who later confessed) but also pinpointed his residence.

AI systems that perform crime prediction use two approaches to profiling. The less controversial one is based on actions that are typical precursors to specific crimes: for example, a system is being trialled by the Chinese government that tracks the movements and behavior of individuals (using facial recognition from surveillance cameras) and highlights suspicious behavior [13]. A spokesman for the developers said, "Of course, if someone buys a kitchen knife that's OK, but if the person also buys a sack and a hammer later, that person is becoming suspicious".

The other approach is to use known factors about a person to predict their likelihood of committing crime. Controversy arises because any such system based on data mining inevitably includes demographic information about age, race and residence. An investigation into one such system – COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), which is used in several state jurisdictions in the USA – found the system claiming that black people were almost twice as likely as white people to reoffend [14].

5.2 Approach taken in BOBBIN

The approach taken in BOBBIN is to use knowledge about typical precursors to crimes. The precursors are identified by the 'department agents' and reasoned about by the supervisor's agent. If the department agent identifies a sufficient number of precursors, it passes that information to the supervisor which recommends further action and/or carries out appropriate information sharing with other departments. Table 1 lists some of the typical indicators / precursors of online purchase fraud.

Indicators / precursors of purchase fraud
Not mentioning the exact product that is for sale or other indications they haven't read the advert properly (e.g. requesting a 'best asking price' when the price is listed in the advert).
Being ready to pay 'without delay'.
Offering more than the asking price.
Item to be collected by a shipping company rather than the purchaser.
Paying through a third party
Bad grammar.

Table 1: Indicators/ precursors of purchase fraud [15]

BOBBIN’s supervisor agent links the knowledge about typical precursors to “fraud plans” [11]. This approach represents generic types of fraud and shows how the different generic steps are specialized to create different types of fraud. Figure 2 shows a generic plan for purchase fraud.

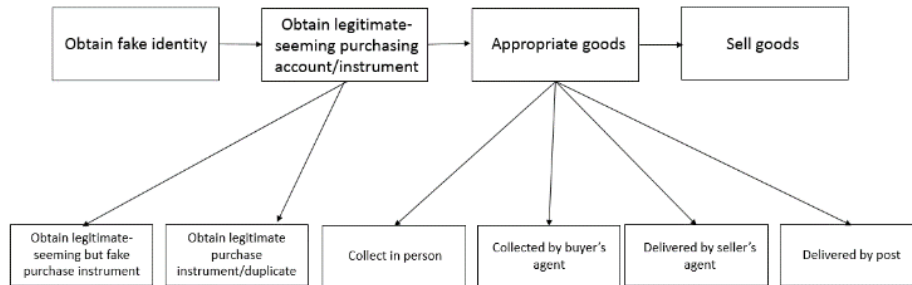


Fig. 2: Generic plan for purchase fraud (from [12])

For example, assume that a motorcycle has been offered for sale online at a price of \$5000. A potential purchaser who claims to live in a different US state replies as follows:

“I really want this as gift for my brother Can you give best asking price I will pay through Paypal account without delay and my shipping agent will collect.”

If the purchaser feeds this message into BOBBIN, it will be sent to the Sales Department agent where the list of fraud precursors is stored. BOBBIN does not yet support identification of these precursors from natural language; instead it is able to present the request to a member of staff in the sales department along with a list of precursors and then to ask if any of the precursors appear to be present. Given that five of the six precursors listed in Table 1 are present, this information is sent to the supervisor agent which recommends sending a message to the buyer stating that the price has been increased (i.e. to test if the final, most damning precursor is also present). It could also notify other supervisor agents in other instances of BOBBIN (if suitable secure connections were available) to scan for messages from the same person/domain.

Precursors can be linked to specific steps in a crime plan. The use of a shipping company and bad grammar (from someone who claims to be a US resident) are linked to “Obtain fake identity”; paying through a third party (Paypal) is linked to “Obtain legitimate seeming purchase instrument”. Others are linked to general characteristics of crime: failing to mention the exact product may be due to a fraudster re-using the same message in many attempts at fraud rather than tailoring them to individual targets, while “without delay” is typical of almost all frauds – for a fraudster, time is always vital because they want to complete their plan before the fraud is discovered. Using such a structured approach helps to explain why precursors are linked to offence types; to recognize unknown actions as being variants on previous precursors; to prompt human experts for further precursors; or to classify patterns of precursors found through machine learning.

6 Real Time Response in the Smart City

The architecture can be adapted to take input not only from Named Person agents but from other sources. These sources might include information websites about subjects such as transport, shopping, public events or weather, and social media sources.

Much of this information will be publicly available, but some may be regulated or otherwise restricted. For example, it is possible that a ‘smart city’ might include a social media network amongst shopkeepers. If a known shoplifter appears in one of their shops, they have agreed to share this information amongst themselves. Since publicising a claim that certain person is a known criminal could easily lead to a suit for defamation, they are careful to restrict that information from public availability; however, they may agree to share it, or a summary of it, with the police via an appropriately restricted channel. The BOBBIN architecture provides a suitable channel.

The primary benefit obtained from these sources is information that is updated in near real time. There are also secondary benefits in localization; some sources will address specific localities within a city.

An illustrative scenario assumes that a city contains:

- A neighbourhood watch group that has its own page on social media where they share their observations of anything they deem anti-social, from rude teenagers to properties with overgrown gardens. They are only too happy to give the police access to their page.
- The local rail operator runs an opt-in alert service over social media through which travellers can be told of delays or cancellations to trains.

The scenario is:

It is early Saturday afternoon in Smarton. Most city council officials are off duty for the weekend, although their social media monitoring team has one person on duty. The police, in contrast, are out in force in the main shopping centre and at the football stadium where a big match is taking place in a couple of hours. The first supporters are already arriving at the ground.

The rail operator announces through its web service that a train from London will be arriving late. The train will be carrying many supporters of the away team. The police chief’s tablet computer sends him a message reminding him that he is entitled to delay the kick-off of the football match if necessary but recommending that he does not do so based on current information about the train’s arrival.

A message appears on the local neighbourhood watch group complaining about a group of young men walking purposefully along the middle of a suburban road (and temporarily blocking it) while singing rowdy songs. The police’s social media monitoring team enter this into their intelligent software system. The software system recommends that since the group is moving rather than congregating and is not committing or attempting any criminal activity no immediate police presence is necessary.

The software assesses why such a group might have appeared. It is programmed to consider various possibilities including political rallies; organised fights; ‘steaming’ (mass theft/robbery); rioters; flashmobs; charity collectors; and missionaries/carol

singers. Given the time of day, the time of year and the reported membership and behaviour of the group, it decides that the highest likelihood is that they are football supporters heading for the match. It calculates their likely routes.

An alert flashes on the police chief's tablet device. The information about the rowdy group has been combined with information about the late-arriving train. The police chief reads that the rowdy group is likely to pass the railway station at almost exactly the time that the train full of away fans will arrive.

He quickly deploys two officers to re-route the rowdy group down a different route that keeps them away from the station.

The architecture of such a system can be seen in Figure 3. The current demonstrator version of BOBBIN lacks the capability to calculate routes or to access train timetables but is able to accept input from humans monitoring other information sources; to reason about why a large group of young men might have appeared; and to pass relevant information between agents. It uses a rule-based architecture for its reasoning; common groups are represented as objects with attributes such as gender, age, clothing, time of day and time of year and the rule-based system asks questions to distinguish groups based on those attributes.

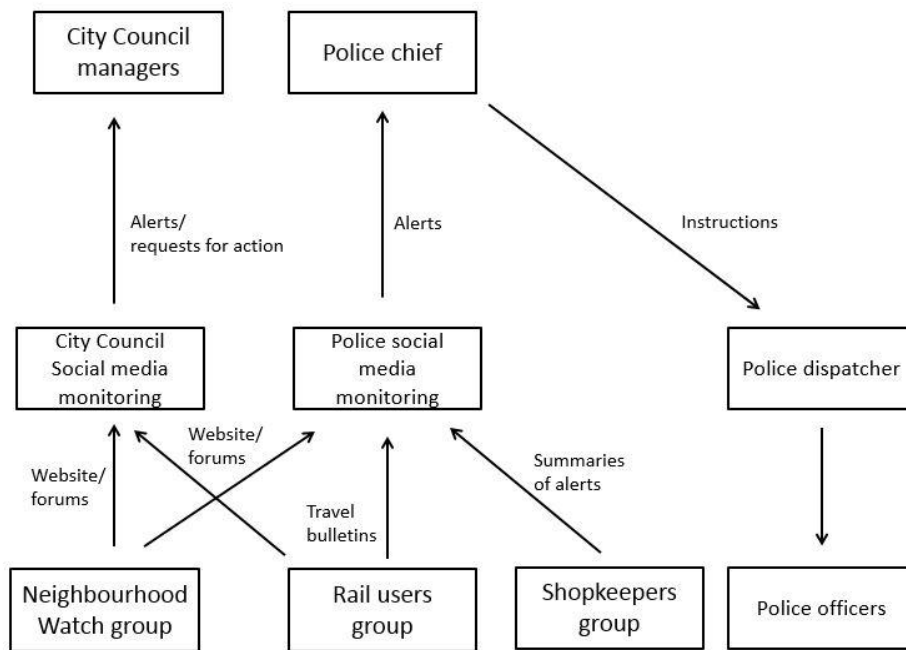


Fig. 3: Real Time Response Architecture for a Smart City

7 Future Work

Possible enhancements to the systems described above could include:

7.1 Natural language understanding

One possible enhancement to this system would be to implement text mining on the web pages/messages of the various information groups, so that instead of a human being having to monitor the groups all the time, automated alerts could be sent when something important is posted by one of these groups.

7.2 Conflict identification and resolution

With many different information sources contributing to it, it is inevitable that BOBBIN will eventually be supplied with contradictory information about the same event, or contradictory cases that address the same ‘problem’.

BOBBIN is well placed to identify such conflicts, since all information that is shared passes through the supervisor agent; all that is needed is a few rules to detect occasions where an event of the same type has been reported at the same time and/or location but with other details being different, or where two or more cases offer contradicting solutions to a ‘problem’.

As for conflict resolution, the quickest method is to ask the system’s user to resolve it. The next quickest is to establish a priority of sources; it may come down to trust (e.g. believing the rail operator over a local resident), although this should be treated as a heuristic rather than a firm rule.

A third method would be to send a message back to the agents who have supplied the contradicting/contradicted information reporting the contradiction, asking them to either change or explain the information that they have supplied. Experience suggests that often both information providers may well have been well-meaning but made different assumptions about the context surrounding the event.

A fourth approach is to treat both information items as provisionally true, and to mark any deductions made on the basis of these items as assumption-based. Assumption-based truth maintenance systems (ATMS) [16] are a recognised technology in Artificial Intelligence and can be used to reason with hypothetical or unproven information, with all such reasoning being retracted if the initial information is proven to be untrue.

7.3 Custom-built workflow

A further knowledge management application that could be built within BOBBIN would be to use the blackboard architecture to develop customised workflows. In this scenario, the ‘agents’ would be different individuals or departments who could offer specialised workflows that solved particular problems. A problem would be posted on the blackboard (a set of events and a request to analyse or transform those events to achieve a particular goal) and these would be matched against the various specialised

workflows. If one of those specialised workflows applied, it would either simulate processing of the events or would actually be carried out and its results would be posted on the blackboard, which might achieve the goal or which might be used to trigger another specialised workflow that advances towards the goal. Finally, the set of workflows used to achieve the goal is assembled into a sequence that can be applied to achieve, or has already been used to achieve the real goal.

For example, imagine that a company had become aware that one of its employees might be mis-using company computers. If the company's various departments were running BOBBIN's custom workflow application, the company could post a description of the events (employee details) and goal (determine if any breach of company policy has occurred) to BOBBIN's supervisory agent. They might expect the Human Resources agent to be first to respond, but it would not do so because the HR agent only has workflows to deal with an employee who is known to be misbehaving. In fact it would be IT Services who would offer two workflows: one to run an automated keyword analysis of the employee's email traffic which is held on the central server; the other to create a disk image of the employee's computer that might be used by a later workflow that conducts forensic analysis of disk images.

The company is informed of IT's response via BOBBIN and authorises them to do the email analysis. The results will be posted onto BOBBIN. If it is found that the employee has emailed confidential documents to his personal email account, then one or more of HR's workflows would apply, and there might also be good reason to carry out the time-consuming forensic analysis; which was done first would depend on other priorities.

References

1. Abbas A, Khan SU. (2014) A review on the state-of-the-art privacy-preserving approaches in the e-health clouds. *IEEE Journal of Biomedical and Health Informatics*. Jul; 18, 4: 1431-41.
2. Barrett B. (2016) The Apple-FBI Battle is over, but the new Crypto Wars have just begun. *Wired*, March 30th. <https://www.wired.com/2016/03/apple-fbi-battle-crypto-wars-just-begun/>. Downloaded 19 August 2016.
3. Naomi Campbell v Mirror Group Newspapers (2002-2004). [2002] EWHC 499 (Q.B.); [2002] IP&T 944 (CA); [2004] IP&T 764 (HL), [2004] UKHL 22.
4. Data Protection Act 1998. <http://www.legislation.gov.uk/ukpga/1998/29/contents>.
5. Hayes-Roth, B. (1985). A blackboard architecture for control. *Artificial Intelligence*. 26, 3: 251–321.
6. Friedman-Hill E.J. (1997) Jess, the Java Expert System Shell. Sandia National Laboratories Report SAND98-8206.UC-405. <http://prod.sandia.gov/techlib/access-control.cgi/1998/988206.pdf>. Downloaded 19 August 2016.
7. Get It Right For Every Child initiative (2016) Scenario 4: Emma’s story. Scottish Government. ISBN 978-1-78652-292-4. <http://www.gov.scot/Resource/0050/00502049.pdf>. Downloaded 19 August 2016.
8. Albino V., U. Berardi and R.M. D’Angelico (2015). Smart Cities: Definitions, Dimensions, Performance and initiatives, *J. Urban Technology* 22, 1: 3-21.
9. Barrionuevo J.M., P. Berrone and J.E. Ricart (2012). Smart Cities: Sustainable Progress, *IESE Insight* 14, 50-57.
10. Bunyan, Nigel (8 May 2012). "Rochdale grooming trial: gang convicted for sex traffick-ing". *Daily Telegraph*. London. Retrieved 8 May 2012.
11. J. Kingston, “Representing, reasoning and predicting fraud using fraud plans” Proceedings of Research Challenges in Information Science (RCIS-2017), University of Brighton, 10-12 May 2017.
12. D.J. Icove, “Automated Crime Profiling”. FBI National Centre for the Analysis of Violent Crime, December 1986. <http://www.icove.com/images/crimeprofiling.pdf>
13. Y. Yang, Y. Yang and S.F. Ju, “China seeks glimpse of citizens’ future with crime-predicting AI”, *Financial Times*, July 23 2017, <https://www.ft.com/content/5ec7093c-6e06-11e7-b9c7-15af748b60d0>
14. J. Angwin, J. Larson, S. Mattu and L. Kirchner, “Machine Bias”. ProPublica, May 23 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing/>
15. N. Provo, “Craigslist Scams Revisited – 5 Main Scam Indicators”, May 30 2010, <https://mollermarketing.com/2010/05/craigslist-scams-revisited-5-main-scam-indicators/>
16. R. Reiter and J. de Kleer, “Foundations of assumption-based truth maintenance systems: preliminary report”. In Proceedings of the Sixth National conference on Artificial intelligence (AAAI’87), Vol 1, pp. 183-188. Seattle, Washington. AAAI Press.