This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Responsibility modelling for risk analysis

PLEASE CITE THE PUBLISHED VERSION

PUBLISHER

© CRC Press

VERSION

AM (Accepted Manuscript)

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Lock, Russell, Tim Storer, Ian Sommerville, and Gordon Baxter. 2019. "Responsibility Modelling for Risk Analysis". figshare. <https://hdl.handle.net/2134/8811>.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

Responsibility Modelling for Risk Analysis

Russell Lock, Tim Storer, Ian Sommerville & Gordon Baxter

School of Computer Science, University of St Andrews, United Kingdom

This paper proposes a novel approach to systems modelling based on responsibilities. The approach is designed to help users identify and analyse the hazards and associated risks that can arise in complex socio-technical systems that span organisational boundaries. Through a case study this paper shows how the technique can identify the vulnerabilities that may arise because human, organisational or technological agents fail to discharge the responsibilities assigned to them.

1 INTRODUCTION

Existing risk analysis techniques commonly focus on the interaction of technical aspects of systems. However, we argue that within complex systems socio-technical factors, including the interaction of people with technical components, and the effect of the environment on those interactions provides a different perspective on the risks associated with a system. This view is especially relevant to safety and mission critical systems, where significant consideration is needed to the risks associated with system change and evolution.

For example, elections are considered mission critical to the government, which is bound by law to conduct them at set periods. They are also mission critical to the councils which are bound to collect and collate ballots to acceptable levels of accuracy to the populace. Both can involve significant amounts of technology, but it is the risks of the technology combined with the level of training, and the capabilities of the staff involved that influence the success of any given election process.

We believe *Responsibilities* are a natural form of expression for risk analysis within complex socio-technical systems. We have developed the technique of responsibility modelling, an approach which allows stakeholders to explore the hazards and associated risks of a given system in a structured and logical manner. These models can then be used to mitigate or avoid the risks associated with misunderstandings, and provide support for the analysis of potential process change.

This paper puts forward an approach which augments responsibility modelling with additional hazard and risk data. We argue that this technique is useful in

both the design of new systems and in the analysis, evolution and reassessment of existing systems. The modelling process is one of collaborative working, in that several stakeholders, perhaps not co-located, use the technique to build and analyse responsibility models augmented with additional risk data. Dependent on the domain the stakeholders could be developers, managers, end-users or any stakeholders who have a need to understand and reach a collaborative agreement on either the way a system needs to work, or the way in which it already does.

We have developed tools to provide both a stand alone, and a web based version of the responsibility modelling. These allow multiple users to work collaboratively on editing and analysing a given model. Our tools provide a complete graphical environment to support the description and modification of keywords attached to responsibility model entities and relationships. We support the ability to compare models of the same responsibilities, in a ‘before and after’ style which allows users to benefit from system overviews that can be used to easily identify fundamental changes in a visual and more usable format than written documentation alone.

The structure of the paper is as follows. Section 2 provides background on responsibility modelling. Section 3 reviews prominent risk / hazard modelling techniques. Section 4 explains how we have extended responsibility modelling to encompass the benefits of these techniques. Section 5 provides an overview of our work on evaluating the technique using a case study based on the Scottish Elections in 2007. Finally, section 6 outlines future work in this area and draws conclusions.

2 RESPONSIBILITY MODELLING BACKGROUND

Responsibility modelling has been proposed by several authors as a useful construct for analysing the dependability of socio-technical systems (Blyth et al. 1993; Dobson and Sommerville 2005; Strens and Dobson 1993). In addition to achieving system goals, both social and technical entities contribute to the broader dependability of a system. The notion that human agents in a system, if employed appropriately, can contribute positively to the dependability of a technical system is one that is often missed in discussions of software dependability (Besnard and Baxter 2003; Besnard and Greathead 2003).

For our purposes, we define a responsibility as:

A duty, held by some agent, to achieve, maintain or avoid some given state, subject to conformance with organisational, social and cultural norms.

The term ‘duty’ in this context refers to more than simply a statement that a given task should be completed. It also encompasses aspects of accountability. The phrase organisational, social and cultural norms relates to the inherent nature of responsibilities; that systems are adapted to fit the culture they operate in, that processes have to work within the social framework of both legal and domain standards. Responsibilities are rarely broken down to individual instructions, as they represent higher level constructs encompassing a remit for initiative. Initiative is bounded by professional conduct, from an organisational perspective as well as the wider social and cultural ones.

We use responsibilities within a graphical modelling environment that encompasses Responsibilities, Agents and Resources, connected by relationships. The following sections only provide an overview of the responsibility modelling technique, focusing instead on the addition of Hazard Analysis. For more information on Responsibility modelling in general we recommend reading (Lock et al. 2009).

Graphical models of responsibility were first proposed in the ORDIT methodology (Blyth et al. 1993), a notation for describing the responsibilities that agents hold with respect to one another. Strens, Dobson and Sommerville have argued for the importance of analysing responsibility and the need to view roles with respect to the responsibility relationships that exist between them (Dobson 1993; Dobson and Sommerville 2005; Strens and Dobson 1993). Dewsbury and Dobson (Dewsbury and Dobson 2007) describe much of the research undertaken on responsibility as part of the DIRC project¹, presenting analysis of inappropriate responsibility allocation in socio-technical systems.

¹<http://www.dirc.ac.uk>

Similar in intent, goal based modelling approaches, such as i^* and KAOS are intended to expose high level dependencies between objectives in a given system (Darimont et al. 1997; Yu 2002). Goals can be hierarchical and achieved through the fulfillment of some or all sub-goals.

Despite some similarities, responsibility modelling differs from goal based techniques. Whilst the notion of responsibility modelling may be viewed as incorporating the specification of objectives to be achieved, there is also an acknowledgment that in complex socio-technical systems, the achievement of an objective (i.e. the discharge of responsibility) is subject to a range of constraints and that even with the best efforts of an agent, a goal may not be achieved. These constraints are difficult to explore and model using a goal-based approach which focuses principally on what has to be achieved.

In contrast to goal-based approaches, there are circumstances in which an authority may judge that a responsibility has been appropriately discharged, despite the fact that a associated goal has not been achieved. Woods (Woods 2005) has noted (in the context of accountability and learning in health care organisations) how actors are required to cooperate with regard to individual responsibilities in order for broader organisational responsibilities to be discharged. The notion of responsibility embodies an assumption that it is how an agent acts and not just what is achieved that is important. For example, a doctor who has carried out the correct procedures may have successfully discharged their responsibility for patient care, even though a patient dies.

For the purposes of this paper the entities and relationships we are dealing with are outlined in Figure 1 and match the key below:

Responsibility: A stated activity or abstract concept

Information Resource: For example a report or database

Physical Resource: For example a piece of equipment

Human Agent: For example, an election clerk

Organisational Agent: For example the government

Responsible For: The allocation of an agent to a responsibility

Has: The allocation of resources to agents or responsibilities

Subordinate To: To model organisational hierarchies

Figure 1: Responsibility Modeling Key

Acts As: For example Bob acts as an election clerk

Association: Used to annotate models with relationships of a domain specific type. These could be anything for example, for example to show A cannot occur at the same time as B.

3 RELATED RISK ANALYSIS TECHNIQUES

This section discusses two prominent risk analysis techniques that use keywords to guide analysis. In building a responsibility modelling technique for risk analysis we build upon many of the concepts described here.

HAZOPS (Kletz 1999) is a goal driven method originally developed by ICI for the chemical industry. It focuses on the identification of potential hazards using keywords and associated risks through in-depth descriptions of the system in question, with a focus on technical operability and efficiency. HAZOPS keywords are used to construct tables examining the effect of deviation from the norm for a given process. For example: given a specific deviation for a given process, (something occurring early, late, never, in reverse, to much etc); what are the consequences; what actions could be taken to mitigate the consequences; what safeguards could be put in place; what is the risk of the occurrence of the deviation etc. HAZOPS is applied predominantly at the mechanical, rather than socio-technical systems level. The HAZOPS approach is a recognition that the use of codes of practice and standards can only partly eliminate the risks associated with the implementation of systems, and that many failures are anticipatable and avoidable given appropriate contingency planning.

The key benefits of the HAZOPS approach include;

- The promotion of systematic understanding of all processes and resources within the system, and examination of the environment in which the system operates. This in turn can be used to find

and evaluate hazards and ameliorate risks associated with operation of the system.

- By examining the consequences of different scenarios in a systematic manner it is possible to determine the effect of failures on other parts of a given system.
- HAZOPS has a demonstrable effect (Pully 2004) in reducing the number of ‘snagging’ issues during the running-in of complex technical systems.

HAZOPS is, however, not suitable for the wider socio-technical system domain due to its reliance on completeness, and focus on low level technical activity sequences. This prevents HAZOPS from analysing human behaviour, and also makes it too complex a technique to be used without considerable effort on the part of investigators.

A number of software tools for HAZOPS have been developed (eg: Dyadem²) but these are targeted at chemical applications and are not available on an open source basis.

Hollnagel’s Cognitive Reliability and Error Analysis Method (CREAM) (Hollnagel 1998) is described as a second generation human reliability analysis method, because it unifies best practice from the fields of human reliability and the cognitive sciences. The former is often more concerned with predicting human behaviour, usually in a quantifiable way, whilst the latter is more concerned with understanding and explaining human behaviour with a view to being able to predict it under given circumstances. CREAM can be used retrospectively to analyse the possible causes of an accident that has happened, or prospectively to identify the possible failures that could lead to accidents. In CREAM performance takes place in a context defined by the interaction between three high level factors: the human operator; the technology; the wider organisation (including the physical environment in which the system is located). CREAM categorises erroneous actions using a small set of possible error modes or effects (Hollnagel calls them logical phenotypes): timing, duration, force, distance, speed, direction, object and sequence, each of which can manifest itself in only a few ways. So, for example, within the timing category, an action can be erroneous because it is too early, too late, or omitted.

CREAM can be applied to socio-technical domains, however it relies on sets of predefined criteria and types which have to be followed rigorously to achieve results. Again, the role of CREAM investigator is not one that can be carried out without significant training. As with HAZOPS, there is a lack of tool support to ease the introduction of CREAM to a wider audience.

²<http://www.dyadem.com/products/phapro/>

4 AUGMENTING RESPONSIBILITY MODELLING FOR RISK ANALYSIS

We provide coverage for both entity and relationship types in what we term ‘risk clauses’. These outline the entity involved, its associated hazards and context, combined with the risk of occurrence and severity. More specifically a risk clause defines the following data:

Target The entity / relationship to which the clause refers. For example the entity in question could be a responsibility, or a resource expected to be used. A relationship could include the allocation of a resource to a group for use in a situation.

Hazard Using a restricted set of keywords we aim to focus discussions by giving a clear checklist of hazard categories to consider. The hazard keywords we use are adapted from HAZOPS and are outlined below:

- **Early** Occurrence of the entity/relationship before required
- **Late** Occurrence of the entity/relationship after required
- **Never** Non-occurrence of the entity/relationship
- **Incapable** If the occurrence of the entity/relationship could not take place even though planned
- **Insufficient** Occurrence of the entity/relationship at an incorrect level
- **Impaired** Occurrence of the entity/relationship in an incorrect manner

Condition The potential conditions that could arise with relation to the hazard occurring. Where multiple conditions exist, these should be separated out for individual consideration.

Risk We define risk in this context as a combination of the probability of the hazard and the severity of the hazard occurring. We currently use qualitative terms to allow prioritisation using terms such as low, medium and high.

Consequences The potential effects of the hazard manifesting itself in the wider system.

Recommended Actions The cause(s) of action, either mitigation or avoidance, that could be taken to deal with the situation in question. Whether a given course of action should be taken is tempered by economic, organisational and political factors and as such is not explored in more depth through responsibility modelling. It instead provides a starting point for further deliberations.

Responsibility models can be used to represent existing organisational structures where the allocation of resources and personnel to responsibilities is something which has already occurred. We also believe that responsibility models can be used dynamically to represent evolving situations. For example, a new piece of equipment is supplied to an agent. In this case risk clauses are useful to determine not only the effect of that piece of equipment being used too late, or not used at all, or used incorrectly etc in the furtherance of a responsibility; but also whether the relationship of allocation itself occurs too late or early etc.

Within our models we separate these static and dynamic aspects by associating allocation risk clauses with the relationships, and by associating usage risk clauses with the entities. The difference between the use of the same clause on allocation and use is illustrated by the following example:

The equipment is allocated too late (a clause attached to the ‘has’ relationship between a responsibility/person and a resource)

vs

The equipment is used too late (a clause attached to the ‘resource’ entity for the equipment in question)

Although both may have the same consequence, accountability could differ. A person cannot be held accountable for a situation where they have been allocated the right equipment too late to make use of it. If however they already have the right equipment they may be held accountable for not using it.

5 CASE STUDY

During 2007 our research group acted as observers for the Scottish Elections accredited by the Electoral Commission. The 2007 Scottish elections involved the use of E-Counting, something that had not previously been attempted in Scotland. E-Counting in this instance involved the use of paper ballots that were then machine read, tallied and stored. Parts of the voting process also used STV, a form of voting where preferences are assigned numerically against candidates, again for the first time.

The election itself caused considerable embarrassment to the government. Problems emerged with ambiguities in the new ballot forms; including issues related to the quality of the printing of ballots, which in turn led to many problems in the scanning process. The company responsible for the counting process had no previous experience of the STV system, and did not understand the role of the political representatives at the counts. As such the process that was put

in place was relatively opaque to political and other observers, with limited information on the ongoing count progress provided by LCD screens. The unreliability of the update screens themselves was another cause for embarrassment.

From a socio-technical perspective the election is interesting as it contained several processes, machines and people, many of whom had limited training, which were dealing with new ballots, new processing, and new tallying mechanisms.

Our accreditation allowed us access to preparatory literature, training materials and briefings prior to the election, plus access to the polling stations and counts. We were able to contribute to the discovery of issues relating to the election (Lock et al. 2008), including input to the government Gould Report (Gould 2007). By constructing responsibility models from the data we had available to us we believe many of the problems that occurred could have been discovered and discussed before the event. This section describes some of the models we have constructed that show how responsibility modelling could have helped in collaborative discussions.

One of the issues that faced those working at the count was related to the E-counting machines themselves. The elections used electronic readers to scan normal paper ballots and automatically determine the voter's intentions. Several problems emerged. Figure 2 outlines the responsibilities of three different agents involved in the counting process.

Firstly, the DRS operators (a largely de-skilled position), responsible for refilling the hopper with un-scanned votes and monitoring for problems (most notably paper jams). The diagram uses the 'allocated to' arrow to indicate that the DRS operative has access to a DRS machine to aid them in fulfilling the responsibilities they hold.

Secondly, the DRS technicians who were more highly skilled, and capable of rectifying problems discovered with the counting machines. The diagram states that the DRS technician is 'responsible for' the DRS machine. This is shorthand for inserting an additional responsibility, of maintaining the resource, to which the machines would be 'allocated to', and the technician made 'responsible for'.

Thirdly, the political party officials who were responsible for estimation. Prior to the 2007 election the political officials had provided a check against corruption by performing their own count of votes as the process went on. In a normal election these unofficial counters have managed to tally results to within 5-10% and as such provide an important feature of British elections. The replacement of hand counting by a DRS counting machine however meant that the speed of scanning made hand counting impractical. Interestingly this had not become clear to the politi-

Figure 2: Partial DRS Counting machine responsibility model

cal officials until the process was actually underway. The responsibility model would have illustrated the fact that although the DRS and party officials would not under normal circumstances meet, they needed to be aware that they were both trying to operate on the same resources, albeit from a passive context for the party officials.

The model produced of these interactions could have been used to promote discussion of the strategy for dealing with impaired processing speeds of the counting machines, and the subsequent demands on the overloaded DRS technicians. Figure 4 shows some examples of risk clauses derived from this scenario. Notice that the consequence of one given situation, that of a breakdown (impaired operation) of a given machine makes calls on another resource, the DRS technician, who is therefore under greater load.

Figure 3 outlines the adjudication process that took place on ballots. The adjudication process consisted of those ballots that could not meet the automated counting systems threshold for decision making. First, an adjudication process involving the council authority staff dealt with those ballots that were easily resolvable (denoted as Normal Adjudication in the diagram). These often included ballots where stray lines, caused by the scanning of folded paper produced enough doubt in the system to pass these to adjudication. Second, the returning officer, and their deputy returning officers in consultation with the party officials dealt with those that remained at 'returning officer adjudication'. This process used a separate area containing a workstation and accompanying data projector to allow close observation by multiple people. Both adjudication processes used the 'adjudication station' resources to interface to the 'Ballot Database' (denoted as a information resource on the diagram) which stored information relating to the stage each ballot in the system had reached.

One of the major issues that would have revealed with modelling is that the returning officer, and their

Figure 3: adjudication & database interaction

deputies responsibilities remained broadly the same between manual and E-count elections, whereas their workload increased dramatically. Problems related to folded ballots etc, caused adjudications that the system, and the council staff could not deal with alone. Council staff had no access to the complete ballot, only to parts of it. Hence lines across the sheet were difficult to interpret and dismiss. The resultant load on the returning officers slowed the count in many cases, and led to returning officers and their deputies needing to frequently address different responsibilities in different parts of the count halls (including Adjudication, Supervision and Declaration as indicated on the diagram). Given the risks associated with introducing new technology it could have been anticipated that a heavier than normal load on those in managerial control would result. A more delegated responsibility structure would have been considerably more robust.

The responsibility model also illustrates the inability of the socio-technical system to deal with problems caused by the ballot database, which acted as a single system, integrated into many of the processes involved in the count. Delays in processing queries from the adjudication process, and from other processes caused by database issues contributed to much of the inefficiency seen at the counts. At one count centre problems with the database caused the count to be closed down temporarily for repairs, reopening the next day. In this case the processing of large numbers of adjudications prevented the system from keeping up with ballot indexing, eventually causing the system to crash. Given that the ballot database was a critical resource relatively little consideration was given

to it in preparatory material or in the discussions on deployment of the E-Counting system beforehand.

Figure 5 shows a subset of the risk clauses associated with the returning officer / deputies and the ballot database. In particular, this example shows the importance of context to the analysis of risk, as different levels of failure were observed during the process. The result of persistent overloading of both returning officer / deputies and the ballot databases shows that the procedure for dealing with low levels of occurrence were not in place, and that nothing could be done at the time to mitigate the problems caused by higher levels of occurrence.

Figure 6 shows the risks associated with the process of adjudication itself. Both party officials and local authority / returning officers had the responsibility of overseeing the stages of adjudication, but the speed of adjudication brought on by the introduction of an electronic system impaired the process. This occurred as those operating the adjudication terminals became familiar with the motions they needed to go through. This led to an increase in speed which led to party officials lacking the time necessary to consult, or in some cases even think about given cases. This was compounded by system design decisions which prevented operators from going back more than one adjudication to make a change. As the speed of adjudication ran at around 2 seconds per case party officials often lost the ability to query decisions before they could stop the process. As such, the two safety factors in adjudication, having two operators, and having party oversight failed due to the unforeseen consequences of speeding up the process. Throughout the count issues relating to the relegation of political representatives was prevalent, and it was clear that insufficient consideration had been given to their role at the design stages.

6 CONCLUSIONS AND FUTURE WORK

This paper has proposed a novel approach combining existing notions of responsibility modelling with hazard / risk based data to allow improved analysis of socio-technical systems. In doing so we have developed a technique that goes beyond the capabilities of many of the existing technically focused analysis techniques, and applied this in a topical case study.

We are currently developing tools to analyse the data stored by our models. In particular we are interested in highlighting critical entities and relationships. Such an approach could for example provide visual cuing of the most connected processes, agents and resources. We believe visualisation of such concepts, in a similar form to that currently performed within many network analysis tools could be used to further reinforce the recognition of potential trouble spots in a model. We are also developing ‘what if’

Target	Hazard Keyword	Condition	Consequence	Risk Li / Sev	Recommended Actions
DRS Machine	Impaired	DRS Machine Breakdown.	DRS technician required	High / Low	Improve machine reliability
DRS Technician	Late	Machines out of service.	Count slowed	Low / Medium	Provide more technicians
Estimation	Incapable	Party official redundant to this process	Decreased oversight of election	High / High	Slow down scanning or remove requirement

Figure 4: DRS machine analysis

Target	Hazard Keyword	Condition	Consequence	Risk Li / Sev	Recommended Actions
Returning Officer	Insufficient	Returning officer not available to deal with a query	Potential slow down in other activities involving returning officers	High / Low	Delegate more to deputies
Returning Officer	Insufficient	Returning officer unable to publish results by legal deadline	Failure of the process	Low / High	Increase the number of deputies next time
Ballot Database	Insufficient	Slow down in database interaction	Increases returning officer load	High / Low	Raise awareness of database loading with staff (to reduce calls to a minimum)
Ballot Database	Insufficient	Database crash	Failure to publish results by legal deadline	High / Low	Increase database system capabilities next time

Figure 5: adjudication and Database interaction analysis

Target	Hazard Keyword	Condition	Consequence	Risk Li / Sev	Recommended Actions
Local Authority Staff	Impaired	Staff do an adjudication too quickly	Incorrect level of oversight of parts of the adjudication process	High / Low	Increased supervision of activities
Local Authority Staff	Impaired	Staff do adjudications too quickly	Significant failure of adjudication the process	Low / High	Halt count if necessary
Party Officials	Impaired	Party officials cannot effectively monitor that adjudication	Incorrect level of oversight of parts of the adjudication process	High / Low	Increased supervision of activities
Party Officials	Impaired	Party officials cannot effectively monitor adjudications	Significant failure of adjudication the process	Low / High	Halt count if necessary

Figure 6: adjudication oversight analysis

style capabilities to allow users to simulate responsibility failures and understand the consequences across a model of a given scenario. Such functionality would be useful in simulated exercises of the type used in the civil emergency domain in order to stress plans to eliminate problems and improve training.

REFERENCES

- Besnard, D. and G. Baxter (2003, November). Human compensations for undependable systems. Technical Report CS-TR-819, School of Computing Science, Newcastle upon Tyne.
- Besnard, D. and D. Greathead (2003). A cognitive approach to safe violations. *Cognition, Technology & Work* 5(4), 272–282.
- Blyth, A. J., J. Chudge, J. E. Dobson, and M. R. Strens (1993). ORDIT: A new methodology to assist in the process of eliciting and modelling organisational requirements. In S. Kaplan (Ed.), *Proceedings on the Conference on Organisational Computing Systems*, Milpitas, California, USA, pp. 216–227. ACM Press.
- Darimont, R., E. Delor, P. Massonet, and A. van Lamsweerde (1997, May). GRAIL/KAOS: an environment for goal-driven requirements engineering. In W. R. Adrion (Ed.), *ICSE'97: Pulling Together, Proceedings of the 19th International Conference on Software Engineering*, Boston, Massachusetts, USA, pp. 612–613. ACM Press.
- Dewsbury, G. and J. Dobson (Eds.) (2007, June). *Responsibility and Dependable Systems*. Springer-Verlag London Ltd.
- Dobson, J. (1993). New security paradigms: what other concepts do we need as well? In *NSPW '92-93: Proceedings on the 1992-1993 workshop on New Security Paradigms*, Little Compton, Rhode Island, United States, pp. 7–18. ACM Press.
- Dobson, J. E. and I. Sommerville (2005, October). Roles are responsibility relationships really.
- Gould, R. (2007). Independent review of the scottish parliamentary and local government elections. Technical report, Commissioned by the Electoral Commission.
- Hollnagel, E. (1998). *Cognitive Reliability and Error Analysis Method: CREAM*. Elsevier.
- Kletz, T. (1999). *HAZOP and HAZAN-Ref to Icheme*. CRC Press.
- Lock, R., I. Sommerville, and T. Storer (2009). Responsibility modelling for civil emergency planning. awaiting publication. *Reliability Engineering and System Safety Journal*.
- Lock, R., T. Storer, N. Harvey, C. Hughes, and I. Sommerville (2008). Observations of the scottish elections 2007. *Transforming Government: People, Process and Policy*, vol 2. *Emerald Insight* 2008.
- Pully, A. (2004). Utilization of hazard and operability studies in a petroleum refinery. *Process Safety Progress* 12(2), 106–110.
- Strens, R. and J. Dobson (1993). How responsibility modelling leads to security requirements. In *NSPW '92-93: Proceedings on the 1992-1993 workshop on New security paradigms*, New York, NY, USA, pp. 143–149. ACM Press.
- Woods, D. D. (2005). Conflicts between learning and accountability in patient safety. *Depaul Law Review* 54, 485–505.
- Yu, E. S. K. (2002, May). Agent-oriented modelling: Software versus the world. In M. Wooldridge, G. Weiß, and P. Ciancarini (Eds.), *AOSE*, Volume 2222 of *Lecture Notes in Computer Science*, Montreal, Canada, pp. 206–225. Springer.