Preventing Elderly from Falls: The Agent Perspective in EPRs

Sebastian Ahrndt, Johannes Fähndrich, and Sahin Albayrak

DAI-Labor, Technische Universität Berlin, Germany
Ernst-Reuter-Platz 7
10587 Berlin, Germany
sebastian.ahrndt@dai-labor.de (Corresponding author)

Abstract. This work presents an approach combining multiple electronic patient records (EPR) to a self-learning fall risk assessment tool. We utilized the agent-perspective to model the system, to address privacy issues and to evaluate different distributed information fusion and opinion aggregation techniques towards there applicability to the addressed domain. Each agent represents a single patient negotiating about unknown fall risk influences in order to adapt the fall-risk assessment tool to the population under care. In addition, we will outline the planned real-world case study.

1 Introduction

Injuries or disabilities contracted by falls constitute not only as a cause of suffering for the elderly, but also lead to substantial economic burdens. For example, the total amount of fall related costs amount to 0.85% - 1.5% of the yearly German health care costs [14]. Here, the increasing development of electronic patient records (EPR) offers new opportunities for the prevention of falls as EPR contain a treasure trove of data [10]. Although, professionals gain benefits from the consolidation of such data sources, the available feature space is to huge to conceive for humans (e.g. the ICD-10 comprise more then 68,000 codes). In order to not overwhelm the user, EPRs provide standardized tools – such as fall-risk assessment tools – utilising a subset of the feature space in order to assist the user. In this work, we want to introduce such a fall-risk assessment tool which is agent-based and self-learning in order to reveal some kind of personal healthcare based on huge amount of data. To start with, a literature research done to identify well-researched fall-indicators and their influences on patients fall-risk were conducted [21]. Here, we were able to identify 25 fall-risk indicators. However, during the course of this work we applied the agent perspective to negotiate about unknown fall-risk influences in order to adjust the fall-risk assessment tool to the population under care. This requires population-based data, which can not be provided by studies as those conceived during the literature research as it would be impractical to conduct as much studies as populations exist [10]. Nevertheless, studies still outline the starting point for personalised health care. Consequently our approach utilised the identified fall-risk indicators and their

influences as part of the initial knowledge of each agent – where each agents represents a single patient (See Section 2). Further, we applied the results of the literature review as part of the evaluation of our approach (See Section 3) and as default version of the assessment tool used during the planed case study (See Section 3.1).
The contribution of the work is threefold. To start with, we underline the applicability of contemporary agent-frameworks within EPR, due to their support for data privacy mechanisms. Germany in particular has strict rules when it comes to patient data, as such, privacy issues are not considered a nice extra, but required by law. Further, we show that constituting features of agents (e.g. the sensor-effector metaphor or cooperation) can be utilized to adjust EPR components to the population under care pro-actively. Subsequently, we compare different pooling methods for distributed information fusion and opinion aggregation in order to identify ways in which such adjustments can be established and, further, to identify the most-fitting method for the addressed domain.

2 Approach

In order to obtain the goal of personalized health-care there exist the need to observe population-based data and to utilise the observation results to adapt the original conceived health-care [10]. One can imagine, that this is a common procedure in the daily routine of heath-care professionals but a rather hard task for computers. During the course of this work, we want to outline an agent-based approach which achieves this for a fall-risk assessment tool as part of a mobile EPR [1]. Such tools enable health-care professionals to determine the individual fall-risk of the patients based on several fall-risk indicators and as a consequence to initiate suitable retaliatory actions. The basic idea of the intended approach is to enable the patients to negotiate about arising fall-risk indicators which are not yet available in the assessment tool. Here, each health-care professional takes care of several patients. We utilised the agent-based approach to model such a system as illustrated in Fig. 1. Consequently each patient is represented by a single agent, further, multiple patient agents run on a single platform representing a health-care professionals tablet. The whole environment consists of several tablets. However, Fig. 1 also illustrates that the negotiation process consists of four different stages to determine whether an adaption is necessary or not: The occurrence of a fall, the local information fusion, the global opinion aggregation and the notification stage. In the following we will explain each of the stages in detail. Afterwards we will present a comparison of different distributed information fusion techniques and outline the most fitting one for the addressed problem.

2.1 Elaborating Fall-Risk Indicators

Whenever a patient happens to fall, the incident is documented by the health-care professional. The health-care professional adds the observed fall-event $d \in D$

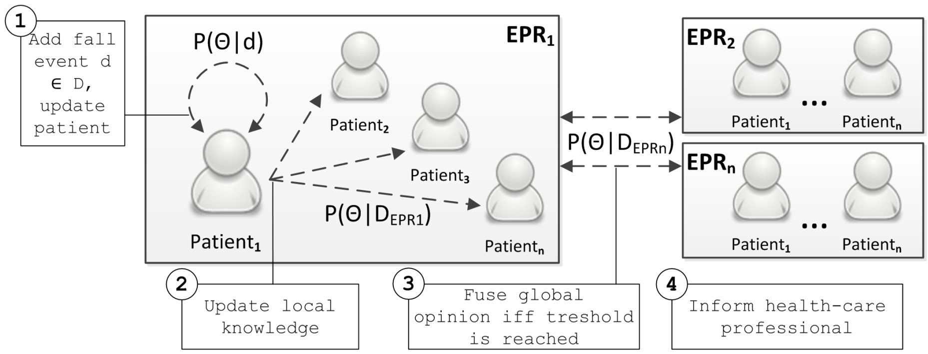

Fig. 1. Illustrating the approach where the patient agents at a single EPR negotiate which feature needs to be evaluated in negotiation with other EPRs and its appropriated patient agents. The process starts with the occurrence of a fall-event where the appropriated patient agents updates the probabilities of all features conceived (1). Afterwards the agent aggregates its new personal opinion with the prior local opinion of the node and propose the result to the other patient agents available through the node (2). If one feature reach a given threshold the agents aggregates the global opinion requesting all local opinion of the available nodes (3). If the threshold of the feature is still exceeded the EPR notifies the health-care professional about the new perceived fall-risk indicator (4).

to the EPR of the associated patient agent. With each observation d, new evidence is collected to elaborate possible fall-risk indicators in the feature space Θ, where Θ are all the quantities of interest on which the group wants to elaborate an opinion. In our domain the Θ represents the set of impact factors on a patients fall-risk, with $\Theta = \{\theta_i | i = 1..n\}$. For each patient the patients agent $p \in P$ elaborates at each observation the evidence for all fall-risk factors. The agent then updates its believe by using Bayesian information fusion [15] as illustrated by Eq. 1.

$$p(\Theta|d) = \frac{p(d|\Theta)p(\Theta)}{p(d)} \propto p(d|\Theta)p(\Theta) \qquad (1)$$

The a-priori $p(\Theta)$ represents the prior knowledge of the quantities of interest. As mentioned above we received this prior knowledge from a literature overview. Here we might exclude e.g. surnames of a patient record to be a fall-risk factor by setting $p(\theta_{surname}) = 0.0$. The likelihood $p(d|\Theta)$ represents the impact of the observation of d on the belief state of the agent. The a-posteriori $p(\Theta|d)$ can be calculated with these two inputs [2]. In this work the probability distribution $p(\Theta|d)$ can be seen as an experts opinion. As the observations $d \in D$ are assumed to be conditional independent, the expert opinion represents the 'degree of believe' of a single patient agent [11]. Each expert $p \in P$ has an degree of believe on which fall-risk factor has influenced the fall observed in the observation d and formulates this opinion in $p(\Theta|d)$. After each observation the expert updates its believe over Θ. One can easily imagine the vast amount of update functions possible. For example, Eq. 2 represents one update function which whether doubles the θ_i if the feature changed between this observation and the last one or halved it otherwise.

$$f(\theta_i) = \begin{cases} 2\theta_i & \text{if } d^t = d^{t+1} \\ \frac{1}{2}\theta_i & \text{if } d^t \neq d^{t+1} \end{cases} \tag{2}$$

With such an update function the likelihood of each feature can be calculated as shown in Eq. 3.

$$p(d|\theta) = \frac{f(\theta_i)}{\sum\limits_{1}^{n} f(\theta_i)} \tag{3}$$

The fusion of the new evidence with each preceding observation is then integrated into the agents believe by using Bayesian information fusion as show in Eq. 4.

$$p(\Theta|d^{t+1}) \propto p(d^t|\Theta)p(\Theta|d^t) \tag{4}$$

After the agent has a new degree of believe $p(\Theta|d^{t+1})$ the first stage is completed. Now the local information fusion updates the group opinion of all patient agents on the same node. Here, in contrast to the update function of a single agent, if between agents the same θ is found, its impact factor should increase to exchange the effects of Eq. 2. Since the communication on one tablet is cheap and secure, the agents on each tablet are enabled to reassess the fall-risks after each observation of a fall locally. We need to introduce this distinction between intra and inter Tablet communication because the region the health-care professionals are working in is not fully covered with wireless Internet connections.

After the probability of a possible fall-risk indicator exceeds the experimentally established threshold shown in Eq. 5 the global opinion aggregation starts. Here the patient agents on one tablet fuse their node opinion with the more global opinion of all devices using the same method as in the local information fusion phase. If the threshold of the feature is still exceeded the suggested feature is forwarded to the health-care professional which has to decide whether this feature is a fall-risk indicator for the population under care or not. The result of this decision is communicated to all available nodes and if necessary updates the fall-risk assessment tool with the new feature. Each θ which was handled by a health-care professional is removed from the opinion elaboration by setting $p(\theta_k) = 0.0$. A side effect of the distributed information fusion is anonymisation. Since the probability distribution has been created through the aggregation of multiple opinions the communication between the different tablets during the global opinion aggregation becomes secure. The messages contain only the probabilities but not the values of the features as the values are not needed anymore at this stage.

$$\frac{1}{|\Theta \setminus \{\theta_k | p(\theta_k) = 0.0\}|} \tag{5}$$

Now each patient agent represents some kind of fall-risk expert which is able to express its opinion about the fall-risk influences of its patient. Following *Roback* and *Givens* the arising issue 'is to pool opinions in a rational way that yields a single probability distribution from which inference can be made' [19]. Here, we might use methods of information fusion to establish a group opinion [22].

2.2 Aggregating the Group Opinion

As usual in the information fusion we need some kind of aggregation method to fuse these local expert opinions together to a generalized view [23]. Since there are multiple options, we have to choose an appropriate method (a so called pooling method) to combine multiple opinions into a group opinion. Pooling methods might represent different voting strategies like an dictatorship of one opinion, a democracy where every agents opinion has the same weight or an strategy which is based on the reliability of the experts. The interested reader is referred to *Faehndrich* [4].
However, to see how those pooling methods work, we will look at some of them, and evaluate their usefulness in our multi-agent system. Some of the best known pooling methods are the 'Linear Opinion Pool' [7] (LinOp), the 'Logarithmic Opinion Pool' [9] (LogOp) and the 'Supra-Bayesian Pooling Method' [8]. Each of them profits from a growing body of evidence [3]. There are two ways on evaluating a pooling method: On the one hand, we can evaluate their theoretical properties like a 'non dictatorship' or an 'unanimity' [18]. On the other hand one can measure their performance in a real world example. In this work, we want to evaluate how those theoretically well-researched pooling methods can be used in the real world task of elaborating fall-risks as a group decision in a multi-agent system.

Table 1. Classification of the examined pooling methods and there applicability to the domain we address. With n being the amount of patient agents and m being the number of nodes (devices) available.

Method	**Space**	**Time**	**Communication**
Supra-Bayesian	$O(1)$	$O(1)$	$O(m+n)$
Linear Opinion Pool	$O(m*n^2)$	$O(m*n^2)$	$O(m+n)$
Logarithmic Opinion Pool	$O(m*n^2)$	$O(m*n^2)$	$O(m+n)$

Table 1 classifies the examined pooling methods utilizing several criteria. Here we point out the time, space and communication complexity of the pooling methods. Although, the opinion aggregation using LinOP and LogOP requires the mean value and therefore requires to aggregate the opinion of every agent and node depending on the voting round, the communication complexity of all pooling method is equal. This is due the fact, that using the Supra-Bayesian the agents have to broadcast the new aggregated opinion to all other agents and nodes as illustrated in Fig. 1. However, the time and space complexity differs. Since we have two voting rounds, one local on the node and another one between nodes, the space and time complexity rises quadratically using LinOP or LogOP. Here, the Supra-Bayesian Pooling Method can reuse the last available pooling result generated in the prior voting round to update the node and/or group opinion. This reduces the time and space complexity to be linear.

3 Evaluation

In order to evaluate the described approach we developed a prototype using the agent framework JIAC V [12]. Here, we decided to conduct a simulation in order to test the algorithm. Each patient agent where set-up with the initial set of well-researched fall-risk indicators which we received from our literature research [21]. All utilized features such as age, sex or diseases are available in contemporary EPR and can be ordered through their impact on the fall-risk of a patient. For example, a physical deficit in the lower extremities influences a patients fall-risk more than the age of the patient (factor 4.4 vs. 1.7) [20]. We implemented the patient agents as described in Section 2 and migrated the data model of the EPR developed within the agnes^{zwei} project [1]. Further, we implemented an agent (SimAgent) which simulates the health-care professionals role. The SimAgent adds a fall-event each 100ms to the multi-agent system and decides whether a suggested feature is a fall-risk indicator for the population under care or not. As fall-events do not occur randomly the SimAgent decides which patient should fall based on the patients current fall-risk. Further, the SimAgent changes the features of the patient record as it would be during the use of the EPR system. The simulation of the aging process of a patient is based on the research results of prior studies with risk-equivalent patients [5, 6, 13]. We utilized this behavior to evaluate the system. In the following we will discuss the simulations result of each of the three implemented pooling methods. For each pooling method we observed 50 features during the first 80 voting rounds and carried out several simulation runs.

Despite that the simulation always starts with the same initial set of patients, the aging process is non-deterministic and varies between the simulation runs. Therefore the subsequently presented figures can not be compared directly.

Fig. 2 illustrates the probability distribution for the LinOP method. To aggregate the opinions of the patient agents LinOP requires weights for each expert. During the simulation we set the weights to $\frac{1}{|Agents|}$ which can be interpreted as a fair democracy. Here we observed that the probability of the observed features grows slowly during the first 80 voting rounds. LinOP did not exceed a maximum of 0.14 for a single feature making it necessary to observe a greater number of falls. In our real-world problem we do not expect this high rate of falls.

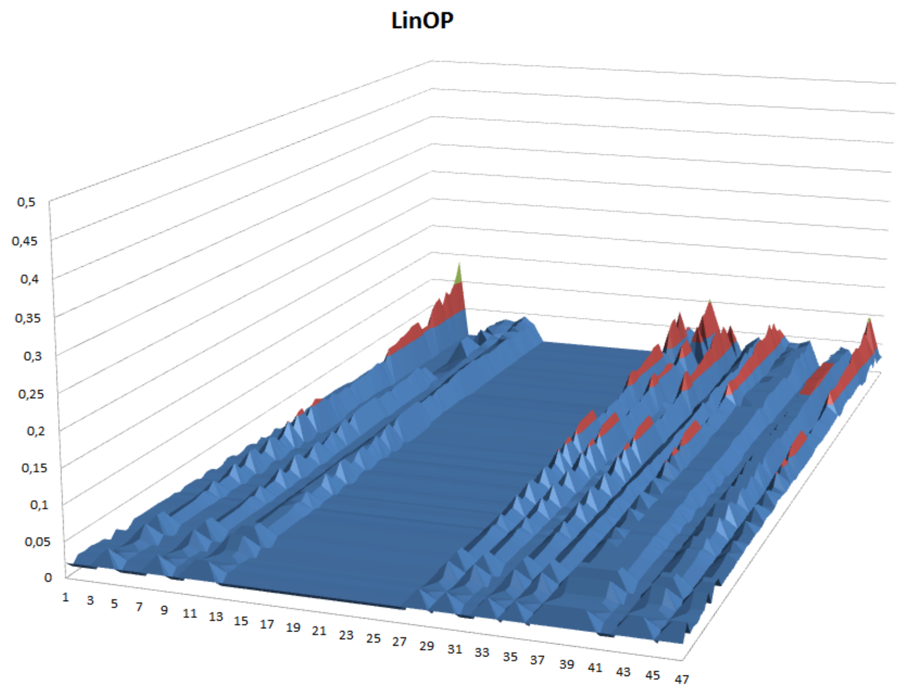

Fig. 2. Progress of the probability distribution using the LinOP method. The x-axis shows the feature, the y-axis the probability and the z-axis the voting round.

Fig. 3 illustrates the probability distribution for the LogOP method. For the simulation with LogOP we applied the same weights as used in the LinOP simulation. During the first voting rounds LogOP produces similar results as the LinOP does. The group decision behavior changes at the moment the first features reach a zero probability. Since with a growing body of evidence single features get excluded. One must notice, that in the case of LogOP a single expert vote with $p(\theta) = 0.0$ suffices to exclude a single feature disregarding all other expert opinions. This is based on the multiplicative nature of the pooling method making single feature disproportional likely.

Fig. 4 illustrates the probability distribution for the Supra-Bayes Pooling method. Using Supra-Bayes it is difficult to decide which likelihood function to use after an observation d. For the simulation we applied the arbitrary function shown in Eq. 2. Supra-Bayes requires no weights presenting a rational way of pooling a democratic group opinion. In contrast to LogOP no single expert can use its vote to disregard all other votes. In addition, Supra-Bayes profits much faster from a growing body of evidence. The first feature exceeds the threshold after only a few voting rounds. This can be seen as advantage and disadvantage. For

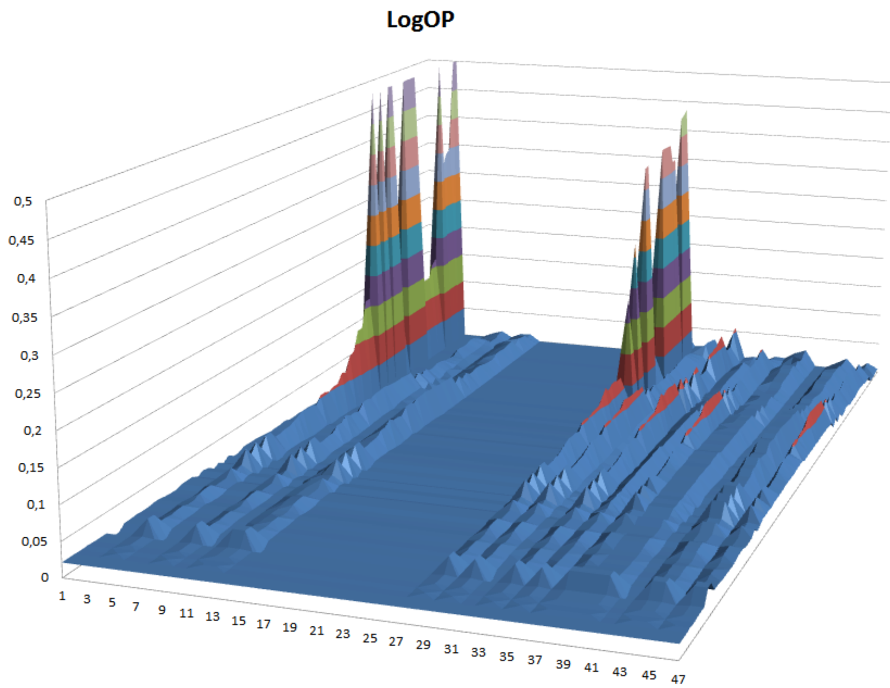

Fig. 3. Progress of the probability distribution using the LogOP method. The x-axis shows the feature, the y-axis the probability and the z-axis the voting round.

our real-world problem we profit from the fast convergence since less falls have to be observed.

3.1 Field Test

The simulation results presented above showed that our approach seems valid and is applicable to the problem it addresses. In addition, we are currently facing the launching date of the real-world field test. The first phase of the field test includes six home-visiting nurses engaged in the wide-spreaded area of Brandenburg, Germany. This nurses will treat between 180 to 300 patients which are at least 65 years old and multi-morbid (more then three diseases). The first phase is scheduled for three month and we expect a noticeable number of fall events as prior studies with risk-equivalent patient groups show that the addressed patients have a rather high fall-risk (up to 60% per year [5]). As this rate only addresses the self-reported fall events of the patients it will be interesting to see how many falls actually occur.

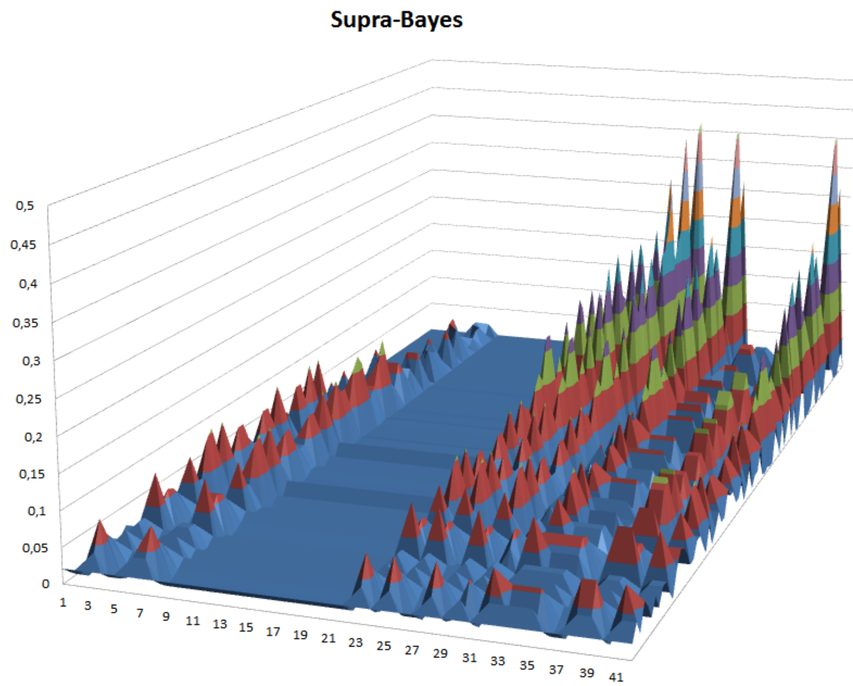

Fig. 4. Progress of the probability distribution using the Supra-Bayesian method. The x-axis shows the feature, the y-axis the probability and the z-axis the voting round.

3.2 Selecting a Pooling Method

To conclude, for our approach we finally decided to use the Supra-Bayesian method through the following reasons.

During our simulation run we analysis the first 80 voting rounds where each voting round was triggered by a single fall observation. During the field test we expect less falls. As mentioned above the health-care professional will treat between 180 to 300 patients. Therefore, we can calculate that we will observe approximately 30 to 50 falls. Only the Supra-Bayesian Pooling method is able to produce meaningful results under this circumstance. Hence, one reason is the expected fall-rate. Another reason is the space and time complexity as the addressed platform are tablets and the system should act resource-efficient. Even the communication complexity remains equal for all considered pooling methods, the Supra Bayesian method outperforms the others in time and space complexity here (as shown in Tab. 1). In addition, we can underline the finding of *Roback* and *Givens* [19] that the implementation of LinOP and LogOP in real problems was not simple as weights for the group members or the likelihood function have to be chosen. As this might be an interesting opportunity for future work (weighting of experience), this effort is not necessary for the current field test.

4 Conclusion

In this paper, we introduced an agent-based approach for a self-learning fall-risk assessment tool. This tool as part of a mobile EPR enables health-care professionals to determine the individual fall-risk of the patients based on several fall-risk indicators. Usually this fall-risk indicators are based on studies which are not able to capture population based data. In order to adapt the assessment tool to the population under care, we enabled the patients to be part of the elaboration of arising fall-risk indicators. Here each patient is represented by a single agent where each mobile EPR contains multiple agents. The whole environment consists of multiple EPRs. In order to aggregate all opinions we introduced an approach consisting of four steps: The occurrence of a fall, the local information fusion, the global opinion aggregation and the notification stage. During the local information fusion and the global opinion aggregation each patient votes as a fall-risk expert. Here, we applied three different information fusion methods to pool a rational single opinion, where a single opinion represents a probability distribution over all features in the feature space. The evaluation of the fusion techniques emphasises that one method outperforms the others under the special circumstance of our application. Here, we showed that the fall frequency we expect is not large enough for the other methods to produce meaningful results. Also we require a method which is resource-efficient in time and space complexity as the addressed platform are tablets.

4.1 Future Work

The presented approach utilises a great amount of information available through the EPR. However, at this point of development the information fusion is limited to data sources which can be easily computed by the agents. Even the software engineers and designers try to standardize and automate most of the binary input in EPRs, there will be always the need to provide free text fields for the health-care professionals as the addressed working environment is to manifold to normalize it. Hence, in the next stage we want to make accessible the full potential of EPRs using Natural Language Processing techniques to evaluate the free text fields. Furthermore, we plan to integrate external influences into the risk decision process. Which means that we want to examine if weather conditions or the actual season of the year also influence the fall risk of elderly. Another interesting research focus in information fusion for EPRs is the integration of sensor data into health records (e.g. *Mohomed et al.* [16] and *Moulton et al.* [17]). Although the aggregation of sensor data is not in the focus of this work it provides interesting aspects for the future with upcoming Bluetooth 4.0 health devices. Extending the research on pooling methods, different pooling methods have to be evaluated as well as their parameters. Analyzing different update function and impact on the learning algorithms will be further focused.

References

1. Ahrndt, S., Rieger, A., Albayrak, S.: Entwicklung einer mobilen elektronischen Patientenakte für die ambulante Versorgung in ländlichen Regionen (development of a mobile electronic patient record for ambulatory health care in the countryside). In: Goltz, U., Magnor, M., Appelrath, H.J., Matthies, H., Balke, W.T., Wolf, L. (eds.) INFORMATIK 2012. pp. 1167–1181. No. 208 in Lecture Notes in Informatics, Gesellschaft für Informatik, Braunschweig, Germany (2012)
2. Bernardo, Jos M. ; Smith, A.F.M.: Bayesian theory. Wiley series in probability and statistics, Wiley, Chichester [u.a.], repr. edn. (2004), http://www.ulb.tu-darmstadt.de/tocs/133503178.pdf
3. Cooke, R.: Experts in Uncertainty: Opinion and Subjective Probability in Science. Oxford University Press (January 1991)
4. Fähndrich, J.: Analyse von Verfahren zur Kombination von Expertenwissen in Form von Wahrscheinlichkeitsverteilungen im Hinblick auf die verteilte lokale Bayes'sche Fusion. Diploma thesis, Karlsruhe Institut of Technology (May 2010)
5. Fiss, T., Dreier, A., Meinke, C., van den Berg, N., Ritter, C.A., Hoffmann, W.: Frequency of inappropriate drugs in primary care: Analysis of a sample of immobile patients who received periodic home visits. Age and Ageing 40(1), 66–73 (September 2010)
6. Fiss, T., Ritter, C.A., Alte, D., van den Berg, N., Hoffmann, W.: Detection of drug related problems in an interdisciplinary health care model for rural areas in germany. Pham World Sci 32(5), 566–574 (July 2010)
7. Genest, C.: Pooling operators with the marginalization property. The Canadian Journal of Statistics/La Revue Canadienne de Statistique 12(2), 153–163 (1984), http://www.jstor.org/stable/3315179
8. Genest, C., McConway, K.K., Schervish, M.M.: Characterization of externally bayesian pooling operators. The Annals of Statistics 14(2), 487–501 (1986), http://www.jstor.org/stable/2241231
9. Genest, C., Weerahandi, S., Zidek, J.: Aggregating opinions through logarithmic pooling. Theory and Decision 17, 61–70 (1984), http://www.springerlink.com/index/XT152131QJ43G001.pdf
10. Goth, G.: Analyzing medical data. Communications of the ACM 55(6), 13–15 (June 2012)
11. Halpern, J.: From statistical knowledge bases to degrees of belief: an overview. In: Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. p. 113. ACM (2006), http://portal.acm.org/citation.cfm?id=1142367
12. Hirsch, B., Konnerth, T., Heßler, A.: Merging agents and services – the JIAC agent platform. In: Bordini, R.H., Dastani, M., Dix, J., Amal, E.F.S. (eds.) Multi-Agent Programming: Languages, Tools and Applications, pp. 159–185. Springer (2009)
13. Hoffmann, W., van den Berg, N., Thyrian, J.R., Fiss, T.: Frequency and determinants of potential drug-drug interactions in an elderly population receiving regular home visits by GPs – results of the home medication review in the AGnES-studies. Pharmacoepidemiology and Drug Safety 20(12), 1311–1318 (December 2011)
14. König, H.H.: Gesundheitsökonomische Aspekte der Sturz- und Frakturprävention (March 2012), http://www.sturzpraevention2012.de/programm_abstract.php?No=9, last visit: 29.10.2012
15. Lindley, D.: Theory and practice of bayesian statistics. The Statistician 32(1), 1–11 (1983), http://www.jstor.org/stable/2987587

16. Mohomed, I., Misra, A., Ebling, M., Jerome, W.: HARMONI: Context-aware filtering of sensor data for continuous remote health monitoring. In: Sixth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom 2008). pp. 248–251. IEEE Computer Society (March 2008)
17. Moulton, B., Chaczko, Z., Karatovic, M.: Data fusion and aggregation methods for pre-processing ambulatory monitoring and remote sensor data for upload to personal electronic health records. International Journal of Digital Content Technology and its Applications 3(4), 120–127 (December 2009)
18. Pennock, D., Wellman, M.: Graphical models for groups: Belief aggregation and risk sharing. Decision Analysis 2(3), 148–164 (2005), http://ai.eecs.umich.edu/people/wellman/pubs/decisionanalysis05.pdf
19. Roback, P.J., Givens, G.H.: Supra-bayesian pooling of priors linked by a deterministic simulation model. Communications in Statistics – Simulation and Computation 30, 447–476 (February 2007)
20. Rubenstein, L.Z., Josephson, K.: Falls and their prevention in elderly people: What does the evidence show? Med Clin North Am. 90(5), 807–824 (September 2006)
21. Schenk, A., Ahrndt, S., Albayrak, S.: Predicting fall risks in electronic patient records. In: Goltz, U., Magnor, M., Appelrath, H.J., Matthies, H., Balke, W.T., Wolf, L. (eds.) INFORMATIK 2012. pp. 1194–1198. No. 208 in Lecture Notes in Informatics, Gesellschaft für Informatik, Braunschweig, Germany (2012)
22. Torra, V., Narukawa, Y.: Modeling decisions: information fusion and aggregation operators. Springer (2007)
23. Wark, E.B.J.S.: Concepts, Models, and Tools for Information Fusion. Artech House (2007), http://www.amazon.com/Concepts-Models-Tools-Information-Fusion/dp/1596930810