Towards a Knowledge (Experience)-Based Recommender System for Crisis Management

Elsa Negre

Towards a knowledge(experience)-based recommender system for crisis management

Elsa Negre
LAMSADE, Université Paris-Dauphine, France
Email: elsa.negre@lamsade.dauphine.fr

***Abstract*—An early warning system can be defined as a chain of information communication systems comprising sensor, detection, decision, and broker subsystems, in the given order, working in conjunction, forecasting and signaling disturbances adversely affecting the stability of the physical world; and giving sufficient time for the response system to prepare resources and response actions to minimize the impact on the stability of the physical world.**

In this paper, we present a framework for a recommender system for crisis management.

This framework uses the actions already implemented to manage former crises to enhance the management of a given crisis. The main idea is to recommend the actions already implemented in those former crises that are similar (the similarity between two crises is based on some indicators such as the gap (hurricane, tsunami, ...)) as the actions to be implemented.

Finally, this paper proposes to exploit the knowledge gained from past experiences to make the best decision (i.e. the best actions to implement) in order to better manage a crisis ready to occur.

***Keywords*-Knowledge, Experience, Recommender systems, Crisis management, Early warning systems, Decision support.**

I. Introduction

Early warning systems are contextualized and used in many fields such as school [1], finance [2], [3], environment [4], [5], humanitarian [6], [7], [8], biology [9], ...

An early warning system can be defined as a chain of information communication systems comprising sensor, detection, decision, and broker subsystems, in the given order, working in conjunction, forecasting and signaling disturbances adversely affecting the stability of the physical world; and giving sufficient time for the response system to prepare resources and response actions to minimize the impact on the stability of the physical world [10]. Thus it is a set of tools to predict disasters, dropouts, ... [11], [12].

Early warning systems face two fundamental problems. First, there is the informational problem of obtaining both the necessary quantity and quality of intelligence in a reliable and timely fashion. Second, there is the analytic problem of avoiding misperception or other faulty analysis (likelihood of diffusion and/or escalation of the conflict, potential risks, ...) [7].

In fact, [13] explains that "Natural hazards, such as storms, droughts, volcanic eruptions, or earthquakes, become disasters only if a community or population is exposed to the natural hazard and cannot cope with its effects. Torrential rain in the middle of an ocean will not cause a disaster, but the same heavy rainfall on a vulnerable population say a shanty town on the side of a hillside stripped of trees - may result in landslides and a huge loss of life. A minor drought may cause a famine if a region's agricultural production is already highly stressed by civil war. A community that lacks an early warning system for volcanic eruptions will be devastated when volcanic ash clouds bear down upon them. Vulnerability is the potential additive that mixes with natural hazards to cause disasters".

[11] points out the fact that former experiences have to be taken into account in order to improve the decisions to make or the actions to implement as part of the crisis management, in particular via the use of early warning systems.

Consequently, in this paper, we propose to put in place a recommender system based on former experiences / knowledge in order to guide the decision-makers when warnings appear, i.e., to recommend actions to implement / decisions to make using the ones implemented / made during former similar situations (circumstances).

The paper is organized as follows: Section 2 presents related work, Section 3 introduces our generic framework. We discuss future work in Section 4.

II. Related work

Early warning systems are specific to the context / field for the one they are implemented. However, eight early warning system characteristics have been defined [14]: continuity in operations, timely warnings, transparency, integration, human capacity, flexibility, catalysts, apolitical. The ones that are interesting in our context are:

- flexibility: this characteristic can be improved,
- integration: an early warning system is a part of a whole, some other parts can be added, for example, a recommender system,
- apolitical position: we propose to recommend former actions that have maybe been implemented by different political parties.

According to [11] and [13], a complete and effective early warning system comprises four elements: risk knowledge, monitoring and warning service, dissemination and

communication, and response capability. Among these four elements, three are more interesting in our context:

- Risk knowledge: [11] indicates that there exists a problem of safeguarding and use of data, especially for old data. Having a tool integrating former experiences / knowledge can be an improvement.
- Monitoring: [11] indicates that there is a lack of data exchange and procedure sharing. Having a tool integrating former experiences / knowledge coming from different sources can be an improvement.
- Communication: [11] indicates that it is difficult to integrate lessons from the past (experiences or warnings). Having a tool integrating former experiences / knowledge coming from different sources and using what happens in the past to make a decision in the present can be an improvement.

Moreover, [15] and [16] highlight the importance of taking into account current knowledge [15] but also past knowledge [16].

Additionally, [17] and [18] emphasize the need of putting in place early warning systems that are efficient, relevant and user-centered and that allow interactions between all tools, persons and other available actors. As a result, international collaborations [12] and information sharing mechanisms [17] have to be set up.

Finally, early warning systems produce warnings. But, a simple warning is often not enough, in particular for a politician. In case of crisis, one wants solutions, i.e., actions to implement / decisions to make [11], [19].

To overcome these shortcomings, we propose to use a recommender system (see [20] for a survey). Recommender systems are a particular form of information filtering designed to present information items (movies, music, books, images, web pages, ...) that may interest the user.

Recommender systems have been studied in many fields, cognitive science, information retrieval [21], [22], web [23], [24], e-commerce [25], web usage mining [26], [27] and many others. The problem of recommendation can be summarized by the problem of estimating scores for items that have not been seen by a user. Indeed, the number of items and the number of users of the system can be very important, it is, therefore, difficult for each user to see all items or that each item is evaluated by all users. It is therefore necessary to estimate the scores for items not yet evaluated.

Intuitively, this valuation is usually based on the scores given by a user to other items and other information that will be formally described below. When it is possible to estimate the scores for items not yet evaluated, then the items with the highest scores may be recommended to the user. More formally, [28] formulates the problem of recommendation in the field of e-commerce as follows.

Definition: Recommendation for e-commerce
Given P the set of all users and M the set of all possible items that can be recommended (such as books, movies, restaurants, ...). Given u a function measuring the utility of an item m for a user p, i.e., $u : P \times M \rightarrow \mathbb{R}$. Then, for each user $p \in P$, we want to choose the item $m' \in M$ that maximizes the utility for the user: $\forall p \in P, m'_p = argmax_{m \in M} u(p, m)$.

In recommender systems, the utility of an item is usually represented by a score that indicates how a particular user liked a particular item. For example, the user $Michel$ gave the score 3 (the maximum score being 10) to the movie "Harry Potter".
Example:
In this example, items are movies that the users Elsa, Camille, Michel and Nicolas have given a score. We obtain the matrix $P \times M$:

u(p,m)	Harry Potter	Ice Age	Ice Age 2	Hulk	Transformers
Elsa		8			7
Camille	9	8			6
Michel	3	5		5	5
Nicolas	5	3		3	3

Note that a cell (i, j) of this matrix corresponds to the utility score given to the movie j by the user i.

The central problem of recommender systems is that the utility u is not usually defined on the full $P \times M$ space, but only on a subset of it. This means that u must be extrapolated to the entire $P \times M$ space.

In recommendation systems, the utility is typically represented by the scores and is first defined over the items previously rated by users. Therefore, the recommendation engine should be able to estimate / predict the scores of item / user unevaluated combinations and to propose relevant recommendations based on these forecasts.

Once the unknown scores are estimated, actual recommendations of an item to a user are proposed by choosing the highest score among all scores provided for the user, according to the formula given in the previous definition (Recommendation for e-commerce).

A recommendation in e-commerce, as defined previously, is the item $m \in M$ (set of all items (movies, books, ...)) such as the utility for a user $p \in P$ (set of all users) is maximum.

To the best of our knowledge, this is the first work dealing with the problem of recommending actions in crisis management and early warning systems context. The idea of using what the other users (called, collaborative approach) did to generate recommendations is very popular in Information Retrieval [20], and Web Usage Mining [29]. Our contribution is to adapt these existing techniques to our context.

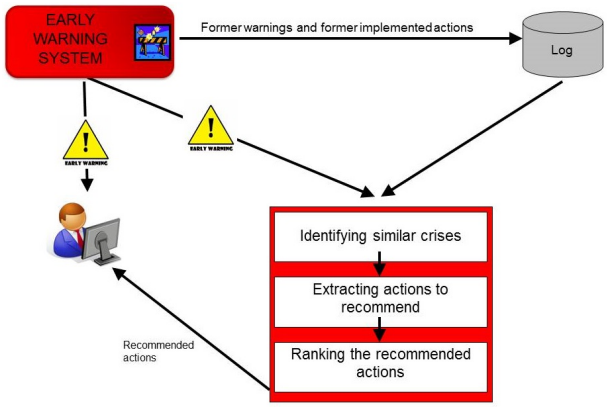

Figure 1. Overview of the recommendation framework

By analogy, we can define a recommendation for warnings as an action $a \in A$ (set of all possible actions) to implement such as its utility for a warning $w \in W$ (set of all possible warnings) is maximum.

Definition: Recommendation for warnings
Given A the set of all possible actions and W the set of all warnings, given a log of warnings and the corresponding indicators and a current triggered warning and given u a function measuring the utility of an action a for a warning w, i.e., $u : A \times W \rightarrow \mathbb{R}$. Then, for each warning $w \in W$, the recommended action $a' \in A$ is the one that maximizes the utility for the warning: $\forall w \in W, a'_w = argmax_{a \in A} u(w, a)$.

Example:
In this example, we just illustrate the obtained matrix $W \times A$ where a score indicates that the action has been implemented and if this action has been considered as efficient:

u(w,a)	Action 1	Action 2	Action 3	Action 4	Action 5
Warning 1	8	7			
Warning 2	9		3		
Warning 3	3	5		5	5
Warning 4	5	3		3	3

Note that a cell (i, j) of this matrix corresponds to the utility score given to the action j by the warning i.

III. Recommendation process

In this section, we describe our generic framework for recommending actions. The framework uses both the characteristics of the current warnings / crises, and the log of former warnings / crises, i.e., the actions implemented during each former crisis. It consists of the three following steps, as illustrated in Figure 1:

1) The first step consists in identifying / selecting into the log, warnings similar to the current triggered warning,
2) The second step consists in extracting the actions that were implemented to manage the former similar warnings,
3) The last step consists in ranking the candidate recommended actions.

The simplified algorithm 1 represents this process. Each step of this process can be parametrized with one or more functions. By changing these parameters, the way how recommendations are computing changes.

Algorithm 1 RecoEWS(L, W_c, Match, Extract, Rank, Default, Clean, $\prec$)

Require:
L: The log of former triggered warnings,
W_c: The current triggered warning,
Match: A match function between two warnings,
Extract: A function extracting actions,
Rank: A function ranking actions,
Default: A function returning a default recommendation,
Clean: A function deleting duplicates in a given set,
$\prec$: An action ranking.
Ensure: An ordered set of recommendations

```
SimW ← ∅ // for the similar warnings
CandAct ← ∅ // for the candidate actions
for each warning w_i ∈ L do
    SimW ← SimW ∪ Match(w_i, W_c)
end for
for each warning w_j ∈ SimW do
    CandAct ← CandAct ∪ Extract(w_j)
end for
Clean(CandAct)
if CandAct ≠ ∅ then
    return Rank(CandAct, ≺)
else
    return Default(L)
end if
```

A. Data sources

First of all, we give an overview of what we consider as our sources.

Each time a warning is set off by the early warning system, then the decision-makers have first to decide if this warning is sufficient major to be taken into account. If this warning is major, the decision-makers have to set up some actions in order to face the crisis / problem pointed out by the early warning system. Thus, the warning triggered by the early warning system can generate, or not, the implementation of a certain number of actions. From our point of view, in both cases, the triggered warning is important and has to be logged.

Each warning is associated to a set of indicators. These indicators have also to be logged. In fact, these indicators are thresholds, values intervals, characteristics, ... For example, the Stockholm International Peace Research Institute (SIPRI) identifies 1260 indicators of early warning ([30]) that can be divided into indicator categories (and subcategories) such as justice and human rights, socio-cultural factors, geopolitical setting, military and security, ... (see [31] for more details). Unfortunately, it seems that indicators are specific to each kind of risk / crisis.

In [32], the authors try to organize the large volume of indicators present in 30 different models and show the diversity of indicators and the special repartition among the different categories. Their review can be helpful in answering questions about: (i) What information the models use to anticipate future events, (ii) How are the early warning models alike, (iii) How are they different, ...

Thus, the authors highlight that comparing indicators and so, warnings from different early warning systems is an hard task.

Our work is a first step towards a recommender system for early warning systems. So, this problem is not the topic of our proposition but it will be a challenge to rise in future work. Finally, in addition to the warning and the corresponding indicators (and values), we consider that the actions implemented to answer to the triggered warning have also to be logged.

More formally, the warning triggered by the early warning system can be seen as a 3-uple containing the warning description, the set of pairs of an indicator and its value and the set of implemented actions. Note that the set of implemented actions can be empty if the decision-maker do not decide to implement some actions when a warning is triggered by the early warning system.

So, we have, for each warning W_i, $\forall i \in \mathbb{N}^{+*}$:

$$\begin{gathered} W_i = \langle\ Description_i, \\ \{\ \langle Indicator_i^1, Value_i^1 \rangle, ..., \langle Indicator_i^n, Value_i^n \rangle\ \}, \\ \{Action_i^1, ..., Action_i^m\}\ \rangle \\ \text{where } m \in \mathbb{N}^+ \text{ and } n \in \mathbb{N}^{+*}. \end{gathered}$$

Note that, in this paper, we consider that indicators are the same for each warning (only the corresponding values are different).

Finally, our data are composed of a log of former triggered warnings and the current triggered warning.

B. Identifying similar crises

Using the log of former triggered warnings and the current triggered warning, this first step consists in identifying similar warnings. We propose to use a match function. This function is used to find a set of warnings matching a given warning. In fact, this match function is used to search among the set of former warnings which ones are matching the current triggered warning. This function outputs a set of warnings similar to the current warning. Note that this returned set can be empty.

A first example of the match function simply consists in comparing, for each logged warning, each pair $\langle Indicator_l, Value_l \rangle$ to each pair $\langle Indicator_c, Value_c \rangle$ of the current warning where $Indicator_l = Indicator_c$. It comes to compare $Value_l$ and $Value_c$.

Because the values of the indicators can be quantitative or qualitative, some similarity measures can be used and / or combined. One purpose easy to implement is to use the cosine similarity which is a (well-known and often used) measure of similarity between two vectors of an inner product space that measures the cosine of the angle between them. One of the advantages of this measure is that it is normalized, i.e., the returnined similarity measure between the two vectors is bounded between 0 and 1.

In fact, for each warning, the indicator vector $\vec{w_i}$ is composed by the indicator values ($Value_i^k, \forall k \in [1..n]$) such as (according to the notation given Section III-A):

$$\vec{w_i} = (Value_i^1, Value_i^2, ..., Value_i^n)$$

The formula of cosine similarity is:

$$sim(\vec{w_1}, \vec{w_2}) = \frac{<w_1, w_2>}{\|w_1\|\|w_2\|} = \frac{\sum_k w_1^k w_2^k}{\sqrt{\sum_k w_1^{k^2}}\sqrt{\sum_k w_2^{k^2}}}$$

Some more sophisticated and more appropriated similarity measures have to be proposed in future work such as, for example, semantic similarity for qualitative indicators.

C. Extracting actions to recommend

The match function of the previous step outputs a set of similar logged warnings, $SimW$, matching the current triggered warning. The goal of this step is to extract a set of actions that will be the basis for the recommendation. Our purpose consists in extracting for each warning of $SimW$, the corresponding set of actions, without duplicates. The obtained set of actions, $CandAct$, is the set of unordered recommendations which is returned.

D. Ranking the recommended actions

In the previous step, a set of recommendations is obtained. The purpose of this next step is to select the most relevant one w.r.t a satisfaction criterion expressed by the user. To this end, an action ranking is needed, that orders the candidate recommendations. Again, there are many ways of ranking the candidates, from very basic to sophisticated ones. We list here just few:

- Ranking the candidates according to their number of occurrences
- Ranking the candidates according to a user profile. For example, some actions sometimes can not be implemented in a given context.
- Ranking the candidates according to the order of the corresponding warning into the log, i.e., the most recent implemented actions will be rank first. For example, diffusion channels evolve over time.

E. Default recommendation

As previously noted, the set of candidate recommendations can be empty: when no similar warning is founded or when the set of similar warnings is empty. In such a case, some recommender systems are able to provide the user with a default recommendation. Of course, various default recommendations can be proposed to the user. Unfortunately, considering our context of crisis management and the impacts

and stakes of such a context, and according to the status of this work, we consider that it is better to recommend anything as default recommendation. Proposing an efficient default recommendation is an interesting research issue.

IV. Example of instantiation

In this section, we display a possible instantiation of our framework to illustrate the applicability of our algorithm.

Suppose the early warning system triggers a warning, W_c and our log L of former triggered warnings is composed by three warnings W_1, W_2, W_3 such as:

- $W_c = \langle\ Desc_c, \{\ \langle I_1, V_1\rangle, \langle I_2, V_2\rangle, \langle I_3, V_3\rangle, \langle I_4, V_4\rangle\ \}, \emptyset\ \rangle$
- $W_1 = \langle\ Desc_1, \{\ \langle I_1, V_1\rangle, \langle I_2, V_2\rangle, \langle I_3, V_3\rangle, \langle I_4, V_5\rangle\ \}, \{A_1, A_2, A_3\}\ \rangle$
- $W_2 = \langle\ Desc_2, \{\ \langle I_1, V_6\rangle, \langle I_2, V_7\rangle, \langle I_3, V_8\rangle, \langle I_4, V_9\rangle\ \}, \{A_4, A_5\}\ \rangle$
- $W_3 = \langle\ Desc_3, \{\ \langle I_1, V_{10}\rangle, \langle I_2, V_2\rangle, \langle I_3, V_3\rangle, \langle I_4, V_4\rangle\ \}, \{A_1, A_6\}\ \rangle$

The corresponding utility matrix could be:

u(w,a)	A1	A2	A3	A4	A5	A6
W_1	6	5	6			
W_2				7	8	
W_3	9					6
W_c						

First, by representing each warning with an indicator vector, we obtain:

$\vec{W_c} = (V_1, V_2, V_3, V_4)$ $\quad \vec{W_1} = (V_1, V_2, V_3, V_5)$

$\vec{W_2} = (V_6, V_7, V_8, V_9)$ $\quad \vec{W_3} = (V_{10}, V_2, V_3, V_4)$

Suppose now computing the cosine similarity shows that $sim(\vec{W_c}, \vec{W_1}) = sim(\vec{W_c}, \vec{W_3})$ and $sim(\vec{W_c}, \vec{W_2}) >> sim(\vec{W_c}, \vec{W_1})$. So, the first step of our process returns $SimW = \{W_1, W_3\}$.

Then, our extract function returns the actions corresponding to each warning of $SimW$, i.e. $CandAct = \{A_1, A_2, A_3, A_1, A_6\}$. This set is then cleaned. So, the second step of our process returns $CandAct = \{A_1, A_2, A_3, A_6\}$.

Finally, our rank function orders the candidate actions. Here, we combine the two of the three options we propose Section III-D. In this toy example, from our point of view, the most relevant actions are the ones that frequently occur, then the ones which warnings are more recent. Compare to the other actions, A_1 occurs twice. And W_3 is the more recent warning of the log. So, the last step of our process returns the ordered set of actions: $\{A_1, A_6, A_2, A_3\}$.

The corresponding utility matrix could become:

u(w,a)	A1	A2	A3	A4	A5	A6
W_1	6	5	6			
W_2				7	8	
W_3	9					6
W_c	9	5	5			6

V. Conclusion and perspectives

In this position paper, we propose challenges and research issues for recommending actions to implement when a warning is triggered by an early warning system for crisis management. We turn our framework towards a generic framework, in the sense that it can be instantiated to change the way recommendations are computed. We give few examples of how it can be instantiated.

Challenges and research issues to rise include (but are not limited to):

- The implementation of our framework and the conduction of experiments in order to better assess the quality of recommended actions, as well as the assessment of various instantiations of our framework in order to determine to what context they are better adapted.
- The problem of the indicators has been mentioned in this article. Indeed, a large number of indicators exists. They are different depending on the type of risk but also they can be different for the same type of risk. However, no standard exists. A standardization of indicators, for example, by type of risk (or other), seems to be a promising research issue.
- The investigation of other instantiations of our framework. For example, we would like to investigate how to compute the more efficient distance between indicators. As another example, we would like to propose other Match and Rank functions. More precisely, the Match function which is at the heart of the candidate generation, could be instantiated with more sophisticated distances.
- When no recommendation can be computed, considering our context of crisis management and the impacts and the stakes of such a context, the purpose of a smart default recommendation is an interesting research issue.
- The integration of the user knowledge has to be taken into account in our framework. Finding a way in that sense, for example, with the user profile, is an other research issue.
- The longer term objective is to provide a generic platform of recommendations that can adapt to the needs of users, types of crises and many other parameters.

This work in progress, which uses the knowledge gained from past experiences to make the best decision in the present, is a first step towards the enhancement of early warning systems and crisis management.

References

[1] J. Heppen and S. Therriault, "Developing early warning systems to identify potential high school dropouts," National High School Center, Tech. Rep., 2009.

[2] N. Cheang, "Early warning system for financial crises," *Macao Monetary Research Bulletin*, no. 11, pp. 61–77, 2009.

[3] A. Gaytàn and C. A. Johnson, "A review of the literature on early warning systems for banking crises," Central Bank of Chile, Working Papers Central Bank of Chile 183, Oct. 2002.

[4] M. Angermann, M. Guenther, and K. Wendlandt, "Communication architecture of an early warning system," *Natural Hazards and Earth System Science*, vol. 10, no. 11, pp. 2215–2228, 2010.

[5] German Committee for Disaster Reduction (DKKV); Humanitarian & Development Network (HDN); Platform for the Promotion of Early Warning (UNISDR PPEW), "Emerging challenges for early warning systems in context of climate change and urbanization," 2010.

[6] M. Buchanan-Smith, "Role of early warning systems in decision making processes," in *an Expert Group Meeting held September 5-7, in Lisbon, Portugal.*, 2000. [Online]. Available: http://drought.unl.edu/monitor/EWS/ch2_Buchanan-Smith.pdf

[7] E. Bakker, "Early warning by ngos in conflict areas," in *Bas Arts*, M. Noortmann and B. Reinalda, Eds. Non-State Actors in International Relations, 2001, pp. 263–277.

[8] International Federation of Red Cross and Red Crescent Societies (IFRC), "Community early warning systems: guiding principles," 2013.

[9] P. Genovesi, R. Scalera, S. Brunel, D. Roy, and W. Solarz, "Towards an early warning and information system for invasive alien species (ias) threatening biodiversity in europe," EEA (European Environment Agency), Tech. Rep., 2010.

[10] N. Waidyanatha, "Towards a typology of integrated functional early warning systems," *IJCIS*, vol. 6, no. 1, pp. 31–51, 2010.

[11] United Nations, *Global Survey of Early Warning Systems: an Assessment of Capacities, Gaps and Opportunities Toward Building a Comprehensive Global Early Warning System for All Natural Hazards: a Report Prepared at the Request of the Secretary-General of the United Nations*, 2006.

[12] B. E. Joseph E. Quansah and G. L. Rochon, "Early warning systems: A review," *Journal of Terrestrial Observation*, vol. 2, no. 5, 2010.

[13] P. E. R. I. (PERI), "Early warning systems - interdisciplinary observations and policies from a local government perspective - four elements of people centered early warning systems," Public Entity Risk Institute (PERI) Symposium, 2010.

[14] M. Glantz, *Early Warning Systems: Do's and Don'ts*, ser. Usable science. The Center, 2004.

[15] MEA (Millennium Ecosystem Assessment), "Ecosystems and human well-being: A framework for assessment," Washington DC: Island Press, p. 245, 2003.

[16] P. Schrodt and D. Gerner, "The impact of early warning on institutional responses to complex humanitarian crises," Paper presented at 3rd Pan-European International Relations Conference, 1998.

[17] P. Hall, "Early warning systems: Reframing the discussion," *Austral. J. Emergency Manag.*, vol. 22, no. 2, 2007.

[18] Hyogo Framework for Action, "2005-2015: Building the resilience of nations and communities to disasters," extract from the final report of the World Conference on Disaster Reduction, UN ISDR, 2005.

[19] H. Lavoix, *Developing an early warning system for crises*. From early warning to early action? The debate on the enhancement of the EUs crisis response capability continues. Brussels: European Commission, 2008, pp. 365–382.

[20] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, 2005.

[21] R. A. Baeza-Yates and B. A. Ribeiro-Neto, *Modern Information Retrieval*. ACM Press / Addison-Wesley, 1999.

[22] G. Salton, *Introduction to Modern Information Retrieval (McGraw-Hill Computer Science Series)*. McGraw-Hill Companies, September 1983.

[23] R. A. Baeza-Yates, C. A. Hurtado, and M. Mendoza, "Query Recommendation Using Query Logs in Search Engines." in *EDBT Workshops*, ser. Lecture Notes in Computer Science, W. Lindner, M. Mesiti, C. Trker, Y. Tzitzikas, and A. Vakali, Eds., vol. 3268. Springer, 2004, pp. 588–596.

[24] R. W. White, M. Bilenko, and S. Cucerzan, "Studying the use of popular destinations to enhance web search interaction," in *SIGIR*, 2007, pp. 159–166.

[25] J. B. Schafer, J. A. Konstan, and J. Riedl, "E-commerce Recommendation Applications," *Data Mining and Knowledge Discovery*, vol. 5, no. 1/2, pp. 115–153, 2001.

[26] J. Srivastava, R. Cooley, M. Deshpande, and P.-N. Tan, "Web Usage Mining: Discovery and Applications of Usage Patterns from Web data," *SIGKDD Explorations*, vol. 1, no. 2, pp. 12–23, 2000.

[27] R. Baeza-yates, C. Hurtado, M. Mendoza, and G. Dupret, "Modeling user search behavior," in *In LA-WEB 05: Proceedings of the Third Latin American Web Congress*. IEEE Computer Society, 2005, p. 242.

[28] G. Adomavicius and A. Tuzhilin, "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 17, no. 6, pp. 734–749, 2005. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1423975

[29] J. Srivastava, R. Cooley, M. Deshpande, and P.-N. Tan, "Web usage mining: discovery and applications of usage patterns from web data," *SIGKDD Explor. Newsl.*, vol. 1, no. 2, pp. 12–23, Jan. 2000.

[30] Stockholm International Peace Research Institute (SIPRI), "Sipri early warning indicators database," 2009. [Online]. Available: http://apps.sipri.org/ewi/db_new.phtml?ind_search=&word=All&submitButtonName=Submit

[31] O. Walton, "Helpdesk research report: Early warning indicators of violent conflict," Governance and Social Development - Resource Centre (GSD-RC), Tech. Rep., 2011.

[32] F. Barton, K. von Hippel, S. Sequeira, and M. Irvine, "Early warning? a review of conflict prediction models and systems," Center for Strategic and International Studies (CSIS), Tech. Rep., 2008.