

Automatic Security Classification by Machine Learning for Cross-Domain Information Exchange

Hugo Hammer, Kyrre Wahl Kongsgård, Aleksander Bai, Anis Yazidi, Nils Agne Nordbotten and Paal E. Engelstad

Abstract—Cross-domain information exchange is necessary to obtain information superiority in the military domain, and should be based on assigning appropriate security labels to the information objects. Most of the data found in a defense network is unlabeled, and usually new unlabeled information is produced every day. Humans find that doing the security labeling of such information is labor-intensive and time consuming. At the same time there is an information explosion observed where more and more unlabeled information is generated year by year. This calls for tools that can do advanced content inspection, and automatically determine the security label of an information object correspondingly. This paper presents a machine learning approach to this problem. To the best of our knowledge, machine learning has hardly been analyzed for this problem, and the analysis on topical classification presented here provides new knowledge and a basis for further work within this area. Presented results are promising and demonstrates that machine learning can become a useful tool to assist humans in determining the appropriate security label of an information object.

Index Terms—Security, classification, labeling, cross-domain information exchange, machine learning.

I. INTRODUCTION

SECURITY labels are used by the military, government agencies, international organizations and private corporations to associate security attributes to a specific information object [1]. These labels are intended to convey for instance the sensitivity of the contents of the information object. Traditionally, the information objects were typically *paper documents* with printed security markings indicating the confidentiality classification of the documents. In a military setting, examples of such security markings include the labels "Unclassified", "Restricted", "Confidential", "Secret" and "Top Secret". The security label mandates how the information in the document shall be treated according to the governing security policy. For example, the "Unclassified" security label might mean that the information in the document does not need any particular protection, while the other markings might indicate that the information is classified and that the information in the document must be handled accordingly.

Document submitted on 25. April 2015.

H. L. Hammer, A. Bai, A. Yazidi and Paal Engelstad are with the Oslo and Akershus University College of Applied Sciences (HiOA), Oslo, Norway (e-mails: hugo.hammer, aleksander.bai, anis.yazidi, paal.engelstad@hioa.no).

K. W. Kongsgård, N. A. Nordbotten and P. E. Engelstad are with the Norwegian Defense Research Establishment (FFI), Kjeller and University of Oslo (UNIK), Norway (e-mails: kyrre-wahl.kongsgard, paal.engelstad, nils.nordbotten@ ffi.no).

This work was partially funded by the University Graduate Center (UNIK).

In modern environments the information objects are more likely to consist of digital information, e.g., Word documents, text messages and e-mails. If the purpose of the security labeling is concerned with preserving confidentiality, a security label such as the XML Confidentiality Label [2] can be used. The security label can be digitally attached and bound to the data object, e.g., by using a cryptographic mechanism such as a digital signature. This will also protect the integrity and authenticity of the security label (and its associated data object) during transportation and storage. However, the digital signature does not guarantee the correctness of the originally assigned security label.

Within a military setting, there are many information domains (e.g. networks, information systems, etc.) that have not implemented a mechanism to attach digital security labels to the digital information objects residing within the domain. The security attributes of the information objects will then typically be determined implicitly by the context, e.g., that "all digital information objects residing within Domain X shall be treated as Secret (S)" (Figure 1).

The first problem with this approach is that it can easily introduce inconsistent classification and massive overclassification of the information objects within the domain, which is considered a practical challenge today [3]. For instance, Domain X in the example above might contain only a small number of information objects that require the security label Secret (Fig. 1). However, to preserve the confidentiality of the few Secret information objects, the remaining vast majority of information objects in the domain, which all deserves a lower-level security label, have to be considered as Secret as well.

The second problem of not having an explicit mechanism for assigning security labels is that it easily renders the information domain into an inflexible and isolated information silo. Even though Domain X in the example above contains mostly security information of low security classification, the information objects within the domain cannot easily be exchanged with systems that are not accommodating Secret information or be accessed by systems or personnel without a corresponding security clearance.

In the advent of the information age, these information silos are exactly what most organizations have been struggling to avoid and to move away from over the past 10-20 years. For instance, within this period of time the NATO allies

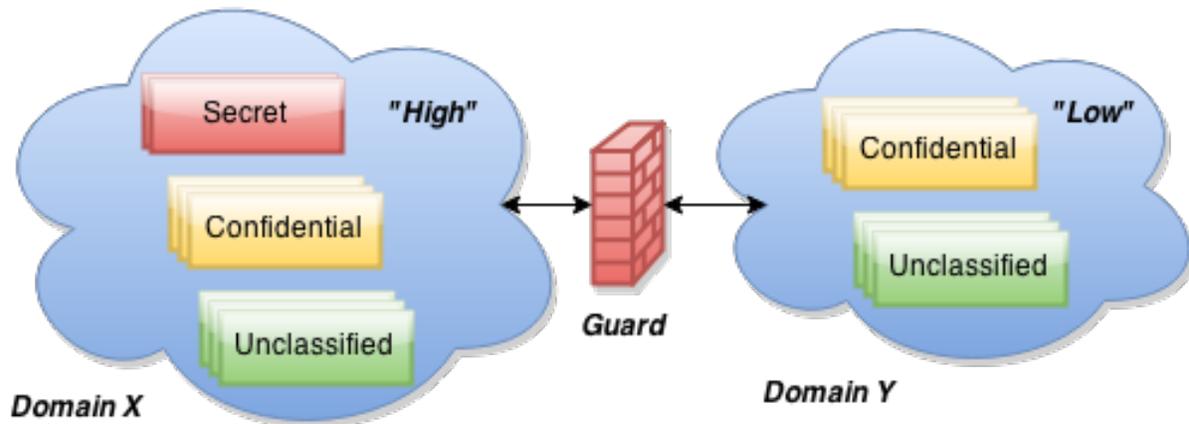


Fig. 1: Example of cross-domain information exchange across a Guard between Domain X and Domain Y, where both domains contain various Unclassified (U) and Confidential (C) information objects. Only Domain X contains Secret (S) information objects. Here, the Guard should ensure that the Secret objects in Domain X are not leaked into Domain Y.

have identified the need to move from the old paradigm of "need-to-know" to a new paradigm of "need-to-share" [4], [5], [6]. This means shifting away from single-domain information silos and move towards cross-domain information exchange (Fig. 1), without losing the ability to protect the information appropriately. Information sharing is considered a strategic capability and a steppingstone to obtain information superiority by ensuring that all allies have the newest and most pertinent information at hand at any time [7].

Two-way cross-domain information exchange is usually accommodated by a guard that is responsible for the information flow control between two domains (Figure 1). The guard inspects the security label of any information object that is requested moved from one domain to the other. The guard will only permit the information object to be passed to the other domain if it carries a security label that authorizes this action. The use of guards will be outlined in further detail in Section II.

A critical point is to determine which security label to assign to an information object in the first place. A simple approach is to attach an explicit security labels based on the implicit security classification of the origin of the information object. In many cases this is acceptable. For instance, it may be known that a specific sensor only produces data of a specific type with a given classification. However, often the origin of the object will not give an accurate classification (e.g., if the origin is Domain X in the example above). Thus, in many other scenarios, one would prefer to assess the actual content of the data object instead, to determine the correct attributes of the security label.

The traditional way of assigning security labels is based entirely on human judgment, e.g., the security classification of a document is determined solely by the author of the document or by another evaluator performing a review. While security labeling is paramount to cross-domain information exchange and future information superiority in the military domain, undertaking the actual labeling is a tremendous challenge. The

most optimistic advocate the introduction of fully automated tools. However, one may argue that tools are primarily needed to support, assist and partly offload the humans in their effort of assessing the security classification of an information object. Results presented in this paper are applicable to both approaches.

The work in this paper is applicable to both human-assisted (semi-automatic) and a fully automatic the security labeling process. The three main use-case scenarios include:

- 1) Proactive labeling (human-assisted): Security labeling is undertaken by a combination of human effort and automated tools. The tools can be used proactively, e.g., the software is assisting the human in determining the correct security label for the information object.
- 2) Reactive labeling (human intervention): In a setting where all the data has been manually assigned a security label, we still want to detect if data, when it leaves a secure domain, has been previously mislabeled either due to human error or malicious agents. These reactive label checkers can trigger actions depending on the policy, such as denial of access or a notification that calls for human intervention.
- 3) Fully-automated labeling: A fully automated solution would be desirable, and could be used in some settings. However, in many military settings and other security-dependent scenarios, it might not be realistic to expect it implemented in the near future.

While we acknowledge that a completely automated security labeling process might be unrealistic in near future, solutions for automatic labeling is a corner stone also for the scenarios above that includes human assistance or intervention.

Even though cross-domain information exchange could benefit from methods for automatic security labeling, little has been published on the topic. The contribution of this paper is putting this topic on the agenda, and exploring techniques for content analysis that is more sophisticated than the dirty word scanning techniques that are currently deployed. In

this paper, we employ a machine learning approach to the problem, and to the best of our knowledge, this paper is together with our previous works in [8], [9] and [10] the only published work addressing this issue. Other relevant available information we have found is a Master thesis from 2008 [11], as well as a military (thus, difficultly accessible) technical report [12]. All three works make comments on the surprising lack of published data. As noted in Section II on related work below, there is a need for specific methods tailor-made for the problem and commercially available products are not sufficient to this end.

In summary, methods and algorithms on the specific problem of security labeling need to be explored in the open literature, and their accuracy in determining the correct security classification of an information object should also be assessed openly. This paper attempts to bring the problem space and the technology development into the open literature domain.

II. RELATED WORK

The following sections provides a brief overview of related work in the field.

A. Cross-Domain Information Exchange and Guards

Figure 1 shows a simple scenario requiring information exchange between different domains. As indicated in the figure, it relies on placing a data guard between the two domains. A number of commercial cross-domain information transfer solutions/guards, e.g., Lockheed Martin's Radiant Mercury (RM), BAE's DataSync Guard and Boeing's eXMeritus Hardware Wall, have been certified and officially approved for use by the Department of Defense (DoD) in the US [13]. These data guards facilitates the secure information transfer between networks operating at different levels of classifications and with different security policies. Common for most of the commercially available guards is that they perform basic checking of the format and metadata of the information objects, such as checking that the information object is formatted according to required specifications or that it carries the required metadata/label. In terms of content scanning, the guards typically support some type of basic "dirty word" checking. As discussed in Section I there is a need for more advanced content scanning techniques. The security labels that are checked by the guard before it permits information to be released, can be first set by the guard itself (e.g. in a separate pre-processing software module), by a module on the clients, or by a labeling gateway (e.g., see [14]). If there are less stringent requirements for assurance, e.g., for scenarios involving relatively low risk, labeling may also be performed by the user on a commodity system. Titus [15] offers several applications for this, e.g., enabling a user to label files on Microsoft Windows and providing plug-ins for use within Microsoft Office applications. SMHS [16] provides a similar plug-in for Microsoft Outlook, utilizing a security label service from Isode [17].

B. Frameworks for Security Labeling of Information Objects

Kongsgård et. al. [18] provide a security labeling framework for determining what security label attributes are to, and should not, be included within a given data objects security label according to policy. They present a solution for the use of attribute based access control (ABAC) principles to the process of information labeling. In particular, the framework provides support for pluggable attribute modules (e.g., content checkers) whose output serve as input to the policy decision. The work conducted herein is as such highly relevant to the work in [18], by providing a potential attribute module.

A variation of ABAC, Content-based Protection and Release (CPR) [19], has been proposed for future use in NATO. In CPR attributes within a content label are used to convey the properties of an information object. Access decisions are then based on protection and release policies effectively expressing requirements (in terms of attributes) on the user and her terminal and/or environment in order to be granted access to information objects with such properties. CPR depends on the ability to assign content properties to information objects, and the work proposed in this paper is therefore relevant. The CPR paper [19] also presents the NATO Metadata Binding Service (NMBS). NMBS can be used to bind specified metadata (e.g., a security label) to an information object, using a binding mechanism of particular strength (e.g., digital signature). Similar services are also available as commercial products, e.g., [20].

C. Content Scanning for Automatic Security Labeling

Since the content scanning of existing cross-domain information exchange solutions is typically limited to some form of "dirty-word checking", there is a need for exploring more advanced scanning techniques. A review of commercially available content tools is undertaken in [12]. The conclusion is that there might exist commercial content analysis tools that could be appropriate for the task, but these are proprietary. Without knowledge of their implementations or reports on their performance for this problem, they are of little use in an academic setting. Note that a product that is designed for one general task (e.g. determining the topic of a document) might not perform well in a specific task (e.g. determining the security classification of a document). Furthermore, most research in document categorization focuses on identifying the topic of a document (topical classification), while security classification is non-topical, and usually a more challenging problem. Moreover, often the methods need to be tailor-made for the specific tasks (e.g. for the type of security attribute addressed) and optimized for a specific context (e.g. for the type of information object or for the exact topic of the information content). Thus, knowledge of specific methods is more important in the long run than availability of generic proprietary products or research results from related or more general problems. It was not before 2008 that it was suggested in a Master thesis to use general machine learning techniques to the problem [11]. The thesis did not make attempts to apply machine learning, but proposed an architecture for the

problem. In the architecture, the document is first checked for compliance with policy and classified by topic, before the actual security classification is undertaken. General machine learning methods were first applied to the problem in a military technical report from 2010 [12]. The three general classifiers - Nearest Neighbor (NN), Naive Bayes (NB), and Support Vector Machine (SVM) are applied for security classification and compared. The corpus, which is collected from the Digital National Security Archive [21], is manually separated by topic in advance. Different pre-processing methods commonly used for text analysis, such as word stemming, term weighting and dimensionality reduction, are applied, and the effects of applying them are also analyzed. Among different contributions, this paper investigates the effects of topic classification prior to security classification, as only assumed in previous works [11], [12]. To the best of our knowledge, the works in [8], [9] and [10] were the first published papers on automatic security classification and the application of machine learning techniques to this problem.

III. EXPERIMENTS, RESULTS AND DISCUSSIONS

A. Lexical Corpus

The Digital National Security Archive contains the most comprehensive collection of declassified US government documents available to the public [21]. We base our analysis on documents from the same collection that were used in [8], [9] and [10]. This collection was originally chosen because it contains a mix of both classified and unclassified documents from three unrelated domains:

- AF, Afghanistan: The Making of U.S. Policy, 1973-1990
- CH, China and the United States: From Hostility to Engagement, 1960-1998
- PH, The Philippines: U.S. Policy during the Marcos Years, 1965-1986

Of the 5853 documents available within these three topics, we do not use duplicate documents and documents classes that are very small or have an unsuitable classification [8], [9]. Then we remove documents with 30 words or less (after having also removed some keywords as explained below), and we end up with 2805 documents in total (Table I).

To simplify the analysis, we reduce the security classifications into two main classes used for the further analysis. The first class is *Unclassified*, which comprises 1079 documents. The second class is *Classified*, which is an aggregation of the document classes "Confidential", "Secret" and "Top Secret" in Table I [8], [9] [10]. However, then we remove some of the latter documents to get an approximately equal share between classes, ending up with around 2158 documents. This is only approximate, because we pick randomly documents from the aggregate "Classified" class with an expectancy of remaining with 1079 documents.

The documents on the DNSA website are scanned pdf documents of poor quality, and the content text extracted from OCR (optical character recognition) has many errors. However, each document has also a meta-data in plain text format containing

TABLE I: Documents used in the experiments (after removal of short documents and before balancing the corpus).

	Total docs	Unclassified.	Confidential	Secret	Top Secret
<i>AF</i>	834	333	395	102	4
<i>CH</i>	948	322	247	286	93
<i>PH</i>	1023	424	514	85	0
<i>Sum</i>	2805	1079	1156	473	97

further information about the pdf documents. Many of the documents have an abstract (i.e. extract of the content of the document), i.e. the abstract is given in plain text format in the meta-data attached to the pdf documents. The abstracts are short texts that are assumed to contain high quality information about the document content. Such meta-data was used in [12], but details are not specified, except that the limited number of abstracts used indicates that they selected a small subset of the abstracts for their analysis.

B. Data processing and machine learning

To go further in our analysis we extracted the raw textual contents using OCR techniques, using the OCR service provided by Abbyy [22].

For all the experiments we resorted to the simple bag-of-words model [23], in which any word order was discarded and a document was represented merely by a vector of term frequency-inverse document frequency (tf-idf) weights.

For our "base case" analysis we utilized spell checkers and auto-correction to mend many OCR errors and performed some analysis on this processed material. The processed/auto-corrected text material is more concise, and less verbose than the raw text (i.e. mis-spelled words are merged into the same word for the analysis). However, some important information might be lost, such as abbreviations that might be significant for the classification. Thus, we also undertook analysis on the raw uncorrected text material. (The latter analysis, which is not part of the "base case", is referred to as "raw" later.)

The raw material has the advantage of containing all information, while the quality of the information is poorer, e.g. in terms of many words with spelling errors etc. This trade-off indicates that in a real scenario where one does not have to rely on OCR, results would be generally better than presented in this paper.

We have not performed word stemming before machine learning. It turns out that this does not affect the performance considerably in neither positive nor negative way, but it reduces the size and sparsity of the vectors and gives easier computation.

Furthermore, as part of the pre-processing in our "base case" analysis we remove/ignore any keywords of the type "SECRET", "UNCLASSIFIED" etc. from all the textual contents prior to training the machine learning algorithm. If we were to leave these types of words in the text, it would potentially result in the classifiers yielding artificially good results, that effectively would only determine the security label based on

absence or presence of such words. However, we also do some non-base-case analysis on material where these keywords are not ignored (referred to as the "keyword" analysis later).

As for the actual machine learning we apply Support Vector Machine (SVM) for classification of classified vs unclassified documents. 70% of the documents are used for training set, while the remaining 30% constitutes the test set. As a starting point, we try to avoid any historic effect in terms of change within the topic over time. Thus, in our base-case analysis documents are not ordered chronologically, so both the training set and the test set contain documents that span the entire time period for the given topic.

C. Analysis of Base Case

As outlined above, the base case comprises full-text OCR'd documents that are post-processed with various spell checking and auto-correction techniques. Classification-related keywords, such as "Secret" and "Unclassified", are removed from the text, and all documents are shuffled into a non-chronological order before the documents are split into a training set and a test set that is analysed with SVM. Results are summarized in the first line entry of Table 1.

A result of 78% classification accuracy (in the "All docs together"-column) indicates that the application of machine learning is a promising solution to the problem addressed in this paper. (The other columns will be discussed later in relation to clustering.)

D. The effect of keywords

In the second line entry we use the same documents as in the base case, but do not remove words like "Secret" and "Unclassified". We see that this has a clear effect on the results where the accuracy jumps from 78% to 87%. The second line entry indicates that our precautionary concern about ignoring keywords (such as "Secret" and "Unclassified") to avoid unrealistically positive results, was well-founded.

E. Chronological Order

In a real setting the classification of new documents will be based on machine learning performed on historical documents. To test this we arranged the documents in chronological order, and performed training on the first part of the range. Then we tested on the second part. Note that this is an almost unrealistically difficult task. The training documents are used to classify documents that are created 10 to 15 years after the end of the training period.

We see that chronological aspect has a clear effect on the results with the accuracy dropping from 78% (base case) to 71%.

F. Analysis of Raw Text

The "Raw text (not auto-corr.*)" row in Table 1 shows results for the base case machine learning undertaken on the raw output of the OCR, i.e. that is not subject to auto-correction. By performing auto-correction many words are corrected to the correct spelling (e.g. Chima is change to China) and one

may expect that this is important. On the other hand, one might also reduce the amount of data and eliminate important part of the information (e.g. such as abbreviations).

The results show that the accuracy increasing from 70% to 78% when auto-correction is performed. Auto-correction is very important for the classification performance of scanned documents. By visual inspection of the effect of the auto-correction, this is not a very surprising result.

In a real scenario where OCR is not needed, these results indicate that machine learning will probably perform even better than demonstrated in this paper, because there will not be such a large amount of spelling errors, and at the same time no information need to be dropped to fix them.

G. Analysis of Abstracts Only

To analyze machine learning on more sparse information, Table 1 summarizes also results for machine learning undertaken only on the abstract (meta-data) of the documents. The abstracts are high-quality information that is available in text format without requiring OCR (i.e. few spelling errors seen), assumed to be formulated "to the point" and containing a brief overview of the essential information of the documents.

The results indicate a little drop in performance. This may indicate that the amount of information is often important for the applicability of machine learning. It might give an indication that the machine learning approach is well applicable in scenarios where large information objects (e.g. text documents) are exchanged, while less applicable to scenarios with smaller information objects (e.g. exchange of short emails or other text messages).

H. Clustering

The fourth column of Table 1 shows results where documents are split per country (manually split into 3 clusters), before the SVM is applied. In the sixth column we have used k-means clustering to split into three clusters automatically, disregarding the fact that documents are divided into three parts from the start (countries). The three lower entry lines shows results for other number of clusters, and also the use of two topical clusters within each of the three manual country clusters.

These results indicate that manual clustering (see the results for "split per country" in Table 1) is not necessarily better than automatic clustering (in the lower part of the table). We also observe there are great potential in clustering, as separate security classes might form separate clusters.

The sparse amount of previous work has assumed that manual topical clustering (here: per country) should be undertaken before the security classification without supporting these claims. Our results indicate, however, that this is not necessarily the case. This indicates that often the advantage of learning from more documents might outweigh the disadvantage of learning from other topics that are less relevant.

IV. CONCLUDING REMARKS

This paper is aiming to put machine learning on the research agenda for cross-domain information exchange. Experiments

TABLE II: Main results of base case analysis and different variants of this

Analysis case	All docs together	95% conf.int.	Split by country	95% conf.int.	Three clusters	95% conf.int.
<i>Base case</i>	0.78	(0.75, 0.81)	0.80	(0.76, 0.83)	0.78	(0.75, 0.81)
<i>With keywords included</i>	0.89	(0.87, 0.92)	0.88	(0.85, 0.91)	0.90	(0.88, 0.92)
<i>Chronologically ordered</i>	0.75	(0.71, 0.78)	0.66	(0.62, 0.70)	0.70	(0.67, 0.73)
<i>Raw text (no auto-corr.)</i>	0.87	(0.84, 0.89)	0.84	(0.80, 0.86)	0.85	(0.82, 0.88)
<i>Short text abstracts</i>	0.72	(0.66, 0.78)	0.72	(0.66, 0.78)	0.69	(0.61, 0.76)
<i>Base case w/ 2 clusters</i>	0.79	(0.76, 0.82)	0.79	(0.75, 0.82)		
<i>Base case w/ 4 clusters</i>	0.76	(0.73, 0.80)				
<i>Base case w/ 8 clusters</i>	0.77	(0.73, 0.80)				

confirm that machine learning approaches perform well and might become a valuable tool in this setting. Even though the performance can probably be increased further from an accuracy of 86% achieved in this paper, the remaining 14% inaccuracy indicates that without further research and improvements the method is applicable as an automatic tool targeted at assisting humans in determining the appropriate label or at detecting potential mislabeling of data objects.

The paper provides a number of learning-points compared to the sparse amount of previous work. In [12] it was assumed that manual topical clustering (here: per country) should be undertaken before the security classification, while the work in [11] assumed automatic topical clustering. None of the works support these claims. Our results indicate, however, that this is not necessarily the case. This indicates that often the advantage of learning from more documents might outweigh the disadvantage of learning from other topics that are less relevant.

Furthermore, our results indicate that manual clustering is not necessarily better than automatic clustering. We also observe that there is a great potential in clustering, as separate security classes might form separate clusters. Exploring more along the line of more fine-grained topical classification is an interesting issue for further work.

We also argue that the chronological aspect should be taken into consideration in an analysis, since this comes natural with a practical use of automatic security classification. Our results confirms that the temporal order of the documents does make a difference, but the reduction in the performance was limited, indicating that machine learning still is a applicable and promising method. The development of a machine learning method that takes the temporal aspects into consideration is an issue for further investigation.

Finally, this paper has been limited to the use of classification of two security classes. Expanding the analysis presented in this paper to a scenario with more classes is a natural next step.

REFERENCES

- [1] R. Kissel, *Glossary of Key Information Security Terms*. DIANE Publishing Company, 2011. [Online]. Available: <http://books.google.no/books?id=k5H3NsBXIsMC>
- [2] A. Eggen, R. Haakseth, S. Oudkerk, and A. Thummel, "XML confidentiality label syntax," FFI-rapport 2010/00961, 2010.
- [3] U. S. G. P. Office, *Too Many Secrets: Overclassification as a Barrier to Critical Information Sharing*. Hearing Before the Committee on Government Reforms, US House of Representatives, 2004.
- [4] C. Cavas, *Petraeus: US must share more info with allies*. [Online], <http://www.defensenews.com/story.php?i=4623591>, 2015. [Online]. Available: <http://www.defensenews.com/story.php?i=4623591>
- [5] P.-P. Meiler and M. Schmeing, *Secure Service Oriented Architectures (SOA) supporting NEC, (Technical Report TR-IST-061)* NATO, 2009.
- [6] U. Wolf, "Does NATO meet the challenge of the information era?" in *23rd International Workshop on Global Security, Berlin, Germany*, 2006.
- [7] N. Brown, *Statement for the Record Before the 108th Congress Committee on Armed Services*. US House of Representatives, 2003.
- [8] P. E. Engelstad *et al.*, "Automatic security classification with lasso," *Proceedings of The 16th International Workshop on Information Security Applications (WISA 2015)*, Jeju Island, Korea, August 20-22, 2015.
- [9] P. E. Engelstad, H. L. Hammer, A. Yazidi, and A. Bai, "Advanced classification lists (dirty word lists) for automatic security classification," *Proceedings of The 7th IEEE International Conference on Cyber-enabled distributed computing and knowledge discovery (CyberC, 2015), Cyber Security and Privacy (CSP), Xian, China, Sept 17-19, 2015*.
- [10] —, "Analysis of time-dependencies in automatic security classification," *Proceedings of The 7th IEEE International Conference on Cyber-enabled distributed computing and knowledge discovery (CyberC, 2015), Cyber Security and Privacy (CSP), Xian, China, Sept 17-19, 2015*.
- [11] K. Clark, "Automated security classification," Master's thesis, Vrije Universiteit, 2008.
- [12] J. D. Brown and D. Charlebois, "Security classification using automated learning (scale)," DRDC Ottawa CR, Tech. Rep., 2010.
- [13] UCDMO. Cross domain wiki. <http://www.crossdomain.org>. Accessed: 2015-03-26.
- [14] R. Haakseth, N. A. Nordbotten, Ø. Jonsson, and B. Kristiansen, "A high assurance guard for use in service-oriented architectures," *International Conference on Military Communications and Information Systems*, 2015.
- [15] Titus classification. "<http://www.titus.com/>". Accessed: 2015-03-26.
- [16] SMHS-Ltd and Braid, [Online]. [Online]. Available: <http://www.smhs.co.uk/>
- [17] Isode, "Isode security label server." [Online]. Available: <http://www.isode.com/products/security-label-server.html>
- [18] K. W. Kongsgård, N. A. Nordbotten, and S. Fauskanger, "Policy-based labelling: A flexible framework for trusted data labelling," *International Conference on Military Communications and Information Systems*, 2015.
- [19] K. Wrona and S. Oudkerk, *Content-based protection and release architecture for future NATO networks*. in *Proc. Military Communications Conference (MILCOM)*, 2013.
- [20] Infodas, "SDoT labelling-service: XML security labels for cross-domain information exchange." [Online]. Available: http://www.infodas.de/download/SDoT_Labelling_Service_eng_130422_II.pdf
- [21] Digital nation security archive. "<http://nsarchive.chadwyck.com/home.do>". Accessed: 2015-03-26.
- [22] Abbyy. "<http://www.abbyy.com/>". Accessed: 2015-03-26.
- [23] R. Baeza-Yates, B. Ribeiro-Neto *et al.*, *Modern information retrieval*. ACM press New York, 1999, vol. 463.