

Gaining insight in domestic violence with Emergent Self Organizing Maps

Jonas Poelmans¹, Paul Elzinga³, Stijn Viaene^{1,2}, Marc M. Van Hulle⁵, Guido Dedene^{1,4}

¹K.U.Leuven, Faculty of Business and Economics, Naamsestraat 69,
3000 Leuven, Belgium

²Vlerick Leuven Gent Management School, Vlamingenstraat 83,
3000 Leuven, Belgium

Amsterdam-Amstelland Police Department, James Wattstraat 84,
1000 CG Amsterdam, The Netherlands

⁴Universiteit van Amsterdam Business School, Roetersstraat 11
1018 WB Amsterdam, The Netherlands

⁵K.U.Leuven, Laboratorium voor Neuro- en Psychofysiologie
Campus Gasthuisberg O&N2, Bus 1021, Herestraat 49
3000 Leuven, Belgium

{Jonas.Poelmans, Stijn.Viaene, Guido.Dedene}@econ.kuleuven.be
Paul.Elzinga@amsterdam.politie.nl
marc@neuro.kuleuven.be

Abstract

Topographic maps are an appealing exploratory instrument for discovering new knowledge from databases. During the past years, new types of Self Organizing Maps (SOM) were introduced in the literature, including the recent Emergent SOM. The ESOM tool is used here to analyze a large set of police reports describing a wide range of violent incidents that occurred during the year 2007 in the Amsterdam-Amstelland police region (the Netherlands). This article aims to demonstrate that the ESOM tool provides a valuable exploratory instrument for examining unstructured text in police reports. First, it is shown how ESOM was used to discover a range of new features that better distinguish domestic from non-domestic violence cases. Second, it is demonstrated how this resulted in a significant improvement in classification accuracy. Third, the ESOM tool facilitates an in-depth investigation of the nature and scope of domestic violence, which is particularly useful for the domain expert. Interestingly, it was discovered that the definition of domestic violence employed by the management was much broader than the definition employed by police officers. Fourth, the ESOM tool enables an accurate and automated assignment of either a domestic or a non-domestic violence label to unclassified cases. Finally, ESOM is a highly accurate and comprehensible case triage model for detecting and retrieving wrongly classified cases.

Keywords: Emergent Self Organizing Map (ESOM), domestic violence, exploratory data analysis, knowledge discovery in databases, text mining

1. Introduction

Policing is a knowledge intensive affair. Over the past fifteen years or so, voices have risen to make a shift from a more traditional reactive, intuition led style of policing to a more proactive intelligence led approach [25]. Intelligence Led Policing (ILP) promotes the use of factual, evidence based information and analysis to inform the police management and to guide police actions at all levels of a police organization. Specifically, the goal is to complement intuition led police actions with information coming from analyses on aggregated operational data, such as crime figures and criminal characteristics [26]. While over 80% of all information available to police organizations is in text form, analysis has to date been primarily focused on the structured portion of the available data. Only recently have the first steps for applying text mining in criminal analysis been taken. Although text mining has been identified as a promising area in the formal framework for crime data mining by Chen et al. [27], this work has hardly found its way into the mainstream scientific literature. One of the important exceptions is the paper by Ananyan [28] in which historical police reports were analyzed to identify hidden patterns.

In the case we are studying, we aim at automating the detection of domestic violence from the unstructured text in police reports. In 1997, the Ministry of Justice of the Netherlands conducted a first inquiry into the nature and scope of domestic violence [2]. It showed that 45% of the population had experienced non-incidental domestic violence at some point. For 27% of the population, the incidents even occurred on a weekly or daily basis. These gloomy statistics put the issue at the top of the political agenda. By consequence, taking firm action against this phenomenon became one of the pivotal projects of the administration of Prime Minister Balkenende of the Netherlands. ¹

Pursuing an effective policy against offenders is currently one of the top priorities of the Amsterdam-Amstelland region police department in the Netherlands [22]. Of course, to this end, being able to swiftly recognize cases of domestic violence and to label reports accordingly are of the utmost importance. Still,

¹http://www.regering.nl/Het_kabinet/Eerdere_kabinetten/Kabinet_Balkenende_II/Regeerakkoord#internelink4

this has proven to be problematic. In the past, intensive audits of police databases related to filed reports have shown that many reports tended to be wrongly classified.

To cope with this problem, several attempts have been made to develop a technique that automatically classifies cases as domestic or non-domestic violence. A multi-layer perceptron and an SVM were used, but unfortunately classification accuracy was only around 80%. Moreover, these techniques did not provide any insight into the performed classification, since they are black-boxes.

In the current paper, this problem will be tackled using a special class of topographic maps [4] called Emergent Self Organizing Maps (ESOM) [5], which are particularly suited for high-dimensional data visualization. From a practitioner's point of view, topographic maps are a particularly appealing technique for knowledge discovery in databases [12]. It performs a non-linear mapping of the high-dimensional data space to a low-dimensional one, usually a two-dimensional one, which enables the visualization and exploration of the data [9]. An Emergent Self Organizing Map (ESOM) is a very recent type of topographic map [5]. According to [17] "Emergence is the ability of a system to produce a phenomenon on a new, higher level." In order to achieve emergence, the existence and cooperation of a large number of elementary processes is necessary. An Emergent SOM differs from a traditional SOM in that a very large number of neurons (at least a few thousand) are used [6]. In the traditional SOM, the number of nodes is too small to show emergence. It will be demonstrated that, from the unstructured text in police reports, essential knowledge regarding domestic violence is obtained by using ESOM. In addition, it will be shown that an efficient, comprehensible and highly accurate automated classification model can be constructed using an ESOM.

The remainder of this paper is as follows. In section 2, an overview of the domestic violence problem area will be given. In section 3, we shall cover the essentials of topographic map theory, and in particular of the Emergent Self Organizing Maps. In section 4, the dataset used will be discussed, after which, in section 5, the ESOM application to the domestic violence problem is demonstrated. Finally, section 6 concludes the paper.

2. Identifying domestic violence

2.1. Domestic violence

According to the department of Justice and the Netherlands police, domestic violence can be characterized as serious acts of violence committed by someone from the domestic sphere of the victim. Violence includes all forms of physical assault. The domestic sphere includes all partners, ex-partners, family members, relatives and family friends of the victim. Family friends are those persons who have a friendly relationship with the victim and who (regularly) meet the victim in his/her home [1].

For the year 2007, the database of the Amsterdam-Amstelland police included more than 8000 cases that contained a statement made by the victim of a violent incident. Because it is physically impossible for any individual to process this amount of information, applying text mining technology seems a natural approach. Text mining has been defined as “the discovery by computer of new, previously unknown, information by automatically extracting information from different written resources” [15]. A pivotal step in this process is the text analysis phase. This approach has been tried out in the past, but the results were not convincing enough. This was largely due to the lack of a good thesaurus, the vague definition of domestic violence, the classification errors made by police officers and the lack of a tool that facilitates in-depth exploration of the data [29,30,31].

Documents related to certain types of crime receive corresponding labels. Immediately after a crime is reported, police officers have the opportunity to judge whether or not it is a domestic violence case. If they believe it is, they can assign the label “domestic violence” to the report. However, not all domestic violence cases are recognized as such by police officers and, by consequence, many police reports are wrongly assigned the “non-domestic violence” label (i.e. false negatives).

One of the conclusions was that there is a need for an efficient and effective case triage system that automatically filters out suspicious cases for in-depth manual inspection and classification. The in-place case triage system retrieves the filed reports that have not been assigned a domestic violence label, and in

which the perpetrator and the victim live at the same address, or that contain terms like “ex-husband”, “son”, “mom”, etc. The instrument was introduced about five years ago to substantially reduce the number of domestic violence cases that were not recognized as such. However, a large number of these retrieved cases are wrongly selected for further in-depth analysis and classification. Going back to 2007, only about 20% of the 1091 retrieved cases were reclassified as domestic violence. Given that it takes at least 5 minutes to read and classify a case, it is clear that a more accurate case triage model results in major time savings.

2.2. Human-centered knowledge discovery in domestic violence cases

The reason why ESOM was chosen is threefold. First, in the literature, the need for exploratory data analysis has often been described [14]. When beginning the analysis of a new dataset of which very little is known a priori, the first step is to explore the data. According to [18], data mining should be primarily concerned with making it easy, convenient and practical for organizations with a large number of users to explore very large databases, without requiring years of data analysis training. Unfortunately, according to [19], much attention and effort has been focused on the development of data mining techniques but only minor attention has been given to the development of tools that support the analyst in the identification process. The authors argue for a more human-centered approach. According to [20], a significant contribution to the art of data mining is the user’s intuition with respect to the tools. One of the major advantages of ESOM is that it can be used to create intuitive graphical map displays of high-dimensional datasets. These map displays can easily be used and understood by the business users - in our case the police officers - who are not likely to be experts in data analytics and statistics, but who want to improve the results of some business process along one or more dimensions [21].

Second, two types of exploratory searches can be distinguished, i.e. the *learning* and the *investigative* search. Both involve multiple iterations and return sets of objects (such as in our case: statements made by a victim to the police) that require cognitive processing and interpretation. The objectives of a learning search are: knowledge acquisition, comprehension of concepts and the discovery of the boundaries of

meaning for key concepts. Investigative search mainly focuses on discovering gaps in the existing knowledge. Both learning and investigative searching require strong human participation in a continuous and exploratory process. To effectively support the full range of search activities, humans should be given a more active role in the search process. Human-centered KDD stresses the involvement of the user in the knowledge discovery process. In most real-world knowledge discovery processes it is the analyst who explores the data and sifts through the raw data to become familiar with it and to get a feeling for what the data may cover [19]. This can only be achieved with tools offering highly interactive user interfaces that continuously engage human control over the information seeking process [13].

Finally, when a victim of a violent incident makes a statement to the police, the police officer has to judge whether domestic violence is involved. If he/she considers it to be domestic violence, he/she can assign the domestic violence label to the case. Because it is a very costly task to classify cases and to verify whether or not the performed classifications are correct, the introduction of an automated classifier would result in major savings. It is clear that high overall accuracy, a low false negative rate and a comprehensible classification model are key requirements. Comprehensibility is essential, and implies that the user understands the motivations behind the model's prediction [23]. In the domain of police investigations, the lack of comprehensibility is a major issue; it causes reluctance to use a classifier or even leads to a complete rejection of the model. This is largely due to the very high cost of classifying a case incorrectly as non-domestic violence. Clarity and explainability of the performed classification are major constraints. Comprehensibility measures the "mental fit" of the classification model [24].

3. Emergent SOM

It is claimed by Ultsch and co-workers that the topology preservation of the traditional SOM projection is of little use when the maps are small: the performance of a small SOM is argued to be almost identical to that of a k -means clustering, with k equal to the number of nodes in the map [5]. It is argued to be especially useful for visualizing sparse, high-dimensional datasets, yielding an intuitive overview of its

structure [7]. Using large numbers of neurons as in ESOM permits one to observe data at a higher level capturing the overall structures, disregarding the elementary ones and allowing the consideration of structures that would otherwise be invisible.

An ESOM map is composed of a set of neurons I , arranged in a hex-grid map structure. A neuron $i \in I$ is a tuple (w_i, p_i) consisting of a weight vector $w_i \in W$ and a position $p_i \in P$ in the map. The input space D is a metric subspace of R^n . The training set $E = \{x_1, \dots, x_k\}$ with $x_1, \dots, x_k \in R^n$ consists of input samples presented during the ESOM training. The training algorithm used is the online training algorithm in which the best match for an input vector is searched for, and the corresponding weight vectors, and also those of its neighboring neurons of the map, are updated immediately.

When an input vector x_i is supplied to the training algorithm, the weight of a neuron $n_i = (w_i, p_i)$ is modified as follows, let $\eta \in [0,1]$, then:

$$\Delta w_i = \eta \times h \times (bmi, n_i, r) \times (x_i - w_i)$$

The best-matching neuron of an input vector $x_i \in D$

$$D \rightarrow I : bmi = bm(x_i)$$

is the neuron $n_b \in I$ having the smallest Euclidean distance to x_i :

$$n_b = bm(x_i) \Leftrightarrow d(x_i, w_b) \leq d(x_i, w_b) \forall w_b \in W .$$

Where $d(x_i, w_j)$ stands for the Euclidean distance of input vector x_i to weight vector w_j . The neighborhood of a neuron

$$N_i = N(n_i) = \{n_j \in M \mid h_{ij}(r) \neq 0\}$$

is the set of neurons surrounding neuron n_i and determined by the neighborhood function h . The neighborhood defines a subset in the map space of the neurons K , while r is called the neighborhood radius.

ESOM can be used to detect clusters and maintains the neighborhood relationships that are present in the input space. It also provides the user with an idea of the complexity of the dataset, the distribution of

the dataset (e.g. spherical) and the amount of overlap between the different classes. The lower-dimensional data representation is also an advantage when constructing classifiers. Finally, only a minimal amount of expert knowledge is required for the user to be able to use ESOM effectively for exploratory data analysis. An additional advantage is that it can be trained directly on the available dataset without first having to perform a feature selection procedure [8]. ESOM maps can be created and used for data analysis by means of the publicly available Databionics ESOM Tool [16]. With this tool, the user can construct ESOMs with either flat or unbounded (i.e., toroidal) topologies.

In this paper, it is demonstrated that ESOM is one of those rare tools that meet the key requirements of Human-centered KDD. It is shown that it provides a highly interactive user interface that moves the exploratory search process beyond predictable fact retrieval.

It is also demonstrated that the ESOM tool provides the ideal exploratory instrument for supporting the critical text analysis phase. Because of its highly interactive user interface, human participation and effective use of their expert prior knowledge in the search process is promoted.

Finally, one of the most important advantages of a nearest neighbor classifier based on an Emergent Self Organizing Map is the comprehensibility of the performed classification. To answer the question why a police report was classified as domestic or as non-domestic violence, one simply needs to inspect the cases that belong to the best matched neuron(s) of the police report.

It shall be demonstrated that a nearest-neighbor classifier based on the ESOM tool meets all the requirements stated in section 2.2 and outperforms other more complex classifiers such as the SVM, Naïve Bayes, and multi-layer perceptron.

4. Dataset

The dataset consists of a selection of 4814 police reports describing a wide range of violent incidents that occurred in 2007. All domestic violence cases from that period are a subset of this dataset. This selection

came about partly by filtering out those police reports that did not contain a crime report filed by a victim, which is necessary for establishing domestic violence. This happens for example when a police officer includes their findings in a report after having been sent to an incident, , while the victim did not make an official statement to the police. The follow-up reports referring to previous cases were also removed. From the 4814 police reports contained in the dataset, the person who reported the crime, the suspect, the persons involved in the crime, the witnesses, the project code and the victim statement were extracted. Of these 4814 reports, 1657 were cases of domestic violence; the others were not. These data were used to generate the 4814 html-documents that were used during our research. An example of such a report is displayed in Figure 1.

The validation set consists of a selection of 4738 cases describing a wide range of violent incidents from the year 2006 where the victim made a statement to the police. Again, the follow-up reports were removed. 1734 of these 4738 cases were classified as domestic violence by police officers. In 2006, the in-place case triage system retrieved 1157 police reports, containing a victim statement, which had to be manually classified by police-officers. 318 reports were classified as domestic violence, while 839 were classified as non-domestic violence.

Title of incident	Violent incident xxx
Reporting date	31-03-2008
Project code	Domestic violence against ex-partner
Crime location	Amsterdam Wibautstraat yyy
Suspect (male) Suspect (18-45yrs)	Zzz
Address	Amsterdam Waterlooplein yyy
Involved (male) Involved	Neighbours

(>45yrs)

Address Amsterdam Wibautstraat www

Victim (female) Victim
Uuu

(18-45yrs)

Address Amsterdam Waterlooplein vvv

victim statement

Yesterday morning I was taking a bath. Suddenly my daughter ran into the bathroom followed by her ex-boyfriend. She screamed for help. He had a gun in his hand and he was clearly under influence of beer or drugs. He yelled that he couldn't live without her. He threatened to kill me and my daughter if she wouldn't come back to their house. The neighbours who were alarmed by all the noise came to help. Meanwhile another neighbour phoned the police. I jumped out of my bath and tried to push him to the floor. During this fight I got serious injuries on my back etc.

Figure 1. Example police report.

The initial thesaurus – a collection of terms – was obtained by performing frequency analyses on these police reports. The relevant terms that occurred most often were retrieved and added to the initially empty thesaurus. This resulted in an initial set of 123 terms. In the dataset , it is indicated for each police report which of these terms is present. An excerpt of this dataset is displayed in table 1.

Table 1. Excerpt of the categorical dataset used during the research.

	kicking	Dad hits me	Stabbing	cursing	scratching	maltreating
Report 1	X	X				X
Report 2			X	X	X	
Report 3	X	X	X	X	X	
Report 4						X
Report 5				X	X	

An initial analysis of the data revealed that the two cases (domestic and non-domestic violence) do not appear in separate clusters; hence, the use of a clustering method followed by a labeling of the high density peaks is not a viable approach (see Figure 2 where the high density regions (darker pixels) do not correspond to separate labels). Topographic maps such as ESOM may overcome this since they not use the labels but rather approximate the data manifold.

In a first step, an ESOM map with a toroidal topology of the neurons as well as a flat topology were trained using this dataset, in order to capture the distribution of the dataset. To simulate ESOM, we used the Databionics software and its standard parameter settings. We did not attempt to optimize them. A SOM with a lattice containing 50 rows and 82 columns of neurons was used ($50 \times 82 = 4100$ neurons in total). The weights were initialized randomly by sampling a Gaussian with the same mean and standard deviation as the corresponding features. A Gaussian bell-shaped kernel with initial radius of 24 was used as a neighborhood function. Further, an initial learning rate of 0.5 and a linear cooling strategy for the learning rate were used. The number of training epochs was set to 20. In the map displayed in Figure 2, the best matching (nearest-neighbor) nodes are labeled in the two classes for the given test data set (red for domestic violence, green for non-domestic violence). The red squares in all figures represent neurons that mainly contain domestic violence reports, whereas the green squares represent neurons that mainly contain non-domestic violence reports.

An analysis of the ESOM map, based on the categorical dataset, led us to conclude that it was spherically distributed. It can be seen that there is one large domestic violence cluster located at the center of the map, and a domestic violence cluster running upward and to the left of the map. The latter continues over the edge of the map (note that the map is actually toroidal) and has an outlier on the right of the map. Moreover, when the flat ESOM map was compared to the toroidal ESOM map, the toroidal map provided a much better visualization of the dataset. The border effect was clearly present in the flat map resulting in undesired distortions of the map. Most of the observed clusters were located at the border of the map, which made them smaller in area, and the large group of domestic violence cases was less clearly

demarcated from the non-domestic violence cases. Therefore, it seemed more natural to use a toroidal ESOM for visualizing this dataset.




Figure 2. Toroidal ESOM map trained on the categorical dataset with all features

The map displayed in Figure 2 was trained directly on the entire dataset with 123 features. Figure 2 clearly shows that the density profile of the ESOM map does not match the uniform distribution of the labeled data vectors. Moreover, there is no ridge in the map that separates the domestic- from the non-domestic violence cases. Therefore, the ‘watershed’ technique [6] will not lead to a correct identification of the classes. We have also applied feature selection to reduce the input space dimensionality, prior to applying the ESOM tool. A heuristic feature selection procedure, known as minimal-redundancy-maximal-relevance (mRMR), as described in [11], was considered. The aim was to select the 65 most relevant features. To obtain the optimal feature set, an SVM, a Neural Network, a kNN (with $k=3$) and a Naïve Bayes classifier were used to measure the classification performance for an increasing number of features. The classification performance is plotted as a function of the number of features in Figure 3.

We opted to retain the best 44 features. A toroidal ESOM map was trained on this dataset with a reduced number of features and was compared to that of Figure 2. However, the density problem was not solved by lowering the number of features.




Figure 3. Classification performance

5. Results

5.1 Domain exploration

An interesting result is that the map allows for a relatively clear demarcation between domestic and non-domestic violence cases. Many of the best matched nodes dominantly contain either domestic or non-domestic violence cases. This indicates that there is only a limited amount of overlap between them. The observed overlap probably indicates that the feature set is not sufficiently refined to discriminate between

the two classes. However, it should be noted that some cases might have been wrongly classified by the police officers.

It is interesting to observe that some red squares are located in the middle of a large group of green squares and vice versa. In-depth manual inspection of the police reports corresponding to these red outliers led to some interesting discoveries. Surprisingly, only a small part of the police reports turned out to have been incorrectly classified as domestic violence cases. Astonishingly, many of these reports contained multiple new, important features that were lacking in the user’s understanding of the problem area. An example of such a newly discovered feature was a homosexual relationship. The initial feature set used to train the map dominantly contained features that were specifically attuned to heterosexual relationships. The result was that police reports describing domestic violence incidents between homosexual men were located in the middle of the non-domestic violence cluster. This is important knowledge for building a highly accurate classifier. Another important feature that was discovered is pepper spray. In about 80% of the domestic violence incidents, the perpetrator is a man. By consequence, the feature set contained many terms like “kicking”, “stabbing”, “maltreating”, etc. These violent acts are mostly performed by men. The weapons that are typically used by female offenders were a blank spot in the user’s knowledge of the problem area. Again by investigating these red outliers, we were able to discover that pepper spray is one of the most frequently used weapons of female aggressors. The most important features we discovered are displayed in table 2 and 3.

Table 2. Newly discovered features by studying the domestic violence outliers in the ESOM map.

Pepper spray
Homosexual relationship, lesbian relationship
sexual abuse, incest
Alternative spelling of some words (e.g. ex-boyfriend, exboyfriend, ex boyfriend)
Violence terms lacking in the thesaurus: abduction, choke, strangle, etc.

Weapons lacking in the thesaurus: belt, kitchen knife, baseball bat, etc.

Terms referring to persons: partner, fiancée, mistress, concubine, man next door, etc.

Terms referring to relationships: romance, love affair, marriage problems, divorce proceedings, etc.

Reception centers: women's refuge center, home for battered women, etc.

Gender of the perpetrator: mostly male

Gender of the victim: mostly female

Age of the perpetrator: mostly older than 18 and younger than 45

Age of the victim: mostly older than 18 and younger than 45

Terms referring to an extramarital affair: I have an another man, lover, I am unfaithful, etc.

Table 3. Newly discovered features by studying the non-domestic violence outliers in the ESOM map.

Places of entertainment: dancing, disco, bar, etc.

Crime locations: on the street, on a bridge, under a viaduct, on a crossing, etc.

Public locations: metro station, bus stop, tram stop, etc.

Reception centers: refugee center, shelter for the homeless, relief center, etc.

Drugs: drug abuse, drug joint, etc.

Addresses of youth institutions, prisons, etc.

Hotel: hotel room, hotel, etc.

Description of suspect's origin: Turkish descent, white man, north-African descent, etc.

Description of suspect's body: 1.75 meters tall, 119 centimeters tall, muscular appearance, etc.

Description of suspect's hair: curly haired, blond hair, redhead, etc.

Description of suspect's clothes: black jacket, leather shoes, blue pants, jeans, etc.

Description of suspect's face: beard, moustache, facial hair, etc.

Description of suspect's accent

Unknown person is involved in the crime

Street robbery

Burglary
Car theft
Bicycle theft
Attack by unknown person
Moped theft
Neighborhood quarrel

The reason why some of the non-domestic violence cases, containing one or more of the features of table 4, were located among the domestic violence cases was because they often contained sentences like “I was walking with *my husband*, when we were suddenly *attacked by an unknown person*”. These sentences contain terms that regularly appear in domestic violence reports. By introducing the features presented in table 4, these cases can be better distinguished from the domestic violence cases. As a consequence, a larger number of these cases will be located among the non-domestic violence cases in the ESOM map. A new toroidal ESOM map was trained on the dataset based on the refined thesaurus. The resulting map is displayed in Figure 4.




Figure 4. Toroid ESOM map trained on the categorical dataset with all features

When the ESOM map of Figure 2 is compared to that of Figure 3, it is clear that the amount of overlap between the two classes is much lower for the map based on the refined thesaurus. Moreover, the classification accuracy of the SVM, Neural network, Naïve Bayes and kNN classifiers were significantly better after the newly discovered features were added to the thesaurus. For example, for the SVM, the best classification accuracy on the initial dataset was around 83%, while the best classification accuracy on the dataset with the refined thesaurus was around 89%.




Figure 5. Classification performance

Another interesting observation was that the outlier neurons often contained cases that cannot be uniquely classified as either domestic or as non-domestic violence on the basis of the definition. For example, in one case a twelve-year-old girl had had a relationship of three days with a boy of the same age. After the girl ended the relationship, the boy continued running after her for months. He often went to

her home and started banging on the front door; he kept sending messages to her mobile phone, etc. This case was classified as domestic violence by a police officer. However, one may question whether this is correct. Firstly, can we speak of ex-partners if the persons involved had a relationship that lasted only three days? Second, the girl herself was not physically harmed. Only her father sustained minor injuries after he was attacked by the boy. This example makes it clear that it is not always easy to distinguish domestic from non-domestic violence cases. This is often due to the vagueness of the terms contained in the domestic violence definition. Therefore, a pivotal step in the research consisted of clearly defining how broadly these terms should be interpreted and clearly demarcating the borderline between domestic and non-domestic violence. After presenting those doubtful cases for which the definition does not provide a unique classification to the board members responsible for the domestic violence policy, it became clear that the decisive factor in classifying an incident as domestic violence should be the presence of a dependency relationship between the perpetrator and the victim.

Taking this classification guideline into account, the cases that were outliers in the ESOM map were further investigated. This led to the discovery of some regularly occurring situations in which there is a clear dependency relationship between the perpetrator and the victim but that were typically classified as non-domestic violence by police officers. An excerpt of these circumstances is listed in table 4.

Table 4. Circumstances in which the offender abuses the dependency relationship with the victim, but that are not recognized by police officers as domestic violence.

Circumstance	Dependency relationship
Lover boys	The victim is in love with the lover boy, who abuses this dependency relationship to force her into prostitution
Extramarital relationship	If the mistress of an adulterer blackmails him, for example by threatening to reveal their affair to his wife, the mistress abuses the dependency relationship that exists between herself and the man.
Violence between a caretaker and an inhabitant of an institution	If the caretaker threatens or maltreats the inhabitant (for example, a nurse who maltreats an elderly woman in an old folks' home), the latter is often helpless because she depends on the caretaker.
Violence between colleagues	If two colleagues had a relationship and one keeps stalking the other, this is domestic violence between ex-persons.
An ex-boyfriend attacks the	This is considered to be domestic violence because the ex-boyfriend often aims at emotionally

new boyfriend

hurting his ex-girlfriend.

Third person makes statement to the police for somebody else

Police officers regularly fail to recognize cases in which a third person makes a statement to the police for somebody else (e.g. a father who makes a statement about the sexual abuse of his daughter by her stepfather) as domestic violence.

It was also interesting to observe that some of these outlier cases described incidents that in themselves can not be classified as domestic violence, but might be an early warning indicator for impending domestic violence. For example, one incident described an ex-boyfriend who threw a stone through the window of his ex-girlfriend's car. His ex-girlfriend was not threatened, nor physically assaulted by him, because she was not there when the incident took place; the neighbors saw it happen. This isolated incident is not domestic violence; however, it may be a prelude to an escalation of violence between the two ex-partners. Interviews with a representative number of police officers revealed that the majority would not assign a domestic violence label to this type of situation. However, according to the board members responsible for the domestic violence policy, this should be classified as domestic violence. This exposed the mismatch between the management's conception of domestic violence and the classification performed by the police officers. The definition employed by the management turned out to be much broader.

5.2 ESOM classification

To build an optimal classifier, it was necessary to verify whether or not the dataset is stationary, i.e. whether or not there are seasonal influences playing a role in the classification performed by police officers. An ESOM map was trained on the police reports from the year 2007. A kNN classifier was built for this ESOM map and k was set to 1 and 2 consecutively. In order to obtain the misclassification error of the ESOM map, the Euclidean distance of each input vector to each weight vector was measured. For each weight vector (corresponding to a neuron of the map) it was calculated how many of the domestic and non-domestic violence cases had this weight vector as a best match. If the node dominantly contained

domestic violence cases, it was labeled as a domestic violence node and the non-domestic violence cases that best matched this node were considered to be wrong classifications.

This map was used to classify the police reports from the four quarters of 2007. The results of the nearest neighbor classifiers based on the ESOM map are displayed in table 5 and 6.

Table 5. Classification accuracy of the 1 nearest neighbor classifier applied on the map trained on the dataset of the year 2007.

	Overall accuracy	False Positive Rate	False Negative Rate
Year 2007	88.3%	9.4%	16.0%
1 st quarter 2007	92.4%	10.9%	4.2%
2 nd quarter 2007	90.6%	8.3%	12.0%
3 rd quarter 2007	88.1%	10.1%	15.4%
4 th quarter 2007	89.7%	8.9%	13.0%

Table 6. Classification accuracy of the 2 nearest neighbor classifier applied on the map trained on the dataset of the year 2007.

	Overall accuracy	False Positive Rate	False Negative Rate
Year 2007	85.2%	9.9%	23.9%
1 st quarter 2007	87.6%	15.5%	9.3%
2 nd quarter 2007	86.1%	12.3%	17.3%
3 rd quarter 2007	82.5%	13.7%	24.9%
4 th quarter 2007	85.9%	10.5%	22.1%

From table 5 and table 6, one may conclude that the overall accuracy of the 1NN classifier based on the ESOM map is better than the overall accuracy of the 2NN classifier based on the same ESOM map. It is clear that there are only minor differences in the classification accuracy on the four datasets. Therefore, it is a logical choice to put the datasets of the four quarters of 2007 together in one dataset consisting of 4814 police reports.

An interesting result is the difference in performance of the traditional kNN classifier on the original data (around 83%) and the kNN classifier based on the toroidal ESOM map (around 90%). This is due to the topographic map being a model of the data distribution: it forms an approximation of the data manifold, offering interpolating facilities, and it spends more neural hardware at clusters in the data, leading to a modeling of the local density.

The ESOM's modeling of the data manifold did not vary very much for different dimensionalities of the dataset. When the classification accuracy of the NN classifiers based on the ESOM map with a reduced number of features was compared to the classification accuracy of the NN classifiers based on the map trained on the dataset with all features (displayed in Figure 4), it turned out that the performances were very similar. Moreover, when the maps were visually inspected, the amount of overlap between the different classes was very similar for both maps.

It should be noted that more complex classifiers such as the SVM did not perform better than the ESOM, and that the previously developed system were multi-layer perceptrons, which did not provide any insight into the problem (since it is a black-box), and their performance was only around 80% .

5.3 Improvement of case triage system

We aimed at automatically classifying the output of the in-place case triage system using ESOM. The approach used to tackle this problem was as follows. First, a toroid ESOM map was trained on the entire dataset of the year 2007, including the retrieved cases. The latter were taken into account during training in order to make sure that the data distribution was modeled correctly. Then, the non-retrieved cases were shown as green and red dots. Finally, the cases retrieved by the triage system were made visible on the map by assigning a blue color to the neurons containing one or more of these cases.

When the cases retrieved by the in-place triage system were made visible on the ESOM map, it became clear that many of them hit a map neuron that did not contain any other cases. The assignment of a label to the police reports could be performed in three different ways; first, by using a kNN classifier based on the ESOM map. Second, the classification could be based on the map coordinates of the neuron containing retrieved cases. The number of domestic and non-domestic violence cases in the neighboring neurons is counted and the most frequently occurring label is assigned to the retrieved case. The advantage of these two approaches is that they can be performed automatically. However, in comparison with the first approach, the second one makes better use of the neighborhood relationships available in the map. Third, the label assignment can be performed based on the visual display of the map. For example, if a retrieved case is located in a region that dominantly contains domestic violence cases, the user can decide to classify the case as domestic violence. This approach has the important advantage that if a case

is located in a region containing a mixture of class labels, it can be left unclassified, resulting in a lower misclassification error.




Figure 6. Toroid ESOM map with retrieved cases made visible

The classification results are displayed in table 7. It is clear that the classification accuracy of the nearest-neighbor classifier and the classifier based on map coordinates were very similar. For both classifiers, the optimal classification accuracy was only around 76%. Unfortunately, a misclassification rate of 24% was not acceptable. Therefore, the classification behavior of the ESOM tool will be analyzed in more detail in the next section, resulting in a more fine grained and highly accurate case triage model.

Table 7. Classification accuracy of the nearest neighbor classifiers and classifiers based on map coordinates applied on the cases retrieved by the case triage system in the year 2007.

	Overall accuracy	False Positive Rate	False Negative Rate
1 nearest neighbor	69.3%	35.6%	18.3%
2 nearest neighbor	68.8%	37.1%	16.5%
3 nearest neighbor	72.0%	31.1%	20.1%
4 nearest neighbor	73.3%	29.6%	19.2%
5 nearest neighbor	73.7%	28.2%	21.5%

6 nearest neighbor	75.8%	26.1%	19.2%
7 nearest neighbor	75.6%	26.1%	19.7%
4 map coordinates	68.3%	37.3%	17.8%
8 map coordinates	70.7%	33.3%	19.2%
12 map coordinates	72.0%	31.3%	19.7%
16 map coordinates	73.5%	29.4%	19.2%
20 map coordinates	75.1%	27.4%	18.8%
24 map coordinates	75.1%	26.9%	20.1%
28 map coordinates	75.7%	26.3%	19.2%
32 map coordinates	75.9%	25.7%	20.1%
36 map coordinates	75.5%	26.5%	19.7%

Finally, the results were validated on a test set from the year 2006. This showed that the maps and classification performances were similar.

5.4 Risk analysis

Although the false negative rate is acceptable, it should be stressed that false negatives are critical. In practice a false negative is much worse than a false positive. Therefore, the next step consisted of developing a more fine-grained nearest neighbor classifier. Because a comprehensible visualization of the classification is of the utmost importance, this risk analysis was embedded in the ESOM map. The optimal classification accuracy based on the map coordinates was obtained using a range of 32 neighboring neurons. It was chosen to use this range for the construction of the risk analysis map. For each neuron, the number of domestic and non-domestic violence cases contained in the neuron and the 32 surrounding neurons was counted and used to calculate the probability that a police report that has this neuron as a best match described a domestic violence incident. For the visualization, a color scale

consisting of 5 shades of color was used. Red indicates a 90-100% probability of domestic violence, orange a 70-90% probability, yellow a 30-70% probability, green a 10-30% probability and dark green a 0-10% probability. The labels of the cases retrieved by the in-place case triage system were not used to construct this risk analysis map. The resulting map for the dataset of the year 2007 is displayed in Figure 5.




Figure 8. ESOM map with risk analysis included

This map allows a more fine-grained nearest neighbor classification. The cases retrieved by the in-place case triage system were projected onto the risk analysis map. Cases that were reclassified by police officers as domestic violence were shown as black dots, while the cases of which the original non-domestic violence label was not changed, were shown as light blue dots. The resulting map is displayed in Figure 9.




Figure 9. ESOM map with risk analysis included and retrieved cases made visible

It was remarkable to observe that many of the retrieved cases were located in the red and orange area of the map, but were not reclassified by police officers as domestic violence. About 16.5% of the retrieved cases were located in the red area and about 15.5% in the orange area of the map displayed in Figure 9. About 27.2% of the cases located in the red area and 49.6% of the cases located in the orange area of the map were classified as non-domestic violence by police officers. In-depth analysis of these police reports revealed that the majority of these cases should have been classified as domestic violence. On the other hand, only a small percentage of the cases located in the dark green and green areas of the map were reclassified as domestic violence by police officers (8.9% and 7.5% respectively). Inspection revealed that all of these reclassified cases actually described non-domestic violence incidents.

Table 9. Distribution of retrieved cases over different map areas

Domestic violence probability	Map area color	% of retrieved cases located in map area	% classified as domestic violence	% classified as non-domestic violence
0-10%	dark green	20.7%	8.9%	91.1%

10-30%	green	26.4%	7.5%	92.5%
30-70%	yellow	20.9%	24.6%	75.4%
70-90%	orange	15.5%	50.2%	49.6%
90-100%	red	16.5%	72.8%	27.2%

Based on the map displayed in Figure 9, a correct label can be automatically assigned to 79,1% of the retrieved cases (i.e. the cases located in the dark green, green, orange and red areas of the map), while the remaining cases (i.e. the cases located in the yellow area of the map) have to be classified manually. This is a major improvement over the previous situation where each retrieved case had to be dealt with manually.

6. Conclusions

Intensive audits of police databases revealed that many police reports tended to be wrongly classified as domestic or as non-domestic violence cases. In this paper, it has been shown that the ESOM tool is an ideal instrument for in-depth analysis of domestic violence. It was used to discover new features that better distinguish domestic from non-domestic violence cases resulting in a higher classification accuracy. Moreover, it proved to be a useful tool for analyzing the domestic violence definition. A mismatch was found between the management's conception of domestic violence and the classification as performed by police officers. We found that police officers generally employed a much narrower domestic violence definition than the management. Additionally, using the ESOM tool, a number of regularly occurring situations that were often wrongly classified as non-domestic violence by police officers (e.g. lover boys, etc.) were found. Subsequently, the ESOM tool was used to build an accurate, comprehensible and automated classifier. The tool was used to classify the cases retrieved by the in-place case triage system. Finally, in order to improve classification performance, a risk analysis map was constructed to better understand and refine the performed classification. In addition, a large number of incorrectly classified cases were detected and corrected. Topics for future research include the optimization of the ESOM tool's parameter settings.

7. Acknowledgements

The authors would like to thank the Amsterdam-Amstelland region police department and in particular Deputy Chief Reinder Doeleman and Chief Hans Schönfeld for supporting this research. The authors are grateful to the Amsterdam-Amstelland Police Department for providing us with the data. JP is supported by a research chair given to SV by the Amsterdam-Amstelland Police Department.

8. References

- [1] Keus, R., Kruijff, M.S. (2000) Huiselijk geweld, draaiboek voor de aanpak. Directie Preventie, Jeugd en Sanctiebeleid van de Nederlandse justitie.
- [2] Van Dijk, T. (1997) Huiselijk geweld, aard, omvang en hulpverlening (Ministerie van Justitie, Dienst Preventie, Jeugdbescherming en Reclassering, oktober 1997)
- [3] Ritter, H. (1999) Non-Euclidean Self-Organizing Maps, pages 97–109. Elsevier, Amsterdam.
- [4] Kohonen, T. (1982), “Self-Organized formation of topologically correct feature maps”, *Biological Cybernetics*, Vol. 43, pp 59-69.
- [5] Ultsch, A., Moerchen, F. (2005) ESOM-Maps: Tools for clustering, visualization, and classification with Emergent SOM. Technical Report Dept. of Mathematics and Computer Science, University of Marburg, Germany, No. 46
- [6] Ultsch, A., Hermann, L. (2005) Architecture of emergent self-organizing maps to reduce projection errors. In Proc. ESANN 2005, PP1-6
- [7] Ultsch, A. (2004) Density Estimation and Visualization for Data containing Clusters of unknown Structure. In proc. GfKI 2004 Dortmund, pp 232-239
- [8] Ultsch, A. (2003) Maps for visualization of high-dimensional Data Spaces. In proc. WSOM'03, Kyushu, Japan, pp. 225-230
- [9] Ultsch, A., Siemon, H.P. (1990) Kohonen's Self Organizing Feature Maps for Exploratory Data Analysis. Proc. Intl. Neural Networks Conf., pp305-308
- [10] Van Hulle, M. (2000) Faithful Representations and Topographic Maps from distortion based to information based Self-Organization. Wiley: New York
- [11] Peng, H., Long, F., Ding, C. (2005) Feature Selection Based on Mutual Information: Criteria of Max-Dependency, Max-Relevance, and Min-Redundancy. *IEEE Transactions on pattern analysis and machine intelligence*, Vol. 27, no. 8.
- [12] Argamon, S., Olsen, M. (2006) Toward meaningful Computing, *Communications of the ACM*, Vol. 49, no. 4
- [13] Pednault, E.P.D. (2000) Representation is everything. *Communications of the ACM*, Vol. 43, no. 8
- [14] Marchionini, G. (2006) Exploratory search: from finding to understanding. *Communications of the ACM*, Vol. 49, no. 4
- [15] Fan, W., Wallace, L., Rich, S., Thang, T. (2006) Tapping the power of text mining. *Communications of the ACM*, Vol. 49, no. 9
- [16] <http://databionic-esom.sourceforge.net/>
- [17] Ultsch, A. (1999) Data mining and knowledge discovery with Emergent SOFMS for multivariate Time Series. In Kohonen Maps, pp. 33-46
- [18] Fayyad, U., Uthurusamy, R. (2002) Evolving data mining into solutions for insights. *Communications of the ACM*, Vol. 45, no. 8
- [19] Brachman, R., Anand, T. (1996) The process of knowledge discovery in databases: a human-centered approach. In *advances in knowledge discovery and data mining*, ed. U. Fayyad, G. Piatetsky-Shapiro, P. Smyth and R. Uthurusamy. AAAI/MIT Press
- [20] Smyth, P., Pregibon, D., Faloutsos, C. (2002) Data-driven evolution of data mining algorithms. *Communica*

- tions of the ACM, Vol. 45, no. 8.
- [21] Kohavi, R., Rothleder, N.J., Simoudis, E. (2002) Emerging trends in business analytics. *Communications of the ACM*, Vol. 45, no. 8.
 - [22] <http://www.politie-amsterdam-amstelland.nl/get.cfm?id=86>
 - [23] Kodratoff, Y. (1994) The comprehensibility manifesto. *KDD nuggets* (94:9)
 - [24] Maimon, O.O., Rokach, L. (2005) *Decomposition methodology for knowledge discovery and data mining: the theory and data mining: theory and applications (machine perception and artificial intelligence)*. World scientific publishing company.
 - [25] Collier, P.M. (2006) Policing and the intelligent application of knowledge. *Public money & management*. Vol. 26, No. 2, pp. 109-116.
 - [26] Collier, P.M., Edwards, J.S. and Shaw, D. (2004) Communicating knowledge about police performance. *International Journal of Productivity & Performance Management*. Vol. 53, No. 5, pp. 458-467
 - [27] Chen, H., Chung, W., Xu, J.J., Wang, G., Qin, Y., Chau, M. (2004) Crime data mining: a general framework and some examples. *IEEE Computer*, April 2004.
 - [28] Ananyan, S. (2002) Crime Pattern Analysis Through Text Mining. *Proceedings of the Tenth Americas Conference on Information Systems*, New York, New York, August 2004.
 - [29] Raaijmakers, S.A., Kraaij, W., Dietz, J.B. (2007) Automatische detectie van huiselijk geweld in processen-verbaal. TNO-rapport 34293.
 - [30] Elzinga, P. (2006) Textmining by fingerprints. *Onderzoeksrapport huiselijk geweld zaken*. IGP project Activiteit 0504
 - [31] Kraaij, W., Raaijmakers, S., Elzinga, P. (2008) Maximizing classifier utility for a given accuracy. *TNO Information and Communication technology*.