HOSTED BY
ELSEVIER

Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/jides

Using neural networks to aid CVSS risk aggregation — An empirically validated approach

Alexander Beck[a], *Stefan Rass*[b,*]

[a] *VW Financial Services AG, Gifhorner Strasse 57, Braunschweig, Germany*
[b] *Universität Klagenfurt, Universitätsstrasse 65-67, 9020 Klagenfurt, Austria*

HIGHLIGHTS

- A method for automated CVSS risk aggregation is proposed.
- The aggregation can be tailored/trained to domain expertise and uncertain knowledge.
- Results have been verified along an empirical study.
- A method to reduce answer variability and ambiguity in empirical CVSS risk assessments is described.

ARTICLE INFO

Article history:
Received 14 July 2016
Accepted 31 October 2016
Published online 23 November 2016

Keywords:
Risk management
Neural network
Data aggregation

ABSTRACT

Managing risks in large information infrastructures is often tied to inevitable simplification of the system, to make a risk analysis feasible. One common way of "compacting" matters for efficient decision making is to aggregate vulnerabilities and risks identified for distinct components into an overall risk measure related to an entire subsystem and the system as a whole. Traditionally, this aggregation is done pessimistically by taking the overall risk as the maximum of all individual risks, following the heuristic understanding that the "security chain" is only as strong as its weakest link. As that method is quite wasteful of information, this work proposes a new approach, which uses neural networks to resemble human expert's decision making in the same regard. To validate the concept, we conducted an empirical study on human expert's risk assessments, and trained several candidate networks on the empirical data to identify the best approximation to the opinions in our expert group.

1. Introduction

Risk management is among the core duties of the general steering in large companies. While financial risk management enjoys a comprehensive set of helpful tools and methods, security risk management until today appears to widely rely on heuristics, (subjective) human expertise and common practice knowledge. Likewise, compiling vulnerabilities, known problems and security issues of components into a concise risk report for decision making

Peer review under responsibility of Qassim University.
* *Corresponding author.*
E-mail addresses: alexander.beck@vwfs.com (A. Beck), stefan.rass@aau.at (S. Rass).

http://dx.doi.org/10.1016/j.jides.2016.10.002

is a matter of simplifying and quantifying a situation and its impact, to make security manageable. Especially the quantification is herein a central and crucial issue, as security is a cost-benefit consideration, and quantitative measures of security are hard to define soundly. Most of the related difficulty comes from the inherent complexity of contemporary information and communication technology (ICT), which makes a hierarchical decomposition of a system into smaller subsystems necessary. Upon such a decomposition, a risk assessment can be applied, either top-down (in which case the overall risk is broken up into individual risks for subsystems), or bottom-up, when risks that are specific for limited scopes are put together into a risk picture of the bigger system. This aggregation is then iterated along the hierarchical decomposition up to the top, where the final result on the risk can be reported to decision makers for the daily business of risk control. Unfortunately, the precise process of how to aggregate risks is rarely well documented nor comprehensively studied or understood (from a psychological perspective), so most of this labor is done using rules-of-thumb. More importantly, the specific ways in which risk is aggregated is often quite context dependent. Today, these dependencies have led to a large volume of best-practices relating to many diverse domains. Risk management standards are in their core a compilation of such best practices that have been abstracted to make it amendable to the specific situation at hand. This work is an extended version of Beck and Rass [1], where a first step towards a general and flexible risk aggregation rule has been proposed. One of the few related existing such general rules to aggregate risks is the "maximum principle" (cf. section 4.3.3. in BSI [2]), which prescribes to take the vulnerability of a (sub) system as the maximum vulnerability of any of its components (herein, "vulnerabilities" are quantified as likelihoods for failure upon any attack from a known and a-priori identified set of threats).

Obviously, this approach is wasteful on information and pessimistically overestimates the risk, so that risk experts tend to refine a so-obtained first guess using their own expertise and experience. The problem that motivated this research, was an automated aid for risk assessment and decision support by "approximating" human decision making. We propose doing so by using neural networks (alternatives are discussed in Section 1.2). Our contribution is a concrete neural network (NN) trained on empirical findings from a study that queried risk experts on several scenarios, asking for their informed opinion about the overall risk as they would assess it in a real process.

1.1. *Motivation by example*

As a simplified example, consider a subsystem in an enterprise infrastructure model, composed from two representations, given as Figs. 1 and 2. First, we have a physical dependency model of applications on components (Fig. 1), which is augmented by the logical dependency model of applications on one another (Fig. 2). The risk analysis is usually done in a bottom-up fashion. That is, the vulnerability of application A is influenced by the security of its (indirect) ancestor nodes VM_1, VM_2 and their parent AS_2. Normally, we need

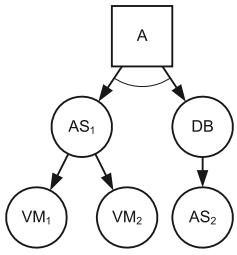

Fig. 1 – Dependencies of applications on physical components.

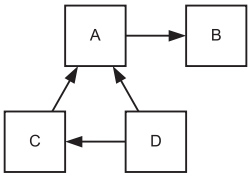

Fig. 2 – (Logical) Interdependencies between applications.

to account for "and/or"-dependency relations, if an application depends on any ("or") or all ("and") shown components. Various industrial standards can help with the assessment, and our pick in this work is the common vulnerability scoring system (CVSS; see first.org[3]). Let CVSS(X) denote the 12th dimensional (real-valued) scoring assigned to component X that results from the expert rating the CVSS criteria related to component X in terms of CVSS.[1] So, the risk assessment on application A would start with $\text{CVSS}(VM_1)$, $\text{CVSS}(VM_2)$. These two vectors would then go into the assessment $\text{CVSS}(AS_1)$. However, the assessment *cannot* straightforwardly take the maximum of the children's assessments (in a naive attempt to model the "OR-branch" of AS_1 into VM_1, VM_2), since the expert has to take into account switching times between the working and the fallback virtual machine, as well as characteristics of AS_1 that are intrinsic to the application server itself. Therefore, the assessment $\text{CVSS}(AS_1)$ only partially but not exclusively depends on $\text{CVSS}(VM_1)$ and $\text{CVSS}(VM_2)$. At this stage, most standard risk management methods hit their limits and leave the consideration of the relevant information up to the expert. In our case, this means casting the scores $\text{CVSS}(VM_1)$, $\text{CVSS}(VM_2)$ and the information known about AS_1 into a scoring $\text{CVSS}(AS_1)$. Normally, this is a nontrivial and fuzzy process.

Abstractly, the risk expert's task is traversing the graph bottom-up, where at node AS_1, his duty is to evaluate $\text{CVSS}(AS_1) = f(\text{CVSS}(VM_1), \text{CVSS}(VM_2))$, additional information about AS_1, where the function f here represents her/his expertise, experience and general/personal method to assess the vulnerability for the application server AS_1. This process is nontrivial to automate, since it assumes the graph to be acyclic, and a straightforward bottom-up aggregation would implicitly assume each node to appear exactly once in the

[1] Note that CVSS does only address confidentiality, integrity and availability. Accounting for Authenticity and other security goals is up to a manual addition to the risk management process that we do not discuss here.