Visual Risk Specification and Aggregation

Jasmin Wachter, Thomas Grafenauer, Stefan Rass
Institute of Applied Informatics, System Security Group
Universität Klagenfurt
email: {jasmin.wachter, thomas.grafenauer, stefan.rass}@aau.at

***Abstract*—Quantitative risk assessments are commonly based on estimates of impacts and likelihoods regarding threats. Both quantities are usually uncertain, subjective and therefore difficult to estimate objectively and reliably. To ease the matter, assessments are often done in categorical terms, which avoids the issue of finding numeric figures where there is typically no accuracy, but at the same time makes an expression of uncertainty more difficult. If, for an impact or the likelihood, two categories apply (not necessarily to an equal extent) or neither of the offered options is a good match, how can an expert express this kind of uncertainty or fuzzyness? Moreover, how should we deal with multiple diverging opinions on the same risk? We propose a graphical approach to tackle both issues on a single ground, by casting a common visual risk representation form into a visual risk specification system. The proposed method aids the specification of risk parameters under uncertainty, as well as opinion pooling based on the so-obtained results.**

Keywords–uncertainty representation; expert elicitation; risk assessment; opinion pooling.

I. INTRODUCTION

The quantitative specification of risks typically involves stating beliefs about impact and likelihood of a given incident. Both such specifications strongly depend on domain expertise and can usually not be described in fixed terms. Instead, the recommended way of quantifying likelihoods and impacts is based on a few (commonly three to six) categories whose textual description is matched against the current incident or threat description. Treating impact and likelihood categories as defining a cartesian coordinate system, we arrive at the well-known risk matrices, which help prioritizing risks along the +45 degrees diagonal from lower risks (events with low impact and low likelihood) up to high priority risks with significant impact and large likelihood. An example of this technique is displayed in Figure 1.

Mostly, these pictures appear in later stages of a risk management process, at the risk evaluation stage when the relevant threats have been identified and classified in both dimensions. The specification of impacts and likelihood is done a priori, and not regulated to happen in any particular form by any standard (as ISO31000 [2], or its relatives [3] [4]). Neither are matters of consensus finding and opinion pooling subject of a deeper discussion or detailed recommendations. A suitable method for such data aggregation is the second contribution of this work.

While using an illustration like Figure 1 as an output format, why not use the same form of graphical display to *input* the same values in first place? In other words, when an expert is polled regarding its opinion about a given threat, this person will see which category describes best the threat regarding its impact and likelihood, and utter the respective categories as the risk assessment. It can hardly be expected that the ultimate choice is perfect, and there may be an almost equally good alternative category to describe the matter. The idea put forth in this work is letting the domain expert not point to a single category, but rather allow marking a whole range along both axes, to express uncertainty, or (in a different view), an "overlapping" membership to the categories at hand.

Such a flexible specification appears beneficial for several reasons:

- The expert is not forced to choose a specific category, possibly issuing a caveat regarding other alternative choices,
- The expert has an intuitive way of expressing uncertainty in the overall opinion, regarding both dimensions.

Organisation of this work: Section II puts this work in the context of selected existing risk management literature. Section III describes the visual method to specify risks, and Section IV develops an algorithm to compile several assessments (based on the previous input method) into a single risk estimate. Conclusions and an outlook to future work are given in Section V.

II. RELATED WORK

Though purely quantitative risk assessment is sometimes discouraged [5], an assessment in qualitative (categorical) terms is nevertheless standard in almost all risk management approaches (as [3] [2] [4] and many more). A typical issue with any such assessment is the specific domain [6] [7], different a priori knowledge of the involved experts as well as their risk attitudes, incentives [8] and personal history that all play a strong role in how risk is perceived (and hence assessed). Interestingly (though perhaps not too surprisingly),

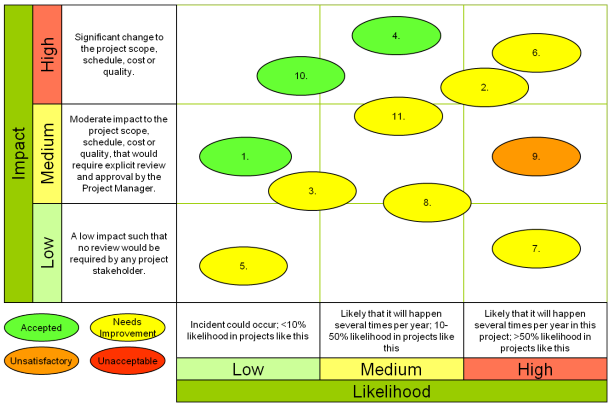

Figure 1. Example of Risk Bubble Chart [1]

 ISBN: 978-1-61208-582-1

the personality itself has only a relatively minor impact on how risk is assessed, as some empirical studies investigated [9]. Designing good questionnaires for empirical investigations is a challenging issue on its own, but left mostly unconstrained and without much explicit recommendations in risk management applications and standards. Likewise, the problem of consensus finding and compiling multiple opinions into a representative value received interest as an isolated problem [10] [11] [12], but should be an intrinsic part of the risk management process [13]. This is the gap that this work aims to fill, by proposing a first step towards a graphical way of risk specification as an alternative to existing textual and discussion based ways of getting these values. This step is mostly left open and a degree of freedom in the instantiation of various risk management methods [14] [15] [2]. Our work is intended as an auxiliary tool when using such standards.

III. Visual Risk Specification

To put this idea to work, we directly cast Figure 1 into an input system for risks, where the expert – upon speaking about a given threat – can simply draw a rectangle within 2D-area spanned by the categorical axes, where the projections onto the horizontal and vertical axis mark the matching categories. The extent of coverage expresses the degree of match, and the width/height of the rectangle corresponds to the uncertainty in the assessment (in both dimensions). Figure 2 shows an example of this technique.

Naturally, this process results in not only two but four values, which we denote as $impact_{\min}, impact_{\max}$ and $likelihood_{\min}, likelihood_{\max}$, and abbreviate as $i_{\min}, i_{\max}, \ell_{\min}, \ell_{\max}$. Both define ranges in which the expert considers the respective quantity to fall into. Constructing a statistical model from this information is straightforward: for analytic convenience, let us suppose that the expert's assessment and uncertainty is expressible by a Gaussian distribution, then based on these four values, the risk assessment would come to two Gaussians, denoted as X_I for the impact, and X_L for the likelihood, with distributions

$$X_I \sim \mathcal{N}\left(\frac{1}{2}(i_{\max}+i_{\min}), \frac{1}{3}(i_{\max}+i_{\min})\right), \quad (1)$$

$$X_L \sim \mathcal{N}\left(\frac{1}{2}(\ell_{\max}+\ell_{\min}), \frac{1}{3}(\ell_{\max}+\ell_{\min})\right), \quad (2)$$

where $X \sim \mathcal{N}(\mu, \sigma)$ denotes the distribution of the random variable with mean μ and standard deviation σ. Our choice makes the well-known 99.73% of probability mass of the Gaussian distribution fall into the given range, leaving a small residual inaccuracy allowance in the assessment. The overall uncertainty in the risk assessment is reflected in the area of the specified box; the larger the box, the less certain is the risk assessment.

Outlier Elimination

When compiling a risk picture, it is often useful to apply occasional corrections when risks are implausibly assessed relative to each other. Manually, this can be done by placing all boxes into the same picture to see outliers or do a fine-correction of risks in light of one another. Figure 3 shows an example.

IV. Pooling Several Expert Opinions

When considering several domain experts' opinions it can be a complex and tiresome task to agree upon a common risk quantity. Especially when data are sparse and risk assessments do not coincide, aggregating the final risk parameters can be challenging. Communicative methods, such as the Delphi technique or time-consuming meetings with discussion often do not lead to a consensus. Instead, mathematical pooling functions and formulas are employed to merge the opinions to a single value. This method called mathematical opinion pooling has a long tradition in statistics concerning forecast combination as well as decision making. There exist a large number of approaches and opinion pooling formulas, which can be found in [10] [11].

The easiest and most straight forward way of opinion pooling is done by simply averaging over all values, i.e., by computing the arithmetic mean. This approach is widespread and in practice often implemented blindly, as many decision makers are not aware there exist severe drawbacks of the arithmetic mean when dealing with expert opinions.

First of all, the arithmetic mean is very sensitive to outliers – especially when the sample size is small. A single extreme data point might cause a remarkable shift in the aggregated value and hence might distort the final result. Therefore, depending on the data, robust approaches and/or outlier detection and correction prior to risk aggregation should be considered. Secondly, when data are sparse smoothing might lead to more stable estimates and should thus not be neglected when aggregating data. Thirdly, the different levels of expertise and knowledge of the individual experts and their level of assurance or uncertainty regarding their risk quantity statements need to be taken into account. The arithmetic mean lacks all of the above points, yet they are crucial to the validity of the pooled result and thus need to be considered when aggregating individual expert opinions.

We, therefore, propose an intuitive iterative opinion pooling scheme that considers all aspects mentioned before. We remark that opinion pooling is generally a lossy form of data aggregation, in opposition to *lossless aggregation*, where the full data defines a whole distribution object. Decision theory in this generalized setting rests on stochastic orders, and comes with the appeal of inherently avoiding the aforementioned problems of consensus finding. Expanding this alternative branch of theory is, however, beyond the scope of this work (see [16] [17] for example).

A. Iterative Opinion Pooling Method

The input system for risk assessment described in Section III serves to specify the parameters of two Gaussian distributions – one for the impact X_I, and one for the likelihood X_L – with parameters $\mu_i = \frac{1}{2}(i_{\max}+i_{\min})$, $\sigma_i = \frac{1}{3}(i_{\max}+i_{\min})$, and $\mu_\ell = \frac{1}{2}(\ell_{\max}+\ell_{\min})$, $\sigma_\ell = \frac{1}{3}(\ell_{\max}+\ell_{\min})$ respectively. Thus, after all N experts contibuted with their risk assessment, four vectors of length N are obtained: $\boldsymbol{\mu}_i = (\mu_{i1}, \ldots, \mu_{iN})$, $\boldsymbol{\sigma}_i = (\sigma_{i1}, \ldots, \sigma_{iN})$ regarding the impact, and $\boldsymbol{\mu}_\ell$, $\boldsymbol{\sigma}_\ell$ for the likelihood respectively. For simplicity reasons, we will now drop the subscripts i or ℓ, as impact and likelihood will be pooled separately. The aim is to separately aggregate the impact and likelihood estimates, i.e., to obtain two Gaussian distributions representing the final risk distribution for the impact and the likelihood of a certain threat.

 ISBN: 978-1-61208-582-1

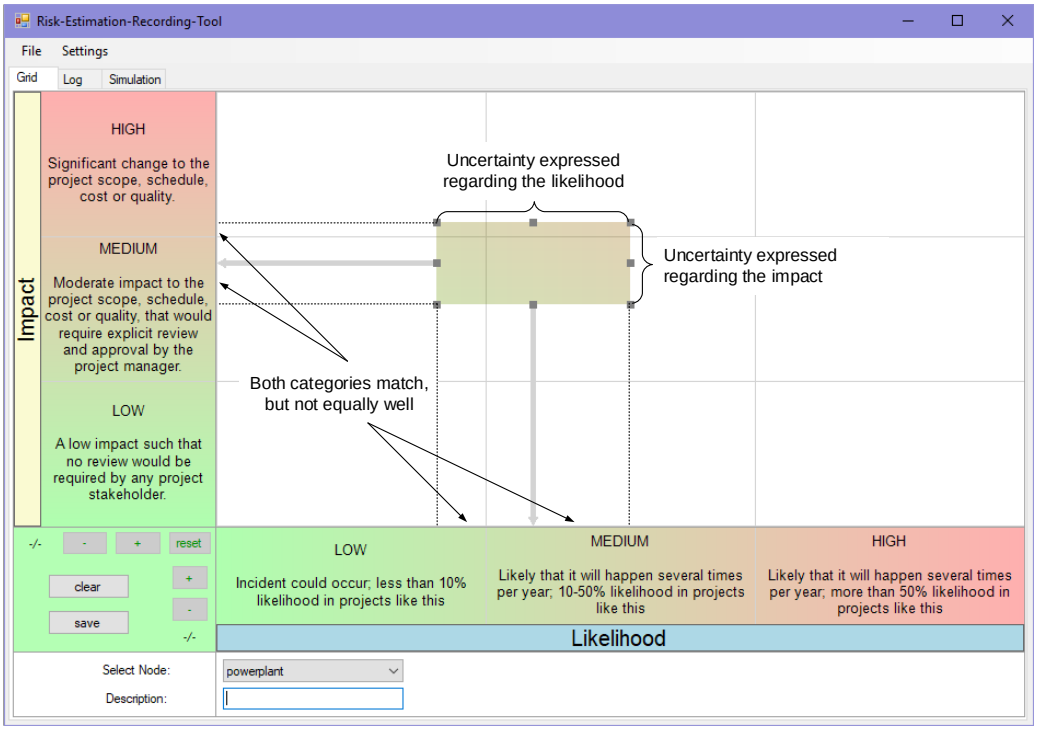

Figure 2. Graphical Risk Specification

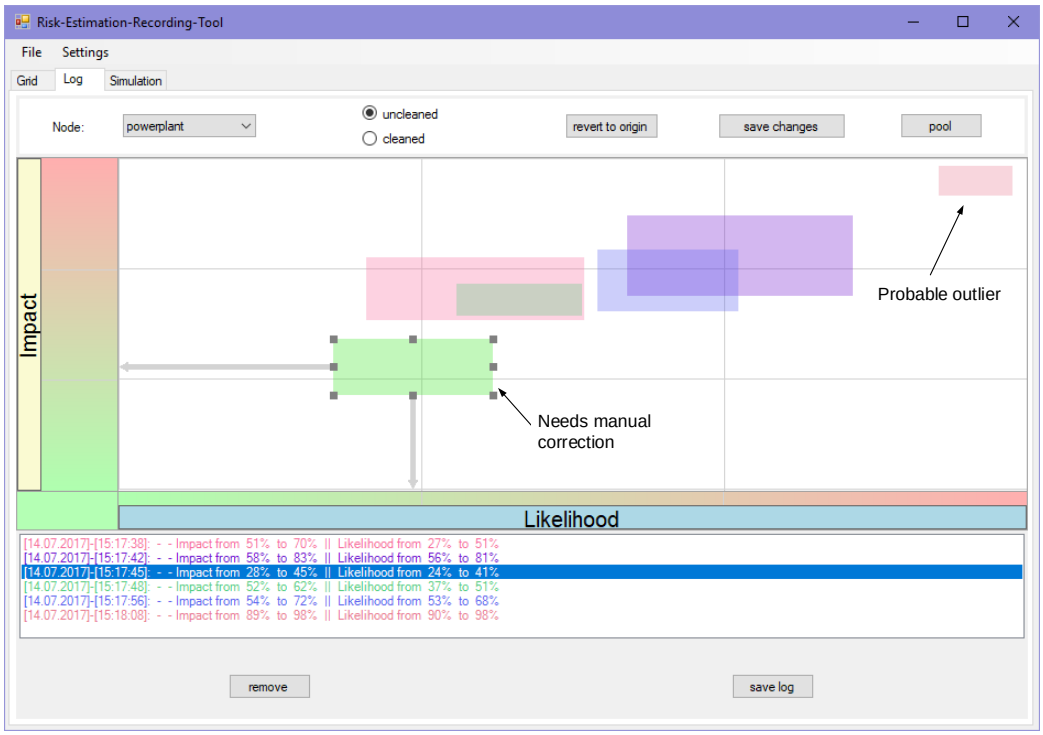

Figure 3. Manual Corrections

A possible solution to this is to consider the situation in a Bayesian framework: each expert $j \in 1, \ldots, N$ regards his estimates of the parameter of interest (e.g., impact) as prior knowledge of the parameter. Thus, the prior distribution hyperparameters are $\mu \overset{prior}{\sim} \mathcal{N}(\mu_j, \sigma_j)$. Expert j interprets the remaining experts distributions $(\mu_k, \sigma_k)\,,\ k \in \{1, \ldots, j-1, j+1, \ldots N\}$ as independent observations which make up the likelihood function. Applying Bayes rule, the posterior distribution of μ has the parameters

$$\sigma^p = \frac{1}{\sigma_j} + \sum_{k \neq j} \frac{1}{\sigma_k}, \tag{3}$$

$$\mu^p = \mu_j \cdot \frac{\sigma^p}{\sigma_j} + \sum_{k \neq j} \mu_k \cdot \frac{\sigma^p}{\sigma_k}, \tag{4}$$

and $\mu \overset{posterior}{\sim} \mathcal{N}(\mu^p, \sigma^p)$. Note that for symmetry reasons all σ^p and μ^p are the same, no matter which expert rating $j \in \{1, \ldots, N\}$ is chosen as prior distribution. Thus, the expert's posterior distribution represents the aggregated distribution for the quantity of interest. This way, each expert's (un)certainty regarding their risk assessment is incorporated in the pooling process. Hence, it is ensured that risk estimates with very high levels of assurance are given more weight than those having very low levels of assurance.

Although this method is quite intuitive and possesses many convenient mathematical properties, it does not incorporate any kind of smoothing to the data.

An alternative method, which is described as consensual opinion pooling in [12], iteratively smooths the data with a discrete inverse distance kernel until convergence to the same value. Epistemically, their procedure can be interpreted in the following way: in every iteration, each expert updates their belief about the unknown parameter by incorporating information of all experts (including themselves). Therefore, in every iteration t, each expert $j \in \{1, \ldots, N\}$ updates their belief μ_j on μ as a linear combination of all risk assessments: $\mu_j^{(t)} = \sum_{k=1}^{N} c_{kj}^{(t)} \cdot \mu_k^{(t-1)}$ with $c_{kj}^{(t)}$ inversely proportional to the distance of $\mu_k^{(t-1)}$ and $\mu_j^{(t-1)}$,

$$c_{kj}^{(t)} = \frac{\alpha_j^{(t)}}{\epsilon + d(\mu_k^{(t-1)}, \mu_j^{(t-1)})} \quad \text{with} \quad \alpha_j^{(t)} = \frac{1}{\sum_{k=1}^{N} c_{kj}^{(t)}} \tag{5}$$

and $\epsilon > 0$. This way, each expert assigns more weight to those experts, whose risk assessment are close to their own, than to experts whose risk assessments deviate strongly from their own. After a number of iterations a "consensus" among all experts is reached. While this method is very intuitive, it does not include any weighting of the experts' estimates regarding their assurance. Therefore, we suggest an adapted iterative method, which interpolates between the two above mentioned methods.

In the algorithm shown in Figure 4, in each step the risk statements are smoothed based on a discrete inverse-distance kernel and updated according to Bayes rule. This way, the data are not only smoothed, but the assurance of each expert about their risk judgement is considered too. The Bayes update ensures that risk estimates with very high levels of assurance are given more weight than those having very low certainty, while smoothing gives more weight to expert

Data: $\boldsymbol{\mu}^{(0)} = (\mu_1^{(0)}, \ldots, \mu_N^{(0)})$, $\boldsymbol{\sigma}^{(0)} = (\sigma_1^{(0)}, \ldots, \sigma_N^{(0)})$, $\epsilon > 0$, $\delta > 0$
Result: μ^p – the pooled value for μ; σ^p – the pooled value for the standard deviation; w – the vector of weights assigned to each expert.

$W \leftarrow I_N$;
while $\|\boldsymbol{\mu}\|_{\max} > \delta$ **do**
  for $j = 1$ *to* N **do**
    for $k = 1$ *to* N **do**
      $c_{kj}^{(t)} \leftarrow \frac{\alpha_j^{(t)}}{\epsilon + d(\mu_j^{(t-1)}, \mu_k^{(t-1)})}$;
    end
    $\sigma_j^{2^{(t)}} \leftarrow \left(N \cdot \sum_{k=1}^{N} \frac{c_{kj}^{(t)}}{\sigma_k^{2^{(t-1)}}} \right)^{-1}$;
    $\mu_j^{(t)} \leftarrow \sigma_j^{2^{(t)}} \cdot N \cdot \sum_{k=1}^{N} \mu_k^{(t-1)} \frac{c_{kj}^{(t)}}{\sigma_k^{2^{(t-1)}}}$;
  end
  $\tilde{W}^{(t)} \leftarrow \left(\sigma_j^{2^{(t)}} \cdot N \cdot \frac{c_{kj}^{(t)}}{\sigma_k^{2^{(t-1)}}} \right)_{j=1,\ldots N,\ k=1,\ldots,N}$;
  $W \leftarrow \tilde{W}^{(t)} \cdot W$;
  $\boldsymbol{\sigma}^{2^{(t)}} \leftarrow \frac{\boldsymbol{\sigma}^{2^{(t)}}}{\sum_{j=1}^{N} \sigma_k^{2^{(t+1)}}}$;
  $t \leftarrow t + 1$;
end
$\mu^p \leftarrow \mu_1^{(t)}$; $\quad w \leftarrow W_1$; $\quad \sigma^p \leftarrow \sqrt{\sum_{k=1}^{N} \sigma_k^{2^{(0)}} \cdot w_k^2}$;

Figure 4. Iterative Opinion Pooling method with weights

opinions that are located in the center than to extreme data points. This procedure is iterated until all risk statements have converged to one value, which yields the aggregated risk. Note that the pooling algorithm (Figure 4) interpolates between the two above mentioned methods: if $\epsilon \to \infty$ this method coincides with the Bayes update, whereas equal variances, i.e., $\sigma_1^2 = \cdots = \sigma_N^2$, yield the consensual opinion pooling.
Note that $\left(\sigma_j^{2^{(t)}}\right)_{t \in \mathbb{N}}$ is a monotonically decreasing null sequence for all $j = 1, \ldots, N$, which may lead to numerical instability in the computation process. To avoid this, we added the command $\boldsymbol{\sigma}^{2^{(t+1)}} \leftarrow \frac{\boldsymbol{\sigma}^{2^{(t+1)}}}{\sum_{j=1}^{N} \sigma_j^{2^{(t+1)}}}$ to normalize the sum of the variances in each step. It can be shown that the procedure converges and that $\lim_{t\to\infty} \mu_1^{(t)} = \cdots = \lim_{t\to\infty} \mu_N^{(t)}$ holds. In every iteration, the entry w_{jk} in matrix of weights W corresponds to the weights expert j has so far assigned to all experts $k = 1, \ldots, N$. Note that W converges to a matrix with equal rows. Thus, the final weights of all experts coincide. By default, we choose the first row of W, $w = W_1$ as a final weighting vector.

B. Numerical Example

Assume four experts were asked to quantify the likelihood of a certain threat. Let $\boldsymbol{\mu^{(0)}} = (0.26, 0.255, 0.43, 0.315)^T$ and $\boldsymbol{\sigma}^{(0)} = (0.03, 0.018\dot{3}, 0.0\dot{3}, 0.028\dot{3})^T$ respectively. In figure 5 the densities are depicted. By choosing $\epsilon = 1$, we then obtain $\mu^p = 0.2922$ as pooled mean value for the likelihood, $\sigma^p = 0.0127$ as standard deviation and the weight vector $w = (0.1798, 0.4804, 0.1386, 0.2012)^T$.

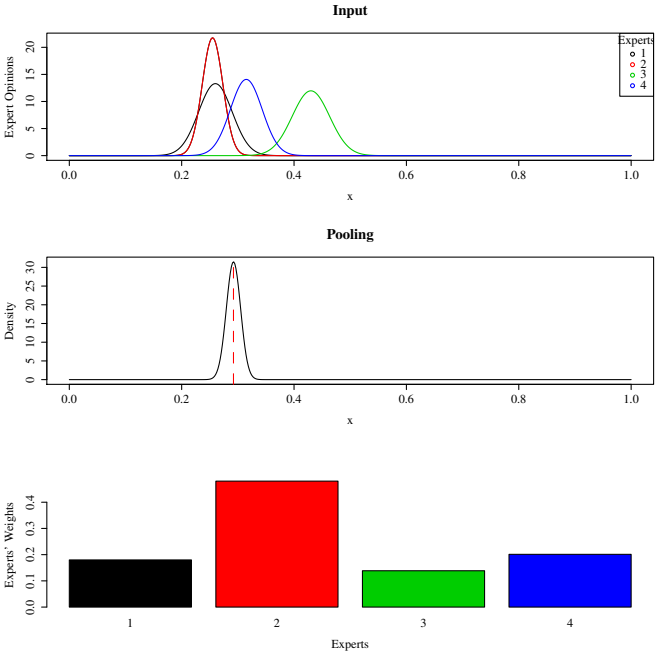

Figure 5. Numerical Example – Opinion Pooling of four Opinions

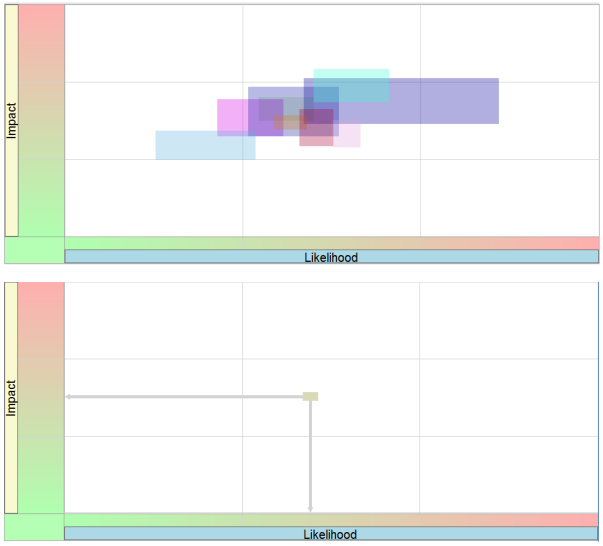

Figure 6. Graphical Risk Specification (top) and Aggregation (bottom)

TABLE I. POOLED VALUES DEPENDING ON THE CHOICE OF THE TUNING PARAMETER

ϵ	μ^p	σ^p	w_1	w_2	w_3	w_4
0.001	0.2792	0.0130	0.211	0.564	0.083	0.142
0.01	0.2833	0.0129	0.201	0.534	0.099	0.166
0.022	0.2852	0.0128	0.195	0.519	0.106	0.179
0.1	0.2887	0.0127	0.187	0.496	0.120	0.197
1	0.2922	0.0127	0.180	0.480	0.139	0.201
5	0.2929	0.0127	0.179	0.478	0.143	0.200
10	0.2931	0.0127	0.178	0.478	0.144	0.200

Note that the choice of the tuning parameter ϵ has a strong impact on the result. Depending on the desired degree of smoothness ϵ can be increased or decreased. Table I illustrates how different values for ϵ result in different outcomes regarding the opinion pooling. We suggest, however, not to oversmooth the data, and thus keep the size of ϵ reasonable. A handy approach is to use a modified version of Silverman's rule of thumb, i.e., $\epsilon \approx 1.06 \cdot \bar{\sigma} \cdot N^{-1/5}$, where $\bar{\sigma}$ denotes the arithmetic mean of $\boldsymbol{\sigma}$. In the given numeric example Silverman's rule yields $\epsilon \approx 0.02209$.

We suggest the opinion pooling method to be implemented in the visual risk specification. This way, the whole process from data collection, data correction and smoothing, to risk aggregation is combined in a single tool. In Figure 6 it is depicted how the experts' assessments are compiled into a single value.

V. Conclusion and Future Work

Specifying risks is in any case a matter of dealing with subjectivity and uncertainty. The application of statistical methods is especially challenging in this field, since "risk" is not an observable property of some physical process (as common elsewhere when statistics or probability theory is applied). Nonetheless, the issue is one of reasoning under uncertainty, and specifying this uncertainty in first place should be consistent with how the results are presented. This brings us to the proposed method of turning a risk presentation mechanism into an input system, and framing important tasks like data correction and opinion pooling into this approach. The techniques put forth here straightforwardly apply for one-dimensional quantities, such as when only likelihood or only impact should be elicited. Future steps mainly concern outlier analysis and way to automate outlier elimination. This entails in particular an analysis of bias and non-inferiority of the outlier-corrected risk data sets, and a more detailed stochastic model of risk estimation, where the risk is an unknown quantity, about which only correlated (independent, yet not necessarily identically distributed) random quantities can be measured (i.e., the subjective estimates).

Acknowledgment

This work was done in the context of the project "Cross Sectoral Risk Management for Object Protection of Critical Infrastructures (CERBERUS)", supported by the Austrian Research Promotion Agency under grant no. 854766.

References

[1] S. De Bock, "Effective risk management for complex it projects," https://sdb-plus.com/2012/05/15/effective-risk-management-for-complex-it-projects/, May 2012, [retrieved: July 14th, 2017].

[2] I. S. Organisation, "Iso/iec 31000: Risk management – principles and guidelines," 2009.

[3] Information Systems Audit and Control Association, "Cobit 5," 2012, [retrieved: August 11th, 2017]. [Online]. Available: http://www.isaca.org/cobit/pages/default.aspx

[4] C. J. Alberts and A. Dorofee, Managing Information Security Risks: The Octave Approach. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc, 2002.

[5] I. Münch, "Wege zur Risikobewertung," in DACH Security 2012, P. Schartner and J. Taeger, Eds. syssec, 2012, pp. 326–337.

[6] E. U. Weber, A.-R. Blais, and N. E. Betz, "A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors," Journal of Behavioral Decision Making, vol. 15, no. 4, 2002, pp. 263–290.

 ISBN: 978-1-61208-582-1

[7] C. S. Weber, “Determinants of risk tolerance,” International Journal of Economics, Finance and Management Sciences, vol. 2, no. 2, 2014, p. 143.

[8] C. F. Camerer and R. M. Hogarth, “The effects of financial incentives in experiments: A review and capital-labor-production framework,” Journal of Risk and Uncertainty, vol. 19, no. 1/3, 1999, pp. 7–42.

[9] J. Brenot, S. Bonnefous, and C. Marris, “Testing the cultural theory of risk in france,” Risk Analysis, vol. 18, no. 6, 1998, pp. 729–739.

[10] F. Dietrich and C. List, “Probabilistic opinion pooling,” October 2014, [retrieved: August 11th, 2017]. [Online]. Available: http://philsci-archive.pitt.edu/11349/

[11] A. O’Hagan and J. J. Forster, “Kendall’s advanced theory of statistics, volume 2b: Bayesian inference,” 2004.

[12] A. Carvalho and K. Larson, “A consensual linear opinion pool,” in Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, ser. IJCAI ’13. AAAI Press, 2013, pp. 2518–2524.

[13] S. Rass, J. Wachter, S. Konig, and S. Schauer, “Subjektive Risikobewertung – Über Datenerhebung und Opinion Pooling,” in DACH Security 2017, P. Schartner and J. Taeger, Eds. syssec, 2017.

[14] h. . https://www.coso.org/Pages/ermupdate.aspx. y. . . m. . D. n. . r. Committee of Sponsoring Organizations of the Treadway Commission, title = COSO Enterprise Risk Management – Integrated Framework Update.

[15] J. Chittenden, J. van Bon, and S. Polter, Risk Management: A Management Guide based on M_O_R, ser. Best Practice. Zaltbommel: Van Haren Pub, 2006.

[16] S. Rass, S. Konig, and S. Schauer, “Decisions with uncertain consequences-a total ordering on loss-distributions,” PLoS ONE, vol. 11, no. 12, 2016, p. e0168583.

[17] M. Shaked and J. G. Shanthikumar, Stochastic Orders. Springer, 2006.

 ISBN: 978-1-61208-582-1