

3

The RUCP system and the RAC engine in the Autonomous Province of Trento

This chapter presents the functioning of the RUCP with the RAC engine in Trento Province. Originally conceived as a single register holding data from different inspectorates, the RUCP is now evolving into a system designed to perform risk analysis by segregating businesses as per risk classes and levels. RUCP uses the internationally accepted scorecard tool to estimate risk-based ratings of operators. The chapter also delves into the technical realm and explains the system architecture and the process that goes behind RUCP's output delivery. The chapter illustrates the system application through the example for an environmental risk assessment for a business holding environmental authorisation. Three regulatory streams i) environment ii) labour iii) agriculture payments, show the areas in which the RAC engine can be developed and also the areas in which data-based analysis can be used. Finally, the chapter notes current limitations in data and recommendations for the future.

Introduction

The RUCP was created to be the single register of the inspections conducted in the Autonomous Province of Trento. It is not yet fully operational, but a pilot on agriculture payments inspections has been performed. It is meant to function with data collected by provincial agencies and inspectorates regarding their enforcement activities, and data related to objective characteristics of businesses. To do so, the RUCP draws information from the chamber of commerce database (C.C.I.A.A., Chamber of Commerce Registry) and other sources related to specific activities such as construction (i.e., CNCPT database¹) and goods and services certifications (for instance, Accredia²).

When first conceived by the regulator, the RUCP was envisaged to support inspectors on providing reliable data to better plan activities. Additionally, it was meant to prevent overlapping and duplications among inspectorates. Yet, following an initial assessment within the RAC Project, it was concluded that a punctual definition of the RUCP and its functions was needed. For this purpose, an Inception Deck methodology was implemented.³

As resulted from the Inception Deck, it was decided that the RUCP would become a system – not only a register – with the RAC engine inside, capable of i) delivering a risk-based rating of companies; ii) validating existent risk parameters; iii) defining new risk parameters. Accordingly, the RUCP will be able to deliver the following outputs: i) display for each enterprise the risk classes calculated according to fields in which it operates; ii) represent the impact-likelihood of each company, possibly by showing the risk matrix; iii) extract a list of businesses organised according to their risk level; iv) develop a function that enables an easy selection of business groups organised according to their risk level.

The RAC engine is able to deliver risk-based rating of businesses if existent risk parameters in the form of scorecards are first provided (see Box 3.1). Once the parameters are provided, the tool will process the existing knowledge and perform predictive modelling (OECD, 2019^[1]) (OECD, 2020^[2]).

Box 3.1. Using scorecards on data-based risk analysis

Relevant international experiences show that the most used tool aimed to define a risk profile is the “scorecard”, i.e., a list of – legal and technical – parameters formulated to assess businesses’ compliance weighted on the risk they represent. This way, scorecard allows to associate the potential risks for a company with a score in a very simple way. Since risk assessments have no straitjacket approach and need to be tailor made in the context of the specific risk and risk factors, activities and policies when assessed and scored against predetermined parameters can help identify high risk operations and set priorities.

The scorecard is widely used not only because of its advantages in defining companies’ risk profiles but also due to their simplicity. Beyond defining risks, scorecards are also remarkably useful for their benefits on machine learning processes aimed at calculating probability of risks. Scorecards indeed facilitate the procedure of calculating a risk score by providing a defined risk used by the ML algorithms to multiply the probability that the risk becomes concrete and the potential impact (damage) that may occur.

In order to better utilise the European Structural and Investment (ESI) funds and to protect it from potential fraud and corruption, the OECD has provided targeted guidance to authorities in Slovak Republic through data driven approaches to risk assessment and a systematic approach to managing these risks. The importance of enhancing the use of Arachne risk-scoring tool by authorities to refine risks has also been prioritised.

Based on the scorecards inserted into the system, the engine applies the algorithms aimed at identifying businesses characteristics and assigning them a non-compliance level of risk. Once the features are identified and risk labeled, they will be analysed through another type of algorithm (a linear algorithm) to produce a risk prediction based on the probability estimated with ML. During this process, particular attention has to be paid to i) the features interaction in order to avoid bias⁴ (e.g. linear independence where there is none); ii) being flexible on the prediction of likelihood of non-compliance. The algorithms, thus, must be customised and validated to obtain more certain and accurate predictions. With the support of the Department of Information Engineering and Computer Science of Trento University initial research is being done to define a sort of validation mechanism of the algorithms (see Chapter 4).

The system architecture: the risk calculation process and the risk assessment

Risk calculation process

As mentioned, the RAC engine delivers a rating of businesses according to a given level of risk by applying algorithms to data and scorecards available on the RUCP. The risk calculation is the process devoted to such output, and consists of different phases: exportation of the data, data preparation, prediction, UX which is the actual usage of the risk classification:

EXPORTATION → PREPARATION → PREDICTION → UX

Exportation: The exportation concerns the transfer of information from available databases to the RAC Engine. The IT systems that contain the data needed to perform the risk assessment are extremely heterogeneous. A good part of these systems is managed by Trentino Digitale (TD, IT public company of Trento), but there are also solutions acquired from different suppliers, in some cases even outside the TD datacenter.

The process requires to convert data into csv file⁵ in order to prepare it for cleaning, as described in the next step. Such a host will organise information to create scripts (bash, sql, python, etc.). The scripts are then processed by a pipeline software to enable them to download the csv file, clean, aggregate/organise, classify the risk, and provide an output. The proposal mainly focuses on scripts for two reasons: make the system development process simpler without implementing advanced algorithms by data scientists; grant more flexibility in developing the model. The model allows direct access to data when it is managed by TD. Otherwise, it should be foreseen an exportation, for example, by HTTP. The output of this phase is the disposal of data in the system ready to be used by the algorithm.

Data preparation: This phase is dedicated to clean and aggregate data in a single csv file. Here are included all the lines concerning the characteristics to consider in the companies' risk assessment. This stage also covers data denormalisation⁶ from multiple tables and the dataset reconstruction in a convenient format for the algorithm. The output of this phase will therefore be a single dataset per area, with a row for each company, showing all and only the characteristics used for the calculation of the risk class.

Prediction: Finally, an algorithm for the risk evaluation combines the companies' characteristics providing a numeric result. Such a number is then converted in a risk class by using thresholds. To do so, the algorithm takes the csv dataset file of the characteristics as input, and for each row calculates a numerical value that can be interpreted as a risk level. Through defined thresholds this is converted into a label (e.g., low, medium, high). The work is delegated to a script that will perform these steps: i) Reading the characteristics of the company from the csv file; ii) Multiplication by the weight vector that represents the algorithm; iii). Identification of the risk class based on the defined thresholds. The output of this phase is a risk value associated with the company in the RUCP system.

In some circumstances, Machine Learning techniques have been applied for the construction of the algorithm, which may be refined from time to time with new data, according to a frequency that could be annual. The algorithm is the result of the analysis of company data, and of the parameters suggested by the experts.

UX: The risk classification is ready to be used by the operators (inspectors). To this end, the inspector shall be able to access at least the following information: the risk class assigned to a company for each field in which it operates; extract a list of companies classified by their level of risk and field. The output of this phase is a sample to planning inspections.

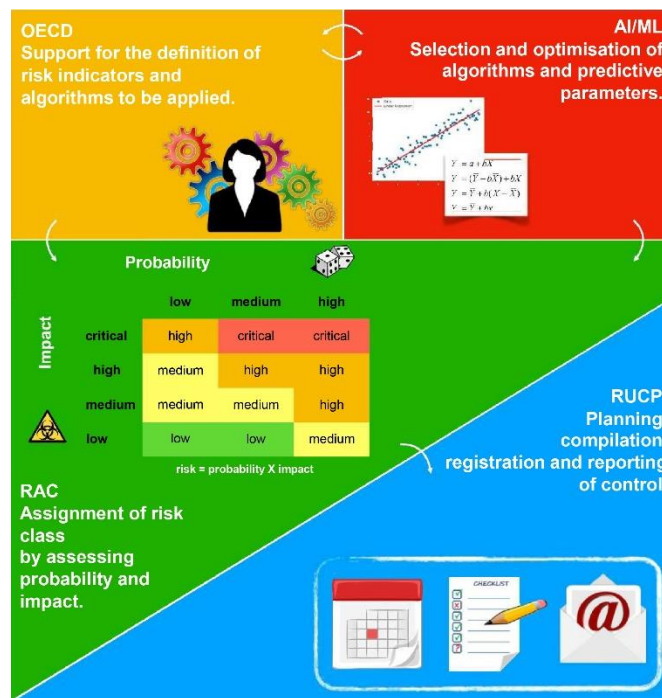
Risk assessment

For the system to retrieve the businesses characteristics, assess their risk, and execute the risk prediction two algorithms are needed. One which provides a risk rating list of the companies. This algorithm could consist of a software code developed within the RUCP system itself or with an external service. The second, driven by the rating provided by the first one, should extract and propose a list of businesses to investigate. Also, this algorithm should process systematic criteria.

Businesses’ risk is assessed by evaluating their historical behaviour based on previous inspections and considering the fields in which businesses operate. In this regard, the calculation of the risk class of a particular company is different for each area in which it is involved. Therefore, the same company may have an assessment under the environmental point of view and, at the same time, another evaluation for the work safety one, for instance.

The RAC Engine will provide a final risk assessment without performing any differentiation among the two components: i) the historical behaviour of the business and the probability that the way it acts is liable to non-compliance; ii) the objective risk characteristics related to the activity itself. A sum combines these components. The mathematic operation used keeps the system considerably simple. Figure 3.1 represents the RAC Engine.

Figure 3.1. RAC Engine



When assessing businesses' risks three aspects are considered: business characteristics, probability and impact characteristics, systemic criteria.

Business characteristics: A rating tool must be able to make predictions by leveraging several business characteristics. For convenience, these have been divided in three categories: general features, which are common to all enterprises (i.e., size of establishment or number of employees); specific factors linked to the type of activity of the enterprise (i.e., treating toxic or polluting products); management characteristics related to "how" the enterprise has been managed (i.e., certifications or control assessments).

Those data on characteristics need to be automatically collectable and updatable, using the abovementioned web services where possible (i.e., Chambers of Commerce).

Probability and impact characteristics: These include both characteristics that will have a predictive weight for the probability and impact of the risk. For a risk assessment on knowledge-based weights (i.e., a scorecard proposed by an inspector), this second classification is not relevant. If Artificial Intelligence (AI) and ML criteria can increase predictive efficiency, it is necessary to address the following issue. The target characteristics to train supervised algorithms are limited to the aspect of probability: there are precise measures of non-compliance, but there are none related to the impact that such "non-compliance" may produce if an accident occurs. What is relevant for the implementation of the ML algorithms is that the features that would be used in the predictive model of the probability of non-compliance will be included in the trained model and will inherit their weights from it. Moreover, the impact-related features will continue to have weights defined by the expertise of specialists (inspectors). Furthermore, the determination of the risk rating will depend on an additional combination criterion i.e., a linear combination or a risk matrix.

Systemic criteria: The systemic criteria help the RAC Engine to select companies to be investigated. Those are not necessarily related to the determination of risk. The most effective choice to perform a certain number of controls and find the most significant number of nonconformities is to inspect the companies with the highest rating.

Box 3.2. Example of the system application

Below is a representation of an environmental risk assessment for businesses holding environmental authorisation (called "AIA") to release specific quantities of pollutants into a reservoir. The hypothesis assumes the used of the following values:

Toxicity of the substance (T) Quantity of pollutants allowed (V) Reservoir critical issues (C)

The designed model foresees that potential risk is calculated using this formula:

$$\text{Potential Risk} = C + T \times V$$

The system uses two thresholds to set up three risk levels (Low/Medium/High).

To evaluate the probability component of non-compliance of a business, we take into account two values:


- The average of times in which it resulted in non-compliance during previous controls (R);
- Last checks differential (D)

In this case, we foreseen a linear algorithm derived from a regression (machine learning) to set up the potential risk class. Again, we assume three risk levels (Low/Medium/High).

$$\text{Probability of non-compliance} = w0 + R \times w1 + D \times w2$$

Finally, let us assume that the final assessment is calculated by applying this simple risk matrix (3x3)

Figure 3.2. Model of a risk matrix



Issues in the selection of companies such as: non-compliant selection,⁷ uneven distribution, not compatible with the resources needed, and inconvenience may occur. Those circumstances inevitably lead to a loss of effectiveness. Thus, they must be anticipated. Yet this cannot be addressed by an artificial intelligence algorithm, and must therefore be developed in software, following a more traditional methodology.

Key findings


Regarding the RUCP system and the development of the RAC Engine the pilot initially involved the assessment of five areas of inspections and their IT tools: OSH, food safety, environmental protection, agriculture payments and labour law. The research looked at the available data and its potential to be processed by the RAC Engine. The analysis revealed that just three of them (environment, agriculture payments and labour law) would be able to dispose their data to activate the process. While the data from these three inspectorates was indeed managed by TD (Trentino Digitale, public IT company from Trento) the other two agencies (i.e., OSH and food safety) have their own databases with privacy (legal) restrictions to export data into RUCP. The RAC Engine is therefore first oriented to deal with the inspectorates with data managed by TD. The export phase in the future, cannot be managed in a uniform manner for all these systems, the strategy must guarantee a certain uniformity, while imposing the least number of technical constraints. For the remaining inspectorates the exportation phase of the calculation risk process becomes easier to perform. This way inspectorates can fully integrate, streamline and digitalise the audit activity.

Yet, even in these cases, some relevant information remains outside the process. Unfortunately, information from control results are not always properly digitalised and systematised but continue to be pure text files, difficult to analyse. Because of this, just a few of the existing data is useful to elaborate risk ratings and to perform risk assessment through ML techniques at this stage.

However, limited availability of data has not prevented the pilot from achieving significant results. In the environmental protection stream, a risk assessment is being implemented using insights from the IMPEL model⁸, and available data regarding the objective characteristics of the businesses and the intrinsic risk of their activity (i.e., the potential to produce environmental harm given the toxicity of the substances managed by a given company, the correspondent volumes by it spilled, and the reservoir where the spillover take place). The behavioral history of businesses is not being considered in the assessment since, for now, information that can be used as data to activate the RAC Engine is not available. In the future, the service intends to collect inspections record in a more systematic way, which will allow to use the “behavioural history” of the business as one of the risk criteria in the RUCP. In the meanwhile, the possibility of using information related to fines and sanctions is being assessed. Besides this situation, current data allows a basic risk assessment and is useful enough to provide a risk-based rating of businesses.

Figure 3.3 shows risk assessment for a given company in terms of water pollution. The X-axis represents the reservoir in which the discharge occurs, while the Y-axis represents the volume of water exposed to pollution. Once the information related to the level of discharge and to the volume of the toxic substance is submitted into the system and given specific parameters of risk and scorecards, the RAC Engine applies the algorithms and performs a risk analysis delivering a specific level of risk for the given company. In the graph it is represented by the little blue point, which according to the calculation, is located at a medium level of risk represented in the figure by a yellow stripe.

Figure 3.3. Risk assessment on environmental protection



The same data used for the rating is being analysed with ML techniques to define the accuracy of the existing scorecards or to produce new risk parameters. Both are later used to choose samples for planning inspections, and to elaborate checklists for on-site inspections. On these bases the pilot is proving to deliver important results to improve regulatory and enforcement systems on environmental controls.

On the labour law stream, data-based analysis is mainly being directed to explore a validation mechanism of the algorithms. To assure the accuracy of a predictive method such the RAC Engine, it is necessary to agree on a sort of validation mechanism. In the first instance, the validation is generally to be found in the expertise of the inspectors who provided the scorecards. Drawing on this, AI and ML techniques are


applied seeking to answer the following questions: i) How a chosen algorithm can be validated? ii) How its effectiveness can be measured? iii) How it can be optimised? The answers come from the data analysis. Using historical data, the scorecard is analysed to determine its degree of effectiveness. Interestingly, much more can be done. Through data preparation works, new characteristics of the business – which were not taken into consideration before – can now emerge and provide potential information regarding the companies’ level of risk.

The inspectorate in charge of labour law enforcement provided a significant set of historical data, later used by the Department of Information Engineering and Computer Science of Trento University to initiate a research, through a degree thesis work, on ML applied to prediction: With automatic learning, new combinations of characteristics and weights can be searched and measured in effectiveness, until they converge to the optimised algorithm. Then this algorithm can be also used within the sample extraction software, e.g., RUCP (or even with pen and paper) to define the companies that will be inspected according to their risk profile. Over time it is also possible to continue supervising the effectiveness of the algorithm predictions, giving positive or negative reinforcements depending on the success or failures of the predictions. Periodically (for example every year) it is possible to reprocess the data by adding new ones, to refine or update the algorithm.

The research developed with Trento University analysed the features already present inside the labour law inspectorate’s application (called SISL in Italian) aiming at studying which combinations of companies’ characteristics are more effective for predicting compliance levels. In doing so, also many algorithms were tested and their performance were compared. The thesis was therefore very useful to provide a first analytical glimpse of the original dataset. Drawing upon these results, the OECD Team produced some preliminary data interpretations that, after being shared and discussed with the Labour law inspectorate, gave place to the creation of better indicators starting from historical non-conformities. With an approach partly similar to the one developed for OHS in Lombardy, the OECD studied the impact that the compliance trend of the last previous inspections has, and found the linearity expressed in Figure 3.4, where C correspond to compliance and N to non-compliance.


Figure 3.4. Compliance trend analysis on labour law

Second-last	C	-	N	C	-	-	N
Last	C	C	C	N	N	-	N
%	5.95%	19.88%	25.84%	38.32%	46.02%	55.84%	64.89%
Score	0	0.1666	0.3333	0.5	0.6666	0.8333	1



The algorithm provides a positive prediction, and indicates how reliable it is, i.e., it indicates its degree of certainty on the prediction's accuracy. On well-calibrated models, the "algorithm confidence" varies in the same way that the probability of the prediction's correctness does. Thus, whenever predictions with "greater confidence" are taken, the likelihood of correct predictions is greater as well. The analysis for the labour law inspectorate allowed thus to create a strong feature which, combined with the others in a logistic regression algorithm, increased the precision of the non-conformity control prediction to almost 50% (see Figure 3.5).

Figure 3.5. Algorithm calibration



Due to the characteristics of the chosen (calibrated) model it was possible to rely on the confidence algorithm to extract a given percentage of companies to be controlled, on which the precision rose to 61%, extracting 10% of the companies, and to 66% extracting 5% of businesses.

Though the results are not yet finalised, and the OECD Team continues to develop this pilot, evidence show the advantages of customised algorithms to correct predictions and review risk parameters. It has been possible to define the general characteristics of the machine-learning model, and its integration into the system. Due to the impossibility to inspect all the companies, and the need to identify a limited sample of businesses at risk, it becomes imperative to assure the highest degree of "confidence" of the algorithm. For this reason, it is advisable to use calibrated classification models.

Notes

¹ The National Commission of Territorial Joint Committees managed a database related to construction sites and occupational safety measures.

² Accredia is the body designated by the Italian government to certify the competence, independence and impartiality of the bodies and laboratories that verify the conformity of goods and services to the standards.

³ The OECD Team adopted a tool called Inception Deck to map, along with the IT company from Trento, Trentino Digitale, all the project's boundaries and the RUCP functionalities. Such a method foresees a set of ten slides defining the proposal's expectations by the stakeholders' points of view, the focus points, the criticalities, a roadmap and a first blueprint of the system's architecture.

⁴ For these purposes bias are assumptions in the machine learning process that generates results systemically prejudiced.

⁵ CSV File: Comma-separated values is a text file in which all values are separated by a comma. The csv is a standard and wide used format for data in Machine Learning.

⁶ Meaning having a single dataset with all the information replicated that can be used by the algorithm.

⁷ The optimal sample from the risk profile may not coincide with legal prescriptions aimed at defining the sample in a certain way.

⁸ The Impact Management Planning and Evaluation Ladder (IMPEL) model is the most extended method used by European Union member states to perform risk-assessment on the environmental protection stream.

References

OECD (2020), *OECD Public Integrity Handbook*, OECD Publishing, Paris, [2]
<https://dx.doi.org/10.1787/ac8ed8e8-en>.

OECD (2019), *Tackling Fraud and Corruption Risks in the Slovak Republic: A Strategy with Key Actions for the European Structural and Investment Funds*, OECD Public Governance Reviews, OECD Publishing, Paris, [1]
<https://dx.doi.org/10.1787/6b8da11a-en>.



From:
Data-Driven, Information-Enabled Regulatory Delivery

Access the complete publication at:
<https://doi.org/10.1787/8f99ec8c-en>

Please cite this chapter as:

OECD (2021), "The RUCP system and the RAC engine in the Autonomous Province of Trento", in *Data-Driven, Information-Enabled Regulatory Delivery*, OECD Publishing, Paris.

DOI: <https://doi.org/10.1787/ced63529-en>

This work is published under the responsibility of the Secretary-General of the OECD. The opinions expressed and arguments employed herein do not necessarily reflect the official views of OECD member countries.

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area. Extracts from publications may be subject to additional disclaimers, which are set out in the complete version of the publication, available at the link provided.

The use of this work, whether digital or print, is governed by the Terms and Conditions to be found at <http://www.oecd.org/termsandconditions>.