

ABSTRACT: Nonchalant design of a society's "safety system" is stupid. Luckily, it's also avoidable.

**By considering how to deploy our human resources most effectively through the introduction of a decision support tool (while imbedding an ethic of transparency), a public safety regulator has embraced constraints and delivered abundance.**

We used to think that there is a linear relationship between safety, and the number of experts who confirm things are safe. In other words, if you wanted better safety, that meant ensuring there were more people checking everything, all the time. A future-world expectation of decreasing risk tolerance, coupled with fewer and fewer deployable technical resources, has derailed that equation.

With limited resources it is not possible (nor is it desirable) to do 100% inspection of all hazardous equipment across BC – it is only possible to physically assess a fraction of technical equipment assets. This raises the question of how to design a safety system that is complete. BCSA has developed a decision support tool, called the Resource Allocation Program, which allows us to embrace and leverage our human resource limitations by prioritizing tasks based on risk. This algorithm has been designed to continuously learn, by receiving data from every site assessment.

By using our limited human resources' insight, experience, knowledge, and training, we can automate and populate a database that can help identify hazards, improve public safety, and deliver safe technical systems everywhere.

In evaluating the tool, we are designing in ethical considerations, so that there is a shared understanding of what risk is acceptable and what is unacceptable. We are also learning from the tool's limitations: If our human technical resources make errors in the ranking of samples, is it more important to deal with false positives or false negatives, especially when a false negative result has a higher cost impact? From an ethical standpoint, which correction is better and deserves greater focus given our limited resources? Additionally, over time as the database matures and become highly populated with our experts' knowledge, our human resources will not have as many opportunities to showcase their experience in selecting the safety priorities – how do we balance professional knowledge with machine-automated priority lists? Do we allow discretion? And to what extent can we allow machine learning to replace humans? The design of a safety system, and its implicit output as providing for a better human existence, must also take into account the tools used to deliver that safe outcome.