Improving Expert Elicitation on Non-Traditional Chemical Threat Agents Using a Mixed-Method Approach

Given there are potentially millions of non-traditional chemical agents, how can the security community down-select and choose on which to focus our efforts?

Laurel MacMillan, MS; Froggi Jackson, PhD; and Ryan Ritterson, PhD

Gryphon Scientific, LLC 6930 Carroll Ave., Suite 810, Takoma Park, MD 20912

Abstract

Historically, expert elicitation (EE) methodologies have been critical to informing strategies to combat weapons of mass destruction (WMD), such as the Chemical Weapons Convention (CWC). These expert elicitations are likely to become increasingly important to understand which scientific discoveries change the threat landscape, as scientific and technical capacities continue to advance. However, expert elicitations often suffer from a lack of available expert time and the challenge of identifying who among a set of potential experts are truly the most qualified. Here, we undertook a novel, mixed-methods approach to the elicitation of expert consensus to assess which non-traditional chemical agents (i.e., those not previously used for chemical warfare) are of greatest urgency to address. Our approach used supervised machine learning (ML) augmented with a statistical method called Bayesian Truth Serum (BTS).

Our proof-of-concept approach demonstrates several advantages, including: (1) a generalized approach that transfers of the burden of the assessment from the experts to the assessors, (2) a methodology to determine who among the experts elicited is most expert (and how to weigh expert input based on those insights), and (3) a machine learning algorithm that, once trained on a small set of agents informed by expert opinion, can estimate how concerned experts would be about additional chemical agents without needing to solicit those experts again. Using this approach, the CBRN stakeholder community could begin to assess the enormous body of potential non-traditional agents to identify those of most concern, as well as rapidly assess any new potential hazardous chemical agent without reliance on additional expert elicitations.

Challenge with Traditional Elicitation	Machine Learning Solution
Labor-intensive for experts	Ask experts to assess only a small number of chemicals as of high, medium or low concern
Can't know who is most expert	"Bayesian truth serum" reveals relative strength of experts
Too many agents to assess	Tool can predict ranking without requiring new expert input

Bridging ML and EE Methods

Machine Learning	Expert Elicitation
• Uses a **fixed** set of definite *features* to predict classification • Must have feature data about agent to make prediction • All rankings must derive from same features	• Uses any available information and expertise to derive ranking • Can use different information for different agents • Can use indefinite information

We built a proof-of-concept approach and tool to address challenges with **chemical threat/risk assessments** that require **expert elicitation and subjective judgement**

Implicit in our ML approach is an assumption that a fixed set of chemical characteristics correlate to how experts rank chemicals **even if experts don't use those characteristics explicitly in deriving their own rankings**

Mapping Features to Concern

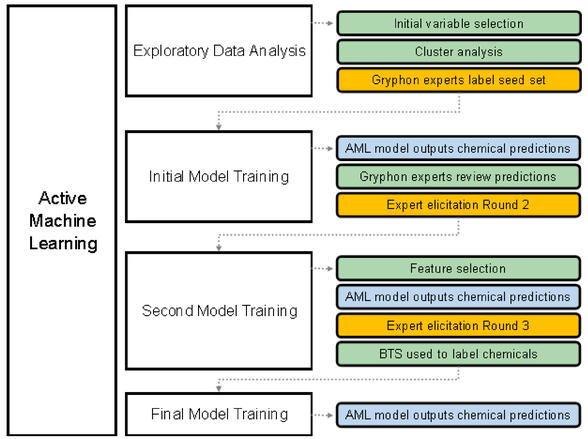

Figure 1. The model building process used. Green boxes indicate analysis and model building steps, orange boxes indicate steps where chemicals were labeled prior to model building, and blue boxes indicate the chemical prediction output yielded from each model training step.

An ML algorithm was developed that can predict how experts would rank, or label, chemicals in terms of their risk, or "level of concern".

The ML algorithm was developed using descriptive information – ~40 different data types ("features") – collected from widely available data sources for 166 chemicals coupled with multiple rounds of model refinement based on expert labeling of a subset of these chemicals. Thus, the ML algorithm was underpinned by a large dataset, and required only simple, limited label inputs from the experts to "train" it, reducing their effort and avoiding language barriers.

The resulting algorithm can predict how experts would label any chemical in terms of a rating of low, medium, or high level of concern and could stand-in for the experts when the risk of a new chemical needs assessing, given the availability of similar descriptive data for that chemical.

Risk Assessment – Features vs Concern

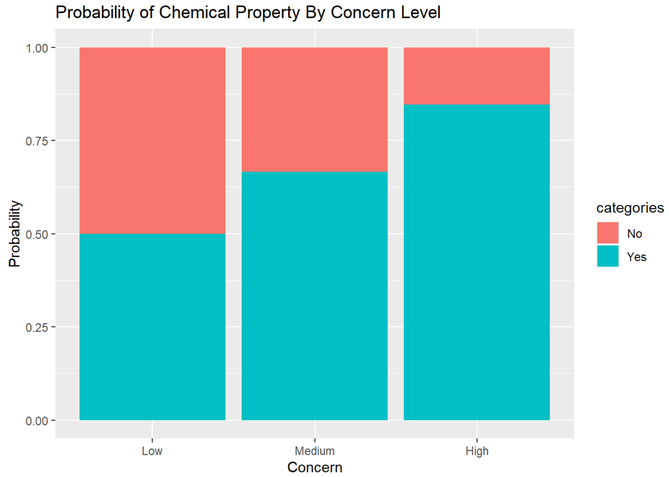

- In the final model, each chemical feature has a set of probabilities relating the feature values to the level of concern classes
- In the example above, we show a feature, specifically whether the chemical has a colored appearance, that can take one of two values: yes or no.
- Our model predicts this feature has a positive correlation with level of concern – there is a greater probability that the chemical is of "high" concern if the chemical has a colored appearance
- Note: strong correlations do not mean that experts should use this as the sole method for determining level of concern

Top 10 Chemical Features

An analysis of the most informative features from one final model revealed the following top 10 chemical features for predicting level of concern (some expected, some not):

Toxidrome
LD50 Oral Value
Chemical causes cancer
Chemical is toxic to aquatic life
Chemical is in liquid state under ambient conditions
Colored Appearance
Extinguish Fires using Dry Chemical Powder
Extinguish Fires using Foam
Chemical is available in the Americas
Chemical is available in Africa

GRYPHON SCIENTIFIC