Uncertainty-aware Artificial Intelligence for More Effective Decision Making

Lance Kaplan, Jin-Hee Cho, Murat Sensoy, Feng Chen, and Paul Sullivan

Summary:

This work addresses the determination of uncertainty in AI&ML reasoning when the training data is sparse. Such reasoning should output low uncertainty when the observational sample under test are representative of the training data, but the uncertainty must be high when the test sample is not representative. The characterization of uncertainty is established for low level (i.e., neural networks) and high level (i.e., probabilistic graphical models) reasoning systems.

Abstract:

Many information processing systems have been proposed to provide situational awareness (SA) for a decision making based upon observations from physical sensors and/or human generated reports. Advancement of such systems are crucial for the Army's modernization plans. For instance, next generation combat vehicles (NGCVs) will rely on SA to decide when it is necessary to engage their active protection systems or when to alter routes to avoid unnecessary (but possible) kinetic engagements. All of these automated information processing systems are designed to work under certain conditions, and the algorithm developers do not necessarily understand the bounds of the operational space outside of which these systems cannot perform reliably. It is crucial for an information processing system to express its degree of confidence on its reported information about the area of interest so that the decision maker can digest this report with his/her own observations to form a proper SA picture. If the information system provides wrong information instead of stating that it does have enough evidence to reliably make an inference, the decision maker can easily make a fatal mistake.

Many of information processing systems exploit artificial intelligence and machine learning (ML) technologies to understand how to interpret patterns in the observations. A common approach is for observations to feed into ML algorithms such as (deep) neural networks (NNs) to detect and classify objects in the scenes. Then, a higher level reasoning system such as a Bayesian network (BN) can be used to fuse these detections and classifications into a threat assessment. The NNs must be trained to learn the values for its weights to effectively classify and detect objects. Likewise, the BN requires past evidence to determine the conditional probabilities that characterize the network. This historical evidence can come from either 1) a domain expert who uses his past experiences or 2) historical data to formulate the probabilities. In the military domain, the past evidence about a particular area of interest is usually very sparse. Therefore, the parameters describing the NNs and BN models cannot be determined precisely. This lack of precision affects the outputs of the information processing systems that are inferred by the observables and the trained models, and it becomes crucial to characterize the uncertainty of these outputs.

This works builds upon subjective logic to establish the theoretical foundation to realize uncertainty aware processing components of the information processing systems. Subjective logic is a framework for probabilistic reasoning under uncertainty. It connects belief mass assignments to second-order distribution knowledge of probabilities. Specifically, the beliefs and uncertainty values, which constitutes a subjective opinion in subjective logic, map to the parameters of a Dirichlet distribution. This work summarizes our recent efforts to probabilistically reason over uncertain probabilistic graphical models

where knowledge of the parameters for the modes are expressed as subjective opinions, i.e., uncertain probabilities. This includes subjective BNs where uncertain knowledge of the conditional probabilities is represented as subjective opinions. We have developed an efficient method to infer subjective opinions for latent variable conditioned on the values of the observed variables. Similarly, collective subjective logic expands probabilistic soft logic (PSL) with uncertain reasoning. We have developed the inference method for subjective opinions of latent variables in light of the opinions of observed variables and the set of probabilistic first order logical rules inherent in PSL. Finally, we are currently training uncertainty-aware NNs by developing loss functions that interpret the output layers as the parameters of a Dirichlet distribution.

We can evaluate the utility of uncertainty-aware information processing in terms of how well the uncertainty characterizes the disparity between an expected inference and ground truth. Specifically, we report the divergence of the desired confidence bounds for the higher-level reasoning methods, i.e., subjective BNs and collective subjective logic. The uncertainty-aware NNs are evaluated based upon how uncertainty values change as the test data is extracted from similar or different sets in relation to the training data. We also demonstrate that on the testing set, classification performance increases as more uncertain test samples are pruned.

Currently, all the real data for quantifiable evaluations are common ML datasets that are not necessarily military relevant. We will use a military relevant vignette for route planning of a NGCV to show how an end to end uncertainty-aware information processing system is able to exhibit confident inferencing when the observations are normal but indicate high uncertainty once the observations do not represent the training cases.