

SIMULATION-BASED EVALUATION OF INFORMATION FUSION METHODS

Mikołaj Roszczyk, Daniel Paczesny

Institute of Telecommunications and Cybersecurity,
Warsaw University of Technology,
Nowowiejska 15/19, 00-665 Warszawa, Poland

The paper addresses the problem of combining uncertain information from multiple sources into a coherent decision. The problem is considered within the second layer of a multi-layer decision system, where multimodal information and uncertainty estimates obtained from the preceding layer are fused. Several information fusion approaches are implemented and evaluated, including Bayesian inference, voting-based methods, and methods based on Dempster-Shafer theory. A simulation environment is used to investigate the behavior of these methods under varying levels of uncertainty and agreement between information sources. The obtained results provide a basis for comparing the considered fusion approaches and analyzing their ability to produce reliable decisions from uncertain data. The presented study constitutes an initial step toward a broader investigation of uncertainty propagation and information fusion in multi-layer decision systems.

1 Introduction

Decision support systems rely on information describing the current state of an observed process or object. Based on this information, a decision-making mechanism can determine an appropriate action or provide a recommendation for a user. Such systems can be applied, for example, to monitoring and control of industrial processes, where the state of an object must be continuously observed and decisions must be made based on it.

In many such applications, the state of the observed object cannot not be determined directly. Instead, it should be estimated from measurements obtained from multiple sources. Individual measurements may be uncertain, incomplete, or inconsistent with one another. Combining information from multiple sources is therefore an important step in providing a reliable representation of the current state for the subsequent decision-making process [1].

Information fusion has been studied using a variety of mathematical frameworks [1]. Among the classical approaches are probabilistic methods, belief-based methods, and methods based on fuzzy set theory. Probabilistic approaches represent uncertainty using probability distributions, with Bayesian inference being one of the fundamental examples. Belief-based approaches provide a framework for representing and combining uncertain or incomplete information; the Dempster-Shafer Theory (DST) and Dezert-Smarandache Theory (DSmT) are examples of this family. Fuzzy set-based methods, in turn, are used to represent imprecise information through degrees of membership.

The designed system is organised as a multi-layer architecture. The first layer is responsible for data acquisition and for estimating the state of the observed object and the associated certainty based on individual sensor measurements. The second layer performs information fusion using selected approaches. The third layer, which is not discussed in this paper, implements an actual decision-making based on the observed state.

This work focuses on the implementation of the information fusion algorithms in the second layer and their evaluation in a simulated environment. The simulation models the first and the second layer of the described system, while providing a known underlying state that can be used as a reference for evaluating the fusion results. This allows different environmental scenarios and sensor configurations to be tested and the performance of chosen fusion methods to be compared.

The source code, progress and additional resources are being published on GitHub repository *DecisionSystemFeeder* [2].

2 Information fusion

2.1 Process of state estimation

Let X be the set of possible states that the observed object might be in and let $x \in X$ denote the actual current state of the observed object.

The first layer consists of multiple sensors, each of which may observe different characteristics of the same object. Each sensor processes its measurement using its state estimation logic and produces an estimate of the current state. Let $Y_i \subseteq 2^X$ denote the set of possible state estimates produced by sensor S_i , and let $y_i \in Y_i$ denote the state estimate produced by this sensor [3]. Sensors return a tuple (y_i, c_i) , where c_i is the certainty value.

The second layer receives the outputs of multiple individual state estimators and combines them using an information fusion function to determine the most probably actual state of the object. The resulting representation may be provided to the decision-making layer.

The full state estimation process is presented on Figure 1 [3].

2.2 Information fusion approaches

Several information fusion methods are considered in the second layer of the decision system. The first group consists of voting-based methods [4], which combine decisions from multiple sources by determining the outcome based on the votes provided by individual state estimators. The considered methods include majority voting, cumulative voting, approval voting, and the Borda count. A particular focus is placed on Nanson's method, which combines the Borda count with iterative elimination and satisfies the Condorcet criterion. Voting-based approaches provide a simple and computationally efficient way of combining decisions from multiple sensors without requiring an explicit probabilistic model




Figure 1: Process of state estimation and information fusion

of the environment. Such methods have also been applied in sensor networks, where Silent Negative Voting has been proposed as a power-efficient data fusion mechanism for improving the reliability of decisions while taking into account the limited resources of sensor nodes [5].

The second group consists of probabilistic information fusion methods, which represent uncertainty using probability distributions. Bayesian inference is one of the fundamental approaches in this group [3]. The Bayesian inference requires prior information about the probability of occurrence of each possible state and a confusion matrix, which contains the conditional probabilities of obtaining a sensor observation y_i given that the actual state is x_i ($\mathbb{P}(y_i|x_i)$). Bayesian inference is based on Bayes' theorem and provides a probabilistically justified estimate while taking into account both prior knowledge about the environment and the reliability of individual sensors. Bayesian information fusion has been widely applied in Internet of Things systems [6].

Another group consists of belief-based information fusion methods, which provide a framework for representing and combining uncertain, incomplete, or imprecise evidence. One of the most established approaches in this group is the Dempster-Shafer Theory (DST) [7]. Let Θ denote the frame of discernment, consisting of all mutually exclusive and exhaustive hypotheses describing the possible states of the observed object. Unlike classical probabilistic approaches, DST allows evidence to be assigned not only to individual states but also to sets of possible states, making it possible to represent information that cannot be precisely attributed to a single hypothesis. A basic probability assignment m assigns a mass value to each subset of Θ , where the assigned mass represents the amount of evidential support committed to that proposition. Information from multiple sources can then be combined using Dempster's rule of combination (Eq. 3).

$$m_1 \oplus m_2(C) = \frac{\sum_{A \cap B = C} m_1(A) \times m_2(B)}{1 - \sum_{A \cap B = \emptyset} m_1(A) \times m_2(B)} \quad (1)$$

The resulting evidence can be evaluated using the Belief (Eq. 3) and Plausi-

bility (Eq. 3) functions, which provide lower and upper bounds for the support of a given hypothesis [7].

$$Bel(A) = \sum_{B \subseteq A} m(B), \quad A \subseteq \Theta \quad (2)$$

$$Pl(A) = \sum_{B \cap A \neq \emptyset} m(B), \quad A \subseteq \Theta \quad (3)$$

Belief-based methods have found applications in safety-critical systems such as autonomous vehicles [8], where information from multiple sensors must be combined despite uncertainty, incompleteness, and conflicting observations. Unlike probabilistic approaches, DST does not require complete prior knowledge of the environment, such as prior probabilities of the possible states, which makes it suitable for situations where such information is unavailable or difficult to obtain.

Fuzzy set-based methods represent imprecise information using degrees of membership rather than binary assignments or classical probabilities [1]. Fuzzy logic and fuzzy set-based approaches have been applied to information fusion in a variety of domains, such as IoT [6], including the aggregation of trust and decision-level fusion under uncertainty. Although fuzzy-based methods are not considered in the current simulation, they constitute an important class of information fusion approaches and may be considered in future extensions of the proposed framework.

3 Simulation

3.1 Simulation environment

The developed simulation environment is designed to reproduce the first and second layers of the system, corresponding to data acquisition and information fusion. The purpose of the simulation is to generate controlled observations of an environment, process them using sensor models, and evaluate the resulting state estimates produced by different information fusion methods.

The simulated environment is defined by a set of possible states X and their corresponding *a priori* probabilities. The current state of the environment can be either manually specified or sampled according to the defined *a priori* probabilities. The environment can also contain measurable parameters defined by user-specified functions, e.g. temperature, humidity, tilt. Such parameters can be used to model properties of the environment that affect sensor measurements or to construct specific experimental scenarios.

Sensors are represented by configurable models that observe the current state of the environment and produce state estimates. Different sensor models can be used to represent different measurement characteristics and levels of uncertainty.

A simulation consists of a specified number of iterations, where each iteration represents a single observation of the environment. During an iteration, the

current state is sampled, sensor measurements are generated, and each sensor produces its corresponding state estimate. The estimates are then provided to the selected information fusion method, which produces a fused state estimate. The fused estimate can subsequently be compared with the actual state.

The performance of an information fusion method is assessed by comparing its fused state estimates with the corresponding actual states of the environment using predefined evaluation criteria. This provides a basis for comparing different fusion methods under reproducible conditions. The separation between the environment, sensor models, and fusion algorithms also allows the sensor interface to be extended in the future to support data from physical sensors.

3.2 Implemented sensor types

At current stage of the development, three sensor types are implemented to reflect different operational characteristics.

Scalar Sensor models sensors measuring continuous environmental parameters (e.g., temperature, pressure, humidity, vibrations). The measurement uncertainty u is calculated dynamically from the scalar measurement v using relative (ϵ_{rel}) and absolute (ϵ_{abs}) uncertainty parameters (Eq. 4).

$$u = v \cdot \epsilon_{\text{rel}} + \epsilon_{\text{abs}} \quad (4)$$

The probability assigned to a state corresponding to an interval $[L, U]$ is calculated using the cumulative distribution function of a normal distribution $\mathcal{N}(v, u^2)$ (Eq. 5).

$$\mathbb{P}(y) = \Phi\left(\frac{U - v}{u}\right) - \Phi\left(\frac{L - v}{u}\right) \quad (5)$$

Computer Vision Sensor represents vision-based object detection systems. The estimation process consists of two stages: extracting bounding boxes (characterized by spatial coordinates, class IDs, and model confidence scores) from the environment, and analyzing them via a dedicated pipeline to estimate state probability distributions.

Simple Sensor represents categorical decision components characterized by a constant error parameter $p_{\text{mistake}} \in [0, 1]$. Given a primary decision function, it assigns mass of $1 - p_{\text{mistake}}$ to the chosen state while assigning the remaining mass p_{mistake} to the complement set of unselected states.

3.3 Evaluation

The simulation framework provides a procedure for evaluating the performance of an information fusion method against a known ground truth. Since the actual state of the simulated environment is available during the simulation, the fused decision can be directly compared with the state that was used to generate the sensor observations. This makes it possible to evaluate fusion methods under controlled and reproducible conditions.

For each simulation iteration, the current state of the environment is first determined. The configured sensors then obtain measurements from the environment and produce their respective state estimates together with the associated uncertainty representation. The outputs of all sensors are passed to the selected information fusion method, which produces the final fused decision. The fused decision is then compared with the actual current state of the environment.

The evaluation procedure distinguishes between cases in which the fusion method returns a single state and cases in which multiple possible states are returned. If the fused decision contains only the actual state, the result is considered a fully correct decision. If the actual state is included among multiple states returned by the fusion method, the result is considered partially correct. This distinction allows the evaluation to account for methods that preserve uncertainty by returning a set of possible states rather than selecting a single state.

Three metrics are calculated during the evaluation. The first metric, *including correct*, measures the proportion of cases in which the actual state is included in the fused result. The second metric, *only correct*, measures the proportion of cases in which the actual state is returned as the only state. The third metric, *correct parts*, provides a partial score when multiple states are returned. In such a case, the contribution of a correct result is calculated as the reciprocal of the number of states included in the fused decision. For example, a decision containing two possible states, one of which is correct, contributes 0.5 to this metric, whereas a single correct state contributes 1.0.

The metrics are calculated both for the complete set of simulation iterations and separately for each possible environment state. The per-state results can be used to construct a confusion matrix and identify which states are most frequently confused. This is particularly important for probability-based fusion methods, such as Bayesian inference, which may perform worse when detecting less likely states. In many use cases, the *idle* state is expected to occur most frequently, while other states may be considerably less common. Therefore, the quality of a fusion method, and ultimately of the entire decision system, should not be assessed solely based on overall accuracy, but should also consider its ability to detect less likely states.

In addition to the aggregated evaluation metrics, the simulation framework can record detailed logs for individual iterations. These logs contain the actual state, the outputs produced by individual sensors, and the final fused decision. The recorded logs can be used for further analysis of individual cases without requiring the simulation to be run again.

4 Results and future work

This work presented a simulation-based framework for the evaluation of information fusion methods in the second layer of a multi-layer decision system. The developed environment enables different sensor models, uncertainty levels, environmental states, and fusion approaches to be combined into reproducible

simulation scenarios. This provides a basis for investigating how different information fusion methods behave under varying levels of uncertainty and agreement between information sources.

The current implementation should be considered a work in progress. The main objective of the present stage was to develop and make available the simulation environment and the implemented information fusion methods, rather than to provide a comprehensive experimental comparison.

Detailed experiments, including systematic variation of sensor reliability, source agreement, and other simulation parameters, are planned as future work. These experiments are expected to provide a more comprehensive comparison of the considered fusion approaches and to investigate the propagation of uncertainty through the multi-layer decision system.

Work is also ongoing [2] on a real-world application of the described system: bird detection at a bird feeder. The smart feeder contains several types of sensors, including a motion sensor, a vibration sensor, an RGB camera and an infrared camera with computer vision-based detection, and a microphone with audio analysis. Data obtained from these sources are intended to be fused using one of the methods described in this paper.

The source code and the current state of development are available in the project repository [2].

References

- [1] Xinde Li, Fir Dunkin, and Jean Dezert. “Multi-source information fusion: Progress and future”. In: *Chinese Journal of Aeronautics* 37.7 (2024), pp. 24–58. ISSN: 1000-9361. DOI: <https://doi.org/10.1016/j.cja.2023.12.009>. URL: <https://www.sciencedirect.com/science/article/pii/S1000936123004247>.
- [2] Mikołaj Roszczyk. *DecisionSystemFeeder Repository*. 2025-2026. URL: <https://github.com/Roszczyk/DecisionSystemFeeder> (visited on 08/20/2026).
- [3] Tadeusz Pietkiewicz. “Fuzja danych”. Military University of Technology. 2025.
- [4] Piotr Pałka. “Grupowe podejmowanie decyzji: negocjacje, głosowanie”. Warsaw University of Technology. 2026.
- [5] M.Umashankar and C.Chandrasekar. “Power Efficient Data Fusion Assurance Scheme for Sensor Network using Silent Negative Voting”. In: *International Journal of Computer Applications* 1.4 (2010). ISSN: 0975 – 8887.
- [6] Farag Azzedin and Mustafa Ghaleb. “Internet-of-Things and Information Fusion: Trust Perspective Survey”. In: *Sensors (Switzerland)* 19 (2019). DOI: <https://doi.org/10.3390/s19081929>.
- [7] Glenn Shafer. *A mathematical theory of evidence*. Princeton, NJ, USA: Princeton University Press, 1976.

- [8] Boquan Yang, Jixiong Li, and Ting Zeng. “Review of Environmental Perception Technology Based on Multi-Sensor Information Fusion in Autonomous Driving”. In: *World Electric Vehicle Journal* 16.20 (2025). DOI: <https://doi.org/10.3390/wevj16010020>.