# Sensor Synergetics: The Rationale of Sensor Fusion

*Ariën J. van der Wal*

## PROLOGUE: Why study sensor fusion?

In this prologue we will introduce the subject and the motivation of this study through a simple example illustrating some of the problems associated with sensor fusion. Sensor fusion is the process of combining the individual information streams acquired with a number of sensors in order to achieve "better" situation awareness. How to combine these streams depends on both the specific goal of the observation process and, having defined a performance measure for fusion, how this measure is affected by external factors and intrinsic sensor limitations.

Let us imagine the following situation: An autonomous vehicle (a robot) is programmed to move from position A to position B. Positions A and B are located in the same horizontal plane in which the robot can move. In the plane there are a number of obstacles that have to be avoided by the robot. The robot is equipped with a number of sensors, say 3, in such a way that the robot can sense any obstacles that lay ahead in its path.

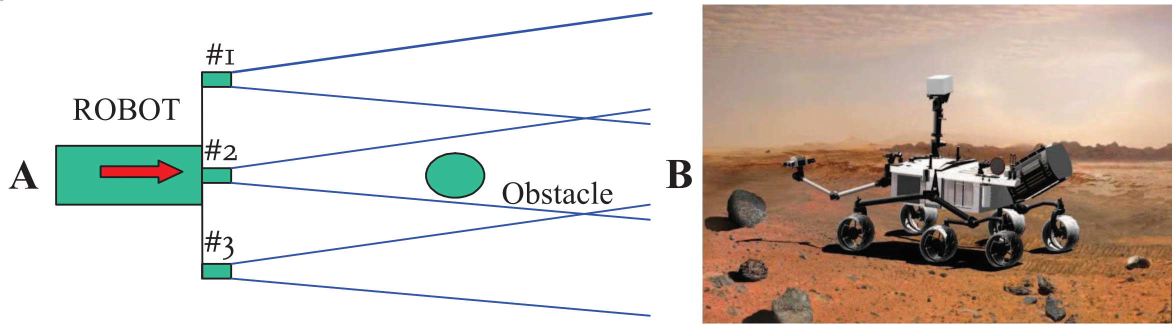


*Figure 1. Top view of a robot with 3 identical forward-looking sensors moving from A to B using sensor fusion to avoid collision with an obstacle. The blue lines indicate the field of view (FOV) of each sensor. At the right an artist impression of such a robot is shown.*

Now the following situation could arise: (cf. Fig. 1) Sensors #1 and #3 indicate that the way ahead is free, whereas sensor #2 "sees" an obstacle, say a tree, ahead. Basically there are two fundamental choices possible for a rudimentary sensor fusion mechanism: a linear and a non-linear scheme. The first would consist of a superposition (= linear combination) of the partial decisions made by the individual sensors, viz. "If no obstacle is sensed, continue route along original path, else change heading". A joint decision based upon the partial decisions with this rule and using the principle of a majority vote would result in a collision with the tree, since two out of three sensors arrive at the conclusion to continue with the original heading.

Also in the case that a layered (i.e. *delayed*) decision scheme would be used, problems could arise, as illustrated by the following: Suppose that the partial decisions of the sensors #1 and #2 are derived via the rule: "If sensor #1 indicates no obstacle and sensor #2 detects an obstacle, then turn left, else continue route along original path", and similarly a partial decision is derived for sensors #2 and #3, we are faced with the situation where the partial decisions of the pairs (1,2) and (2,3) are to be combined. If we use superposition, again we end up with an undesirable result, because linear combination of the decisions "turn left" and "turn right" effectively results in steering straight ahead, thereby colliding with the obstacle.

Only if we combine the partial decisions in a non-linear way, the combination of three sensors could help to avoid a collision with the obstacle and therefore adds value to the robotic observation system, e.g. by postulating the sequence of following rules:

- R1: If all three sensors do not detect an obstacle, move ahead.
- R2: If sensor #1 (#3) sees an obstacle turn right (left).
- R3: If sensor #2 sees an obstacle, throw a coin and decide to turn right (left) in case of heads (tails).

With this very simple example in mind it will be evident that linear combinations of sensor inputs have only limited use for making decisions. For this reason we direct our attention to *non-linear* combinations of partial decisions. In the present article we first discuss why effective sensor fusion necessarily needs to be a non-linear process and next sketch the use of partial and soft decisions and the way to aggregate these via non-linear schemes, in such a way as to arrive at useful and meaningful final decisions, based upon the available raw information from individual sensors. The mathematical formalism of fuzzy logic provides a versatile and adequate means to formally describe sensor fusion.

## Introduction

Numerous research papers have been published dealing with the application of multi-sensor data fusion, also referred to as distributed, or "network enabled", sensing combined with high-level fusion, especially in the domain of military observations [1-6]. Although intuitively appealing, one may conclude that data fusion has not yet brought about the expected breakthrough. Several explanations for this can be given, such as the particularity of the application domain, the limited availability of general methods for fusion, and finally the quality of the primary 'raw' sensor data. Another problem may be the unrealistic expectations of the virtues of the synergy of multiple sensors.

In the absence of a general way to approach the subject, many ad hoc experiments and simulations have been published. In the following we will shortly review the history of fusion, define sensor fusion as a field of research in its own right, and next discuss the problem of how to model sensor fusion and suggest some directions for answering some of the pertinent questions in this field using concepts from soft computing. Especially the use of fuzzy measures looks promising as a way to model the sensor fusion process quantitatively.

## Historical overview

Historically the idea of sensor fusion is not new: As early as in the sixties multi-radar trackers have been in use by the military for air traffic control and air defence. Multi-sensor data fusion seeks to combine information generated by multiple sensors or multiple samples from one sensor to achieve goals that would be very hard or impossible to achieve with a single sensor, or a single sample. From the point of view of efficiency, scheduling, accuracy, and redundancy it seems intuitively obvious that several sensors should be 'better' than a single sensor.

Nowadays data fusion is a well-accepted method for making superior inferences in the field of industrial automation (e.g. for controlling a power plant, an oil refinery, a cement kiln; for a review on industrial applications, see e.g. [7, 8]), or even a nuclear reactor [9, 10], and for carrying out real-time pattern recognition in industry using a variety of sensors. Especially since the advent of soft-computing methods, such as fuzzy logic, data fusion has become a widely accepted successful fusion technology in industry. We note however that the success of such methods is primarily due to their ability to model human behaviour or expertise in supervisory control. Sensor fusion also endeavours to mimic cognitive processes in humans by absorbing the signals of the human observation system, i.e. our five senses, from the real world and integrate, or 'fuse', these signal streams to arrive at a coherent picture of our environment. As such, sensor fusion is concerned with lower abstraction levels and much higher information rates. It requires therefore faster response than the data fusion used in supervisory control systems. This forms also the key problem in applying soft computing methods to this field: in controlling complex industrial or organizational processes over relatively long timescales human operators have accumulated ample experience over the years. In contrast, there is only limited insight in the way a human being builds an environmental picture, his awareness, from continuous multi-sensate observations. It may therefore be a useful approach when studying sensor fusion to have a close look at how the human cognitive system works. Although cognition is still far from understood in detail, a few global characteristics are apparent: the human recognition system consists of a massively parallel processor that merges vague, qualitative inputs and a priori knowledge, acquired by learning from experience into a more or less consistent picture of the environment. It consists of a large number of hierarchically ordered decision processes running concurrently, simultaneously inferencing on the same set of input data at different levels of granularity, both in feature space and in space-time. We will not discuss these points in detail here, since they are outside the scope of this article.

Although sensor fusion is important to virtually all phenomenological sciences and engineering disciplines, most research has been done in the field of *defence* research. One reason for this can be understood as follows. In *analytical* approaches, e.g. in a physics experiment, the measured quantities or interactions are often so small that the experimental setup has to be designed in such a way as to make sure that the desired quantity or effect is optimally measurable. If the quantities to be measured are small, the experiment is repeated many times and ergodicity and statistics are used to arrive at average values with low standard deviation. Especially in the case where one tries to prove or disprove the correctness of a theoretical model, this often is a good approach. A final point to note here is that – apart from intrinsic physical real-time aspects – such experiments very often can be repeated many times and that real-time constraints are not a bottleneck.

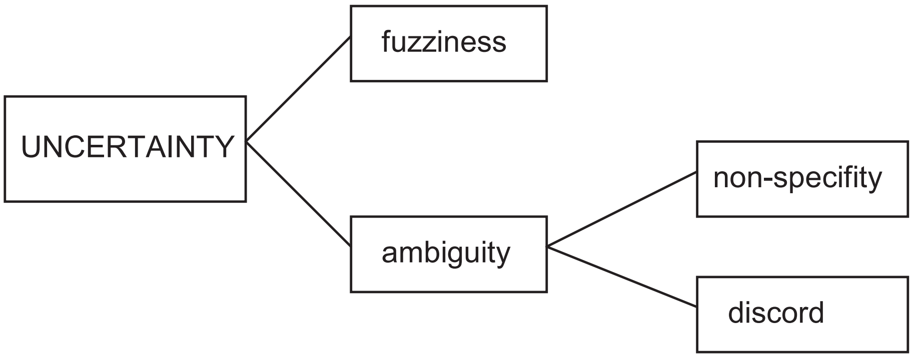


*Figure 2. Taxonomy of different types of uncertainty*

In engineering approaches the use of sensors is more *synthetic*, as illustrated e.g. in the field of factory automation. Here one deals with a well-defined problem such as the quality control of products on a manufacturing line, e.g. checking the soldering joints on a printed circuit board with an automated vision system. This problem certainly has real-time aspects, but the optimization can be done offline and the observation circumstances, like in the physics experiment, can be optimized offline, e.g. by testing the best combination of sensors, the proper cameras and illumination and parallel operation with more than one quality control station if the speed of production requires so.

In contrast, in military observation systems we deal with a situation that is far less comfortable than the situations described above: generally speaking it is necessary to assess in real-time an often complex situation, that almost certainly is outside one's complete control. Handling such observations requires the modelling of *uncertainty*. Apart from the ordinary problems such as noise and clutter, radar and electro-optical sensors operate also under adverse weather and atmospheric conditions, without any possibility to improve the circumstances of the experiment, or to repeat the experiment, under strict real-time constraints, with sometimes enormous consequences of false classification and even more serious penalties for non-detection. In addition, by the nature of the military métier, most targets of interest move at high speeds, try to actively or passively avoid detection or mislead sensors by jamming or using decoys and are designed in such a way as to present a minimal scattering cross section to commonly used sensors and thus to be virtually invisible ('stealth').

Under such circumstances it is clear that doing military observations invariably implies the modelling of uncertainty. Classically this is often done by applying statistical methods, notably Bayes' theorem to formulate a (multi-) hypothesis testing problem [11]. It is however clear that statistical uncertainty can only model part of the uncertainty. The different types of uncertainty, whose measures are now well established in classical set theory, fuzzy set theory, probability theory, possibility theory and evidence theory [11] are schematically summarized in Fig. 2. The breakdown distinguishes *fuzziness*, or vagueness due to a lack of definite and sharp conceptual distinctions and *ambiguity*, the situation where we are dealing with one-to-many relationships in the information obtained from sensors, yielding *non-specifity* in the case that the data leaves two or more alternatives unspecified, or even *discord*, i.e. disagreement in choosing from among several alternatives.

Methods that explicitly deal with ambiguity and partially overlapping hypotheses such as Dempster Shafer theory [12, 13] and the application of belief functions instead of probability densities have become popular. Of more recent date is the application of general fuzzy measures [14]. The difficulty inherent to making accurate observations in military applications and the lack of measurement statistics are the prime motivations to improve single sensor observations by merging (partial) inferences/conclusions from one sensor with inferences from the another one. Recent history shows that the nature of military operations changes rapidly: Although sensors are vital to the success of any military mission, it becomes at the same time much more difficult to interpret these observations. This can be explained by the introduction of stealth technologies (radar), the subtleties of 'peacekeeping' missions compared to classical, full scale warfare scenarios, and finally the complexities and greater vulnerability of navy vessels operating close to shore ('littoral warfare' or 'brown water operations'). Finally it should be noted that there is a genuine need to fuse sensor-generated information, at least at the higher levels of command and control: the throughput of the man-machine interface being the limiting factor. Although new sensors have been developed (e.g. GPS) and accuracy and resolution in space and time of most existing sensors have greatly increased over time, the bandwidth of the man-machine interface has *not*. The situation of having to deal with more information than one can process in a certain time is not dissimilar to the situation where a *lack* of information exists. Both situations involve taking decisions in the presence of uncertainty and would benefit from intelligent data reduction techniques, such as sensor fusion.

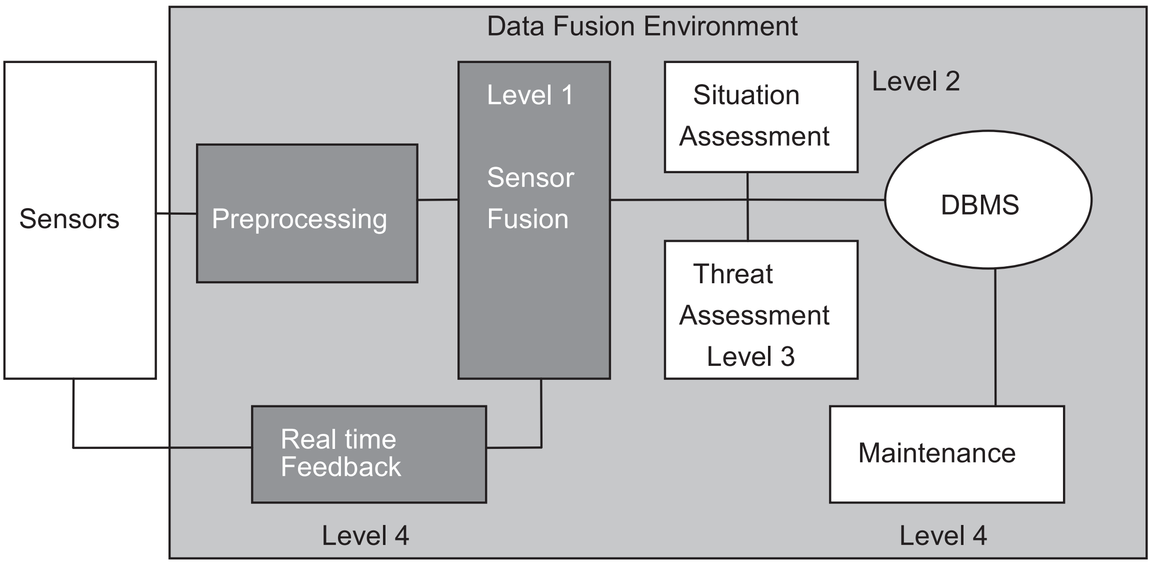


*Figure 3. The JDL process model for data fusion distinguishes 4 levels of data processing. The darker areas indicate the scope of this paper. The data base management system (DBMS) provides the environmental information.*

**Sensor fusion vs. data fusion**

Following the definition of the functional model of the data fusion process that is widely accepted in the military research community, as e.g. presented by Hall [15], and using the terminology as agreed by Joint Directors of Laboratories (JDL) of the Data Fusion Subpanel, multi-sensor data fusion is defined by the JDL as *"A continuous process of dealing with association, correlation and combination of data and information from multiple sources to achieve refined entity position and identity estimates, and complete and timely assessments of resulting situations and threats, and their significance."*

In this paper we restrict multi-sensor fusion to level 1 processing, with basic processes: data alignment, association and correlation, positional and identity fusion, complemented by the real-time part of level 4 processing ("maintenance") (Fig. 3). The reason for including part of level 4 perhaps seems strange at first sight, but is immediately apparent when maintenance is interpreted as the assessment of the status of each sensor *at short time scales* in order to keep it optimally tuned. Apart from optimizing individual sensors the monitoring of sensor performance makes it possible to perform '*sensor management*', i.e. to optimize groups of sensors, which is e.g. important for military observation systems where a large number of sensors co-operate in a coherent way and where part of the sensors may be damaged during operation (cf. a phased array radar system). The short-time sensor management system contributes therefore directly to the *robustness* of a system.

Early sensor fusion can be viewed as a two-step process, a direct fusion step followed by a complimentary fusion step (Fig. 4). The direct fusion process acts immediately on *raw* sensor data, after a possible preprocessing stage for alignment. This type of fusion is in practice limited to combining signals from similar sensors. In the next stage, in the complementary fusion process, very different types of sensors can be fused. In this stage it is possible that a considerable data reduction is achieved and that the information can be represented as a vector in feature space. Features such as range, position, orientation, effective cross section, shape, colour, etc. are extracted from the different sensors and combined in qualitative or quantitative ways. Combination of information from complimentary sensors can thus be seen as augmenting the dimensionality of the feature space. After this fusion step all sensor information has been fused and next one needs to combine feature vectors with existing, a priori information about the environment, collected from previously measured data or intelligence. This more abstract fusion process is typical for levels 2 and 3 of the JDL model and will not be considered here.

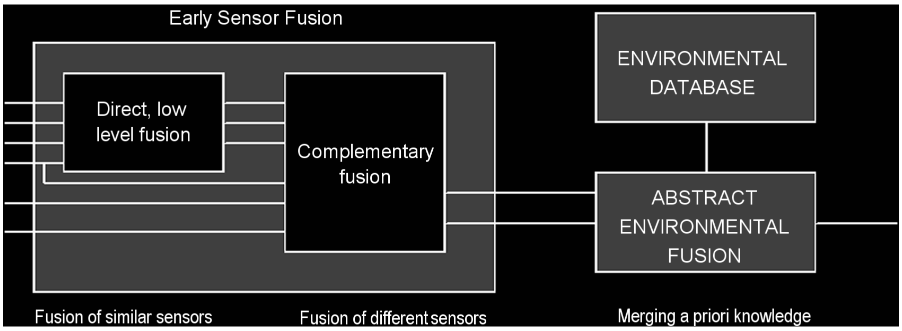


*Figure 4. Global information flow in sensor fusion: the level of abstraction in the fusion process increases to the right, whereas the data rate increases to the left reflecting the data reduction caused by the fusion process*

## Sensor fusion

*Motivation for applying sensor fusion*

The general motivation why multi-sensor fusion is pursued is generally answered in terms of "to improve the combined observations of different sensors and thus create a better situational awareness". From the operational point of view possible benefits include greater user friendliness because of data reduction, a greater robustness of the resulting observation system, higher reliability due to a higher plot rate and therefore a better performance of e.g. a tracking algorithm, and thus a better observation, even under adverse circumstances. Scientifically however one immediately faces difficulties with this formulation of the motivation, because it introduces subjectivity: How can one measure "situational awareness"? An even more complicated question to be answered is *what* exactly can be improved in conventional, single sensor observation and *how* different sensors can benefit from each others measurements. Answering these two questions is by no means trivial and they can in fact only be answered by considering a specific application. Before doing so, it is worthwhile to take a moment and review the various abstraction levels, methods, hardware and software implementation methods of observation systems and the different ways that they can interact and co-operate. The basic difficulty is the absence of a unique theoretical framework for objectively combining the information streams generated by the various sensors. The difficulty lies essentially in adequately describing the information content in each sensor stream. The amount of *useful* information in a data stream created by a sensor is of course dependent on the ultimate goal of the total of observations. Whenever this goal cannot be formulated in a clear, transparent and unambiguous way, it will be extremely difficult to develop synergy between the sensors and to compare the performance of the various fusion algorithms.

*Benefits and limitations of sensor fusion*

We first note that most of the benefits quoted in literature, see e.g. Waltz [16], are benefits that are exclusively associated with the presence of *multiple* sensors; they are *not* the benefits of sensor fusion *per se*. Most of these benefits are qualitatively and intuitively immediately clear. Globally we can distinguish three types of benefits:

- The first type of benefit of multiple sensors is an extended spatial, temporal, or spectral coverage of the observed phenomenon.
- A second type of benefit follows from statistical arguments: multiple sensors increase the measurement accuracy and from this an increased confidence may be derived, or at least a reduction in the number of hypotheses about the targets and thus an improved detection rate, c.q. a shorter detection time. Only in the last two cases a sensor fusion step is needed.
- Finally multiple sensors create overlap in observations and thus redundancy. If this redundancy is properly exploited in the system design, the maintenance (level 4) module will optimize the sensor scheduling and will result in the graceful degradation of system performance if sensors breakdown.

A quantitative aim of sensor fusion is to improve the accuracy of the observation, e.g. the position of a target. An example of this is e.g. the combination of a forward-looking IR sensor (FLIR) and a radar sensor. The inaccuracy in azimuth and elevation of the radar sensor is compensated by the more accurate measurements of the FLIR sensor, while the pulsed radar accurately determines the range. This example illustrates how a radar can initially detect a target, because of its wider field of view. Subsequently the FLIR can be cued using the inaccurate coordinates of the radar to initiate the FLIR measurement. Together they determine a small region of interest around the target, so that the benefit of fusion is an improved estimate (or reduced uncertainty) of the position of the target.

An interesting attempt to illustrate in a quantitative way the virtues of sensor fusion is described in [17]. In the article an odd number $N$ of identical sensors are fused with the aim to classify an observed phenomenon following a majority vote rule. The sensors are assumed to be statistically independent and the a priori probabilities are taken equal to $1/N$, corresponding to the principle of maximum entropy, equivalent with a minimum of a priori knowledge. Although the example is an idealized model case of identical, independent, unbiased sensors, all following the same statistics, using binary classification and a majority vote scheme as fusion aggregator, a number of qualitative results are worth mentioning here:

- Fusing data from multiple sensors (each with an individual probability of correct detection c.q. classification of less than 0.5) results in a decrease in performance in going from a single sensor classification to the multiple sensor fused result.
- If the individual sensors are very accurate (probability of correct detection larger than 0.95) sensor fusion cannot significantly improve the results of the inference process.
- The relative improvement in performance of an $N$ sensor fusion process over single sensor performance increases as a function of $N$ levelling off at about $N = 10$. Adding more, identical sensors does not pay off beyond this number (cf. Fig. 5).
- The maximum relative improvement of $N$ sensor performance for $N \to \infty$ compared to a single sensor is of the order of 15-25 %, depending on the fusion scheme. The maximum marginal gain is reached if the single-sensor probability of correct detection is in the range between 0.60 and 0.75.

Of course the numbers mentioned above should be treated with care because they depend on the type of aggregation operator chosen to represent the fusion process. Moreover these conditions refer to the simplified case of *identical* sensors, i.e. same positioning, calibration statistics, biases, sampling rates, bandwidths, sensitivities, dynamic behaviour and the same measured quantities. If a new type of sensor is added to the sensor suite, the dimensionality of the observation is increased and a substantial increase in information content may be expected, depending on fusion objective.

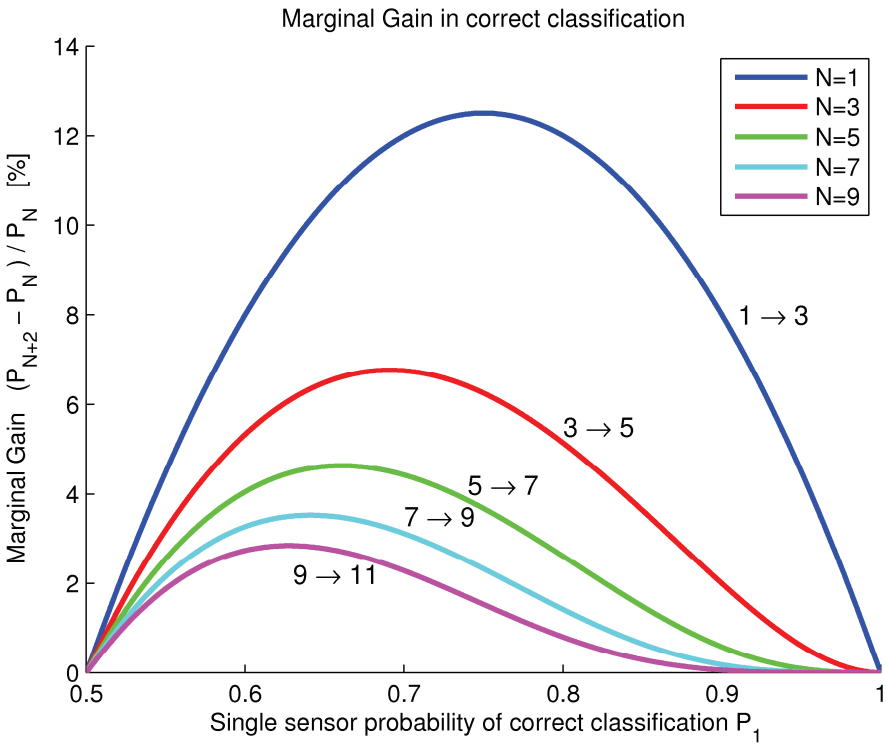


*Figure 5. Marginal gain in the probability of correct classification by adding sensors in going from 1 to 3, 3 to 5, until 9 to 11 sensors as a function of the single-sensor probability of detection P, as calculated in [17]*

In another recent study [18] the dependence of decision reliability on different fusion strategies for the case of 2 different sensors (i.e. sensors that are not identical) has been investigated based upon simple Boolean-type operators, along the lines of a probabilistic detection system (two hypotheses that are mutually exclusive and span up the whole universe of discourse). In a way this study extends the validity of the rules formulated above to more subtle fusion strategies than the 'majority vote', e.g. fuzzy decision strategies [19]. Also from these results we may conclude that for fusion to be successful the single sensor detection probability must be higher than 0.50 and that, in the case of two-sensor fusion, increases in detection performance of maximum 11% (from a single-sensor probability of detection $P_1$ of 0.80 to a two-sensor fused probability of detection $P_2$ of 0.90), or 17% (an increase from $P_1 = 0.80$ to $P_2 = 0.941$) have been calculated, depending on the fusion method.

Although these results do in no way preclude a substantial increase in the combined sensor performance brought about by a suitably chosen sensor fusion strategy, these numbers indicate that if measurements from identical, hard decision-making sensors are combined under the assumption of statistical independence, the marginal gain in performance will be limited to a percentage of a few tens, say 10-20%.

*Sensor fusion: synergy*

A general concept that is intimately connected to the idea of fusion is that of *ergodicity*, i.e. the concept that the outcome of an observation is unique, independent of the fact that one makes a series of consecutive measurements with one system, or that one makes $N$ parallel setups and combines the $N$ different outputs at one time. In the macroscopic physical world the concept of ergodicity generally is assumed to hold without exceptions, although in microscopic physics some examples have been found in spinglasses, where ergodicity does not hold. For a review the reader is referred to [20]. Throughout this paper we will assume that ergodicity holds. However it should be pointed out that this statistical principle is sometimes difficult to apply in the real world, such as in military observations, because of the fact that we are mostly dealing with *single,* non-repetitive isolated events.

The fusion of sensor observations, i.e. the combination of observations obtained from the same sensor at different times (temporal fusion), or the combination of simultaneous observations taken by a number of equivalent sensors (repetition), or the combination of sensor observations with a priori information obtained from previously measured data, is the focus of attention in the present work. The key concept in fusion is how to take advantage of *non-linearity*.

We are not so much concerned with increasing the accuracy of an observation by repeating a measurement a number of times and thus reduce the statistical variation of the average. Rather it is our objective to extract additional information out of this data set (reduction of information) by correlating (*not* superimposing) the measurement with observations from other sources. Moreover in the case of fast moving targets it is generally impossible to obtain a sufficient number of samples to apply statistics.

Two different types of measurements are of interest in military observation systems:

1. Determination of presence, position, orientation and speed of a target.
2. Identification and recognition of a target (type, identification friend vs. foe etc.).

Although identification clearly is an entirely different characteristic of a target compared to its position and speed and although the latter can generally be determined at much larger distances than those at which identification can the accomplished with reasonable confidence, identification can help improve the accuracy with which speed can be determined and vice versa. In particular the identification of a target may be helped through a wealth of observations, whereas establishing the position and speed of a target can only be accomplished by the few sensors. The identification of a target will be accomplished more easily, because of the higher dimensionality of the 'feature vector', provided that a good database of properties is available. An example is in underwater acoustics where non co-operative target recognition of vessels by means of their acoustic signature is standard practice.

The basic problem in recognition is to exploit the high dimensionality and representing data in such a way that differentiation between various possibilities becomes easier. Therefore the task of sensor fusion is the combination of, possibly incompatible, measurements and to try and construct from these an improved environmental picture.

In this modelling one needs to include the confidence level of the new measurement, as well as *how* to combine this information with the already existing picture. Various schemes have been proposed in the past: e.g. Bayes' rule of combination from statistics, belief measures, and Dempster Shafer theory. Although most of these methods have a sound theoretical basis, their application sometimes lacks theoretical justification or simply yield non-intuitive results in specific situations. This makes it difficult to compare results that are obtained with different methods.

*A systematic approach to Level 1 processing*

In this section we will outline a practical approach to sensor fusion in military observation, following the theoretical framework referred to as Level 1 data fusion in the JDL model. Two generic tasks are of importance in almost all observation processes, viz. detection and recognition. Despite the fact that different sensors and methods are commonly used to accomplish these tasks, it is obvious that the successful completion of one task will almost certainly have a positive effect on the other one. If however the identification and recognition tasks are considered to be "hard" decisions that result from the independent processing of separate sensors, interaction of the two processes can only take place *after* the first process has reached a decision. In executing more complex tasks we have already indicated in the section where we discussed generic sensor fusion that it may be more advantageous to allow for partial, delayed, or "soft" decisions which may offer a way to separate the various goals and thus allow us to break down a complex task into a hierarchy of relatively simple decisions. Partial or "soft" decisions can be combined at earlier moments in the processing chain without discarding too much information and thus offer a method for applying early fusion of sensor streams. Before discussing in more detail how soft computing methods can be used to achieve sensor fusion, we will first review the signal processing steps that are necessary to benefit from sensor fusion.

The classical way in which fusion is applied is by transforming a physical measurement into some hard decision (e.g. a 'plot', 'track', 'identity', etc.) that is communicated to the user via the man-machine interface, generally an optical display. The fusion process of the information shown on a number of different displays then takes place *in* the mind of the operator, who assesses the situation and makes a threat analysis. All these fusion processes take place in the human mind, *after* the sensor signals have been processed completely (Fig. 6a).

A first step towards true multi-sensor fusion is 'late' fusion (Fig. 6b): the construction of a special, goal-oriented architecture that fuses on the level of images, with the goal to enhance the image (e.g. combining IR and visible light images using some false colour scheme), or to fuse the consecutive plots of moving targets into a single track by taking into account some type of assumed target dynamics, or the fusion of tracks generated by different sensors (e.g. two radars or a radar and an electro-optical sensor).

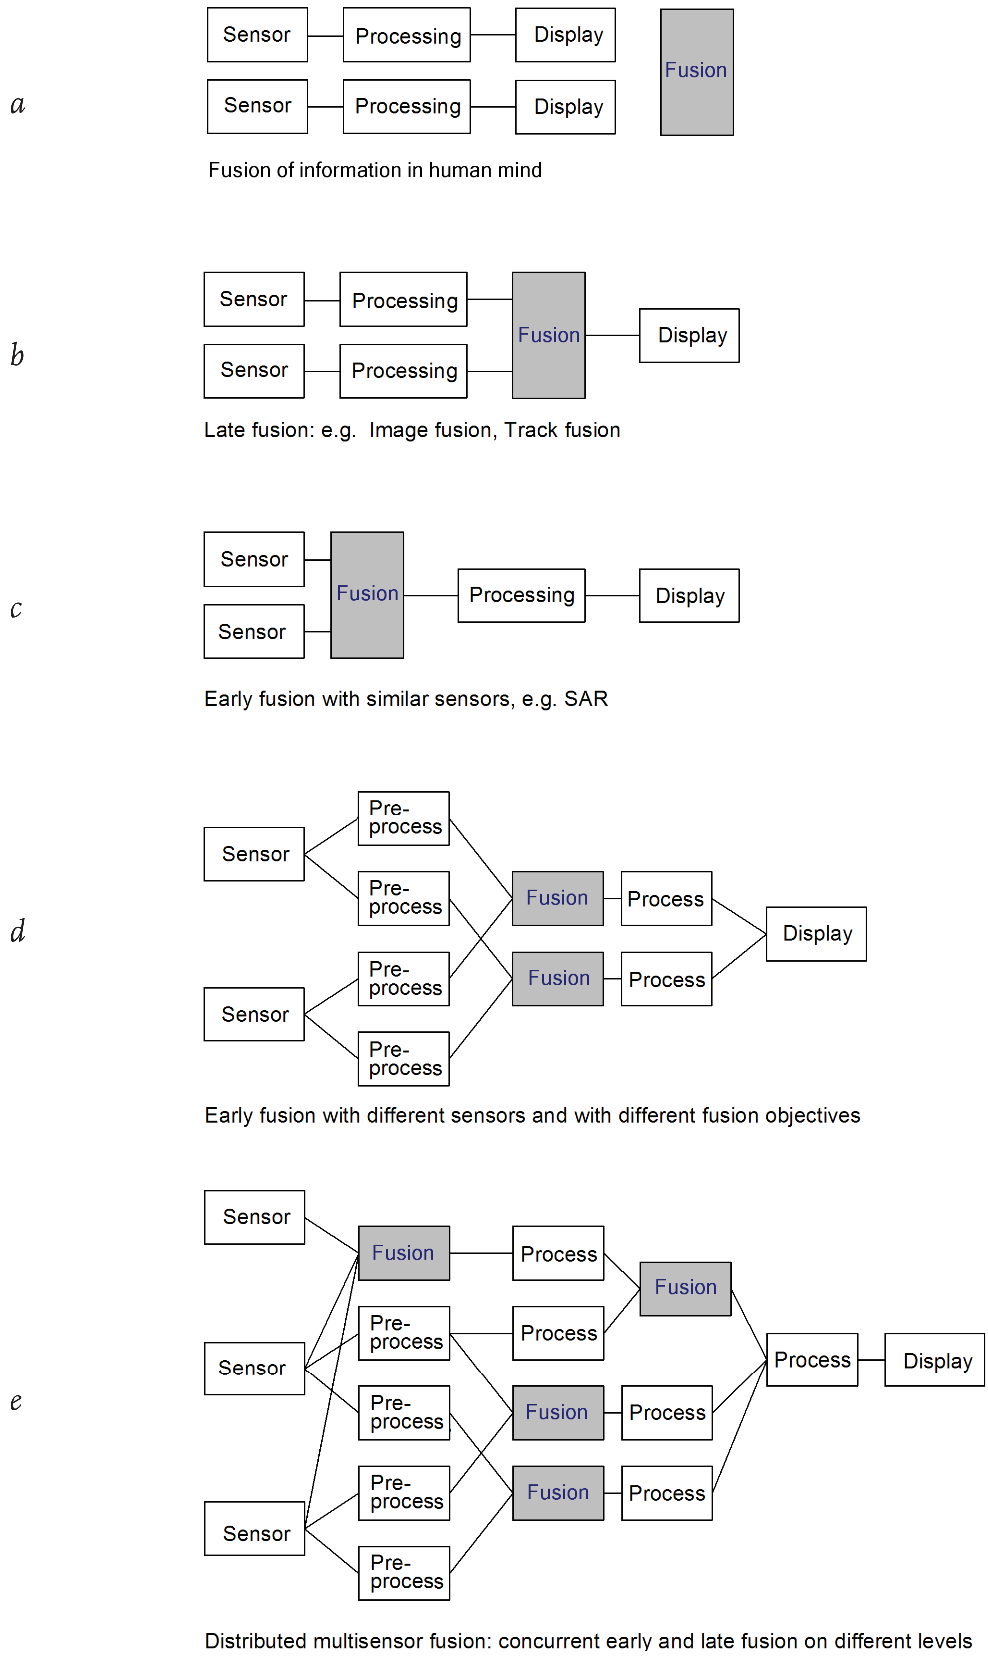


*Figure 6 a-e. Different fusion architectures starting from fusion in the human mind (a) and late fusion (b), via early, multi-sensor fusion with the same (c) and with different sensors (d) to the general case of multi-sensor multiobjective fusion on different levels concurrently (e)*

The results are qualitatively better, especially if the sensors are on different platforms, or, alternatively, under difficult circumstances (adverse weather, jamming) when one sensor can be supplemented by the other. In the case of *early* fusion with many sensors of the same type, one can process directly on the raw data without much preprocessing, possibly except for bias removal. In this sense one may view SAR (synthetic aperture radar) also as a 'temporal' fusion process, combining different samples of single sensor to simulate a kind of multi-sensor phased array with a much larger virtual aperture than that of the actual sensor (Fig. 6c).

The early fusion process that we would like to discuss is schematically given in Fig. 6d: a collection of interacting sensors, each contributing to one or more decision trees of positional c.q. feature declarations. Each tree represents a hierarchical decision process, building up from fast, low level, noisy decisions, up to well-founded, abstract decisions that require some time to take. The sensor fusion architecture outlined in Fig. 6e illustrates the ultimate goal of sensor fusion: mixed early and late fusion processes on different levels of abstraction, each subgoal benefiting from its own support set of sensors and with the results presented on a single display to the human observer in a representation that is highly informative, indicating at the same time alternatives, as well as the associated confidence levels.

We have already noted that application of multiple sensors in general can be beneficial to the accuracy of an observation, even though a low resolution sensor directs a sensor with a higher resolution ("cueing") without actually sensor *fusion* occurring. For making hard detection statements explicit knowledge of the clutter is required, and in order to improve single-sensor detection capabilities, one needs to carefully analyze the model assumptions that have gone into the system design. New, more detailed clutter models corresponding to the state of the art in sensor technology may be needed for the detection of harder targets. On top of that layer of processing sensor fusion may be applied. But it should be noted that sensor fusion can never make up for poor clutter modelling.

For any sensor fusion process to be successful, one has to properly prepare the raw signals originating from the single sensors. From a system point of view one has to determine the stage at which fusion has to take place (ranging from 'early' to 'late') in relation to the goal that has to be achieved by the fusion process. Once this has been decided the first step in preparing the sensor signals for fusion is *alignment*, to guarantee that the fields of view (FOV) have maximum overlap. Early fusion can only be successful if there is an *overlap* between the FOVs. The second step is the proper correlation of the same objects in the FOV of the sensors. This task can in practice be quite laborious when many targets are observed simultaneously (large FOV; scanning sensors). In addition attention has to be paid to the optimization of performance of the isolated sensor by *removal of the biases* in the observed quantities, before any of these preprocessing transforms and associations can be carried out.

*Fusion*

As we have seen in the previous section, it is worth to select the proper sensors and it is also necessary to spend sufficient effort in the preparation of the signals before they can be fused. It should be noted that fusion is *not* a magic way to reduce the quality or the price of a sensor and still get the same observation performance. Making accurate observations always requires a great deal of study, modelling and experimentation; the successful application of sensor fusion initially means more work than applying isolated sensors.

Depending on what the goal of the fusion process is, there are different time scales to consider:

- The maximum sampling rate, limiting the highest instantaneous bandwidth of a signal, which is essentially a measure for the sensor resolution: in a pulsed radar it is the range resolution and in a camera system it is the transverse 2D spatial resolution.
- The frame rate. This rate is important in extracting information from a time sequence of camera images, e.g. with optic flow analysis. By using the temporal correlation of objects in the pictures when the platform is moving, it is possible to make estimates of the distance of each of these objects.
- The batch processing rate: If no batch signal processing is performed as in the case of Kalman filtering, one is not forced to make a priori assumptions and may therefore be more accurate than Kalman filtering in the early stages of the signal processing chain. A disadvantage is that batch processing is considerable costlier than Kalman filtering in terms of processor (CPU) time.

Although we will here focus only on sensor fusion, many more aspects in modelling need to be considered, before an attempt to set up sensor fusion should be made. We mention here only a few:

- Clutter modelling: the type of statistics, correlation times and correlation lengths.
- Construction of an 'a priori' environmental data base, necessary to make (partial) decisions on identity and position.
- The modelling of the target dynamics if the target is moving and its significance with respect to improving classification.

If the observables of the fused sensor suite are mutually 'orthogonal', complementary fusion will invariably yield more information then each of the separate sensors can provide. It is therefore conceptually the simplest way to demonstrate the benefits of fusion in practice. In this context one could make an analogy between a single-sensor observation of the real world as a (stochastic) *projection* of the real world onto a sensor observation space. In this analogy, fusion can be seen as (partially) reconstructing the real world, representing it as the direct product space of all observation spaces of sensors that participate in the fusion suite. Effectively the dimensionality of the observation space increases by adding up the dimensions of complementary sensor spaces. Adding sensors of the same type through the Ergoden hypothesis basically improves the statistics of the observation in the particular sensor space. However the *dimensionality* of the single-sensor space does *not* increase by adding more sensors of the same kind.

In case that the observables of the sensor suite are *not* 'orthogonal', the fusion process can increase the speed with which the accuracy or resolution of the observation is achieved by acting as a smart scheduler, via cueing. The shorter response time is realized by first determining areas of interest via the sensor with the lower resolution or accuracy and the largest FOV, and then focusing attention on these areas using the high accuracy sensor, instead of scanning the entire area with a high accuracy sensor with a small FOV. In addition this type of sensor fusion increases the robustness: if the cued sensor fails or is jammed, the other sensor can take over, although with lower resolution or accuracy *(graceful degradation)*.

An example of multi-sensor fusion with different sensors is the combination of a radar measurement and an optical image: if an airplane is observed by radar, the range from observer to the airplane is accurately measured, while azimuth and elevation can only coarsely determined. In contrast, an optical measurement provides azimuth and elevation with relatively high accuracy, whereas the uncertainty in range is high. Combination of the two sensor types can considerably diminish the absolute uncertainty in the position of the airplane in 3D space, which is a natural consequence from the two complementary measurement principles.

Finally we remark that from a system theoretical point of view we can express the expected effect of the fusion process symbolically as: $S_1 \oplus S_2 \geq S_1 + S_2$ ,where $\oplus$ represents the fusion operator and $S_i$ is a quality measure associated with sensor $i$.

*Fusion with fuzzy aggregation operators as a way to reduce complexity*

In our study we have concentrated on the synergy of sensors at the signal level ("*early fusion*"). Although this does by no means preclude the use of a priori information or taking into account any human-generated inputs and feedback, our study focuses specifically on sensing, because the signals are not yet distorted or corrupted by signal analysis operations, and because there is relatively little room for subjectivity. The attractiveness of this approach is of course that by operating close to the primary sources of information, one expects to be able to significantly enhance the detection and recognition processes by applying sensor fusion.

A sensor fusion system that receives raw signal inputs from all sensors retains control over all primary sensors and has, at least theoretically, a number of advantages over secondary (level 2 and higher) fusion. Apart from the larger information content of raw information, it should however be noted that each fusion step requires a certain processing time and in early fusion it is effectively the slowest sensor in the fusion suite that determines this latency, even if we neglect the execution time for the fusion process itself. In addition the latency is increased because fusing information from autonomous, asynchronous and dissimilar sensors requires synchronization. A designer of a sensor fusion suite should be aware of this and take precautions to ensure that the pile-up of latencies does not degrade the real-time performance of the overall fusion system, or jeopardize the quality of the fusion process, e.g. by constructing a deficient situation awareness picture.

It is relatively straightforward to illustrate these ideas by the improvement of operation of a target tracker during the observation of a manoeuvring target in cluttered areas. A variety of different sensors can be used to generate plot reports and these can be combined on the basis of a simple confidence criterion that is based on the presence of clutter for a particular sensor. In [21] this has been illustrated. However, though a useful idea, this example basically supports the idea of increasing robustness by increasing the number of different sensors. Our goal is more ambitious: we would like to improve the quality of the single-sensor conclusions in such a way that

$$P(S_1 \cup S_2) \geq P(S_1) + P(S_2),$$

or if this condition appears to be too strong, at least

$$P(S_1 \cup S_2) \geq \max(P(S_1), P(S_2)),$$

where $P(S_i)$ indicates the performance or 'added value' (the effective added information accumulated over time) of stream $S_i$, measured by sensor $i$. From this formulation it is clear that in order to model sensor fusion, we will need *nonlinear* operators.

The earliest attempts to combine measurements from multiple sources are by Bayes [22]. He introduced the notion of conditional probability $Prob\,(A|B)$, defined by:
$Prob(A|B) = \dfrac{Prob(A \cap B)}{Prob(B)}$ i.e. the probability of $A$, given that event $B$ has occurred. This definition is easily extended to $n$ observations obtained by $n$ sensors. There are a number of difficulties connected with the application of the Bayesian sensor fusion formula:

- difficulty of assigning a priori probabilities;
- complexity when there are multiple hypotheses and/or multiple conditional events;
- requirement that hypotheses have to be mutually exclusive and exhaustive;
- absence of uncertainty modelling.

In trying to find an appropriate way to model fusion and take advantage of the nonlinearity of the process, Dempster and Shafer (DS) created a generalization of Bayesian theory that allows the incorporation of uncertainty by using (overlapping) probability *intervals* and uncertainty modelling to determine the likelihood of hypotheses based on multiple evidence [23]. The essential generalization of DS theory is that not all hypotheses need to be mutually exclusive as in the Bayesian theory. In DS fusion evidence is assigned both to single and more general propositions, instead of assigning directly a probability to hypotheses as in Bayesian theory.

Noting that belief and plausibility measures are both examples of Sugeno's [24] λ-fuzzy measure $g_\lambda$, the question arises whether it is possible to combine the intuitive ideas on sensor fusion and the properties of $g_\lambda$. We will show that in contrast the basic probability assignment in DS theory, fuzzy $g_\lambda$ measures can indeed be utilized for the problem under consideration. We will take a closer look at this in the following and propose to view the multisensor fusion process in terms of a synergy between (sets of) sensors that are grouped in such a way as to support a certain decision or hypothesis. Instead of attempting to make a decision (detection or classification) in one step, either by a single sensor, or by a linear combination of a group of sensors, it is proposed to combine supporting evidence for a hypothesis in a hierarchical way by building a tree structure that combines at the lowest level clusters and in the next levels aggregates the outputs of several initial clusters in superclusters and so on. At each level in the tree decisions need to be made from different sources with different weights. This is conveniently modelled by the fuzzy λ-measure $g_\lambda$ ($0 \le g_\lambda \le 1$). In particular we have in the absence of relevant information towards the classification/detection goal: $g(\emptyset) = 0$ and $g(A) \le g(B)$ if $A \subseteq B$. This coincides with the intuitive feeling that if the evidence support is larger (i.e. if we observe the same scene with more sensors), that then the information content should also increase. In addition the following property holds for all $A, B \subset X$ with $A \cap B = \emptyset$ :

$$\exists \lambda > -1 \quad g(A \cup B) = g(A) + g(B) + \lambda\, g(A) g(B) .$$

This again supports the intuition that adding more independent data ($A \cap B = \emptyset$) co-operates towards an increase in confidence about the final decision. In addition both intuitive features about the fusion of two independent sensors are reproduced, viz.

$$\lambda \ge 0 \qquad g(A \cup B) \ge g(A) + g(B),$$

i.e. fusion is more than superposition and

$$-1 < \lambda \le 0 \quad g(A \cup B) \ge \max(g(A), g(B)),$$

implying that even if the level of confidence is larger than -1 (as reflected by the negative λ) then it may still be fruituitous to apply sensor fusion. In the event that λ = 0, i.e. the case where all sensors have the same importance and completely cover the universe of discourse, the degree of importance $g_\lambda$ towards the final decision becomes additive and coincides with the definition of a probability measure.

Following ideas put forward in [25], sensor fusion may also be modelled using the concept of fuzzy integration. For a review on the role of fuzzy integrals in the framework of multiple criteria decision-making see [26]. A fuzzy integral may be interpreted as an aggregation functional of subjective evidence, where the subjectivity is expressed in the fuzzy measure, and integration is defined over measurable sets [27]. In contrast to normal (Lebesque) integrals, fuzzy integrals are *non-linear* functionals. It is exactly this nonlinearity and the possibility to include a fuzzy measure $g_\lambda$ that is attractive in the context of fusion. Formally Sugeno's fuzzy integral is defined in the following way: Let X be a set of elements (e.g. sensors, features or classifiers) and let h(x): X→[0,1] denote the confidence value belonging to element x∈X (e.g. the class membership of data determined by a specific sensor (classifier)), then the *fuzzy integral* of h(x) over a subset E of X with respect to the fuzzy measure $g$ can be calculated. The evaluation of the fuzzy integral may be interpreted as evaluating the degree of agreement between objective evidence h(x) and the expected observation outcome (the hypothesis). We will not discuss the properties of this fuzzy fusion operator here, but note that it is ideally suited to combine information from different sources *without* having to deal with the combinatorial explosion.

A similarity between fuzzy fusion (FF) aggregation and the way DS theory fuses data from different sources is that both make use of fuzzy measures: DS uses the belief measure exclusively, whereas the FF operator uses the $g_\lambda$ measure. For $\lambda \geq 0$ this measure is equivalent to the DS belief measure. The conceptual difference between both methods is twofold:

1. The frame of discernment (the universe of discourse) is different in both methods.
2. There exists a clear separation of objective and subjective uncertainty in the case of FF.

We will illustrate these points in the following: For the FF scheme the frame of discernment contains the information sources (the sensors) related to the hypothesis under consideration, whereas in DS theory the universe of discourse contains *all* possible hypotheses. In combining the different information streams, the fuzzy fusion aggregator fuses all sources according to their relative a priori importance as well as to the degree to which each sensor supports the hypothesis under consideration. In contrast, the fusion process in DS theory associates with each knowledge source a belief function that is defined over the power set of the set of hypotheses and combines these in the fusion process. The evaluation of Dempster's rule of combination therefore has exponential complexity $O(2^N)$, where $N$ equals the number of hypotheses under consideration. In contrast, in FF one fuzzy integration has to be calculated, which implies that $g_\lambda$ has to be calculated $nN$ times, where $n$ is the number of sensors. The evaluation of the fuzzy integral can then be carried out in $O(n)$ steps. The second advantage of fuzzy aggregation is that both the weighting with the degree of support by which a sensor supports a certain hypothesis, as well as the weight of importance of a certain sensor, reflecting a subjectivity or an a priori confidence in the particular sensor, are explicitly modelled.

We therefore conclude that the formalism of fuzzy measure theory offers an opportunity to model the process of sensor fusion in a natural, intuitive and adequate way, allowing arbitrary sensors to be fused and allowing different ways to weigh various combinations of observations. As an example of the application of the belief measure $g_\lambda$ consider three sensors, labelled 1-3 with belief measures 0.1, 0.3, and 0.2, respectively. If sensor 2 ("the most decisive sensor in support of the hypothesis under consideration of the three sensors") is combined with one of the other two sensors (1 or 3), the combined evidence as reflected by $g_\lambda$ must be larger than the sum $g_2+g_3$. This is indeed the case, as follows from the definition of $g_\lambda$ with $\lambda = 3.109$: $g_{23} = 0.687 > g_2 + g_3 = 0.3 + 0.2$ and also $g_{23} > g_{21} > g_{13}$, since $0.687 > 0.493 > 0.362$, which we would also expect intuitively. In addition we have that $g_{123} = 1$ and $g_\varnothing = 0$, thus having consistency within the powerset of the three sensors.

## Conclusions

In this article we reviewed the added value of sensor fusion in military observation systems. Sensor fusion is motivated by the expected qualitative and quantitative improvement of observations and thus of situation awareness. We have focused on early sensor fusion and found that the performance enhancement due to extending one sensor to a suite of identical sensors and assume a majority vote is limited to a few tens of a percent. Early sensor fusion offers the best perspective to maximally benefit from multiple sensor observations, but at the same time demands extensive data acquisition efforts. The real-time constraints and the need to (re)use intermediate fusion results in different decision processes suggest that in early fusion soft decisions are more effective than hard decisions. Finally it can be concluded that from the theoretical point of view key concepts of fuzzy logic, such as fuzzy belief and plausibility measures and Sugeno's fuzzy integral, provide us with suitable mathematical tools to combine soft decisions and describe the type of synergy behaviour expected of a fusion process.

**References**


[1] J. Llinas and E. Waltz, Multisensor Data Fusion, ArtechHouse, Norwood, MA. 1990.

[2] D. Hall and R. Linn, "A Taxonomy of Multi-Sensor Data Fusion Techniques", Proc. 1990 Joint data fusion symposium, vol 1, pp 593-610, 1990.

[3] C.B. Weaver, Ed., "Sensor Fusion", Proc. of SPIE, vol 931, Orlando, FL, 1988.

[4] P.S. Schenker, Ed., "Sensor Fusion, Spatial reasoning and scene interpretation", Proc. of SPIE, vol 1003, 1988.

[5] C.B. Weaver, "Sensor Fusion II", Proc. of SPIE, vol 1100, Orlando, FL, 1989.

[6] P.S. Schenker, "Sensor Fusion II, Human and Machine Strategies", Proc. of SPIE, vol 1198, Philadelphia, PA, 1989.

[7] A.J. van der Wal, "Application of fuzzy logic control in industry", Fuzzy Sets Syst 74, pp 33-41, 1995.

[8] A.J. van der Wal, "Fuzzy Logic: Foundations and Industrial Applications", ed. D. Ruan, Kluwer, Chapter 14, 275-311 (1996).

[9] D. Ruan, Z. Liu, L. Van den Durpel, P. D'hondt, and A.J. van der Wal, "Progress of Fuzzy Logic control applications for the Belgian Nuclear reactor BR1", Proceedings EUFIT '96, Aachen vol 2, 1237-1241 (1996).

[10] A.J. van der Wal and D. Ruan, "Controlling the output power of a nuclear reactor with fuzzy logic", JCIS 97, vol 1, 136, 1997.

[11] G. J. Klir and B. Yuan, Fuzzy sets and fuzzy logic, Prentice Hall, New Jersey, 1995.

[12] A.P. Dempster, "Upper and lower probabilities induced by a multivalued mapping", Ann. Math. Statistics, 38, pp 325-339, 1967.

[13] G. Shafer, A mathematical theory of evidence, Princeton University Press, Princeton, 1976.

[14] K. Leszczynski, P. Penczek, and W. Grochulskki, "Sugeno's fuzzy measures and fuzzy clustering", Fuzzy Sets Syst., vol 15 pp 147-158, 1985.

[15] D.H. Hall, Mathematical techniques in multisensor data fusion, Artech House, Boston, 1992.

[16] E.L. Waltz, "Data fusion for C3I: A tutorial", Command, control, communications and Intelligence handbook, EW Communications Inc, Palo Alto, CA, pp 217-226, 1986.

[17] P.J. Nahin and J.L. Pokoski, "NCTR plus sensor fusion equals IFFN", Proc. IEEE Trans. Aerospace Electronic Systems, Vol AES-16, pp 320-327, 1980.

[18] B.V. Dasarathy, "Fusion strategies for enhancing decision reliability in multisensor environments", Opt. Eng. 35 (3), pp 603-616, 1996.

[19] L.A. Zadeh, "The linguistic approach and its application to decision analysis", in Y. C. Ho and S. K. Mitter, eds., Directions in large scale systems, Plenum Press, New York, pp 339-370, 1976.
[20] A.S. Mikhailov and A.Yu. Loskutov, Foundations of Synergetics II; Chaos and Noise, Springer, Berlin, 1996.
[21] E. Sviestins, "Multi-sensor tracking for air traffic control and air defense", ATC Systems, pp 10-16, 1995.
[22] T. Bayes, "Essay towards solving a problem in the doctrine of chances", Philos. Trans. Royal Soc. London, vol 53, pp 370-418, 1763.
[23] S.C.A. Thomopoulos, "Sensor integration and data fusion", J. Robotic, Syst. Vol 7(3), pp 337-372, 1990.
[24] M. Sugeno, "Fuzzy measures and fuzzy integrals: A survey", in Fuzzy Automata and Decision Processes, North Holland, Amsterdam, pp 89-102, 1977.
[25] J.M. Keller, H. Qiu, and H. Tahani, "The fuzzy integral and image segmentation", Proc. NAFIPS June 1986, pp 324-338, 1986.
[26] M. Grabisch, "On the equivalence classes of fuzzy connectives-The case of fuzzy integrals", IEEE Trans Fuzzy Syst.3(1), pp 96-109, 1995.
[27] W.F. Pfeffer, Integrals and measures, Marcel Dekker, New York, 1977.