1

Multisensor Data Fusion

1.1. Issues at stake

Why would anyone seek to combine multiple sensors while this inevitably increases cost, complexity, cumbersomeness and weight, etc.?

The first reason that often comes to mind is that we can use multiple identical sensors to improve their performances. Yet, if n sensors provide the estimation of the same value with the same signal-to-noise ratio (SNR), at best, the joint use of those n sensors will lead to a gain of $\sqrt{n}$ in relation to that SNR, while multiplying by a factor close to n all the material factors of the resulting system (cost, weight, bulk etc.). Additionally, in such cases, there are often simpler and more effective solutions available – particularly solutions based on temporal integration of the data from a single sensor.

This example highlights the fact that combining multiple sensors is only irrefutably advantageous in the production, in specific conditions, of information, which a single sensor (whatever its type) would be unable to provide. In practice, in order to identify the situations where it is helpful, we consider three categories of objectives that a multisensor approach may serve. Each of these categories can be illustrated by looking at a few situations, where observation and surveillance systems are used.

The first major benefit of multisensor systems is their robustness in any observation context, which is usually a decisive factor in the choice to use such systems. For example, the system may be less vulnerable to disturbances – whether intentional (counter-measures specifically targeted at a particular wave form or wavelength, but that do not affect those of the other sensors), or natural (atmospheric phenomena that adversely affect one sensor but not the others, such as multiple trajectories to a low site, and the effect of an evaporation duct on radar, or atmospheric transmission in optoelectronics). Other examples include the ability to function in an environment or conditions of observation that impede the operation of a single sensor, but do not have the same effect if a variety of appropriate observation devices are used simultaneously. Thus, various types of weather-related disturbances, geometrical masking effect, problems of spatial or radiometric resolution, or limitations in detection range may render one of the sensors (though not always the same one) non-operational. In the same vein of ideas, there is also the problem of representativeness of certain data used to train a given sensor to later recognize specific objects, in relation to the reality on the ground. If the training data used are not representative, the only way to recognize the target objects is by cross-referencing the data from different sensors.

The second point of superiority of multisensor systems is the acuity and richness of the information gleaned. For example, one sensor might discriminate between targets independently of their size on the basis of the features of their rotating parts, while another sensor, which is not capable of observing these features, distinguishes them by their size. The combination of the distinguishing capabilities of these sensors will, obviously, help to refine the taxonomy finally generated. Similarly, the relevant association of a radar – which provides good distance – and Doppler resolution with a passive optical device with good angular resolution will generate a fine-grained analysis in a four-dimensional space – those dimensions being the site, the bearing, the distance and the Doppler. Partial non-availability of data to one sensor (unobservable measurements, non-availability of training data, etc.) can also be compensated for by data from another sensor.

The third great capability of multisensor systems is a better reaction time when presented with the most complex requests, because they can share out the required tasks between the different sensor components used. Indeed, each of the different sensors can, in parallel, focus on dedicated functions, which are appropriate to their capabilities. The synergy of the work of acquisition and processing then optimizes the reactiveness of the whole system. For example, a radar can quite easily perform a quick "pre-screening" of the space – a survey with a high detection rate but also a high false alarm rate – with a simple wave form, in order to provide a small number of potential targets for detailed analysis with an optoelectronic identification system.

To begin with, it is useful to note that for these three major categories of benefits reaped with the multisensor approach, the expected gain can only be obtained by appropriate complementarity of the sensors used and their processing. Hence, above all else, the quality of a multisensor system is dependent upon the diversity of its components in the face of the problem at hand. Consequently, the functional specificity of each of these components, the diversity of the data they provide, and the exponential increase in the volume of data to be processed are all unavoidable complexifying factors for the design and deployment of multisensor data fusion modules.

In addition, combining multiple sensors only makes sense, correlatively, to carry out functions that a lone sensor of any type would be incapable of performing, in any and all foreseeable circumstances. This means that the system's performances hinge upon the capabilities of one or other of the sensors at different times. (The same sensor will not always be fully functional, and different sensors will perform better at different times; otherwise we would only need to look at one sensor – we would have no need for the others). What follows from this is that we must constantly fuse relevant data with defective data. Yet, as we will see, blithely combining good and bad data always yields an inaccurate result, as the bad data "pollute" the good. Therefore, we need to constantly use all of the available information, both exogenous and previously collected, to assess and qualify the observations feeding from the different sensors, and exploit those observations on the basis of their relevance. Of course, this further increases the diversity and volume of the information needing to be integrated, which in turn further increases the complexity of the processing, because at all levels, this qualitative dimension needs to be integrated in detail.

In view of this significant increase in the complexity of the system and its processing, its real-time operation necessitates objectives in terms of reactivity, and therefore rapidity, often associated with constraints in terms of "on board ability". A crucial objective in terms of data fusion processing, therefore, is to find a compromise between the complexity needed to ensure the desired benefits and the simplicity needed to be compatible with the operational constraints.

1.2. Problems

In practice, the combination of different sensors may be useful for two types of goals:

– Distinguishing hypotheses in a discrete set: this is the case for the functions of detection, extraction, classification, recognition, identification, counting or diagnostics more generally.

– Estimating variables in a continuous set: of particular note here are the functions of localization, tracking, navigation or, more generally, metrology (quantification of descriptors on the basis of observations).

In both cases, the fusion algorithms must not only exploit the richness of all the available information as best they can, but also satisfy the expression of high-level operational requirements imposed by the pooling of different means of observation in increasingly complex systems.

As a support and as a reference for the coming discussion, consider the expected evolution of a generic classification system. Figure 1.1 illustrates the traditional structure of such a system, where the objective is to find the class of objects O_i which an observed object most closely resembles, choosing from an exclusive and exhaustive set of possible classes. These objects will be entities in the broadest sense of the term: vehicles, types of ground occupation, infrastructures, states and generic situations, etc.

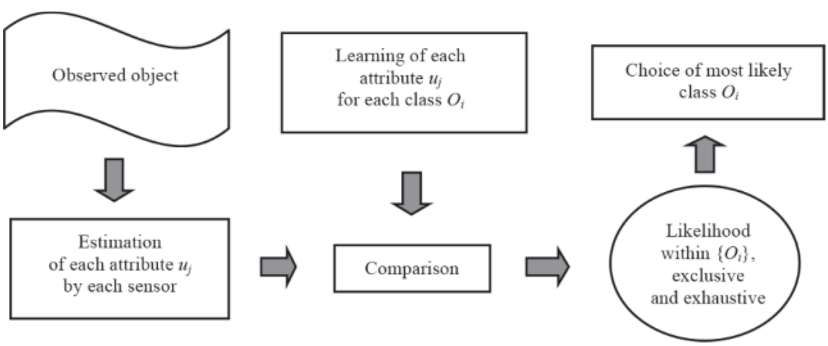

Figure 1.1. *Usual approach in classification*

In this process, for each class O_i, the system undergoes prior training, learning the possible values of a number of discriminating factors or attributes u_j (e.g. descriptors of size, shape or kinematics), for an object belonging to that class. These values are then compared, for each class, to the observations of those same distinguishing attributes on the object needing to be classified. The resulting measurement of the resemblance gives the likelihood that the observed object belongs to each of the classes O_i in turn. By maximizing this likelihood, it is possible to identify the class to which the observed object actually belongs.

The necessary integration of this classification function in complex systems where a number of very diverse components interact, requiring specific uses to be made of the available dataset, leads to the general approach presented in Figure 1.2.

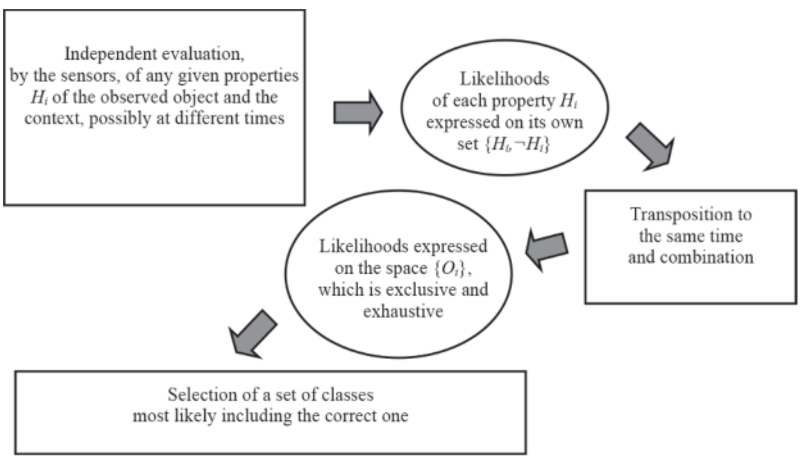

Figure 1.2. *Evolution of the concept of classification*

In this system, the sensors all give independent evaluations of a number of properties H_i of the object, or of the context of the observation, and possibly at different times. These properties are unrelated to one another, and in no way constitute an exclusive or exhaustive set, unlike the set of classes discussed above. For example, the system may evaluate whether or not a land vehicle has caterpillar tracks, whether or not an aircraft has rotating parts, whether a terrain is irregular with relief features, etc. As with the traditional system, the properties of interest can be evaluated by comparing the discriminating values u_j characterizing them against the same values observed on the object being analyzed.

This yields a certain number of likelihood functions, each expressed on its own set $\{H_i, \neg H_i\}$. These likelihood functions must be adjusted to the same moment of a decision, particularly using models of the evolution of the properties, and combined in order to evaluate the likelihood that the observed object belongs to one of the classes O_i previously listed, which, for their part, constitute an exclusive and exhaustive set. This combination must, of course, integrate a previously acquired description of each class O_i in terms of the properties being examined. Yet, it must be borne in mind that in practice, the relations describing the object classes O_i in terms of the properties H_i are usually poorly defined, because of imperfect compatibility between the useful information and the available information.

At this stage, to do as we did above and determine the single most likely class may prove senseless, if not actually dangerous: two classes may have similar likelihoods without a significant difference, but they may lead to extremely different actions, so the consequences may be drastic if a mistake is made. One example would be confusing an enemy target with an allied or civilian vehicle. The goal at this level, therefore, is actually to identify a set of classes that is as small as possible, but which is most likely to contain the correct class, and where the uncertainties of discrimination are brought to the attention of the decision-maker so as to grant him/her a better understanding of the consequences of the choices.

A common use of this scheme is to directly identify a class O_i with a property H_i. It highlights two of the fundamental advantages of this approach in comparison to that presented in Figure 1.1: first, each class can be evaluated on the basis of different attributes specifically appropriate to it and therefore more effective; second, the different classes can be processed separately depending on the availability of information, which means that we can exploit an incomplete fragment of knowledge, or enrich it gradually.

Of course, the scheme in Figure 1.2 can be extended to the situation where certain properties have a number of states greater than the two discussed here, H_i and $\neg H_i$.

This example illustrates the need to manage, in detail, the uncertainty, the distinct sets, the evolutions over time, the combination of information fragments with complex relations between them and the principles of decision-making.

More generally, the requirements in terms of functional development relate to the major areas introduced below.

1.2.1. *Interpretation and modeling of data*

The data input in the fusion processes are obviously the output from the sensors, such as measurements, signals or images, but also all of the knowledge that helps to draw full benefit from those data – e.g. databanks, expert knowledge, previously learnt features, or models identified previously or online – be they dynamic, statistical, descriptive or behavioral.

Consequently, these data are extremely varied, first in terms of their nature and secondly in terms of the use that can be made of them, but also, above all, in terms of the disparity of their points of insufficiency. The goal of data fusion is to exploit this diversity as fully as possible so as to gain the greatest possible benefit from the relevant available information, without it being polluted by the imperfections. Therefore, it is crucially important to correctly interpret the potential contribution of each piece of information, and thus model it in the theoretical framework, which corresponds most closely to its peculiarities. The difficulty then lies in jointly processing the different theoretical frameworks involved in the same form.

The most challenging of these imperfections are uncertainty and imprecision. Uncertainty expresses a lack of knowledge regarding the occurrence of an event (e.g. it may rain), while imprecision characterizes a value that is not accurately known (e.g. estimated speed of a sea current). For example, uncertainty would be caused by insufficient or inappropriate training of the system, or by atmospheric conditions, which reduce the perceptive capacities. Imprecision typically arises from insufficient resolving power or approximate descriptions. These problems can be taken into account due to uncertainty theories.

The data are also usually incomplete, because the system has not had all of the necessary training, or because of temporary non-observability events of interest. Certain desired characteristics can therefore not be directly evaluated on the basis of appropriate observations, and must be approximated as closely as possible on the basis of any other available information. Hence, the idea is to reduce the initial uncertainty as far as possible with an appropriate processing architecture, developed in the context of uncertainty theories.

The reliability of the gathered data is certainly one of the most sensitive points, as the main aim of data fusion is to compensate for the deficiencies of one sensor by using one or more others. It is therefore helpful to formalize the reliability of each piece of information and thus model its impact in terms of the uncertainty induced about the observations, using theories capable of handling this uncertainty.

Finally, apart from the observations, the fusion system must adequately exploit all of the previous exogenous or contextual knowledge accessible to it. This knowledge, usually gained from human assessment or interpretation, is of course tainted with subjectivity, which must be accounted for in terms of the uncertainty and imprecision caused.

In addition to the diverse nature of the information taken into account, we also need to consider the heterogeneity of the respective imperfections in the different information fragments. These fragments must therefore be able to be processed jointly in the same overarching theoretical framework.

1.2.2. *Reliability handling*

As the main goal of data fusion is to compensate for the deficiencies of one sensor by using one or more other sensors, the process must, at all times, be robust when faced with a loss of reliability of one or more of the pieces of information being processed – that is, it must *a minima* ensure that the good-quality pieces of information are not polluted by the erroneous ones. This is crucially important, because when a good and a bad piece of information are carelessly fused, the result usually inherits the poorer of the two quality levels.

Thus, this objective can only be served if the system has sufficient knowledge of the relative reliability of the different sources, and is capable of exploiting that knowledge effectively. This poses the problem of evaluating that reliability as accurately as possible on the basis of additional information, either compiled beforehand or acquired in real time as regards the context and the environment, or possibly provided by exogenous sources. This additional information must be processed in an appropriate theoretical framework, which is capable of handling the uncertainty regarding the more or less pertinent knowledge that the new information provides, as well as the uncertainty caused in the process of exploiting the observations.

Furthermore, it is important that the fusion process ensures that information about reliability is integrated into the processing of the observations. With this in mind, we need to define an appropriate process architecture, and employ fitting operators to adjust the knowledge drawn from the observations on the basis of the relevance of each knowledge fragment.

1.2.3. *Knowledge propagation*

Whichever theoretical set we use, the rules usually put forward for data fusion assume that the input sets and output sets are all the same. Yet in practice, this is hardly ever the case. To begin with, the inevitable diversity of the input data (both in terms of type and quality), as discussed above, means that in modeling those data, we have to use sets which are adapted to the particular distinguishing potential of each type of data and which are therefore necessarily distinct. Additionally, in accordance with the need expressed in terms of the expected decision, we have to fit the available knowledge into a set appropriate for that decision, which must necessarily be higher level than the input sets. Furthermore, in complex systems, a number of resources are pooled, and interact with one another on different levels. The same piece of information may be used for different purposes, and if so it will need to be expressed in different sets.

Also, in order to implement operational systems, it is necessary to take account of the observations delivered at different times with regard to situations likely to evolve in the meantime, and consequently to deliver conclusions at very specific moments, which are, themselves, different from the instants of observation. It is therefore useful to be able to look at a piece of knowledge available at a certain time in a given set, and transpose it to a later time in the same set, using a model of the possible evolution of the situation over time.

The implication of all this is that there is a need to develop the capability to transpose a given piece of knowledge in one set into a second, different set. This transformation is, of course, possible only if the relations linking the elements of the second set to those of the first are known. However, in general, the definitions of these relations given by the available expertise are uncertain or imprecise, and account needs to be taken of this in the processing performed. Also, the relations in question must integrate any inter-dependency between the knowledge fragments used.

1.2.4. *Matching of ambiguous data*

For the reasons of complementarity discussed above, a single object observed by multiple sensors is usually analyzed by each of them in a set specific to the object. Thus, the different sensors used acquire different views of the object and by comparing and contrasting these views, the system is able to gain a more accurate picture of it. Matters become more complicated, however, when numerous objects are being observed simultaneously. In this case, the difficulty lies in correctly combining the observations taken of the same object by each sensor, ensuring that only the data relating to that particular object are being fused. This very common issue, which affects all possible types of sets, can be exemplified in detail for two classic contexts.

The first problem is that of matching spatially ambiguous data, which is better known as "deghosting". For example, imagine that two remote passive sensors observe a target in a plan passing through both the sensors. Each sensor then reports an azimuth at which it is detecting a target, and the target being observed is localized by triangulation, i.e. the intersection of the two directions reported by the sensors. Now, if two targets are present in the plane, each sensor reports two azimuths, so triangulation finds four intersections. Two of those intersections correspond to the actual positions of the targets, whereas the two others are artifacts, also known as "ghosts". The system then needs to try to eliminate the two artifacts so as to unequivocally determine the positions of the two targets.

The second problem is that of fusing temporally ambiguous data. This time, imagine that any sensor detects two nearby moving targets at a given moment in time, and another sensor detects the same targets asynchronously, i.e. at two different times. The positions detected by the second sensor will, obviously, be different to those detected by the first sensor, because the targets have moved in the interim. The problem then becomes one of determining which detections from the first and second sensors correspond to the same target.

Generally speaking, ambiguous data fusion requires us to examine the available information to identify the data likely to characterize the similarity of the observations, with a view to matching them. In general, unfortunately, the available data are insufficient to reliably determine the correct association when data fragments are considered in isolation, and the process can only work by using numerous imperfect fragments of information, jointly. The resulting uncertainty must be taken into account when modeling these data, and processed when matching them.

1.2.5. *Combination of sources*

Combination of sources is, of course, the heart of the data fusion process. Hence, naturally, it is the focal point of the main difficulties. First, as we have seen in section 1.2.1, the diverse nature and quality of the data taken into account force us to model each piece of information using the most appropriate theoretical formalism. Therefore, we now need to combine data expressed in different theoretical frameworks. Consequently, in each case, it is useful to find the formalism that is capable of encapsulating all of the issues at stake, while minimizing the complexity induced.

Additionally, as introduced in section 1.2.3, the data being fused are usually expressed in different sets, and the result of the fusion, in turn, needs to be expressed in a different set from the input sets. For example, for a classic problem of classification such as that discussed above, the input sets are those peculiar to each distinguishing attribute, and the output set is the set of classes of objects. Therefore, it is helpful to be able to simultaneously fuse and propagate the data, while ensuring as "optimal" as possible an exploitation of their "useful" content.

A crucially important point for a combination operator is the definition of the underlying logic, and the expression of that logic in terms of axioms needing to be satisfied. The logic might, for instance, be that of a conjunction (consensus), disjunction (plurality), etc., and the axioms typically the definition of the neutral element, monotony, commutativity, associativity etc. Naturally, the aim when choosing the logic is to satisfy the requirements imposed on the fusion process. As the desire is usually to maximize the amount of information output, conjunction is generally the first candidate to be considered.

However, the underlying logic also needs to compensate for the pitfalls which may occur in certain particular situations. Of these, the recurrent problem of conflict between sources is a major concern, which can render the conjunction utterly meaningless. If, for example, one source gives a set “A” of solutions and a second source produces a set “B” which is totally separate from “A”, the conjunction of these opinions yields a null set of solutions! An in-depth analysis of these situations of complete discord between sources shows that they necessarily correspond to the use of a theory or method in conditions, which violate the axioms or principles of that approach. For example, the set of solutions considered is not exhaustive, or not exclusive, or not all the sources are reliable, etc. In this case, the best approach is to analyze the conflict, identify its cause and, having duly rectified it, repeat the modeling of the problem. If this proves insufficient or impossible, the only option is to look for the formulation of the combination which exploits only the consistent portion of the available data, ignoring data which are not mutually validated. This is often a tricky task, both in terms of fitting into a rigorous theoretical framework and ensuring pertinent implementation.

Another major challenge for data fusion is taking account of the dynamic aspects, linked particularly to the fact that the sources do not all deliver their data at exactly the same time. Thus, in general, the combination is referenced to the moment the result of the data fusion is delivered. Therefore, it is usually necessary to extrapolate the knowledge from each source to that moment, by modeling the temporal evolution of the objects. In addition, certain sources may yield information which is more or less frequent, more or less up to date, etc. In particular, this covers problems of prediction, updating, revision of knowledge, etc.

As well as the other unavoidable difficulties, it must be remembered that the processing of the data delivered by the sources also has to include all the contextual or expert information needed to deal with the imperfections in those data. This is a key point in the performance of data fusion, discussed above. It is therefore necessary to put in place the formalism which effectively positions the information about the quality of the data when modeling those data.

One final point, which is not overly easy to deal with, relates to the fact that generally, the sources being combined are actually not independent, as they are assumed to be by most conventional combination laws. If ignored, the interdependency relations between the data being fused may cause undue confirmation or undermining of certain points of view. Therefore, the effect of such relations needs to be modeled and taken into account in the processing so as to prevent any harmful effects on the conclusions of the process. On the other hand, in certain cases, the correlations between the data may, in fact, provide additional useful information, and therefore necessitate a particular appropriate exploitation.

1.2.6. *Decision-making*

This step in data processing is the final operation, which actually produces the required intelligence on the basis of the observations carried out, or directly the actions required in view of the observed situation. Decision-making may be involved at different levels, and the task can prove difficult in practice, depending on the nature of the problem at hand. To introduce the different types of difficulty encountered, Figure 1.3 illustrates the paths taken by the available data for decision-making for different major types of systems.

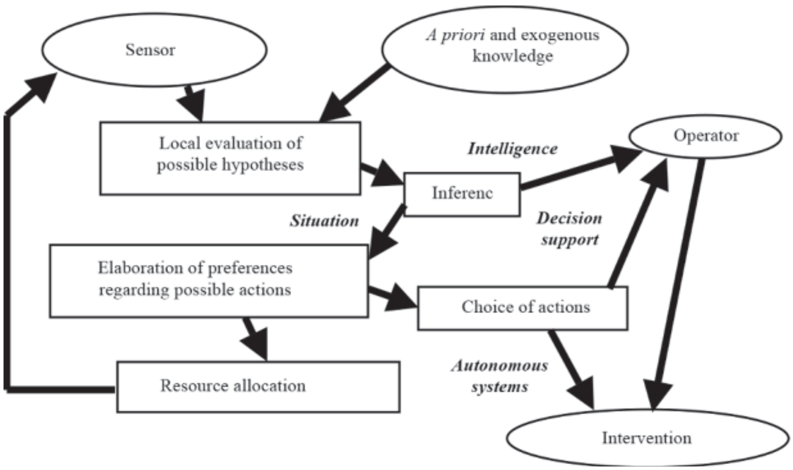

Figure 1.3. *General process of decision-making on the basis of data*

The first step is to look at all of the data, whose nature we have already discussed above, to evaluate the likelihood of different hypotheses pre-determined as relevant to the particular problem at hand. Fusion of those likelihoods then generates an estimation of the situation, which can either be transmitted directly to an operator to serve in decision support or exploited by automated processing to help an autonomous system. In the latter case, preference criteria need to be defined for each of the possible actions, and optimized in order to determine which action to perform. These criteria must express the mechanisms of choice, which an operator would use to take a decision in the context of the situation produced by the fusion of the observations. Their formalism, therefore, must be compatible with an imperfect knowledge of that situation, subjective preferences that are often difficult to express, complex mechanisms of comparison, and compromises needing to be found between contradictory objectives. The action thus determined will be directly implemented by a fully autonomous system such as an unmanned, non-linked vehicle, or suggested to an operator for approval in the context of decision support.

Finally, the process will usually improve the system's decision-making capacity by using a resource allocation function that sends requests back to the sensors in order to obtain the information likely to enrich the discerning capacity of the decision step, as quickly as possible.

The first difficulty in such a process relates to the greater or lesser compatibility that it is possible to ensure between the informative content of the input data and that of the required conclusions. The models of the information at all levels of the chain, and the underlying decision-making principles, therefore need to be defined in order to produce only legitimate conclusions in regard to the only available knowledge, both in terms of their nature and their acuity. The more or less complex decision-making principles which serve that aim must, correlatively, be able to be expressed rigorously in terms of operators defined in the theoretical framework adopted for processing the information, which generally requires specific developments, and the formalism of the conclusions must be capable of expressing strictly the available knowledge as accurately as possible. For example, as mentioned at the beginning of section 1.2 with regard to the illustrative problem of classification, it may prove pointless or even dangerous to determine the single most likely hypothesis if it is based on an insignificant difference between the most likely hypotheses, while the consequences may be drastic in case of error. In this case, the need is actually to identify as small a set of classes as possible, which is most likely to contain the correct class, but where the uncertainties in discrimination are brought to the attention of the decision-maker to facilitate a better grasp of the consequences of their choices.

The second major difficulty of decision-making processes is handling the inevitable inconsistencies. This type of situation is, for example, illustrated by Condorcet's paradox, which involves three decision-makers D1, D2 and D3, charged with choosing between three possible actions A1, A2 and A3. Suppose the decision-makers' preferences are as follows:

– D1: A1 > A2 > A3;

– D2: A2 > A3 > A1;

– D3: A3 > A1 > A2.

If the actions are compared two by two in order to establish the consensual preferences by majority vote, the result is that which is illustrated by Figure 1.4 – in other words, an intransitive set of equal preferences which renders any conclusion impossible.

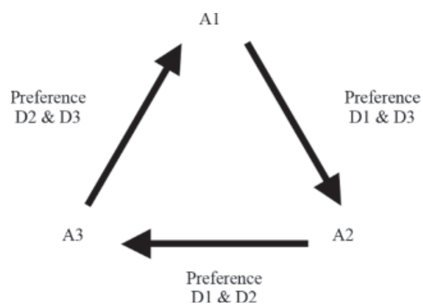

Figure 1.4. *Intransitiveness of Condorcet's paradox*

More generally, the properties of collective decisions were examined by Arrow, on the basis of five axioms responsible for their consistency [ARR 63]:

– Unrestricted domain: all individual choices can be catered for.

– Unanimity: if x > y for all voters, then x > y for the group vote.

– Pairwise independence: the collective ranking of two options depends only on the individual rankings of those two options alone.

– Completeness: all pairs can be ranked (indifference is a possibility).

– Transitivity: if x > y and y > z then x > z.

Arrow demonstrates that the only decision-making rule capable of satisfying all five of these axioms is dictatorship, i.e. only taking one opinion into account, ignoring all the others! Thus, any practical solution must be the result of a compromise between rationality (expressed by Arrow's axioms), effectiveness (to reach a conclusion whatever the circumstances), and consensus (which respects the plurality of opinions). All of these notions therefore need to be accounted for in detail in the decision-support algorithms.

The third difficulty is to approximate the behavior of human decision-makers, as closely as possible, with all of their peculiarities, relating particularly to the subjectivity of perception, to knowledge, intuition, greater or lesser temerity, level of wisdom, etc. Evidently, this has a direct impact on the interpretation and modeling of the information being manipulated, on the plurality of the criteria used, on the logic underlying the process of decision-making (the extent to which it is conjunctive, disjunctive or consensual; complete or partial aggregation, etc.), and on the architecture of the process – particularly in terms of centralization/distribution of the decision.

In practice, the perimeter which is of interest to us in our coming discussion of multisensor observation systems is that of situation elaboration, as the choice of actions to be performed stems from the operation of those systems. Thus, we will, on the one hand, be dealing with decision-making for extraction of useable intelligence, and, on the other, with expression of information in a formalism, which is compatible with the techniques of decision-making for choosing actions.

1.3. Solutions

Evidently, there is no universal, ready-to-use solution available to deal with all of the problems presented above. On the other hand, there are a number of theoretical frameworks, each specifically designed to deal with a different aspect of the requirements expressed, and which provide good coverage when used together. Generally, these theories were not initially designed for data fusion, but they exhibit a satisfactory potential for this purpose. The objective of this book is therefore to put forward a set of original tools exploiting, first, the specificities of each of the theories in order to deal with a particular aspect of the problem, and second, all the synergies which can be established between those theories to ensure the overall consistency of the chain of processing in which they constitute the different links.

1.3.1. *Panorama of useful theories*

The theories and techniques potentially concerned by the process of multisensor data fusion as defined above are graphically represented in Figure 1.5. Areas of overlap between the boxes indicate the links that can be formalized between the theories in question, with a view to their joint exploitation.

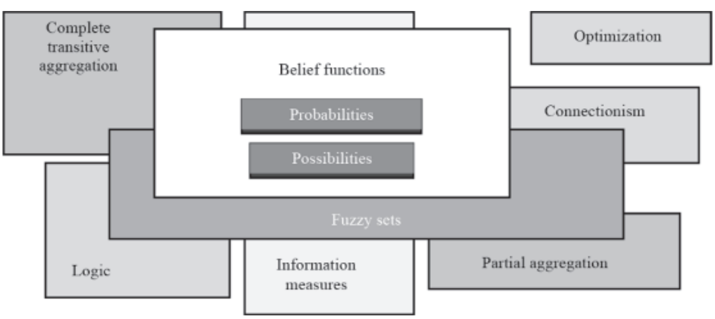

Figure 1.5. *Diagram of the main theoretical frameworks concerned by multisensor data fusion*

More specifically, they include:

– theories designed to deal with uncertainty: the theory of belief functions, with its two intrinsically distinct particular cases – probability theory and possibility theory, which we will examine later on; these theories are obviously at the heart of the need expressed above regarding the multiple imperfections of the data being manipulated;

– theories designed to deal with imprecision: essentially the fuzzy sets theory; quite apart from the duality between uncertainty and imprecision which fuzzy sets theory helps to exploit, in conjunction with possibility theory, it can easily be combined with uncertainty theories to jointly deal with uncertainty and imprecision; we will see later on that it is possible, for example, to determine the probability of a fuzzy event; it therefore also fits into the working model outlined above;

– measures of information, which can be used to evaluate the degrees of uncertainty and imprecision conveyed by the processing, in particular for the purpose of evaluating the data fusion processes; such measures constitute an invaluable addition to the aforementioned theories, but relate more closely to the aspects linked to the engineering of the process or real-time resource management; hence they are beyond the strict context of data fusion methods *per se* in which we are interested here;

– the different methods of multi-criterion aggregation, which can be classified into two main categories. The first category relates to complete transitive aggregative methods, which synthesize a single criterion which needs to be optimized in the decision space. The archetype of these methods is the *multiple attribute utility theory*; uncertainty theories can, fairly easily, give rise to methods in this category (e.g. see sections 6.3.5 and 6.4), or be coupled with other methods, such as utility functions. Therefore, with these methods, we are well equipped to deal with the imperfections in the data conveyed in the decision-making processes. The second category includes partial aggregation methods, which compare all the solutions two by two before drawing the conclusion about the preferences obtained. The ELECTRE family of methods is certainly the most representative of this approach; in addition, most partial aggregation methods have a "fuzzy" version to regulate their behavior (for example ELECTRE 3 for the aforementioned ELECTRE family of methods). Fuzzy sets also have their own approach to partial aggregation, based largely on exploiting fuzzy order relations. Uncertainty theories can also deliver a partial aggregation type approach on the basis of binary comparisons (section 7.4). However, multi-criterion aggregation methods are designed for problems stemming from the choice of actions to perform, for which they are able to find appropriate solutions quite easily, rather than for intelligence extraction. Therefore, they will not be discussed further in this book;

– mathematical logics, which facilitate high-level reasoning processes. These logics can advantageously be combined with imprecision and uncertainty theories to integrate the imperfection of the knowledge (fuzzy logic, possibilistic logic, etc.); they are more closely linked to a specific exploitation of the information produced by the multisensor data fusion than to the system itself which we are interested in;

– connectionist approaches – particularly neural networks. The idea is to repeat a behavior that has been directly learnt from a sufficient number of real cases. Thus, it is a useful support (especially for complex learning processes), but one which must necessarily be based, from a methodological point of view, on analytical approaches such as those mentioned above, to overcome the problems of generalization on the basis of imperfect learning. With this in mind, for example, it is relevant to mention neuro-fuzzy approaches and certain analogies which have been established with Bayesian approaches, but the connectionist aspect is not at the heart of the breakthroughs likely to serve the requirements expressed previously;

– robust optimization methods, which are crucial in searching for solutions in large spaces using complex cost functions, at all levels of the process; however, in this case the need is fairly generic and disconnected from the concept of data fusion per se.

In conclusion, in view of the above remarks, the coming discussion will focus on uncertainty– and imprecision theories, with the aim being to discover the tools capable of serving the requirements expressed.

1.3.2. *Process architectures*

The recurrent problem in this area is the problem of the level of fusion, i.e. the position of the fusion operator in the chain of processing between the raw data from the sensors and their high-level exploitation, and correlatively that of the centralization or distribution of the processing. In fact these two aspects are closely connected, as data fusion close to the point of output from the sensors necessitates centralized processing of those data, while fusion of the data at a higher semantic level facilitates local processing of each measurement, which is generally exploited to compress the useful information and thereby decrease the throughput needed in data transmission.

To begin with, the type of fusion that produces the richest result is that which takes place closer to the sensors, when the data are least compressed, and can therefore be compared in greater detail. However, this common-sense principle may be incompatible with other requirements or constraints. In particular, it may prove senseless to fuse the data at a very early level, and it may be useless or even damaging to the quality of the result, depending on the granularity and the intention of the desired conclusions.

For example, the interpretation of perfectly registered multispectral spatial images to determine soil occupation would be based on the fusion of pixels, because they are naturally and easily associable and correspond to the spatial resolution of the information being sought. On the other hand, the extraction of particular objects in airborne optoelectronic and RADAR images would focus on the fusion of attributes of objects estimated on both sides, because it makes little sense to fuse pixels of different size and geometry, and additionally the nature of the final characterization must be pertinent to the level of the objects.

In addition to this, we may come up against a certain number of operational constraints such as the limitation or vulnerability of communications for a delocalized function, or the volume and time of the processing with regard to requirements such as reactivity, time restrictions or onboard capability.

Another problem relating to the architecture of the processing is the need to respect the hierarchical ranking of the information fragments, which may cover very different forms. For example, not all the sensors used necessarily deliver information of the same semantic level, and the fusion of such data must begin with the lowest semantic levels, working up gradually to the highest semantic level with the processes of extraction and dissemination of the usable information, in accordance with the ontology provided by the application.

A different type of hierarchization relates to the inclusion of quality information that may be gathered about the sensors (acuity, reliability and usefulness, etc.), and serves to help manage the observations. This information thus needs to be integrated into the formalism of processing these observations to usefully modify their impact, by properly exploiting the difference in quality between the different sensors. The quality data can, of course, be fused themselves, at a different level to that of the observations, while respecting the particular effect that each measurement is intended to produce.

In the process of fusion, we must also rank the effect of the different pieces of information on the final conclusions of the processing, on the basis of their (more or less specific) utility for the problem at hand, in view of the potential for that information to evolve (context, requests, etc.).

Finally, the hierarchization may be linked to the particular relations that exist between certain pieces of information, starting with statistical dependencies and, as before, to the operational constraints, relating to the distribution of the sensors or the processing capacity, for example.

Of course, the architecture of the data fusion process is also guided by the desire to create synergy between the different analytical functions. For example, we will see the advantage in having a global approach to target extraction for surveillance (detection, numbering, classification and tracking, etc.), which leads to these different functions being implemented simultaneously, rather than sequentially, as happens in single-sensor mode. The different sensors may also be led to cooperate with one another, to mutually enrich their respective capacities. Finally, judicious sharing and parallelism of the tasks usually helps optimize the effectiveness of each component, so the yield of the whole system is enhanced.

What emerges from this brief overview of existing fusion architectures is that it would be ill-advised – dangerous, even – to attempt to set a universal methodology in stone. Every application requires a solution specific to it, which can only come from an in-depth analysis of the peculiarities of that application, in keeping with the few common-sense rules mentioned above. The important thing for our coming discussion is to be aware of these different architectural problems in providing developers with all the processing tools they need to deal with the variety of situations they are likely to encounter.

1.4. Position of multisensor data fusion

Before discussing the development of the tools necessary for multisensor data fusion, it is helpful to situate this issue within the broader framework of data fusion in general, and identify the intended uses of the aforementioned tools.

1.4.1. *Peculiarities of the problem*

Data fusion actually covers a very broad range of problems, depending on the nature of the information being exploited and the goal of the procedure, as shown by the discussion presented in [BLO 01]. With regard to the information being exploited, four major categories can be distinguished, *a priori*:

– The observations captured by the sensors.

– The knowledge available in the form of databases, expert knowledge bases, information, intelligence, etc.

– The preferences used in multi-criterion decisions, with multiple decision-makers, etc.

– The multiple regulations, the conflicts and inconsistencies between which need to be resolved in order to determine the rights, responsibilities, etc., of all the actors in all cases.

The output, for its part, may lead to the development and/or updating of two types of model:

– A model of the real world, of which we are seeking to form an estimation on the basis of an imperfect perception of it; this approach stems from what is usually called an “inverse problem”.

– A model of the ideal world which we wish to create, e.g. by way of a decision which satisfies several points of view, or by balancing several regulations.

Clearly, in this panorama, the input to multisensor data fusion comprises observations and knowledge (contextual, *a priori*, exogenous, etc.), and the objective is the development and updation of models of the real world.

The perimeter we are interested in for our discussions is even, more specifically, that of sources providing concurrent information fragments, which mutually enrich one another when compared. In particular, this excludes signal – or image – processing to reconstruct a particular physical value, e.g. the processing of networks of RADAR antennas (beam forming by calculation, etc.) or stereo-vision. Indeed, in this case, first, the processing methods are highly specific and well known, and second, the set of sensors and processing constitute a single sensor yielding an original physical measurement.

1.4.2. *Applications of multisensor data fusion*

While there are, as yet, few implementations which truly draw the full benefit from the techniques presented above, the range of applications is still very broad. Of course, defense systems are the main applications – particularly with regard to tactical situation assessment, cooperative multiplatform engagement, aerial defense systems, surveillance and alarm systems, recognition systems and intelligence. The requirements in terms of data fusion are increased, in particular, by the networking of all the means of observation, command and intervention.

Another sector of interest, which is highly similar to the previous one, is that of global security, be it in the prevention, alert, intervention or resilience phase. Here, again, all of the available resources are pooled. The objective may be the protection of persons, property or interests, and the requirements in terms of data fusion are very similar to those in the area of defense.

A number of other domains should also be mentioned, though. The extent of the requirements in these domains is not yet fully defined, but they have a high potential for investment. They include:

– information systems in general, which are intended to handle varied datasets, and often designed for decision support;

– autonomous vehicles, such as drones, which exploit and respond to numerous measurements of their environment;

– robotics in general, where data captured by different sensors are used to automate functions of greater or lesser complexity;

– agile multisensor perception systems, used particular for observation of the environment;

– non-invasive diagnostic means – notably in the medical or engineering field;

– and more generally, cooperative smart systems.

In the discussion to come, the different techniques presented are, as far as possible, illustrated in terms of the implementation in one of these domains – most often that of defense, because of its richness and its advances, but with a view to facilitating transposition to the other domains as soon as possible.