1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM'99) • Atlanta, Georgia • Sept. 19-22, 1999

MULTISENSOR INDUSTRIAL INSPECTION AND GRADING USING ELSA

Michael D. Naish
Computer Integrated Manufacturing Laboratory
Dept. of Mechanical and Industrial Engineering
University of Toronto
Toronto, Ontario, Canada M5S 3G8
naish@mie.utoronto.ca

Elizabeth A. Croft
Industrial Automation Laboratory
Department of Mechanical Engineering
University of British Columbia
Vancouver, B.C., Canada V6T 1Z4
ecroft@mech.ubc.ca

ABSTRACT

The Extended Logical Sensor Architecture (ELSA) has been developed for industrial applications, particularly, the on-line grading and classification of non-uniform food products. This architecture addresses a number of issues specific to industrial inspection including modularity, scalability, and design by non-expert users. To address these needs, the sensors are encapsulated by a logical sensor model, providing robustness and flexibility. The construction methodology is based upon the object model which represents object classifications through combinations of primary features weighted by fuzzy membership functions. The features guide the selection of sensors and processing routines; the classifications determine the rulebase used by the inference engine for process decisions.

1. INTRODUCTION

The ability to consistently produce high-quality products is important to the success of manufacturing and processing operations. Traditional quality assurance methods have often relied on human operators who use visual cues in order to determine product quality. Such methods are tedious, time-consuming, and inconsistent. Many products, especially biologically formed items, exhibit non-standard and non-uniform characteristics. For example, products such as fish, chicken, tomatoes, and other types of produce may, even within a single classification or grade, vary widely in appearance. The problem is compounded by variations in the product characteristics within a species, region, or industry. As a result, industry has turned to automation to address these grading and quality assurance needs.

Much of the success of machine vision and multisensor systems is dependent upon the ease of use of these systems for industrial users who may understand the process but not the details of the sensor technology. However, such systems tend to be highly specialized for each industry, and, at least partially, developed or contracted "in house." Therefore it is important that the development and operation of a multisensor system is orderly, comprehensible, and simple.

The architecture and approach presented herein provides a methodology for the design and implementation of multisensor systems for industrial automation; it is particularly suited to the on-line grading and classification of non-uniform food products. Although inspection is the focus of this work, it is intended to be applicable to a variety of automation tasks which may benefit from a multiple sensor system.

1.1 Multisensor Integration

Automated systems which make multi-factored decisions about non-uniform products on the basis of information from a single sensor have had limited success. Often, there is simply inadequate data for a proper product assessment. The transition to multiple sensors can extend the capabilities and improve the robustness of existing systems.

A system which employs multiple sensors may enjoy several advantages over single sensor systems [1]. Redundant and complementary information increase accuracy and enable features undetectable with individual sensors to be perceived in less time and with less cost.

The majority of the work in the area has been conducted for mobile robot navigation [2] and military target tracking [3, 4]. Much less published work has considered sensor integration architectures for industrial applications.

1.2 Industrial Applications

In the area of quality assessment and assurance, machine vision is often used to gather the bulk of the required information, especially for the grading or classification of non-uniform (biological) products. Other sensors, such as scales, mechanical measurement devices, and ultrasound are employed to gather information that is used to enhance the machine vision data.

There have been a number of vision-based multisensor systems developed for quality assurance and assessment over the past decade. Examples include systems for poultry [5], fruits and vegetables, [6], herring roe [7], printed circuit boards [8], and baked goods [9]. In most of these applications, ad-hoc methods are used to develop a sensor integration system to monitor the process. Such systems lack a formal architecture, and are typically designed by experts in machine vision and/or systems integration. This can result in difficulties with the use and maintenance of the system for the everyday user. Additionally, upgrading the system to change or add additional sensors and/or requirements often requires the system to be redesigned. This is a problem for industrial users whose requirements in terms of speed, feature recognition, accuracy, and other process monitoring parameters invariably change over time.

2. SYSTEM ARCHITECTURE

The basic structure of the Extended Logical Sensor Architecture for multisensor integration may be decomposed into three groups, according to the following tasks:

1. **Sensing:** The acquisition of information from the environment which is used as the data for inference and decision making.

2. **Inference:** The combination of the sensory information with information contained in a knowledge base to infer decisions.
3. **Action:** The conversion of decisions into commands and signals which control process machinery.

The structure of ELSA is illustrated in Fig. 1. An object-oriented approach to the system configuration has been adopted. The encapsulation of the primary components leads to a scalable and flexible system which is particularly suited to industrial grading tasks. The system may be easily reconfigured to adapt to advances in sensor and processing technologies or changing market demands. This paper shall focus primarily on the sensing group; details of the architecture may be found in [10].

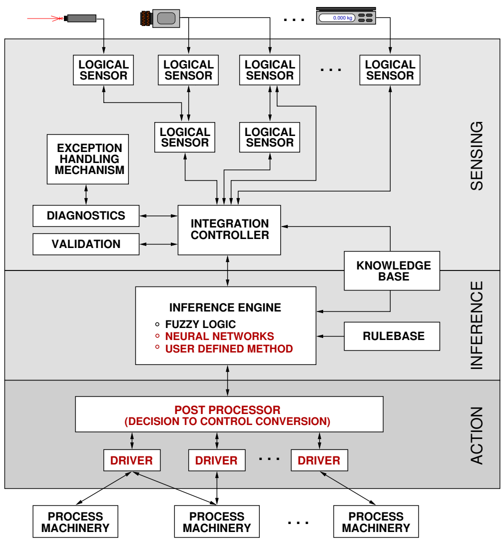

Fig. 1: Extended Logical Sensor Architecture structure.

2.1 Logical Sensors

A logical sensor (LS) is an abstract definition for a sensor. Logical sensors were first defined by Henderson and Shilcrat [11] and later broadened to include a control mechanism by Henderson, Hanson, and Bhanu [12]. This definition provides a uniform framework for multisensor integration by separating physical sensors from their functional use within a system. Logical sensors are used to encapsulate both physical sensors and processing algorithms. This encapsulation defines a common interface for all sensor types allowing the straightforward addition, removal, and replacement of sensors within the architecture.

The logical sensor model has been extended for a model-driven open architecture. As shown in Fig. 2, the proposed Extended Logical Sensor (ELS) is comprised of a number of different components. The components are object-oriented by design; each is responsible for a single task within the sensor. The ELS strongly encapsulates the internal workings of each logical sensor while allowing the modification of the sensor's operating characteristics.

The logical sensor hierarchy structures data in a bottom-up manner. The raw data collected by the physical sensors

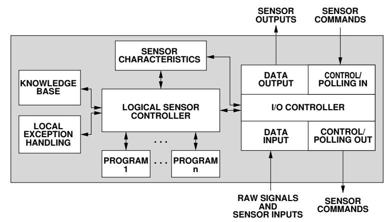

Fig. 2: Basic Extended Logical Sensor components.

is processed through different levels of logical sensors to produce high-level representations of sensed objects and features. To higher-level sensors, each antecedent logical sensor appears as a single entity with a single output, regardless of the scope of *its* antecedents. This approach offers considerable flexibility. High-level tasks may be implemented without regard to the specific sensing devices.

The implementation of an ELS requires an understanding of signal processing. This is knowledge that most industrial users will not possess. They will understand *what* they would like the ELS to do, but not necessarily *how* to accomplish it. This limitation is overcome to some degree by the provision of an ELS library containing logical sensors for common signal processing operations. However, for these users, it will be necessary to have others implement specialized ELS units for inclusion in the library.

3. OBJECT MODELLING

Given a data set provided by a set of sensors to perform some level of object recognition, an object model is used to classify the data. The model provides a generalized description of each object to be recognized. For industrial grading applications, the object model must represent the important features which designate the 'grade' or value of an object. Ideally, the model is simple to construct.

ELSA utilizes feature-based recognition, using distinguishing features to recognize and differentiate objects without discriminating every object feature. This uses the theory of recognition-by-components (RBC) proposed by Biederman [13] which suggests that objects are recognized, and may therefore be represented, by a small number of simple components and the relations between them. A brief review of ELSA's object model [14] is provided herein.

Objects are represented by a connected graph structure similar to that proposed by Tomita and Tsuji [15] for the recognition of objects from texture features, Fig. 7. The object model is a top-down representation of an object. Object nodes represent salient features of an object. The object itself is represented by the classification layer, which is the highest level of abstraction. An object whose relevant features are invariant or which does not require classification may have a single node in the classification layer.

Below the classification layer lie nodes representing the primary features upon which classifications are made. ELSA utilizes fuzzy links which relate primary object features to produce classifications. Here, fuzzy logic has been chosen as a method to incorporate human expertise [16]. Primary features are composed of a number of subfeatures.

The hierarchical structure minimizes the disturbance to the model should a feature used for classification require modification. This structure provides a representation for real-world objects that allows for the quantification of deviations from an ideal model. The following section demonstrates how this structure guides the construction of a multisensor system.

4. CONSTRUCTION METHODOLOGY

To maximize system robustness and usability, the construction of an industrial sensing and processing system using ELSA follows a set procedure. An overview of this methodology is presented in Fig. 3. The following subsections detail the phases of the process. The methodology will be illustrated by an example in Section 5.

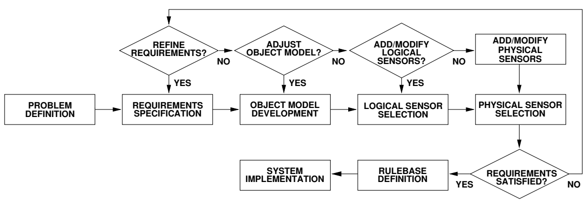

Fig. 3: Overview of construction methodology.

4.1 Problem Definition/Requirements Specification

The first phase of the design process involves the recognition of the needs of the particular industry or process. From these needs, a clear statement of the problem to be solved may be formulated. This problem definition is more specific than the general needs; it must include all of the specifications for what is to be designed. The designer must consider what the capabilities of the system should be. Following the general principles for system design outlined in [17], a set of minimum functional requirements is specified. The articulation of these requirements is used as a guide for subsequent phases. If any of the requirements are left unsatisfied, the design is inadequate. The requirements also serve to keep the design focused.

4.2 Object Model Development

Object model development for ELSA is a two-stage process. First, based upon the requirements of the system from the previous phase, the primary features or characteristics upon which classifications are to be made are identified. As discussed in Section 3, it is advantageous to keep the size of this set to a minimum. Typically, the features in this set are at a high level of abstraction. They occupy the top of the feature layer of the model (Fig. 7). Following the methodology outlined in Fig. 4, each primary feature is decomposed into a set of subfeatures.

Object classifications are defined by combining primary features with fuzzy links. By adjusting the combination of features and their corresponding fuzzy descriptors, object classifications may be distinguished. The object model may be refined by adjusting the classifications and/or the primary (and subordinate) features. The classification layer of the object model (relevant features in combination with relative weights) serves as a template for the inference engine which, in practice, makes classification decisions based on the information extracted by the logical sensors.

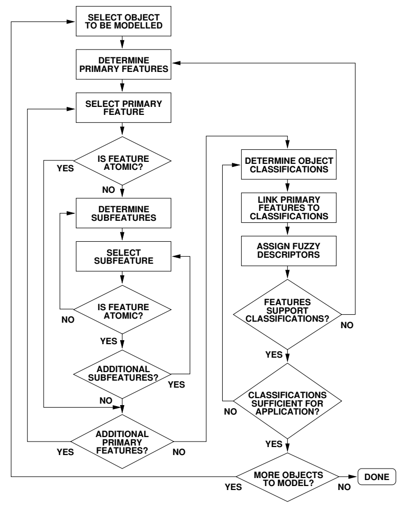

Fig. 4: Object model development methodology.

4.3 Logical/Physical Sensor Selection

The selection of logical sensors is driven by the primary, intermediate, and atomic features that have been identified as necessary for the object model. Sensor selection starts with the primary features. Each feature has a corresponding ELS which packages the information from lower-level sensors (logical or physical) into the representations used for object classification. Many of the low-level logical sensors are selected from a reusable ELS library. The logical sensors contained within the library perform standard image and signal processing operations. The algorithm for constructing the ELS hierarchy is given in Fig. 5.

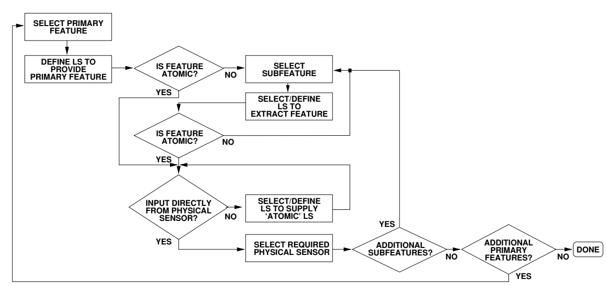

Fig. 5: LS hierarchy development methodology.

The selection of physical sensors requires the consideration of both the input requirements and the capabilities of available transducers. A feature that is beyond the range or capabilities of a single sensor may be accommodated by a LS which fuses the data from multiple sensors that cover the feature space. A single sensor may also be used for multiple tasks; separate cameras are not necessarily required to extract size, colour, and shape information.

4.4 Rulebase Definition

The rulebase defines both rules for object classification and rules to infer the appropriate system output from these classifications. It is generated directly from the object classifications contained in the object model.

The classification rules are based on the fuzzy descriptions of each classification. Each rule expresses a degree of certainty in the classification of the object based on the detection of the primary features. The rules for each classification are combined using the compositional rule of inference to determine to what degree the object belongs in each classification.

Action rules are defined using the certainty of each object classification as the antecedent(s); the appropriate action(s) forms the consequent. If an object classification is certain, the appropriate action is straightforward.

The membership functions required for the definition of these rules are automatically generated so that users are not required to have an in-depth understanding of fuzzy logic. Classification of features which are not easily quantified use the membership functions *low*, *average*, and *high* to express the confidence in the detection of the feature. These functions span the range 0 to 1; *average* is centred over 0.5. This is intended to provide users with an intuitive feel for the specification of classifications. The user does not consider values or fuzzy membership, but rather the linguistic variables *low*, *average*, and *high*.

For features that are easily quantified, such as length and mass, the system prompts the user to supply information about the universe of discourse (range of expected values), and linguistic variables for the classifications over this universe. Triangular membership functions centred at the mean values of each variable are used, since with sufficient representation the membership function shape is not critical [18]. Expert users may by-pass the system, allowing direct definition and tuning of the membership functions.

Inferring action from the object classifications uses a similar methodology. The membership functions *uncertain*, *unsure*, and *certain* are used to span the universe of certainty. Again these functions are defined over the range 0 to 1, with *unsure* centred over 0.5.

5. APPLICATION EXAMPLE

5.1 Background

Herring roe is an important part of the B.C. economy, with an annual value of $200 million dollars. A herring roe skein is a sac of tiny herring eggs. Being a natural product, it exhibits many non-uniform characteristics. Roe is a particularly challenging product due the large number of classifications. Each classification is dependent on the presence or absence of a number of features. Weight, appearance, and texture of the salted herring roe are the primary factors influencing price. Proper classification allows processors to offer improved value to their customers.

The roe grades are subject to change each season, due to the customer driven nature of the industry. Currently, there is no standardization of the various grade specifications. Distortions of the roe are commonly described using linguistic terms — the interpretation of which varies among expert graders. This inconsistency makes the quantification of product quality difficult. Examples of good (Grade 1) and distorted (Grade 2-H) roe are presented in Fig. 6.

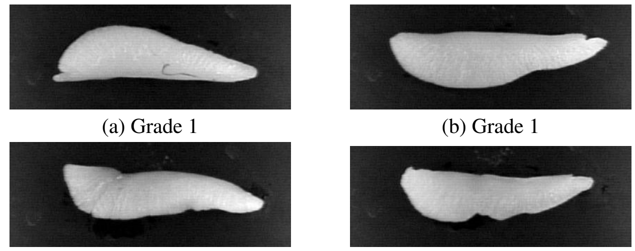

(a) Grade 1 (b) Grade 1

(c) Grade 2-H (d) Grade 2-H

Fig. 6: Examples of herring roe classification grades.

5.2 Problem Definition/Requirements Specification

The shape of the roe has been the most difficult to access using machine vision [19]. Human graders look for roe to be "well formed," or free from defects. The current industrial prototype [7] is limited to the use of a single image without intensity information, making it difficult to consistently classify roe with various defects. The original system was designed as a two-classifier — separating good roe from bad — and while Grade 1 roe are identified consistently, attempts to subclassify the defective roe have met with limited success. To address these limitations, additional information is required. Three dimensional and texture information would provide the system with features that are essential to proper classification. The following example demonstrates how ELSA can be used to implement such a multisensor system. Development of this prototype is continuing.

The requirements of the multisensor system for the grading of herring roe skeins are as follows:

1. Accurate determination of skein length (± 1 mm).
2. Estimation of skein weight (5–50 g $\pm$ 0.5 g).
3. Estimation of roe thickness as ratio of width.
4. Detection of parasite bites.
5. Detection of broken skeins (broken head or tail).
6. Detection of depressions (> 12 mm^2).
7. Detection of twists (> 12 mm^2).
8. Detection of proper yellow colour.
9. Estimation of roe firmness as a measure of maturity.
10. Detection of bumps and 'cauliflower' deformations.
11. Detection of cracks in the roe skein (> 1 mm).

5.3 Object Model Development

From the requirements, there are eleven salient features that can be identified as necessary for classification. These include weight, length, thickness, firmness, presence/absence of parasite bite(s), breaks, cracks, twists, depressions, cauliflower, and proper colour. Many of these can be assessed on the basis of 2D (overhead) visual information.

Thickness, twists, and depressions require information about the three dimensional profile of the roe skein. The 3D profile is usually represented as a surface map. Due to the variability of roe, thickness is represented as a ratio between the depth of the roe (as estimated by the 3D profile) and the width of the roe at the minor axis.

Weight is not measured directly, but is estimated using a linear regression model based on the peripheral length, area, and thickness of the roe [7]. Firmness is assessed using ultrasonic echo imaging. The strength of the echo signal is directly dependent on the structure, uniformity, and firmness of the object region which generates the echo. Therefore, the echo image contains features correlated with the firmness of the roe. These features may be extracted and used as an indirect measure.

None of the primary features are atomic. Each is broken down into the various atomic components which permit the detection of the feature. The primary features and corresponding subfeatures for Grade 2-H roe are shown in Fig. 7. The complete object model details the other object classifications similarly.

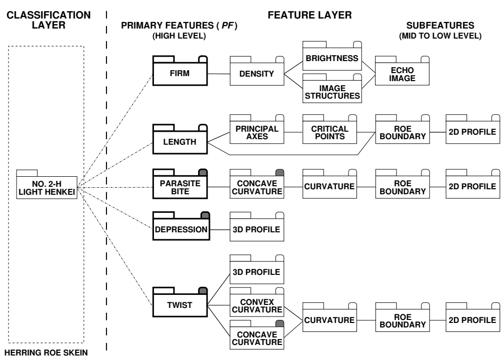

Fig. 7: Object model for Grade 2-H roe.

5.4 Logical/Physical Sensor Selection

From the object model, a logical sensor hierarchy is constructed, Fig. 8. Each of the features identified during the development of the object model is associated with an ELS which can extract the required feature. Many of the primary features share common subfeatures, simplifying the logical sensor hierarchy. There are three physical sensors required: two CCD cameras and an ultrasonic probe; the details of each follow:

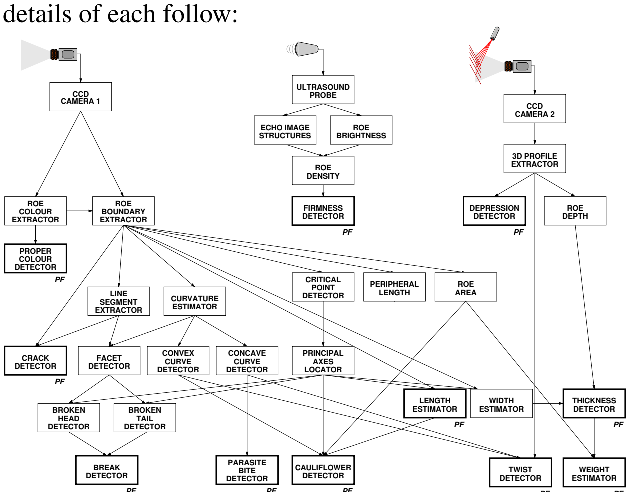

Fig. 8: Sensor hierarchy for herring roe grading.

The low-level subfeatures RGB Image and 2D Profile, from which a number of other subfeatures are derived, may be provided by a single colour camera with evenly distributed, diffuse lighting.

The Echo Image subfeature is provided by a 10 MHz mechanical sector ultrasound probe. This provides input to the ELSs responsible for estimating the firmness of a roe skein. Upon extracting the relevant image features, these are passed to the Firm ELS which uses fuzzy logic to estimate firmness.

The 3D Profile Extractor ELS utilizes a second CCD camera in combination with a structured laser light. Using knowledge of the image geometry, the 3D profile of the roe skein may be reconstructed.

5.5 Rulebase Definition

The rulebase generation follows from the object model. The object classifications outlined in 5.3 are used as the basis for the classification rules, Fig. 9. For clarity, only rules related to the classification of Grade 2-H roe are shown.

IF Firm IS *high* AND Length IS *normal* AND ParasiteBite IS *high* THEN No2-H = *high*
IF Firm IS *average* AND Length IS *normal* AND ParasiteBite IS *high* THEN No2-H = *high*
IF Firm IS *high* AND Length IS *normal* AND ParasiteBite IS *average* THEN No2-H = *low*
IF Firm IS *high* AND Length IS *normal* AND Depression IS *high* THEN No2-H = *high*
IF Firm IS *average* AND Length IS *normal* AND Depression IS *high* THEN No2-H = *high*
IF Firm IS *high* AND Length IS *normal* AND Depression IS *average* THEN No2-H = *low*
IF Firm IS *high* AND Length IS *normal* AND Twist IS *high* THEN No2-H = *high*
IF Firm IS *average* AND Length IS *normal* AND Twist IS *high* THEN No2-H = *high*
IF Firm IS *high* AND Length IS *normal* AND Twist IS *medium* THEN No2-H = *high*
IF Firm IS *high* AND Length IS *normal* AND Twist IS *low* THEN No2-H = *no*
IF ParasiteBite IS *low* AND Depression IS *low* AND Twist IS *low* THEN No2-H = *no*

Fig. 9: Rules used to identify Grade 2-H herring roe from primary features.

Once the roe have been classified, a decision is made about which of six bins it should be ejected into; unclassified roe are allowed to fall off the end of the conveyor into a seventh bin. Fig. 10 presents the rules which are used to infer this decision. The fuzzy membership functions associated with these rules are shown in Fig. 11.

IF Grade1-3L IS *high* THEN Decision = *bin1*
IF Grade1-2L IS *high* THEN Decision = *bin1*
IF Grade1-Large IS *high* THEN Decision = *bin1*
IF Grade1-Medium IS *high* THEN Decision = *bin2*
IF Grade1-Small IS *high* THEN Decision = *bin3*
IF Grade1-Pencil IS *high* THEN Decision = *bin7*
IF Grade2 IS *high* THEN Decision = *bin4*
IF Grade2 IS *low* THEN Decision = *bin4*
IF Grade2-C IS *high* THEN Decision = *bin5*
IF Grade2-C IS *low* THEN Decision = *bin5*
IF Grade2-H IS *high* THEN Decision = *bin6*
IF Grade2-H IS *low* THEN Decision = *bin6*
IF Unclassified IS *high* THEN Decision = *bin7*
IF Unclassified IS *low* THEN Decision = *bin7*

Fig. 10: Rules used to determine decisions about how roe should be handled based on object classifications.

5.6 Discussion

In this simple application, the advantages of the ELSA approach to system design are demonstrated. By formalizing the design process, a system can be designed to meet the specified functional requirements in a systematic and comprehensible way. Each stage involves the extraction and utilization of the user's expert knowledge about the process and desired outcomes. Specification of the requirements leads to the identification of primary features and object classifications. These are expanded into subfeatures. The features themselves determine the algorithms (encapsulated by logical sensors) and physical sensors that are required by the system. Decisions are inferred directly from the object classifications.

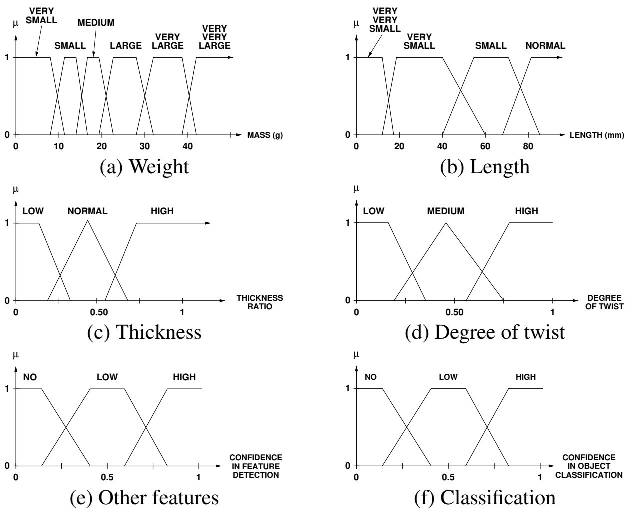

(a) Weight (b) Length (c) Thickness (d) Degree of twist (e) Other features (f) Classification

Fig. 11: Membership functions used for classification of herring roe grades.

The specification process effectively separates the expert domain knowledge from the technical programming knowledge required to develop an ELS. Should a required sensor be unavailable from the library, an industrial user may still specify the system while the ELS development is deferred to a technical expert.

The structure of the architecture ensures that should additional capabilities be desired, they may be added without affecting the existing components. The object model is simply expanded to include the additional features and/or classifications. Any required logical sensors are added to the ELS hierarchy. This is accomplished without disturbing the remainder of the system.

6. CONCLUSIONS AND FUTURE WORK

ELSA was developed to provide an organized approach to the development of industrial-based sensor systems. It addresses the need for scalable, modular, and structured sensor systems, replacing current ad hoc approaches. The construction methodology enables domain experts, who lack signal processing knowledge, to design and understand a sensor system for their particular application.

The object model used by ELSA is particularly suited to the representation of non-uniform products, or any object for which classification is desired. Logical sensors are chosen to provide each of the features defined by the object model; this in turn determines what physical sensors are required by the system. The classification layer of the object model directly specifies how primary features are combined to determine object classifications.

Future work will include the development of a library of Extended Logical Sensors that is suitable for a variety of inspection and grading tasks. This will assist in the development of ELSA-based systems for applications such as the grading of potatoes, blueberries, and other produce. Other applications include assembly, material handling, and machining operations.

ACKNOWLEDGEMENTS

The financial support of the Natural Sciences and Engineering Research Council of Canada, the Gordon M. MacNabb Scholarship Foundation, and the Garfield Weston Foundation is gratefully acknowledged.

REFERENCES

[1] R. C. Luo and M. G. Kay, "Data fusion and sensor integration: State-of-the-art 1990s," in *Data Fusion in Robotics and Machine Intelligence* (M. A. Abidi and R. C. Gonzalez, eds.), pp. 7–135, Academic Press, 1992.

[2] R. R. Murphy, "Biological and cognitive foundations of intelligent sensor fusion," *IEEE Trans. on SMC-A*, vol. 26, pp. 42–51, Jan. 1996.

[3] L. A. Klein, *Sensor and Data Fusion Concepts and Applications*, SPIE Optical Engineering Press, 1993.

[4] S. S. Iyengar, L. Prasad, and M. Min, *Advances in Distributed Sensor Technology*. Prentice Hall PTR, 1995.

[5] W. Daley, R. Carey, and C. Thompson, "Poultry grading/inspection using color imaging," in *Proc. SPIE – Machine Vision Applications in Industrial Inspection*, vol. 1907, pp. 124–132, 1993.

[6] J. Calpe, F. Pla, J. Monfort, P. Díaz, and J. C. Boada, "Robust low-cost vision system for fruit grading," in *Proc. 8th Mediter. Electrotech. Conf.*, vol. 3, pp. 1710–1713, 1996.

[7] E. A. Croft, C. W. de Silva, and S. Kurnianto, "Sensor technology integration in an intelligent machine for herring roe grading," *IEEE/ASME Trans. on Mechatronics*, vol. 1, pp. 204–215, Sept. 1996.

[8] G. Brown, P. Forte, R. Malyan, and P. Barnwell, "Object oriented recognition for automatic inspection," in *Proc. SPIE – Machine Vision Applications in Industrial Inspection II*, vol. 2183, pp. 68–80, 1994.

[9] Dipix Technologies Inc., "QualiVision product information," 1998. `http://www.dipix.com`.

[10] M. D. Naish and E. A. Croft, "An open architecture for intelligent multisensor integration in industrial applications," in *Proc. SPIE – Architectures, Networks, and Intelligent Systems for Manuf. Integr.*, vol. 3203, pp. 33–41, 1997.

[11] T. C. Henderson and E. Shilcrat, "Logical sensor systems," *J. of Robotic Systems*, vol. 1, no. 2, pp. 169–193, 1984.

[12] T. C. Henderson, C. Hansen, and B. Bhanu, "The specification of distributed sensing and control," *J. of Robotic Systems*, vol. 2, no. 4, pp. 387–396, 1985.

[13] I. Biederman, "Recognition-by-components: A theory of human image understanding," *Psychological Review*, vol. 94, no. 2, pp. 115–147, 1987.

[14] M. D. Naish and E. A. Croft, "Data representation and organization for an industrial multisensor integration architecture," in *Proc. IEEE SMC*, vol. 1, pp. 821–826, 1997.

[15] F. Tomita and S. Tsuji, *Computer Analysis of Visual Textures*. Kluwer Academic Publishers, 1990.

[16] L. A. Zadeh, "Fuzzy logic = computing with words," *IEEE Trans. on Fuzzy Systems*, vol. 4, pp. 103–111, May 1996.

[17] J. E. Shigley and C. R. Mischke, *Mechanical Engineering Design*. McGraw-Hill, 5th ed., 1989.

[18] C. W. de Silva, *Intelligent Control: Fuzzy Logic Applications*. CRC Press, 1995.

[19] L. X. Cao, C. W. de Silva, and R. G. Gosine, "A knowledge-based fuzzy classification system for herring roe grading," in *Proc. ASME Winter Annual Meeting on Intelligent Control Systems*, vol. DSC-48, pp. 47–56, 1993.