Book Review

Jonathan Legg

Context-Enhanced Information Fusion: Boosting Real-World Performance with Domain Knowledge

Lauro Snidaro, Jesus Garcia, James Llinas, and Erik Blasch, editors

Springer, 2016
DOI 10.1007/978-3-319-28971-7

This is an exciting book. It is large (703 pages and 25 chapters), broad in scope, and covers a lot of material to a useful depth. It is also well organized and describes the state-of-the-art in methods for incorporating background knowledge into information fusion problems.

Information fusion is the process of taking multiple pieces of information, from single or multiple sources, and combining them to estimate the underlying processes of interest. These could be the trajectories of vehicles seen by radars, or a developing situation involving ground forces, for example. Estimation problems such as these combine information from sensors with a priori background information, such as a model of how vehicles move. But this motion model could be usefully constrained through knowledge of the likely types of vehicles being tracked, potentially leading to greater tracking accuracy.

This book discusses the use of context in information fusion across both the Joint Directors of Laboratories (JDL) spectrum of information fusion (for example, [1]) and across hard and soft data fusion (that is, fusing "physics-based" data such as from radars with "human-sourced" data such as from Twitter). While it provides several explanations of "context" in information fusion, the term essentially refers to the use of "structural knowledge of the scene, known a priori relationships between the entities and the surrounding environment, dynamic scenarios necessary to interpret or constrain the system output, and user preferences, social norms and cultures when estimating the situations of interest for the domain." In slightly more depth, there are discussions of "context of" ("background context, which provides a more general and stable environment") and "context for" ("secondary characteristics, which can be more dynamic") (chapter 2). Other examples of context include road networks and driver behavior, which represent information that should have an effect on the output of the information fusion system, but would not normally be explicitly factored into a textbook Kalman filtering problem.

Intuitively, contextual information should be particularly valuable for improving the performance of a fusion system when you only have access to a finite volume of sensor data, because it can be thought of as another source of measurements. In the extreme, it may also help solve problems that otherwise could not be solved due to a lack of information. "Challenging scenarios where context can play a role include: object labelling, track correlations/stitching through dropouts, and activity recognition" (chapter 18, p. 479).

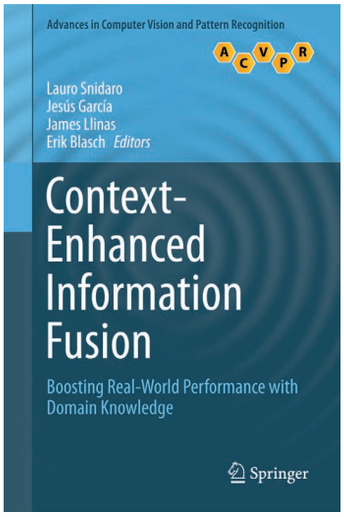

I find there are two dangers with books where each chapter is written largely by different authors: that each chapter introduces the topic in its own way, and that it is unusual for many chapters to go beyond hand-waving, providing enough information for the enthusiastic reader to apply the material. This book does a good job of defining the problem and maintaining consistency across the various perspectives (although one feels that the definition of "context" depends on the context). Some chapters provide an overview of some algorithms, but this is not a "how-to" book. It does, however, provide a useful catalogue of references to the literature.

An outline of the book follows, highlighting some areas that I find particularly interesting:

- **Foundations** (chapter 1), where the editors attempt to impose order through definitions and background in a book with many disparate contributors and applications;
- **Concepts of context for fusion** (chapters 2–6), with background material that includes patterns of life and anomaly detection as context (noting that this topic is also mentioned elsewhere in the book), and chapter 3 discussing uncertainty in contextual information by treating it as a source of uncertain data;
- **Systems philosophy of contextual fusion** (chapters 7–10), with fusion system architectures, in particular middleware to handle context-related information exchanges (chapters 8, 9);
- **Mathematical characterisation of context** (chapters 11–15), with context incorporated into low-level tracking, such as the incorporation of knowledge of road networks (section 12.4.3, with ground target tracking also discussed in chapters 4 and 22); and learning approaches that estimate the context from the data (chapter 15). A variety of approaches for handling context are discussed in these chapters.
- **Context in hard/soft fusion** (chapters 16–19) additionally discusses multilevel fusion (in a JDL sense), with

context as a binding element (chapter 16), and chapter 17 suggests Battle Management Language to communicate both data and contextual information between data fusion levels, with some discussion of uncertainty representations;

- **Applications of context approaches to fusion** (chapters 20–24) discusses a logic-based knowledge base to assist with video tracking (chapter 23); and
- **Context in robotics and information fusion** (chapter 25) appears to be a very good summary of sensor fusion in the world of robotics, with some discussion of context representation approaches including logic-based and probabilistic approaches (section 25.2.2).

I also have an interest in a context-enhanced reasoning system that operates through interactions between a Bayesian process and a knowledge base, or vice versa. (This is the motivation for the "middleware" referred to in chapters 8 and 9.) Such ideas lead to the question of how uncertainty could be incorporated into knowledge bases, and this is a topic that I feel could have been expanded on further. The book mentions Markov logic networks (e.g., section 4.2 on contextual tracking, and 16.4 on design directions for context-aware systems) and Bayesian networks (e.g., section 3.4 regarding ontologies, section 25.2.2.3 on probability-based representations), with details left to the references, which are generally in the fusion literature rather than comprehensive texts (e.g., [2], [3]). The broader area of probabilistic programming has been undergoing dramatic changes in recent years (e.g., [4]), so its absence is not a major concern.

If you would like to understand contextual information and how your data fusion system might benefit from it, this book should make very interesting reading.

REFERENCES

1. Steinberg, A. N., Bowman, C. L., and White, F. E. Revisions to the JDL data fusion model. In *Proceedings of the SPIE 3719, Sensor Fusion: Architectures, Algorithms and Applications III*, 12 Mar. 1999. DOI: 10.1117/12.341367.
2. Domingos, P., and Lowd, D. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan & Claypool, 2009.
3. Getoor, L., and Taskar, B. *Introduction to Statistical Relational Learning*. Cambridge, MA: MIT Press, 2007.
4. Wikipedia. Probabilistic programming language, https://en.wikipedia.org/w/index.php?title=Probabilistic_programming_language&oldid=840447199, last access May 9, 2018.