

Sensor Modeling and Simulation: Can it pass the Turing Test?

Mel Siegel
Robotics Institute

Carnegie Mellon University
Pittsburgh, PA 15213-2605, USA

Phone: +1 412 268 88025, Fax: +1 412 268 5569

Email: mws@cmu.edu, URL: <http://www.cs.cmu.edu/~mws>

Abstract – *The VIMS workshop encompasses both virtual instruments and “virtual sensors”, but historically the latter are hardly represented in the program. This paper is not so much a research report as a friendly challenge to the VIMS community to begin thinking actively and creatively about this complementary aspect of our theme. In particular, it raises the question of what and how much we need to build into a virtual sensor agent for it to pass something analogous to the Turing Test, i.e., for a virtual sensor, as a “black-box”, to be indistinguishable from a real sensor.*

Keywords – *sensor modeling, sensor simulation, virtual sensor, wished-for sensor, Turing Test*

I. INTRODUCTION

The basic premise of this paper is that – in contrast to artificial intelligence, which cannot yet create systems that pass the Turing Test [1] – artificial sensing, i.e., sensor agents built to “stand in” for missing or unavailable real sensors, probably can already be built with sufficient realism that they could pass an analogous test. Its goal is to challenge the VIMS community to flesh out this hypothesis and define an agenda that will lead eventually to agreement on how we will know when we have succeeded.

VIMS workshop papers typically address issues relating to real sensors that are monitored, processed, interpreted, and displayed by virtual instruments, i.e., by software systems and *tabula rasa* displays that simulate in software and graphics all but the lowest-level portions of the hardware that is normally found in traditional signal conditioning, processing, fusing, and reporting instruments. The complementary set of topics relating to “virtual sensors” is also within the scope of the VIMS workshop theme, but with only a few exceptions it has been absent from past programs [2][3][4]. Yet with increasing capability of virtual instruments, it becomes increasingly important to be able to do a convincing job of sensor simulation, e.g., to satisfy the need, temporarily or during some interim period, to substitute software agents for temporarily unavailable or not yet developed (“wished-for”) sensor hardware. This will be especially necessary during periods when large systems and the new sensors their functionality demands are developed in parallel. Such large systems are now under intense discussion and in some cases prototype development for diverse applications such as wide area environmental monitoring, detection of clandestine attack with

biological agents, monitoring remote areas for illegal or hostile activity, etc.

This speculative paper reports no experimental research results. It is rather the author’s attempt to define the problem and to invite and challenge the VIMS workshop community to start thinking about, arguing about, and working intensively on this heretofore neglected component of the Virtual and Intelligent Measurement Systems theme. In particular, we work toward answering the question of whether a sort of “Turing Test” is appropriate for evaluating the effectiveness of sensor simulation agents, i.e., “virtual sensors”.

Section II elaborates briefly on the background issues. Section III summarizes briefly previous work on virtual sensor topics. Section IV raises the question “is there a Turing Test for virtual sensors?”. Section V discusses generalized classes of data sources and data sinks, e.g., sensors, databases, and reporters. Section VI discusses the practical issues involved in synthesizing an interim working model. Finally, Section VII briefly recites some conclusions and opportunities.

II. BACKGROUND

Sensor modeling and simulation is required to support development and verification of, for example, large automated systems for numerous continuous monitoring tasks, e.g., those already mentioned above, and other such as small-scale weather extremes, environmental resource monitoring and protection, large-scale fundamental research such as measurement of continental drift, detection of gravitational radiation, and many others. These systems typically involve three main components:

- One or more real-time databases;
- One or more live sensor systems; and
- A system controller with appropriate human interface.

Sensor modeling and simulation is required in this context because large developmental systems must be validated, and cutting-edge deployed systems must be challenged, by some or all of:

- A number of sensors substantially larger than is currently actually available;

- Sensor readings that are outside the range that is currently available;
- Sensors that are too valuable to divert, e.g., large diagnostic or imaging instruments; and
- Sensors that are still “wished for”, i.e., they are plausible but not yet actually available, thus they should be provided for in advance in the overall system plan.

We anticipate work toward simulating these “virtual sensors” with software agents that implement physically correct models of:

- The transduction elements *per se*;
- Low level signal conditioning;
- Analog-to-digital conversion;
- Linearization, calibration, and digital signal processing;
- Sensor fusion and signal-to-symbol abstraction; and
- Network communication management.

While there are a few extant examples, summarized in Section **Error! Reference source not found.**, in which “virtual sensors” are discussed [2]-[10], the virtual sensors described in these references are either agents that reformat (perhaps taking into account some noise and error issues) data from real sensors, or they are *ad hoc* efforts that were handcrafted to meet the momentary needs of a specific narrow application.

In contrast, I would like to challenge the VIMS community to articulate a higher-level framework that encompasses all of these, and that provides machinery for systematically addressing more-or-less arbitrary applications. This leads me to ask and examine exactly what must be the characteristics of a virtual sensor agent for it to be able to pass something analogous to the Turing Test, i.e., for it as a “black-box” to be indistinguishable from the real sensor that the agent emulates.

III. PREVIOUS WORK

In work reported around 1990, Muir [5] at Sandia (previously and later at Carnegie Mellon University) used the term “virtual sensor” in a way that is quite different from the sort of sensor-simulating agent that is suggested herein. Muir used the term to describe an intermediate stage of data fusion and analysis that could usefully be regarded as a stand-in for a hypothetical sensor. In [5] Muir productively applied this approach, theoretically and experimentally, to the problem of robot kinematic identification. In this article he describes an approach that combines and processes, with appropriate consideration of measurement errors and noise, the output of several real but indirect sensors (or multiple separate measurements reported by a single indirect sensor when the robot is in different kinematic configurations) to synthesize the

would-be outputs, errors, and noise characteristics of another set of unavailable but conceptually desirable direct sensors.

In previous editions of the VIMS workshop series (including its predecessors with different names in the same evolutionary lineage) and in some proceedings of the parent IMTC conferences we find a few examples of the term “virtual sensor” used in more-or-less the same way intended herein, although still more oriented toward abstraction of real sensor outputs into the hypothetical output of an imagined sensor than toward complete bottom-up simulation of an unavailable sensor. In particular, this approach is seen in several reports by Groen et al of the Universiteit van Amsterdam Instituut voor Informatica [2], wherein they addressed several quite different application areas, for example, an automatic debiting system [3], and a model-based robot vision [4]. It is especially useful to quote from [3], where (at the beginning of Section II, The Virtual Sensor Model) the virtual sensor is carefully defined by the following paragraph:

In any goal-directed sensory system, a *virtual sensor* is a (conceptual) device whose output can be modeled in terms of the relevant characterizing parameters, and the outputs of other virtual sensors. The virtual sensor modules should be chosen at the highest level of abstraction that enables a sufficiently accurate characterization of the total system behavior, but at which the interactions between various virtual sensor modules are (relatively) simple, both in their statistical (in)dependence and in their causal relationships. In a simulated system, we have the additional demand that the virtual sensor models should be amenable to being *validated*.

Atkinson et al at the University of West Virginia report a real-time automobile emission monitoring system employing a neural-network based virtual sensor [6], and Oosterom et al [7] at the Technical University of Delft report a virtual sensor for aircraft acceleration used in a fault detection and identification system using fuzzy logic methods. The latter is particularly interesting in that it the virtual sensor can be switched into the monitoring system in place of one of several redundant sensors in a voting arrangement. The approach aids identification of faulty real sensors.

Dixon et al at Carnegie Mellon University report a virtual environment for modeling multiple mobile robot systems that includes the capability to combine real and virtual sensors [8]. The term “virtual sensors” is actually used in two ways, first, to include in the system model an agent that faithfully represents a real sensor in order to perform a cost-benefit analysis before actually purchasing and installing it, and second, to include in the system model an agent that represents a “wished-for” sensor. This inclusive definition is quite consistent with the one applied herein.

Haley et al at Purdue University report a “virtual sensor” for the organism *Listeria monocytogenes* that contaminates proc-

essed meat products [9]. Based on a model that describes the thermal death rate kinetics for the organism on specific type of processed meat surfaces and a numerical heat transfer model for pasteurizing the surfaces, the model returns a prediction of the amount of residual organism as if a specific *Listeria monocytogenes* were available.

Kestell et al, University of Adelaide (Australia) describe active noise control under that employs virtual sensors [10]. Summarizing the pros and cons of the approach, authors report "in general the 'virtual energy density sensor' outperforms the actual energy density sensor, the actual microphone and the virtual microphone in terms of centering a practically sized zone of local control around an observer who is remotely located from any physical sensors; the virtual sensor algorithms however, are shown to be sensitive (by varying degrees) to short wavelength spatial pressure variations of the primary and secondary sound fields".

The previous paragraphs are far from a comprehensive review of the field to date, but they do provide a fair sampling of the range of approaches and applications that have been considered.

In ending this section, it may provide us with some perspective to note, only slightly tongue-in-cheek, that if the VIMS community had as much chutzpah as many of our academic brethren, we could at this point lay claim to the field of computer graphics as a subset of our general sensor modeling and simulation endeavor. After all, what is computer graphics but the calculation of the expected signal given a geometrical world model probed by a given illumination model and observed by one or more virtual cameras with specified world coordinates? Much of the prior work mentioned above is quite analogous to recently highly publicized achievements in "virtualized reality": real sensors (cameras) capture data (images from multiple real perspectives), the data are compiled and compressed via construction of a world model, and arbitrary new perspective views are then rendered according to the above defining prescription for doing computer graphics.

IV. A TURING TEST FOR VIRTUAL SENSORS?

By now all modern real and virtual sensors can, of course, be presumed to be digital, networked with a system controller and one or more databases, and capable of two-way communication, i.e., all modern sensors can be assumed to be both "talkers" and "listeners". Simulation success can thus be measured – presumably by a kind of Turing Test – by examining the data emanating from a sensor node, analyzing the node's responses to queries, and deciding from the character of this communication whether the sensor is real or virtual. The procedure is analogous to the famous Turing Test in artificial intelligence, wherein it is generally agreed that a computer can reasonably be called intelligent if a human interact-

ing with it, e.g., via a terminal, cannot decide whether he or she is communicating with a computer or a human [1].

A *valid* virtual sensor is thus an agent that we cannot distinguish by communication alone from a real sensor. It is important to say "by communication alone" because we can easily imagine circumstances under which we might surreptitiously change some actual environmental parameter, e.g., the temperature, in a region from which a virtual sensor purports to be transmitting, and the virtual sensor would, of course, fail to respond appropriately. This sort of physical test will of course always unmask a pretending agent, both in the classical artificial intelligence context and in the present virtual sensor context.

Where do we imagine a virtual sensor trying to conceal its unrealness might be less than convincing? It seems most likely that it will be at the lowest levels of the transduction - conditioning - conversion - processing - fusion - communication pipeline, i.e., at the level where there is real hard-to-simulate physics at work; at the higher levels there is only software that processes simulator-generated data no differently from the way it processes transducer-generated data.

How do we imagine a virtual sensor trying to conceal its unrealness might fail? It seems most likely that the giveaway will be that *the simulator-generated data will be too good*, i.e., the noise model will be recognizable as overly predictable. It is again useful to reason by analogy with the well know computer graphics field. How do we discern that a "photo realistic" image rendered by a modern high quality computer graphics engine is in fact a fake? Only a few years ago the answer would have been that it is too simple: polygon-approximating surfaces too big and too few, lighting too simple, shadow-colors wrong, inter-reflection ignored, background and foreground in unrealistically good focus, etc. But today the answer is quite the opposite: we recognize the fake because it is too good! There is no dirt, there are no scratches, no nicks, no scuffmarks, no fingerprints; in the time domain, realistic frame-to-frame fluctuation due to the thermal and quantum noise characteristics of real image recording sensors is missing or unconvincing.

Analogously, a plausible starting point for discussion, setting priorities, and designing future system-building experiments holds that inadequate models of noise originating in the virtual transducer per se, in the virtual environment as detected by the virtual transducer, and in the model of the low-level signal processing simulation will be the key weakness that enables recognizing virtual vs. real sensors.

V. SENSORS, DATABASES, AND REPORTERS

Perhaps the simplest definition of a *sensor* is "a source of new data". In contrast, a *database* is a source of *historical*

data and also a sink for both historical and new data, i.e., data received in the first instance from other databases and in the second instance from sensors. These broad definitions of database and sensor keep open the potentially useful future possibility that we might want to regard, within the general model, a person (a “*reporter*”) as either a sensor or as a database in an integrated automated observation, analytical thinking, and action-taking monitoring and control system. Note that here “automated” doesn’t mean “no people”, but rather that in the cooperative functioning of man-machine systems with people continuing to make the most critical high level decisions, but perhaps serving interim roles, filling in for missing sensors and actuators at lower levels until such time as suitable electromechanical replacements become available.

However, to avoid at this early stage the possibility of getting bogged down in long, pedantic, and probably useless discussions, we limit the present discussion to *digitally communicating electronic sensors* and to *digital databases*.

Thus, as intimated previously, we presume that all *sensors* transmit data to a network as packets that are encapsulated according to an agreed protocol, that each packet is suitably labeled with a legitimate recipient-*database’s* address, and that every database correctly receives suitably encapsulated and addressed packets. In other words, we presume that transparent software utility layers check and enforce the legitimacy of sensor-to-database communications. And, as we also intimated, we also presume that sensors listen as well as talk, e.g., they listen to and obey suitably validated messages from the system controller that configure their operating mode and parameters, initiate autocalibration, request data transmissions, etc.

VI. SYNTHESIZING AN INTERIM WORKING MODEL

A sensible interim working model for initial experiments toward the virtual sensor goal might be the NIST IEEE 1451 standard [12], which provides a uniform bi-directional interface between transducers (sensors and actuators) and networks. The standard specifies a STIM (standard transducer interface module) that provides the sensor support elements bulleted above, and a TEDS (transducer electronic data sheet), a small onboard database that characterizes the individual sensor’s pedigree and personality. With the emergence of sensor communication standards like NIST IEEE 1451 implemented on top of standard network communication protocols like TCP/IP, we can focus our attention on the content of the data packets without being sidetracked by the details of their transport.

The content of sensor data packets spans a potentially enormous dynamic range. To take an example from the domain of a population being monitored at a low level for exposure to clandestine low level biological attack, the lower end of the range might be just the simple output of a wirelessly net-

worked pocket thermometer that sends a monitored individual’s temperature, her ID number (accurately read from her smart wristband), and a timestamp, all the way to the complex multidimensional output of some “wished for” future three-dimensional imaging instrument that instantaneously reports pathogen, drug, and blood constituent identity and concentration voxel-by-voxel, as well as the instrument’s mode and operating parameters, detailed case background information about the patient, and the patient ID and data timestamp.

In the heterogeneous environment in which an effective bio-surveillance system like this would have to operate, it would be essential for sensors to report not just data, but also realistic dynamic estimates of the current validity of the data. This requirement becomes increasingly important as the sensors and the data become more sophisticated, i.e., as the reports being to provide information at increasingly higher levels in the signal-to-symbol pipeline.

Let us enumerate and discuss in a preliminary sort of way the levels of this pipeline via a concrete example. Consider an analytical instrument that delivers a spectrum-like signature of the sample. It makes hardly any difference whether the sample is environmental, e.g., air, water, soil, etc., or clinical, e.g., saliva, urine, blood, etc., nor does it make any fundamental difference whether the dispersion is by optical wavelength, molecular mass, ion mobility, column retention time, etc: at the lowest level, every such instrument delivers simply an electrical signal as a function of a dispersion parameter.

Viewed at the elementary raw signal level, the uncertainty in the signal, for example, expressed in terms of an electrical signal current measurement in an appropriate location, is modeled by the shot noise (proportional to the square root of the product of the signal and its amplifier’s integrating time) plus an estimate of the technical noise due to imperfect engineering in the signal collecting environment. At the level of the elementary spectrum, the signal is typically digitized and binned by an algorithm whose output is a sequence of (peak area, dispersion parameter) pairs. The digitization adds quantization noise, and depending on the nature of the analyte and how well the algorithm anticipates the possibilities, the actual binning may or may not be sensible.

If this algorithm happens to successfully model what the analytical measurement instrument is really doing, then each of the (peak area, dispersion parameter) pairs represents the sum of signature contributions from each constituent of the sample mixture in the integrating window around the reported value of the dispersion parameter. So after low level signal amplification and digitization, a matrix inversion algorithm is typically applied to extract the concentration of each constituent. A good implementation, e.g., singular value decomposition, delivers both the individual concentrations and confidence estimates. The confidence estimates have components related

to the observed fluctuations in the output when the input remains nominally constant and other components relating to best estimates from other sources regarding systematic errors that can, with some uncertainty, be compensated.

This is where, e.g., medical lab reports typically end: the physician is told how many micrograms of molecule X were purportedly found per deciliter of sample of body fluid Y, and she folds this report into her diagnosis and treatment plan. It now seems universally accepted that one more layer of algorithms implementing artificial intelligence technology, e.g., expert systems (or their less dated descendents) will in the not-too-distant future do a better job than the run-of-the-mill physician at this last step. Some practical implementations, notably fuzzy logic, also provide quantitative estimates of their confidence in themselves. These measures of self-confidence at each level of abstraction are all straightforward to incorporate in the corresponding sensor modeling and simulation agents. Of course, as is usual, the result will be meaningful only to the extent that pitfalls are knowledgeably and comprehensively incorporated in the models.

In concluding this section it seems appropriate to point out explicitly that if we succeed in creating virtual sensors that can pass a "Turing Test", then we must require of ourselves that we will build them and the systems into which they fit in such a way that they will tell us, perhaps via a dedicated bit in every message packet, whether they are real or virtual sensors. This is essential because, if they do pass the test, then we have no alternative but to trust their veracity.

VII. CONCLUSIONS AND OPPORTUNITIES

We have presented not a conventional experimental or theoretical research paper but rather an airing among friends and colleagues in the VIMS workshop community of the initial results of some general thinking, along with an invitation and a challenge to the community to participate in the anticipated ongoing process as individuals and as a group. The challenge is to recognize, and the invitation is to begin paying increased attention, to the topics of sensor modeling and simulation. In particular, we need to address the growing demand by emerging large-scale applications for the incorporation of sensor models and simulations in real-time monitoring and control systems. The practical need is for model-based simulations to function embedded in these systems as "virtual sensors".

The paper introduces the topic and its background, and summarizes previous work, including the chaotic evolution of what people mean by the term "virtual sensor". It then asks, in essence, "how will we know when we are there?". It tentatively answers this question with the suggestion that there is probably a sensing equivalent of the "Turing Test" for artificial intelligence: if a purported virtual sensor can successfully masquerade as a real sensor, then it is a real virtual sensor. In

the transduction - conditioning - conversion - processing - fusion - communication pipeline it seems natural to surmise that the test's most difficult challenge will be in the first one or two steps, where insufficiently realistic simulation of fundamental and technical noise would serve to unmask an imposter. The similarity is noted to the current situation at the cutting edge of computer graphics technology, where a graphic is now best recognized as not a photograph because it looks "too clean". Finally, some perhaps useful generalizations are drawn regarding sensors, databases, and "reporters" (people) as network data sources and sinks, and some of the general issues are illustrated via an example drawn from the perceived need to do large-scale population monitoring to detect a clandestine attack with biological agents.

REFERENCES

- [1] Turing's classic paper proposing this test, "Computing Machinery and Intelligence" (A. M. Turing, in "Mind, a Quarterly Review of Psychology and Philosophy", Vol. LIX, No.236, p.433ff, October, 1950), can be read on-line at <http://www.abelard.org/turpap/turpap.htm>.
- [2] Doerst, L., F. Groen, and H. Yakali, *Project Virtual Sensors*; see http://www.science.uva.nl/research/ias/princip/virtual_sensor.html.
- [3] Doerst, L., F.C.A. Groen, F.C.A., J.M. van den Akker, J.M., J. Bree-man, A. Hoekstra, J. Lagerberg, A. Visser, H.H. Yakali, and L.O. Hertzberger, L.O. *The virtual sensor concept applied to simulation: the evaluation of Automatic Debiting Systems*, 1997 Workshop on Emergent Technologies and Virtual Systems for Instrumentation and Measurements, Niagara Falls, Ontario, Canada, May 15-17, 1997, pp. 118-124.
- [4] Groen, F.C.A., P.P.J. Antonissen, and G.A. Weller, *Model Based Robot Vision*, in IMTC/93 IEEE Instrumentation and Measurement Technology Conference, pages 584-588.
- [5] Muir, P., *A Virtual Sensor Approach to Robot Kinematic Identification*, Proceedings of the 1990 IEEE International Conference on Systems Engineering, August 1990, pp. 440-445.
- [6] Atkinson, C.M., E. Hanzevack, T.M. Long, and M. Traver, *A Real-Time Neural Network-Based Intelligent Performance and Emissions Prediction System for On-Board Diagnostics and Engine Control*, see <http://www.cemr.wvu.edu/~virtsens/>, neural network based virtual sensor.
- [7] Oosterom, M., R. Babushka, *Virtual Sensor for Fault Detection and Isolation in Flight Control Systems Fuzzy Modeling*, Delft University of Technology, Faculty of Information Technology and Systems, Control Engineering Laboratory, Delft, The Netherlands, see http://dutera.et.tudelft.nl/~oosterom/papers/2000/cdc00_inv5806.pdf, fuzzy logic based virtual sensor.
- [8] Dixon, K., J. Dolan, W. Huang, C. Paredis, P. Khosla, *RAVE: A Real and Virtual Environment for Multiple Mobile Robot Systems*, in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'99), Kyongju, Korea, October 17-21, 1999.
- [9] Haley, T. et al, *Virtual Listeria Sensor*, <http://foodsafety.agad.purdue.edu/research/virtualsensor/>.
- [10] Kestell, C.D., B.S. Cazzolato and C.H. Hansen *Active noise control in a free field with virtual sensors. Part 2: Experimental Verification*, submitted to the Journal of the Acoustical Society of America, see <http://www.mecheng.adelaide.edu.au/anvc/publications/techreports/1999/virtual2.html>, and the same URL with .../virtual1.html for the accompanying theoretical paper.
- [11] Kanade, T., P. Rander, and P.J. Narayanan, *Virtualized Reality: Constructing Virtual Worlds from Real Scenes*, IEEE MultiMedia, vol.4, no.1, Jan. - Mar. 1997, pp.34-47.
- [12] IEEE 1451 is summarized at <http://www.motion.aptid.nist.gov/> and <http://www.ic.oml.gov/p1451/p1451.html>