NISTIR 7060

A Review and Implementation of Algorithms for Fast and Reliable Fire Detection

Walter W. Jones

NIST
National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce

NISTIR 7060

A Review of Algorithms for Fast and Reliable Fire Detection

Walter W. Jones
Fire Research Division
Building and Fire Research Laboratory

October 2003

U.S. DEPARTMENT OF COMMERCE
Donald L. Evans, Secretary
TECHNOLOGY ADMINISTRATION
Phillip J. Bond, Under Secretary of Commerce for Technology
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
Arden L. Bement, Jr., Director

Abstract

The purpose of detecting fires early is to provide an alarm when there is an environment which is deemed to be a threat to people or a building. High reliability detection is based on the supposition that it is possible to utilize a sufficient number of sensors to ascertain unequivocally that there is a growing threat either to people or to a building and provide an estimation of the seriousness of the threat. The current generation of fire detection systems is designed to respond to smoke, heat, gaseous emission or electromagnetic radiation generated during smoldering and flaming combustion. Smoke is sensed either by light scattering or changes in conductive properties of the air, heat by thermocouples and thermistors, the electromagnetic spectrum by photodiodes, and gas concentrations by chemical cells. There is much additional work in progress to use solid-state and electrochemical sensors for oxygen, hydrogen, water vapor, carbon dioxide, chlorine, hydrogen sulfide. The full gamut of fire detection is possible utilizing currently available sensor technology. This includes very early detection as well as fire following. It has been shown to be possible to detect fires early and reliably using the analog signal of the current generation of fire detectors. The best combination for early detection has been shown to be the complement of ionization, photoelectric, carbon monoxide and temperature. This is "best" in the sense that it is possible, using current day sensors, to see characteristic signatures very early, as well as to deduce quantitative information beyond the normal tenability limits. This paper will demonstrate that low level sensing can achieve the goal of producing early detection, while improving reliability. The example we use is a neural network trained with a model of fire growth and smoke spread. This allows us to reduce the time to detection as well as reduce the error rate for both false alarms as well as missing fires.

Acknowledgment

This work is funded by the NIST/NEMA consortium for developing an advanced fire fighter interface panel, and includes funds from the consortium as well as appropriated funds.

Background

The Advanced Fire Service Interface project is a systematic approach to providing useful and consistent information to building managers and emergency response personnel. The intent of this project is to provide better information, faster and more reliably from systems that monitor the environment in buildings, to those responsible for emergency response.

There are three parts to this project:

> The first is to develop a consistent visual display, including a set of symbols (icons), which scales from high resolution displays to personal information displays (handheld computers);
>
> The second part is high reliability detection and feature extraction, through information processing and presentation;
>
> The third part is to provide a basic "look ahead" capability to provide interpretation of building sensor information in a way that is efficient and easy-to-understand.

The systems must be scalable, reliable, and robust, and must allow for new sensors and new algorithms used in the assessment of the environment resulting from unwanted fires and similar threats. The information should be available whenever and where ever it is needed, and to whomever will benefit from the knowledge.

The **first part** of this project, the standard for the graphic annuciator panel, has been incorporated into the 2002 Edition of the NFPA Fire Alarm Code, 72. A full scale demonstration of the capability and usefulness of the system was demonstrated at NIST on April 3, 2002[1] with a full response to what would have been a flashover fire had the NIST fire department not responded with the critical information at the appropriate time.

In order to implement **the second part** we examined the type of information needed to make good decisions rapidly. Answering that question devolved into understanding two levels of sensing and concomitant use for that information. The two levels of sensing are low level for early detection, and high level for fire following. Low level sensing is needed for high reliability early detection and assured fire safety. It should also provide for graceful degradation as sensors become unavailable. Sensors can be unavailable for many reasons: physical degradation if it fails during a fire which would affect fire-following, lack of maintenance from dust accumulation, or extreme environments such as water damage. High level sensing is needed for feature extraction from transducers, *e.g.* the heat release rate (fire size).

The **third part** is extrapolating current conditions into the future. In order to extrapolate from current conditions, one must have an understanding of how a fire is behaving in terms of a model of the environment. The goal is to estimate the environment in a building and to provide this information to all interested parties in a timely manner. Figure (1) shows the effect of measured data on simple extrapolation techniques and the necessity of using a physical model for such parameter estimation.

An example implementation of the visual interface with simple rules for reliability and fire following was highlighted on April 3, 2002 with a full scale demonstration of showing heat release rate derived from heat sensors[2]. Examples of prediction and "what if" (part three) were presented at the 1998 Fire Suppression and Detection Symposium[3].

This paper will concentrate on low level sensing with the aim of producing early and reliable detection. The context of implementation for this study and analysis in which the information is to be displayed is shown in figure (2). This is an example of a graphic annunciator panel as specified in the 2002 Edition of NFPA 72. Of special interest to the fire service is the text item labeled "Likelihood," and is the focus of this paper. It is important to emphasize that this study is aimed at finding a means to indicate the seriousness of a fire, given that an alarm has been sounded. Within the context of detector acceptance and use, and operating procedures used by municipalities, all alarms must be treated as actual fires.

Introduction

The purpose of detecting fires early is to provide an alarm consistent with an environment which is deemed to be a threat to people or a building. High reliability detection is based on the supposition that it is possible to utilize a sufficient number of sensors to ascertain unequivocally that there is a growing threat either to people or to a building and provide an estimation of the seriousness of the threat.

The current generation of fire detection systems[4] is designed to respond to smoke, heat, gaseous emission or electromagnetic radiation generated during smoldering and flaming combustion. Smoke is sensed either by light scattering or changes in conductive properties of the air, heat by thermocouples and thermistors, the electromagnetic spectrum by photodiodes and photovoltaic cells, and gas concentrations by chemical cells[5]. There is much additional work in progress to use solid-state and electrochemical sensors for oxygen, hydrogen, water vapor, carbon dioxide, chlorine and hydrogen sulfide. The use of these detectors, what signals they must (and should not) respond to is mandated by various acceptance tests from Underwriters Laboratories Inc[6], and FM Global Corporation[7].

An important facet of the present work is utilization of sensors which are currently in use in fire detection systems, as well as those available from other systems, such as energy management and security. While the alarms are based on specific criteria, there is additional information available in modern systems. The information from the sensors themselves is analog data, measuring temperature, obscuration, species density, heat flux and other characteristics of the environment. What is needed is a means to provide earlier warning, and more useful information before and after alarm using these sensor suites.

Algorithms

There are three principle ways to approach the detection problem: set point and rate of rise (or a combination), principal component analysis, and neural networks. In all cases, the intent is to signal the presence of a fire or other threat, while not responding to signals which occur in a natural, or non-threatening, environment. In logic terms, the intent is to have no false negatives (missing an alarm) and no false positives (alarm without cause). A difficulty with picking any algorithm is the fact that many different environments can generate similar signals, and any threat can generate many different signals.

We are not trying to classify fires into small or large, nuisance or real. Rather, any situation that might be a threat to people or a structure should be recognized. Signals which show conditions suitable for rapid growth of fire, or possible transition from non-threatening to a threatening environment must be bracketed. The use of multi-criteria based detection technology continues to offer the most promising means to achieve both improved sensitivity to real fires and reduced susceptibility to nuisance alarm sources.

The overarching concern with any improvement are the issues of scalability and backward compatibility. Any system proposed for the future must work with current sensing technology and in existing buildings. It must degrade gracefully as components become inoperable, and must allow for quantifiable improvements in detection time and reliability as new transducers and other technologies are introduced. At the present time, the most common detectors are heat, ion and photoelectric. The addition of carbon monoxide (already exists for other systems) and carbon dioxide detectors are candidates in the near term, and total volatile carbon compounds seem to be viable in the not too distant future. Thus, any study should include this set of transducers. Our effort does not include new means of sensing, such as cross-correlated infrared detectors[8] or video imaging. Such systems might hold potential for the future, but are not common or well characterized at the present time. We will also concentrate on single head, multi-sensor spot type detectors, with future work examining cooperative transducer interaction[9].

There has been a great deal of research into early and reliable detection of fires. Most of the effort has been to reduce the alarms from nuisance signals (Table 1) while responding to those signals which emanate from actual fires (Tables 2 and 3)[10,13,15,27]. These tables were developed from the work focused on combat ships. They have been modified for common civilian situations. In particular, welding puts a normal building at high risk. While it is reasonable to provide a warning rather than an alarm for this type of source, in the uncontrolled environment in common buildings, caution is important.

Rather than try to classify fires with certainty, we are trying to indicate very early detection with the likelihood that the source will constitute a threat in a short time. In these terms, we would classify a nuisance environment as a low likelihood of fire, a smoldering condition as a medium likelihood fire and flaming as a high likelihood fire. Likelihood connotes the relative time to serious conditions, and how well we can confirm these conditions.

Taxonomy for response indicators

Nuisance	(Low Likelihood) - of concern, but no immediate danger; contained or low toxicity; this includes non-combustion sources of particulates, steam and controlled combustion which do not constitute a threat to health or safety
Incipient	(Medium Likelihood) - long time (½ hour or more); can lead to prompt fires over minutes to hours; uncontrolled combustion
Fires	(High Likelihood) - will lead to critical conditions on time scales of seconds to minutes

We actually have two (consistent) meanings for likelihood, which we will use interchangeably. The first is a statistical statement, that a fire really does exist given a set of sensor signals. The second is that given a set of sensor signals, we can know that the environment will become untenable in a long, medium or short length of time. While these are different mathematically, for our purpose they are interchangeable. We do not have a clear enough definition of what constitutes a fire as opposed to a hazard, and not enough data from testing to make a distinction. From the perspective of alerting occupants or the fire department to respond, the end result (time to catastrophe) is the essence. The Underwriters Laboratories Inc acceptance criteria does this implicitly by specifying the time to alarm for various conditions. Their source of the fire is representative of the range of fuels and combustion modes in normal applications.

Set Point and Rate of Rise

"Set Point" and "Rate of Rise" are the methods used in present day sensors. Modern systems also use simple correlations to combine multiple sensors to accomplish the aim of high reliability. Set point values and rate of rise criteria are listed in Table 4. The values cited are the ones used initially in this paper. The actual rules for detector acceptability are more complex.

Used with highly sensitive sensors (*e.g.* smoke detectors), these are good techniques for discriminating signals indicative of fires. The difficulty in their use lies with the relatively high nuisance alarm rate inherent in the techniques. Non-fire signatures look similar to those of fire at low levels of detection. Changes in time-to-detection will be explored using various criteria.

The next level of detection utilizes rules designed to eliminate the low level signals which arise from purely environmental causes. For example, decreasing the level for alarm of a photodetector with a high rate-of-rise in temperature. These rules are really simplified forms of fuzzy logic: "if sensor one is in this range, and sensor two is in this range then it is likely that ..." It is an appealing approach. While such an approach has not yielded great results as yet, this is a reasonable choice as the number of available sensor signals increases. If one examines the test conditions and acceptance criteria enumerated in UL 268 and 521, they lend themselves to such rules, the basic criterion being the time to alarm. It is recognized that fire is somewhat chaotic and so the test conditions are written as rules for effluents from the various test fires. Several such rules will be examined, with a view towards improvement in detection and noise rejection.

Principal Component Analysis

Another technique is principle component analysis (PCA). It is a statistical engineering approach[11], which looks at signals and "judges" which is the most indicative of a fire by examining the matrix of response of detectors to various fires. The essence is to convert the sensors being used to a linear combination which explains most of the variation in transducer output.

If one has a set of instruments (fire detectors in this case), then the set of data listed over time forms a matrix of n (number of instruments) x m (number of measurements). This matrix, which is a function of the vector of instruments, can be changed to another form, where the new instrument vector consists of linear combinations of the original instrumentation. In mathematical terms, one is taking the basis vector (the instruments) and changing them to an orthogonal set of basis vectors, which will have the same resulting matrix when there are identical sensor readings. These new vectors, the principal components, can then be ordered by relative size. If the transducers are independent, all components will be identical and the combination of instrumentation will not improve the predictive capability. If there is some correlation among the sensors, then the principle components will vary in size, depending on their relative importance in explaining the results. If all sensors were completely correlated, then there would be only one principle component.

By using the largest few principal components, rather than the full set of original sensors, one can analyze a much smaller data set. An extreme example would be if one were to find that a single number could be used to show the onset of untenable conditions in an entire building.

There are a number of problems with the approach that renders this method unsuitable for this application. PCA analysis depends on the size of the instrument readings themselves, both from a absolute value and from a deviation perspective. This is not possible *a priori* for fire detectors. It is not known, for example, what the temperature in a compartment will be, nor how much it might vary, what the maximum and minimum might be and so on. Normalizing such data automatically introduces an unknown weighting. Further, such an approach is computationally expensive, and not suitable for real-time analysis. For these reasons, among other, the method has not been successful to date[12] in application to early detection.

On the other hand, an appropriate application in detection would be to consider the historical (hours to days) data from sensors, to determine whether an individual sensor was out of an appropriate range. Given readings over days or weeks, normalization is straightforward and this could be a useful technique for fault detection and sensor re-calibration for fire protection signaling systems. The major alarm manufacturers use various techniques to solve the problems that arise from lowered sensitivity, or changing *ambient* conditions. For example, the Simplex[a] Tru Alarm measures variations from a 256 point moving average of the sensor output signal for that device. The advantage to a PCA analysis over such approaches is that it elucidates the most important sensing components as a part of the analysis. The shortcoming of the PCA analysis is there is not a simple physical interpretation to the few components that result, *i.e.*, what is the meaning of combining a smoke detector signal with a temperature measurement, other than they both correlate with the onset of a fire?

Curve Matching

A further improvement comes with curve matching. Unwanted fires in buildings generally have a characteristic growth rate which is exponential. Usually, the second term in the mathematical expansion of the growth (the square of time), dominates the heat release curve. For this reason, fires are often referred to as "time squared" phenomenon. The short hand expression is t^2. While threats can be present from fires which either have reached a steady-state smoldering condition or arise from some other source, such as carbon monoxide from automobile exhaust or from a residential furnace, the type of phenomena we are trying to detect have characteristics of at least t^2. Curve matching adds a filtering layer to the basic data detected by set point and rate of rise algorithms. Curve matching comes in two flavors: neural net and feature extraction. This paper will focus on neural net implementations which are more suitable for low signal to noise ratio regimes. As will be seen, although it is necessary to use many points on a curve to train a network to recognize the shape of the curve, it is possible to discern the shape with just a few of the early measurements, allowing for very early detection. The error that arises is analogous to the number of points needed to (uniquely) define a curve - an infinite number of parabolas pass through 2 points, but only 1 through three.

Neural net applications are most suitable to the type of filtering we are interested in for likelihood analysis. It is strictly a mathematical technique and provides no information about the underlying phenomenon. Feature extraction through functional analysis allows one to get quantitative information about the fundamental measure, such as heat release rate, but usually requires a higher level of signal. For example, although one can derive the fire size from thermocouples, these same sensors are not sensitive to an incipient fire.

There are two approaches for neural net implementation, depending on the amount of training and testing data available. If one is limited in the number of training scenarios, then the probabilistic neural network (PNN) technique works best. This is for systems which train quickly, though they may run more slowly. The latter problem can be solved by increasing processing power. The most definitive work to date is that of Rose-Pherson *et al.*[13] in which they have used a probabilistic neural network approach. This improvement provides a further reduction of the nuisance and false alarm reaction without missing real fires.

The extension of this idea, explored in this paper, is a full neural network[14]. The difficulty in this approach is the necessity of a large set of training and testing data. Often this becomes too great a hurdle, simply because it is impractical to find the required data. Whereas the PNN approach

---

[a] Certain commercial equipment, instruments, or materials are identified in this document. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available for the purpose.

can be done with hundreds to thousands of test cases, a full implementation can require from thousands to tens of thousands of examples, covering a full range of possibilities for all parameters. Fire modeling has advanced sufficiently that generating a large training set is possible.

While both techniques allow one to further improve nuisance signal rejection, both are inherently "black box" oriented techniques, and so the actual cause for signal matching cannot be well elucidated. They do, however, substantially reduce the false alarm rate, and this can be quantified. The advantage the full neural network approach has is that one is assured that it will reject a higher fraction of signals deemed not to be a threat, yet not miss any threat which can be quantified as a problem. The disadvantage is the much larger number of training sets needed. And as pointed out by Williams and Gottuck[15], one must be careful to separate the data sets used for training a neural network from those used for testing such systems. Otherwise, the system will detect what it already knows, rather than being checked against cases which it should detect but does not know. This is similar to being able to answer a question, given prior knowledge about the question and the answer. There is a distinct advantage to the PNN approach: one can estimate the probability of actual alarms from the mathematics, whereas to do this for the full NN approach requires a great deal more analysis; however, one must classify all signals (no allowance for negative) whereas the latter only requires a prior adjudication as to whether a signal is from a real threat, or must be judged only a nuisance signal.

This study is intended to examine means to provide very early detection for conditions which lead to unacceptable threats. Thus we will be looking for algorithms which can reduce the time to detect a particular threat.

Experiments for Examples of these Algorithms

In order to give examples of current detection strategies and to illustrate advanced detection strategies, we have utilized data from two series of full-scale fire tests. The first series is the Home Smoke Alarm Tests (Part 1, Phase II[16]) and the second is the Smoke Toxicity Tests[17].

Smoke Detector Tests

The smoke alarm tests were run in a manufactured home and were designed to look at changes in efficacy of detection since the original smoke detector tests were performed three decades ago[18]. Fires were set in different areas of the home and the detectors were moved around for each fire area. Signals were analyzed from inside the burn room when the fire was in the bedroom or kitchen and signals outside the burn room when the fire was in the living room (the detectors were located in the hallway outside the remote bedroom). The photoelectric and ionization signals were only available for the tests when the fire was in the living room and the fire source was outside of the burn room. In those instances where there was obviously a detector malfunction (no signal, out of range or alarm prior to the fire) the signal was not considered for analysis.

The physical layout and location of sensors is shown in figure (3) and the location and fuel type of fires in Table 5.

The data measured for this series were temperature, carbon monoxide, carbon dioxide, oxygen, opacity, ionization and photoelectric detectors. The opacity measurements were done with 30 cm laser signal attenuation probes. That is a different technique than either the photoelectric detectors (scattering) or the ionization detectors (current).

Toxicity Tests

The toxicity tests were run in a single room with an extension to a corridor. and all measurements (gas, temperature, velocity) were available for each of the tests. The focus of this research was to ascertain the toxic gases emanating from various fuels as fires progressed from incipient to full room involvement. The intent was to provide a sound basis for the measures cited in the ISO document 13571[19].

A typical layout is shown in figures (4) and (5) and the fuel types and locations in Table 6. The data measured for this series were temperature, carbon monoxide, carbon dioxide and oxygen. No detector or opacity measurements were made.

Existing Algorithms

Set-Point and Rate of Rise

Set-point and rate-of-rise measures are quantified in Underwriters Laboratories Inc 268 and 521standards. The values for carbon monoxide detectors are enumerated in 2034. Although this latter standard is for hazard warning for the presence of carbon monoxide, it does give a lower limit of what is acceptable for a fire test. Pfister[20] has shown that even lower levels of CO detection should give comparable results to smoke detection (we use the 2034 criterion). For the experiments used in this study, temperature, carbon monoxide, carbon dioxide, opacity (from NIST smoke meters), ionization current (not an ionization chamber but output from commercial ionization detectors), and light scattering as seen by photoelectric detectors.

Table 4 enumerates the (initial) values used in this study. Generally, these values are those in the various Underwriters' Standards, though others such as the CO/CO_2 ratio are used from other sources, *e.g.*, Milke[21,22].

An example of detection time based on these values is shown in figure (6). This test was a set of 8 cushions in the form of a loveseat, set against the back wall of the burn room. This is one of the sets of data discussed later which will be used to examine various detection strategies.
An observation of differences between photoelectric and ionization detectors is that the former tend to be better for smoldering fires and the latter for flaming fires, for a given sensitivity. A common strategy for improving the response of photoelectric detectors is to lower its sensitivity when the temperature is rising. For example, if we use the alarm point (shown in Table 4) from 0.05 m^{-1} to 0.025 m^{-1} for the photodetector when the rate-of-rise of temperature is greater than 7 °C /min (a commonly used value), as can be seen in figure (7), a detection time comparable to the ionization detector is obtained.

Correlation of Multiple Signals

Another algorithm which shows a great deal of promise is a carbon monoxide - smoke correlation[23]. In this case, a product of

$$\text{CO (ppm)} * \text{Obscuration (\% per meter)} > 10$$

signals a fire. In tests run so far, it has shown a higher nuisance signal rejection and a shorter time to detection, without missing what are considered real fires.

An example of this algorithm for the home smoke alarm tests (#37 and #38) is shown in figure (8) and figure (9) for a smoldering mattress fire (sdc 37) and a flaming mattress fire (sdc 38).

This compares the CO*Ion signals in units of "ppm*%/m". Generally photodetectors do the best job of detecting smoldering fires whereas ionization based detectors do well with flaming fires. As can be seen, the CO*Ion algorithm does as well as the photodetector in both cases, and much better than the ionization detector for the smoldering fire. This would indicate, as the authors argue, that some additional correlation is needed, and they discuss other permutations, such as adding the CO or ionization signals linearly.

This type of correlation is the simplest form of curve matching, which we will discuss next. It is considerably better than discrete rules such as lowering the sensitivity of one detector based on a reading of a second. However, as discussed in the paper, the choice of correlations is not obvious. It is a shortcoming of trying to choose a correlation based on *a priori* knowledge, and not having a sufficient range of experiments on which to base the choice.

Qualey and Seymouri[24] discuss another such correlation using photo-detectors and heat-detectors instead of the product carbon monoxide - ionization. The intent is to maintain the advantage of photodetectors for smoldering fires but incorporate the same sensitivity as ionization detection for flaming fires. In their paper, they demonstrate that when there is a high correlation between a photodetector and a heat sensor (a multiplier as the CO*Ion shown above), the response time can be similar to an ionization detector. Unfortunately, while results were stated for the reduced time to detection of flaming fires for the photo/heat combination, no guidance was provided on the degree of correlation for such multi-criteria detectors, nor what would constitute a minimum level of detectability.

Curve Matching Algorithms

Curve matching covers a wide range of mathematical techniques, from functional analysis to neural networks. Functional analysis is most useful when the signal to noise ratio is high[25] and one can match the signal to a specific curve of interest, for example, relating a t^2 signal to a heat release rate. Neural network analysis is useful when only the general shape of the curve is known and detail is not justified by the available signal. The regions 1, 2 and 3 in figure (10) show conceptually such a delineation. For all three regions, a pattern can be discerned. However, pattern matching is most usefully applied to the early, noisy signals in region 1 which does not lend themselves to definite statements of functional form, that is, when the signal-to-noise ratio is not high enough to provide a measure of the environment, typically S/N~2 to 4. Region 2 is the current range of available detection when point measurements provide sufficient signal to alarm, typically S/N~3 to 5.[26] Region 3 is appropriate for signal extraction for fire following when the signal to noise ratio is typically greater than 10.

Returning to the earlier discussion, we can also label these regions as low-likelihood (nuisance), medium-likelihood (incipient) and high-likelihood (fires). We want to push detection capability into region 1, yet classify it correctly in terms of advice to the fire service or occupants.

Classification of fire types into low, medium and high likelihood consequences has implications for both fire service as first responders, and building maintenance personnel who might be able to fix problems before they rise to emergency status.

Figure (11) shows a typical sensor reading from a fire, carbon monoxide in this case. Detecting the presence of a fire traditionally has been to measure such signals, and provide an alarm when some condition is reached, for example, when the opacity is high or the carbon monoxide too high. Shown in the figure are alarm points for several detection strategies, an ionization detector, a photoelectric detector, and the CO*Ion algorithm discussed previously. The example is a surrogate for the range of signals which might be used for detection of fires[4,27]. Currently, temperature (T), opacity (OD), ionization (Ion) and carbon monoxide (CO) are the core signals we will focus on. In addition to these, carbon dioxide (CO_2), volatile organic hydrocarbons (VOC), nitrogen-oxygen compounds (NO), oxygen(O_2) and water concentration (RH) are possible future signals to incorporate.

An example of using pattern matching is discussed in the paper by Rose-Phersson et al.[28] The focus of the paper was the use a probabilistic neural network to combine signals from several transducers to reduce the likelihood of both false positives and false negative responses from detector systems. While this is similar to what we will use to reduce the time delay, the focus was on more reliable detection. The goal behind their work was to automate response to fires (e.g. sprinkler activation), so very high reliability is even more important than early detection. They demonstrated the optimal sensor set to be ionization, photoelectric, carbon monoxide and carbon dioxide, with temperature providing the best confirmation signal. In our case, we will work from the premise that the patterns we see will result in an alarm condition from the installed alarm base, so we want to respond as early as possible to these signals or patterns, in order to reduce the response time of the firefighters.

An Example of Implementation of a Neural Net Algorithm

An artificial neural networks (ANN) is a collection of mathematical models that emulate some of the observed properties of biological nervous systems and draw on the analogies of adaptive biological learning. The key element of the ANN paradigm is the structure of the information processing system. It is composed of a large number of highly interconnected processing elements that are analogous to neurons and are tied together with weighted connections that are analogous to synapses.

Learning in biological systems involves adjustments to the synaptic connections that exist between the neurons. This is true of ANNs as well. Learning typically occurs by example through training, or exposure to a "truthed" set of input/output data where the training algorithm iteratively adjusts the connection weights (synapses). These connection weights store the knowledge necessary to solve specific problems.

ANNs are being applied to an increasing number of real- world problems of considerable complexity. They are good pattern recognition engines and robust classifiers, with the ability to generalize in making decisions about imprecise input data. They offer ideal solutions to a variety of classification problems such as speech, character and signal recognition, as well as functional prediction and system modeling where the physical processes are not understood or are highly complex. ANNs may also be applied to control problems, where the input variables are measurements used to drive an output actuator, and the network learns the control function. The advantage of ANNs lies in their resilience against distortions in the input data (blips, noise and such) and their capability of learning. They are often good at solving problems that are too complex for conventional technologies (e.g., problems that do not have an algorithmic solution or for which an algorithmic solution is too complex to be found) and are often well suited to problems that people are good at solving, but for which traditional methods are not (you know a fire when you see it!).

There are multitudes of different types of ANNs. Some of the more popular include the multilayer perceptron (what we used) which is generally trained with the backpropagation of error. Some ANNs are classified as feedforward while others are recurrent (i.e., implement feedback) depending on how data are processed through the network. Another way of classifying ANN types is by their method of learning (or training), as some ANNs employ supervised training while others are referred to as unsupervised or self-organizing. This work used supervised training. An improvement over the results discussed below would include a regime of unsupervised training (see for example the work of Phersson et al.[13]) for classification of signals rather than simple decision aids.

The study of fire occupies a unique niche in the world of science and engineering because an unwanted fire is considered a failure in the sense that it is not a desirable outcome and is to be avoided. Detection and suppression are thus posed as means to avoid failure, which can be well characterized. For detection in particular, we have well defined failures which can be tested fairly reproducibly. In this highly regulated environment, in order for detectors to be approved for use they must detect fires as defined in UL 268 and EN 54 tests. In addition, there are nuisance criteria when the detectors should not alarm. While these latter are well recognized (dust, for example[29]), there are no formal tests, though a simple negative (no fire) should in no case produce an alarm (a false positive). For the UL tests, there is a time prior to when the alarms should not activate.

The biggest difficulty in training neural networks is the extent of the training scenarios available. In fire research, the work has been limited to experimental data sets, for example of work of Rose-Phersson discussed earlier. Typically the training set consists of tens to hundreds of scenarios, while ANNs need tens of thousands to produce highly reliable classification. Using the fire model, CFAST, we can generate a very large set of training and testing scenarios.

For this example, we consider the use of single head (multisensor) detector in a single compartment. The use and limitations of such detection are covered by NFPA 72, the National Fire Alarm Code. We have a model for fires which has been extensively tested, CFAST[30]. We used this model to generate training and testing scenarios which cover a very fine delineation of the event to be detected. Using such a model allows us to generate the tens of thousands to hundreds of thousands of examples necessary to provide sufficient training for a network.

The base case used

Standard atmosphere of 101,300 Pa
A single compartment of 13x13x2.4 m
Two cracks (one vertical, one horizontal) to account for leakage
One door of (0.9 x 2.3) m and One window of (0.9 x 1) m.

Starting with this base case, variations of the base case scenario were generated based on

(3) Ambient conditions: outside to inside temperature the same or ±15 °C
(3) Wind: none, into door (away from window if present) or away from door
(3) Fire size: (1, 10, 100) kw - note: no fire at all is a special case
(3) Position of fire: floor, and 0.5 m, 1.0 m above the floor in the center of the room
(4) Door width: open, ½, 1/4, 1/8 width
(2) Window: open or closed (0.9 m x 2.3 m)
(4) CO: (0.0, 0.001, 0.01 and 0.05) kg/kg or (0%, 0.1%, 1% and 5%) by fraction
(3) Smoke yield (optical depth): (0, 0.01, 0.05) kg/kg or (0%, 1% and 5%) by fraction
(2) Hydrogen carbon ratio in the fuel: (0, 0.2) kg/kg

(Note that the hydrogen carbon ratio is for the fuel, whereas CO and smoke are kg/kg for the combustion products).

This is 20,768 variations, which were then used to train the neural network. The scenarios were 300 second calculations with a time slice every 30 seconds. While this suite is sufficient to demonstrate the feasibility of training multisensor networks, a somewhat more comprehensive set of scenarios might include

(9) Room size: 2 m' to 5 m ceiling, 3x3 m to 15x15 in 3 increments each
(2) Geometry: rectangular parallel piped, aframe
(3) Radiative fraction: 0.1, 0.2 and 0.3

which would increase the number of scenarios, calculation and training time about an order of magnitude.

Three training exercises were performed
1) a subset of the parameter space comprising 5000 scenarios, and 5000 for testing
2) a complete set of scenarios (20726), and a small subset for testing (42) (total of 20768), and
3) preconditioning to supplement training for those cases when a fire is known to exist.

The sensor suite used were four sensors: oxygen, carbon monoxide, opacity and temperature. A more complete characterization would consider each sensor separately, as well as all combinations. This would provide a sense of the effect of losing a sensor (fault detection).

There are two aspects of detection which are important if we are to consider the algorithms to be fast and reliable. In order to be considered fast, the detection scheme must be at least fast as current detection algorithms. For high reliability, we are looking for means of seeing all real fire (no false negatives), and not responding to those deemed to be nuisances (no false positives). A metric for the former will be discussed as part of the analysis of results. The metric for false positives (nuisance alarms in the present context) and false negatives (missing a real fire), the scenarios are either fires or nuisance signals. Except for the base case of no heat release, which by definition is not a fire, the remainder are classified as real or nuisance by whether they pose a threat at any point in the curve to people or property. The classification is based on the Hazard I[31] methodology and the ISO Toxicity Specification[32]. For exercises 2 and 3, of the total scenario space, 15 916 cases were fires and 4 852 non-fires. These latter (23%) are nuisance signals in the present context. A more complete classification scheme would further classify these according to Tables 1 through 3.

Mathematically, a neural network is a set of weight matrices which multiply sensor signals, and use a function (in our case a linear ramp) to combine the results. This provides a classification of data. Schematically, it is shown in figure (12), where **p** represents the measurement points, a vector of length R (in our case, this is the number of sensors), **b** a bias vector for the algorithm (always set to zero in our training), **w** the weight matrix (the answer so to speak). In the following training cases, we used R=4, but typically, it can range from 1 (a single sensor) to 9 (see ref. 10) which would be a very general multi-criterion sensor head.

The end point of such a system is a weight matrix which when multiplied by the sensor suite (**p**) produces a classification number; we used a simple classification of true or false (fire or non-fire). We trained a network with a single hidden layer of 10 neurons, and a single output layer using a linear transfer function. Thus we have only one matrix which needs to be adjusted. The training method used was Levenberg-Marquardt[33]. We have a set of four sensors, with 31 points (30 intervals). The data were presented to the learning algorithm, which modified the weight matrix (**w**) until a (defined) error level was reached.

We applied this technique using the Matlab[34] simulation tool, with the Neural Network Toolbox. Each data set was presented to network, and it adjusted the weight matrix. After completing the training, the network was presented test data, and classified the new sensor readings as a fire or non-fire event. Since we are concerned with a binary decision, the results were descritized to 0 or

1. In actuality, the data was a spectrum and additional training could be provided to further refine the classification scheme to non-fire, nuisance or significant event.

For the first case, there were no false positives or false negatives. That is, all fires were detected and no alarms when a fire did not exist. The time to alarm was generally the same for conventional detection and the trained network. The time to do the CFAST calculations was approximately 45 minutes, and the training time approximately 1 hour.

The time/temperature curve shown in figure (13) has the alarm points overlaid. The solid lines are example 1 and the dashed lines example 2. The vertical ticks are the corresponding detection time for conventional detection (green) and the neural net with training (red).

For the second example, all 20768 scenarios were used. In order to test the network, 42 of the 20768 scenarios were used for testing and not used for training. This then constituted a sampling of data which the network should be able to recognize. Of the forty two tested, there were no false positives (nuisance alarms), that is no fire detected when a fire did not exist; however, there was one false negative, not showing an alarm when a fire was present. This is about a 2% failure rate. The scenario which failed is marginal for the network, and to improve performance, the scenario suite needs to be extended to provide a finer resolution. In actual commercial detection systems, false negatives occur (3 to 20)% of the time[35] and false positives (30 to 50)% of the time[23], so we have improved on the detection capability as well as reduced the time to detection.

This training was done with a 10 neuron system. A systems with 20 neurons and two hidden layers was tried as well, without improvement. The time to detection for this second training example was always as early as, and usually earlier, than conventional detection, as shown in figure (14). The same notation is used in this figure as earlier. The time to do the CFAST calculations was approximately 2 hours, and the training time approximately 3 hours. The two cases shown, 007051 and 017658, are randomly picked from the 42 test cases.

For the third training example, the truth vector (when the fire exists) was preconditioned for those cases we know a fire will exist. For example, for the 100 kW source, it will at some time be considered a fire. For these cases we can set the training vector to "true" at after the first interval. Once again, there were no false positives and a single false negative (same case as before). The time labeled "preconditioned" in figure (14) was the response for the two cases shown in the figure for the example 2 testing regimen, 007051 and 017658, thus showing the value of using additional information in the training regimen.

This third training example takes advantage of the fire problem. We start with the scenarios. These produce curves of time, temperature, co, and so on. At some point we decide there is a fire. At the simplest level, used in 1 and 2, it is done the based on commercial detection schemes or the toxicity assessment discussed earlier. However, we can add to that information base, by noting that certain scenarios are going to be classified as fires, and tell the system from the beginning. For example, a 100 kW fire will must be detected, as must a 5 % CO condition. So for certain scenarios, one tell the system that it is a fire after the first interval. That gets factored into the weight matrix so that curves of similar shapes trigger an alarm very early. And even ones that are close do so. It is because we are matching curves (high precision) and not trying to get detailed information (high accuracy) that this technique is so appealing in this application.

There is additional work which needs to be done before this can be used in actual sensor suites: final testing for this case needs to include an example experiment such as the Smoke Detector Tests[36]. In addition, the standard qualification tests and a set of nuisance signals must be included. This latter will require an instrument transfer function, which can be measured using the FE/DE test apparatus[37]. Finally, the training suite ought to be extended to include the wide range of geometries which exist in practice rather than just those used for qualification testing. The training of a neural network should allow this extension and would improve the robustness of detection systems. This then allows one to include cases which currently cause alarms, such as steam, but are clearly not fires.

A further extension would be to go beyond the simple alarm/no-alarm classification we have done here and report on nuisance alarms as distinct from fires. Interestingly, a cursory inspection of the testing scenarios shows that the network is doing a reasonable characterization of the scenarios in terms of the type of fires. It is likely that this work could be extended to classification according to Tables 1 through 3. This is important in that a nuisance signal is often a precursor to more serious conditions. The prime example is the case of an oven (and even more commonly a toaster oven) which can develop the right conditions (and measurable effluent) but has a low level fire until a door is opened.

Conclusions

The full gamut of fire detection is possible utilizing currently available sensor technology. This includes very early detection as well as fire-following. It has been shown that it is possible to detect fires early and reliably using the analog signal of the current generation of fire detectors. The best combination for early detection has been shown to be the complement of ionization, photoelectric, carbon monoxide and temperature. This is "best" in the sense that it is possible, using current day sensors, to see signatures very early, as well as to deduce quantitative information beyond the normal tenability limits.

The most useful of the algorithms studied is the curve matching concept embodied in neural network methods. In training such algoritjms, it is important to use a sufficiently large set of training and testing samples so that that the algorithm is robust. We would expect a single experiment to provide very early detection for that single response curve. However, as the number of training sets is increased, incorporating variations in geometry and insult, the time to reliable detection increases. As the number of sensors used increases, we expect the detection time to decrease. The trade-off is in the necessity for using large (more than 10,000) sample sets. With a judicious use of modeling and experimental testing, this should not be a burdensome exercise. We have demonstrated the training of a neural network to shown that it is possible, including very early detection. Although we find a 2 % error rate with the present training regimen, this is still considerably better than current detection (3 to 30)% as well as methods proposed to date (2 to 10)%.

References

1. http://www.bfrl.nist.gov/bfrlnews/monthly_highlights/june_2002.htm

2. Spring Meeting and Exposition of the National Fire Protection Association, Building Fire Safety Systems, Minneapolis (2002).

3. Using Sensor Data to Predict the Environment in a Building, W.W. Jones, R.D. Peacock, G.P. Forney and P. A. Reneke, Proceedings of the Fire Suppression and Detection Research Application Symposium, Orlando (1998).

4. A Review of Measurements and Candidate Signatures for Early Fire Detection, W.L. Grosshandler, NIST Internal Report 5555 (1995).

5. The Use of Surface and Thin Film Science in the Development of Advanced Gas Sensors, S. Semancik and R. Cavicchi, Applied Surface Science *70/71*, 337 (1993).

6. Underwriters' Laboratories, 333 Pfingsten Road , Northbrook, IL 60062-2096 USA

7. Factory Mutual Global, 1301 Atwood Avenue, Johnston, Rhode Island 02919 USA

8. Novel CCTV Fire Detection Technologies for Industrial Applications, O. Spector, E. Jacobson, and J. Walker, Fire Suppression and Detection Research Application Symposium, Orlando (2001).

9. Distributed Sensor Fire Detection, T. Cleary and K. Notarianni, International Conference on Automatic Fire Detection AUBE '01, Gaithersburg (2001).

10. Development of an Early Warning Multi-Criteria Fire Detection System: Analysis of Transient Fire Signatures Using a Probabilistic Neural Network, R.E. Shaffer, S.L. Rose-Pehrsson, C. Barry, D.T. Gottuk, and F.W. Williams, NRL Memorandum Report, NRL/MR/6110-00-8429 (2000).

11. *Principal Component Analysis*, I.T. Jolliffe, Springer-Verlag, New York (1986).

12. Using Multivariate Statistical Methods to Detect Fires, T.J. McAvoy, J. Milke and T.A. Kunt, Fire Technology, *First Quarter*, 6 (1996).

13. Early Warning Fire Detection System Using a Probabilistic Neural Network, S. Rose-Phersson, S.J. Hart, T.T. Street, F.W. Williams, M.H. Hammond, Fire Technology (in review).

14. *Advanced Algorithms for Neural Networks*, T. Masters, John Wiley, New York (1995).

15. Multi-Criteria Fire Detection: A Review of the State-of-the-Art. D.T. Gottuk and F..W. Williams, NRL Letter Report Ser 6180/0472 (1998).

16. Home Smoke Alarm Project, Manufactured Home Tests, R.D. Peacock, J.D. Averill, R.W. Bukowski, and P.A. Reneke, NIST Report of Test FR 4016, June (2002).

17. Smoke Component Yields from Room-scale Fire Tests for Validation of Bench-scale Smoke Toxicity Apparatus, R.G. Gann, J.A. Averill, E.L. Johnsson, M.R. Nyden, R.D. Peacock, and D.W. Stroup, NIST Technical Note xxx (2002).

18. Detector Sensitivity and Siting Requirements for Dwellings, R. Bukowski, W. Christian, and T. Waterman, T.E., Nat Bur Stand, Washington DC, NBSGCR 75-51, NFPA SPP-43 (1975).

19. "Fire Hazard Analysis – Estimation of Time Available for Escape from a Fire," Proposed Draft Technical Specification, ISO CD 13571, 14th Draft. International Organization for Standardization. (To be published.)

20.Detection of Smoke Gases by Solid State Sensors - A Focus on Research Activities, Fire Safety Journal 6, 165 (second quarter, 1997).

21. Analysis of Fire and Non-fire Signatures for Discriminating Fire Detection, J.A. Milke and T. J McAvoyy, Proceedings of the Fifth International Symposium on Fire Safety Science,

Melbourne (1997).

22. Development of a Discriminating Fire Detector for Use in Residential Occupancies, MS Thesis, College Park, Fire Protection Engineering, University of Maryland (1993).

23. Advanced Fire Detection Using Multisignature Alarm Algorithms, D. Gottuk, M. Peatross, R.J. Roby and C. L. Beyler, International Conference on Automatic Fire Detection, AUBE '99, Duisberg (1999).

24. Development of a Multisensing Detector, J. Qualey and R. Seyouri, Proceedings of the Suppression and Detection Sympoisum, Orlando (1998).

25. *Functional Data Analysis*, J.O. Ramsay and B.W. Silverman, Springer-Verlag, New York (1997).

26. Real-time Probabilistic Neural Network Performance and Optimization for Fire Detection and Nuisance Alarm Rejection, S.L. Rose-Pehrsson, S.J. Hart, T.T. Street, P.A. Tatem, F.W. Williams, and M.H. Hammond, International Conference on Automatic Fire Detection, AUBE '01, Gaithersburg (2001).

27. Identification of Fire Signatures for Shipboard Multi-criteria Fire Detection Systems, D.T. Gottuk, S.A. Hill, C.F. Schemel, B.D. Strehlen, R.E. Shaffer, S.L. Rose-Pehrsson, P.A. Tatem, and F.W. Williams, NRL Memorandum Report, NRL/MR/6180-99-8386 (1999).

28. Multi-Criteria Fire Detection Systems using a Probabilistic Neural Network, S.L. Rose-Pehrsson, R.E. Shaffer, F.W. Williams, D.T. Gottuk, S.A. Hill and B.D. Strehlen, Sensors and Actuators, *B 69*, 325 (2000).

29. Smoke Detector Response to Nuisance Aerosols, T. Cleary, W. Grosshandler, A. Chernovsky, Proceedings of the 11th International Conference on Automatic Fire Detection, Duisburg (1999).

30. Technical Reference for CFAST: An Engineering Tool for Estimating Fire and Smoke Transport, W. Jones, G. Forney, R. Peacock and P. Reneke, NIST TN 1431 (2000).

31. Hazard I, The Fire Hazard Assessment Method, R. Peacock, W. Jones, R Bukowski, and C. Forney, NIST Handbook 146, Vols I and II (1991). An Update Guide for Hazard I, Version 1.2, R. Peacock, W. Jones, G.Forney, R. Portier, P. Reneke, R. Bukowski and J. Klote, NIST Technical Report 1510 (1994).

32. ISO/TS 13571:2003, Life Threatening Components of fire - Guidelines for the estimate of time available for escape using fire data, International Organization for Standardization 1, rue de Varembé, Case postale 56, CH-1211 Geneva 20, Switzerland.

33. D. Marquardt, An algorithm for least squares estimation of non-linear parameters, SIAM J., 11:431 – 441, 1963; K. Levenberg. A Method for the solution of certain nonlinear problems in least squares. Quart. Appl. Math, 2:164 – 168, 1944.

34. The Matlab Development Environment, The MathWorks, Inc., Apple Hill Drive, Natick, MA 01760.

35. Analysis of Fire Signature Data Toward the Development of an Advanced Fire Detector, Progress Report on NIST Grant 1D0076 by J. Milke (September, 2001).

36. Performance of Home Smoke Alarms, Analysis of the Response of Several Available Technologies in Residential Fire Settings, R. Bukowsi, J Averill, R. Peacock, T. Cleary, N. Bryner, W. Walton, P. Reneke and E. Kuligowski, NIST TN xxxx (in press).

37. Fire Emulator/Detector Evaluator: Design, Operation, and Performance, T. Cleary, M. Donnelly, and W. Grosshandler, NIST SP 965, Proceedings of the 12 International Conference on Automatic Fire Detection, AUBE '01, Gaithersburg (2001); Towards the Development of a Universal Fire Emulator-Detector Evaluator, W. Grosshandler, Fire Safety Journal *29*, 113 (1997).

Table 1. Nuisance signals (low likelihood)

Hairspray, Nail polish remover , bleach, furniture cleaning agents, disinfectants
Toaster effluents - except as can be classified as incipient fires, for example in toaster ovens
Ovens
Boiling water, coffee, showers and other steam sources
Dust and sawdust, concrete dust, overcooked popcorn and other microwave products
Propane and kerosine heaters and stoves
Candles
Cigarettes and matches
Heating systems (furnace)

Table 2. Incipient (long time to disaster) fires

Toaster oven effluents
Welding torch and arc welding
Cook-top effluents, frying bacon
Smoldering mattress, chair or other cushion furniture: cotton, down

Table 3. Fires (prompt)

Open cellulose fires (crumpled newspaper)
Flaming mattress, chair or other cushion furniture: cotton and foam
Liquid pool fire (heptane, gasoline, alcohol, paint thinner, acetone, vegetable oil)
Wood (wood based) furniture such as bookcases
Smolder mattress, chair or other cushion furniture: foam
Power and signaling cables
Interior wall coverings such as wallpaper

Table 4. Set point and Rate of Rise Criteria (single point measurements)

The values shown in this table are those used in the study. The actual values which can be used are somewhat more complex. The applicable tests are Underwriters Laboratories[6] 217, 268, 521, 529, 2034. Specific values in common usage are given in Ref. (15) except where noted.

Ion smoke detectors:	(2 to 6.4) % obscuration per meter
Typical is 4 % which corresponds to an optical density of 0.017 m^{-1}	
Photoelectric detectors:	(2 to 12) % obscuration per meter
Typical is 11 % which corresponds to an optical density of 0.051m^{-1}	
Temperature:	57 °C.
Temperature rate of rise:	7 °C per minute
Carbon monoxide:	50 ppm (Pfister[20] suggests 25)
Carbon dioxide[22]:	1.5%
Oxygen:	17%
Ratio of carbon monoxide to carbon dioxide[22]:	0.01

Table 5 - Test in the Smoke Alarm Series - fuel type and location
SD 37, Smoldering Mattress, Bedroom SD 38, Flaming Mattress, Bedroom SD 39, Flaming Mattress, Bedroom SD 40, Smoldering Mattress, Bedroom SD 41, Cooking Oil Fire, Kitchen

Table 6 - Tests in the Toxicity Series - fuel type and location

Test series	Fuel	Location
BW	2 bookcases	Rear wall
CW	8-12 cushions	Rear wall
BW	2 bookcases	Rear wall
PW	Cable	Rear wall
CW	14 cushions	Rear wall
BP	2 bookcases/PVC	Center
CC	8 cushions	Rear wall
BW	2 bookcases	Rear wall
BP	2 bookcases/PVC	Center

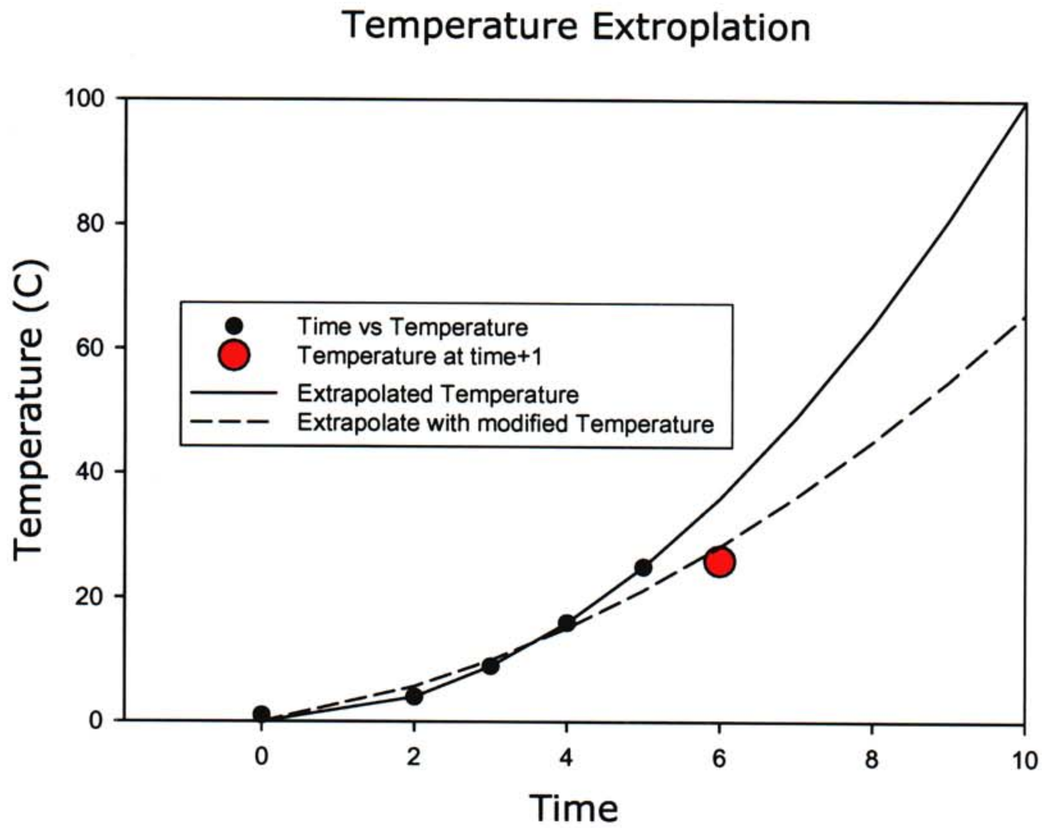

Figure 1. Extrapolating temperature based on t^2 curve fit.

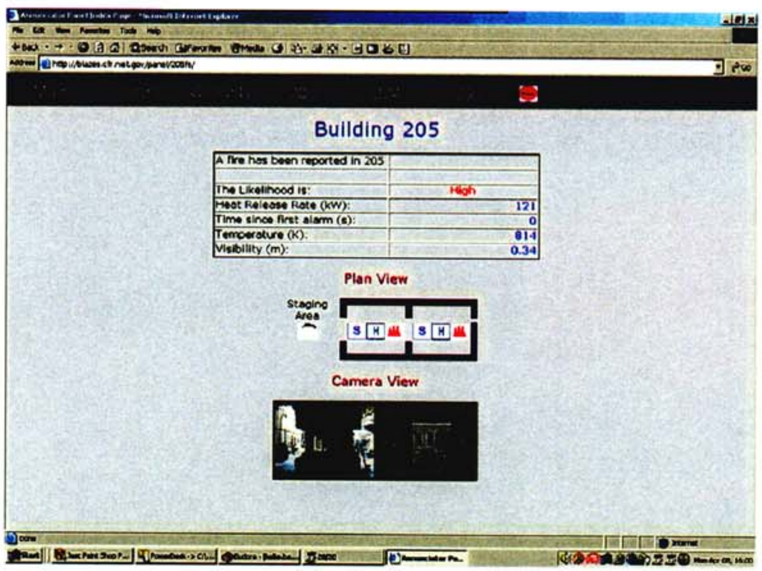

Figure 2. Sample implementation of NFPA 72 Annex A (http;//panel.nist.gov/).

Mattress Fire in the Bedroom

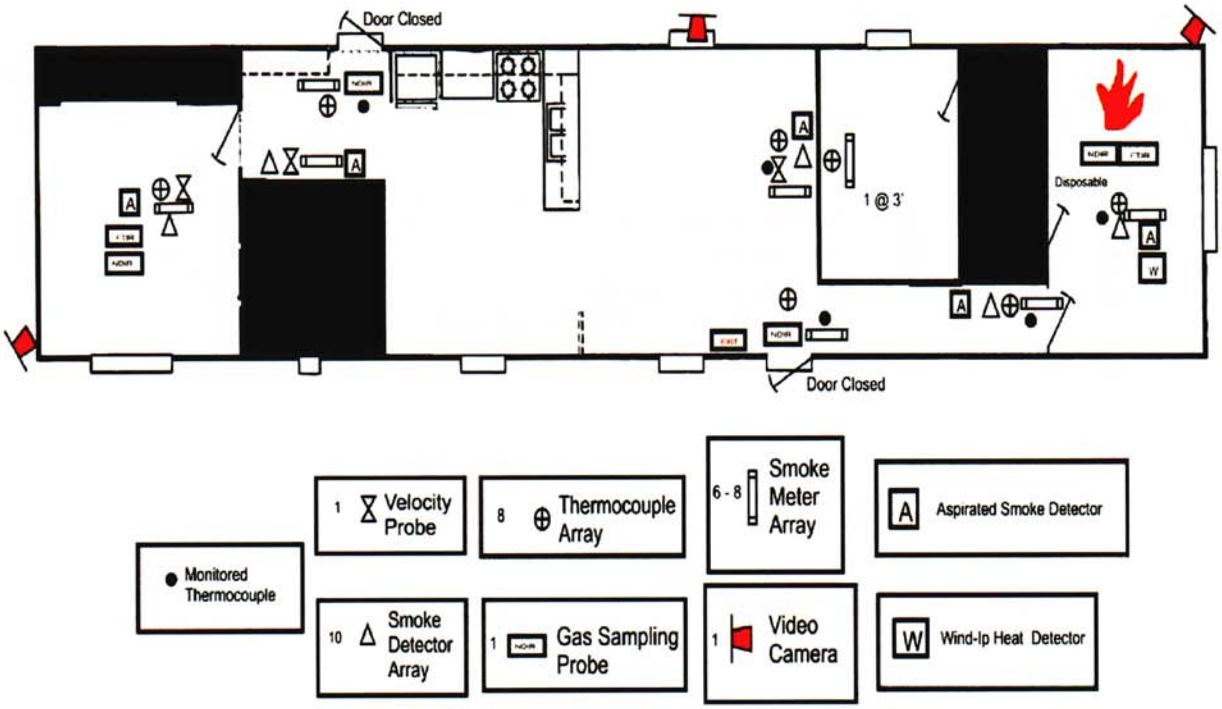

Figure 3. Instrument locations for smoke detector experiments.

Figure 4. Cushion fire prior to flashover

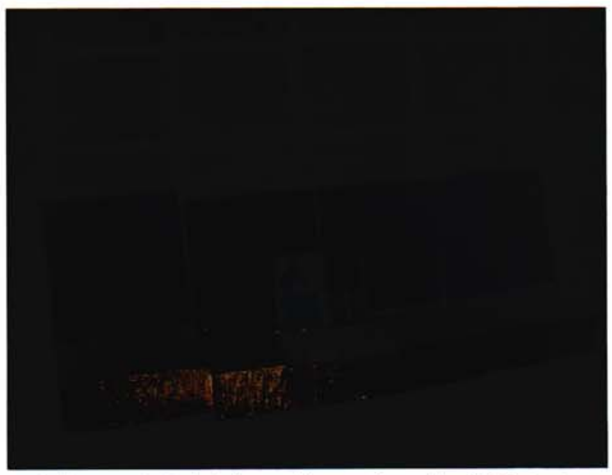

Figure 5. Initial configuration for CW test series

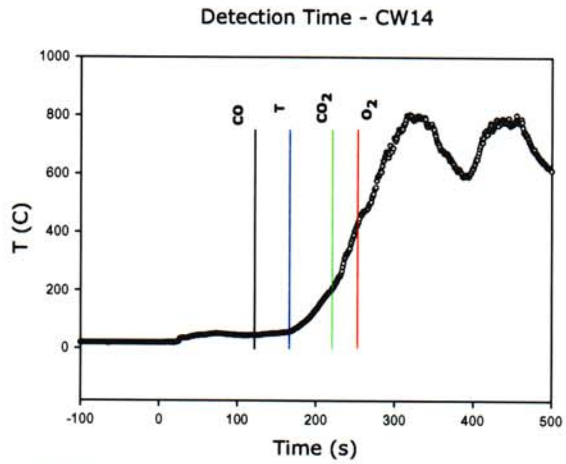

Figure 6. Comparison of detection schemes using specified set points.

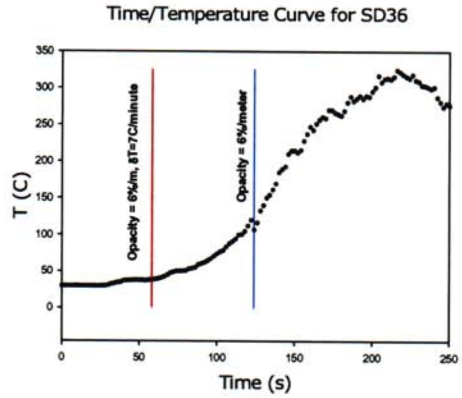

Figure7. Augmenting opacity set point (–) with a temperature rate of rise (–).

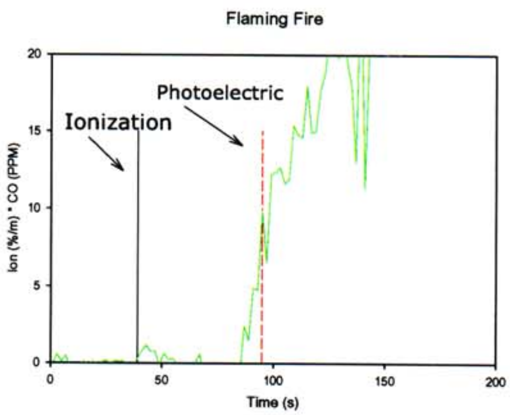

Figure 8. Comparison of detection time for photoelectric (–), ionization (–) and the CO*Ion algorithm (–)for a flaming mattress fire.

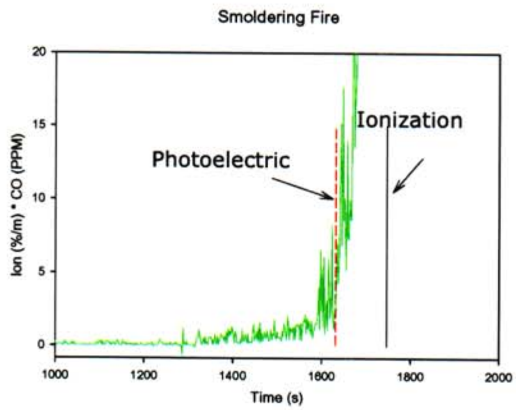

Figure 9. Comparison of detection time for photoelectric (–), ionization (–) and the CO*Ion (–) algorithm for a smoldering mattress fire.

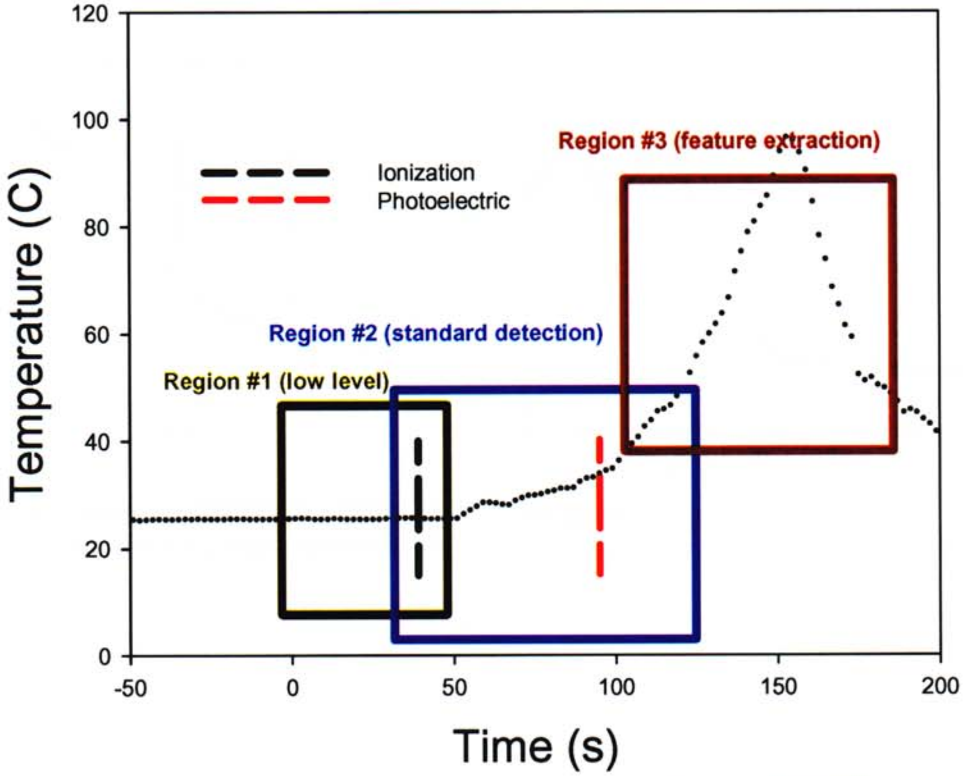

Figure 10. Delineation of detection regions from a flaming fire.

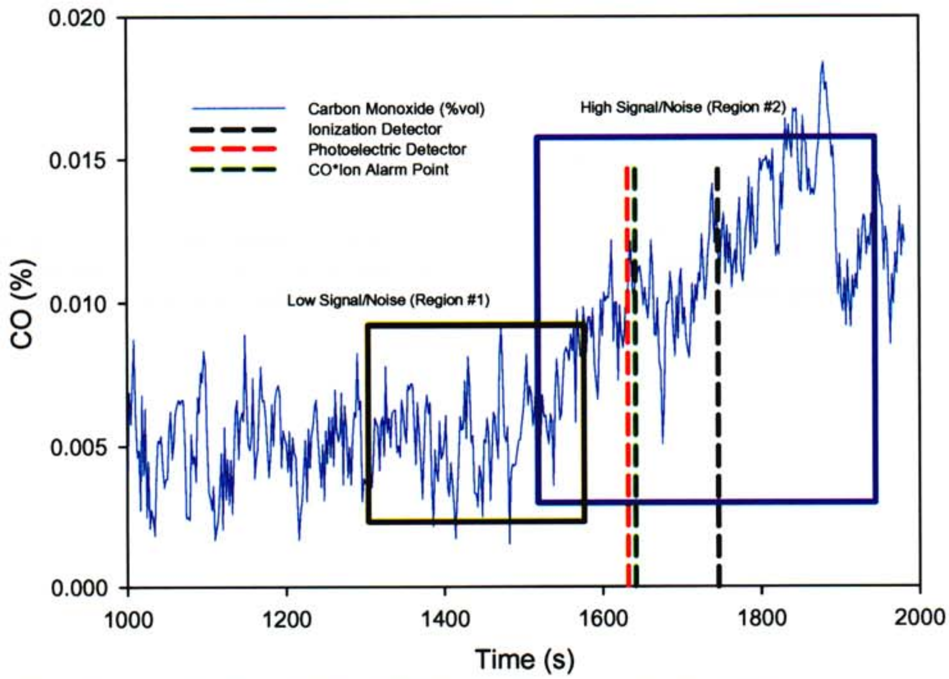

Figure 11. Carbon monoxide signal in SD 37, showing regions for curve matching algorithms.