

Survey of autonomous gas leak detection and quantification with snapshot infrared spectral imaging

Nathan Hagen 

Utsunomiya University, Department of Optical Engineering, Center for Optical Research and Engineering (CORE), Utsunomiya, Japan

E-mail: nh@hagenlab.org

Received 19 April 2020, revised 5 August 2020

Accepted for publication 24 August 2020

Published 22 September 2020



Abstract

We survey the current state of the art in industrial gas sensing, explaining how existing sensing techniques estimate gas concentration, column density, total volume, and leak rate. We especially focus on the application of snapshot infrared spectral imaging to gas detection and quantification, as this is the newest of the available techniques. After discussing the strengths and weaknesses of each measurement technique for deployment in autonomous sensing, we explain how autonomous detection systems are important to improving sensing capabilities in industrial facilities, and advancing the state of the art.

Keywords: gas detection, infrared imaging, leak rate, emissions quantification, autonomous sensing

(Some figures may appear in colour only in the online journal)

1. Introduction

Two important challenges face the oil and gas industry: an increasing number of explosions at industrial facilities caused by gas leaks [1], and tightening controls on gas emissions. These challenges point to a need for improved equipment and methods for gas leak detection and quantification [2], not only for locating fugitive emissions (caused by malfunctioning parts) but also for quantifying emissions by equipment that release gas as part of their design, such as flares and pressure relief valves [3]. Current leak detection and repair (LDAR) protocols as well as emission quantification efforts are implemented by human engineers, so that evaluation is time-consuming and expensive. This high cost of current methods presents a barrier to improvement.

Autonomous gas sensing methods can advance the current state of affairs by removing humans from the most labour-intensive detection and monitoring tasks. Autonomous systems can operate continuously night and day, can log and transmit measurements in real time, and have been demonstrated to not just detect leaks but quantify them at the same time. Moreover, they can provide ancillary information such as

long-term monitoring of equipment and environmental conditions, detecting trespassers, and more.

For an autonomous system to be useful, it must have a low false detection rate, be reliable at alerting when a dangerous situation occurs, and should be effective at quantifying gas plume sizes and leak rates—all without human involvement. In the discussion below, we explain how snapshot infrared spectral imaging is the key to making this approach workable. In order to show how autonomous instruments can be used to overcome existing labour-intensive methods, we survey the capabilities and limitations of this rapidly developing approach and compare with alternative instrumentation. After reviewing the theory for detection and quantification of gas clouds using passive spectral imaging, we show measurements from early field deployments.

2. The gas detection status quo

At chemical processing facilities, the standard protocol to monitor for gas leaks in the United States is EPA Method 21 [4], which requires that designated components in the facility (valves, flanges, pumps, etc) be leak-tested at regular intervals.

An LDAR engineer implementing the monitoring procedure carries a portable gas leak detection instrument (a ‘gas sniffer’) and slowly passes the instrument’s probe head around the component. The instrument draws ambient air up through the probe head and into a measurement chamber that quantifies the gas concentration in the air. For gas leaks whose measured concentration exceeds a given threshold for the component under test, the protocol demands that the component be scheduled for repair.

While the Method 21 procedure is straightforward, we can also see that the subjective behaviour of the engineer plays a substantial role in the outcome. For instruments with a slow response time, the engineer must be careful to hold the probe head at the source of the leak for a time period exceeding the instrument’s 90% response time. Meanwhile, wind conditions and component accessibility affect the extent to which the engineer can locate and accurately quantify the leak.

Method 21 also mentions an alternative procedure in which a soap solution is sprayed onto the components under test, a procedure sometimes referred to as ‘soaping’. The user then looks for the appearance of bubbles. If none appear, then we presume that there is no significant leak; if bubbles appear then the protocol calls for returning to the location of the leak with a gas sniffer.

For quantification of the volumetric flow rate of a leak, Method 21 suggests two approaches [5]. The more accurate approach is to seal the leak source with a plastic bag and allow the trapped gas to pass through a flow meter and a gas concentration detector, together which can be used to give the leak rate in units of L/min or g/s. While this ‘bagging’ procedure provides an accurate measure of the size of the leak, it is extremely labour intensive, so that an engineer can quantify only about 10 leaks per day this way. Method 21 gives a second leak rate quantification method, which is simply to multiply the measured gas concentration near the leak, using established scaling factors that vary according to the item under test. The result, however, can vary over three to four orders of magnitude [6, 7], so that the quantification uncertainty in this latter approach remains frustratingly large.

Recognising the advantages of using infrared cameras for gas detection, facilities have increasingly turned to infrared imaging for LDAR monitoring. This process involves briefly staring at the components under test with an infrared camera and looking for any signs of gas leaks. Since a user can examine many components within a single image, can image components that are difficult to access with a sniffer, and can evaluate a component for leaks more quickly than with a sniffer, it is generally acknowledged that optical gas imaging (OGI) leak detection is about 10 times more efficient than sniffer-based detection [8].

In 2008, the EPA approved an alternative work practice, effectively an amendment to Method 21, allowing the use of OGI in place of Method 21’s use of gas sniffers [9, 10]. In 2016, the EPA took this even further and issued new rules [11] requiring the use of OGI as part of a facility’s LDAR protocol. However, while infrared cameras have proven useful in detecting leaks, their use in *quantifying* leaks has only recently been analysed, and is the subject of ongoing research.




Figure 1. An example detection using an infrared camera viewing a propylene gas release, with the gas column density overlaid onto a visible camera image.

What we show in sections 6 and 7 below is that quantification with infrared imaging is now viable, but that users—and algorithms—need to be better trained to understand how to achieve the conditions necessary for good results. Finally, while the algorithms and methodology are still too new to make their way into the EPA’s rules, researchers are busy developing methods to quantify gas leak rates using infrared camera video data.

While the discussion above has centred on regulation developments within the United States, European regulatory agencies are following a similar trend, with several countries still requiring sniffing, and other countries allowing either OGI or sniffing to satisfy monitoring requirements [12].

2.1. Emissions monitoring

In contrast to the detection-centric goal for safety LDAR, emissions monitoring focuses on quantifying gas emissions. These two goals are of course closely connected, since one cannot quantify what one cannot detect, but they are also less closely related than one might at first expect. For autonomous detection, an essential part of the task is to keep false detections to a minimum, and for this it is necessary to use tight constraints on the detection algorithm. For a system aiming for accurate *quantification*, however, false detections are a lesser concern if they do not have a strong impact on the total amount of gas detected. A system can often improve emissions measurement accuracy by using looser constraints on the detection algorithm. A second difference in the two modes of operation is that emissions monitoring takes place over a much longer timeframe than will an algorithm focused on safety and locating leaks. Thus, monitoring timescales are often days or months as opposed to the second- or minute-response times of safety monitoring.

When quantifying gas emissions, users are often attempting to characterise the efficiency and environmental impact of operations. For safety and LDAR, intentional emissions are

treated very differently than fugitive emissions, but for emissions monitoring both categories are treated on an equal footing. The difference is significant because the two categories have been shown to be similar in magnitude for industrial chemical processing plants. [13]

Emissions monitoring suffers from an even larger technology gap than gas leak detection does. Because existing tools for estimating the quantity of gas emissions are largely human-driven, it is labour-intensive and costly to monitor variable and intermittent emissions sources. The recent literature on emissions from natural gas production, for example, contains a lot of hotly contested numbers, stating that conventional EPA estimates [14] are excessive [15], about right [16], or underestimate total emissions [17, 18], depending on the source [19]. While a significant portion of the controversy arises from differences between ‘bottom-up’ and ‘top-down’ measurement methods, in which top-down methods typically estimate much larger emissions, it is clear that better measurement methods are needed for resolving these estimates. While infrared imaging cameras have had a significant impact on reducing fugitive emissions, their impact on emissions quantification so far has been limited.

2.2. Leak rate estimation

The third mode of operation in gas sensing is estimating the volumetric flow rate of a leak, in which an engineer wants to characterise the size of an already detected leak in flux units (volume per unit time) such as L/min or g/s. While leak rate estimation is closely connected to emissions quantification, it differs from the other modes in that current technology for doing this is already quite accurate. The problem to solve, therefore, is not to improve the accuracy of existing methods but to make them more practical and less labour intensive.

The established method for measuring leak rate is to wrap the area of the leak in an air-tight enclosure (‘bagging a leak’) and connect a flow rate meter and gas concentration sensor to the enclosure outlet. The problem with this standard method, however, is that it is labour intensive, and it also directly exposes the engineer to the leaking gas, which is a safety problem. An automated solution that can avoid the safety hazard is a goal that the industry has been pursuing.

3. Autonomous methods for gas sensing and quantification

Although infrared imaging has been shown to be a more efficient method for LDAR than sniffing, it remains dependent on the person viewing the camera display to evaluate whether or not a gas leak is present. Thus, these methods are still problematic for tasks such as continuous monitoring or detection of intermittent leaks. *Autonomous* methods do not rely directly on human labour, and thus allow continuous monitoring capabilities that would be prohibitive with existing protocols. This can involve using existing sensors in new ways, such as installing sensors on mobile robot platforms, or developing entirely new sensor types, and require measurement methods

that are robust enough for users to depend on them to make important decisions.

There are four types of sensors currently used for gas detection and quantification: point sensors, line sensors, infrared imagers, and snapshot infrared spectral imagers. While each technology has its own strengths and weaknesses, such that an autonomous detection and monitoring framework that combines multiple sensor types seems likely to develop, snapshot infrared spectral imaging is an essential tool to make autonomous approaches practical and effective.

3.1. Point sensor networks

Point sensors are gas detectors that analyse gases passing directly through them, with the most common sensor types being chemi-resistive and infrared [20]. Chemi-resistive sensors use a film whose electrical resistance changes in response to chemical environment changes, while infrared optical sensors measure the absorption signature of gas passing through a small chamber. Point sensors themselves and their associated electronics are typically quite small, but satisfying the strict safety specifications required by many industrial facilities means that they are packaged in durable explosion-proof casings that makes them larger and heavier than one might expect.

Point sensors report the detected gas in concentration units (parts per million, or ppm), have typical response times on the order of seconds to tens of seconds, and can often detect concentrations of < 1 ppm, depending on the technology used and the type of gas detected [21, 22]. They typically cannot discriminate among different gas types, and generally have different response amplitudes for different gases, so that the gas type must be known *a priori* for accurate concentration measurement. Point sensors used in fixed emplacements detect the gas passing through their immediate vicinity, while sensors attached to workers’ uniforms detect gas present in the area through which a worker walks. They have also been attached to robots or moving platforms [23].

Because point sensors detect the gas concentration at one location, by themselves they cannot distinguish between a low-level leak nearby, a large leak far away, or even the direction of the leak. Industrial chemical processing facilities have point sensors distributed at key locations, such as in figure 2, so it is natural to try networking the sensors together to localise leaks. With multiple sensors operating as a network, localising the leak becomes more effective as the density of the network nodes increases. The complex 3D behaviour of gas plume and natural wind motion, however, make this a challenging inverse problem. Even at short distances from a leak, point sensor detections can be erratic. Averaging the detected concentration over longer measurement periods effectively smooths out the most complex details of the motion, improving the problem’s invertibility [24–26], but any sensor layouts near leak sources still require fluid dynamical models to determine the relationship between air flow patterns and detected concentrations. Chraim *et al*, for example, demonstrated experimentally how a sensor grid with nodes spaced 4 m apart allowed for localising leaks to within 5 m of the source, for measurements averaged over a 100 s period [27]. If we consider an



Figure 2. Conceptual image of a point sensor network at an industrial site, where each red circle indicates the location of a sensor.

example facility that is 1 km² in area, however, this spacing requires a network of 62,500 sensors, and that is only for the two-dimensional area of the facility and so ignores the need for some vertical stratification of nodes as well, since most gases are significantly lighter or heavier than air.

The most common way of getting out of these measurement difficulties is to change the scale of the model in both space and time. By averaging over long periods of time, and by analysing the concentration far downwind of the leak—far away from the disturbing effects of nearby equipment on airflow—the time-averaged behaviour of the gas mixing with the ambient air comes to approximate a kind of diffusion transport under certain atmospheric conditions. Diffusing fluids have well-behaved Gaussian distributions of concentration that are easily inverted by estimating a small set of model parameters. This is the Gaussian plume model of gas emission [28].

In a Gaussian plume model, the 3D gas concentration distribution $\rho(x, y, z)$ downwind of a constant leak source can be constructed for either smooth wind flow or turbulent flow, by incorporating different diffusion constants [29, 30]. In smooth wind flow, ρ represents an instantaneous snapshot of the gas concentration, while in turbulent flow ρ represents a *time average* of the gas concentration, with the decorrelation time of the plume flow giving the minimum length of time to average over. Appropriate time scales range from 10 min to 10 hr depending on the extent that turbulence causes changes in wind direction [25, 31]. Figure 3 shows a plume model overlaid on an example gas release.

A widely-used version of the Gaussian plume model gives the downwind gas concentration ρ (g m⁻³) by [31]

$$\rho(x, y, z) = \frac{\phi}{2\pi u \sigma_y \sigma_z} \exp\left(-\frac{y^2}{2\sigma_y^2} - \frac{[z-h]^2}{2\sigma_z^2}\right) \quad (1)$$

where x is the distance downwind from the leak source, z is the vertical displacement from the ground, ϕ is the gas leak rate (g/s), and u is the x -direction wind speed (m s⁻¹). The plume concentration profile width parameters are given by $\sigma_y(x)$ and $\sigma_z(x)$ in the horizontal and vertical directions, while h is the vertical position of the middle of the plume. While this simple




Figure 3. Gaussian plume model.

model of (1) does not account for gas buoyancy, it can easily be adapted to include it [32].

With the Gaussian plume model and knowledge of the instantaneous (for smooth flow) or time-average (for turbulent flow) wind speed u and direction together with the general atmospheric conditions, we can take any three or more gas concentration measurements from point sensors within the plume to determine the (x, y, z) location and intensity ϕ of a leak from the Gaussian parameters h , σ_y , and σ_z . While this means that only three sensors, or a single sensor transported to three locations, are in principle needed to localise and quantify a leak, the difficulty of accurately satisfying the model assumptions and the extended time required to average out turbulence effects can be problematic. Attempts to verify measurement accuracy have run into difficulties with model specification (e.g. defining the dependence of σ_y , and σ_z on distance x) and the need to accurately measure low concentration values so far from the leak source [33, 34].

As implied by figure 3, the Gaussian plume model is only appropriate when the air flow is unobstructed, so that the model is not useful for locating leaks within the perimeter of complex industrial facilities but only far from leak sources. Large-scale plume vortices may even dominate dispersion of plumes 5 to 10 km downstream of the source [35]. Closer to the leak source, more accurate fluid dynamical models of the plume are needed, constructed from knowledge of the 3D shape of the structural environment [22, 36].

3.2. Line sensors

Line sensors are the second type of gas detection technology. A typical line sensor system uses a laser light source and a photodiode detector, with the laser tuned to a wavelength matching a strong absorption peak of the target gas. The detector commonly uses a narrowband spectral filter so that it is insensitive to any light other than the illumination laser. For gas detection, the laser passes along a line through the region to be monitored, at the end of which is either the detector or a retroreflector that returns the laser back towards the detector (i.e. bistatic or monostatic configurations) [37].




Figure 4. Three types of gas sensing instruments: point sensor, line sensor, and imager.

A modified implementation of the line sensor involves using two lasers, or a single laser emitting two wavelengths, where one wavelength is tuned to a gas absorption peak while the other is tuned to an absorption minimum. This helps to distinguish between absorption by the target gas and absorption by an interference agent such as steam or dust. Tunable diode laser absorption spectroscopy (TDLAS) uses a similar technique of tuning the laser wavelength across the gas' absorption spectrum profile [38].

Line sensors have a significant advantage over point sensors in that a single instrument can monitor for gas along all points intersecting the laser, rather than at just one location. The use of specific gas absorption wavelengths makes the line sensors robust to false detections, while also maintaining a high sensitivity. A major difference between these instruments and point sensors, however, is that they use a less familiar unit of measure. Rather than gas concentration ρ (units of m^{-3}), the unit of measure is the gas column density ζ (units of m^{-2}) — the gas concentration integrated along the laser path length,

$$\zeta = \int \rho(z) dz,$$

where z is the propagation direction of the laser. While the column density units can take more effort to get used to than concentration does, ζ is actually more useful than ρ for leak rate quantification, as shown in later in section 6. (The symbol ζ is appropriate for the column density in that its shape somewhat resembles both a 'c' and a 'd' for 'column density'.)

3.3. Gas imaging techniques—active sensing methods

The third gas sensing technology is gas imaging, implementation methods for which can be separated into two broad categories of active and passive techniques: either requiring a specialised light source or using the natural environment for measurement [39]. Like line sensors, active imagers generally use a laser for illumination, but scan the illumination beam across the field of view in order to estimate the spatial distribution of gas. Thus, many are scanning-based extensions of line sensing methods.

Differential absorption LIDAR (DIAL) uses a pulsed tunable laser operating at two discrete wavelengths, one at a strong absorption peak and the other at an absorption minimum. Using the pulses for time-gating, the DIAL system looks for the backscattered signal and uses the return time to determine the distance to the gas. Thus, unique among the gas imaging techniques discussed here, DIAL's unit of measure is the gas concentration ρ rather than its integral $\zeta = \int \rho(z) dz$. An example DIAL instrument [40] was used to measure the gas concentration in a $50 \text{ m} \times 500 \text{ m}$ vertical plane over a period of 5–7 min.

Backscatter absorption gas imaging (BAGI) illuminates a scene with laser light at a wavelength tuned to an absorption peak of the target gas, analysing changes in backscattered light from the scene for presence of gas [41, 42]. When the laser wavelength is strongly absorbed by the target gas, the gas plume thus appears as a black cloud against the more brightly illuminated equipment background. This makes for a compact imaging instrument, but suffers from reduced light intensity, and therefore sensitivity, in comparison to the line sensing methods.

The scanning laser illumination methods of active gas imagers means that they are generally too slow to provide video-rate measurements, even when using high-power laser sources. As a result, they have difficulty measuring the wisps and swirls of turbulent gas that occur at small scales, so they are best suited to large-scale leaks or smooth flow conditions.

3.4. Gas imaging techniques—passive sensing methods

Passive imaging instruments use only the natural signals from the environment to detect gas plumes. The basic principle that they all employ is that when gas passes between the camera and the background scene, it reduces (or enhances) the radiance of the background if the gas is cooler (or warmer) than the background [43–45]. This difference between gas and background radiance is generally referred to as the 'thermal contrast' or 'radiance contrast'. While this measurement principle has the advantage of not needing an artificial light source, it has the disadvantage that it must rely on sophisticated algorithms to separate natural changes in the scene radiance from any changes due to gas. Sections 4–7 below discuss the theory in detail and show example experimental data.

Gas correlation imaging is a method that has been around for over two decades now, and works by viewing a scene

simultaneously through one spectral filter and one transparent window, either using a split pupil or two separate cameras [46–48]. The spectral filter is constructed as a cylinder filled with the target gas—an effective way to optimise measurement sensitivity, since the filter transmission function is exactly matched to the target gas. The camera behind the filter experiences a larger relative change in signal when the target gas appears than will the camera with the broadband filter, and this difference can be used to calculate the column density of gas present. Sandsten *et al* showed example 136×136 images of methane, ethylene, and ammonia captured at 15 Hz frame rates [47].

Filtered-camera gas imaging has the simple hardware setup of inserting a narrowband spectral filter in front of a cryo-cooled thermal camera [13]. While the narrowband filter is not fully optimised to the target gas, as it is in correlation imaging, this method requires only one camera/pupil and works for a wide range of gases. A viewer watching the scene through one of these cameras looks for changes in the scene radiance to indicate the presence of gas moving across a scene. The simplicity and effectiveness of this approach has made it the most popular choice for LDAR use, to the point that many now refer to it as ‘smart LDAR’.

A disadvantage of the filtered-camera approach is that it is not gas-specific, and thus can only be used for quantification if the gas species is known *a priori*. An alternative is to allow the camera to use one of a number of different filter bands [49]. Since gas spectra are more distinct in the longwave spectral band (LWIR, 8–12 μm) than the midwave range (MWIR, 3–5 μm) usually used by filtered cameras, this approach operates in the former range, and employs a filterwheel to select among a set of different spectral filters for a specific scene.

While these passive methods have a large speed advantage over active imaging techniques, they generally have difficulty in accurately estimating *absolute* column densities rather than relative column densities. The measurements are typically obtained as the difference in radiance between the current scene and a reference scene. In order to estimate absolute column densities, these systems have to make assumptions about the gas present in the measurement and reference scenes—assumptions which can make the measurements susceptible to error.

Imaging Fourier transform spectrometry (IFTS) uses an imaging Fourier-transform spectrometer to measure the infrared radiance spectrum at each pixel of the scene [50]. The instrument spectral resolution is typically chosen as a tradeoff of spectral information detail for measurement speed, so that a user can reduce the speed of measurement in order to get high spectral resolution, and therefore higher sensitivity to gas. A strength of this approach is that the detailed spectrum allows an algorithm to estimate the quantities of multiple species of gases simultaneously. This makes measurements robust against interference gases such as steam, and allows the quantification of gas mixtures, even when the gas mixture proportions vary from pixel to pixel. Sabbah *et al* demonstrated the use of an imaging FTIR set to 4 cm^{-1} resolution over the $1050\text{--}1150\text{ cm}^{-1}$ spectral range (i.e. 8.3 nm average resolution over $8.70\text{--}9.52\text{ }\mu\text{m}$) imaging methane emitted from industrial

stacks and also a large plume of SF_6 gas viewed from a distance of 800 m [51]. At this resolution, however, the imaging rate achieved was only 1 frame per 15 seconds (0.067 Hz).

3.5. Gas imaging techniques—snapshot spectral imaging

Snapshot spectral imaging is the newest among gas sensing techniques. It consists of a method that augments the passive infrared imaging techniques with recently-developed instrumentation to measure the infrared spectrum at each pixel of the scene in real time. [52] This maintains the ability to image scene dynamics at video rate, while also keeping the robustness advantages of capturing the spectrum. Since smaller gas plumes are inherently dynamic, with rapidly moving swirls and eddies, this video measurement capability is essential for better sensitivity and locating smaller leaks. Moreover, since many industrial facilities are permeated with steam, and with pipe structures placed in dust-prone environments, using the radiance spectrum to remove false detections due to these interference agents provides a more robust platform for autonomous detection. Finally, as with IFTS, having the spectrum at each pixel in the scene also allows the user to discriminate between gas types and to measure gas mixtures.

The recent development of snapshot spectral imagers takes advantage of the confluence of four modern technology advances: fast computing, large-pixel-count detector arrays, high-speed data transmission, and precision manufacturing. Each of these advances has been essential to making the technology work, so that it is only in the past decade that snapshot spectral imaging methods have become practical for industrial sensing. Although these systems can be difficult to build, they basically operate like N different gas imaging cameras, all imaging the same scene at the same time, using N different spectral transmission bands. Thus, unlike with scanning systems such as IFTS, the snapshot approach does not have to sacrifice light efficiency in order to collect spectral data, and collects the data at a much faster rate, though at a lower spectral resolution than IFTS [53]. Additional advantages from an algorithmic point of view are that because the data is collected in snapshot form, there are no motion artefacts due to scanning over a moving scene, and algorithms can make use of the a complete 4D (x, y, λ, t) dataset for measuring spatio-spectro-temporal correlations to detect gas. For example, a gas plume emanating from a leak will produce a distribution of gas near the leak that will, in the next frame of the video, expand out into the neighbouring pixels. Therefore the spatial elements in one frame are closely correlated to nearby spatial elements in the subsequent frame—a *a priori* information that algorithms can put to good use to improve detection sensitivity.

Currently the only snapshot spectral imaging systems actively deployed in industrial gas sensing are based on uncooled longwave infrared detectors, so that they have different measurement characteristics than the midwave (3–5 μm) filtered cameras used in ‘smart LDAR’. Midwave cameras view the ‘functional group region’ of gas absorption spectra, where many hydrocarbon gases have similar spectral shapes. Cameras that work in the longwave (8–12 μm) spectral region, however, operate in the ‘fingerprint region’ for these gases.

In this region, the spectral shapes of the various hydrocarbon gases are more distinct from one another, so that speciation is more robust.

Two driving factors determine that there will be large differences between practical snapshot spectral imager detection and quantification algorithms and those of the existing literature: the need for autonomous measurement, and the need to operate on a live video stream. The former precludes methods that require user input and supervision, while the latter precludes any algorithm that cannot operate live on a data stream of >5 megavoxels per second. In addition, the need to collect video-rate data in real time precludes the use of high-resolution spectroscopy, as this both increases the data rate even further and reduces the signal-to-noise ratio of the data.

3.6. What each technology means for autonomy

Autonomous sensing for gas LDAR and emission monitoring has long been a goal for instrumentation developers, but has remained out of reach. For point and line sensors, the sensing regions are often too sparsely sampled to give sufficient information for autonomous action. For infrared imagers, the primary difficulty has either been the vulnerability to interferences such as steam, or failure to discriminate among gas species. For scanning-based spectral imagers such as IFTS, it has been difficult to overcome the low data rate and poor SNR. Snapshot spectral imagers have a much higher light throughput [54] than traditional scanning methods can achieve, and allow for video-rate imaging [55]. Since gas clouds are dynamic phenomena and require video analytics to properly detect them, these are important advances for improving sensitivity, reducing false detections, and accurately quantifying the gas.

The discussion so far has focused on fixed emplacements of point, line, and imaging sensors, but another option is to attach sensors to mobile platforms—vehicles, mobile robots, or facility workers themselves—so that the gas leak detection can be pursued *actively* [56]. Mobile sensor platforms can use chemotaxis, the principle by which animals locate the sources of scents for food etc to locate gas sources by moving in the direction of highest concentration gradient [27, 57]. Future facilities may see a network of mobile detection robots deployed, each monitoring for leaks and, upon detection, collectively searching for leak sources.

Mobile options for LDAR have not yet seen much development, possibly due to the safety difficulties of operating moving equipment in strict safety conditions and potentially flammable environments. However, a number of researchers have developed mobile platforms for emissions monitoring and quantification, such as by attaching a series of point sensors to a truck that passes through a thin-concentration (non-flammable) gas cloud [58, 59].

4. Gas imaging detection model

Infrared gas imaging uses radiance changes in a scene to look for absorption or emission signatures generated by gases




Figure 5. The measurement geometry for a single pixel in a gas cloud imager: the camera line of sight views the background infrared radiation (1) through a gas cloud layer (2) and a foreground atmospheric layer (3). Each layer has a spectral radiance L , transmission τ , and temperature T .

passing between the camera and the background [60]. We can model the effect of a gas cloud passing across a scene with a three-layer radiative transfer system (figure 5), in which

- Spectral radiance $L_b(\lambda)$ is generated within a source region which can be either an opaque object such as the ground, or the atmosphere itself, such as when viewing a cloudless sky, or a combination of the two.
- The source spectral radiance traverses the gas cloud layer, and is attenuated/increased by absorption/emission of gases located there.
- The radiation passes through an atmospheric layer to reach the camera.

This three-layer model does not require any assumptions on the background spectrum $L_b(\lambda)$ behind the gas cloud [61].

In the discussion below, we use the following variable definitions:

L_f, L_g, L_b	spectral radiances originating from foreground, gas, and background layers
M	radiant flux at the camera pupil
τ_f, τ_g	transmission of foreground and gas layers
ϵ_f, ϵ_g	emissivity of foreground and gas layers
T_f, T_g	temperatures of foreground and gas layers

The absorption spectrum of a gas is defined by

$$\begin{aligned} \alpha_g(\lambda) &= 1 - \exp \left[- \left(\int_0^\ell \rho(z, \lambda) dz \right) \right] \\ &\approx \sigma(\lambda) \rho \ell, \end{aligned} \quad (2)$$

where ℓ is the path length through the gas cloud. The approximation used here assumes that the gas cloud is homogeneous (so that the integral becomes a simple multiplication) and that the absorption is small ($\alpha \ll 1$) so that the exponential can be approximated as $e^{-x} \approx 1 - x$ so that $\alpha \approx x$. This is the ‘thin gas’ approximation.

For a pixel in the scene observed by the camera, we can write the radiative transfer equation of a ray along the line of sight to give an at-pupil radiative flux M of

$$M_w = \int_{\lambda_1}^{\lambda_2} \left[\epsilon_f(\lambda) B(T_f, \lambda) + \tau_f(\lambda) \epsilon_g(\lambda) B(T_g, \lambda) \right]$$

$$\begin{aligned}
& + \tau_f(\lambda) \tau_g(\lambda) L_b(\lambda) \Big] d\lambda \\
& \approx \Delta\lambda_w \left[\epsilon_{f,w} B_w(T_f) + \tau_{f,w} \epsilon_{g,w} B_w(T_g) \right. \\
& \quad \left. + \tau_{f,w} \tau_{g,w} L_{b,w} \right]. \quad (3)
\end{aligned}$$

where λ_1 and λ_2 give the passband of the w th spectral channel, and a w subscript indicates an average across the passband. Also, $B(T)$ is the Planck blackbody spectral radiance function for temperature T , τ is the spectral transmission of a layer, and ϵ the spectral emissivity of a layer. For a broadband imager, λ_1 and λ_2 are given by the spectral range of the optical filter on the camera. Using Kirchhoff's law [62], we can equate the emissivity spectrum with the absorption spectrum, $\epsilon = \alpha$ and further relate the absorption and transmission spectra, $\alpha = 1 - \tau$. Thus, using (3), we can model the radiance measured with (M') and without (M) gas present as

$$\begin{aligned}
M_w &= \Delta\lambda_w \left[(1 - \tau_{f,w}) B_w(T_f) + \tau_{f,w} L_{b,w} \right], \\
M'_w &= \Delta\lambda_w \left[(1 - \tau_{f,w}) B_w(T_f) + \alpha_{f,w} \tau_{g,w} B_w(T_g) \right. \\
&\quad \left. + \tau_{f,w} \tau_{g,w} L_{b,w} \right].
\end{aligned}$$

The difference $\Delta M_w = M'_w - M_w$ gives

$$\Delta M_w = \tau_{f,w} \alpha_{g,w} \Delta\lambda_w [B_w(T_g) - L_{b,w}]. \quad (4)$$

Figure 6 shows $\alpha_g(\lambda)$ (green curve) for ammonia gas, together with an example $M(\lambda)$ (red curve) and $M'(\lambda)$ (black curve) for a 30°C gas passing in front of a background blackbody at 50°C. Re-arranging (4) to solve for the gas absorption, we obtain

$$\alpha_{g,w} = \frac{M'_w - M_w}{\tau_{f,w} \Delta\lambda_w [B_w(T_g) - L_{b,w}]} \quad (5)$$

The denominator here is a difference of radiances (the 'radiance contrast') $\Delta L = B_w(T_g) - L_{b,w}$ that is closely related to the thermal contrast $\Delta T = T_b - T_g$. The sign of the denominator indicates whether the cloud is observed in emission ($\Delta L > 0$) or absorption ($\Delta L < 0$), but the estimated α_g using (5) is positive in both situations. Note that the denominator can come close to zero, so that the estimate for α_g should be regularized, for example by adding a small constant value to the denominator, or by using a Wiener filter.

If the reference measurement M contains no gas, we can easily solve for the background radiance $L_{b,w}$ in each spectral channel:

$$L_{b,w} = \frac{1}{\tau_{f,w}} \left(\frac{M_w}{\Delta\lambda_w} - [1 - \tau_{f,w}] B_w(T_f) \right). \quad (6)$$

It is common to assume that the temperature of the gas and the ambient air are approximately the same, $T_g \approx T_f$, under the assumption that the gas quickly entrains into the local air. Inserting this into (6) and substituting the result into (5) gives

$$\alpha_{g,w} = \frac{M'_w - M_w}{[B_w(T_g) \Delta\lambda_w] - M_w}. \quad (7)$$

Note that $\tau_{f,w}$ has disappeared from the expression: the foreground layer plays no role in the absorption estimate, except




Figure 6. An example measurement for ammonia gas (column density $\zeta = 5500$ ppm.m), showing a simulated measurement spectrum $M'(\lambda)$ (black curve) when the background source is a blackbody at $T_b = 50^\circ\text{C}$ (red curve) and the gas temperature is at $T_g = 30^\circ\text{C}$. The absorption spectrum (green curve) is given at an arbitrary scale, for a thin gas cloud.

insofar as it reduces the measurement SNR. All of the variables on the right hand side of this equation are quantities that we can measure or estimate. Equation (7) gives the gas absorption separately for each spectral channel in the system, from which we can calculate its correlation with the known spectral shape of various library gases. Together with an estimate of the noise, this correlation allows us to calculate the probability that the measured spectrum indicates presence of the gas. If the probability exceeds a selected threshold, the pixel is labelled as a probable detection. The detection at this step is labelled as only probable because a more sophisticated algorithm can next take advantage of spatio-temporal correlations in the gas motion to improve the probability estimate. For example, if gas is detected at a given pixel, then it is likely to be detected in its neighbouring pixels as well, and the same gas parcel is likely to be found nearby in the following data frame as well. Incorporating these spatio-temporal features, and knowledge of gas fluid dynamics, into the detection algorithm produces an iterative procedure that improves the detection sensitivity.

Discussions of (7) typically refer to gas absorption, implying that the gas is cooler than the background source. However, an important case where observation of gas is seen in *emission* is that of a sky background. Not only does the sky occupy a large portion of the field of view in many situations, but it can also have a large radiance contrast. A clear blue sky background can provide over 50°C of thermal contrast for many spectral bands across the 8–12 μm spectral range, if the local humidity is low [63].

4.1. Background-foreground separation methods

The most difficult step in equations (4)–(7) above is to determine the spectrum of the background layer $M(\lambda)$ — a variant of the well-known 'background-foreground separation problem' from the image processing literature [64]. Perhaps the simplest approach for separating background M from foreground ΔM begins by assuming that in a given video sequence the background is present much more often than the foreground at any given pixel. If we simply take a running mean

or running median image, this mean or median can then be used as a rough estimate of the background. The foreground (the gas cloud, in our case) is then obtained by taking the difference between the current image and the estimated background.

A more sophisticated approach may use an adaptive Gaussian mixture model. In this algorithm, at each frame of a video sequence, the spectrum at each pixel is modelled as belonging to either of two Gaussian distributions, representing the background spectrum and gas spectrum. As more frames follow, the estimates of the two distributions improves, with the two means representing the background and foreground spectra at that pixel. This approach can be configured as an adaptive online algorithm that is not computationally expensive, so that the calculation can run in real-time [65].

A third approach assumes that the background is a greybody, so that its spectrum follows a set of easily parametrized curves [66]. Once we fit the data to the nearest greybody curve, we can take the remainder of the spectrum to represent the gas layer. Since there are many situations in which these assumptions do not hold well, researchers have also tried to relax the greybody requirement and instead assume that the background radiance $L_b(\lambda)$ is smooth: $L_b(\lambda)$ varies more slowly with respect to λ than the gas absorption cross-section spectrum $\sigma(\lambda)$ [67, 68]. Making use of this assumption, however, generally requires a significant spectral resolution, making measurement times long. It is also more effective on gases that have narrow spectral features, such as methane and ammonia, than those with broad spectral features, such as propane and propylene [69].

A fourth approach to the separation problem is to decompose the spatial-spectral scene into its principal components, so that if gas pixels comprise only a small number of pixels in the scene, then the primary principal components will be unaffected by gas spectra [70, 71]. Because of the complexity of the decomposition, however, it is hard to predict the reliability of the algorithm for any given scene, and spatially complex structures can cause the algorithm to become unstable. Moreover, scenes in which gas clouds comprise a significant fraction of the image make the algorithm insensitive to gas [72]. This is a dangerous property for an algorithm to have, as it makes the system blind during the most critical situation. In addition, the principal components approach involves a heavy computational load that can cause difficulties for video-rate spectral imaging.

4.2. Gas mixtures and interferents

In the presence of multiple gas species, each of which is weakly absorbing, the overall absorption can be written as a linear sum over the component absorbances of each species [73]. Thus, for N gases, in the thin gas and homogeneous layer approximations,

$$\alpha_{\text{mix}} = 1 - \exp \left[- \sum_{i=1}^N \sigma_i \left(\int_0^\ell \rho_i(z) dz \right) \right] \approx \sum_{i=1}^N \sigma_i \rho_i \ell. \quad (8)$$

(Wavelength dependence in each parameter is left implicit.) If the ratios of the various constituent species within the gas mixture are known and remain constant—a common situation for industrial sites and if the gases are of similar densities—then the gas mixture can be characterised with single effective cross section: $\alpha_{\text{mix}} = \sigma_{\text{mix}} \rho_{\text{mix}} \ell$. Thus, while one may use spectral unmixing [74] and matched filtering [75] algorithms to estimate the column density images of multiple gases simultaneously, this simple solution is sometimes all that is needed.

Another complication of measuring in industrial facilities is the ubiquitous presence of steam and sometimes dust. Although these are aerosols, both of these act as interference gases that change the measured spectral radiance, but which are of no interest to the final measurement. As aerosols, these two interferents also change spectral shape depending on particle diameters [76, 77], so that their spectral profiles belong to families of shapes rather than a single fixed spectrum [78]. If any member of these continua of shapes has a high correlation with one of an instrument's target gases, then discrimination between them becomes difficult using infrared alone. Instruments that employ visible cameras registered to the same scene as the infrared images, however, have a convenient means of separating the two: steam and dust plumes are visible, whereas most target gases are invisible.

5. Gas image column density quantification

Once gas is detected within a pixel, estimating its column density ζ is straightforward. Equation (4) relates the absorption of the gas layer to the measured changes in radiance at the sensor. Using the Beer–Lambert–Bouguer law, the transmission is related to the concentration ρ and path length ℓ through the gas cloud as

$$\tau_g(\lambda) = \exp \left[- \sigma(\lambda) \int_0^\ell \rho(z) dz \right] \approx \exp \left[- \sigma \rho(\lambda) \ell \right]. \quad (9)$$

The absorption cross-section $\sigma(\lambda)$ (units of m^2) is known *a priori* from the type of gas detected, and can be measured directly or obtained from a spectral library, such as the NIST Infrared Database or the PNNL infrared spectral library [79, 80]. We use $\rho(z)$ for the gas concentration (number density) along the line of sight and ρ (with no explicit z -dependence) as the average concentration along the line of sight. Because ρ and ℓ cannot be separated from one another, they are grouped together into a single quantity—the column density $\zeta = \rho \ell$ —the typical units for which are ppm.m (parts-per-million times meter).

When the gas concentration is low (the thin gas approximation), equation (9) can be written as $\alpha_g \approx \sigma \zeta$, and so the gas column density becomes a linear function of the absorption. Going from the estimated absorption to estimated gas column density therefore requires only scaling by the cross-section:

$$\zeta = \alpha_g / \sigma. \quad (10)$$

The column density has SI units of m^{-2} . We can convert this to units of ppm.m by multiplying by the density of molecules




Figure 7. A detector pixel projected to the location of a gas cloud.

in air ($\rho_{\text{air}} = 2.687 \times 10^{-25} \text{ m}^{-3}$ at STP), and scaling the result by 10^{-6} [81]. An example of applying this estimation to each pixel in an image is shown in figure 1 for a controlled methane gas leak.

5.1. Estimating the total quantity of gas in a plume

From an image of the gas column density, we can also estimate the total amount of gas in a cloud. That is, if we scale each pixel to the projected area A_{proj} at the location of the gas cloud (figure 7), then multiply the projected area by the detected column density, we obtain the total volume of pure gas in the pixel (units of L, m^3 , or ft^3):

$$V_{\text{gas}} = A_{\text{proj}} \zeta.$$

The projected area of a pixel at the gas cloud is easily obtained by multiplying the known dimensions of the detector array pixels (p_x, p_y) by the image magnification, given by the lens focal length f and the distance z to the gas cloud:

$$A_{\text{proj}} = (p_x z / f) (p_y z / f). \quad (11)$$

If we wish to estimate the mass of gas within the cloud, we can first calculate the total number of molecules N in the cloud by multiplying V_{gas} by the number of molecules per cubic meter ρ_{air} at the ambient atmospheric temperature and pressure. The total mass of the gas cloud (g, kg, or lbs) is then obtained by multiplying N by the gas molecular weight.

While the pixel size and the lens focal length are known as a result of the hardware design, the gas cloud distance z from the camera varies and must be estimated. One method for doing this is to establish distances when setting up a gas cloud imaging camera at a facility. After fixing a camera's position, the installer measures the distances from the camera to the primary items to be monitored—items such as a collection of pipes, the wall of a gas container tank, the relief valves of a gas separator tank, etc. Thus the distance at which leaks are likely to occur can be established *a priori*. Other methods include using triangulation with multiple cameras, or using fluid dynamical models to compare estimated speeds of gas motion with measured speeds.

This ability to directly estimate the total volume or mass of a gas cloud is one of the advantages of measuring in ppm.m

units instead of concentration. From the total volume, we get a view of the overall scale of the leak, which helps to decide whether to issue a safety alarm. This can be something as simple as logging a measurement when the volume of gas is small, issuing a warning when the gas cloud is large, and causing an alert when the cloud has reached a dangerous size.

5.2. Estimating the volumetric gas concentration

From a safety perspective, it is important to know the volumetric concentration $\rho(x, y, z)$ of a flammable gas cloud, as it indicates whether the gas is capable of igniting. Methane, for example, ignites if in concentrations of between 5% and 15% (the lower and upper explosion limits, LEL and UEL), when mixing in air at STP [82, 83]. A cloud of methane containing concentrations everywhere below the lower limit will not ignite, even when exposed to an ignition source such as a spark or flame. Clouds of methane gas that are too rich (above the upper limit of 15%) also do not ignite. However, high-concentration clouds quickly diffuse into surrounding air, so that their outer edges contain a flammable border region. Thus, high-concentration plumes are a safety hazard.

For estimating concentration, point sensors have an advantage in that they measure concentration directly, but have the disadvantage that they can only sample the concentration at one point. While a point-sensor measurement giving a value above the lower explosion limit directly indicates a hazard, a value below the LEL is ambiguous. Without a physical model for the gas cloud, one cannot know whether the measurement is low because the overall cloud concentration is low or because the leak source is too far from the sensor.

Although gas imagers detect in units of column density ζ rather than concentration ρ , they can obtain rough estimates of the concentration distribution within the gas in real-time using heuristics, as we describe below. Imagers can also form accurate estimates using fluid-dynamical models of the detected gas cloud [84, 85], or by simultaneously collecting projections through the cloud from multiple points of view [86, 87], but these methods are typically too computationally intensive to run in real-time. In practice, a rough estimate (via heuristics) can be useful as a quick initial guess for the concentration distribution $\rho(x, y, z, t)$.

The first step in the heuristic method is to model the gas cloud measurement as the projection of a cylindrical shape with its axis oriented in the direction of the gas flow, as shown in figure 8. If the axis of the gas plume motion is purely transverse (does not have a component along the view direction, z), then the estimated depth dimension of the plume becomes the same as its height dimension. For example, in figure 9, the transverse gas motion allows us to say that, at the position in the plume indicated by a red border, where the height of the gas cloud is 0.52 m (23 pixels in the image), we can estimate the depth of the plume to be approximately 0.52 m.

The second step involves a weighted backprojection of the 2D data into the 3D volume, using a Gaussian-shaped weighting function such that the center of the cloud has a higher estimated concentration than the edges of the cloud. Figure 8 shows the projection of a 1D slice from the data into a 2D




Figure 8. The backprojection of a column density gas measurement into a cylindrical concentration volume. The vertical column of data is taken from a slice through the estimated gas cloud shown in figure 9(b).

volume in order to demonstrate the procedure clearly. One can tune the degree to which the concentration is focused in the center versus the periphery by adjusting the width of the Gaussian function which sets the relative weights. The resulting tubular volume gives an estimate for the concentrations within the cloud, where integrating through the volume reproduces the column density image.

Because the concentration estimates using the above approach employ simple heuristics rather than a physics-based model, the results give only a rough estimate of the concentration. However, it has the important advantage that it can operate in real time. For a more accurate model, it is necessary to build a 3D fluid-dynamical model that is constrained to give the same projection through the cloud as was measured by the camera. Doing this accurately under the weak constraint of a single measured image is difficult, but the model can be improved through the use of multiple frames in a video sequence. Getting this to work well is an ongoing subject of research [87].

6. Estimating the volumetric flow rate of a leak

One of the primary advantages of a video-rate imager used for gas sensing is that its sampling is complete in the 3D (x, y, t) space of a video sequence. Thus, pixels in one frame are neighbours of the same pixels in the subsequent frame of the video, so that it becomes possible to design algorithms analysing spatio-temporal correlations within the gas cloud. If a gas is found in one area within the scene, there will be a high probability of finding it in the same place, or close nearby, in the following frame of the video. Thus, temporal correlations can be used to improve the sensitivity of the gas detection algorithm.

Another application for using these spatio-temporal correlations is estimation of the volumetric flow rate—the gas flux or leak rate. Existing algorithms for doing this all incorporate some variant of optical flow techniques to estimate how each parcel of gas moves within the image. That is, multiplying the volume of a parcel of gas times its average rate of flow away from the leak source, and summing over all parcels of gas detected, gives an estimate of the total flux [50, 88, 89].




Figure 9. An example gas flow measurement for methane released from a pipe through a gas flow controller: (a) a simplified physical model for the jet and plume components in the flow; (b) a snapshot of one frame in the gas column density video; (c) the corresponding gas velocity vectors. The red and blue squares in (c) indicate two example locations for placing a boundary used to evaluate the flow rate. The red rectangle shown in (b) indicates the region used in figure 8 as an example concentration backprojection.

When using a pixel-by-pixel estimate of the flow rate, one way of verifying that the algorithm is behaving consistently is to use nested boundaries, as shown in figure 9(c). Since each boundary contains the leak source, each should produce the same time-averaged leak rate.

While we focus on imaging methods for estimating the flow rate, point sensor networks employing Gaussian plume models of the leak can also be used for estimating flow rate [29]. However, these have only been demonstrated as useful for large leaks and smooth flow conditions. In order to expand the range of application, Safitri *et al* demonstrated that one can combine the Gaussian plume modelling with single-frame infrared gas imaging to quantify leak rate [90]. While this allows improved modelling of smaller leaks, it still overconstrains the profile and dynamics of the gas, affecting accuracy. Although it has not yet been demonstrated, video gas imaging should be able to significantly improve on this. By averaging over a long time-sequence of gas detection frames, the data comes closer to the time-average approximation required for Gaussian diffusion.

In the discussion below, we introduce the conventional optical flow algorithm and then show why it is inadequate for most gas imaging data. Fixing the algorithm requires using a more sophisticated implementation that constrains the flow estimate with a fluid-dynamical model.




Figure 10. (a) A camera measuring the light coming from a fluid distribution (e.g. gas cloud). (b) The integral along a line through a fluid of concentration ρ gives a value at the image in column density ζ . The average velocity \bar{u} along the line of sight is derived from a weighted sum of the individual velocities u_i .

6.1. Principles of optical flow algorithms

Optical flow algorithms are based on the concept that when an object moves across the image, its brightness I at each point stays the same. This is often referred to as the ‘brightness constancy’ assumption [91]. In the context of measuring gases, the ‘brightness’ represents the gas concentration, in which case it is more appropriate to call this assumption the *continuity condition*, as it enforces the conservation of mass within a local neighbourhood. The continuity condition allows us to take the brightness within one image frame, $I(x, y, t)$, and track its 2D displacement (u, v) in the next frame by searching for the path satisfying the continuity equation

$$I(x, y, t) = I(x + u \Delta t, y + v \Delta t, t + \Delta t). \quad (12)$$

This path satisfying continuity gives

$$\frac{d}{dt} I(x(t), y(t), t) = 0, \quad (13)$$

which by the chain rule produces

$$\frac{d}{dt} I(x(t), y(t), t) = \frac{\partial I}{\partial x} \frac{dx}{dt} + \frac{\partial I}{\partial y} \frac{dy}{dt} + \frac{\partial I}{\partial t} \frac{dt}{dt} \quad (14)$$

$$0 = u \frac{\partial I}{\partial x} + v \frac{\partial I}{\partial y} + \frac{\partial I}{\partial t}. \quad (15)$$

In vector notation, this is

$$\mathbf{u} \cdot (\nabla \rho) + \partial \rho / \partial t = 0. \quad (16)$$

This expression is commonly called the *gradient constraint equation*, giving the relationship of the displacement vector (u, v) to the spatial and temporal derivatives of the brightness (or gas concentration).

When imaging volumetric fluids, instead of I what we actually measure is the fluid column density ζ . If the output of

the optical flow algorithm is the fluid-concentration-weighted average velocity \bar{u} along the line of sight [92–95], as in figure 10, then we can simply multiply \bar{u} with the total volume V of fluid within a pixel to give a point-by-point estimate of the fluid flux $\phi = \bar{u} V$. Summing this over all pixels along a boundary gives the total amount of fluid flow across it,

$$\Phi = \sum_{x,y} [\bar{u}(x,y) \cdot \mathbf{s}] V(x,y),$$

where \mathbf{s} is the unit vector normal to the boundary, so that $\bar{u} \cdot \mathbf{s}$ gives the projection of the gas velocity vector across the boundary. If the boundary is *closed* (such as the red or blue boundary shown in figure 9(c)) then Φ gives the estimated overall source flux (leak rate), which is often expressed in units of liters per minute (lpm) or grams per second (gps). An example measurement is shown in figure 11.

6.2. Adapting optical flow algorithms to volumetric flow

In principle, adapting the optical flow algorithms originally developed for opaque objects (i.e. (16)) to measuring volumetric fluids seems straightforward: we simply work with ζ instead of I , and multiply each pixel’s column density by the local velocity estimate to obtain the measured flow rate. This is close enough to the truth as to be confusing—a problem which we will see is further reinforced by an unfortunate choice of nomenclature.

The established method for deriving a representation of projected fluid volumes for optical flow algorithms starts with the continuity equation of fluid dynamics, which establishes the conservation of mass [92, 93, 95]:

$$\nabla \cdot (\rho \mathbf{u}) + \partial \rho / \partial t = \mathbf{u} \cdot (\nabla \rho) + \rho (\nabla \cdot \mathbf{u}) + \partial \rho / \partial t = 0. \quad (17)$$

This closely resembles the gradient constraint equation (16), but with the addition of $\rho (\nabla \cdot \mathbf{u})$ — the ‘divergence term’. Integrating (17) along the line of sight, the continuity equation becomes

$$\int [\mathbf{u} \cdot (\nabla \rho) + \rho (\nabla \cdot \mathbf{u}) + \frac{\partial \rho}{\partial t}] dz = 0. \quad (18)$$

If we consider the divergence term by itself for a moment, we can write

$$\int \rho (\nabla \cdot \mathbf{u}) dz = \int \rho \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) dz. \quad (19)$$

which by use of Leibniz’ rule becomes

$$\begin{aligned} \frac{d}{dx} \int_{z_1}^{z_2} \rho u dz &= \int_{z_1}^{z_2} \frac{\partial}{\partial x} (\rho u) dz + [\rho u]_{z_2} \frac{\partial z_2(x,y)}{\partial x} \\ &\quad - [\rho u]_{z_1} \frac{\partial z_1(x,y)}{\partial x} \\ &= \int_{z_1}^{z_2} \frac{\partial}{\partial x} (\rho u) dz = \int \left(\rho \frac{\partial u}{\partial x} + u \frac{\partial \rho}{\partial x} \right) dz. \end{aligned}$$

Therefore, we can write this result as

$$\int \rho \frac{\partial u}{\partial x} dz = \frac{\partial}{\partial x} \int \rho u dz - \int u \frac{\partial \rho}{\partial x} dz.$$




Figure 11. Example leak rate measurement. (a) A snapshot of one frame from a 10 s long gas detection video sequence, with the gas column density colormap superimposed on an infrared thermal image. The concentric squares indicate possible boundaries at which to evaluate the gas flow. (b) The pixel-by-pixel estimate of the flow velocities (shown with green arrows) superimposed on an image showing the gas column density. (c) The time-resolved leak-rate estimate for the video sequence. The horizontal dashed line indicates the overall mean flux 17.2 L/min, while the black circle indicates the time location of the snapshots shown in (a) and (b).

A similar result follows for the vertical component of the velocity, v . Substituting both expressions into (19) gives

$$\begin{aligned} \int \rho \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) dz &= \frac{\partial}{\partial x} \int \rho u dz + \frac{\partial}{\partial y} \int \rho v dz \\ &\quad - \int u \frac{\partial \rho}{\partial x} dz - \int v \frac{\partial \rho}{\partial y} dz \\ &= \nabla \cdot (\hat{u} \zeta) - \int \mathbf{u} \cdot (\nabla \rho) dz, \end{aligned} \quad (20)$$

where \hat{u} is the concentration-weighted-average velocity,

$$\hat{u} = \frac{\int \rho \mathbf{u} dz}{\int \rho dz}. \quad (21)$$

Inserting (20) into (18), we find that the first term of (18) cancels out, leaving the ‘integrated continuity equation’

$$\nabla \cdot (\hat{u} \zeta) + \frac{\partial \zeta}{\partial t} = 0. \quad (22)$$

Thus, the continuity equation (17) and not (16) should be used as the basis for estimating the concentration-weighted-average velocity \hat{u} .

6.3. Compressibility vs. incompressibility

Fluid dynamics often separates fluid flow into two regimes: compressible and incompressible flow. For gases, the compressible regime generally holds for high-speed motion and

large pressure gradients. For incompressible flow the density remains constant within a parcel of fluid that moves with the flow velocity. The limit between these two regimes for air is generally considered to be the point at which the ratio of the flow velocity to the speed of sound in air (the Mach number) is much less than 1. This limit almost always holds for the case of imaging gas leaks, so that for gas imaging we are clearly in the incompressible flow regime.

Section 6.2 showed that once we include the divergence term in the continuity equation, as in (22), we have a method for analyzing volumetric fluid flow that has been shown to work well and provide accurate quantitative results [84, 96]. However, this has produced confusion among algorithm users, since it is also well-known that the divergence term is negligible in the incompressible regime, and these gas flows are clearly incompressible. The difficulty arises from the fact that we are viewing the mixing flow of two different fluids, one of which is transparent. Whereas ‘incompressibility’ refers to the constraint

$$\rho_{\text{gas}} + \rho_{\text{air}} = \text{constant},$$

the gas imaging data contains only ρ_{gas} . Satisfying incompressibility for the target gas alone would mean that the concentration of a gas parcel does not change as it moves along with the flow. However, since the gas diffuses into the surrounding air as it propagates, the gas concentration is clearly not constant, and so the compressible-flow regime is the appropriate choice.

In order to understand the impact of ignoring this distinction, we can re-derive the expression for the weighted-average velocity, but this time hold the divergence term to be negligible, as shown in A. What we find is that whereas (22) (which keeps the divergence term) produces velocity estimates correctly weighted by the fluid concentration along the line of sight, the no-divergence form produces velocity estimates weighted instead by the *gradient* of the fluid concentration along the line of sight. That is, in a naïve implementation of optical flow that ignores the divergence term (13), any layers having larger spatial gradients, even if they have lower amounts of gas, will play a larger role in the velocity estimate. Additionally, any layer in which the velocity gradient is zero (i.e. spatially uniform concentration) will have no influence at all, even if it contains the majority of the gas in the image. This is a disaster for accurate quantification, as the error differs from pixel to pixel, and from frame to frame, in accordance with the characteristics of the fluid flow. Thus, it is necessary to work with (22) rather than (13).

7. Experimental results and comparison to the existing state-of-the-art

Quantitative gas detection involves measurements of gas concentration (ppm units), gas column density (ppm.m), gas cloud volume (L), or gas leak rate (L/min). Evaluating each of these four measurement modes involves a different type of experiment setup. Whereas evaluating the performance metrics of a point sensor in a lab is straightforward, it can be surprisingly difficult to evaluate point sensors in practical outdoor conditions. Typical point sensor metrics include the quantification accuracy of the gas concentration, the lower limit of concentration at which gas can be detected (the ‘sensitivity’), and the response time. For flammable gases such as methane and propane, sensor measurement precisions are often reported in the range of 1–100 ppm, but calibration fluctuations mean that measurement accuracy is worse than this. The sensitivities of commercial sensors, however, are generally 50–150 ppm, for response times of 10–45 s [97, 98]. Although lower numbers are always better for each of these metrics, it is not a simple process to connect them to practical use metrics, such as the minimum size of a leak that a sensor can detect, or the utility for safety alarms. In these situations, the response time of the sensor plays a role, as does the distance of the sensor from the leak, and local wind conditions.

Evaluating gas imagers faces similar challenges. Infrared gas imaging systems are generally stated as having sensitivities in the range of 500–5000 ppm.m for common hydrocarbon gases [49, 99]. Response times are generally on the order of 1 s. While column density is the basic unit of measurement for these systems, the addition of the gas cloud size as part of the unit can make it difficult to compare these numbers directly with the sensitivities of point sensors. If we assume that the length through a typical gas cloud size is on the order of a meter, then the sensitivities appear to be much worse than point sensors. On the other hand, imagers detect leaks directly

at the source whereas point sensors must rely on the gas reaching the location of the sensor. The concentration is much lower at the sensor and so the corresponding sensitivities have to be lower in order to detect the same size leak. Another complication for comparing the two is that, while the sensitivities of both sensor types also suffer with increasing distance and wind speeds [10, 100], infrared imagers have an additional loss of sensitivity on cloudy days and during twilight or dawn. At these times, the air and exterior surface temperatures are closer together, reducing the thermal contrast needed for infrared measurement. Finally, wind turbulence places a fundamental lower bound on the sensitivity of infrared gas imaging [101].

It is also important to keep in mind that the sensitivities reported by imaging systems are typically for measurements taken over short time intervals. For longer measurement times (such as the 10–45 s response times of many point sensors), the imaging sensitivity can be improved through time-averaging [102]. Although very little of these results have been reported in the literature, the author’s experience has been that one can typically improve the sensitivity by a factor of 2–10, depending on measurement conditions. More favorable conditions such as low wind speeds allow for a larger improvement factor.

Response time effects can cause some unexpected behavior. For example, during the hundreds of quantitative experiments that the author has performed using controlled gas leaks or bags filled with 100% gas concentrations, not once did a wearable personal detector set off an alarm. Although the gas concentrations were orders of magnitude beyond the sensitivity limits of these sensors, the presence of the gas was for too short a period to allow the point sensors to respond.

In addition to time-averaging, autonomous sensors gain another advantage by monitoring a scene repeatedly over a long period: the increased chances of ‘lucky’ acquisition, in which the conditions are just right for achieving the highest sensitivity. This may happen when, for example, the wind changes and begins pushing the gas towards the direction of an infrared gas camera, increasing the size of the gas cloud. As a result, imaging detection of methane clouds has been shown for distances of over a kilometer, despite the high attenuation of methane signals with distance due to water-vapor.

Because of the difficulties of evaluating point sensors and gas imagers using their basic measurement units, a common metric used instead is the minimum detectable leak rate (MDLR) of a system. In order to evaluate this metric, a detection system is deployed in a realistic measurement environment and exposed to various leak rates and various measurement conditions. (Table 1 reproduces the MDLR measurements recorded in [10].) A system with a low enough MDLR under a wide range of conditions, and a corresponding low false positive detection rate, will be useful for autonomous implementation. In addition to the MDLR, we can also evaluate the quantification accuracy of these systems, and this is one area where imaging systems have shown capabilities that point sensor networks have difficulty reaching. For example, in a gas volume quantification test, we can fill a bag of known volume with a gas mixture, and release the gas at once within the imager’s field of view. The imaging sensor can then estimate

Table 1. Laboratory measurements of the minimum detectable leak rates (MDLR) in g/hr for gas imaging based on using a filtered MWIR infrared-camera, at various wind speeds (reported in [10]).

Analyte	0 ms ⁻¹	Wind speed	
		0.9 ms ⁻¹	2.2 ms ⁻¹
Butane	0.72	5.8	15.9
Ethane	0.97	5.2	17.8
Methane	3.96	20.8	49.0
Propane	0.76	9.8	19.1
Ethylene	13.9	53.7	104
Propylene	4.37	15.6	59.8

Table 2. Results of methane leak rate experiments using Rebellion Photonics' GCI (see table 3.24 in [39]).

Viewing distance (m)	Actual flow rate (L/min)	Measured leak rate (L/min)	Error (%)
17	1.67	1.85	10.8
17	10	8.95	10.5
17	16.7	17.2	3.0
17	100	81.9	18.1
30	1.67	1.50	10.2
30	10	11.0	10.0
30	16.7	16.1	3.6
30	100	84.9	15.1

the total gas volume from the gas detection video, and this result can be compared with the bag volume. While this volume-estimation capability has been demonstrated in the field, its accuracy has not yet been published in the literature, and so it is not clear what the current status is of this mode.

The fourth quantification mode is the gas leak rate, typically measured in g/hr or L/min. Evaluating the leak rate quantification accuracy involves placing a flow-controlled leak source into the field of view of an imaging camera (or the center of a point sensor network). Table 2 gives a series of measurements taken with the Rebellion Photonics' GCI camera, for tests supervised by the US Environmental Protection Agency (EPA) [39]. Representative gas leak sizes in equipment fugitive emissions are 0.038 L/min (median) or 24 L/min (mean) [103], so the test conditions of table 2 fit into this range. The results indicate that, under good measurement conditions of low wind, good thermal contrast, and the leak source fully in the field of view, the leak rate can be quantified to within an accuracy of $\pm 20\%$.

We can also see that the quantification errors in table 2 are actually larger than the MDLR results shown in table 1. Thus, at the lowest gas leak levels the quantification has too much error to be useful—while we can detect gas at the lowest levels, a higher signal-to-noise ratio in the measurements is needed in order to provide useful quantification.

8. Conclusions

The three gas sensing applications—LDAR, safety, and emissions monitoring—each have different needs, and sensor

designers are working to fulfill these needs with ongoing development. Point sensor networks, filtered infrared cameras, and snapshot infrared spectral imagers are the three primary detection modes that designers are competing to develop for industrial gas detection and quantification, with the snapshot spectral imagers being the newest among them. While historical usage has focused primarily on detection (the binary decision of gas 'present' or 'absent') and gas concentration measurement, recent development has expanded into new forms of quantification: gas column density, plume volume, and leak rate. We have argued above that imagers provide important advantages for these new quantifiers—especially for allowing remote estimation of leak rate.

Autonomous gas detection systems have the potential to reduce the labor burden for LDAR, and to improve our ability to quantify gas emissions from industrial sites. Camera-based systems are undergoing early deployment, and show promising benefits over the networks of point sensors that currently dominate the industry. While camera-based 'smart LDAR' initially had a reputation for being less sensitive than point sensors, industrial users have become familiar with their different strengths and weaknesses, and camera detection has achieved widespread use in LDAR. However, camera-based systems have made fewer inroads into safety and emissions monitoring applications, due less to any fundamental weakness in their capabilities than to the current immature state of algorithm development. We have provided some preliminary testing information in section 7, but much more testing is needed to adequately characterize performance. As development of these new modes matures in the coming years, we will likely see a widespread change in how industry approaches safety and emissions monitoring.

The newest detection mode, snapshot spectral imaging, also provides an improved ability to discriminate among gases and to remove the effects of interferents such as steam and dust. This allows for an improved false-detection rate that is an essential step in the transition to making a gas sensor system truly autonomous.

Finally, we can also keep in mind that there are ancillary benefits to using infrared cameras for detection. In addition to gas sensing itself, infrared camera data can also be used for other purposes such as intruder detection, monitoring equipment for defects or long-term maintenance issues, or even tracking the liquid levels in large storage tanks.

Acknowledgment

The author was formerly an employee of Rebellion Photonics, Inc. (Houston TX, USA) as a researcher developing hardware and software for snapshot spectral imaging devices applied to gas sensing. However, the author is no longer affiliated with the company and declares no conflict of interest.

Appendix. Derivation of the naïve continuity equation

In order to understand the types and magnitudes of errors generated by ignoring the divergence term in the continuity equation (17), we will start by re-deriving the expression for the weighted-average velocity, but this time holding the divergence term to be negligible. Let us start with (18) and neglecting the divergence term, which leads to

$$0 = \int \left(u \frac{\partial \rho}{\partial x} + v \frac{\partial \rho}{\partial y} + w \frac{\partial \rho}{\partial z} + \frac{\partial \rho}{\partial t} \right) dz.$$

According to the divergence theorem, the z -component term $w \partial \rho / \partial z$ can be neglected, leaving

$$0 = \int \left(u \frac{\partial \rho}{\partial x} + v \frac{\partial \rho}{\partial y} + \frac{\partial \rho}{\partial t} \right) dz. \quad (23)$$

The third term inside the integral is simply

$$\int \frac{\partial \rho}{\partial t} dz = \frac{d}{dt} \int \rho dz = \frac{d\zeta}{dt},$$

while the first term inside the integral, expressing horizontal motion, has the form

$$\int u(x, y, z, t) \frac{\partial \rho(x, y, z, t)}{\partial x} dz. \quad (24)$$

If we approximate $u(x, y, z)$ as a series of discrete uniform layers along z , then we can write (24) as the sum

$$\sum_i u_i(x, y, t) \left[\frac{\partial}{\partial x} \rho_i(x, y, t) \right] \Delta z \approx \tilde{u} \sum \frac{\partial \rho_i}{\partial x} \Delta z.$$

where \tilde{u} is now the weighted average velocity formed by weighting the layer velocities u_i by the corresponding gradients $\partial \rho_i / \partial x$. In the integral limit of infinitely thin layers, \tilde{u} becomes the average velocity weighted by the continuous gradient $\partial \rho / \partial x$ along z :

$$\int u(x, y, z, t) \frac{\partial \rho(x, y, z, t)}{\partial x} dz = \tilde{u}(x, y, t) \frac{\partial \zeta(x, y, t)}{\partial t}.$$

The same analysis follows for the vertical velocity v .

Inserting (25) and the corresponding result for v into (23) recovers another form of integrated continuity equation

$$\tilde{u} \frac{\partial \zeta}{\partial x} + \tilde{v} \frac{\partial \zeta}{\partial y} + \frac{\partial \zeta}{\partial t} = 0. \quad (26)$$

The difference between this result and (22) (which keeps the divergence term) is that while (22) produces velocity estimates weighted by the fluid concentration along the line of sight, (26) produces velocity estimates weighted by the *gradient* of the fluid concentration along the line of sight.

ORCID iD

Nathan Hagen  <https://orcid.org/0000-0003-1122-9081>

References

- [1] Atkinson G, Cowpe E, Halliday J and Painter D 2017 A review of very large vapour cloud explosions: cloud formation and explosion severity *J. Loss Prev. Proc. Ind.* **48** 367–75
- [2] Thoma E D *et al* 2005 Open-path tunable diode laser absorption spectroscopy for acquisition of fugitive emission flux data *J. Air Waste Mgmt Assoc.* **55** 658–69
- [3] Lyon D R, Alvarez R A, Zavala-Araiza D, Brandt A R, Jackson R B and Hamburg S P 2016 Aerial surveys of elevated hydrocarbon emissions from oil and gas production sites *Environ. Sci. Tech.* **50** 4877–86
- [4] US Environmental Protection Agency 2017 Method 21—volatile organic compound leaks
- [5] US Environmental Protection Agency 1995 Protocol for equipment leak emission estimates *Technical Report* EPA-453/R-95-017
- [6] American Petroleum I 2004 Smart leak detection and repair (Ldar) for control of fugitive emissions (Washington, DC)
- [7] Cuclis A 2012 Why emission factors don't work at refineries and what to do about it *Emissions Inventory Conf.* (<https://www3.epa.gov/ttnchie1/conference/ei20/session7/acuclis.pdf>)
- [8] Furry D W, Harris G, Ranum D, Anderson E P, Carlstrom V M, Sadik W A, Shockley C E, Siegel J H and Smith D G 2009 Evaluation of instrument leak detection capabilities for smart LDAR application: refinery testing *Environ. Prog. Sustain. Energy* **28** 273–84
- [9] US Environmental Protection Agency 2006 40 CFR Part 60: alternative work practice to detect leaks from equipment
- [10] Footer T L 2015 Optical gas imaging protocol *Technical Report* US Environmental Protection Agency
- [11] US Environmental Protection Agency 2016 40 CFR Part 60 Subpart 0000a: Oil and natural gas sector: emission standards for new, reconstructed, and modified sources
- [12] Kangas P and Vaskinen K 2015 Parallel LDAR campaigns with sniffing and OGI: summary of the field campaign project (www.concawe.eu/wp-content/uploads/2017/01/Presentation_OGI-Concawe_Symposium_FINAL_20.02.2015.pdf)
- [13] Furry D W, Harris G, Ranum D, Anderson E P, Carlstrom V M, Sadik W A, Shockley C E, Siegel J H and White G M 2007 Evaluation of instrument leak detection capabilities for smart LDAR application: chemical plant testing *Environmental Progress* **26** 197–205
- [14] US EPA Office of Air Quality Planning and Standards 2014 Oil and natural gas sector leaks
- [15] Allen D T *et al* 2013 Measurements of methane emissions at natural gas production sites in the United States *Proc. Nat. Acad. Sci. USA* **110** 17768–75
- [16] Alvarez R A, Pacala S W, Winebrake J J, Chameides W L and Hamburg S P 2012 Greater focus needed on methane leakage from natural gas infrastructure *Proc. Nat. Acad. Sci. USA* **109** 6435–46
- [17] Pétron G *et al* 2012 Hydrocarbon emissions characterization in the Colorado Front Range: a pilot study *J. Geophys. Res. D* **117** D04304
- [18] Karion A *et al* 2013 Methane emissions estimate from airborne measurements over a western United States natural gas field *Geophys. Res. Lett.* **40** 4393–7
- [19] Caulton D R *et al* 2014 Toward a better understanding and quantification of methane emissions from shale gas development *Proc. Nat. Acad. Sci. USA* **111** 6237–42
- [20] Feng S, Farha F, Li Q, Wan Y, Xu Y, Zhang T and Ning H 2019 Review on smart gas sensing technology *Sensors* **19** 3760–81
- [21] Dinh T-V, Choi I-Y, Son Y-S and Kim J-C 2016 A review on non-dispersive infrared gas sensors: improvement of

- sensor detection limit and interference correction *Sensors Actuators B* **231** 529–38
- [22] Klein L J, van Kessel T, Nair D, Muralidhar R, Hinds N, Hamann H and Sosa N 2017 Distributed wireless sensing for fugitive methane leak detection *IEEE Int. Conf. on Big Data* pp 4583–91
- [23] Barber R, Rodriguez-Conejo M A, Melendez J and Garrido S 2015 Design of an infrared imaging system for robotic inspection of gas leaks in industrial environments *Intl J. Adv. Robot. Sys.* **12** 60058
- [24] Jones C D 1983 On the structure of instantaneous plumes in the atmosphere *J. Hazardous Materials* **7** 87–112
- [25] Nofsinger G T 2006 Tracking based plume detection *PhD Dissertation* Dartmouth College
- [26] Hodgkinson J, Shan Q and Pride R D 2006 Detection of a simulated gas leak in a wind tunnel *Meas. Sci. Tech.* **17** 1586–601
- [27] Chraim F, Erol Y B and Pister K 2016 Wireless gas leak detection and localization *IEEE Trans. Indus. Info.* **12** 768–79
- [28] Shu L, Mukherjee M, Xu X, Wang K and Wu X 2016 A survey on gas leakage source detection and boundary tracking with wireless sensor networks *IEEE Access* **4** 1700–15
- [29] Visscher A D 2013 *Air Dispersion Modeling: Foundations and Applications* (New York: Wiley)
- [30] Beychok M 2005 *Fundamentals of Stack Gas Dispersion* 4th edn (Irvine, CA: MR Beychok)
- [31] Hanna S R, Briggs G A and Hosker R P J 1982 Handbook on atmospheric diffusion national oceanic and atmospheric administration
- [32] Weil J C, Corio L A and Brower R P 1997 A PDF dispersion model for buoyant plumes in the convective boundary layer *J. App. Meteor.* **36** 982–1003
- [33] Holmes N S and Morawska L 2006 A review of dispersion modelling and its application to the dispersion of particles: an overview of different dispersion models available *Atmos. Environ.* **40** 5902–28
- [34] Korsakissok I and Mallet V 2009 Comparative study of Gaussian dispersion formulas within the Polyphemus platform: evaluation with prairie grass and Kincaid experiments *J. App. Meteor. Climat.* **48** 2459–73
- [35] Thompson V L and Greenkorn R A 1988 Non-Gaussian dispersion in model smokestack plumes *AIChE J.* **34** 223–8
- [36] Isakov V, Sax T, Venkatram A, Pankratz D, Heumann J and Fitz D 2004 Near-field dispersion modeling for regulatory applications *J. Air Waste Manag. Assoc.* **54** 473–83
- [37] Werle P, Mücke R and Slemr F 1993 The limits of signal averaging in atmospheric trace-gas monitoring by tunable diode-laser absorption spectroscopy (TDLAS) *App. Phys. B* **57** 131–9
- [38] Reid J, Garside B K, Shewchun J, El-Sherbiny M and Ballik E A 1978 High sensitivity point monitoring of atmospheric gases employing tunable diode lasers *Appl. Opt.* **17** 1806–10
- [39] Mikel D K and Merrill R 2018 EPA Handbook: optical remote sensing for measurement and monitoring of emissions flux of gases and particulate matter *Technical Report* US Environmental Protection Agency (North Carolina)
- [40] Chambers A K, Stroscher M, Wootton T, Moncrieff J and McCready P 2008 Direct measurement of fugitive emissions of hydrocarbons from a refinery *J. Air Waste Manag. Assoc.* **58** 1047–57
- [41] McRae T G and Kulp T J 1993 Backscatter absorption gas imaging: a new technique for gas visualization *Appl. Opt.* **32** 4037–50
- [42] Kulp T J, Powers P E and Kennedy R 1997 Remote imaging of controlled gas releases using active and passive infrared imaging systems *Proc. SPIE* **3061** 269–78
- [43] Low M J D and Clancy F K 1967 Remote sensing and characterization of stack gases by infrared spectroscopy: an approach using multiple-scan interferometry *Environ. Sci. Tech.* **1** 73–4
- [44] Walter H J and Flanigan D 1975 Detection of atmospheric pollutants: a correlation technique *Appl. Opt.* **14** 1423–8
- [45] Flanigan D F 1986 Detection of organic vapors with active and passive sensors: a comparison *Appl. Opt.* **25** 4253–60
- [46] Sandsten J, Edner H and Svanberg S 1996 Gas imaging by infrared gas-correlation spectrometry *Opt. Lett.* **21** 1945–7
- [47] Sandsten J, Weibring P, Edner H and Svanberg S 2000 Real-time gas-correlation imaging employing thermal background radiation *Opt. Express* **6** 92–103
- [48] Sandsten J, Edner H and Svanberg S 2004 Gas visualization of industrial hydrocarbon emissions *Opt. Express* **12** 1443–51
- [49] Naranjo E, Baliga S and Bernascolle P 2010 IR gas imaging in an industrial setting *Proc. SPIE* **7661** 76610K
- [50] Tremblay P *et al* 2010 Standoff gas identification and quantification from turbulent stack plumes with an imaging Fourier-transform spectrometer *Proc. SPIE* **7673** 76730H
- [51] Sabbah S, Harig R, Rusch P, Eichmann J, Keens A and Gerhard J-H 2012 Remote sensing of gases by hyperspectral imaging: system performance and measurements *Opt. Eng.* **51** 111717
- [52] Hagen N and Kudenov M W 2013 Review of snapshot spectral imaging technologies *Opt. Eng.* **52** 090901
- [53] Hagen N, Kester R T, Gao L and Tkaczyk T S 2012 Snapshot advantage: a review of the light collection improvement for parallel high-dimensional measurement systems *Opt. Eng.* **51** 111702
- [54] Hagen N, Kester R T and Walker C 2012 Real-time quantitative hydrocarbon gas imaging with the gas cloud imager (GCI) *Proc. SPIE* **8358** 83581J
- [55] Hagen N, Kester R T, Morlier C, Panek J A, Drayton P, Fashimpaur D, Stone P and Adams E 2013 Video-rate spectral imaging of gas leaks in the longwave infrared *Proc. SPIE* **8710** 871005
- [56] Azuma S-i, Sakar M S and Pappas G J 2012 Stochastic source seeking by mobile robots *IEEE Trans. Auto. Control* **57** 2308–21
- [57] Ishida H, Nakamoto T and Moriizumi T 1998 Remote sensing of gas/odor source location and concentration distribution using mobile system *Sensors Actuators B* **49** 52–7
- [58] Brantley H L, Thoma E D, Squier W C, Guven B B and Lyon D 2014 Assessment of methane emissions from oil and gas production pads using mobile measurements *Environ. Sci. Tech.* **48** 14508–15
- [59] Lan X, Talbot R, Laine P and Torres A 2015 Characterizing fugitive methane emissions in the Barnett shale area using a mobile laboratory *Environ. Sci. Tech.* **49** 8139–57
- [60] Ben-David A and Davidson C E 2012 Probability theory for 3-layer remote sensing radiative transfer model: univariate case *Opt. Express* **20** 10004–33
- [61] Rodríguez-Conejo M A and Meléndez J 2015 Hyperspectral quantitative imaging of gas sources in the mid-infrared *Appl. Opt.* **54** 141–9
- [62] Weinstein M A 1960 On the validity of Kirchhoff's law for a freely radiating body *Amer. J. Phys.* **28** 123–5
- [63] Eriksson T S and Granqvist C G 1982 Radiative cooling computed for model atmospheres *Appl. Opt.* **21** 4381–438
- [64] Bouwmans T 2014 Traditional and recent approaches in background modeling for foreground detection: an overview *Comp. Sci. Rev.* **11** 31–6

- [65] Stauffer C and Grimson W E L 1999 Adaptive background mixture models for real-time tracking *Proc. IEEE Computer Conf. on Computer Vision and Pattern Recognition* pp 246–52
- [66] Salisbury J W and D'Aria D M 1992 Emissivity of terrestrial materials in the 8–14 μm atmospheric window *Remote Sens. Environ.* **42** 83–106
- [67] Young S J, Johnson B R and Hackwell J A 2002 An in-scene method for atmospheric compensation of thermal hyperspectral data *J. Geophys. Res.* **D 107** 4774–93
- [68] Gross K C, Bradley K C and Perram G P 2010 Remote identification and quantification of industrial smokestack effluents via imaging Fourier-transform spectroscopy *Environ. Sci. Tech.* **44** 9390–7
- [69] Niu S, Golowich S E, Ingle V K and Manolakis D G 2014 New approach to remote gas-phase chemical quantification: selected-band algorithm *Opt. Eng.* **53** 021111
- [70] Gittins C M 2009 Detection and characterization of chemical vapor fugitive emissions by nonlinear optimal estimation: theory and simulation *Appl. Opt.* **48** 4545–61
- [71] Manolakis D G, Golowich S E and DiPietro R S 2014 Long-wave infrared hyperspectral remote sensing of chemical clouds: a focus on signal processing approaches *IEEE Sign. Proc. Mag.* **31** 120–41
- [72] Gerhart T, Sunu J, Lieu L, Merkurjev E, Chang J-M, Gilles J and Bertozzi A L 2013 Detection and tracking of gas plumes in LWIR hyperspectral video sequence data *Proc. SPIE* **8743** 87430J
- [73] Golowich S E and Manolakis D G 2011 Performance limits of LWIR gaseous plume quantification *Proc. SPIE* **8048** 80481F
- [74] Keshava N and Mustard J F 2002 Spectral unmixing *IEEE Sign. Proc. Mag.* **19** 44–57
- [75] DiPietro R S, Manolakis D G, Lockwood R B, Cooley T and Jacobson J 2012 Hyperspectral matched filter with false-alarm mitigation *Opt. Eng.* **51** 016202
- [76] Azarmi F, Kumar P and Mulheron M 2014 The exposure to coarse, fine and ultrafine particle emissions from concrete mixing, drilling and cutting activities *J. Hazard. Mater.* **279** 268–79
- [77] Nyarku M, Mazaheri M, Jayaratne R, Dunbabin M, Rahman M, Uhde E and Morawska L 2018 Mobile phones as monitors of personal exposure to air pollution: is this the future? *PLoS One* **13** 193150–67
- [78] Hagen N 2018 The influence of natural and artificial fogs on visible and infrared imaging *Proc. SPIE* **10765** 1076504
- [79] Chu P M, Guenther F R, Rhoderick G C and Lafferty W J 1999 The NIST quantitative infrared database *J. Res. Natl. Inst. Stand. Technol.* **104** 59–80
- [80] Sharpe S W, Johnson T J, Sams R L, Chu P M, Rhoderick G C and Johnson P A 2004 Gas-phase databases for quantitative infrared spectroscopy *App. Spect.* **58** 1452–61
- [81] Gallagher N B, Wise B M and Sheen D M 2003 Estimation of trace vapor concentration-pathlength in plumes for remote sensing applications from hyperspectral images *Analytica Chimica Acta* **490** 139–52
- [82] Zabetakis M G 1965 Flammability characteristics of combustible gases and vapors *Technical Report* Bureau of Mines Washington, DC
- [83] Cashdollar K L, Zlochower I A, Green G M, Thomas R A and Hertzberg M 2000 Flammability of methane, propane and hydrogen gases *J. Loss Prev. Proc. Ind.* **13** 327–40
- [84] Héas P, Herzet C, Mémin E, Heitz D and Mininni P D 2013 Bayesian estimation of turbulent motion *IEEE Trans. Pattern Anal. Mach. Intell.* **35** 1343–56
- [85] Fu X, Li J, Zhang G, Zhang C and Fan X 2019 Design of a test bench for gas leaks using CFD simulation and IR-thermography detection *Environ. Technol.* **41** 1636144–58
- [86] Grauer S J, Tsang R W and Daun K J 2017 Broadband chemical species tomography: measurement theory and a proof-of-concept emission detection experiment *J. Quant. Spect. Rad. Trans.* **198** 145–54
- [87] Rangel J, Schmoll R and Kroll A 2020 On scene flow computation of gas structures with optical gas imaging cameras *IEEE Conf. on Applications of Computer Vision* pp 174–82
- [88] Harley J L and Gross K C 2011 Remote quantification of smokestack effluent mass flow rates using imaging Fourier transform spectrometry *Proc. SPIE* **8018** 801813
- [89] Sandsten J and Andersson M 2012 Volume flow calculations on gas leaks imaged with infrared gas-correlation *Opt. Express* **20** 20318–29
- [90] Safitri A, Gao X and Mannan M S 2011 Dispersion modeling approach for quantification of methane emission rates from natural gas fugitive leaks detected by infrared imaging technique *J. Loss Prev. Proc. Indust.* **24** 138–45
- [91] Fleet D J and Weiss Y 2005 Optical flow estimation *Handbook of Mathematical Models in Computer Vision* (New York: Springer) ch 15 pp 239–58
- [92] Fitzpatrick J M 1988 The existence of geometrical density-image transformations corresponding to object motion *Comput. Vis. Graph. Image Process.* **44** 155–74
- [93] Wildes R P *et al* 1997 Physically based fluid flow recovery from image sequences *Proc. IEEE Computer Conf. on Computer Vision and Pattern Recognition* eds Wildes R P, Lanzillotto A-M, Amabile M J and Leu T-S (San Juan, Puerto Rico: IEEE) pp 969–75
- [94] Corpetti T, Mémin E and Pérez P 2002 Dense estimation of fluid flows *IEEE Trans. Pattern Anal. Mach. Intell.* **24** 365–80
- [95] Liu T and Shen L 2008 Fluid flow and optical flow *J. Fluid Mech.* **614** 253–91
- [96] Cai S, Mémin E, Dérian P and Xu C 2018 Motion estimation under location uncertainty for turbulent fluid flows *Exp. Fluids* **59** 8–24
- [97] Somov A, Baranov A, Spirjakin D, Spirjakin A, Sleptsov V and Passerone R 2013 Deployment and evaluation of a wireless sensor network for methane leak detection *Sensors Actuators A* **202** 217–25
- [98] Honeycutt W T, Ley M T and Materer N F 2019 Precision and limits of detection for selected commercially available, low-cost carbon dioxide and methane gas sensors *Sensors* **19** 3157–74
- [99] Eastern Research G, Inc. 2014 Spectral testing of gas-imaging cameras and spectral library *Tech. Rep.* ETS-2-1
- [100] Zeng Y and Morris J 2019 Detection limits of optical gas imagers as a function of temperature differential and distance *J. Air Waste Manag. Assoc.* **69** 351–62
- [101] Hagen N 2018 Sensitivity limits on optical gas imaging due to air turbulence *Opt. Eng.* **57** 114102
- [102] Gålfalk M, Olofsson G, Crill P and Bastviken D 2015 Making methane visible *Nat. Clim. Change* **6** 426–44
- [103] Day S, Dell'Amico M, Fry R and Tousi H J 2014 Field measurements of fugitive emissions from equipment and well casings in Australian coal seam gas production facilities CSIRO (Australia)