UNIVERSIDAD
DE MÁLAGA

Doctoral Dissertation

Advances in Gas Sensing and Mapping for Mobile Robotics

Javier González Monroy
2013

Tesis doctoral
Dpt. de Ingeniería de Sistemas y Automática
Universidad de Málaga

UNIVERSIDAD DE MÁLAGA
DEPARTAMENTO DE
INGENIERÍA DE SISTEMAS Y AUTOMÁTICA

El Dr. D. Javier González Jiménez y el Dr. D. Jose Luis Blanco Claraco, directores de la tesis titulada "Advances in Gas Sensing and Mapping for Mobile Robotics" realizada por D. Javier González Monroy, certifican su idoneidad para la obtención del título de Doctor en Ingeniería Mecatrónica.

Málaga, 4 de Septiembre de 2013

Dr. D. Javier González Jiménez

Dr. D. Jose Luis Blanco Claraco

Dept. of System Engineering and Automation
University of Málaga
Studies in Mechatronics

Advances in Gas Sensing and Mapping for Mobile Robotics

AUTHOR: Javier González Monroy

SUPERVISORS: Javier González Jiménez
Jose Luis Blanco Claraco

Thesis defended on 20th November 2013

JURY:
Achim Lilienthal (Örebro University, Sweden)
Lino Marques (Coimbra University, Portugal)
Santiago Marco (IBEC Barcelona, Spain)
Juan Antonio Fernandez Madrigal (Malaga University, Spain)
Alfonso García Cerezo (Malaga University, Spain)

Acknowledgements

First, I would like to express my deep gratitude to my supervisors Prof. Dr. Javier González Jiménez and Dr. Jose Luis Blanco Claraco for giving me the opportunity to join the Machine Perception and Intelligent Robotics Group (MAPIR) at the department of System Engineering and Automation (ISA) at Málaga University. I am also indebted to them for the uncountable hours we have spent together discussing about the different research topics and guiding me through the difficult art of writing scientific papers. I am also grateful to the committee members for having accepted to review this thesis, and to Dr. Patrick P. Neumann and Dr. Marco Trincavelli for their valuable comments and suggestions to improve the quality of this thesis.

Naturally, a research project like this thesis would not have been possible without the great team of people in our lab. Thanks to Raul Ruiz, Ana Gago, Eduardo Fernandez, Francisco Moreno, Mariano Jaimez, Francisco Melendez, Carlos Sánchez, Manuel Lopez and Emil Khatib for being very good friends. Most of the highlights of this last four years would not have happened without them: thank you all for the good moments at work that made it a comfortable place to stay. My gratitude goes also to the senior researchers at MAPIR, in particular to Vicente Arevalo and Cipriano Galindo, for sharing their knowledge and expertise.

During spring 2011 I visited the MR&O lab at the AASS Örebro University. I would like to thank Prof. Achim J. Lilienthal for giving me the opportunity of visiting his team, a group of wonderful and talented people. Special thanks goes for Dr. Marco Trincavelli, who invested a lot of time for helping me with suggestions and fruitful discussions. I also thank the other members of the lab and researchers I had the opportunity to meet there, for being not only colleagues but also good friends.

I would like to acknowledge my close friends Zape, Andrés and Mari for the wonderful and fantastic weekends spent together at Sierra Nevada, Pyrenees, or sharing environmental volunteering activities at my village. They shared with me not only their passion and respect for nature and sport, but taught me a different way to look at everyday life, something that has been of great help along these years of hard research.

Finally, I feel extremely grateful to my family, which has always supported me not only materially but more importantly, believing in me and inspiring me throughout this research career. Last but not least, I am completely indebted to Alex, who always encouraged me in the hardest moments. Yours is the warmest thank you.

Table of Contents

Resumen de la Tesis Doctoral

Introducción

El olfato es uno de los sentidos más directos del ser humano. Para que podamos oler algo, moléculas de ese "algo" tienen que llegar a nuestra nariz. Todo lo que olemos, por lo tanto, está emitiendo moléculas – ya sea un pastel horneándose en una pastelería cercana, perfume, una pieza de fruta podrida en la parte de atrás del refrigerador o una fuga de gas butano bajo la estufa del salón. Esas moléculas son generalmente ligeras, volátiles (fáciles de evaporar), que flotan en el aire hasta llegar a nuestra nariz.

La capacidad para detectar substancias olorosas en el medio ambiente es tan básica y tan importante que todos los organismos desde las amebas unicelulares hasta los seres humanos están dotados de algún tipo de detección química. Todas las especies usan esta capacidad para llevar a cabo uno de los comportamientos más básicos y fundamentales: acercarse y ser atraído por los aromas agradables y potencialmente seguros, así como evitar y ser repelidos por los desagradables o potencialmente dañinos.

Los seres humanos respiramos, en promedio, 20.000 veces al día [13]. Con cada respiración, inhalamos una mezcla genérica compuesta mayormente de nitrógeno, oxígeno, argón y monóxido de carbono, pero también de muchos otros gases en concentraciones mucho más bajas, algunos de los cuales pueden ser tóxicos [14]. Se dice que los seres humanos tenemos la capacidad de distinguir más de 10.000 compuestos diferentes (olores), que son detectados por unas neuronas especializadas que recubren el interior de la nariz [3]. A pesar de esta impresionante cantidad, existe aún una considerable variedad de gases tóxicos, que se encuentran no sólo en el ámbito industrial, sino también en la naturaleza, que son inodoros y, generalmente, incoloros para los seres humanos. La presencia de estos gases tóxicos y el riesgo potencial que presentan para la salud humana, junto con el fuerte interés económico de la industria del perfume y la alimentación (mayoritariamente referente a los procesos de calidad) son las principales causas que han impulsado el desarrollo del olfato artificial.

Dispositivos llamados narices electrónicas o e-noses hicieron su debut en los ochenta con el objetivo fundamental de distinguir una variedad de olores utilizando una matriz de sensores de gas, junto a técnicas de reconocimiento de patrones [120, 119]. Desde entonces, los avances en la electrónica y en la tecnología de los sensores han hecho posible la fabricación de narices electrónicas compactas, permitiendo su integración en plataformas tales como robots móviles o dispositivos inteligentes.

Un robot móvil equipado con una o varias narices electrónicas presenta una serie de cualidades que lo hacen propicio para afrontar una amplia gama de aplicaciones relacionadas con la detección de sustancias olorosas. Entre dichas cualidades, una plataforma móvil con una nariz electrónica puede muestrear continuamente el aire a su alrededor, y decidir en función de esa información olfativa las acciones pertinentes a realizar. Esto hace que los robots sean herramientas ideales para localizar fugas de gas, explosivos, drogas u otras sustancias peligrosas, evitando la exposición de un ser humano o un perro a tales gases. Además, un robot móvil puede beneficiarse de la información proporcionada por otros sensores a bordo (anemómetros, cámaras, escáneres láser, etc.) para complementar el sentido del olfato. Cuando nos encontramos una taza con un líquido oscuro en ella podemos afirmar que se trata de café no sólo por lo que vemos, sino también por lo que olemos. Asimismo, la capacidad de procesamiento de un robot junto con la consideración de técnicas de inteligencia artificial pueden ser usadas para resolver tareas que implican un cierto grado de razonamiento por parte del robot. Cuando percibimos aroma a comida en una casa, inmediatamente asociamos dicho olor a la actividad humana de cocinar e inferimos que alguien debe estar en la cocina. Cuando detectamos un olor a butano no buscamos el posible escape de gas en la sala de estar, sino que vamos directamente a la cocina donde además no inspeccionamos todos los elementos, mas sólo aquellos aparatos que utilizan gas butano (calentador, horno, etc.). Todos estos ejemplos ilustran un comportamiento inteligente y altamente complejo de percibir y actuar en el entorno, haciendo uso de la fusión de datos sensoriales y de conocimientos del mundo a un alto nivel, especialmente relacionados con información semántica.

No obstante, aún nos encontramos en fases muy primitivas del desarrollo de aplicaciones olfativas con robots móviles, necesitando posiblemente varios años antes de poder afrontar escenarios reales y complejos. Esta tesis se centra en esa línea, la investigación y el desarrollo de métodos y modelos que sirvan para avanzar un paso más en el camino hacia la integración de las narices electrónicas en el campo de la robótica móvil.

Ámbito de la tesis

Esta tesis aborda dos temas principales: por un lado el proceso de detección de gases en entornos reales mediante el uso de robots móviles y, por otro lado, el estudio de una aplicación concreta dentro del campo de la robótica olfativa, la construcción de

mapas de olores (Gas Distribution Mapping - GDM). La tesis comienza con una revisión del concepto de nariz electrónica (e-nose) para posteriormente introducir las principales tecnologías de sensores de gas existentes en la actualidad. Luego, las características más relevantes de cada tecnología son descritas, haciendo un énfasis particular en la tecnología MOX (metal oxide semiconductor) ampliamente utilizada a lo largo de esta tesis, describiendo sus ventajas y desventajas en lo referente a su uso en robots móviles. Precisamente, una de sus principales desventajas, la lenta velocidad de recuperación, que limita su uso en aplicaciones donde la concentración de gas puede cambiar rápidamente (como es el caso de la robótica móvil), es la causa que ha inspirado las dos primeras cuestiones que se investigan en esta tesis:

- ¿Existe alguna configuración de nariz electrónica que pueda ayudar a superar los efectos del largo tiempo de recuperación de los sensores MOX?
- Así mismo, y dado que la respuesta al escalón de los sensores MOX es conocida, ¿hasta qué punto podemos hacer uso de este "modelo" para paliar el largo período de recuperación necesario después de cada exposición al gas?

La primera pregunta da lugar al diseño de una novedosa nariz electrónica bautizada como *multi-chamber electronic nose* (MCE-nose). Esta nariz electrónica está compuesta por varios grupos idénticos de sensores de gas, alojados en cámaras separadas que alternan entre los estados de detección (cuando la cámara es expuesta al gas) y recuperación (cuando por ella circula aire limpio). El resultado es un dispositivo capaz de detectar cambios más rápidos en la concentración del gas que las narices electrónicas convencionales.

De la segunda cuestión surge una propuesta basada en la explotación de un modelo doble de primer orden del sensor MOX. A partir de dicho modelo, y haciendo uso solamente de la respuesta transitoria del sensor, se predice la respuesta estacionaria en tiempo real, lo que equivale a acelerar la velocidad de respuesta del sensor. Ambos enfoques, son detallados en el Capítulo 3 de esta tesis.

Dado que las especificaciones legales y reglamentos de seguridad relacionados con los niveles de toxicidad vienen dados en términos de concentraciones absolutas, es de enorme interés que la medida proporcionada por una nariz electrónica venga también expresada en esos términos absolutos y en esas mismas unidades. Sin embargo, no todas las tecnologías de sensores de gas son adecuadas para obtener dichos valores de concentración cuantificados. La tecnología MOX no es una excepción a este problema. Las reacciones que se producen entre las moléculas de un gas objetivo con la superficie del sensor MOX producen una variación en la conductancia del sensor que puede ser medida como una señal eléctrica correlada con la concentración del gas. Esta correlación es, no obstante, no-lineal y además está fuertemente influenciada por la propia dinámica del sensor y por los parámetros atmosféricos del entorno de trabajo, lo cual complica aún más la traducción a niveles absolutos de concentración. Esta problemática en la cuantificación de la concentración de gases ha motivado la siguiente cuestión de esta tesis:

- ¿Cómo pueden obtenerse valores absolutos de concentración de gas a partir de la respuesta de sensores basados en tecnología MOX cuando son utilizados en entornos reales?

Este tema se abordada en el Capítulo 4, empleando *Procesos Gaussianos* (Gaussian Processes - GP) para obtener una estimación de la distribución a posteriori de la concentración del gas dada la respuesta de un conjunto de sensores MOX. Adicionalmente, se presentan dos propuestas para considerar de forma automática la dinámica de los sensores MOX en el proceso de cuantificación, analizando de forma detallada su relevancia en la mejora de la precisión. Este enfoque probabilístico es especialmente conveniente para aplicaciones de robótica móvil olfativa, ya que de la distribución a posteriori de la concentración del gas se pueden obtener intervalos de confianza.

De entre las diferentes tareas olfativas a realizar por un robot móvil, esta tesis se centra en la construcción de mapas de distribución de gas, que aborda el problema de estimar la distribución espacial de sustancias volátiles haciendo uso de un robot móvil equipado con una nariz electrónica. Teniendo en cuenta el hecho físico de que la información proporcionada por los sensores de gas se desvanece con el tiempo debido primeramente a la naturaleza volátil de los gases, pero también a los mecanismos de transporte turbulentos que dominan la dispersión de estos, se plantean las dos siguientes cuestiones:

- ¿Cómo puede tenerse en cuenta el hecho de que los gases se desvanecen con el tiempo en la estimación de su distribución (GDM)?
- Dado que estamos interesados en trabajar en ambientes reales dónde la presencia de obstáculos influye en la distribución de los gases, ¿puede este proceso de construcción de mapas de olor tenerlos en cuenta?

Estas cuestiones se abordan en el Capítulo 5 donde se propone un nuevo enfoque a la construcción de mapas de distribución de gas basado en *Gaussian Markov-Random Fields* (GMRF). Este novedoso método no sólo atiende al "envejecimiento" de las observaciones de gas, sino que también considera la presencia de obstáculos en el entorno de trabajo. Diversos experimentos, tanto simulados como reales, validan este enfoque, proporcionando una comparación cualitativa y cuantitativa con métodos existentes.

Contribuciones

Las aportaciones más relevantes de esta tesis son:

- El diseño e implementación de una novedosa nariz electrónica basada en tecnología MOX para la detección de cambios rápidos en la concentración de

sustancias olorosas. Resultado de este trabajo son una patente y diversas publicaciones [49, 50, 51].

- El desarrollo de un nuevo modelo de sensor MOX [38], el cual mitiga considerablemente una de sus principales limitaciones cuando se emplea junto a un robot móvil: el largo tiempo de recuperación.
- La introducción de un nuevo método probabilístico para la cuantificación de gases con una nariz electrónica basada en sensores MOX. Este método permite obtener valores de concentración absolutos junto con una medida de la incertidumbre (intervalos de confianza), algo que representa un avance importante para las aplicaciones de olfato con robots móviles. Las siguientes publicaciones han surgido de este trabajo [39, 40].
- El desarrollo de un novedoso método probabilístico para la creación de mapas de distribución de gases basado en campos aleatorios de Markov Gaussianos (GMRF). Este método tiene en cuenta por primera vez dos aspectos fundamentales: primero la inclusión de los obstáculos presentes en el entorno de trabajo, obteniendo mapas que son más coherentes con los mecanismos de dispersión del gas; y segundo, la consideración de la "edad" de las observaciones como una medida del desvanecimiento de los gases en entornos reales.
- La colaboración en el diseño y desarrollo de un entorno de simulación para aplicaciones olfativas con robots móviles [36].

Todas las publicaciones derivadas de esta tesis están disponibles en: `http://mapir.isa.uma.es`

Marco de esta tesis

Esta tesis es el resultado de cuatro años de actividad investigadora de su autor como miembro del grupo de investigación MAPIR, el cual forma parte del departamento de Ingeniería de Sistemas y Automática de la Universidad de Málaga. La financiación de este periodo de investigación ha sido proporcionada por la Junta de Andalucía y el Fondo Europeo de Desarrollo Regional (FEDER) en el marco del proyecto TEP-2008-4016. Dicho proyecto abordaba el problema de dotar a un robot móvil con la capacidad olfativa para poder reconocer y estimar la concentración de sustancias olorosas. Uno de los puntos más importantes de este proyecto era la combinación de la información obtenida mediante sensores de gas, con otras modalidades sensoriales tales como cámaras o escáneres láser para alcanzar dicho objetivo. Las múltiples modalidades de sensores incorporados en el robot *Rhodon* a lo largo de sus diferentes etapas de desarrollo (véase Apéndice A), son un claro ejemplo de este compromiso.

El autor completó el programa de doctorado titulado Ingeniería Mecatrónica, coordinado por el departamento de Ingeniería de Sistemas y Automática. Este programa de doctorado otorgó al autor por un lado una visión general del campo multidisciplinario de la mecatrónica, el cual combina las ingenierías mecánica, electrónica, de control e informática, y por otro lado, un conocimiento más profundo sobre el campo de la robótica móvil, algo que ha resultado fundamental a lo largo de estos años de investigación.

Además, el autor complementó su formación académica con la participación en un curso intensivo de invierno (2012), titulado "Análisis de datos, robótica y aplicaciones móviles de sensores químicos" organizado por la Sociedad Internacional para el Olfato y Detección Química (ISOCS), y con una estancia de tres meses en el Centre for Applied Autonomous Sensor Systems (AASS), de la universidad de Orebro (Suecia), junto al grupo de robótica móvil y olfato (MR&O group). Durante dicha estancia, la temática de investigación se centró en el estudio de los sensores MOX y particularmente en cómo obtener valores absolutos de concentración a partir de su respuesta transitoria. Resultados obtenidos durante este período han sido incluidos en el Capítulo 4 de esta tesis.

Asimismo, durante el transcurso de esta tesis han surgido colaboraciones con otros grupos de investigación internacionales que han enriquecido y ampliado el alcance de la misma. Ejemplos son las colaboraciones con el Prof. Michael Biehl de la Universidad de Groningen (Países Bajos) y con el Dr. Frank-Michael Scheif del centro tecnológico de excelencia de interacción cognitiva (CITEC) de la Universidad de Bielefeld (Alemania), trabajando en el desarrollo de algoritmos de clasificación de olores para sistemas de muestreo abierto (open sampling systems – OSS). Aunque los resultados de dicha colaboración han sido presentados, no se encuentran publicados en el momento en que esta tesis ha sido escrita, por lo tanto, sólo los datos que se recogieron como parte de esta colaboración se presentan en el Apéndice C.

Por último, destacar que el marco científico en el que se encuadra esta tesis se caracteriza por la pequeña comunidad científica dedicada a esta temática. Esto es corroborado por los escasos congresos y revistas internacionales, así como por los pocos grupos de investigación dedicados al olfato con robots móviles. Sin embargo, es opinión del autor que el rápido progreso de la robótica móvil, junto con los continuos avances en la tecnología de detección de gases, permitirán en un futuro cercano hacer uso de todo el potencial que un robot móvil equipado con la capacidad de oler posee.

Estructura de la tesis

Con el objetivo de obtener la mención de *Doctorado Internacional* por la universidad de Málaga, el desarrollo completo de esta tesis está escrito en español e inglés. Así, el texto está dividido en dos partes. La primera parte, escrita en español, describe de forma resumida el contenido del trabajo, mientras que en la segunda parte, redac-

tada íntegramente en inglés, se presenta una descripción completa del mismo. Esta segunda parte se compone de los siguientes capítulos:

El **Capítulo 1** introduce la temática de esta tesis y provee una visión general de las contribuciones, del ámbito y de la estructura de la misma.

El **Capítulo 2** introduce de forma general al campo del olfato artificial y la tecnología de narices electrónicas, así como los principales campos de aplicación de un robot móvil con capacidad olfativa. En la primera parte de este capítulo se aborda el concepto de nariz electrónica y se describen las principales tecnologías de sensores de gas. En la segunda parte se revisan los mecanismos físicos que controlan la dispersión de los gases en ambientes reales, y se detallan los principales campos de aplicación de un robot móvil con la capacidad de detectar y medir diferentes sustancias olorosas: localización de fuentes de gas, seguimiento de caminos y construcción de mapas de distribución de gas.

El **Capítulo 3** presenta dos enfoques para mejorar la detección de sustancias volátiles con sensores de gas basados en tecnología MOX (metal oxido semiconductor) en el ámbito de la robótica móvil. El primer enfoque detalla el diseño y la configuración de un nuevo dispositivo de nariz electrónica denominado nariz electrónica de múltiples cámaras (Multi Chamber E-nose – MCE-nose), mientras que el segundo enfoque aborda el uso de un modelo doble de primer orden del sensor MOX para anticipar valores estacionarios de la respuesta del sensor a partir de medidas transitorias. Ambos enfoques son validados a través de diferentes experimentos, demostrando la mejora en el desarrollo de las tareas olfativas con robots móviles.

El **Capítulo 4** presenta un nuevo algoritmo probabilístico basado en procesos Gaussianos para la cuantificación de sustancias volátiles empleando un conjunto de sensores de gas MOX. La primera parte de este capítulo se centra en la matemática subyacente y describe el uso de la herramienta de determinación automática por relevancia (ARD) para seleccionar las características más relevantes de la matriz de sensores. La segunda parte se enfoca en los mecanismos de validación empleados y en las diferentes configuraciones experimentales, mostrando una atención especial a la introducción de la dinámica de los sensores MOX en el proceso de cuantificación.

El **Capítulo 5** aborda una tarea importante de la robótica móvil olfativa: la estimación de la distribución espacial de sustancias volátiles. Este capítulo propone un nuevo método basado en campos aleatorios de Markov Gaussianos (GMRF) para hacer frente a dos problemas fundamentales que no se han tenido en cuenta en enfoques anteriores: la presencia de obstáculos en la zona inspeccionada, y la combinación de las medidas de los sensores de gas tomadas en diferentes instantes de tiempo. Resultados cuantitativos y cualitativos completan este capítulo, así como una comparación detallada con métodos existentes.

El **Capítulo 6** concluye esta tesis, proporcionando un resumen del trabajo presentado y dando una visión de cómo las técnicas y avances propuestos pueden ser extendidos en un futuro.

Conclusiones

En esta tesis se ha abordado el problema de la detección y cuantificación de gases con un robot móvil, y particularmente la construcción de mapas de distribución de gas. En aplicaciones de robótica móvil, los sensores de gas se emplean generalmente en configuraciones llamadas de muestreo abierto (Open Sampling Systems - OSS). Bajo esta configuración, los sensores son directamente expuestos al medio ambiente, sin disponer de dispositivos para el control del flujo de aire, tiempos de exposición o condiciones ambientales (temperatura, humedad, etc.). Esto permite obtener información muy valiosa acerca de la dinámica de la interacción entre los sensores y los gases a analizar, algo que no se puede lograr con un sistema de muestreo cerrado. No obstante, debido también a esa exposición directa, las medidas se ven fuertemente condicionadas por los mecanismos físicos de dispersión de los gases: difusión y advección. La difusión desempeña por lo general un papel mucho menos importante en la dispersión de las moléculas olorosas, incluso en ambientes cerrados sin ventilación, siendo la advección (generalmente de carácter turbulento) la que domina este proceso. Como consecuencia, la distribución espacial de un gas resulta ser generalmente irregular y caótica, y donde además la ubicación de la fuente del gas no se encuentra por lo general en el punto de máxima concentración. Aparte de las dificultades derivadas de los mecanismos de dispersión, una serie de limitaciones en el desarrollo de aplicaciones olfativas con robots móviles vienen impuestas por la actual tecnología de sensores para la detección de gases. Inconvenientes como la falta de selectividad o la lenta recuperación, hacen aún más difícil el desarrollo de robots capaces de llevar a cabo tareas como la localización de una fuente de gas o la estimación de su distribución espacial de forma autónoma.

A lo largo de esta tesis, la tecnología mayormente empleada ha sido la basada en la unión de metal, óxido y semiconductor (MOX), debido principalmente a su alta sensibilidad, amplia disponibilidad comercial y bajo coste. Dos enfoques para paliar en cierta medida uno de sus principales inconvenientes (el largo período de recuperación), fueron propuestos en el Capítulo 3 de la presente tesis. Este largo período de recuperación (hasta decenas de segundos) limita notablemente su aplicabilidad en aplicaciones donde el sensor está expuesto a rápidos cambios en la concentración del gas, como es el caso de las aplicaciones de robótica móvil olfativa. Dicho efecto es particularmente notable en las fases de recuperación, cuando la respuesta del sensor retorna lentamente al nivel de reposo tras haber sido expuesto al gas. Consecuencia directa es la ausencia palpable de valores estacionarios en la respuesta de los sensores

MOX, siendo necesario por tanto la detección y cuantificación de gases empleando medidas pertenecientes al transitorio de las señales.

El primer enfoque sugerido para paliar los efectos de la lenta recuperación de los sensores MOX, se ha basado en la introducción de la nariz electrónica de múltiples cámaras (MCE-nose). Esta novedosa nariz electrónica acomoda conjuntos redundantes de sensores en diferentes cámaras, las cuales alternan entre los estados de detección (cuando la cámara es expuesta al gas) y recuperación (cuando por ella circula aire limpio). En cada instante de tiempo, tan solo una de esas cámaras se encuentra en el estado de detección. La clave de su funcionamiento se basa descartar las medidas de los sensores cuando se detecta el comienzo de una fase de recuperación, delegando la tarea de detección a otra cámara que contenga sus sensores en estado de reposo. Siguiendo este procedimiento, la salida global de la MCE-nose viene dada por la concatenación de las fases de subida de una secuencia de sensores MOX. Aumentar el número de cámaras aumenta indudablemente la posibilidad de disponer un conjunto de sensores en el estado de reposo y, en consecuencia, la posibilidad de detectar cambios más rápidos en la concentración del gas objetivo. Sin embargo, mayor coste, consumo de energía y complejidad son también consecuencias directas de ello, por lo tanto el número óptimo de cámaras depende en gran medida de las características de la aplicación objetivo. En esta tesis, un prototipo de la MCE-nose compuesto de cuatro cámaras ha sido construido e integrado en una plataforma móvil bajo la arquitectura robótica OpenMORA. Además de la validación haciendo uso de entornos simulados, se han presentado varios experimentos reales que corroboran la mejora en la detección de cambios bruscos en la concentración de gases. Finalmente, y a través de esos experimentos, se encontró que aunque pequeñas, las diferencias en los sensores redundantes alojados en las diferentes cámaras (del mismo fabricante y modelo), influyen en el comportamiento de la MCE-nose. Estas diferencias, que se deben no solo al proceso de fabricación, sino también el envejecimiento y al envenenamiento (reacciones químicas que alteran las características de un sensor de forma permanente), requieren por tanto de una calibración previa del sistema para lograr resultados adecuados.

En contraste con la introducción de la MCE-nose, el segundo enfoque, propuesto en la Sección 3.4, compensa la lenta dinámica de los sensores MOX estimando los correspondientes valores estacionarios a partir de una secuencia de medidas transitorias. En general, los modelos de sensor MOX propuestos en la literatura tratan de predecir su respuesta (valores de resistencia) cuando se expone a un cierto perfil de concentración de gas. Este segundo enfoque se inspira en la inversión de dicho modelo: dada una secuencia de medidas de la respuesta transitoria del sensor MOX, se predice el perfil de concentración del gas que ha producido dicha respuesta mediante la estimación de los valores estacionarios de la resistencia del sensor. El modelo empleado en esta tesis se basa en dos sistemas de primer orden (excitación y recuperación) con constantes de tiempo que dependen de la amplitud de respuesta del sensor. Diferentes experimentos han sido presentados para validar dicho enfoque. En primer lugar, y bajo condiciones de laboratorio en la que se controlaron los flujos de aire y la distribución de los gases, un experimento consistente en exponer periódi-

camente una nariz electrónica a una fuente de gas permitió la validación de este enfoque. A continuación, se consideraron dos experimentos de construcción de mapas de distribución de gas (1D y 2D) para enfatizar la aplicabilidad de dicho modelo en aplicaciones de olfato con robots móviles. Con estos experimentos se demostró cómo las "largas colas" producidas por la lenta recuperación de los sensores pueden ser evitadas, algo que conlleva implícitamente una mejora significativa en la detección de gases con estos sensores. No obstante, la mejora introducida por este enfoque no se debe solamente a que se previene la superposición entre las fases de excitación y recuperación (producidas cuando los cambios en la concentración del gas son más rápidos que el tiempo de respuesta del sensor), sino al hecho de que el perfil de concentración estimado por el modelo proporciona valores más consistentes con la excitación real que cuando tan solo se observa la respuesta del sensor. Además, se demostró que ambos enfoques, tanto la MCE-nose como el modelo inverso del sensor MOX, permiten un aumento considerable en la velocidad a la que una base móvil equipada con una nariz electrónica puede inspeccionar el entorno. Esto último conlleva en general una importante reducción en los tiempos de ejecución de las tareas relacionadas con el olfato.

En robótica móvil, muchas de las tareas relacionadas con el olfato requieren poder determinar la concentración de los gases que se están analizando. Ejemplo de ello son las especificaciones legales y reglamentos de seguridad relacionados con los niveles de toxicidad, los cuales vienen definidos en términos de concentraciones absolutas. Además, atendiendo a la naturaleza caótica que domina la dispersión de los gases, es deseable proporcionar, junto con la estimación del valor de concentración, una estimación de la incertidumbre asociada. Este problema fue abordado en el Capítulo 4, donde se presentó una propuesta de aprendizaje supervisado basada en procesos Gaussianos (GP). El problema fue tratado desde un punto de vista probabilístico, estimando una distribución a posteriori sobre la concentración de gas, dada la respuesta de una matriz de sensores MOX. Esto permite no sólo predecir la concentración de gas, sino también la incertidumbre asociada por medio de la varianza de la distribución. Diferentes configuraciones de la matriz de sensores MOX fueron estudiadas, así como la consideración de herramientas para la determinación automática de la relevancia para excluir del proceso de cuantificación los sensores menos relevantes, reduciendo así la dimensionalidad del problema. Por último, se analizaron dos propuestas para introducir la dinámica de la respuesta de los sensores MOX en el proceso de cuantificación: considerar adicionalmente muestras retrasadas en el tiempo (efecto memoria), e incluir la primera derivada de la respuesta de la matriz de sensores. Resultados experimentales demostraron, sin embargo, que ninguna de las dos propuestas producía una mejora significativa.

Habiendo dedicado la primera parte de esta tesis al estudio y la investigación de lo que se pueden considerar problemas a bajo nivel, el resto de la misma ha sido consagrada al estudio de una tarea de alto nivel dentro de la robótica olfativa: la construcción de mapas de distribución de gas (GDM). Desde la perspectiva de la robótica móvil, GDM aborda el problema de la estimación de la distribución espacial de sustancias volátiles utilizando un robot móvil equipado con una nariz electrónica. En

el Capítulo 5 de esta tesis, se presentó un nuevo algoritmo para la construcción de estos mapas de distribución de gas. El problema se abordó desde una perspectiva probabilística, modelando la estimación de la distribución espacial del gas como el cálculo del máximo a posteriori (MAP) sobre un campo aleatorio de Markov Gaussiano (GMRF). Este nuevo método incorporó por primera vez dos aspectos fundamentales que habían sido obviados anteriormente. El primero, relacionado con la validez de las observaciones de gas con el paso del tiempo, viene determinado por la naturaleza efímera de los olores y por el consiguiente hecho de que la información aportada por una observación de gas se desvanece con el tiempo. Para modelar este hecho se introdujo un factor decreciente con el tiempo como medida de la "edad" de cada observación de gas. Cuando dicho factor alcanza un valor suficientemente bajo (observación antigua) la observación asociada deja de tener validez y por tanto es descartada del proceso. La segunda novedad introducida es la de considerar obstáculos tales como paredes o muebles en el proceso de construcción del mapa de olores. Para ello, se propuso modelar la correlación entre celdas vecinas atendiendo a la presencia de estos obstáculos, así, por ejemplo, la concentración de gas de dos celdas que están separadas por una pared son consideradas independientes. Resultados experimentales, tanto simulados como reales, han sido propuestos, proporcionando una comparación detallada con métodos existentes.

Líneas futuras

Además de las propuestas para el trabajo futuro incluidas en los diferentes capítulos de esta tesis, una línea general de continuación del trabajo presentado en esta tesis estaría relacionada con la integración de los avances propuestos en esta tesis, en el campo de la robótica de servicio aplicada a los hogares. La robótica de servicio es actualmente un nicho de mercado en auge debido a los continuos avances en el campo de la robótica (SLAM, navegación autónoma, semántica, etc.) y a la mayor disponibilidad de robots comerciales. Sin embargo, dado que el olfato artificial aún se encuentra en sus primeras fases de desarrollo comercial, la mayoría de los robots comerciales aún no incorporan dispositivos de detección de gases. Se prestará especial cuidado a la colaboración activa entre el olfato y otras modalidades sensoriales, ya que para la mayoría de las aplicaciones reales de robótica el olfato no es por sí mismo el objetivo final, sino uno más de los sentidos del robot para obtener información útil del entorno que le rodea.

Chapter 1
Introduction

Smell is a very direct sense. In order for us to smell something, molecules from that thing have to make it to our nose. Everything we smell, therefore, is giving off molecules - whether it is a cake in the bakery, perfume, a piece of rotten fruit at the bottom of your fridge or a butane leak underneath the stove. Those molecules are generally light, volatile (easy to evaporate), chemicals that float through the air into our nose.

The ability to detect chemicals in the environment is so basic and so important that every organism from single-celled amoebas to human beings are endowed with some kind of chemical awareness. All species use their chemical senses for the most basic and fundamental behavior: approaching and being attracted to pleasant and potentially safe aromas, as well as avoiding and being repelled by unpleasant or potentially harmful ones.

Humans breathe, on average, $20,000$ times a day [13]. With each breath, we inhale a general mixture of nitrogen, oxygen, argon and carbon monoxide, but also many other gases at lower concentrations, some of them which may be toxic [14]. It is said that humans can distinguish more than $10,000$ different smells (odorants), which are detected by specialized olfactory receptor neurons lining the nose [3]. Despite this stunning quantity, there is a considerable variety of toxic gases, found not only in industrial environments but also in nature, which are odorless and generally colourless to humans.

The presence of these contaminants and the potential risk they bear to human health, together with the economical interest of the perfume and food industries (especially for the quality control process) are the major motivations that have boosted the development of artificial olfaction. Devices called electronic noses or e-noses made their debut in the 1980's with the challenge of distinguish a variety of odors using an array of gas sensors together with pattern recognition techniques [120, 119]. Since that debut, advances in electronics and sensor technologies have made possible the manufacturing of compact e-noses, enabling their integration into platforms such as mobile robots or intelligent appliances.

A mobile robot equipped with e-nose devices exhibits several features that make them auspicious to solve a wide range of olfaction applications. Among such features, we can highlight its ability to continuously sample the air, and decide its actions based on the perceived information in a closed-loop manner. This makes robots ideal for locating gas leaks, explosives, drugs, and other dangerous substances, avoiding the risk of intoxication of a human or a dog. Furthermore, a mobile robot may benefit from information provided by other sensors on board (anemometers, cameras, laser scanners, etc.) to complement olfaction. When we face a cup with a dark liquid on it we can assert that it is coffee not only because of what we see but also because of what we smell. In the same way as other sensory modalities may complement olfaction, the computational capacity of a robot together with methods of artificial intelligence can be used to solve tasks that require some degree of reasoning or intelligent behavior. When we perceive the scent of food being cooked we immediately associate it to that particular human activity and infer that somebody must be in the kitchen. When we detect an abnormal butane odor we do not look for the gas leak in the living room, instead we go to the kitchen where we do not inspect everything in there but only those appliances and items that use butane gas (heater, oven, ...). All these examples illustrate an intelligent and complex mechanism of perceiving and acting in the environment, which makes use of sensorial data fusion and high-level world knowledge, especially semantic information.

Although we are still several years away from accomplishing real complex scenarios, the groundwork is essential to forge ahead the fusion of e-noses into robotics and intelligent systems. This thesis contributes to this purpose, bringing closer artificial olfaction and mobile robotics.

1.1 Scope

This thesis concerns about two main topics: the process of gas sensing with mobile robots in natural environments, and one particular task known as gas distribution mapping. It starts by reviewing the concept of electronic nose to subsequently address the most spread gas sensor technologies. Especial attention is devoted to the desired characteristics of gas sensors when used in mobile robots, and particulary, focus is placed on a specific gas sensing technology based on metal oxide semiconductor (MOX), which is by far the most employed gas sensing technology in mobile robotics. The advantages as well as the main drawbacks of this technology, broadly used along the experiments in this thesis, are also detailed. One of their main drawbacks is their slow recovery, which restricts their usage in applications where the gas concentration may change rapidly, as in mobile robotics. This is the motivation that has instigated the first two research questions of this thesis:

- Is there any e-nose configuration that may help overcoming the effects of the long recovery time of MOX sensors?

- Similarly, and since the MOX step response (forward model) its well known, to which extent can we take advantage of this "model" to palliate the long recovery period needed after each gas exposure?

The first question gives rise to the design of a novel e-nose baptized as Multi-Chamber Electronic Nose (MCE-nose). This e-nose comprises several identical sets of gas sensors accommodated in separate chambers which alternate between sensing and recovery states, composing a device able to detect faster changes in the gas concentration than traditional e-noses.

Answering the second question above, a proposal based on exploiting a double first-order model of the MOX-based sensor is presented. From such model, and given the measurements of the transient state signal, a steady-state output is anticipated in real time which is an indirect way of speeding-up the sensor response. Both approaches, hardware and software, are detailed in Chapter 3.

Concentration estimation is a crucial step for realistic gas sensing applications since legal requirements and regulations are expressed in terms of absolute gas concentration, e.g. parts-per-million (ppm). However, not all gas sensing technologies are suitable to obtain such quantified values when working in natural environments. MOX technology is not an exception. The reactions of a target gas with the sensor surface produce a variation of its conductance, which is measured as an electrical signal correlated to the gas concentration. This correlation is non-linear and is highly influenced by the sensor dynamics and atmospheric parameters which makes not straightforward the translation to absolute gas concentrations. This problematic motivated the next research question of this thesis:

- How can absolute gas concentrations be obtained from the response of MOX gas sensors when used in natural environments?

This topic is dealt within Chapter 4, using Gaussian Processes to estimate a posterior distribution over the gas concentration given the response from an array of MOX sensors. Additionally, two different extensions are proposed to automatically account for the dynamics of MOX sensors in the quantification process, analyzing to which extent they can improve the accuracy of the probabilistic quantification. This probabilistic approach is especially convenient for mobile robotics olfaction applications, since from the posterior distribution confidence intervals can be obtained.

From the range of mobile robotics olfaction tasks, in this thesis we focus on the so-called gas distribution mapping (GDM), which deals with the problem of estimating the spatial distribution of volatile substances using a mobile robot equipped with an electronic nose. Given the physical fact that the information provided by gas sensors vanishes with time due to the volatile nature of gases and the dominance of turbulent transport mechanisms in natural environments, the following two research questions emerge:

- Which are the effects of the vanishing nature of gases in GDM? Can this characteristic be taken into consideration to improve the estimation of the gas distribution?

- Since we are interested in working in natural environments where the presence of obstacles such as walls or furniture affect the dispersal of gases, how can the GDM process account for it?

These questions are addressed in Chapter 5 by proposing a new approach to GDM based on Gaussian Markov-Random Field (GMRF). This novel method accounts for both the aging of the observations and the presence of obstacles in the environment. Extensive validation of the proposed method is carried out with simulated and real experiments, providing qualitative and quantitative comparison with other methods. A C++ implementation of this method has been integrated into the Mobile Robot Programming Toolkit (MRPT), and its source code is available online[1] at `http://mrpt.org`.

1.2 Contributions of this thesis

The most relevant contributions of this thesis are:

- The design and implementation of a novel MOX based electronic nose for measuring fast changing gas concentrations. As a result of this work, a patent and some publications have been produced [49, 50, 51].
- The development of a new model for MOX gas sensors [38], which considerably mitigates one of their main limitations when used in mobile robotics: the long recovery time.
- The introduction of a new probabilistic method for gas quantification with an array of MOX gas sensors. This method allows us to obtain absolute concentration readings together with confidence intervals, something that represents an important advance to mobile robotics olfaction. The following publications have arisen from this work [39, 40].
- The development of a novel probabilistic GDM method based on Gaussian Markov-Random Fields which accounts, first, for the obstacles in the environment, obtaining maps which are more compliant with the actual mechanisms of gas dispersion; and second, for the "age" of gas measurements as a way of coping with the vanishing nature of volatile substances. The method has been presented in [37].
- The collaboration in the design and development of a simulation framework for mobile robotics olfaction [36].

[1] In particular, this method is implemented in the class `mrpt::slam::CRandomFieldGridMap2D`, part of the `mrpt-maps` library.

Next, all the publications derived from this thesis are compiled:

Journals

1. *J. G. Monroy, J.L. Blanco, J. Gonzalez-Jimenez*, Time-Variant Gas Distribution Mapping with Obstacle Information, (2013) *submitted*

2. *J. G. Monroy, A. Lilienthal, J.L. Blanco, J. Gonzalez-Jimenez, M. Trincavelli*, Probabilistic Gas Quantification with MOX Sensors in Open Sampling Systems - A Gaussian Process Approach, (2013) in: Sensors and Actuators B: Chemical, 188:0(298 - 312)

3. *J. G. Monroy, J. Gonzalez-Jimenez, J.L. Blanco*, Overcoming the slow recovery of MOX gas sensors through a system modeling approach, (2012) in: Sensors, 12:10(13664 - 13680)

4. *J. Gonzalez-Jimenez, J. G. Monroy, J.L. Blanco*, The Multi-Chamber Electronic Nose - An Improved Olfaction Sensor for Mobile Robotics, (2011) in: Sensors, 11:6(6145 - 6164)

Conference Proceedings

1. *J. G. Monroy, J.L. Blanco, J. Gonzalez-Jimenez*, An Open Source Framework for Simulating Mobile Robotics Olfaction, 15th International Symposium on Olfaction and Electronic Nose (ISOEN), Daegu, South Korea, 2013

2. *J.L. Blanco, J. G. Monroy, J. Gonzalez-Jimenez, A. Lilienthal*, A Kalman Filter Based Approach to Probabilistic Gas Distribution Mapping, Proceedings of the 28th Annual ACM Symposium on Applied Computing, SAC '13, Coimbra, Portugal, pp. 217-222, 2013

3. *J. G. Monroy, A. Lilienthal, J.L. Blanco, J. Gonzalez-Jimenez, M. Trincavelli*, Calibration of MOX gas sensors in open sampling systems based on Gaussian Processes, IEEE Sensors, Taipei, Taiwan, pp. 1-4, 2012

4. *J. Gonzalez-Jimenez, J. G. Monroy, F. Garcia, J.L. Blanco*, The Multi-Chamber Electronic Nose (MCE-nose), Proceedings of the IEEE International Conference on Mechatronics (ICM), Istambul, Turkey, pp. 636-641, 2011

Patent

1. *J. Gonzalez-Jimenez, J. G. Monroy, J. L. Blanco and F. Garcia*, Electronic nose having a high sensing frequency and method for determining the quantitative and qualitative composition of a gas or mixture of gases using same, WO Patent 2012-049341, 2012

1.3 Framework of this thesis

This thesis is the outcome of four years of research activity of its author as a member of the MAPIR research group[2], which is within the Department of System Engineering and Automation[3] of the University of Málaga. For this research, funding was supplied by the Andalucía Government and the European Regional Development Fund (ERDF) under project TEP-2008-4016. This project addressed the problem of providing a mobile robot with olfaction capability to recognize and estimate odor concentrations in real environments. One of the research lines of this project is the combination of the olfactory information with other sensing modalities such as cameras or laser scanners. The multiple sensing modalities incorporated into the research robot *Rhodon* along its different development stages aim at this purpose (see Appendix A).

During the PhD period, the author completed the doctoral program entitled "Ingeniería Mecatrónica" (Mechatronics Engineering) coordinated by the Department of System Engineering and Automation. This doctoral program granted the author both a general view of the multidisciplinary field of mechatronics which combines mechanical, electrical, control and computer engineering, and more importantly a deep knowledge about mobile robotics, something that has proved fundamental throughout these years of research.

Additionally, the author complemented his academic education with the participation in the Short Winter School (2012) entitled "Data Analysis, Robotics and Mobile Applications of Chemical Sensors" arranged by the International Society for Olfaction and Chemical Sensing (ISOCS), and a three months stay at the Centre for Applied Autonomous Sensor Systems (AASS), university of Orebro (Sweden), with the Mobile Robotics and Olfaction (MR&O) group. During this stay, research was focused on studying MOX sensors and particularly how to obtain absolute concentration values from their transient response. The work and results obtained during this period have been included in Chapter 4.

Furthermore, collaborations with other international research groups have been established during this thesis to enrich and extend its scope. Examples are the cooperations with Prof. Michael Biehl at the University of Groningen (Netherlands) and Dr. Frank-Michael Schleif at the Cognitive Interaction Technology-Center of Excellence (CITEC) at the University of Bielefeld (Germany) focusing on the development of high efficient odor classification tools for open sampling systems. Results of this collaboration have been presented but not published at the time this thesis is written, thus, only the data-sets that were collected as part of this collaboration are presented in Appendix C.

Finally, it is worth mentioning that the scientific framework within this thesis stands is characterized by the small research community engaged to this topic. This is actually corroborated by the few international conferences and journals devoted

[2] http://mapir.isa.uma.es
[3] http://www.isa.uma.es

to this field, as well as for the few research groups dedicated to olfaction with mobile robots. However, the author believes that the rapid progress of mobile robotics together with the continuous advances in the gas sensing technology, will, in a near future, unlock the full potential of a mobile robot equipped with the capability to smell.

1.4 Structure of this thesis

The remaining chapters of this thesis are organized as follows:

Chapter 2 gives a general introduction to the field of artificial olfaction and electronic nose technology, as well as to the main odor-related tasks of mobile robotics. The first part of the chapter summarizes the concept of electronic nose and devises relevant aspects of the most common gas sensor technologies, while the second part reviews the mechanisms of gas dispersal in natural environments, and details the principal tasks of a mobile robot with the capacity to detect and measure different volatile substances: gas source localization, trail following, gas distribution mapping and gas discrimination.

Chapter 3 presents two approaches to enhance the sensing of volatile substances with metal oxide semiconductor (MOX) gas sensors in mobile robotics applications. The first approach details the design and configuration of a new e-nose device called Multi-Chamber Electronic Nose (MCE-nose), while the second approach exploits a double first-order model of the MOX sensor to anticipate steady-state values from measurements of the transient state signal. Qualitative experiments are reported for both approaches, demonstrating the improvement achieved.

Chapter 4 introduces a new probabilistic algorithm based on Gaussian Processes for the quantification of volatile substances using an array of MOX gas sensors. The first part of this chapter focuses on the underlaying mathematics and describes the use of automatic relevance determination to select the most relevant features from the array of sensors. The second part deals with the validation mechanisms and the different experimental configurations, giving a especial attention to the introduction of the dynamics of MOX sensors in the quantification process.

Chapter 5 addresses an important task for mobile robotics olfaction: the estimation of the spatial distribution of volatile substances. A new method based on Gaussian Markov Random-Fields is proposed to deal with two important problems that have been disregarded in previous approaches: the presence of obstacles in the inspection area, and the combination of odor measurements taken at different instants of time. Quantitative and qualitative results are reported under different scenarios, as well as a detailed comparison with existing methods for gas distribution mapping.

Chapter 6 concludes the thesis, providing a summary of the presented research work and giving an outlook of how the proposed techniques can be extended further.

Chapter 2
On technology and applications of mobile robotics olfaction

This chapter provides an overview of electronic olfaction and electronic nose technology, with especial interest in their applicability to mobile robots. The chapter summarizes the concept of electronic nose and devises relevant aspects of the most common gas sensor technologies. Then, it reviews the mechanisms of gas dispersal in natural environments to end with a review of the four main odor-related tasks of mobile robotics olfaction.

2.1 Electronic noses

A general accepted definition for an electronic nose was proposed by Gardner et al. [44], after refining the initial concept introduced by Persaud et al. [120] in the early 1980's. According to this definition:

> "An electronic nose is an instrument which comprises of an array of electronic chemical sensors with partial specificity and an appropriate pattern recognition system, capable of recognising simple or complex odors." (Gardner, [44]).

This "electronic nose" term is rather general and consequently may lead to confusion or misleading. In short, an electronic nose is a device designed to artificially

mimic the animal sense of smell by providing an analysis of chemical mixtures. It works as a non-separative mechanism: i.e. an odor is perceived as a global fingerprint, a signal pattern that is used to characterize it.

This pattern is formed by collecting the signal response from each sensor in the array. The absorption of volatile molecules causes physical changes of the sensor, which are measured as electrical signals to compose the fingerprint. In most electronic noses, each sensor is sensitive to a broad range of volatile molecules (low selectivity), but each in their specific way. Thus, the degree of selectivity and the type of odors that can be detected largely depend on the choice and number of sensors in the array.

2.2 Gas sensor technologies

According to the transduction principle applied, gas sensors can be classified into the following families: thermal, mass, electrochemical, and optical [72]. Thermal sensors use the heat generated by the chemical reaction between the analyte and the sensor as the source of analytical information. The general strategy is to place the chemically selective layer on top of a thermal probe and measure the heat evolved in the specific chemical reaction taking place in that layer. For mass sensors, the transduction principle is the detection of the change of mass through the variation in behavior of some oscillator, usually piezoelectric crystals. Change of mass accompanies many interactions of the chemical species with the sensor, thus, mass sensors represent an important segment of the chemical sensing field. However, electrochemical transducers are the largest and the oldest group employed for chemical sensors. This family includes potentiometric sensors, which measure voltage, amperometric sensors, which measure current, and conductometric sensors which measure conductivity. Finally, the measuring principle of optical gas sensors is based on the emission of electromagnetic radiation through the gas sample. Different chemical species exhibit absorption of such radiation at different regions of the electromagnetic spectrum. This property is the basis for their detection and measurement.

Each of these sensor families has advantages and disadvantages over their counterpart and choosing the right technology strongly depends on the type of application. However, within the range of applications related to mobile robotics olfaction, there are a set of desired sensor properties:

- High selectivity - It can be defined as the ability of a sensor to respond primarily to only one analyte in the presence of a mixture. Ideally, the reactive layer of the sensor should completely reject any interfering analyte, and respond exclusively to the desired one.

- High sensitivity - Together with selectivity, they are possibly the most important issues in chemical sensing. Sensitivity is the amount of change in the measurable output magnitude per unit change in the volatile concentration.

- Rapid response - In general, sensors do not change their output state immediately when a change in the analyte concentration occurs. Rather, they change

over a period of time called the response time. This property is especially important when integrating gas sensors with mobile robots which are continuously inspecting an environment. Chapter 3 provides an exhaustive study of this property and its influence in the applicability to mobile robotics.

- Robustness - Understood as stability of the sensor behavior under simultaneous changes in model parameters caused by humidity, temperature, or even physical motion.

- Reliability - It can be defined as the ability of a system to perform and maintain its functions in routine circumstances. This property is closely related to the undesired *drift*, defined as "a gradual change in any quantitative characteristic that is supposed to remain constant" [155]. Thus, a drifting chemical sensor does not give exactly the same response even if it is exposed to exactly the same environment for a long time.

There is no simple answer when looking for the best sensor because of the many different sensing situations and criteria that must be considered. Instead, a general precept must serve: the best sensor is the one that will do the job at a cost which justifies its use [72]. The cost must be viewed in terms of money, time, and ease of use. This is fully applicable to mobile robotics, particularly for out-of-the-lab applications.

A complete survey of the different sensor technologies is out of scope of this thesis and therefore, only a brief overview of the most spread gas sensing technologies used in mobile robotics is provided in the following subsections.

2.2.1 Semiconducting metal oxide gas sensor

Metal oxides (MOX) such as SnO_2, ZnO, Fe_2O_3, and WO_3 are intrinsically semiconductors that at temperatures of $200-500$°C respond to reducible gases such as H_2, CH_4, CO, C_2H_5, or H_2S and increase their conductivity (conductometric family) [119]. Among the different semiconducting materials, SnO_2 doped with small amounts of impurities is the most widely employed for gas sensor fabrication. By changing the choice of impurity and operating conditions such as temperature (temperature modulation), many types of gas sensors can been developed.

MOX competitive advantages include a good commercial availability, high sensitivity, an effective life span from three to five years, and low prices when compared with other sensing technologies. Additionally, they are usually small and light, which make them very appealing for mobile robotic applications. They present, however, different shortcomings, among others:

- The need to be pre-heated at temperatures up to $200-500$°C in order to facilitate the interaction with the target gas, which involves a relatively high power consumption.

- Poor selectivity and relatively low robustness.

- The acquisition cycles are very long because of their slow response, especially when recovering to the baseline level after the exposure to the target gas ends [107]. This baseline level represents the sensor output in absence of target gases and varies with temperature, humidity and among sensors.

These limitations come from the chemical mechanism underlying MOX sensors, related to the semiconductor behavior when exchanging oxygen molecules between the volatile and the MOX film [2, 135].

2.2.2 Photo ionization detector

A Photo Ionization Detector (PID) is a gas detector that is able to measure the concentration of a variety of volatile organic compounds (VOCs) by using photo ionization. This process occurs when an atom or molecule absorbs "light" of sufficient energy to cause an electron to leave and create a positive ion. The PID is comprised of an ultraviolet lamp that emits photons that are absorbed by the compound in an ionization chamber. Ions (atoms or molecules that have gained or lost electrons and thus have a net positive or negative charge) produced during this process are collected by means of an electric potential difference between electrodes. The current generated provides then a measure of the analyte concentration. Because only a small fraction of the analyte molecules are actually ionized, this method is considered nondestructive.

The principal advantage of PIDs compared with MOX sensors is that concentration measurements are almost immediate, that is, PIDs have a very short response time, making them well-suited for mobile robotics applications where fast changing gas concentrations are expected to be faced. However, as stand alone detectors PIDs are not selective; that is, they ionize everything with an *ionization energy* less than or equal to the lamp output. The latter additionally involves that only analytes which have ionization energies similar to or lower than the energy of the photons produced by the PID lamp will be detected.

2.2.3 Surface acoustic waves

Acoustic Wave (AW) gas sensors, also known as Quartz Crystal Microbalance (QCM) gas sensors, are devices that weigh gas molecules (mass sensors family) by measuring the change in frequency of a quartz crystal resonator [131]. The frequency of oscillation of the quartz crystal is partially dependent on the thickness of the crystal. As gas molecules are deposited on the surface of the crystal, the thickness increases; consequently the frequency of oscillation decreases from the initial value. By applying different chemical coatings to the crystal, the QCM sensor can be made responsive to different volatiles. Due to this dependency between thickness and frequency, QCM sensors are sometimes termed as *thickness-shear resonators*.

Between the different modalities of AW sensors, surface acoustic waves (SAW) sensors make use of the fact that the amplitude of an acoustic wave propagating along the surface of a solid material decays rapidly, typically exponentially with the depth

of the material. SAW devices typically apply an alternating electric field to a piezoelectric material covered with a thin film, to generate a SAW. The mass of the film temporarily increases as molecules of the target gas are absorbed, perturbing the propagation of the acoustic waves. This causes a shift in resonance to a slightly lower frequency and thus information about the gas concentration can be obtained. A more detailed description is given by Gardner and Bartlett in [44].

The main advantages of this sensing technology are the low power consumption, the possibility to control the selectivity over a wide range, long term stability, long lifetime and a rapid response. The latter is especially important for applications on mobile systems, requiring in general a shorter recovery time than MOX gas sensors [81]. Disadvantages include comparatively low sensitivity to the target gas and limited robustness to variations in humidity and temperature.

2.2.4 Conductive polymer sensors

The term *polymer* derives from the ancient Greek word *polus*, meaning many/much and *meros*, meaning parts. It refers to a molecule whose structure is composed of multiple repeating units (monomers), with the characteristic property of having a large molecular mass when compared to small molecule compounds. This large mass produces unique physical properties which make polymers an appealing material for sensor fabrication.

Conducting polymers (CP) are a popular sensing technology that is based on measuring the resistance of a thin film polymer. The sensing mechanism is similar to this of MOX gas sensors except because a thin polymer film is used instead of a semiconductor material. The response given by the sensor is created by a chemical reaction that occurs on the surface of a polymer placed between two electrodes. Since the conductivity of pure conducting polymers is rather low, a doping process is usually necessary. The concept of doping is the central theme which distinguish conductive polymers from all other polymers [97]. When the doped polymer comes into contact with particular gases, carriers on the polymer chain become mobile and produce an electrical conductivity which is then measured as the sensor response.

In comparison with MOX gas sensors, polymers can operate at room temperatures (with the consequent energy savings), they show a good sensitivity to a wide variety of organic compounds (but approximately one order of magnitude lower than the MOX sensors [81]), and have short time responses.

2.3 Gas dispersal in natural environments

Far from the complex laboratory setups which use sophisticated sampling systems to keep constant along the olfaction process parameters such as air flow, temperature, humidity and analyte concentration, in most real-world applications, and particulary in mobile robotics, gas sensors are directly exposed to the environment to be analyzed with no control at all. This leads to additional sources of uncertainty, mainly due to

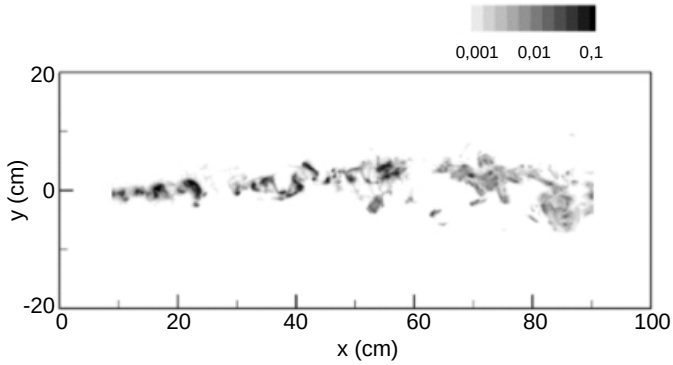

Figure 2.1: Sample of a instantaneous concentration field (figure taken from Webster et al. [154]). The chemical source is located at coordinates (0,0) releasing the gas into a turbulent stream flowing in the positive x-direction.

the mechanisms that rule the dispersal of gases in natural environments: diffusion and turbulent advection.

Diffusion, from Latin "diffundere", means "to spread out". This process causes mass transport without requiring bulk motion. Diffusion can be considered as a result of the "random walk" [74] of particles which are self-propelled by thermal energy. The rate of this movement is a function of temperature, viscosity of the fluid and the size (mass) of the particles. Although diffusion is present under almost all circumstances (as long as the temperature is over the absolute zero), in natural environments advective flow [137] is the process that dominates the dispersal of gases. In average, the diffusion length of typical gas molecules for one hour is only 20 cm, while in both outdoor and indoor environments we usually find airflows with much higher velocities [62]. Thus, in general, the extremely slow process of molecular diffusion can be neglected.

A turbulent flow can be defined as the viscous flow in which fluid particles move in a random and chaotic way within the flow field [141, 139]. Velocity and all other fluid properties vary continuously, with strong concurrent molecular mixing between adjacent fluid layers. In natural environments, turbulence is the dominant mechanism in the mixing and dilution of gaseous releases [22].

Odor patches released by an odor source are mainly transported by the advective turbulent flow of an air stream, forming an odor plume. As the plume travels away from the source, it becomes more diluted due to molecular diffusion and turbulence that mixes the odor molecules with the clean air [99]. Figure 2.1 displays a representative example of the instantaneous concentration of a gas plume in a turbulent flow.

For the case of MOX gas sensors (broadly employed in this thesis), given their slow dynamics and the rapid fluctuations in the gas concentration due to the mecha-

nisms of gas dispersal, they never reach steady state but continuously fluctuate [146]. Thus, olfaction algorithms developed to work in such setups additionally require to extract information about the target gas (the type of analyte, its distribution, the source location, etc.) from the transient of the sensor response.

Understanding how the odor molecules disperse through the environment under naturally turbulent flows is not a trivial task, being necessary the use of computational fluid dynamics (CFD) to numerically approximate the solution. However, CFD applications cannot run online because of their high computational cost, which preclude their use in mobile olfaction.

2.4 Mobile Robotics Olfaction

Gas sensing is a relatively recent research area within the field of mobile robotics. This is in part because the technological progression of compact gas sensors is integral to the solution of detecting odors with mobile robots and there is still much development needed before the gas sensors are satisfactory for real applications [88]. Despite this, significant advances have been reported since the beginning of the 1990s. Four are the main odor-related tasks that have been addressed by the research community: gas source localization, trail following, gas distribution mapping and gas discrimination. In the following subsections a review of the most relevant works on each task is presented.

2.4.1 Gas source localization

Inspired by biology and the fundamental mechanisms of animal olfaction, the task of localizing a gas source has been pursued since the first research works on mobile robotics olfaction.

Researchers started to face the gas source tracing issue using techniques based on gradient-following [126, 134], trying to localize the point of highest concentration as a representative characteristic of a gas source. These methods assumed that diffusion is the dominant short-term method of odor dispersal, leading to a stable and smooth chemical concentrations [77]. Exploiting this "chemical gradient", the robot was expected to move towards the odor source, since on these scenarios the evolution of the chemical concentration along the gas plume is well defined by a continuous function with a peak close to the gas source location [116]. However, in real environments the dispersion of gases is not diffusion-based but it is governed by turbulence (see Section 2.3). The flow contains eddying motions of a wide range of sizes that produce a patchy and intermittent distribution of the gas. The gradient is then time varying, steep, and frequently in the opposite direction to the source [25, 73]. This makes gradient-following techniques not practical in real environments.

Gas source localization under strong constant airflows

Methods designed to operate in the presence of strong and constant airflows are mostly based on "reactive plume tracking" [77]. These methods, largely based on the odor localizing behavior of microbes, insects, and crustaceans, employ reactive control schemes and local sensing to track the plume along its entire length to the source [79, 78, 27, 82, 86, 109].

The main characteristic of these scenarios is that reliable information about the air flow direction is available (anemotaxis), which can be used for upwind navigation together with chemotaxis (gradient based navigation behavior). The first example of an odor localizing robot that utilized both chemical and anemometric sensors was described by Ishida and colleagues [68]. Further developments have been introduced by a number of researchers, most of them based on biological inspired methods.

Worth mentioning is the *Zigzag/Dung Beetle* method [63, 64, 132], which involves moving upwind within the odor plume in a zigzagging fashion. Each time the plume boundary is encountered (detected by measuring the gas concentration gradient), the robot turns back into the plume. Another version of this zigzag approach was implemented for underwater robots and reported by Farrell et al. [29].

Other methods include the *Plume-Centered upwind search* [129, 67, 100], which involves moving towards the center of the plume while tracking upwind, or the *Silkworm Moth* method, which is probably the most studied animal behavior for gas source localization. Male moths can trace a pheromone emitting source (female) even at very long distances in a turbulent environment. The chemotactic behavior of the silkworm moth has been implemented in many research works, as in the case of [80, 78, 101, 121].

It is not until the work of Martinez and Perrinet [102] that visual information was used in collaboration with local chemical sensing. They proposed a visual identification of salient features to locate candidates, reducing the time necessary to find the real gas source. In [95] a similar approach is presented but employing standard image processing techniques to detect cups (the potential odorsources), while in [70, 69] Ishida et al. proposed a subsumption architecture to combine random, visual and olfactory search behaviors.

Another approach that requires a strong and constant airflow condition is to model the shape of the gas plume based on the measurements of the wind vector and chemical concentration. Analytical models of the gas plume [65] or stochastic methods [30, 116] can then be used to predict the gas source location.

Gas source localization under weak airflows

The localization of a gas source in weak airflow scenarios (e.g. unventilated indoor environments) cannot rely on the anemometry information. Under these conditions, an analyte plume can still be detected in most cases [152], however, with different properties concerning its shape, width, concentration profile and stability over time.

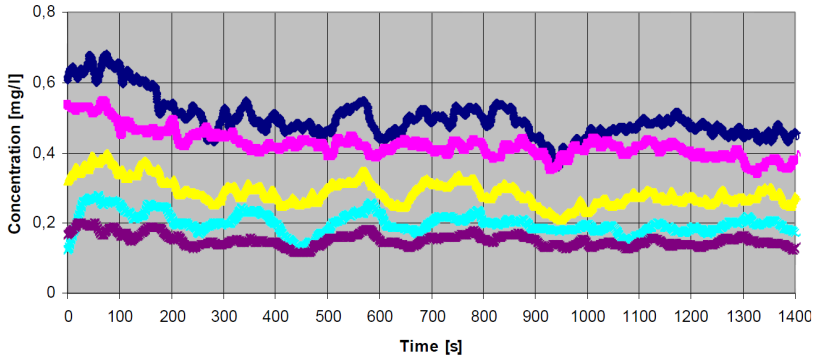

(a)

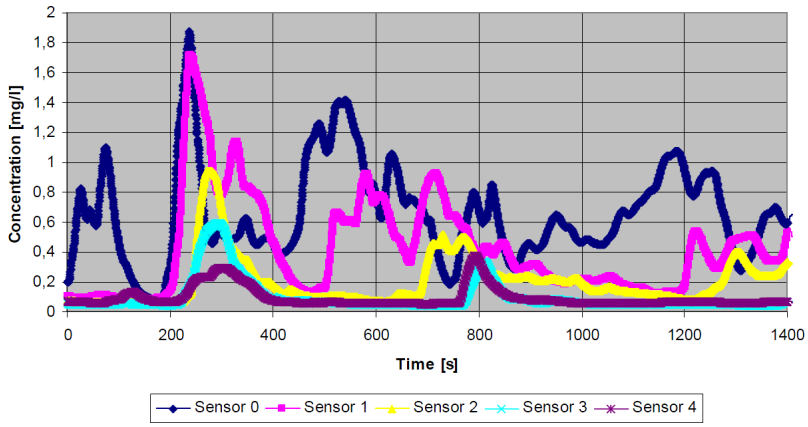

(b)

Figure 2.2: (a) Source characterization with a strong airflow forced by a fan behind the gas source, and (b) without the fan. *Sensor0* is placed at a distance of 30 cm from the source. All sensors are spaced 25 cm, so that the farthest sensor (*Sensor4*) is at a distance of 130 cm. For the case of strong airflow, the sensors are placed inside the plume. Figure taken from Ferri et al. [34]

Few research works have been proposed in this case due to the complexity of gas dispersal in these scenarios: gas plumes are not well formed and the distribution of gases is dominated by turbulence, that is, it is patchy and chaotic as shown in Figure 2.2. The first works were proposed by Lilienthal et al. [92, 91] and Wandel et al. [153, 152], where only chemical sensing was employed to localize the gas source. The importance of these works lies in the raise of questions about the operation of mobile robots in more realistic environments, providing extent sets of experiments.

Further efforts to improve the localization of gas sources under these complex scenarios were reported through the mimic of the repetitive movements of insects

Figure 2.3: An ant following a pheromone trail marked on the ground. Concentrations perceived at the left and right antennae are compared and used to turn back to the trail. Figure taken from Ishida et al. [119]

searching for an odor source [34], or the use of analytical models [90]. The latter approach involves creating a gas concentration gridmap of the environment [84, 83], without tracking the gas plume (if it exists). This approach is further detailed in Section 2.4.3.

2.4.2 Trail following

Inspired by the olfactory-guided behaviors of insects and animals where chemical substances are often used to mark trails or territories, trail following has become an active research area within mobile robotics olfaction. Probably, the most famous example of olfactory-guided behavior is the case of ants. Ants, as many other insects, secrete pheromones to constrain the behavior of other individuals. Pheromone trails serve as a multi-purpose chemical signaling system: it leads members of its own species (and generally same nest) towards a food source or home, while it represents a territorial mark to other species [156]. Once the trail is laid, other members of the nest will recognize the chemical signal and follow the trail [1] (see Figure 2.3). Since chemical marks evaporate with time, each individual following the trail renews the marks on the way back home. While this pheromone is constantly deposited by its members, the chemicals diffuse up into the environment propagating its message. Once the food source runs out, the organisms will simply skip the task of renewing the trail on the way back, thus resulting in the diffusion and weakening of the pheromone [104].

Another example are honey bees. They use chemical markers not to follow a trail, but to increase their efficiency when gathering nectar. After visiting a flower and gathering its nectar, the honey bee marks the flower with a short-life odor to indicate that the nectar of that flower has been recently collected, thus it may be left uninspected for a while [47].

The use of chemical markings for trail guidance or signaling may be of benefit for a number of applications in the field of mobile robotics. Odor trails could provide an inexpensive and more flexible alternative to the metal wires buried under the floor that are often used for industrial automated guided vehicles (AGV) [142]. Odor trails provide higher flexibility since they are easier to lay on the floor, however, they decay

over time as the chemical substance gradually evaporates, which represents one of their fundamental constraints.

Other scenarios in which odor trails may simplify the tasks to be accomplished by mobile robots were proposed by Russell [128, 129]. One example is the use of chemical markings to provide temporary warning signals, for example to indicate areas on the floor that have been cleaned [26]. While this would be particularly beneficial to coordinate the behavior of multiple robots, it could also be helpful in the case of a single robot, because it avoids the necessity for maintaining a consistent spatial representation [81].

To reduce the "odor confusion" effect generated mostly by the advection and diffusion of odor particles, but also due to the slow response of gas sensors, mechanisms to create an air curtain around the sensors has been proposed [127]. This air curtain increases the gas sensing differentiation near the floor by blocking external airflows that may interfere and create confusion in the readings.

Several navigation strategies have been suggested for trail following. The most straightforward implementation is a robot performing tropotaxis (imitating the behavior of ants) with two gas sensors (left and right), as the proposed by Russell et al. [129], or Stella et al. [142]. Further trail following strategies rely on using only one gas sensor, as in [128] where a klinotactic algorithm is employed to follow the edge of a trail, or based on more robust strategies against sensor errors, which involve frequent crossings of the trail along a sinusoidal walk [130].

2.4.3 Gas distribution mapping

Gas distribution mapping (GDM) is the process of creating a representation of how gases spread in an environment from a set of spatially and temporally distributed measurements of relevant variables [9, 89]. Foremost, these measurements include the gas concentration itself, but may also comprise wind, pressure or temperature.

Gas distribution mapping is of great help not only because it can be used to pinpoint the location of a gas source (or of multiple sources) without depending on the environmental conditions (see Section 2.4.1), but also because it provides information of how the gas emissions have spread in the environment, which is crucial in many real olfaction-related applications. For example, lets consider an industrial plant where a leak of a toxic gas has been detected. For safety considerations, it is not enough to locate the room or even the pipe that is leaking, but it is necessary to know which areas of the plant have been affected by the toxic gas to safely prepare the action plan.

Traditionally, the way to create a representation of the gas concentration field is to measure the response of a grid of gas sensors distributed in the environment [150, 71]. The main advantage of these networks of static gas sensors is that the instantaneous gas distribution can be obtained by reading all the sensors in the grid at a time, similar to taking a "picture" of how gases are distributed. Nonetheless, given the dynamic and changing characteristics of the gas distribution in real environments (see Section 2.3)

for many applications it is better to obtain the time-averaged concentration field by averaging the readings over a prolonged time [66].

An important drawback of the sensor network is that it is not scalable when the area to inspect increases, rising considerably the deployment cost and reducing its flexibility. Because of this, the attention shifted to GDM with mobile robots, which using only one electronic nose allows obtaining gas maps with high flexibility. Advantages of this approach include the use of only one sensing device (which may be complex and expensive), the capability of the robot to sample at adaptive resolutions depending on the area being inspected, and the possibility to use additional environmental information gathered by other sensors on board (cameras, laser scanners, anemometers, etc.) [70]. Chapter 5 discusses in detail the advantages of using mobile robots for GDM while providing an innovative approach which accounts for the obstacles in the environment and the dynamic characteristics of gas distribution.

Probably, the first work studying the distribution of gases with mobile robots was presented by Hayes et al. [56], where a group of mobile robots (swarm of robots) worked in a coordinated manner to create a histogram representation of the gas distribution. The histogram bins contained the number of "odor hits", that is, the number of measurements above a predefined threshold. This binary information was collected by all the robots while inspecting the area following a simple random walk pattern. Apart from requiring an even coverage of the environment, this approach also takes a very long time to obtain statistically reliable data, and no extrapolation is performed to areas not inspected. These drawbacks lead to a bad scalability when applied to large environments, fact that makes doubtful its applicability in real scenarios.

Improvements to this approach were reported by Pyk et al. [121], employing bicubic interpolation to extrapolate the gas distribution to zones not directly inspected by the robot. A disadvantage of this method is, however, that no spatial averaging is carried out and therefore fluctuations appear directly on the map.

Nevertheless, the most remarkable works in this field have been reported by Lilienthal and colleagues. In the pioneer work [85] they proposed the kernel-based method, which consists of convolving sensor readings with a Gaussian kernel, thus providing a representation of the gas map without assuming any predefined parametric form for the distribution. This method was later extended for the case of multiple odor sources [96] and to the three-dimensional case [124]. It was further shown how gas distribution mapping methods can be embedded into a Blackwellized particle filter approach to account for the uncertainty about the position of the robot [87]. A deeper review of the works proposed in the field of GDM with ground mobile robots is later presented in Section 5.2.

In the last few years, it's worth highlighting the attention paid by the research community to GDM with unmanned aerial vehicles (UAVs). Quickly deployable, cost-efficient or easy to transport are some of the advantages that flying mobile measurement devices provide when measuring the gas concentration outdoor [113]. Additionally, micro vertical take-off and landing (VTOL) UAVs, such as *quadrocopters*, have the ability to hover over a certain point of interest for a prolonged time, which makes them promising tools for environmental monitoring applications. For exam-

ple, Kovacina et al. [76] proposed a decentralized control algorithm for localizing gas sources and mapping chemical clouds within a region. This approach relied on constrained randomized behaviors and attended to the UAV restrictions on sensors, computation, and flight envelope. Later, Bermúdez et al. [60] investigated the use of blimp-based gas-sensitive UAVs for demining tasks, including strategies for chemical mapping. Recent projects like *AirShield* (airborne remote sensing for hazard inspection by network enabled lightweight drones) [15] investigate the use of autonomous swarm of micro UAVs to support emergency units. For a more detailed review of the state of the art in this field see [113].

2.4.4 Gas discrimination

Gas discrimination deals with the problem of identifying to which of a set of categories a new volatile sample belongs [147]. Traditionally, this process is carried out by hosting gas sensors inside a chamber with controlled humidity, temperature and airflow conditions, as well as regulating the exposure time to the gas and its concentration. Under these conditions, many studies have been proposed on how to classify odours using an array of gas sensors and a pattern recognition algorithm. In [136, 53] the principal methods for chemical classification are reviewed, ranging from classical methods such as k-nearest neighbour (kNN), Mahalanobis linear discriminant analysis, or Bayesian classifiers to most recent artificial neural networks (ANN), cluster analysis with self-organizing maps (SOM) and support vector machines (SVM).

However, when the discrimination is performed with a mobile robot equipped with an e-nose, there is no control over the sensing conditions. This entails that the sensor signals to be processed are noisy and dominated by the signal transient behaviour [145], which entails a number of additional challenges with respect standard analyte identification. Few are the works found in literature that perform classification focussing only on the transient phase of the sensor signals. Probably, the first work addressing this problem was [100], where Marques et al. proposed a feature extraction method based on discrete wavelet transform (DWT). Here, the authors claimed that only 4 seconds of exposure of an array composed of 4 MOX sensors were sufficient to reliably classify from a set of 6 different gas mixtures. Later, Martinez et al. [103] proposed a biomimetic robot for tracking a specific gas plume, dealing with the discrimination of gases by means of a spiking neuronal network. An evaluation for the suitability of different feature extraction techniques for such scenarios was provided in [147], where Trincavelli et al. proposed a preprocessing stage to isolate the relevant parts of the sensor signals that can then be passed to the pattern recognition algorithm. More recently, in [24] a support vector machine was applied to a set of features obtained from changes of the spectral sensor signal characteristics (frequency components, phase shift and energy sums), reporting a substantially increase of the classification performance. For a more detailed review of the field see [145].

Chapter 3
Improving MOX-based gas sensing for mobile robots

Metal Oxide Semiconductor (MOX) gas transducers are one of the preferable technologies to build electronic noses because of their high sensitivity and low price. In this chapter we present two approaches to overcome to a certain extent one of their major disadvantages: their slow recovery time (tens of seconds), which limits their suitability to applications where the sensor is exposed to rapid changes of the gas concentration. The first approach is based in the design of a new e-nose called Multi-Chamber Electronic Nose (MCE-nose), which comprises several identical sets of MOX sensors accommodated in separate chambers (four in our current prototype), alternating between sensing and recovery states, providing, as a whole, a device capable of sensing faster changes in chemical concentrations. The second proposal consists of exploiting a double first-order model of the MOX-based sensor from which a steady-state output is anticipated in real time given measurements of the transient state signal. This approach assumes that the nature of the volatile is known and requires a pre-calibration of the system time constants for each substance, an issue that is also described in this chapter. The applicability of both approaches is validated with several experiments involving rapid sensing of gas concentration in real and uncontrolled scenarios with a mobile robot bearing an e-nose.

3.1 Introduction

The deployment of olfactory sensors is becoming an increasing practice in many industrial and environmental applications due to advances in the gas sensing technology. The exploitation of olfactory sensors can be classified into two main groups according to the level of control over the measurement conditions: Closed Sampling Systems (CSS), where the gas sensors are usually hosted in test chambers with controlled airflow, volatile exposure times, temperature and humidity, *etc.*, and Open Sampling Systems (OSS), with no control over the sensing conditions. Our interest is in the latter, which are more flexible and practical for field applications. Examples of such uses are environmental exploration [148], gas distribution modeling [89], buried land mine detection [122] or pollution monitoring [149]. Some of these applications are usually accomplished with the help of a mobile robot carrying the sensors on board, which makes the sensing task even more challenging.

Within the different technologies and materials available for gas sensor fabrication [119], MOX (Metal Oxide Semiconductor) transducers are one of the most popular and widely employed in mobile robotics olfaction, due to their high sensitivity and low prices. However they present some shortcomings including poor selectivity, response drift (age factor), influence by environmental factors such as humidity and temperature [107] and major limitations in their response speed [49]. These limitations come from the sensing mechanism underlying MOX technology, that is, the exchange of oxygen molecules between the volatile and the metal film [2, 135].

Among these drawbacks, the long duration of the acquisition cycles (up to tens of seconds) is of especial concern for OSS, since inaccurate readings are inevitable when measuring rapid changes of gas concentration, as illustrated in Figure 3.1. Observe how this limitation is particularly noticeable in the decay phase, when the output recovers to the baseline level (the steady output value given by a gas sensor when exposed to clean air). As a consequence of this slow dynamic response and because of the intermittent and chaotic nature of turbulent airflow in OSS [139, 91], steady state values are rarely reached, and therefore gas sensing based on MOX technology must deal with the transient information of the signals [147]. This problem becomes crucial when the sensors are carried on a vehicle (typically a mobile robot) to provide measurements along the way. The adopted solution in such cases is to reduce the vehicle velocity to a few cm/s, such in [68]. This proposal, however, is not acceptable in many applications since the sampling of space must be as quickly as possible to cope with the rapid dynamics intrinsic to gas propagation.

To overcome this shortcoming of MOX-based electronic noses, in this chapter we propose two different approaches: one hardware and one software. The former is based on the design of a new MOX-based e-nose which comprises several identical sets of MOX sensors accommodated in separate chambers. It enables the sensing of faster changes in chemical concentrations by commuting the sensing task to another clean set of sensors when the decay phase is detected. The latter estimates the steady state sensor output from the noisy and distorted transient signal, which corresponds to the search for an inverse dynamical model of MOX sensors. The applicability of

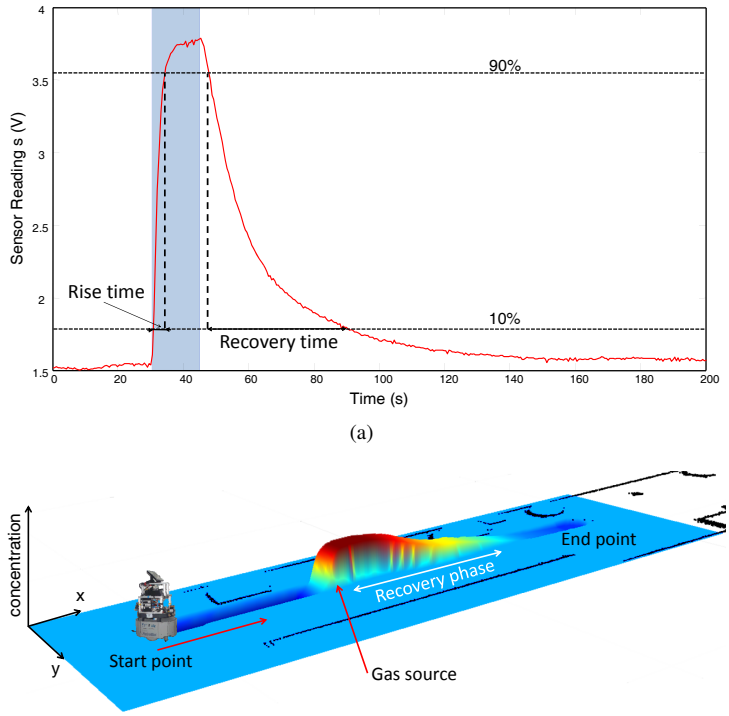

Figure 3.1: Rise and recovery phases of MOX sensor response to a step gas concentration. Subfigure (a) shows a 2D plot of the sensor response over time. The shaded blue region denotes the sensor exposure to the analyte. Subfigure (b) depicts a 3D gas distribution map generated from the readings of an MOX sensor carried by a mobile robot along a corridor. Observe how the recovery phase after the gas exposure is several times longer than the rise one.

both approaches to OSS is validated through different experiments in real scenarios, performed with a mobile robot bearing an MOX-based e-nose.

The rest of this chapter is outlined as follows: after a discussion of the important influence of the e-nose slow recovery in mobile robotic olfaction in Section 3.2, we introduce the proposed sensor configuration, called Multi-Chamber E-nose (MCE-nose, for short) in Section 3.3. Then, Section 3.4 describes the software alternative by introducing the proposed MOX model, giving especial attention to the experimental results. Finally, we end up with some conclusions and discussing future research.

3.2 On the importance of the e-nose long recovery time in mobile robotics olfaction

For a mobile robot intended to accomplish olfaction-related tasks, the problems associated to the slow recovery of MOX gas sensors are manifested, among others, through the following issues:

1. A gas concentration may be masked by another close, stronger one. Suppose two gas sources of different concentrations, separated by a short distance. If the robot trajectory first leads to the lower-concentration gas source, both of them will be probably detected. However, if it happens the other way around, the lower one may be overlooked since it could be hidden below the decay of the stronger concentration. Figure 3.2 displays a simulation of such scenario.

2. Gas concentration maps are not accurate, as a consequence of the integration into the map of unreliable sensed values from the decay phase of the sensor response.

3. Gas source search methods that rely on gradient techniques may not be applicable. These methods require to measure and compare the gas concentration at different points, either successive readings (klinotaxis[1]) or simultaneously sensed intensities from two or more sensors (tropotaxis[2]) as in [86]. For the first case, we cannot trust in the sensor measurement if it is still in the decay phase of the previous sensing.

Thus far, mobile olfaction tasks have managed this limitation in, basically, two ways:

- Slowing down the robot speed up to a few cm/s in order to allow the sensor response to slowly follow the gas distribution even in the decaying phases [68].

- Defining paths that force the robot to pass several times over the same locations but along different directions, in such a way that the decay effect is averaged out over all the measurements. This is a common strategy employed to explore a space with the intention of building a gas concentration map, such as in [96].

Clearly, this type of solutions affect the overall efficiency of the olfactory task and, in many cases, it may be even unacceptable for the robot mission. It is important

[1] Klinotaxis is the achievement of orientation by alternate lateral movements of part or all of a body; there appears to occur a comparison of intensities of stimulation between one position and another and a "choice" between them. Klinotaxis is shown by animals with a single intensity receptor such as the protozoan Euglena, earthworms, and fly larvae.

[2] In tropotaxis, attainment of orientation is direct, resulting from turning toward the less stimulated (negative) or more stimulated (positive) side as simultaneous, automatic comparisons of intensities on two sides of the body are made. No deviations (trial movements) are required. Tropotaxis is shown by animals with paired intensity receptors.

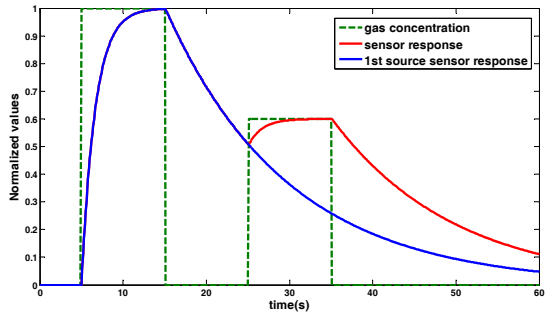

(a)

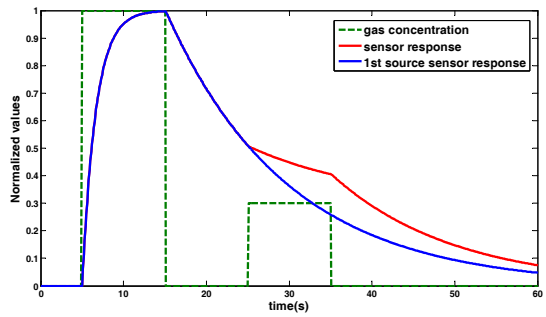

(b)

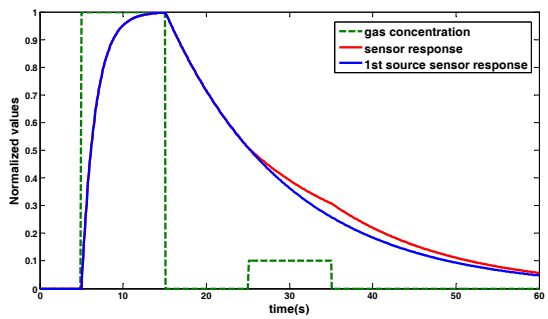

(c)

Figure 3.2: Simulations of the behavior of a MOX sensor when sensing a low gas concentration right after being exposed to a stronger one. The MOX sensor response has been modeled as a double first-order system with time constants 1.7s, and 14.8s, for the rise and decay stages respectively (estimated from system identification techniques). Observe that, when the second gas source is much lower than the first, the response of the MOX sensor (in red) is very similar to that obtained from the first source alone (blue).

to remark that, for most real robot applications, smell is not by itself the ultimate goal for the robot, but just another of the robot senses to gather useful information from the environment (along with vision, range sensing, touch, etc.).

3.3 The multi-chamber electronic nose - an improved olfaction sensor for mobile robotics

It is well known that wine testers have a very developed and well trained sense of smell. In a typical session, to avoid mixing the smells or tastes of different wine samples, they have to clean their mouths and noses by eating a little piece of bread and wiping their noses, for example. Thus, they undergo a "purge/clean" stage between tests and they also stop for a few seconds to ensure their noses are ready to provide new accurate olfactory information. MOX gas sensors behave in a quite similar way, as they require a time (decay phase) to ensure their readings are accurate.

The MCE-nose proposed in this chapter pretends to work in similar way that wine testers, but taking advantage of the reproducibility of electronic devices to avoid the off-time between readings. Thus, the key idea behind the proposed design is to ignore the MOX sensor output when the decay phase is detected and delegate the sensing task to another clean, almost identical sensor. In order to achieve that, we accommodate a set of redundant sensors in different chambers, which are alternatively activated. Thus, the output signal of the whole setup results from the concatenation of the rise phases of a sequence of MOX sensors.

The design of the MCE-nose aims at providing the following characteristics:

- To shorten the cycle of effective sensing as exposed above.
- To recognize a variety of odors by hosting MOX sensors with different selectivity in each chamber.
- To reduce the influence of residuals from previous measurements by scaling down both the chamber room where the sensors are accommodated and the air circuit volume.
- To speed up the interchange of molecules onto the MOX film by feeding a pressured air flow into the chamber by means of a pneumatic pump.

Next, the three main aspects of the MCE-nose design are exposed: mechanics, electronics, and software.

3.3.1 Mechanical design

Figure 3.3 shows a schematic diagram illustrating the interconnections of the different components of the proposed e-nose. The design is conceived to comprise a general number of M chambers with N MOX sensors each. All chambers are identical and

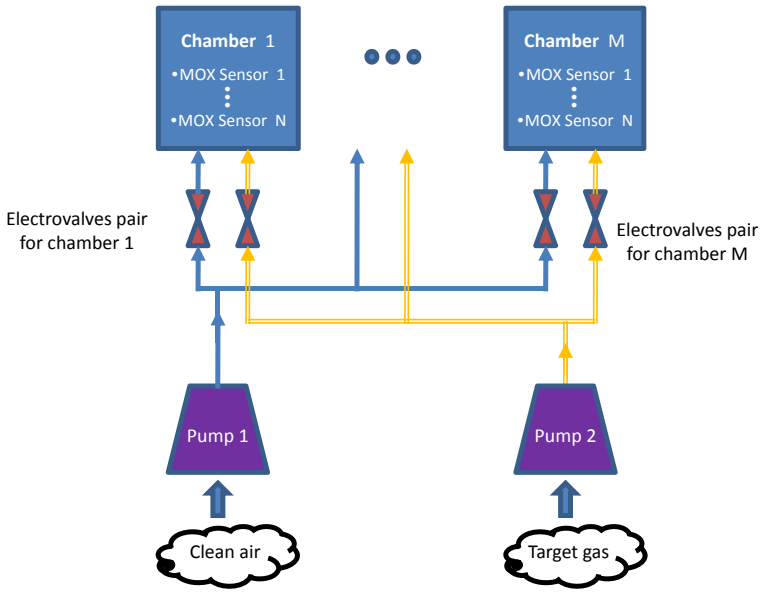

Figure 3.3: A functional schematic diagram of the MCE-nose. There are two pumps: one aspirating clean air and the other the target gas. At each time, only one chamber is receiving the target gas while the other M-1 chambers are being purged with clean air.

contain the same set of sensors. Chambers are also isolated from each other, that is, no airflow circulates between them.

There are two pneumatic circuits: one for clean air and one for the target gas (i.e. odor charged), which are connected to each chamber. Clean and contaminated air flows are taken from opposite sides of the MCE-nose device via two separate pumps. Besides, clean air is forced to flow through an active carbon filter to eliminate possible impurities.

At any given time, only one chamber is fed with the target gas, while the others M-1 are being cleaned. This is done thanks to a set of electro-valves placed at the entrance of each chamber, controlled by embedded software built in the MCE-nose micro-controller, as will be described later in this section.

At any time, each chamber can be found in one of the following three states:

Clean: A chamber is said to be "clean" if all of its MOX sensors are at their baseline level. This may happen because either the chamber has not being used yet for sensing or because it has been injected with clean air long enough.

On-Cleaning: Opposite to a clean chamber, an on-cleaning one is that whose sensors are not completely cleaned (i.e. they have not reached the baseline yet), despite the chamber is being injected with clean air.

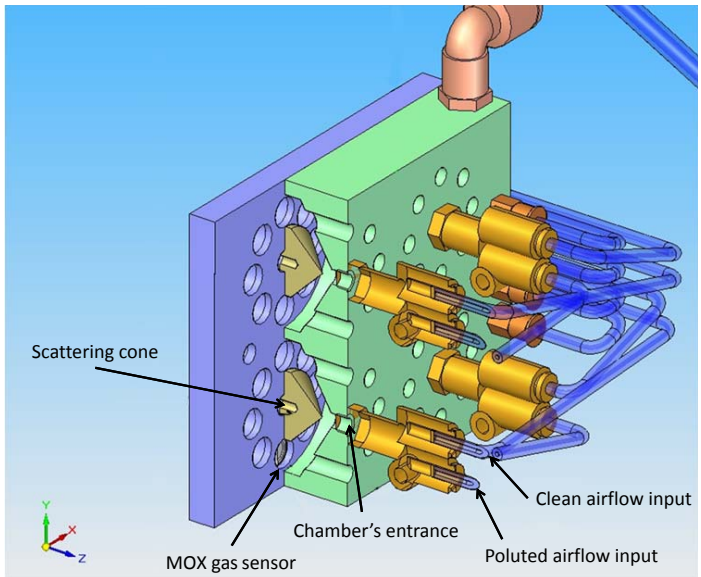

(a)

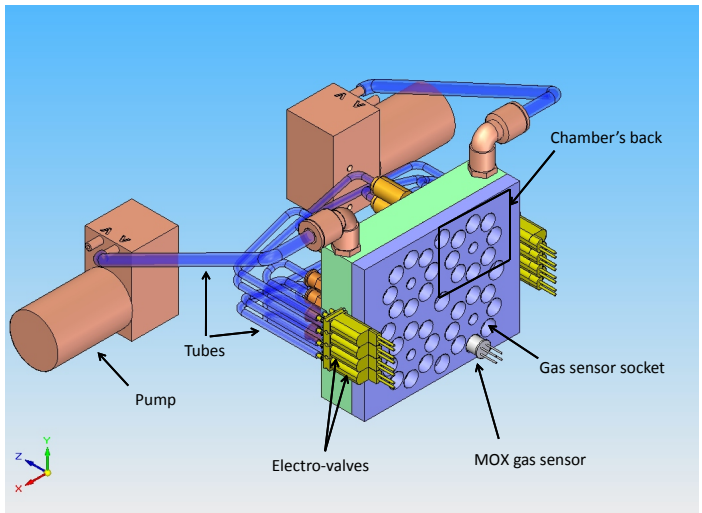

(b)

Figure 3.4: Different views of the 3D model (a) upper view, (b) bottom view, of the pneumatic circuit and the main block containing four chambers which can accommodate up to 8 MOX sensors each.

Active: The chamber is being injected with the target air.

Figure 3.4 shows some of the 3D models created for the current prototype, which consists of four chambers with identical configuration which can accommodate up to 8 MOX sensors each. Our choice of such particular number of chambers obeys to a trade-off between two issues:

- On the one hand, the obvious higher cost and complexity of the device as this number increases: more sensors, valves, A/D converters, etc. as well as problems for dissipating heat on the PCB, power consumption, etc.
- On the other hand, the possibility of having an array of sensors at the baseline level and, consequently, the possibility of sensing at a higher frequency.

The main block, which accommodates the 4 chambers, has been fabricated of resin with a stereolithography machine. Each chamber has a circular array of 8 sockets to lodge MOX sensors of standard size (8 mm diameter). It can be appreciated in Figure 3.4(b) how the sensors are hosted. They are introduced from the bottom side of the main block, leaving the sensing surface inside the chamber and, at same time, facilitating the electronic connections (pin soldering). A cone at the entrance of the chamber scatters the incoming airflow evenly directing it towards the active sensing surface of the sensors. The air is then forced to escape through the upper orifices of the chamber, as illustrated in Figure 3.5.

Each of the 8 sockets can lodge a different sensor. In our case, each chamber contains 7 different MOX sensors, with the extra socket employed for a temperature

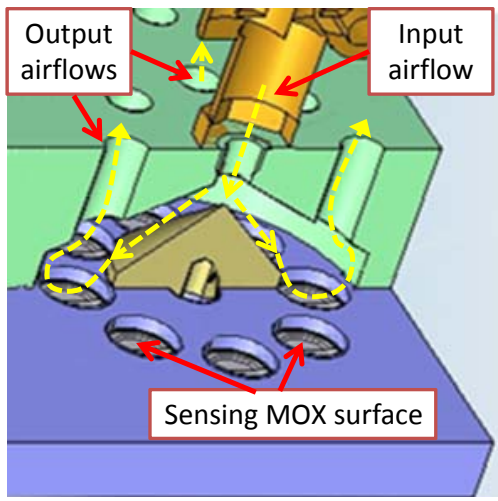

Figure 3.5: Approximate airflow scheme inside each chamber.

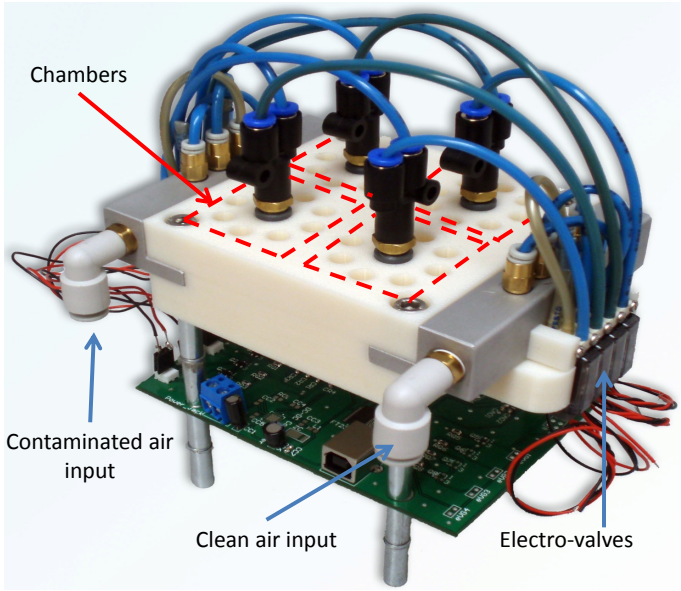

Figure 3.6: The complete MCE-nose. The current prototype contains 4 chambers, hosting 8 different MOX sensors each.

sensor (ADT7301). The 7 MOX sensors were selected with different gas selectivity in order to facilitate odor classification. This amount of sensors has demonstrated to be large enough to allow the recognition of a wide range of odors.

In our prototype, the pumps mounted are EAD NEO IP3 diaphragm pumps: 15V dc, 180kPa maximum attainable pressure, and working flow of 4 lpm. For each of the chambers, two SMC S070C6BG32 electro-valves are used: one for the clean and one for the polluted air flow. To interconnect pumps, electro-valves and chambers, we have used standard pneumatic PVC tubes with diameters of 8 and 3 mm, as well as the required plugs. Figure 3.6 shows a picture of the built prototype.

3.3.2 Electronics

Any conventional MOX-based e-nose requires a minimum of electronics to cope with sensor pre-heating and sensor readings, including signal conditioning and A/D conversion. In our design, the electronic module has to take care also of the synchronization of the pneumatic circuits by controlling the eight electro-valves (one pair for each chamber). As seen in Figure 3.7, such electronics has been mounted on a single printed circuit board (PCB) which is connected to all the components by means of four 16-pins connectors (for the gas and temperature sensors) and eight 2-pins connectors (for the electro-valves).

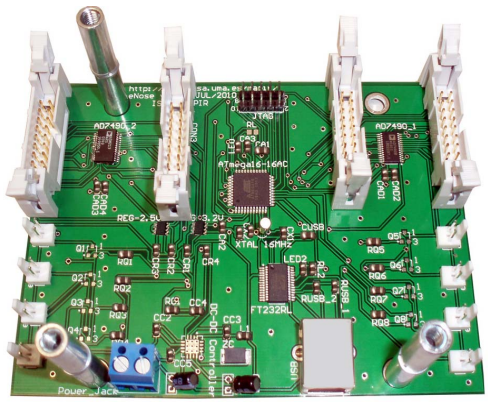

Figure 3.7: PCB where all the electronic components have been mounted.

The core component of the PCB is an ATMega16 8-bit micro-controller at 16 MHz, which provides 32 programmable I/O lines to control two A/D 16-bits converters (connected to the gas sensors), four temperature chips (placed inside each chamber to measure working temperature), and the eight electro-valves. Additionally, the PCB comprises a USB connection to a PC host for easy interfacing and a standard JTAG interface for development.

3.3.3 Embedded software

The firmware we designed for the ATMega16 micro-controller is in charge of controlling the behavior of the MCE-nose components. The operation flow is based on three main stages, described in Figure 3.8:

- The first stage checks if a data frame containing the information about the next active chamber is received from the PC. If this is the case, the appropriate signals are issued such that the electro-valves switch the airflow into the newly selected active chamber. Notice that the switch strategy that dictates the active chamber at any given time has not been embedded into the micro-controller, but it relies on orders from the computer. This decision obeys to our interest in implementing high-level switching strategies that may take into account information from other sensors and the robot task.

- The second stage is in charge of collecting the readings from all the sensors of the MCE-nose (28 MOX and 4 temperature sensors in our case). This is done by means of two A/D 12-bits converters of 16 channels each.

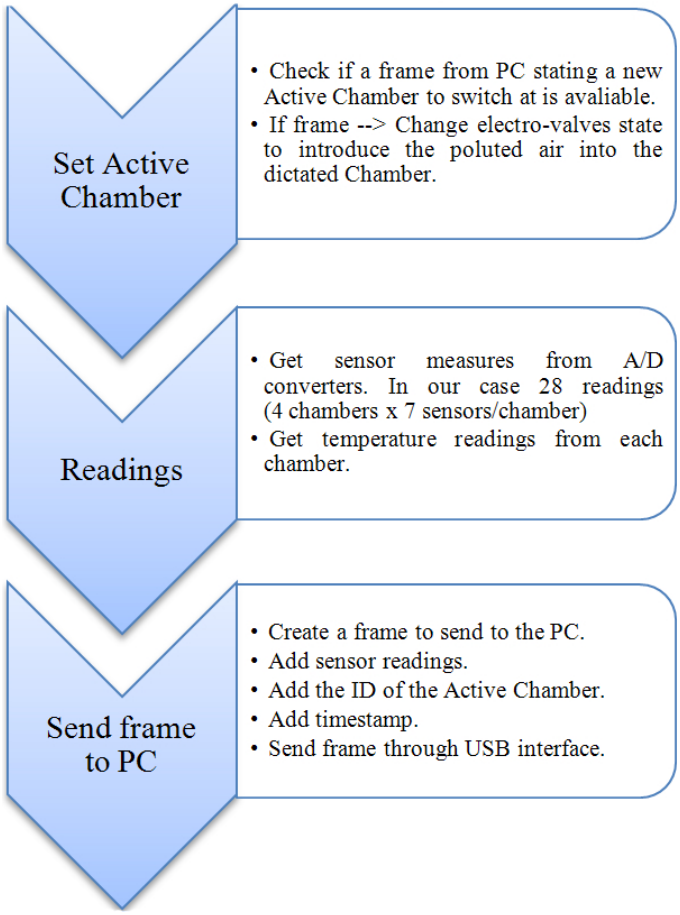

Figure 3.8: Operation flow of the embedded software.

- Finally, all the collected data are packed into one frame, which is assigned a timestamp and the ID of the active chamber. This data frame is then sent to the PC via a USB-to-serial UART interface (FT232RL).

3.3.4 Calibration of gas sensors

As depicted above, the output signal of the MCE-nose results from the concatenation of the rise phases of identical MOX sensors, placed in the different chambers. Nevertheless, in practice, such identical sensors do not response the same and thus, a calibration is required in order to make their responses as similar as possible. For

such calibration, we have to compare the readings of all chambers when exposed to the same concentration.

To ensure that all chambers are flooded with the same gas concentration, the four chambers where individually and sequentially flooded during 60 seconds, allowing their sensors to reach the steady state (see Figure 3.9).

Since only the baseline and the rise phase of each sensor are of interest for the MCE-nose output (as the decay phases are discarded), we compensate outputs of sensors in chamber 1, 2 and 3 to achieve the baseline level and the amplitude of the reference output (chamber 0). Concretely:

- An offset was added to each sensor output to force them to have a similar baseline level when sensors are not exposed to target gases.
- A multiplying factor was estimated for each sensor to ensure identical gain. To account for the non-linear behavior of the sensors we selected an average gain computed from three different concentrations.

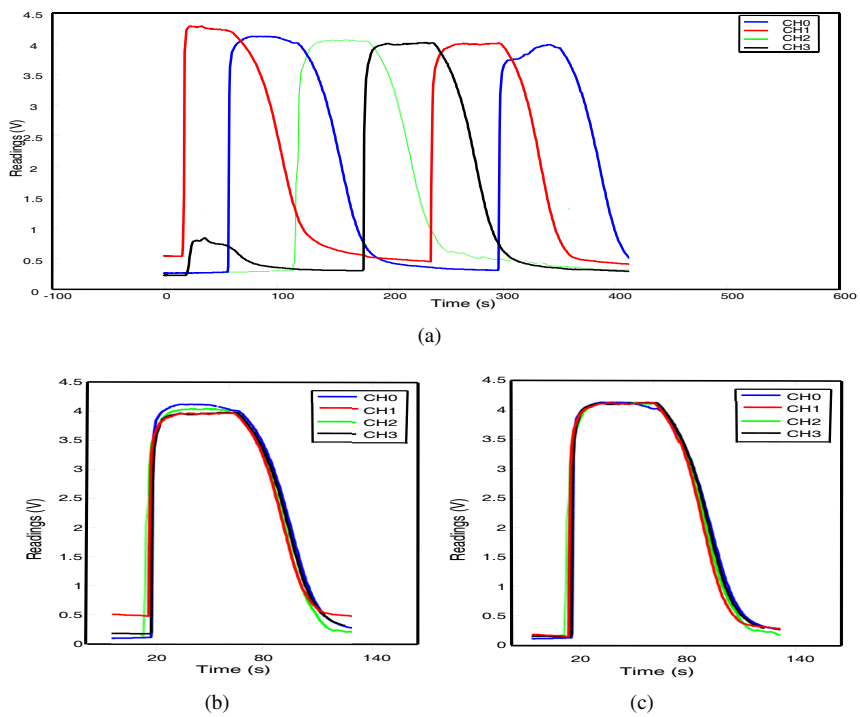

Figure 3.9: (a) Readings of four TGS-2602 sensors placed in each chamber of the MCE-nose prototype during the calibration procedure. (b) Comparison of the four sensor readings before calibration, and (c) after it.

Figure 3.9(a) shows the readings of TGS-2602 sensors placed in each chamber of the MCE-nose prototype during the calibration procedure. It may be notice that even before calibration the readings of the four sensors are all very similar (as reasonably expected). Figure 3.9(c) plots the readings of the same sensors after the calibration has been carried out.

3.3.5 Integration of the MCE-nose into a mobile platform

The MCE-nose presented in this chapter has been designed to be integrated into a mobile robot. Figure 3.10 shows a PatrolBot mobile platform [106] with the MCE-nose already integrated into it. The robot is also equipped with a SICK and a Hokuyo laser range scanners and a sonar ring to provide the necessary functionality for localization and obstacle detection. For a detailed description of the mobile platform and the components onboard, see Appendix A.

One of the main advantages of the MCE-nose is its suitability for mobile olfaction tasks. The mechanical design of the MCE-noise opens a variety of possible configurations:

- It can work either as a MCE-nose (as explained above) or as a conventional e-nose by using only one of the chambers. This may be convenient in some phases of an olfaction task (e.g. odor classification).

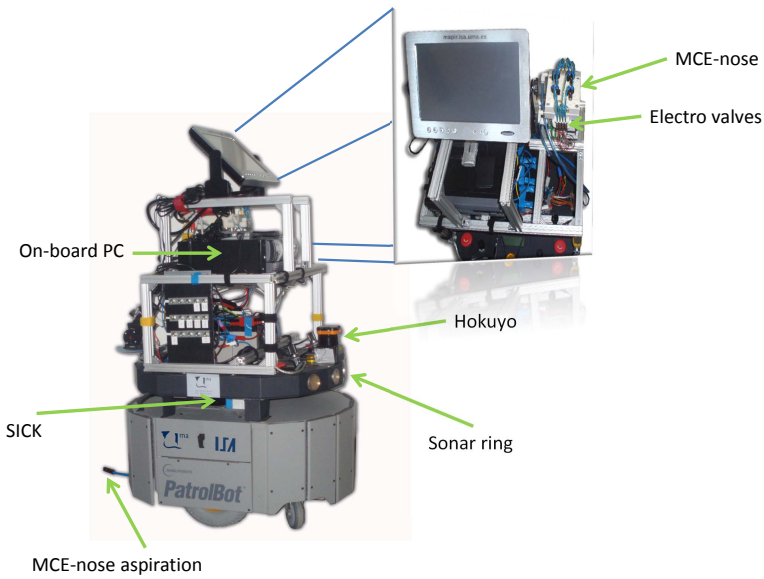

Figure 3.10: The MCE-nose integrated in a mobile platform Patrolbot mobile base.

- Since the aspiration is carried out through a tube, the air input can be conveniently placed at any point around the robot. This allows the MCE-nose to be mounted at any place on the platform, no matter of its shape or size. Also, olfaction strategies that need to compare concentrations from several points around the robot (as gradient techniques) are easily accomplished by just moving the aspiration tube, for example, with a servo motor. Even if no comparison is needed, having such capability bears some advantages: 1) we are not limited by the robot nonholonomic constraints while sampling the workspace, and 2) we reduce the air disturbance caused by the robot movement to a minimum, since we reach the target point with the tube which generates a negligible turbulent airflow.

Considering the possibilities offered by a MCE-nose integrated into a robotic platform, it is necessary to account for high level software able to exploit such potential for any robotic olfaction task. These possibilities include: switching between chambers, focusing only on some specific (more suitable) MOX sensors from the array, taking into account the robot mobility as well as surrounding information from other sensors of the robot (laser scanner, sonar, ...), etc.

Among others, this software has to deal with the following tasks:

1. To detect abnormal level of a gas (probably while accomplishing a non-specific olfaction mission), through a pilot "watchdog" sensor from the MCE-nose. This could be done instead, by a static gas sensor network deployed in the environment.

2. To classify the target gas. MOX sensors have low selectivity, so the multivariate response of an array of chemical gas sensors with broad and partially overlapping selectivity can be used as an "electronic fingerprint" to characterize a wide range of odors or volatile compounds by pattern-recognition means [53]. For this task, typically only one chamber is necessary, thus no chamber switch is required. As an illustrative example, Figure 3.11 shows the responses to a specific odor of seven different MOX sensors within one chamber.

3. Measuring the target gas concentration is crucial for almost all robotic olfaction tasks, including gas source localization and gas mapping. With the purpose of obtaining the best estimation of such concentration, is advisable to select, from the sensors of each chamber, those more sensitive to the target gas. Referring to Figure 3.11, sensors TGS-2620 and TGS-2600 are good candidates for gas concentration purposes due to their high sensitivity to that gas.

4. To control and manage complex switch strategies which could take into account not just the gas sensor readings, but also information provided by other sensors (laser scanner o camera), as well as the olfaction task at hand (e.g. plume detection, gradient following, etc).

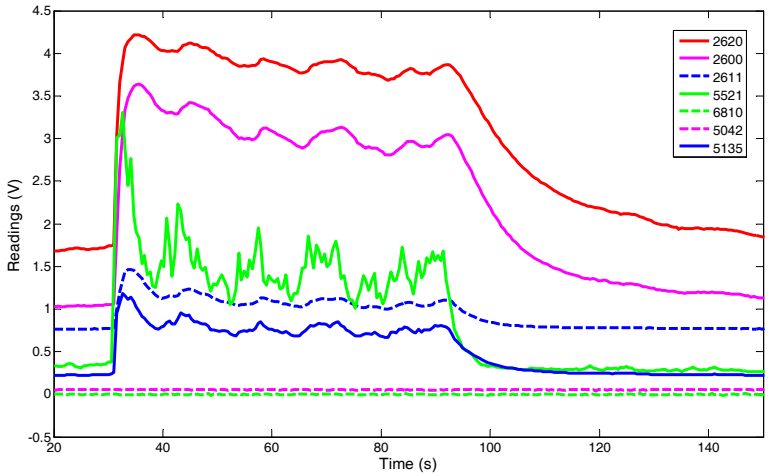

Figure 3.11: Readings from seven different MOX sensors within a chamber when exposed to acetone.

Such software has been implemented under the Open Mobile Robot Architecture (OpenMORA) [52], based on MOOS [114] and MRPT [108]. This architecture allows us to easily control a robot platform and the available sensors as range lasers, cameras or sonar, as well as providing high level functionality as obstacle avoidance, autonomous path planning or localization.

3.3.6 Experiments with the MCE-nose

This section describes different experiments we have carried out to validate the MCE-nose with regard to the improvement in rapid sensing of gas concentrations. The experiments consist of a static smell test, a mobile experiment with multiple gas sources, a mobile test with different gas concentration sources and finally a gas mapping experiment. Since the kind of gas to sense was known a priory, neither odor classification nor sensor selection was required here. The implemented switch strategy is based on two rules for deciding when to switch and what chamber to switch to:

Rule 1: A switch of chamber must happen whenever the sensor readings from the current active chamber (being fed with the input stream) start to decay.

Rule 2: Provided a switching event has been triggered by rule 1, it is necessary to check the state and sensor levels of all the M chambers (clean, on-cleaning and the active one). The one with the lowest sensor readings is chosen to be the next chamber to commute to.

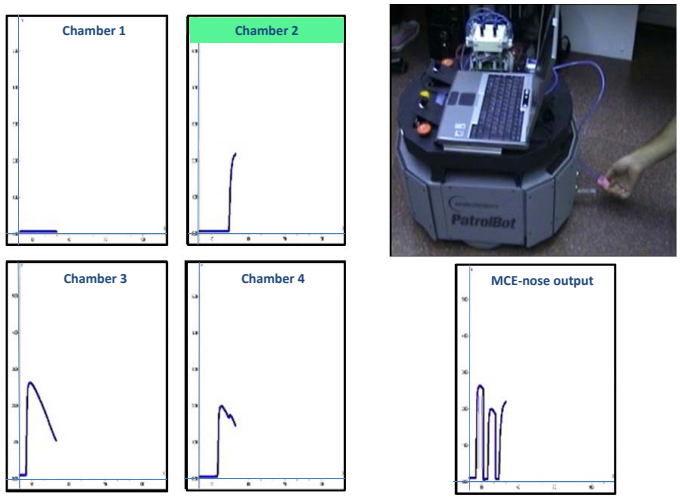

Figure 3.12: Snapshot of the MCE-nose static smelling experiment. The four plots on the left side present the readings of each of the four chambers of our current prototype, while the MCE-nose output is shown on the bottom-right plot. The active chamber is marked in green (chamber 2 in this case).

Static test

In this experiment the robot was kept still, being the gas source (composed by a small cup filled with acetone) the mobile element. The experiment consisted in repeatedly presenting the gas source to the MCE-nose air input, waiting a few seconds and moving it apart. Figure 3.12 shows a snapshot of the experiment and the responses obtained with every chamber (conventional e-nose) and with the MCE-nose. It can be appreciated how the MCE-nose output is able to capture the (three) different exposures by changing to a clean chamber whenever the response of the active one (being odor flooded) starts decaying.

Detecting multiple odor sources

The second experiment was designed to test the behavior of the MCE-nose in the case of multiple gas sources in a more realistic robotic scenario. The scenario consists of a long corridor where three equal-sized small cups filled with acetone were placed at 2 meters from each other. Figure 3.13 displays the experiment setup, and a picture of the MCE-nose integrated in the PatrolBot platform. For the experiment the PatrolBot was commanded to move in a straight line at a constant speed of 20 cm/s. Figure 3.14 illustrates the comparison between the outputs of a conventional e-nose (one chamber) and the MCE-nose.

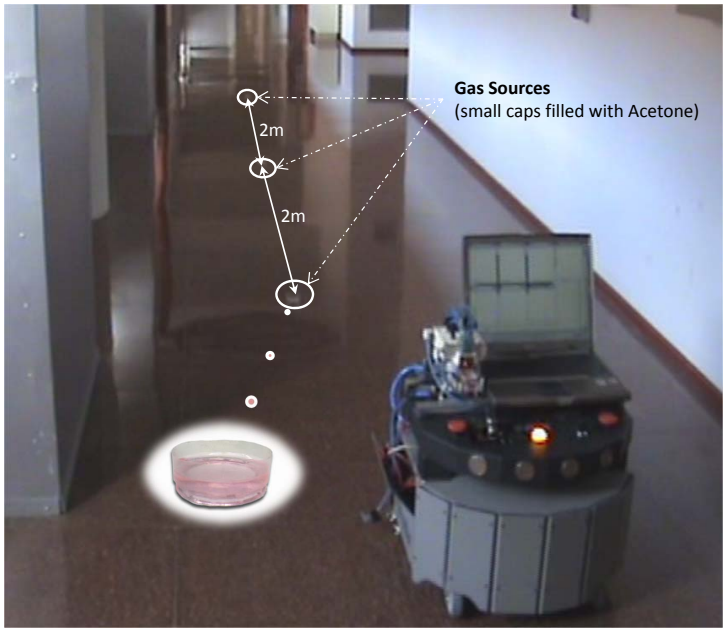

Figure 3.13: Description of the multiple gas source experiment. Three small cups filled with acetone where placed along the robot trajectory to test the behavior of the MCE-nose.

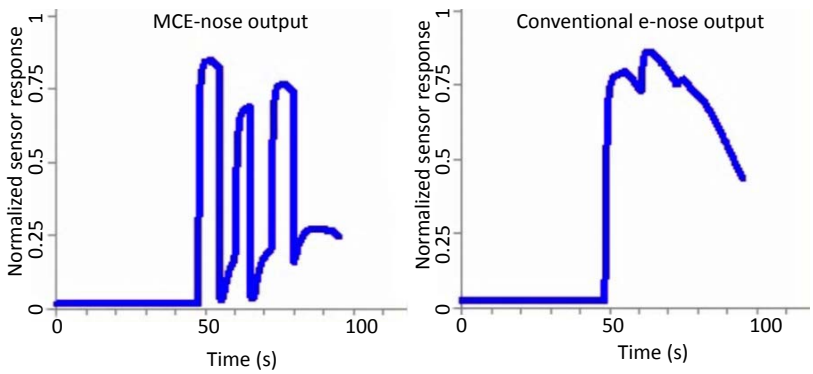

Figure 3.14: Comparison between the readings of a conventional e-nose (right) and the MCE-nose (left) for the multiple gas source experiment. It can be appreciated how the MCE-nose can clearly distinguish the three gas sources presented along the robot path, while the conventional e-nose can hardly detect the second source, and how the third one became completely unnoticed.

Notice that for such a robot speed, the readings provided for a conventional e-nose do not reveal the presence of the three odor sources and the low concentration zones between gas sources are not correctly gauged. The common solution to this problem would be to slow down the robot speed, so the MOX sensors could have time to recover their baseline level, which is not possible or practical in many real robotic applications. Observe, on the other hand, that the MCE-nose is able to provide more accurate measures.

Detecting multiple odor sources of different concentrations

The objective of this experiment is to demonstrate that using the MCE-nose, the problem of disguising lower concentrations or even additional gas sources (as stated in Section 3.2), can be notably palliated.

The experiment was carried out in the same scenario than the previous experiment. In this case, only two gas sources separated one from each other 2m were used. The first one was a wide open vessel (approximately 15cm diameter), while the second one was a small cup (4cm diameter) covered by a grid lid to reduce the gas dissipation. Using this setup, two gas sources of different concentrations were presented to the robot along its path.

Figure 3.15 shows the raw readings of the experiment. These values (after a previous normalization) were then applied to an ICP based SLAM process to generate the map shown in Figure 3.16. We must remark the improvement in the detection of a low concentration source after a high one. From a comparison of the "peak" concentrations from the MCE-nose and the conventional e-nose in Figure 3.16, one may wonder why in the former case the peak seems to extend in a larger area. How-

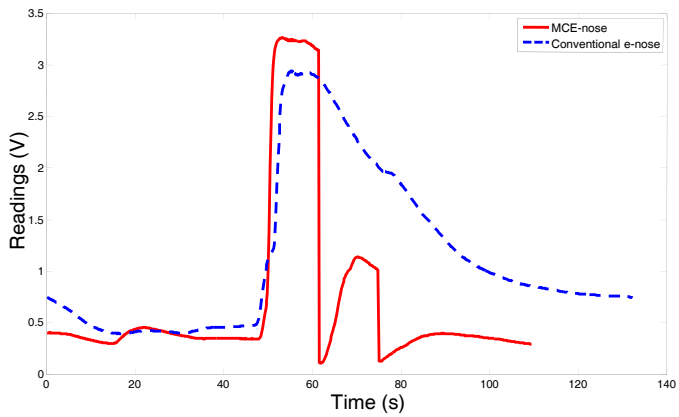

Figure 3.15: 2D comparison of the raw readings between a conventional e-nose (dashed blue) and the MCE-nose (solid red), when faced to two sources of different concentration.

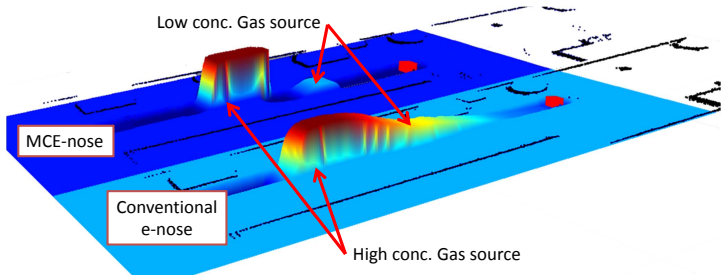

Figure 3.16: 3D comparison of the ICP SLAM generated maps between a conventional e-nose and the MCE-nose when faced to two gas sources of different concentration.

ever, observing the raw readings in Figure 3.15, it becomes clear that the chamber's switch in the MCE-nose takes place as soon as the decay phase starts. Thus, the observed differences are only due to the real differences between experiment repetitions. The MCE-nose switches to a different chamber when the readings from the active chamber decay bellow a threshold. This threshold was set to 0,1volts in the current experiment to avoid miss-switches due to noise or spurious readings. Decreasing the threshold value would mean faster switching after a gas source is detected, but it could then produce non-desired switches due to noise, spurious or because of the small fluctuations inherent in MOX sensors.

Gas distribution mapping

The objective pursued with this experiment is to analyze the performance of the MCE-nose when creating a gas distribution map of a room. A gas source composed by a 10×2 cm container filled with Acetone was placed in a 6×4 meters empty room, next to a wall. The robot was commanded to move following a predefined set of way-points to force the MCE-nose to prove most of the space.

To be able to compare the results obtained in different trials, a methodology was established to ensure similar conditions in the room. Door and windows were kept closed during the experiments and sensors were conveniently preheated before operation. After each trial, the room was purged of residual gases by opening the door and windows, creating a strong airflow of clean air for at least 5 minutes.

Figure 3.17 shows a comparative between the MCE-nose and a conventional e-nose for three different robot speeds. Each map represents the gas distribution estimated in the room at the end of the robot trajectory, making use of the robot positions given by an ICP-based SLAM method and the Kernel DM+V algorithm [89]. It is important to keep in mind that these maps come from different runs of the experiment and, even though we have tried to reproduce the tests in the same conditions, it is inevitable the appearance of some gas patches from one test to another. In our opinion,

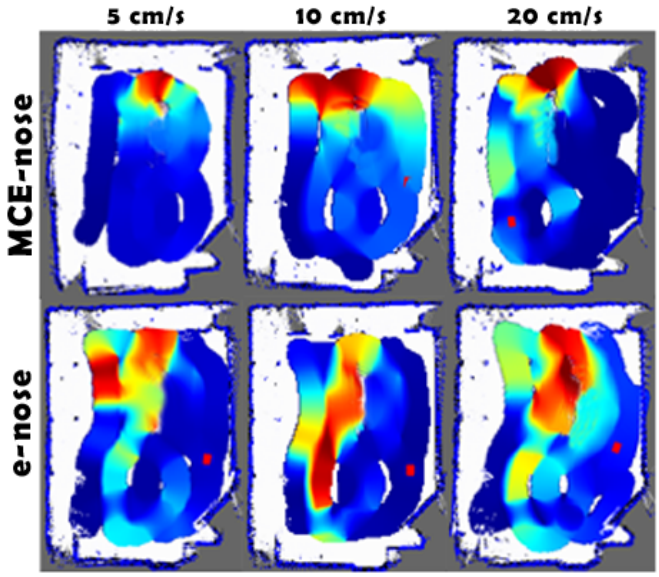

Figure 3.17: Comparison of the ICP SLAM generated maps between a conventional e-nose and the MCE-nose for three different robot speeds.

this explains, for example, the high concentrations near the source when using the MCE-nose at 10cm/s.

In spite of this consideration, it can be seen how the MCE-nose is able to localize the gas source more accurately than a conventional e-nose. This improvement is more apparent when increasing the robot speed, which allows it to perform a simple gas reconnaissance of the environment in a shorter time while obtaining results of high quality.

3.4 Overcoming the slow recovery of MOX sensors through a system modelling approach

In general, a model seeks to represent a system (empirical objects, phenomena, and physical processes) in a logical, objective and simplified way, allowing to predict the output of the system provided the input. Thus, a model of a MOX sensor must predict the sensor resistance (transient and steady state) when exposed to a certain gas concentration profile. Our interest in having such a model is to use it in a reverse way: given a sequence of measurements from the transient response of the MOX

sensor, we seek the exciting gas distribution through the estimation of the steady state sensor resistance.

The modeling of the dynamics of MOX sensors has been addressed in the literature for a variety of purposes. Gardner *et al*. [42] proposed a non-linear diffusion-reaction model to obtain the theoretical transient and steady state responses based on the reactions taking place at semiconductor level. This model is not applicable to our approach since we pretend to obtain the gas distribution from the sensor readings rather than from the electrical and physical properties of the sensors. Later, aiming at increasing the response speed of gas sensors, T. Yamanaka *et al.* [158] reported a two-phases (corresponding to the rise and decay phases) second order linear model to describe the transient response of a semiconductor gas sensor from a visualized gas distribution image. Despite its success, this model requires the use of a CCD camera as a gas detector, which is neither our case.

More recently, E. Llobet [94] reviewed the principal methods for dynamic analysis of the gas sensor response. Interestingly, the main use of these methods is to perform gas classification based on the transient response, from which a feature vector is extracted. For example, in [59] Box–Jenkins linear filters were applied to model an array of MOX sensors in the presence of four alcohols and water vapor with the aim of reducing the effect of the sensor drift in a classification process, and in [133] a multi-exponential transient spectroscopy (METS) method is proposed to improve the selectivity of chemical sensors in the analysis of gas mixtures.

Especial mention deserves some works that rely, as it is our case, on modeling the sensor response to predict steady state values from the initial part of the transient. In [111], a so-called ARMA and multi-exponential models are proposed for reducing the time necessary to calibrate a sensor array. Nevertheless, since the focus is on the calibration of MOX sensors, the dynamic models are only applied to the rise transient signals recorded in Closed Sampling Systems over long time periods (over 800 s), while we aim to predict the gas distribution profile in real time and in OSS. In [117], A. Pardo *et al.* propose and compare different nonlinear inverse dynamic models of gas sensing systems for quantitative measurements. However, the considered dynamic conditions differ from those of OSS. First, a measurement chamber is used to obtain the gas sensor readings, which implicitly modifies the dynamic properties of the measured signals, and second, the acquisition frequency is too low (one sample per minute) to reflect the fast and highly dynamic changes of the gas concentration in OSS.

Based on the multi-exponential model proposed elsewhere [111, 54], and taking into account the differences between rise and decay phases of MOX sensors, in this section we exploit a simplified version of it, where only one exponential is considered to model each phase (see Figure 3.18). According to this model, mathematically expressed in Eq. (3.1), three phases can be considered in the output of a typical MOX gas sensor when exposed to an ideal step in concentration: baseline, gas measurement (rise) and recovery.

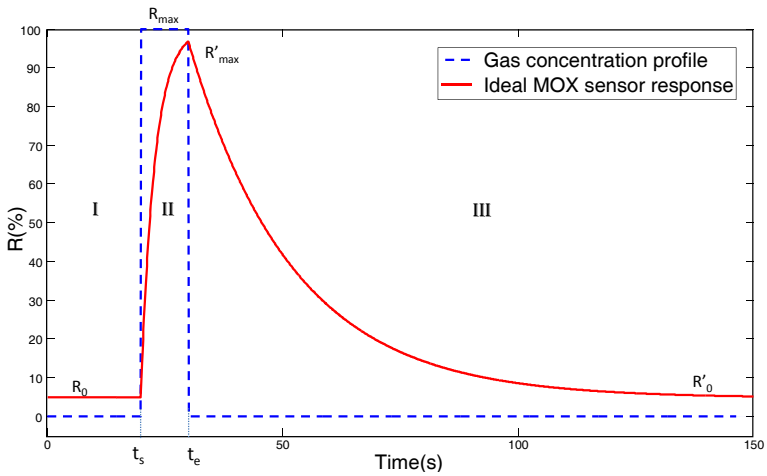

Figure 3.18: Ideal response of an MOX sensor (solid red line) when excited with a step gas concentration (dashed blue line). The curve shows the three phases of a measurement: (I) baseline, (II) gas measurement, and (III) recovery phase.

$$R(t) = \begin{cases} R_0 & t < t_s \\ R_0 + (R_{\max} - R_0)\left(1 - e^{\frac{-(t-t_s)}{\tau_r}}\right) & t_s < t < t_e \\ R'_0 + (R'_{\max} - R'_0)e^{\frac{-(t-t_s-\Delta t)}{\tau_d}} & t_e < t \end{cases} \tag{3.1}$$

where τ_r and τ_d are the time constants for the rise and recovery phases respectively, t_s and t_e represent the starting and ending times of the step excitation, R_0 and R'_0 are the sensor response level before and after the stimulus, R_{max} is the saturation level, and R'_{max} is the maximum response level during the gas measurement phase. Notice that R'_{max} is usually lower than R_{max} for short input pulses, as depicted in Figure 3.18.

This model was used by [81] for characterizing the response of an MOX-based e-nose carried by a robot. In this section we also exploit this model but making use of its inverted form, that is, to predict the gas distribution that the sensor is exposed to from its readings.

3.4.1 The proposed MOX model

As shown in the block diagram of Figure 3.19, three different sub-processes can be distinguished in an MOX-based gas sensing process:

1. A non-linear static system representing the measurement electronic circuit (Ω to V).

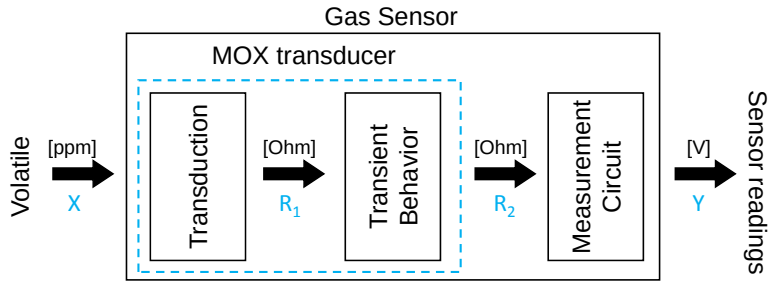

Figure 3.19: Block diagram of a smelling process with an MOX gas sensor. The sensor is excited by a volatile [ppm] producing a variation in the sensor resistance that is measured as an electrical signal [V] by means of a measurement circuit.

2. A transformation (Ω to Ω) that captures the non-linear rate at which the sensor resistance varies over time (although we have separated the transduction and dynamic phases to explicitly denote both functionalities, both stages take place within the MOX transducer and they are most likely coupled).

3. A signal transduction mechanism (ppm to Ω) which results from the chemical interaction between the sensor sensitive surface and the molecules of reducing gases.

Next, we model each of these stages and invert them to come up with a complete inverted MOX sensor model. It is important to remark that the proposed inverted model does not aim to recover the gas concentration (X[ppm]), which corresponds to a sensor quantification problem not addressed here, but only the gas distribution, providing relative results proportional to the gas concentration (R_1).

Measurement circuit

This stage stands for the electronic circuit in charge of measuring the changes in the sensor resistance, which typically consists of a simple voltage divider:

$$Y = V_{RL}(t) = \frac{V_{CC} \times R_L}{R_2(t) + R_L} \tag{3.2}$$

where V_{CC} is the circuit voltage and R_L is the load resistance.

Since the magnitude measured is the output voltage $Y(i)$, we can easily recover the MOX resistance changes $R_2(i)$ by solving Eq. (3.3).

$$R_2(i) = \frac{R_L \times (V_{CC} - Y(i))}{Y(i)} \tag{3.3}$$

Transient behavior stage

As reported by previous authors [49, 81], the transient response of an MOX sensor can be expressed by two first-order systems, as depicted in Eq. (3.1). It clearly resembles a low-pass filter response with the particularity that the filter cutoff frequency ($f_c = \frac{1}{2\pi\tau}$) is different for each phase the sensor is working at (rise or recovery). Equation (3.4) presents the two phases transfer function of this stage in the Laplace domain, where a common static delay (t_d) has additionally been considered in both phases to compensate for the delay introduced by the pneumatic circuit used to draw in the gas and flow it through the sensors:

$$\frac{R_2(s)}{R_1(s)} = \begin{cases} \frac{A}{\tau_r s+1} e^{-t_d s} & \text{,for rise phases} \\ \\ \frac{A}{\tau_d s+1} e^{-t_d s} & \text{,for recovery phases} \end{cases} \tag{3.4}$$

where s is the Laplace variable, A is the filter gain, $\tau_r and \tau_d$ are the filter time constants in the rise or recovery phase respectively, and t_d (in seconds) is the system delay. For nomenclature clarification $\tau_{recovery}$ has been denoted as τ_d, where the subindex d stands for decay.

Since our interest remains in estimating the volatile distribution that the sensor is being exposed to, given the sensor readings, we work out R_1 as a function of the sensor resistance measurements R_2 by applying the inverse Laplace transform to Eq. (3.4). As the transfer function has different expressions according to the phase (rise or recovery) that the sensor is working at, the parameters of the resulting differential equation will have to switch accordingly. Approximating the derivative of the measured sensor resistance R_2' by a backward first order finite difference, such differential equation can be written as:

$$R_1(i-N) \quad \propto \quad R_2(i) + \tau \frac{R_2(i) - R_2(i-1)}{\triangle t} \tag{3.5}$$

where $R_1(i)$ is the unknown steady resistance value for the given gas concentration at the time step i, R_2 is the measured sensor resistance, τ is the time constant for either the rise (τ_r) or the recovery phase (τ_d), N is the number of samples for the system delay, and $\triangle t$ is the time between samples. Notice also that the scale factor A from Eq. (3.4) has been dropped.

As observed in Eq. (3.5), at each time step this dynamic model requires to know the value of parameters τ_r and τ_d, and the phase the sensor is working at. The latter may be determined from the slope of the measured MOX resistance R_2, considering that the sensor is working under rise phase for positive values of the derivative, and under decay phase for negative values of it.

Notice that instead of working with the sensor resistance (which is inversely proportional to the volatile concentration), it may be more convenient to deal with the

sensor conductance ($G_{MOX} = 1/R_{MOX}$), which is proportional to the gas concentration.

For estimating the values of τ_r and τ_d, a common practice in system modelling is that of identifying the parameters of the system upon its step response [115]. The problem with such procedure is that the time constants, in practice, depend not only on the type of volatile but also on its concentration [117]. In this chapter this dependency is made explicit by adjusting a polynomial regression model over a sequence of concentration pulses with different amplitudes for each target gas. In Section 3.4.3, an example of such relation is depicted for the target gas ethanol.

Transduction stage

The transduction stage is commonly defined by the sensitivity characteristics and the temperature and humidity dependencies of the transducer. For the case of Closed Sampling Systems, those characteristics are usually provided by the sensor manufacturer, relating the volatile concentration [ppm] to the sensor resistance ratio R_S/R_0 (sensor resistance in gas over sensor resistance in air) for different target gases and test conditions. Nevertheless, those sensitivity characteristics are obtained by measuring steady state values of the MOX sensor resistance after very long and constant exposure times, which are not applicable to Open Sampling Systems and thus not considered in this chapter.

3.4.2 Signal conditioning and preprocessing

The proposed model-based approach given by Eq. (3.3) and Eq. (3.5) relies on the sensor readings to obtain an estimation of the gas. As can be appreciated, a first order derivative needs to be computed to obtain such estimation, which notably degrades the signal-to-noise ratio and consequently the accuracy in the estimation (see Figure 3.20). Additionally, MOX sensors are susceptible to long and short term drift [119], gradually changing the sensor resistance even if exposed to exactly the same gas concentration under identical environmental conditions.

It then becomes necessary to carry out a signal conditioning to prepare the sensor readings to the posterior estimation process. Initially, for the purpose of drift compensation and dynamic range enhancement, the raw sensor readings $R_2(i)$ are divided by the sensor baseline resistance at $t = 0$, that is, $R_2(0)$. This transformation is known as relative baseline manipulation [119]. Later, in order to mitigate the noise effects on the model, a low pass filter followed by a sub-sampling process are applied to the signal, as depicted in Figure 3.20. The cutoff frequency of the filter and the down-sampling rate have been determined experimentally according to the sampling frequency.

3.4.3 Experimental results

This section presents three different experiments designed with increasing complexity to test how the proposed model can anticipate the steady state values of the sensor resistance from transient measurements in Open Sampling Systems. We start by

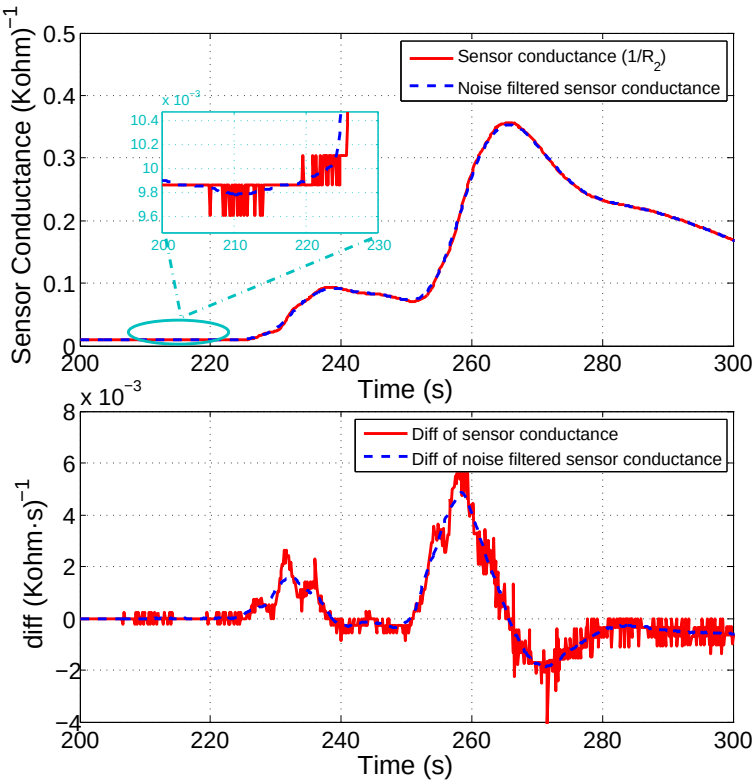

Figure 3.20: Noisy sensor conductance readings and its derivative (solid red line), and the corresponding filtered versions (dashed blue line).

testing our approach in a scenario where airflow and volatile distribution were well controlled. Then, two experiments of gas distribution mapping in unmodified environments are described.

Train of gas pulses in a controlled scenario

This experiment is designed to validate, in the simplest possible way, the ability of the proposed approach, given by Eq. (3.3) and Eq. (3.5), to estimate the volatile distribution that the MOX gas sensors are being exposed to.

Since knowing the ground truth distribution of a volatile in an OSS is a tough task (we can even say utopian), a specific setup was designed to keep its concentration as constant as possible and to confine it in a predefined region, avoiding the dispersion of gas particles to undesired locations. The experiment setup is depicted in Figure 3.21. The gas source, composed by a small cup filled with acetone, was placed inside a cardboard box with a small upper aperture. A fan located at the bottom of the box

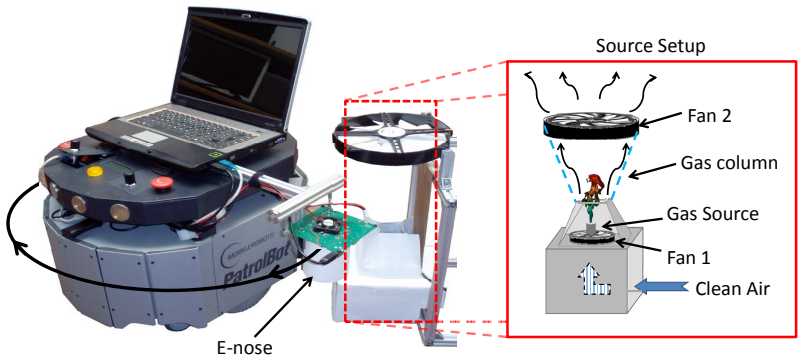

Figure 3.21: Picture of the Patrolbot mobile robot, the e-nose and the scheme of the gas source setup.

(beneath the gas source) was used to generate a constant upstream airflow pushing the volatile through the chamber aperture, while a second fan, placed about 20 cm over the box, sucked up the air from the box. This setup allows us to keep the volatile confined in an approximated gas column between both fans.

A Patrolbot robotic base (see Appendix A) was then placed in front of the gas source and commanded to rotate with constant angular speed ω. The e-nose, placed at the right side of the mobile base, then described a circular trajectory being exposed to the volatile only when passing through the gas column. The resulting excitation signal was then a train of identical duration pulses. The position of the source at each lap (and consequently the duration of each excitation pulse) was precisely obtained from the radial scan provided by a laser range scanner (SICK LMS200) carried by the Patrolbot.

Strictly speaking, because the sensor needs some time to entirely enter into the odor column, the excitation signal will not be a perfect pulse, but the deviation is small and can be neglected.

The pre-calibration of the time constants with the gas concentration, as depicted in Section 3.4.1, has been achieved by a polynomial regression over more than 50 short pulses with different amplitudes of the same volatile, in this case ethanol. The resulting dependency is depicted in Figure 3.22. Please, note that since the gas concentration is not available, we are considering the sensor conductance ($1/R_2$) instead.

Figure 3.23 presents the results of this experiment for a robot angular speed of $30°/s$, that is, the e-nose is exposed to the same pulse every 12 s. The width of the pulse (in seconds) is computed from the robot rotational speed and the angular references detected from the laser scan. The sequence of rise and recovery phases in the sensor output (solid red line) can be observed. It is also clear how the long tails of the recovery phases are interrupted by the next rising phase. (This phase overlap produces an interesting accumulation effect in the sensor output, increasing the sensor reading value on each exposition to the gas source even though the volatile concentration was

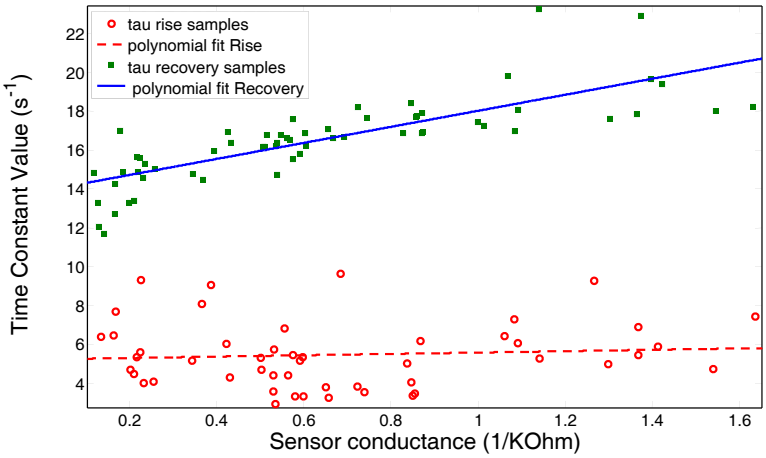

Figure 3.22: Pre-calibration of the time constants τ_r and τ_d with the gas concentration (*i.e.*, sensor conductance). Red circles and green squares represent the values of τ_r and τ_d respectively, while the lines represent their linear regression.

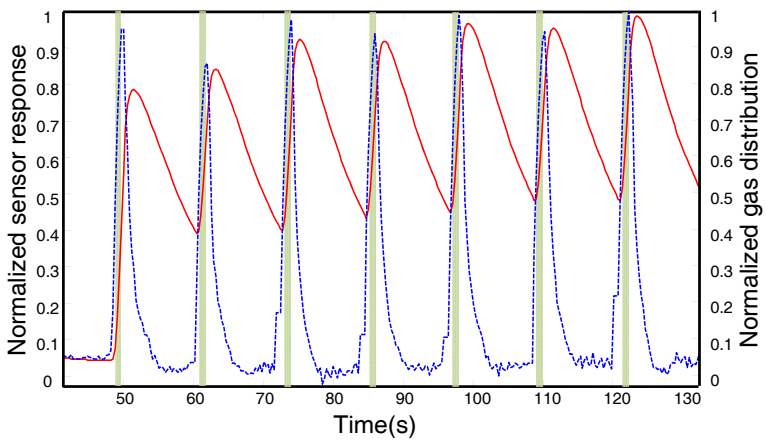

Figure 3.23: Results of the experiment in the controlled scenario for a robot angular speed $\omega = 30^{\circ}/s$. The train of acetone pulses (ground truth) is plotted as green shaded bars, the normalized raw sensor readings (Y) (See Figure 3.19) are plotted as a solid red line and the gas distribution estimated by the proposed approach is plotted as a dashed blue line.

constant along the experiment duration. The study of this effect is out of the scope of this thesis.)

As can be noticed, the estimation of the gas distribution (dashed blue line) as provided by applying the proposed model is a more accurate estimation of the train of pulses than the raw sensor readings. The improvement is significant for the recovering phases, not only because the overlap between phases is avoided (the recovery phase is interrupted by the next rise phase before reaching the baseline), but because our estimation provides values of the gas distribution more consistent with the reality.

1D gas distribution mapping

The motivation of this experiment is to demonstrate the utility of the proposed approach in the generation of gas distribution maps, a challenging problem in robotic olfaction [96, 90].

As described in the introduction, the slow recovery of MOX sensors becomes a serious drawback when the sensors are carried on a vehicle (typically a mobile robot). In such cases, the adopted solution is to reduce the vehicle velocity to a few cm/s, increasing considerably the execution time of the task. We demonstrate here that our estimation of the gas distribution to a certain extent removes the need to reduce the robot velocity for maintaining the map accuracy.

For such a goal, two different gas distribution maps of an unmodified long indoor corridor are generated by driving the robot at two different velocities. Figure 3.24 shows a comparison of the gas maps generated using the Kernel-based method proposed in [85] for a reduced robot speed of 0.1 m/s, while Figure 3.25 represents the results of a similar experiment using a robot speed of 0.4 m/s.

In both cases, the robot travels the corridor twice (round trip), passing over a cap filled with acetone placed in the middle of it. The map evolves as the robot moves and collects new sensory data, so we decided to analyze the map as built, at two different points of the robot path: at the end of the corridor (top row sub-figures), and at the initial position after the task has been completed (bottom row sub-figures).

As can be seen from the results, the *"tail"* effect produced by the slow recovery of MOX sensors leads to a gas distribution map with high concentrations along the robot path once the source has been hit. As expected, this effect is more harmful when increasing the robot speed, one of the reasons why the speed of olfactory robots is usually kept small. Please, note how this effect is substantially palliated when applying the proposed gas estimation (see Figure 3.25(e) and Figure 3.25(f)), even when the robot speed is high.

As a conclusion, our approach not only presents a noticeable improvement in the correct localization of the gas source but also provides confident values of the gas distribution along the robot path. This is important since in most cases it is desirable to not only know the location of the gas source but also know how the volatile emanating from it has spread in the surroundings.

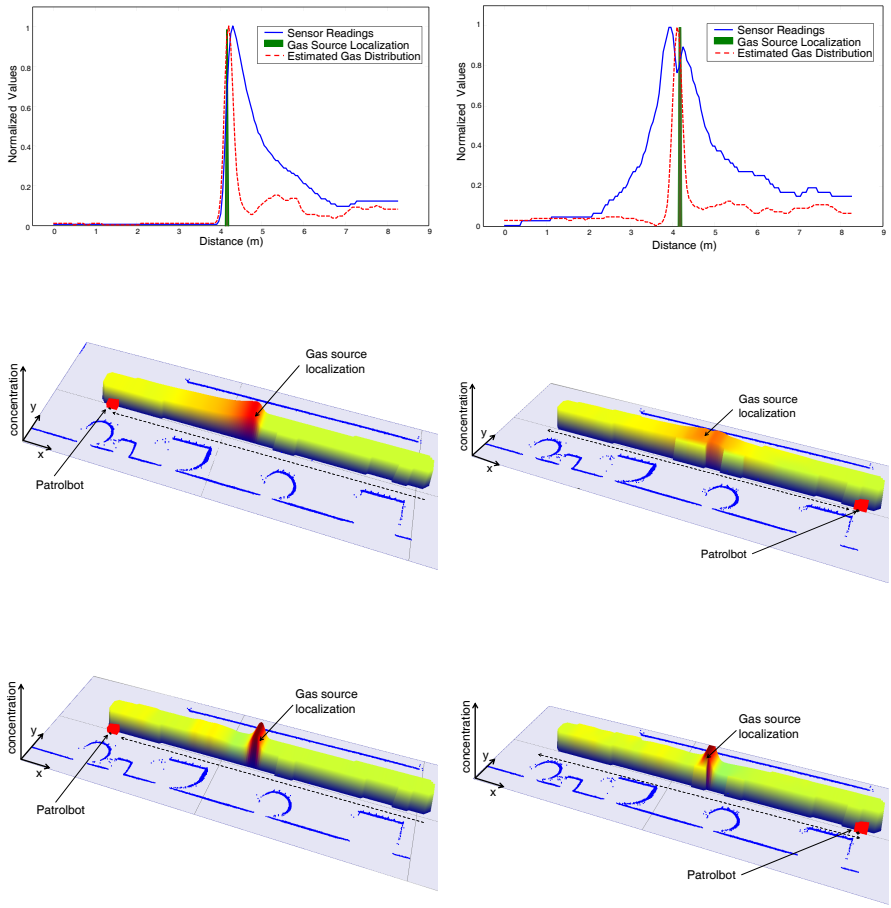

Figure 3.24: Snapshots of the maps generated in the 1D gas distribution mapping experiment at the middle and end of the robot path (left/right column, respectively), for a robot speed of $0.1m/s$. Top row sub-figures illustrate a comparison between the normalized sensor readings and the gas distribution estimated here. Middle and bottom row sub-figures show a 3D reconstruction of the maps generated with the raw MOX readings (middle row), or the estimation provided by the proposed model (bottom row).

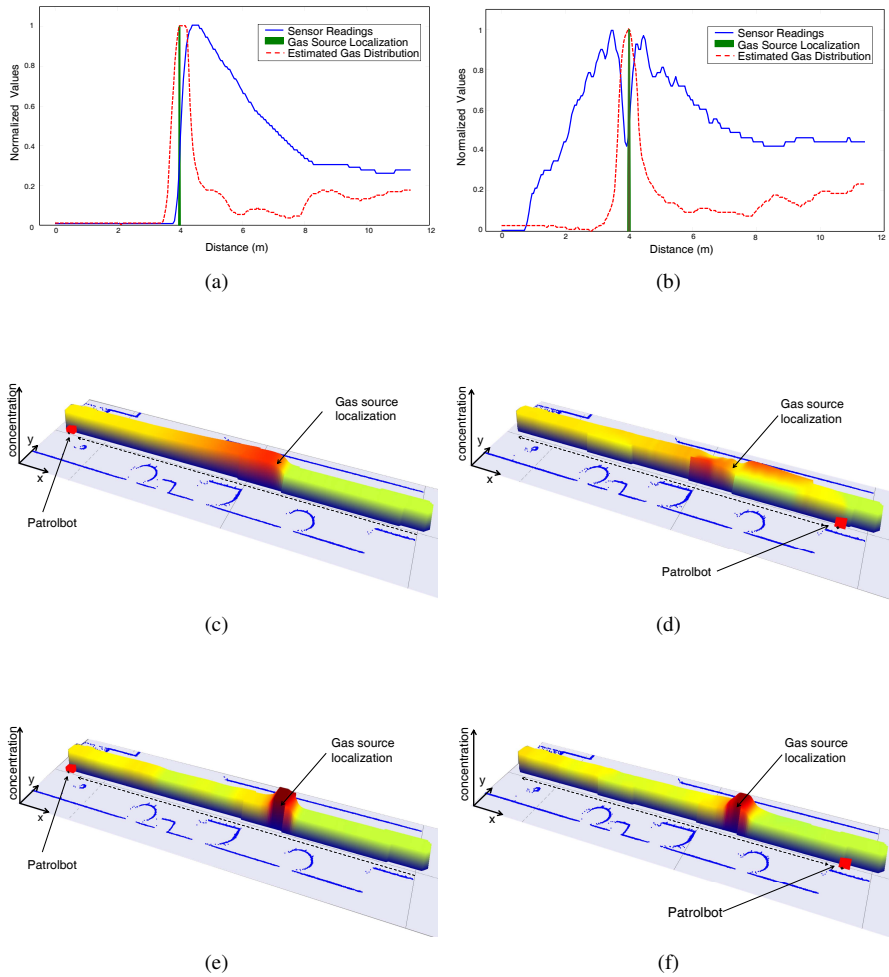

Figure 3.25: Snapshots of the maps generated in the 1D gas distribution mapping experiment at the middle and end of the robot path (left/right column, respectively), for a robot speed of $0.4m/s$. Top row sub-figures illustrate a comparison between the normalized sensor readings and the gas distribution estimated here. The middle and bottom row sub-figures show a 3D reconstruction of the maps generated with the raw MOX readings (middle row), or the estimation provided by the proposed model (bottom row).

2D gas distribution mapping

In this last experiment, we pursue to test our model-based estimation of the gas concentration in a more complex environment. The idea is to accomplish a complete olfaction mission in a realistic scenario where no alterations of the environment are done. Typical examples are the localization of leaks or the declaration of areas of high concentration levels of harmful gases.

The testing scenario consisted of two adjacent rooms communicated through a small corridor (see Figure 3.26). The robot inspected both rooms following the predefined path marked as solid red line, at a speed of $0.3m/s$. The gas source, composed by a cardboard plate impregnated in ethanol, was placed on the floor along the robot path. To be able to compare the maps generated from the raw sensor readings with those provided after applying the proposed MOX sensor model, doors and windows were kept closed to avoid uncontrolled airflows.

Figures 3.27 and 3.28 depict the mean and predictive variance gas distribution maps generated by the Kernel DM+V algorithm [89], when fed with the MOX sensor readings and the model-based gas estimation, respectively.

Focusing on the mean maps, it can be seen that in both cases the maximum concentration falls near the real source location (marked as a white circle). However, when the map is built from the raw sensor readings, a high concentration area appears along the corridor connecting both rooms. The shape reassembles that of a gas plume generated by a dominant airflow, which we know is not possible since doors and windows were kept closed during the experiment. This "fake" plume is attributed

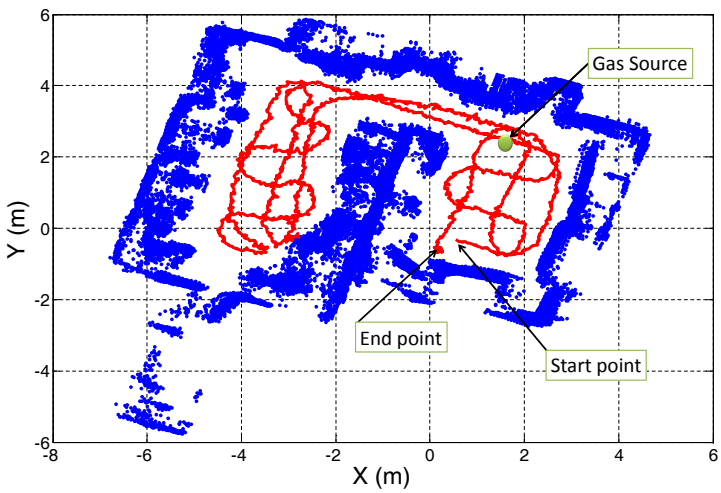

Figure 3.26: Map of the experimental area used in the 2D mapping experiment. Blue points represent obstacles detected by the onboard SICK laser scan, the robot path is marked as red solid line, and the gas source position is pointed with a green circle.

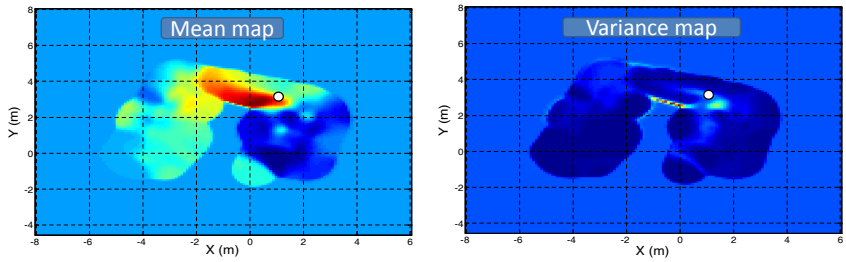

Figure 3.27: Mean and predictive variance gas distribution maps of the inspected area generated by the Kernel DM+V algorithm when fed with the sensor readings (after baseline manipulation and delay correction). The gas source location has been marked as a white circle.

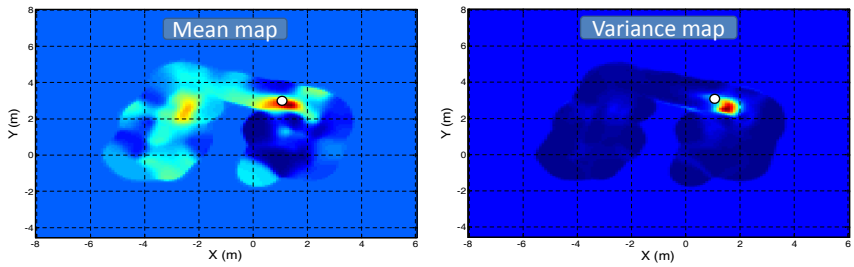

Figure 3.28: Mean and predictive variance gas distribution maps of the inspected area generated by the Kernel DM+V algorithm after applying the proposed model-base estimation of the gas concentration. The gas source location has been marked as a white circle.

to the actual robot path and the slow recovery of MOX sensors. When building the map from the gas concentration estimated by the proposed model (Figure 3.28), this "fake" plume does not appear, correctly representing the gas distribution in the inspected area.

Focusing now on the predictive variance maps, it has been previously reported that the variance of a set of gas concentration measurements has been suggested as a feature that can identify the location of a source of gas [85, 57]. Following this principle, it can be seen that the gas source cannot be precisely located in the first case (Figure 3.27) since there are high variance values not only nearby the real location of the source but also along the corridor. Again, the application of the proposed MOX model leads to a predictive variance map where all high variance values fall near the real source location, allowing a correct estimation of it.

3.5 Conclusions and outlook

In this chapter, we have presented two approaches to deal with the long recovery period of MOX gas sensors. This is a serious drawback for mobile robot olfaction since a rapid cycle of measurement is required in many olfaction-related tasks: source finding, gas concentration mapping, etc.

First, we have introduced a new electronic nose, the MCE-nose, which partially overcomes this problem by accommodating a set of redundant sensors in different chambers. These redundant sensors are alternatively activated, ignoring the sensor output when a decay phase is detected and delegating the sensing task to another clean, almost identical sensor. The output signal of the whole setup results then from the concatenation of the rise phases of a sequence of MOX sensors. A prototype of the MCE-nose has been built and integrated in a mobile robotic platform under the OpenMORA robotic architecture. Then, a modelling approach which compensates the slow dynamic behavior of the MOX sensor, by forecasting the steady state values of the sensor resistance from a sequence of transient measurements has been presented. The exploited model is based on two first order systems (rise and recovery) with time constants that depend on the sensor reading amplitude.

Both approaches have been validated in different scenarios, demonstrating their utility in applications like gas source localization or gas distribution mapping. Additionally, we have proved that these approaches enable a considerable increase in the speed at which a mobile base carrying the MOX-based e-nose can inspect the environment, which directly implies an important reduction in the execution times of the olfaction task.

Future work includes some improvements in the MCE-nose, such as the incorporation of another electro-valve to purge the pneumatic circuit or the enlargement of the tubes section to increase the airflow through the sensor's surface. Related to the proposed model, modeling additional parameters as temperature and humidity to improve the gas distribution estimation in real complex scenarios will be considered. Besides, the study of the transduction stage (quantification) for open sampling systems will be considered to complete the MOX sensor model.

Chapter 4
Gas quantification with MOX sensors in open sampling systems

Gas quantification based on the response of an array of metal oxide (MOX) gas sensors in an open sampling system is a complex problem due to the highly dynamic characteristic of turbulent airflow and the slow dynamics of the MOX sensors. However, many gas related applications require to determine the gas concentration the sensors are being exposed to. Due to the chaotic nature that dominates gas dispersal, in most cases it is desirable to provide, together with an estimate of the mean concentration, an estimate of the uncertainty of the prediction. This chapter presents a probabilistic approach for gas quantification with an array of MOX gas sensors based on Gaussian Processes, estimating for every measurement of the sensors a posterior distribution of the concentration, from which confidence intervals can be obtained. The proposed approach has been tested with an experimental setup where an array of MOX sensors and a photo ionization detector (PID), used to obtain ground truth concentration, are placed downwind with respect to the gas source. Our approach has been implemented and compared with standard gas quantification methods, demonstrating the advantages when estimating gas concentrations.

4.1 Introduction

The problem addressed in this chapter is the estimation of gas concentration from the readings of an array of metal oxide semiconductor (MOX) gas sensors (see Section 2.2) deployed in an open sampling system (OSS). In other words, the problem can be formalized as finding a function that maps the readings of N gas sensors together with any other extra parameters (like temperature and humidity) to a posterior distribution over concentrations.

Concentration estimation is a crucial step for realistic gas sensing applications since legal requirements and regulations are expressed in terms of absolute gas concentration, toxicity levels, etc. For example, it would be of little utility for many applications if we could detect a gas leak in an industrial scenario but we were unable to *quantify* the amount of leaked gas: should an alarm be issued for workers to abandon the area, or is localization of the source and subsequent notification to the maintenance unit enough to handle the problem?

Gas quantification using an array of MOX sensors in an OSS is indeed an important problem. Most of the previous works dealing with OSS do not estimate the gas concentration but work directly with the sensor signal (conductance readings in case of MOX gas sensors). In contrast to gas quantification with sensors within a sensing chamber (where controlled conditions can be imposed), gas quantification in OSS implies additional complications due to the many sources of uncertainty. The most relevant source of uncertainty is the exposition of the sensors to the turbulent airflow that brings the chemical compound in contact with the sensors. As a consequence, given the slow dynamics of MOX gas sensors and the rapid fluctuations in concentration due to turbulent airflow, the sensors never reach a steady state but continuously fluctuate [146].

Instead of trying to find a deterministic function for mapping directly the sensors readings to concentration values, in this chapter we calculate instead a posterior distribution over concentrations c at time t given the measurements of N sensors $\mathbf{r_t} = (r_1(t), r_2(t) \dots r_N(t))$ at times $\mathbf{t} = (t_1, \dots, t_k)$. That is, we search a probability density function:

$$p\left(c(t)|\mathbf{r_{t_1}}, \dots, \mathbf{r_{t_k}}\right) \tag{4.1}$$

where k is the number of past sensor readings considered for the prediction.

Tasks like chemical detection, gas source localization, gas distribution mapping, and odor trail tracking are common tasks for working in OSS. For this kind of systems it is acceptable to have only an approximate estimate of the concentration if there is also an indication of the corresponding confidence. This is exactly what the method proposed in this chapter provides.

We propose to estimate the posterior in Eq. (4.1) using a Gaussian Process (GP) model [123]. Gaussian Processes provide principled supervised machine learning methods which, given a set of samples (i.e. pairs of sensor readings and their corresponding concentration), can estimate the posterior joint probability of the process.

We analyze gas quantification using either a single sensor or the whole sensor array. In the latter context we additionally investigate how Automatic Relevance Determination (ARD) [112, 98] can be used to identify which sensors in the array contribute most to the estimate of the concentration posterior distribution. Finally, we analyze to which extent considering past sensor readings can improve the accuracy of probabilistic quantification.

The rest of this chapter is outlined as follow: After a discussion of related research in Section 4.2, we introduce in Section 4.3 the basics of Gaussian Process regression for gas quantification, giving especial attention to learning a GP from the data. Finally, Section 4.4 presents some experimental results under different configurations of the GP.

4.2 Related research

Estimation of gas concentration from raw readings of MOX-based gas sensors has been traditionally focused on setups where such sensors are enclosed inside a chamber, where environmental conditions, gas exposure times and concentrations are known and controlled. This setup allows the measurement of steady state values, which are used as input to a regressor. Under such controlled conditions, the relation between sensor conductance (Ω^{-1}) and gas concentration (ppm) is usually modelled as an exponential [61]:

$$g_i \equiv \frac{1}{r_i} = A_i \cdot c^{\alpha_i} \tag{4.2}$$

where g_i is the conductance of sensor i within the array (inverse of sensor resistance r_i), c is the gas concentration, and A_i and α_i are the parameters of the exponential model to be estimated during an initial training process.

For concentration estimation with an array of MOX sensors, multivariate linear regression methods like Principal Component Regression (PCR) and Partial Least Squares Regression (PLSR) have been proposed [45, 143, 23]. The main motivation behind the use of these two methods is the strong correlation of the response of different MOX sensors. This allows both PCR and PLSR to reduce the dimensionality of the input space before fitting a regression function, thus reducing the possibility of *curse of dimensionality* related issues. Alternatively, non-linear estimation methods like Artificial Neural Networks (ANN) [55, 12] or kernel algorithms like Support Vector Regression (SVR) [138, 19] have also been proposed.

It is worth noting that some authors consider the transient information for gas quantification with MOX sensors. In [111], a multi-exponential model is used to describe the sensor dynamics and to predict the steady state value of the sensors which is then mapped to a concentration based on the initial values of the transient state. These existing approaches rely on steady state measurements of the sensor. Thus, they are not immediately applicable to OSS because steady state values are almost never reached [146].

The work in [151] addressed gas quantification for urban pollution monitoring. Measurements collected over a long period of time are averaged out and therefore the dynamic information in the sensor response is discarded. The focus of this chapter is instead on applications where sensors are deployed in a highly dynamic environment, where they are exposed to intermittent patches of gas.

Most of the works in mobile robot olfaction avoids quantification issues and use instead the conductance readings of the sensors as an approximate measure of the gas concentration. An exception is [66], where Ishida et al. propose to use steady state calibration to obtain a rough approximation of gas concentration with an OSS.

All the methods mentioned ignore the uncertainty in the quantification, while for OSS it is desirable to provide the uncertainty together with the concentration estimate. The GP-based method detailed in the following section generates an estimate of the uncertainty (as a variance), which can be used, for example, to calculate confidence intervals for the predictions.

4.3 An algorithm for probabilistic gas concentration estimation

This section details our proposal for the concentration estimation in MOX-OSS. Initially, two signal preprocessing methods and their influence on the posterior distribution estimation are described. Next, Gaussian Process regression for the particular case of gas quantification with an array of sensors is summarized, and how Automatic Relevance Determination can be used to select the model parameters. Finally, two different loss functions are proposed for evaluating the results, allowing a comparison between the various proposed configurations.

4.3.1 Signal preprocessing

In the first step, the raw sensor resistance readings $r_i(t)$ are divided by the baseline of the sensor at $t = 0$, that is, $r_i(0)$. This transformation in Eq. (4.3), known as relative baseline manipulation [119],is applied for drift compensation and dynamic range enhancement:

$$\tilde{\mathbf{r}}_\mathbf{t} = [\tilde{r}_i(t)]_{i=1}^{N}, \tilde{r}_i(t) = \frac{r_i(t)}{r_i(0)} \tag{4.3}$$

Next, GP regression is used to predict the gas concentration $c(t)$ from the values $\tilde{\mathbf{r}}_\mathbf{t}$. However, considering the commonly assumed exponential relation between sensor resistance and the concentration, see Eq. (4.2), we will also investigate applying a logarithmic transformation and perform regression between $\log(\tilde{\mathbf{r}}_\mathbf{t})$ and $\log(c(t))$.

In summary, we will compare two regression problems:

$$c(t) = f_1(\tilde{\mathbf{r}}_\mathbf{t}) \tag{4.4}$$

$$\log c(t) = f_2(\log \tilde{\mathbf{r}}_\mathbf{t}) \tag{4.5}$$

4.3.2 Gaussian Process regression for gas concentration estimation

In the general case, the process of inferring the relationship $f : \mathbf{r_t} \longmapsto c(t)$ between the response of an array of sensors and the gas concentration using a training dataset $\mathbf{D} = \{(\mathbf{r}_{t_j}, c(t_j)|j = 1, ..., n\}$ is a supervised machine learning problem. Among the numerous available options, GPs provide a powerful non-parametric tool for Bayesian inference and learning [123]. GPs can be seen as a generalization of the Gaussian probability distribution to distributions over functions. That is, they perform inference directly in the space of functions, starting with a prior distribution over all possible functions and subsequently learning the target function from data samples. Defining a prior over functions corresponds to making assumptions about the characteristics of the function f, as otherwise any function which is consistent with the training data will be equally valid and therefore the learning problem would be ill-defined.

A GP is completely specified by its mean and covariance functions, $m(\mathbf{r_t})$ and $k(\mathbf{r_t}, \mathbf{r_{t'}})$ respectively:

$$\begin{aligned} m(\mathbf{r_t}) &= \mathbb{E}[f(\mathbf{r_t})], && (4.6)\\ k(\mathbf{r_t}, \mathbf{r_{t'}}) &= cov(f(\mathbf{r_t}), f(\mathbf{r_{t'}})) && (4.7)\\ &= \mathbb{E}[(f(\mathbf{r_t}) - m(\mathbf{r_t}))(f(\mathbf{r_{t'}}) - m(\mathbf{r_{t'}}))]. && (4.8) \end{aligned}$$

we denote the GP as:

$$f(\mathbf{r_t}) \sim \mathscr{GP}(m(\mathbf{r_t}), k(\mathbf{r_t}, \mathbf{r_{t'}})). \tag{4.9}$$

To account for noise in the sensor it is assumed that the observed concentration values $c(t)$ are corrupted with an additive i.i.d. Gaussian noise with zero mean and variance σ_n^2, that is:

$$\begin{aligned} c(t) &= f(\mathbf{r_t}) + \varepsilon, \\ \varepsilon &\sim \mathcal{N}(0, \sigma_n^2). \end{aligned} \tag{4.10}$$

It is important to notice that we did not assess experimentally whether the noise is i.i.d. with Gaussian distribution. We rather make this assumption to obtain a closed form solution, and validate the resulting predictions with real sensor data.

In our case of study we consider GPs with zero mean and the commonly used squared exponential (SE) covariance function, that is:

$$\begin{aligned} m(\mathbf{r_t}) &= 0, && (4.11)\\ k(\mathbf{r_t}, \mathbf{r_{t'}}) &= \sigma_f^2 \exp\left(-\frac{1}{2}\frac{\| \mathbf{r_t} - \mathbf{r_{t'}} \|^2}{\ell^2}\right), && (4.12) \end{aligned}$$

where σ_f^2 is the overall variance hyper-parameter and ℓ is the characteristic length scale. It is common but not necessary to consider GPs with zero mean. Note that the mean of the posterior process is not restricted to be zero. The choice of the SE covariance function leads to GP predictions that are *smooth* over the characteristic length scale. That is, if $\mathbf{r_t} \approx \mathbf{r_{t'}}$, then $k(\mathbf{r_t}, \mathbf{r_{t'}})$ approaches its maximum and $f(\mathbf{r_t})$ is strongly correlated with $f(\mathbf{r_{t'}})$. For large distances between $\mathbf{r_t}$ and $\mathbf{r_{t'}}$, $k(\mathbf{r_t}, \mathbf{r_{t'}})$ approaches 0. So, when predicting the concentration value for new data points, distant observations will have a negligible effect. The region of influence, depends on the scale parameter ℓ.

The regression model depends on the selection of the hyper-parameters, which are summarized in a vector $\theta = (\ell^2, \sigma_f^2, \sigma_n^2)$. The optimal hyper-parameters are found by maximizing the marginal likelihood function $p(\mathbf{c}|\mathbf{R}, \theta)$, where $\mathbf{c}$ is a vector of training concentration values, $\mathbf{R}$ is the matrix containing the measurements of the sensor array, and θ are the hyper-parameters. As it is common practice, we minimize the corresponding negative log-likelihood to avoid numerical issues:

$$\begin{aligned} -\log\left(p(\mathbf{c}|\mathbf{R},\theta)\right) &= \frac{1}{2}\mathbf{c}^\top \mathbf{K}^{-1}\mathbf{c} + \frac{1}{2}\log|\mathbf{K}| + \frac{n}{2}\log 2\pi \\ \text{where } \mathbf{K} &= k(\mathbf{r_t}, \mathbf{r_{t'}}) + \sigma_n^2 \mathbf{I} \end{aligned} \tag{4.13}$$

To find the minimum of Eq. (4.13) we use the scaled conjugate gradient method, which requires the calculation of the partial derivatives of the log marginal likelihood w.r.t. the hyper-parameters:

$$\begin{aligned} \frac{\partial}{\partial \theta_j}\log\left(p(\mathbf{c}|\mathbf{R},\theta)\right) &= \frac{1}{2} tr\left[\left(\alpha\alpha^\top - \mathbf{K}^{-1}\right)\frac{\partial \mathbf{K}}{\partial \theta_j}\right] \\ \text{where } \alpha &= \mathbf{K}^{-1}\mathbf{c} \end{aligned} \tag{4.14}$$

The complexity of this step is dominated by the matrix inversion $\mathbf{K}^{-1}$ in Eq. (4.14), which has a complexity $\mathcal{O}(n^3)$ with n being the number of training points. This represents one of the principal inconveniences of GP.

Learning the calibration GP corresponds to the selection of the hyper-parameters. The GP then allows to predict gas concentration values $\mathbf{c}_*$ and a corresponding variance for arbitrary sensor resistances $\mathbf{r}_*$.

The posterior distribution over functions (our prediction) is also a Gaussian, and it is given by:

$$\mathbf{c}_*|\mathbf{R}, \mathbf{c}, \mathbf{R}_* \sim \mathcal{N}\left(\bar{\mathbf{c}}_*, cov(\mathbf{c}_*)\right), \text{where} \tag{4.15}$$

$$\bar{\mathbf{c}}_* \triangleq \mathbb{E}[\mathbf{c}_*|\mathbf{R}, \mathbf{c}, \mathbf{R}_*] = \mathbf{K}(\mathbf{R}_*, \mathbf{R})[\mathbf{K}(\mathbf{R}, \mathbf{R}) + \sigma_n^2\mathbf{I}]^{-1}\mathbf{c}$$

$$cov(\mathbf{c}_*) = \mathbf{K}(\mathbf{R}_*, \mathbf{R}_*) - \mathbf{K}(\mathbf{R}_*, \mathbf{R})[\mathbf{K}(\mathbf{R}, \mathbf{R}) + \sigma_n^2\mathbf{I}]^{-1}\mathbf{K}(\mathbf{R}, \mathbf{R}_*)$$

and where $\mathbf{R} = \{\mathbf{r}_1 ... \mathbf{r}_n\}$ is the $n \times N$ matrix of the n training samples of dimensionality N, $\mathbf{R}_*$ the matrix for the testing inputs, $\mathbf{K}(\cdot,\cdot)$ refers to the matrix with the entries given by the covariance function $k(\cdot,\cdot)$ and $\mathbf{c}$ the vector of the observed concentrations c_i.

Note that the predictive distribution is based on a mean value $\bar{\mathbf{c}_*}$ (our best estimate for $\mathbf{c}_*$), which is a linear combination of the observed values $\mathbf{c}$, and a variance value $cov(\mathbf{c}_*)$ which denotes the uncertainty in our estimation, and does not depend on the observed targets but only on the inputs.

4.3.3 Automatic relevance determination

Automatic Relevance Determination (ARD) is a method based on Bayesian interference for pruning large feature sets with the aim to obtain a sparse explanatory subset. Making use of this powerful tool we can consider different configurations of the input space of higher dimensionality, and then allow ARD to select the most relevant features, which avoids overfitting due to high input dimensionality. In order to introduce ARD Eq. (4.12) can be rewritten as:

$$k(\mathbf{r_t}, \mathbf{r_{t'}}) = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{r_t} - \mathbf{r_{t'}})^\top \mathbf{M}(\mathbf{r_t} - \mathbf{r_{t'}})\right), \tag{4.16}$$

where $\mathbf{M}$ denotes the diagonal weight matrix $\mathbf{M} = \ell^{-2}\mathbf{I}$.

We have seen how maximizing the log marginal likelihood can be used to determine the value of the hyper-parameters. By incorporating a separate hyper-parameter ℓ_i for each input variable [123], i.e. for each sensor in the array, we modify $\mathbf{M}$ to be:

$$\mathbf{M} = diag\left([\ell_1^{-2}, \ell_2^{-2}, \ldots, \ell_n^{-2}]\right) \tag{4.17}$$

where $\ell = [\ell_1, \ell_2, \ldots, \ell_n]$ is a vector of positive values, corresponding to the length-scale of each input variable. This is in contrast to a *global* hyper-parameter ℓ for all input variables.

Since the inverse of the length-scale determines how relevant an input is, the extension (4.17) enables to identify the importance of each different input– if the length-scale has a very large value, the covariance will become almost independent of that input, effectively ignoring its values during the inference.

ARD is performed during the training phase of the GP, specifically during the selection of the covariance function hyper-parameters. When the inputs related to a sensor are discarded by ARD, the corresponding sensors can be excluded from the array. ARD has been applied for two different configurations of the input space: (i) when the whole sensor array is considered at one instant of time in the inference process and (ii) when additional features from previous time steps of the sensor response are considered to account for the dynamics of the signal.

4.3.4 Evaluation of the predictions

To compare different configurations of the GP based quantification proposed in this chapter, among themselves, and with other methods previously proposed in the literature, two performance measures are proposed:

Root Mean Squared Error (RMSE): The RMSE is calculated as the difference between the ground-truth concentration, obtained with the readings from a photo ionization detector - PID (refer to Section 2.2), and the expected value of the predictive distribution obtained from Eq. (4.15).

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(c_i - \bar{c}_{*i})^2} \tag{4.18}$$

Notice that the RMSE takes only into account the predictive mean, while it ignores its uncertainty. However, this indicator allows to compare the predictions of the proposed GP quantification approach with other regression methods, like Partial Least Squares Regression (PLSR) or Support Vector Regression (SVR), which do not provide any estimation of the prediction uncertainty.

Negative Log Predictive Density (NLPD): The NLPD is a standard criterion to evaluate probabilistic models (see Eq. (4.19)).

$$NLPD = -\frac{1}{n}\sum_{i=1}^{n}\log(p(c_i|\mathbf{r_i})) \tag{4.19}$$

It is worth noting that the NLPD considers the whole posterior distribution and not only its expected value. In general, more negative NLPD values indicate better predictions with a small uncertainty.

Considering the two preprocessing methods proposed in Section 4.3.1, two different NLPD formulas arise. In the first case (linear preprocessing - Eq. (4.4)), the posterior distribution of the concentration is a normal distribution, while in the second case, (logarithmic preprocessing - Eq. (4.5)), the posterior distribution of the concentration is a log-normal distribution. Therefore the NLPD is calculated for the two cases respectively:

$$NLPD_{normal} = \frac{\log(2\pi)}{2} + \frac{1}{2N}\sum_{i=1}^{n}\left[\log(\sigma^2(c_i)) + \frac{(c_i - \mu(c_i))^2}{\sigma^2(c_i)}\right] \tag{4.20}$$

$$NLPD_{log-normal} = \frac{\log(2\pi)}{2} + \frac{1}{2N}\sum_{i=1}^{n}\left[\log(c_i^2\sigma^2(c_i)) + \frac{(\log(c_i) - \mu(c_i))^2}{\sigma^2(c_i)}\right]$$

where c_i is the ground truth gas concentration, $\sigma(c_i)$ is the predictive standard deviation, and $\mu(c_i)$ the predictive mean.

4.4 Experimental evaluation and discussion

4.4.1 Experimental setup

The experiments are carried out in a 4.85 m × 3.42 m × 2 m room with an induced artificial airflow of approximately 0.1 m/s. The airflow is created using two arrays of six standard microprocessor cooling fans. The gas source is an odor blender (olfactory display), a device described in [110] that can mix up to 13 gas components from arbitrary recipes using rapidly switching solenoid valves. The odor blender samples from the headspace of the compounds, which are kept in liquid phase. This odor blender enables rapid switching of compound and concentration. The odor blender uses headspace sampling and therefore does not intensify evaporation, contrary to an odor bubbler [128]. The outlet of the olfactory blender is placed on the floor, 0.5 m upwind with respect to an array of 11 MOX gas sensors and a PID[1]. The airflow at the outlet of the odor blender is set to 1 l/min. Figure 4.1 displays the configuration of the experiments. The sensors included in the array are listed in Table 4.1. The sensors are sampled at 4 Hz. The PID is placed next to the array of MOX sensors in order to obtain calibrated measurements in the proximity of the active area of the MOX sensors. This is important since, due to diffusion and advection, the estimation of the gas concentration at the sensors would be very complicated if only the intensity of the gas source would be available. The position of the MOX sensors and the PID has been carefully chosen in order to ensure that the sensors are exposed to a very similar gas concentration. This can be verified calculating the Pearson's coefficient to estimate the linear correlation among the response of the sensors. From the results reported in Table 4.2 it is clear that the response of the MOX sensors and the PID

[1]PID model ppbRAE2000 from RAESystem with a 10.6 eV UV lamp.

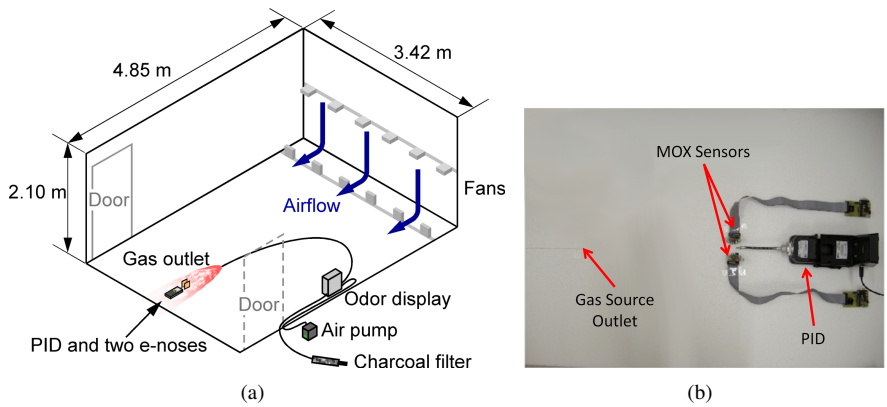

(a) (b)

Figure 4.1: Sketch of the experimental setup (a), and detailed picture of the MOX sensors and the PID used to sample the volatiles (b).

Model	Gases Detected	Quantity
Figaro TGS 2600	Hydrogen, Carbon Monoxide	2
Figaro TGS 2602	Ammonia, Hydrogen Sulfide, VOC (volatile organic compound)	1
Figaro TGS 2611	Methane	1
Figaro TGS 2620	Organic Solvents	1
e2V MiCS 2610	Ozone	1
e2V MiCS 2710	Nitrogen Dioxide	1
e2V MiCS 5521	Carbon Monoxide, Hydrocarbons, VOC	2
e2V MiCS 5121	Carbon Monoxide, Hydrocarbons, VOC	1
e2V MiCS 5135	Carbon Monoxide, Hydrocarbons, VOC	1

Table 4.1: Gas sensors used in the sensor array.

are highly correlated and therefore it can be inferred that the sensors are exposed to very similar concentration profiles. Due to the faster sensor dynamics, the correlation of the PID response with the response of the MOX sensors is in general slightly lower than the correlation between the response of two MOX sensors. The compound selected for these experiments is ethanol, which is heavier than air and, consequently, forms plumes at ground level.

PID	MiCS 2610	MiCS 2710	MiCS 5521-1	MiCS 5121	MiCS 3135	MiCS 5521-2	TGS 2600-1	TGS 2611	TGS 2620	TGS 2600-2	TGS 2602
1,00	-0,75	-0,88	-0,86	-0,89	-0,82	-0,89	-0,82	-0,92	-0,80	-0,81	-0,69
-0,75	**1,00**	0,93	0,93	0,95	0,98	0,91	0,98	0,90	0,98	0,98	0,98
-0,88	0,93	**1,00**	0,94	0,98	0,97	0,95	0,97	0,97	0,96	0,97	0,90
-0,86	0,93	0,94	**1,00**	0,98	0,97	0,99	0,96	0,97	0,96	0,96	0,90
-0,89	0,95	0,98	0,98	**1,00**	0,98	0,98	0,98	0,98	0,97	0,98	0,91
-0,82	0,98	0,97	0,97	0,98	**1,00**	0,96	1,00	0,96	1,00	1,00	0,96
-0,89	0,91	0,95	0,99	0,98	0,96	**1,00**	0,95	0,97	0,94	0,95	0,87
-0,82	0,98	0,97	0,96	0,98	1,00	0,95	**1,00**	0,95	1,00	1,00	0,97
-0,92	0,90	0,97	0,97	0,98	0,96	0,97	0,95	**1,00**	0,95	0,95	0,86
-0,80	0,98	0,96	0,96	0,97	1,00	0,94	1,00	0,95	**1,00**	1,00	0,97
-0,81	0,98	0,97	0,96	0,98	1,00	0,95	1,00	0,95	1,00	**1,00**	0,97
-0,69	0,98	0,90	0,90	0,91	0,96	0,87	0,97	0,86	0,97	0,97	**1,00**

Table 4.2: Cross-correlation coefficients for the MOX sensors in the array and the PID. High values suggest that the sensors are exposed to the same concentrations.

In order to create a dataset that represents a variety of scenarios, four different odor emitting profiles have been used (see Figure 4.2). For all the profiles the gas source does not emit gas for two minutes (this is equivalent to releasing no gas) and the signal of the sensors during this period is assumed as the baseline. Also, at the end of all the experiments the source emits clean air for 2 minutes. Overall, the dataset includes a total of 18 experiments, 3 for the deterministic emission strategies and 9 for the randomized emission strategy.

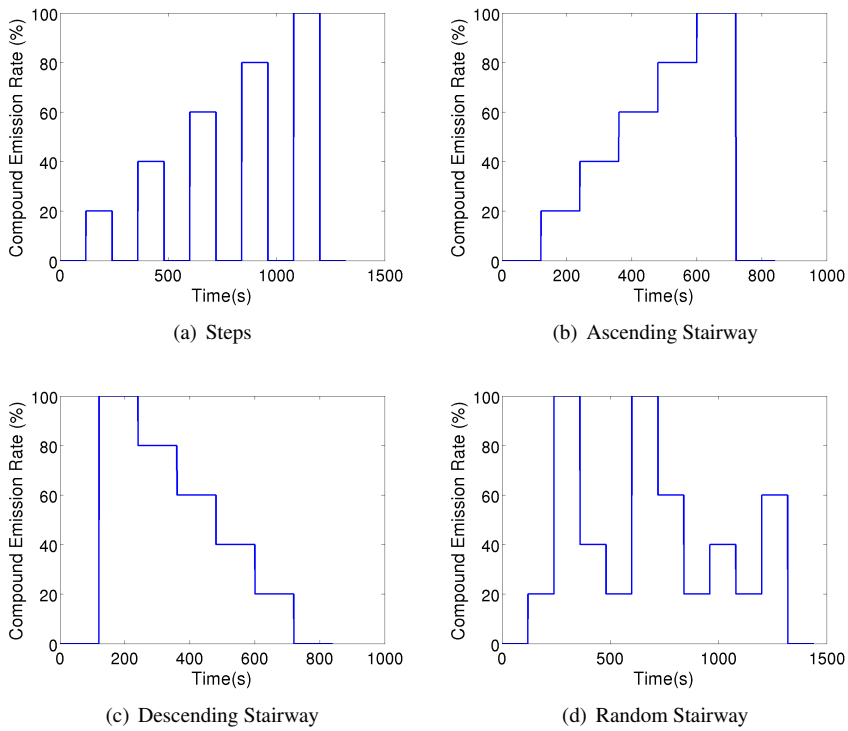

(a) Steps (b) Ascending Stairway

(c) Descending Stairway (d) Random Stairway

Figure 4.2: Gas source emission strategies. For the randomized strategy (d), one exemplary instance is displayed.

4.4.2 Results

In this section we present and compare the gas sensor quantification results obtained with the two preprocessing methods (linear & logarithmic).

We further compare three different sets of input variables. First, we consider the sensors independently. Second, we compare to the case where the whole array is considered as the input to the inference process. Third, we discuss the effect of including the dynamics of MOX sensors by using delayed sensor samples as part of the input space.

For the evaluation we used cross-validation, selecting the folds at the experiment level and not at the sample level. This means that if samples from an experiment have been used during the training procedure, no sample from that experiment was used for calculating the performance measures. In this way an optimistic bias in the results due to evaluation with samples collected in the same trial, i.e. under exactly equal environmental conditions, is avoided. All the experiments have been carried out in the time span of one week and therefore effects due to long term drift are

not considered in this chapter. Furthermore, due to the computational complexity of the training algorithm of the GP (which is dominated by the inversion of the kernel matrix, to be performed at every step of the maximization of the marginal likelihood) a subset of 1000 points from the experiments considered for the training set, was randomly selected for training the GP.

We compare the proposed quantification method with Partial Least Squares Regression (PLSR) and Support Vector Machine Regression (SVR) for all input configurations. Since both methods only provide an estimate of the gas concentration without information about its uncertainty, only the RMSE can be used as a performance indicator for comparison. Please note that PLSR and SVR have been widely used in classical sensor calibration in controlled environments.

4.4.3 Gas quantification using a single sensor

A typical sensor response in an OSS is depicted in Figure 4.3. The PID measurements displayed in Figure 4.3(b) show the fast fluctuations around an average value to which the MOX sensors are exposed when the output of the gas source is steady. These fluctuations are caused by the turbulent airflow and are the reason why the MOX gas sensors do not reach a steady state.

For the case of single gas sensor input, i.e. a univariate function $f : r_i \longmapsto c$, it is possible to plot the relation between sensor resistance and gas concentrations obtained in the experiments (see Figure 4.4). Notice how the distribution of training points (represented by blue dots) corresponds to uncertainty about the measured concentration.

The predictive variance is not necessarily constant across the whole input space, but it depends on the density and dispersion of the training points. If many training points are available in a certain region, then the predictive variance goes down to the global estimate of the signal variance (given by the hyper-parameter σ_f). On the other hand, when few or no training points are available in a region, the predictive variance in that region increases indicating less reliable estimates.

In our case, since we used a uniform distribution to sample the input space for selecting the training points, the posterior variance turns out to be almost constant over the input space (see Figure 4.4(a)). It can be seen in Figure 4.4 that a constant variance does not describe the true signal variance adequately. When the inference is carried out after applying the logarithmic transformation to the sensor resistance and gas concentration, however, the predictive variance represents more accurately the variance in the training points, see Figure 4.4(b). This shows that the process generating the data is indeed not Gaussian and is therefore modelled better by a Log-Normal (non-Gaussian) process. Nevertheless, we can efficiently obtain this non-Gaussian Process by applying a non-linear transformation to the data and then performing Gaussian Process regression.

The estimated gas concentrations for three different gas emitting strategies are displayed in Figure 4.5 and Figure 4.6 for the linear and logarithmic preprocessing respectively. A notable difference exists between the predictive uncertainty in the two

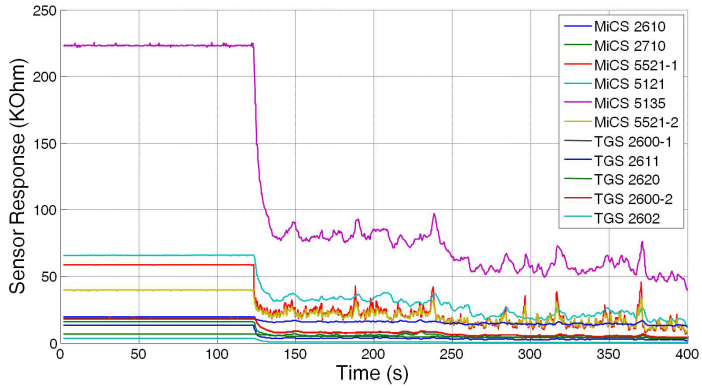

(a)

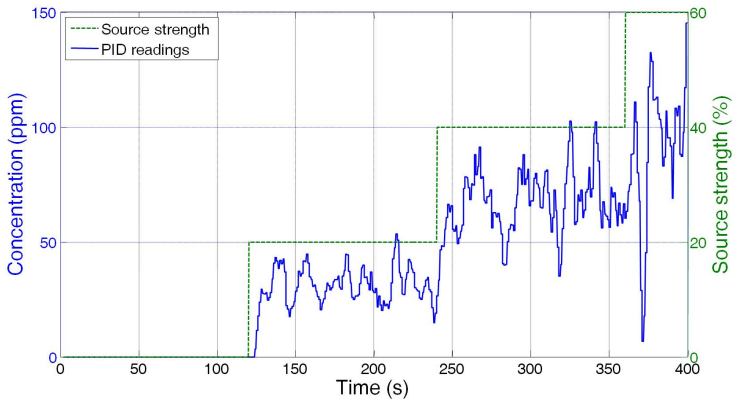

(b)

Figure 4.3: (a) Instantaneous sensor response and (b) PID measurements for a standard open sampling system experiment.

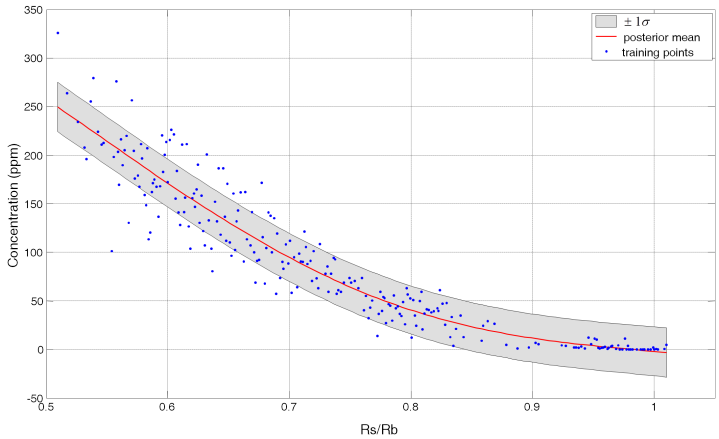

(a)

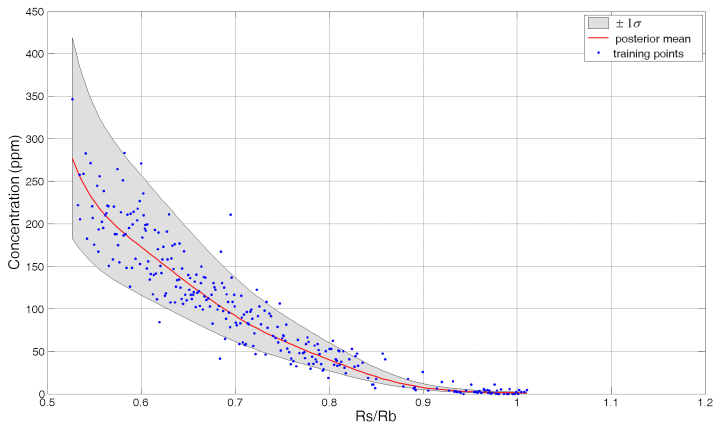

(b)

Figure 4.4: Examples of the learned regression functions for the case where a single sensor (here *TGS-2611*) is considered. Blue dots represent the training data. The solid red line is the mean of the posterior, and its variance is visualized by the $\pm 1\sigma$ confidence interval. (a) A typical calibration function learned when using linear pre-processing, Eq. (4.4). (b) Calibration after applying a logarithmic transformation not only to the input data, Eq. (4.5).

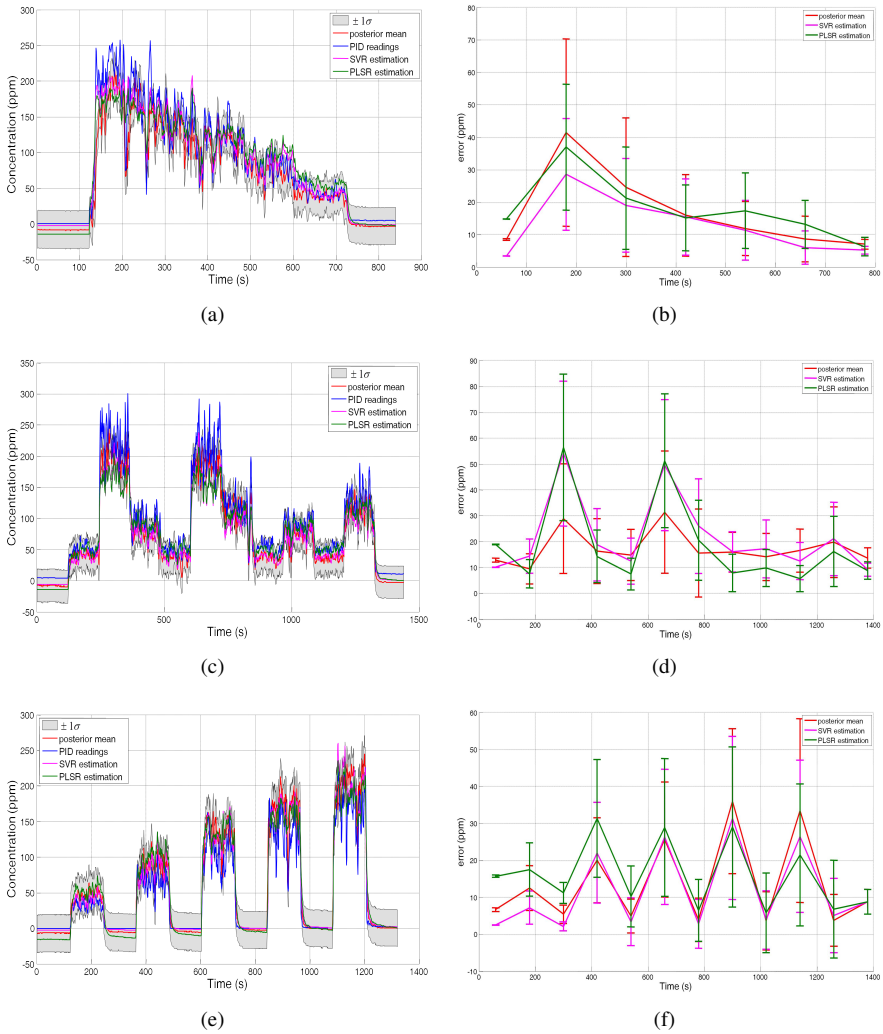

Figure 4.5: Estimation of the gas concentration for three different gas emitting strategies (rows), obtained from sensor *TGS-2611* with linear pre-processing $c = f(\tilde{r}(t))$, and the proposed GP calibration. The left column shows the ground truth (blue) together with the GP calibration estimate (the red line is the posterior mean and the shaded grey region represents $\pm 1\sigma$ confidence interval), the PLSR (green line), and the SVR estimate (magenta line). The right column shows for each scenario the error between the different estimates and ground truth (in the case of GPs only the mean value is taken into account).

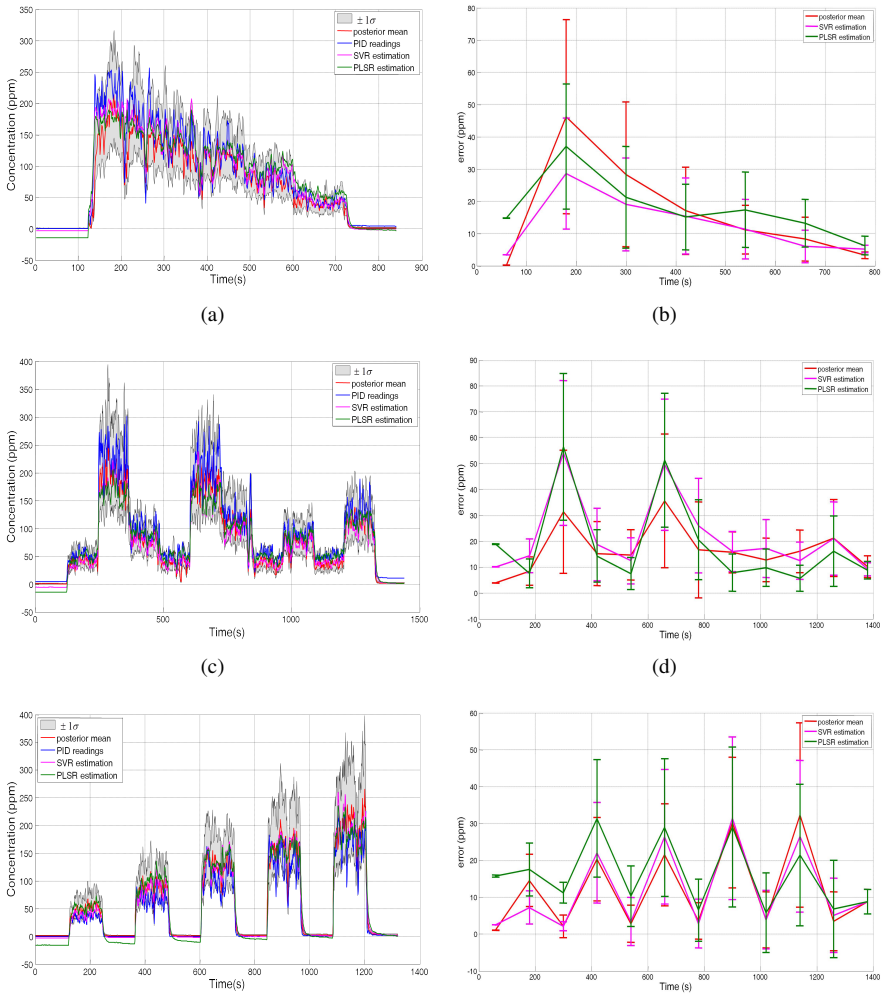

Figure 4.6: The predictions of the same experiment as in Figure 4.5, but after training on log-transformed inputs: $\log(c) = f(\log(\tilde{r}_i(t)))$.

cases. In Figure 4.7 and Figure 4.8 the NLPD and RMSE are plotted for each sensor in the array using a box-plot format. When calibrating using the linear relation we obtain, in average over all the sensors, a RMSE of $35.15 \pm 10.32 ppm$, and a NLPD of 5.16 ± 0.71, while for logarithmic preprocessing the achieved average RMSE is $31.17 \pm 6.27 ppm$ and the NLPD is 4.27 ± 0.18. The results after applying the log transformation in Eq. (4.5) are better according to both performance measures.

From the results in Figures 4.7 and 4.8 we can also see that the sensors *TGS-2611* and *MiCS-5121* perform best for the specific target gas (ethanol).

Table 4.3 summarizes the comparison of the proposed GP quantification with PLSR and SVR. For the same input, the RMSE, averaged over all sensors is $33.77 \pm 5.77 ppm$ for the PLSR approach and $30.94 \pm 6.78 ppm$ for the SVR approach. The PLSR approach performs slightly worse than SVR, and SVR is on par with the GP approach considering only the RMSE. A possible explanation for this result is that PLSR is a linear method while both SVR and GP are non-linear and use an SE (RBF) kernel.

Sensor Model	GPs	PLSR	SVR
MiCS 2610	43.81 ± 22.88	44.79 ± 14.26	46.01 ± 19.44
MiCS 2710	30.38 ± 7.83	31.89 ± 6.13	29.71 ± 6.42
MiCS 5521-1	33.64 ± 9.02	34.26 ± 7.03	34.09 ± 7.56
MiCS 5121	$\mathbf{23.84 \pm 5.58}$	$\mathbf{25.44 \pm 4.20}$	$\mathbf{24.18 \pm 5.91}$
MiCS 5135	32.09 ± 12.00	35.61 ± 8.21	33.37 ± 12.31
MiCS 5521-2	31.11 ± 8.01	32.37 ± 7.47	31.80 ± 8.09
TGS 2600-1	26.05 ± 10.08	31.10 ± 5.42	25.71 ± 8.79
TGS 2611	$\mathbf{21.46 \pm 3.54}$	$\mathbf{26.27 \pm 3.59}$	$\mathbf{21.57 \pm 3.14}$
TGS 2620	34.16 ± 24.02	35.41 ± 7.08	29.61 ± 8.38
TGS 2600-2	28.96 ± 15.13	32.43 ± 5.98	27.23 ± 9.92
TGS 2602	44.64 ± 15.58	41.87 ± 8.21	37.12 ± 12.69
OVERALL	$\mathbf{31.17 \pm 6.27}$	$\mathbf{33.77 \pm 5.77}$	$\mathbf{30.94 \pm 6.78}$

Table 4.3: RMSE values (Mean $\pm$ 1σ) of three different calibration methods when using only one sensor at a time. Sensors *MiCS-5121* and *TGS-2611* provide the overall best performance.

4.4.4 Gas quantification using a sensor array

In OSS applications where the goal is discriminating among several different odors [148], an array of MOX sensors is usually employed instead of a single sensor. For this reason, we investigate whether also gas quantification benefits from using the whole sensor array. In order to do this we apply the GP calibration method with an input space of dimension $d = 11$, and apply ARD (see Section 4.3.3) to automatically select the most relevant inputs, that is, the most relevant gas sensors in the array.

Table 4.4 summarizes the mean and the $\pm$ 1σ confidence interval of the normalized length scale hyper-parameter (l) for the different sensors in the array, after ARD has been computed for the 13 folds used in the cross-validation. As explained

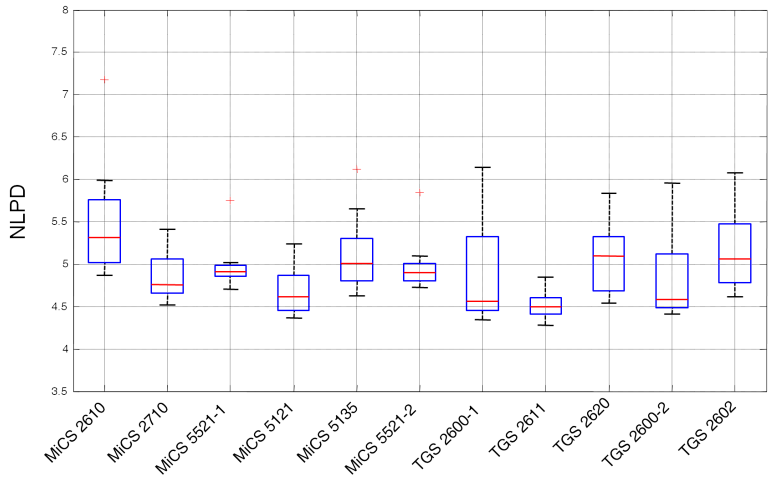

(a)

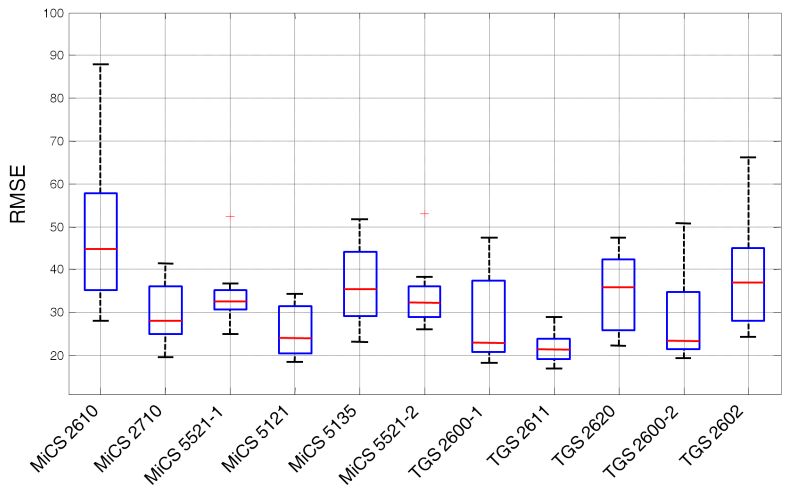

(b)

Figure 4.7: NLPD and RMSE box-plot for the case of one-sensor GP calibration. For both indicators, as lower the value the better the calibration. On each box, the central red mark is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considering outliers, and outliers are plotted individually as red crosses.

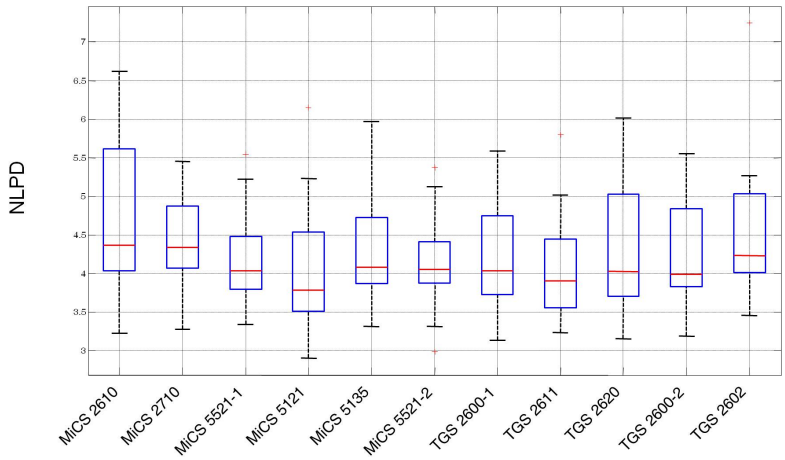

(a)

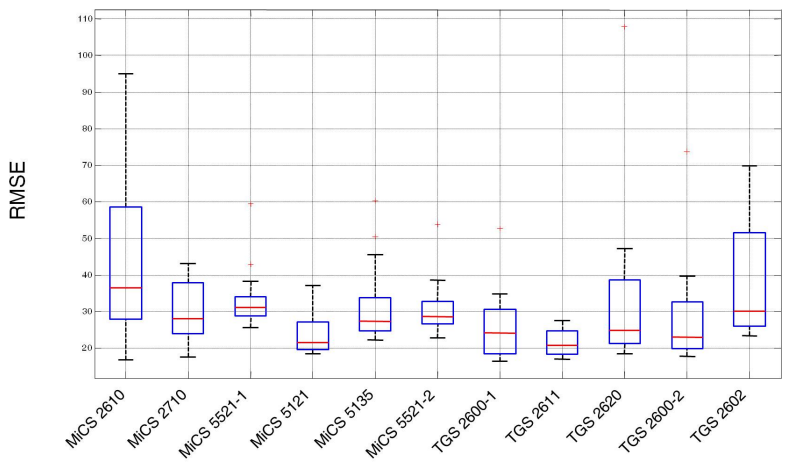

(b)

Figure 4.8: NLPD and RMSE box-plot for the case of one-sensor GP calibration with logarithmic transformation. For both indicators, as lower the value the better the calibration. On each box, the central red mark is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considering outliers, and outliers are plotted individually as red crosses.

in Section 4.3, large values of the length scale corresponds to less relevant sources of information.

Sensor	**Length scale (l)**
MiCS 2610	0.386 ± 0.272
MiCS 2710	0.732 ± 0.337
MiCS 5521-1	0.869 ± 0.145
MiCS 5121	$\mathbf{0.212 \pm 0.175}$
MiCS 5135	0.533 ± 0.312
MiCS 5521-2	0.761 ± 0.305
TGS 2600-1	0.200 ± 0.167
TGS 2611	$\mathbf{0.001 \pm 0.004}$
TGS 2620	0.554 ± 0.387
TGS 2600-2	0.749 ± 0.247
TGS 2602	$\mathbf{0.116 \pm 0.042}$

Table 4.4: Mean $\pm 1\sigma$ interval of the length scale hyper-parameter for the different sensors in the array as obtained from ARD.

Our results show that sensors which perform well in the case of single sensor quantification are in most cases also relevant for array quantification. An exception is the sensor *TGS-2602*, which did not perform well individually, but was found to provide valuable information when considering the whole array (*TGS-2602* is the second most relevant sensor according to ARD but was ranked last in its individual quantification performance, see Table 4.3). Figures 4.9 and 4.10 show gas concentration estimates for three different scenarios when considering the readings of the eleven sensors at once.

In general, the gas concentration estimation obtained using the whole array of sensors outperforms the estimation based on a single sensor. Linear preprocessing provides an average RMSE of $16.44 \pm 3.38 ppm$ and a NLPD of 4.22 ± 0.20, while for the logarithmic preprocessing the results are slightly better: RMSE of $15.97 \pm 2.77 ppm$, NLPD of 3.59 ± 0.76.

In comparison, PLSR applied to the same input achieves a RMSE of $17.55 \pm 2.42 ppm$ which, is slightly worse than the GPs quantification, as in the single sensor case (Section 4.4.3). SVR achieves a RMSE of $16.20 \pm 2.50 ppm$, very similar to the results obtained with the GP approach but without providing the additional information about the uncertainty in the prediction.

The uncertainty estimate is particularly meaningful when the posterior distribution is modelled as a Log-Normal distribution, rather than with a Gaussian distribution. High uncertainty estimates (which can be identified in the left column of Figure 4.10 and Figure 4.8 in the case of predictions performed with a single sensor) correspond to an increased mean and variance of the RMSE (observable in the right column of Figures 4.10 and 4.8). This describes the observed fluctuations in the signal well and thus provides a reliable confidence measure for concentration predictions.

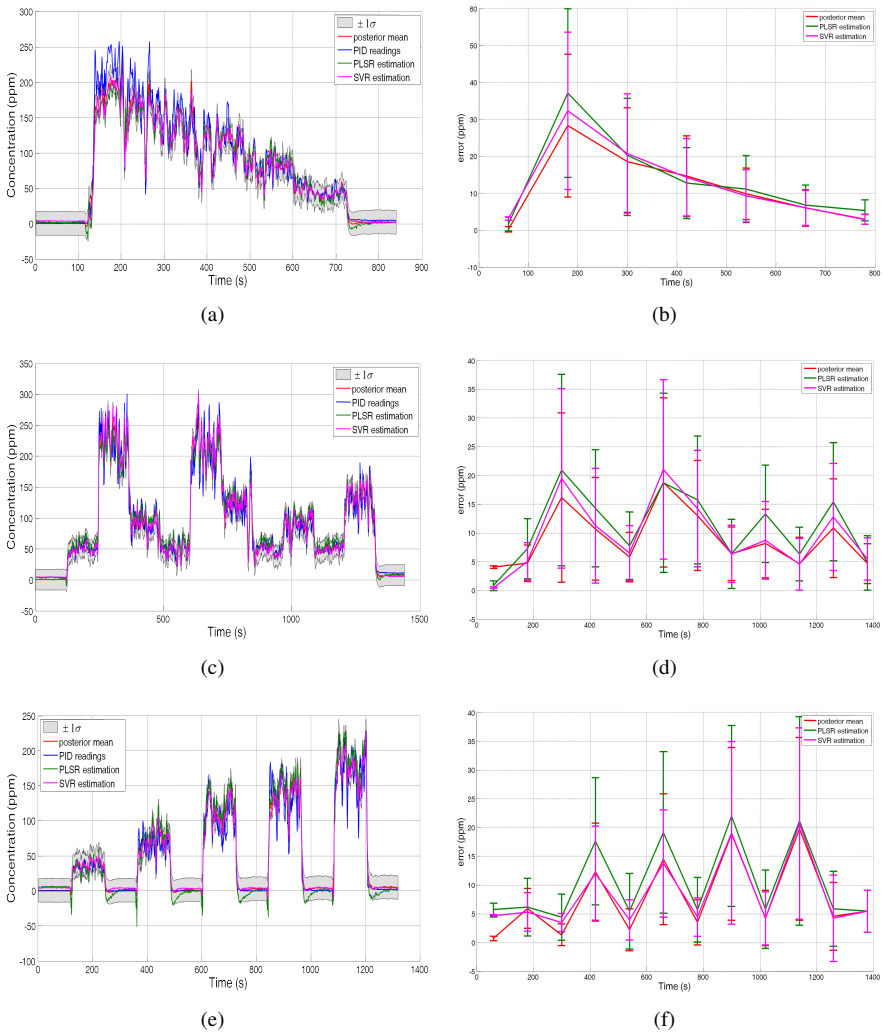

Figure 4.9: Estimation of the gas concentration for three different gas emitting strategies (rows) when considering the readings of the whole array and linear pre-processing $c = f(\tilde{\mathbf{r}}_\mathbf{t})$. The left column shows the ground truth (blue), together with the GP calibration estimate (the red line is the posterior mean and the shaded grey region represents the $\pm 1\sigma$ confidence interval), the PLSR (green line), and the SVR estimate (magenta line). The right column shows for each scenario the error between the different estimates and ground truth (in the case of GPs only the mean value is taken into account).

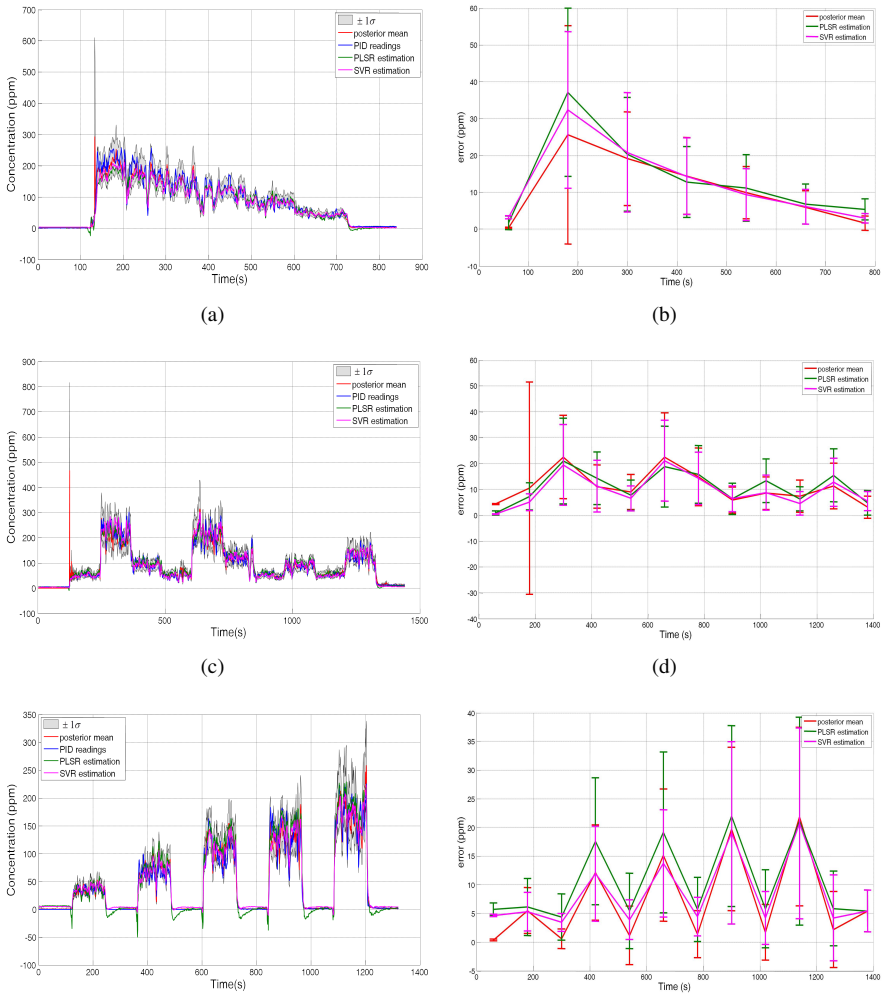

Figure 4.10: The predictions of the same experiment as in Figure 4.9, but after training with the logarithmic relation $log(c) = f(log(\mathbf{rb}(t)))$

Our results also lead to the conclusion that it is recommendable to start using the whole sensor array and select the most relevant sensors with ARD. This allows for better concentration estimate.

4.4.5 Taking into account the dynamics of MOX gas sensors

The signals from MOX gas sensors in an OSS are strongly influenced by the sensor dynamics (see Section 2.3). The uncertainty about concentration estimates is due to the chaotic nature of gas transport in combination with the sensor dynamics, i.e. the non-negligible response and recovery times of the MOX sensors. In this section we propose two different extensions so that the proposed GP quantification can automatically account for the dynamics of MOX sensors. Both methods augment the input signal: the first method ("Memory") by additionally considering delayed samples of every sensor in the input variables (in a so called tapped delay line). The second method ("Derivatives") accounts for the dynamics of each sensor by considering the derivatives of the signal in addition to the signal itself. Figure 4.11 depicts the gas quantification results when considering both alternatives. For the first case ("Memory"), the x-axis $k = 0 \ldots 5$ represents the number of additional delayed samples for each training point, while in the second case it represents the maximum order of derivatives taken into account as additional inputs to the inference process.

The first conclusion from these results is that our GP quantification was not able to infer the dynamics of MOX sensors from delayed samples of the sensor signals, and thus, the quantification results do not improve when increasing k. Furthermore, and possibly due to the increase in the input dimensionality, the results tend to get worse for high values of k, probably due to curse of dimensionality. On the other hand, the "Derivatives" approach yields slightly positive results. The improvement is mainly observable in the RMSE where mean as well as the confidence intervals decrease when increasing k.

4.5 Conclusions and outlook

In this chapter we proposed a new approach for gas concentration estimation using an array of MOX gas sensors in a Open Sampling System (OSS). Despite its importance, this topic has been largely neglected. We addressed the problem in a probabilistic manner and used Gaussian Processes to estimate a posterior distribution over the gas concentration given the response from an array of MOX sensors. This has the advantage of enabling not only predictions of the expected gas concentration but also predictions of the uncertainty of this estimate. This advantage is particularly relevant for OSS applications where typically many sources of uncertainty exist.

In the first part of this chapter, we focussed on gas quantification using a single MOX sensor, and then turned to gas quantification using a sensor array. We found a clearly improved prediction quality with a sensor array compared to using a single sensor. Given the high correlation among different MOX sensors, we used ARD to

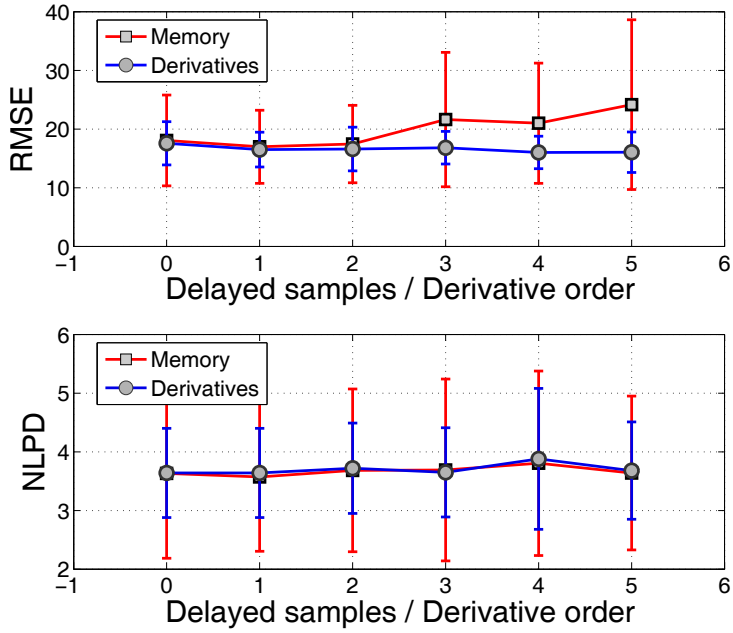

Figure 4.11: Mean and confidence interval ($\pm 1\sigma$) of the NLPD and RMSE of two gas concentration estimation methods, considering: (red) a number of delayed samples and (blue) derivatives with increasing order, as additional inputs.

exclude sensors that are not relevant for estimating the posterior distribution. This proves useful in keeping the dimensionality of the input space low.

We also analysed two data preprocessing strategies, one that performs GP regression directly with the sensor response and ground truth gas concentrations, and a second one that performs GP regression on the logarithms of sensor response and ground truth concentrations. Logarithmic preprocessing has proven advantageous both for the estimation of the expected gas concentration and for uncertainty prediction.

Finally, we studied approaches to mitigate the effect of the slow dynamics of MOX sensors by taking into account past sensor readings in the GP regression. Neither using additional inputs from previous time steps, nor adding the signal derivatives, produced a significant improvement over the concentration estimation algorithm that considers only the current sensors readings.

Future work will include exploring sparse Gaussian Processes like the Relevance Vector Machine (RVM) or Informative Vector Machine (IVM) to improve over the currently random selection of training points, which is a major bottleneck for GP regression. Another interesting aspect to study is the generation of confidence intervals on the prediction and in particular how the Bayesian approach we propose here compares with frequentist approaches like Conformal Prediction (CP). Finally, another aspect to investigate is the use of kernels for time series like the Autoregressive (AR)

or Dynamic Time Warping (DTW) kernel to study if they can efficiently model the dynamics of MOX sensors, and therefore produce even more accurate gas concentration estimates.

Chapter 5
Dealing with obstacles and the ephemeral nature of odors in gas distribution mapping

This chapter addresses the problem of estimating the spatial distribution of volatile substances using a mobile robot equipped with an electronic nose (e-nose). It contributes an effective solution to two important problems that have been disregarded so far: First, obstacles in the environment (walls, furniture, ...) do affect the gas spatial distribution. Second, when combining odor measurements taken at different instants of time, their "ages" must be taken into account to model the ephemeral nature of gas distributions. In order to incorporate these two characteristics into the mapping process we propose modeling the spatial distribution of gases as a Gaussian Markov Random-Field (GMRF). This mathematical framework allows us to consider, for the first time in the gas mapping literature, both: (i) the vanishing information of gas readings by means of a time-increasing uncertainty in sensor measurements, and (ii) the influence of obstacles by means of assumed correlations (and the lack of them) among the different areas. Experimental validation is provided with both, simulated and real-world datasets, demonstrating the out-performance of our method when compared to previous standard techniques in gas mapping.

5.1 Introduction

Gas distribution mapping (GDM) is the process of creating a representation of how gases spread in an environment from a set of spatially and temporally distributed measurements of relevant variables [9, 89]. Foremost, these measurements include the gas concentration itself, but may also comprise wind, pressure or temperature.

In the last decade, GDM is gaining attention in the mobile robotics community because of the advantages a mobile robot offers when compared with the traditional approach based on networks of static gas sensors [150, 71]. To start with, a mobile robot usually carries only one but more expensive and powerful gas sensing device (e-nose), which can analyze more complex compounds, i.e mixtures of different volatiles. Also, the robot can sample at a higher (and adaptive) resolution, while still providing the required accurate localization of each measurement. Moreover, the gas distribution map is created by a robot in an online fashion, allowing decision making to occur depending on such a map, e.g. for exploration tasks. Finally, a mobile robot can leverage environmental information provided by other sensors on board (cameras, laser scanners, etc.) to both enhance the GDM process itself, for example by detecting obstacles, as proposed in this chapter, and to help in any other odor-related task, as can be the identification of potential gas sources [70].

Building a gas distribution model with a mobile robot turns out to be a tough problem for a number of reasons. First of all, and in contrast to most exteroceptive sensors employed in mobile robotics, an e-nose is a *point* sampling device, that is, it only samples the very near air around it. Furthermore, the dispersion of gases is strongly conditioned by the obstacles in the environment, such as walls and furniture. Consequently, when building a GDM, they should indeed be taken into consideration to yield accurate estimations. Finally, but not less important, odors are ephemeral due to the mechanisms that rule gas dispersion, mostly, advection and turbulence (refer to Section 2.3). Thus, we can say that the information conveyed by a given measurement quickly vanishes as time goes by.

The two latest points are pivotal characteristics that have been overlooked by previous works on GDM [96, 9, 89]. Traditionally, the influence of obstacles in the gas distribution has been only considered explicitly for the simulation of the gas dispersal [118, 144], and in some works on plume tracking [99]. However, when facing the GDM problem the information related to obstacles is neglected, leading to maps where nearby areas are always correlated, even when physical obstacles separate them. Moreover, existing approaches to GDM provide the same confidence to all gas measurements regardless of *when* they were taken. As a result, the estimated gas distribution averages out measurements taken at very different moments in time, something that strongly contradicts the vanishing nature of gases (odors).

The example in Figure 5.1, which illustrates the discussion above, shows a robot which is commanded to inspect the different rooms of a house to determine, for example, the possible existence of bad odors. As it moves, the robot collects new gas observations that are incorporated to the GDM. Occasionally the same place is revisited, thus samples from the same location must be somehow combined. Since time-

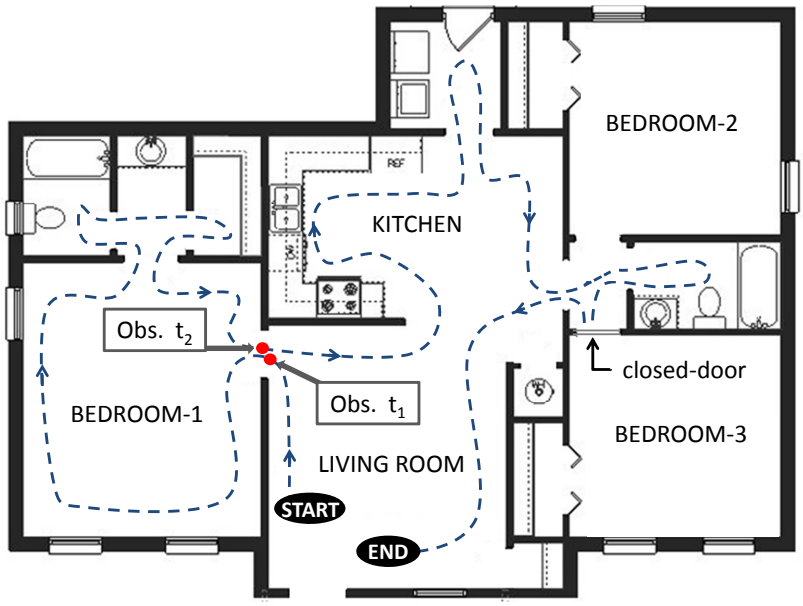

Figure 5.1: An illustrative example where a smelling robot is commanded to inspect an indoor environment by following the predefined path (the blue-dashed line). The robot gathers odor observations as it moves and builds a gas distribution map. Crucial aspects to be considered for such map building include how gas concentrations observed at different instants of time are combined (as for example, t_1 and t_2), and the influence of obstacles such as walls or furniture.

separated observations are gathered at close locations (e.g. samples at t_1 and t_2), the GDM method has to deal with these questions: (i) are observations taken at t_1 and t_2 equally relevant?, and (ii) if not, how do their significance evolve over time? Furthermore, the GDM method has to settle whether or not considering the obstacles present in the environment. That is, to provide an estimation of the gas distribution based only on the observations collected along the covered path, or to additionally account for walls and furniture to model the correlation between the areas they separate (e.g. the presence of bad odors in the kitchen do not imply the same at bedroom-3, since its gas concentration can be considered "independent" given the presence of walls and the closed door).

We can summarize the two contributions of the present chapter as follows. First, we propose accounting for the obstacles in the environment, obtaining maps which are more compliant with the actual mechanisms of gas dispersion. Secondly, we claim that the “age” of a measurement is of relevance in the GDM process. In particular, we propose to associate a time-decreasing weight to each gas measurement, modeling the fact that recent measurements more significantly represent the current gas distribution

than older ones. Thus, observations taken at the same locations and separated in time will be combined according to their respective weights.

As estimation tool for the gas mapping process, we propose to employ a Gaussian Markov-Random-Field (GMRF), which perfectly suits the characteristics of GDM by accounting not only for the information carried by the gas observations, but also for any *prior* knowledge which, in our case, includes both the obstacles in the environment (detected by the robot sensors), and the physics of how gases spatially distribute. A C++ implementation of this method has been integrated into the Mobile Robot Programming Toolkit (MRPT), and its source code is available online[1] at `http://mrpt.org`.

The rest of this chapter is organized as follows. We first discuss the related literature on GDM with mobile robots in Section 5.2, to continue with the introduction of the proposed probabilistic model for GDM in Section 5.3. Then, we show how the maximum a posteriori (MAP) estimation becomes a sparse least squared problem in Section 5.4, and finally, in Section 5.5, we report simulated and real experimental results.

5.2 Related research

We are interested in statistical modeling of gas distributions without making strong assumptions about the environmental conditions (temperature, pressure or airflows). Given that analytical solutions are intractable, it is common practice to divide the space into a regular lattice of cells (gridmap), and then estimate a probability density function (pdf) of the gas concentration at each cell of the grid. Under these circumstances, only a few gas distribution modeling methods have been proposed in the literature.

As mentioned in Section 2.4.3, the most remarkable works in this field have been reported by Lilienthal and colleagues. In the pioneer work [85] they proposed the kernel-based method, which consists of convolving sensor readings with a Gaussian kernel, thus providing a representation of the gas map without assuming any predefined parametric form for the distribution. This method was later extended for the case of multiple odor sources [96] and to the three-dimensional case [124]. It was further shown how gas distribution mapping methods can be embedded into a Blackwellized particle filter approach to account for the uncertainty about the position of the robot [87].

More recently we find approaches that, in addition to providing the most-likely value for the gas distribution, also estimate the uncertainty (via a variance value) for each grid cell of the map. In [89], Lilienthal *et al.* carried out two parallel estimation processes, one for the mean and another for the variance, understanding the latter as the variability of the gas readings, not the uncertainty in the mean estimation which is the standard in probabilistic estimators. In [9], Blanco *et al.* proposed another ap-

[1] In particular, the method described in this chapter is implemented in the class `mrpt::slam::CRandomFieldGridMap2D`, part of the `mrpt-maps` library.

proach, in this case based on a Bayesian interpretation of the problem, which also obtains the variance of each map cell employing a sparsified Kalman filter.

None of these works take into account the constraints imposed by the obstacles of the environment when estimating the gas distribution, neither the physical fact that the information provided by a gas sensor vanishes with time. The latter, however, was pointed out by Assadi *et.al* [5], although, as far as the author knows, no map estimator taking this into account was ever reported. Therefore, the GMRF-based approach proposed in this chapter exploits, for the first time, both concepts.

5.3 Modeling GDM as a Markov Random-Field

In this section we introduce the basis for the estimation of the gas distribution over a 2D lattice of cells using GMRFs. We also describe the highly-sparse structure of the problem which leads to efficient estimates of the problem and, finally, we present our model for observation time-varying uncertainty.

5.3.1 Probabilistic model for GDM

The proposed approach aims at estimating the probability density function of the gas concentration in an environment. As in most previous works on GDM we simplify the problem by estimating a discrete two-dimensional map, dividing the space into a rectangular lattice of *cells*. A map $\mathbf{m} = \{m_i\}_{i=1}^{N}$ is then modeled as a random field where m_i are scalar variables standing for the gas concentration inside the i'th cell with coordinates (x_i, y_i). Let N be the overall number of variables in the map, such that if the map is $N_x \times N_y$ cells, $N = N_x N_y$. Notice that this model resembles occupancy grids in robotics, with the difference of not holding a discrete distribution (occupied vs. free) but a continuous magnitude. Concretely, the gas concentration at each cell is modeled as a univariate Gaussian distribution $\mathcal{N}\left(\mu_i, \sigma_i^2\right)$, with μ_i its mean estimate and σ_i^2 the corresponding uncertainty (see Figure 5.2).

Our goal is obtaining the maximum a posteriori (MAP) estimation of $\mathbf{m}$, along with its uncertainty, given the gas concentrations measured by the robot e-nose (random variables $\mathbf{z}$) and some prior knowledge that includes (i) how the gas spread over the environment, and (ii) how the perceived obstacles affect the propagation of gases between nearby cells. Given the small space sampled by an e-nose (even when employing pumps or fans to aspire the air), this prior is extremely important for inferring the gas concentration at distant locations not subject to direct sensing.

Our proposal is to use a Markov Random-Field (MRF), a tool widely employed in other estimation problems on grids. For example, in image processing, where statistical models are defined for the intensity of image pixels [157]. Notice the strong analogy between problems such as image de-noising or image restoration and the GDM stated here, where gridmap cells play the role of pixels.

According to the *Hammersley-Clifford* theorem [18], the joint probability distribution $p(\mathbf{m}, \mathbf{z})$ can be expressed as a *Gibbs* distribution, that is, it can be factored

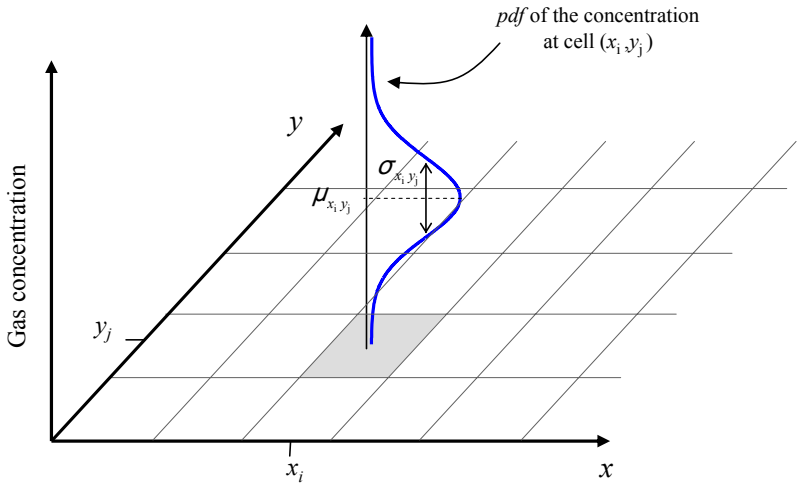

Figure 5.2: The 2D map is represented by a lattice where each cell keeps the estimate of gas concentration by means of a Gaussian density, represented here along the vertical axis.

as the product of the potential functions $\varphi(\cdot)$ for the set of all its maximal cliques ($\mathscr{C}_m$) [6]:

$$p(\mathbf{m},\mathbf{z}) = \frac{1}{Z}\prod_{\mathscr{C}\in\mathscr{C}_m}\varphi_{\mathscr{C}}(\mathbf{n}_{\mathscr{C}}) \tag{5.1}$$

where the proportionality constant Z (called the *partition function*) is not relevant in our problem, $\mathscr{C}$ denotes the different cliques and $\mathbf{n}_{\mathscr{C}}$ the set of variables $(\mathbf{m},\mathbf{z})$ in that clique.

Since we are restricted to potential functions which are strictly positive ($\varphi(\cdot) > 0$), it is convenient to express them as exponentials:

$$\begin{aligned} p(\mathbf{m},\mathbf{z}) &\propto \prod_{\mathscr{C}\in\mathscr{C}_m}\exp\{-E(\mathbf{n}_{\mathscr{C}})\} \\ &= \exp\left\{-\sum_{\mathscr{C}\in\mathscr{C}_m}E(\mathbf{n}_{\mathscr{C}})\right\} \end{aligned} \tag{5.2}$$

where E is the energy function, obtained by adding up the energies of each of the maximal cliques.

An intuitive and convenient way of dealing with the dependencies encoded in a MRF is to consider its *factor graph* [20] as the graphical model from which to derive the optimization equations. In this graphical model, each potential functions $\varphi(\cdot)$

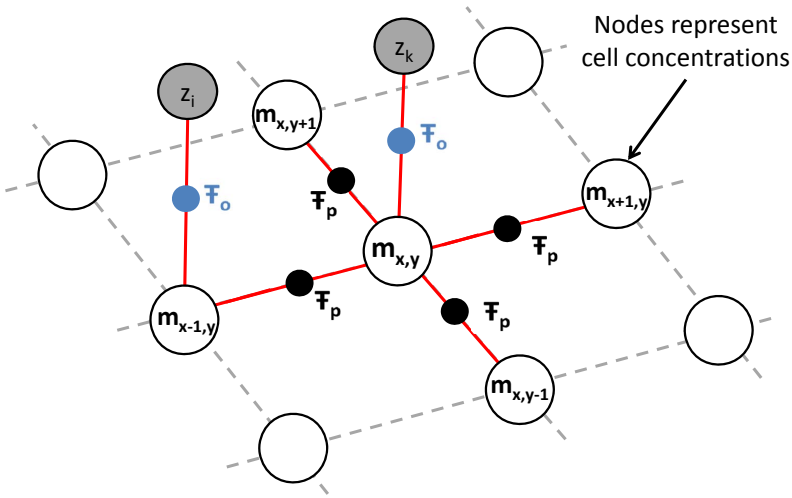

Figure 5.3: Factor graph derived from the MRF employed in our approach. There are two types of nodes: gas concentrations at cells (white circles), and gas observations (grey shaded circles), and two kind of factors: prior factors ($\mathscr{F}_p$), and observation factors ($\mathscr{F}_o$).

over a maximal clique becomes a factor $\mathscr{F}$. As shown in Figure 5.3, this model comprises two kinds of nodes: (i) gas concentrations at cells (unknowns to be estimated), and (ii) gas observations (known data). We also define two distinct sets of factors between nodes: *observation factors* which represent sensor observations and constrain the concentration value of a cell i according to all sensor measurements taken by the robot at that cell, and *prior factors*, which being independent of observations capture our a priori knowledge on how the gas distribution behaves over space.

Attending to the two different set of factors, the joint probability distribution can then be expressed as:

$$p(\mathbf{m},\mathbf{z}) \propto \exp\left\{-\sum_{\mathscr{C}_o} E_o(\mathbf{n}_{\mathscr{C}_o}) - \sum_{\mathscr{C}_p} E_p(\mathbf{n}_{\mathscr{C}_p})\right\} \tag{5.3}$$

5.3.2 Factor parameters

We assume that all the conditional distributions involved in the problem can be reasonably modeled as Gaussians, thus the underlying graphical model becomes a Gaussian MRF (GMRF). This assumption works well in practice, as demonstrated experimentally. Therefore, we need to provide the parameters of each Gaussian distribution that appears in our graphical model in order to have it completely defined.

Observation factors

They encode the observation model, that is, the relationship between an e-nose reading and the true gas concentration of the cell at which it was taken.

Let M be the number of e-nose observations collected by the robot. Each observation consists of a gas concentration value z_k taken at a particular cell i_k at a given instant of time t_k, with $k = 1 \dots M$. Each such observation is assumed to be corrupted with two additive Gaussian errors: one from the inherent sensor noise ($\omega_k \sim \mathcal{N}(0, \sigma_s^2)$) and another time-dependent term that models the potential changes that may have occurred since the sensing time ($\zeta_k \sim \mathcal{N}(0, \sigma_\zeta^2(t - t_k))$). With this last noise, we model the loss of information of a measurement as an increase of uncertainty (variance). Denoting the ideal (noiseless) sensor model as $h(\mathbf{m})$, we have:

$$z_k = h(\mathbf{m}) + \omega_k + \zeta_k \tag{5.4}$$

Given the moderately large size of grid cells in GDM (typically in the range of decimeters) it becomes reasonable to assume that every measurement is associated to one single cell, the one which the robot e-nose is sniffing at, which takes us to the minimalistic sensor model $h(\mathbf{m}) = h(m_{i_k}) = m_{i_k}$.

Under a probabilistic point of view, each observation factor in the graphical model then stands for the conditional pdf:

$$p(z_k|\mathbf{m}) = p(z_k|m_{i_k}) = \mathcal{N}(m_{i_k}, \sigma_s^2 + \sigma_\zeta^2(t - t_k)) \tag{5.5}$$

where we have applied the conditional independence between z_k and the rest of the cells given m_{i_k}.

Then, the energy function associated to the observation factors can be expressed as:

$$E_o = \sum_{\mathscr{C}_o} E_o(\mathbf{n}_{\mathscr{C}_o}) = \sum_{k=1}^{M} \frac{(m_{i_k} - z_k)^2}{\sigma_s^2 + \sigma_\zeta^2(t - t_k)} \tag{5.6}$$

We must stress the novelty of the time-increasing variance $\sigma_\zeta^2(t - t_k)$ above: if two observations from the same cell are combined to estimate its gas concentration, this variable variance gives more weight to the most recent one. This physical-grounded concept has been obviated in previous works.

Eventually, during a GDM process the variance of older measurements will become large enough as to neglect them. Thus, in practice, only a finite set of M observations will account for the estimation, bounding the overall computational complexity of our method for a fixed-size map.

Prior factors

These factors capture the knowledge about how gases distribute spatially. Particularly, we want to model the correlation between gas concentrations of neighboring cells.

We insist in the necessity and relevance of this term in GDM, since gas observations provide us very localized information, i.e. only for the cell at which the sample was taken.

Correlation between cells: Previous approaches based on kernel methods [89] or Kalman filtering [9] have modeled these correlations as Gaussian functions. Similarly, we model the correlation between cells by penalizing the difference in the gas concentration (l_{ij}) between pairs of (vertically and horizontally) adjacent cells:

$$l_{i,j} = \ m_i - m_j \tag{5.7}$$

where m_i, m_j are the gas concentrations at adjacent cells with lattice indices i and j, respectively. Each prior factor then stands for the following pdf:

$$p(l_{i,j}|\neg o_{i,j}) = \mathcal{N}(0, \sigma_p^2) \tag{5.8}$$

meaning that adjacent cells are forced to have the same gas concentration with a tolerance stated by σ_p^2. The meaning of the conditioning on $\neg o_{i,j}$ is explained next.

Obstacles: Environment obstacles and their influence in the gas distribution must be accounted for while modeling the expected difference between adjacent cells in Eq. (5.7). We assume that the probability of their intermediary space to be occupied, $P(o_{i,j}) \in [0,1]$, is readily available in the form of an occupancy grid representation of the environment. Note that the cell size of this grid is not required to match that of the gas map and, in practice, will often be finer.

Denoting as $o_{i,j}$ the fact that an obstacle exists between cells i and j, we can then apply the law of total probability over the two only possibilities (either $o_{i,j}$ or $\neg o_{i,j}$) to obtain:

$$p(l_{i,j}) \quad = \quad p(l_{i,j}|o_{i,j})P(o_{i,j}) + p(l_{i,j}|\neg o_{i,j})P(\neg o_{i,j}) \tag{5.9}$$

where $P(\neg o_{i,j}) = 1 - P(o_{i,j})$ and $p(l_{i,j}|\neg o_{i,j})$ was already given in Eq. (5.8). Regarding our a priori distribution for $l_{i,j}$ in the case of an obstacle blocking the way between two cells, it seems reasonable to assume no correlation at all. Then, a good candidate for $p(l_{i,j}|o_{i,j})$ is a uniform distribution over a sufficiently large interval. Since mixing Gaussians and uniform densities would prevent the formulation of the estimator as a least-squares problem, the following approximation is conveniently proposed:

$$p(l_{i,j}) \quad \approx \quad \mathcal{N}\left(0, \frac{\sigma_p^2}{(1 - P(o_{i,j}))^2}\right) \tag{5.10}$$

which only exactly matches the ideal model in Eq. (5.9) for $P(o_{i,j})$ equals to 0 or 1, while providing a smoothly changing Gaussian model for intermediary values. Since the relevant parts of the map will have occupancy probabilities close to these extremes

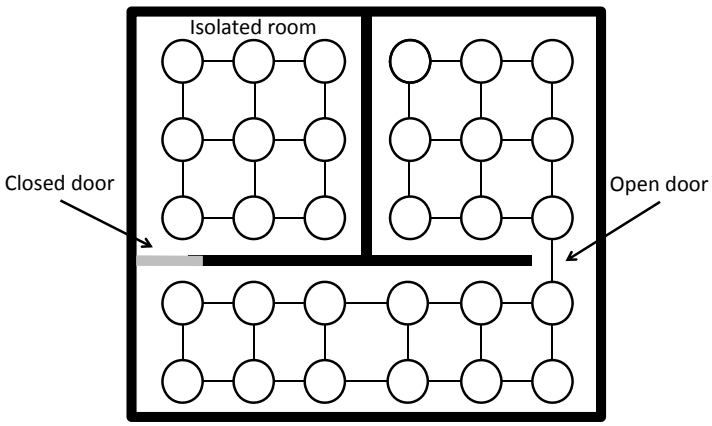

Figure 5.4: Example MRF which considers the physical obstacles in the environment. Nodes represent the gas concentration at cells and edges encode the a priori assumption of correlation between some adjacent nodes. Factors and observations have been omitted here for clarity.

it becomes a minor issue that Eq. (5.10) only poorly models Eq. (5.9) in unexplored areas where $P(o_{i,j})$ is close to 0.5. As depicted in Figure 5.4, when two adjacent cells are physically separated by an obstacle ($P(o_{i,j}) \to 1$), no correlation is assumed between these cells. Note that this leads to infinite variance in $Eq.$ (5.10), which is not problematic since the estimator ultimately handles inverse variances instead.

Finally, the energy function capturing the prior constraints reads:

$$E_p = \sum_{\mathscr{C}_p} E_p(\mathbf{n}_{\mathscr{C}_p}) = \sum_{k=1}^{L} \frac{(m_{i_k} - m_{j_k})^2}{\sigma_p^2/(1 - P(o_{i_k,j_k}))^2} \tag{5.11}$$

with L the number of pairwise cliques of cell nodes in the GMRF, and i_k, j_k the adjacent cells for each such pairwise clique k.

5.4 Maximum a posteriori estimation of the GDM

We show next how the MAP estimation of the GDM becomes a least-squares problem for the proposed GMRF model. We also describe that the uncertainty of the estimated map can be retrieved from its graphical model.

5.4.1 Derivation

We start by conditioning the gas concentration map $\mathbf{m}$ to all available data, that is, to all gas observations $\mathbf{z} = \{z_1, \ldots, z_M\}$. Then, we seek to maximize the posterior $p(\mathbf{m}|z_{1:M}) \propto p(\mathbf{m}, z_{1:M})$ which gives us the MAP estimate $\hat{\mathbf{m}}$.

By taking the negative logarithm over such posterior, the complete energy function becomes the well-known least-squares form of a GMRF inference problem, which in our case reads:

$$E(\mathbf{n}) \quad = \quad \sum_{k=1}^{L} \frac{(m_{i_k} - m_{j_k})^2}{\sigma_p^2/(1-P(o_{i,j}))^2} + \sum_{k=1}^{M} \frac{(m_{i_k} - z_k)^2}{\sigma_s^2 + \sigma_\zeta^2(t-t_k)} \tag{5.12}$$

We can rearrange the terms of the energy function $E(\cdot)$ as a sum of quadratic errors $\mathbf{r}$ weighted by an information matrix Λ, i.e. $E(\mathbf{n}) = \mathbf{r}^\top \Lambda \mathbf{r}$. Errors can be conveniently defined in terms of a prediction function $\mathbf{f}(\mathbf{m})$ such that $\mathbf{r} = \mathbf{f}(\mathbf{m}) - \mathbf{y}$ for some vector of known data $\mathbf{y}$. From the assumed statistical independence between variables and model noises, it follows that Λ is diagonal, leading to:

$$E(\mathbf{n}) = \mathbf{r}^\top \Lambda \mathbf{r} = \sum_{k=1}^{L+M} \Lambda_k (f_k(\mathbf{m}) - y_k)^2 \tag{5.13}$$

where the k subscript denotes the corresponding scalar entry in each matrix or vector. The minimum of the quadratic expression in Eq. (5.13) can be found by solving the *Newton method* equations [31, 21] (refer to Appendix D):

$$\underbrace{(\mathbf{J}^\top \Lambda \mathbf{J})}_{\text{Hessian } \mathbf{H}} \Delta\mathbf{m}^* = -\underbrace{\mathbf{J}^\top \Lambda (\mathbf{f}(\mathbf{m}) - \mathbf{y})}_{\text{Gradient } \mathbf{g}} \tag{5.14}$$

where $\mathbf{J} = \frac{d\mathbf{r}}{d\mathbf{m}}$ is the Jacobian of the error function $\mathbf{r}$, and $\Delta\mathbf{m}^*$ is the increment that leads to the MAP estimate $\hat{\mathbf{m}}$.

It should be emphasized that unlike other mapping problems in mobile robotics, both factor types ($\mathscr{F}_p, \mathscr{F}_o$) are linear with the map $\mathbf{m}$, which implies that $\hat{\mathbf{m}}$ can be solved in closed form, without iterating.

Matching the generic Eq. (5.13) to our particular case in Eq. (5.12) we have:

$$\begin{aligned} \mathbf{f}(\mathbf{m}) = & \quad [\quad l_1(\cdot) \quad \cdots \quad l_L(\cdot) \quad | \quad m_{i_1} \quad \cdots \quad m_{i_M} \quad]^\top \\ \mathbf{y} = & \quad [\quad 0 \quad \cdots \quad 0 \quad | \quad z_1 \quad \cdots \quad z_M \quad]^\top \end{aligned} \tag{5.15}$$

from which in section 5.4.2 we will derive the sparse Jacobian expressions.

Regarding the $(L+M) \times (L+M)$ information matrix Λ, its first L diagonal entries correspond to the prior factors, that is, to the correlation between adjacent cells, and is $\Lambda_{p_k} = (1 - P(o_{i,j}))^2/\sigma_p^2$. The rest M diagonal entries are the weights of e-nose observations, which decrease over time according to their "age", that is, $\Lambda_{o_k} = 1/(\sigma_s^2 + \sigma_\zeta^2(t-t_k))$.

As a important remark for an efficient implementation, the $\mathbf{Ax} = \mathbf{b}$-like system of equations to be solved in Eq. (5.14) is highly sparse due to the strongly local structure of the constraints, leading to a symmetric banded Hessian matrix with a bandwidth of $W = \max(N_x, N_y)$. Then, a sparse $LL^\top$ (Cholesky) decomposition can be

used for efficiently factoring and solving the system in $\mathcal{O}(N_x N_y W^2)$ and $\mathcal{O}(N_x N_y W)$ time, respectively [7, p.220]. This means that, for example, estimating a square map with $N = N_x N_y$ cells has an overall cost that grows with $\mathcal{O}(N^{1.5})$.

5.4.2 Jacobian, Hessian and Gradient

Based on Eq. (5.14) and the particular structure denoted in Eq. (5.15), next we devise the structure of the system matrices:

- Jacobian $\mathbf{J}$: The $\mathbf{J}$ matrix contains the $\frac{d\mathbf{r}}{d\mathbf{m}}$ for every factor in the graph. Rows of the Jacobian derived from *prior factors* contain zeros but for the i'th and j'th columns, corresponding to adjacent cells in the map (see Eq. (5.11)), which have values 1 and -1, respectively. Rows for *observation factors* are all zeros except at the column of the observed cell.

$$\mathbf{J} = \begin{array}{c} \\ 1 \\ 2 \\ \vdots \\ L \\ \hline L{+}1 \\ L{+}2 \\ \vdots \\ L{+}M \end{array} \begin{array}{c} \begin{array}{cccccccc} m_1 & m_2 & \cdots & m_i & \cdots & m_j & \cdots & m_N \end{array} \\ \left[\begin{array}{cccccccc} 1 & -1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ 1 & 0 & \cdots & -1 & \cdots & 0 & \cdots & 0 \\ & & & \vdots & & & & \\ 0 & 0 & \cdots & 1 & \cdots & -1 & \cdots & 0 \\ \hline 0 & 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ & & & \vdots & & & & \\ 0 & 0 & \cdots & 1 & \cdots & 0 & \cdots & 0 \end{array}\right] \end{array}$$

- Hessian $\mathbf{H}$: Since all functions in our problem are linear, the Hessian is exactly $\mathbf{J}^\top \Lambda \mathbf{J}$. The existence of two blocks in the Jacobian matrix, with its upper block corresponding to the prior factors which typically will not change over time, advices us to decompose the Hessian into the sum of two components:

$$\mathbf{H} = \mathbf{H}_p + \mathbf{H}_o \tag{5.16}$$

  The first part, $\mathbf{H}_p$, only contains the following nonzero entries:

  - Each off-diagonal entry $H_p(i,j)$ is the sum $\sum_k J(k,j)\Lambda_{p_k}$ for each *prior factor* k relating cells (i,j). Following the graphical model in Figure 5.3, and attending to the sparse structure of the Jacobian, in our case $H_p(i,j) = -\Lambda_{p_k}$ if cell i is adjacent to cell j, zero otherwise.
  - Each diagonal element $H_p(i,i)$ becomes the sum $\sum_k J(k,i)\Lambda_{p_k} J(k,i) = \sum_k \Lambda_{p_k}$, for each *prior factor* k defined over cell i, that is, for each neighbor of cell i.

The second part, $\mathbf{H}_o$, is exactly diagonal and depends only on the observations. The $i'th$ element of its diagonal $H_o(i,i)$ amounts to $\sum_k \Lambda_{o_k}$, being k the index of each observation taken at cell i.

- Gradient vector $\mathbf{g}$: The gradient vector $\mathbf{g} = \mathbf{J}^\top \Lambda \mathbf{r}$, with length the number of grid cells N, simply becomes:

$$g_i = \sum_k \Lambda_{o_k}(m_i - z_k) + \sum_j \Lambda_{p_k}(m_i - m_j) \tag{5.17}$$

for all the observations k taken at cell i, and all the neighbor cells j of cell i.

5.4.3 Recovering the uncertainty

To obtain the uncertainty of the gas distribution at each cell, we have to compute the diagonal of $\mathbf{H}^{-1}$. Each diagonal element $\mathbf{H}^{-1}(i,i)$ corresponds to the variance of cell i (σ_{ii}^2).

Given that matrix inversion is a computationally expensive operation, and since we are only interested in recovering the diagonal of the inverse Hessian, efficient approximations as the one presented by Golub and Plemmons [48] can be employed. They proposed a method for recovering only the entries σ_{ij} of the covariance matrix Σ that coincide with nonzero entries in the factor matrix $\mathbf{R}$, being $\mathbf{R}$ the upper triangular matrix that results from the Cholesky decomposition of $\mathbf{H}$.

5.5 Experimental evaluation and discussion

This section presents an experiment aimed at validating the performance of the proposed method when estimating the GDM of a time-variant gas distribution in an indoor scenario. The experiment consists of a mobile robot which is patrolling an area with several rooms while building a GDM of the inspected environment.

To provide quantitative results of the accuracy of the estimated gas maps, as well as to allow for a stringent comparison between different GDM methods, we first perform the experiment in simulation in order to have a ground truth (GT) of the gas distribution. Then, we carry out a real experiment in a very similar scenario.

5.5.1 Experiment setup

Both the real and simulated experiment are conducted in an indoor scenario composed of two adjacent rooms connected through a corridor (see Figure 5.5). Each room is inspected through a complete loop by a mobile robot equipped with an e-nose which collects samples along the path (red-dashed line). To go from one room to the other, the robot traverses the corridor (path marked as blue-dashed line).

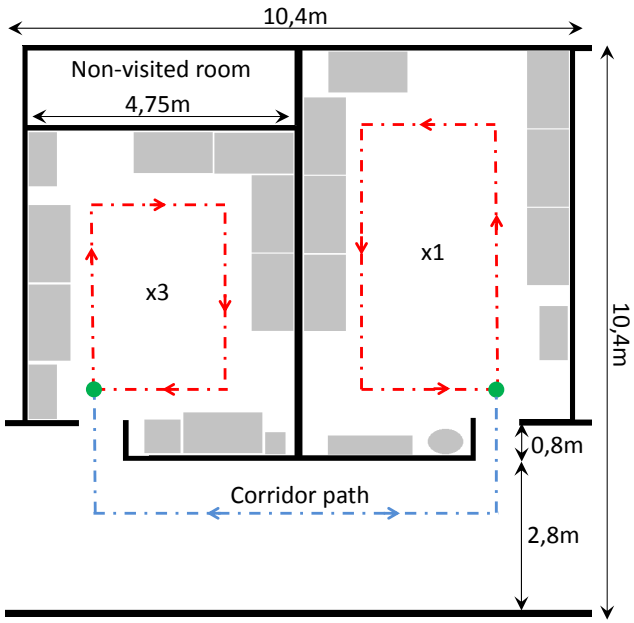

(a)

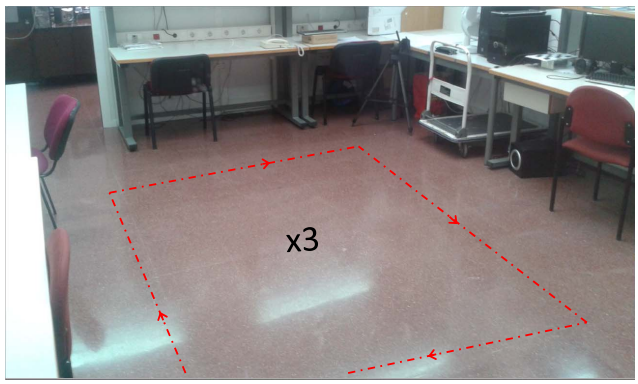

(b)

Figure 5.5: (a) Sketch of the experimental scenario. White color represent free space, grey color obstacles as tables or cabinets, and black color walls. (b) Picture of the leftmost room of the above plan. The path followed by the robot is marked as a dashed line.

To study how different GDM methods behave under time-variant gas distributions, both rooms are alternatively inspected by repeating the pattern "left(x3), right(x1)", under the following three cases: (i) when both rooms are clean of odors, (ii) when an

ethanol leak appears in the leftmost room, and (iii) when the ethanol odor is removed and both rooms become clean again.

We compare the gas maps as estimated by the GMRF approach presented in this chapter, with the Kernel DM+V [89], which is probably the more accurate and effective existing method for GDM. In particular, we take snapshots at the following representative instants of time: (t_0) at the start of the experiment, (t_1) after inspecting both rooms when no odor has yet been released, (t_2) once the ethanol has spread in the leftmost room, and it has been inspected once, (t_3) after the three inspections (loops) in the leftmost room with odor presence, (t_4) after the first inspection of the leftmost room, again clean of odors, and finally (t_5) after the three inspections of the leftmost room with no odor presence.

In both, the simulation and the real experiment, the cell size of the gas gridmap is set to 0.5 m and the robot average speed to 0.25 m/s. Regarding the GMRF parameters, we set $\sigma_p^2 = 2$ and restrict occupancy probabilities to binary values ($P(o) = 0$ for free space, $P(o) = 1$ for obstacles), which gives a value for Λ_p of 0.5 and 0, respectively. We also set $\Lambda_o(\Delta t = 0) = 10$ and decreases it over time at a rate of 0.12336 per second, which corresponds to an observation lifespan of 80 s. For the Kernel DM+V method, we set the kernel width to $\sigma = 0.4$, which governs the amount of extrapolation of individual gas readings, while the cutoff radius for updating close cells is set to 0.7 m.

For each instant of time t_i, the path followed by the robot will be displayed as a grey-thin line, highlighting the last minute of it as a thick-white line. The start point is marked as a circumference, and the current robot localization as a triangle.

5.5.2 Simulated experiment

For the synthetic scenario we have employed the simulation framework presented in [36] and described in Appendix B, which allows an easy comparison between the ground truth of the gas distribution and the maps estimated by different GDM methods.

Figure 5.6 displays the ground truth and the estimated mean and variance maps for the GMRF and Kernel DM+V methods, at the different instants of time (t_i). For an easier interpretation of results we have highlighted obstacles (walls and office furnishings) in the map in white color.

The following conclusions can be drawn from this experiment:

1. At the beginning of the experiment (t_0), both methods initialize their estimate as absence of gases, providing a zero mean map. However, their variance maps differ considerably, highlighting the following important difference between our proposal and the Kernel DM+V:

    - For the GMRF method, the variance map provides the uncertainty about the estimated mean. High variance values indicate that the estimation at such cells cannot be trusted because of the lack of information. This is

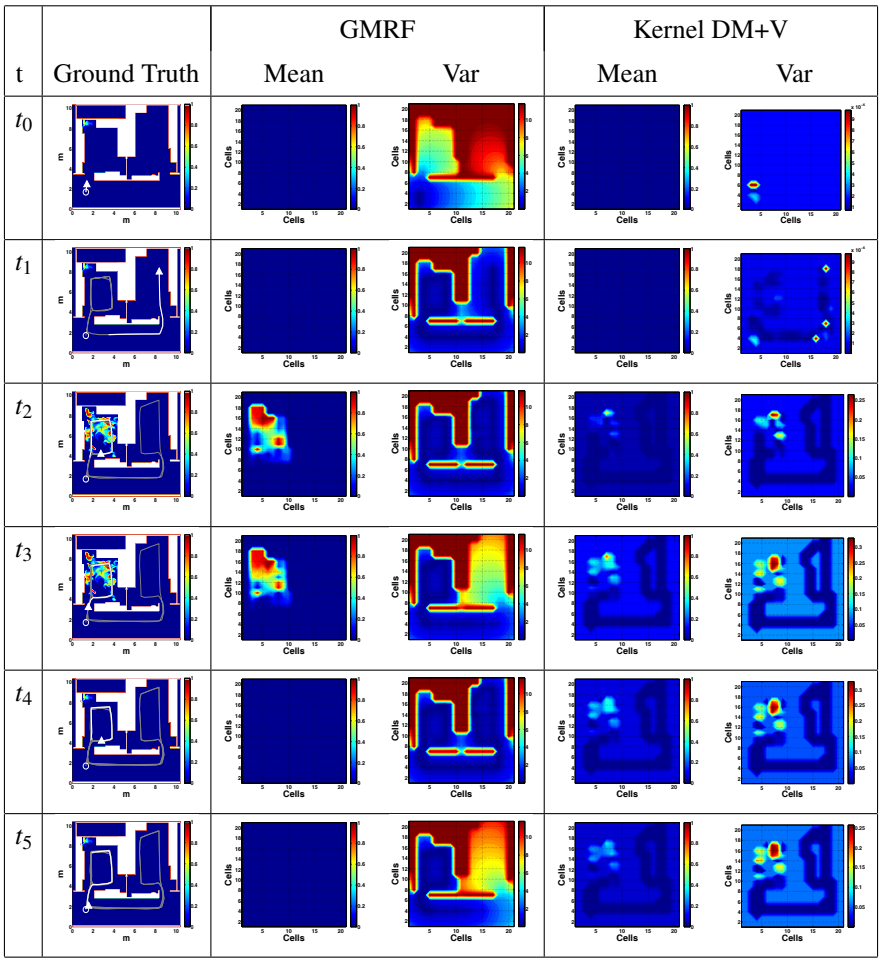

Figure 5.6: Mean and variance gas distribution maps estimated by two different methods (GMRF and Kernel DM+V) at the different time steps t_i of the simulated experiment.

the case of cells distant to the robot path, cells not measured recently, or cells not correlated to those that have been measured (i.e cells separated by obstacles).

- For the kernel DM+V, however, the variance map represents the estimated variability of the gas distribution, that is, cells with different observation values exhibit a high variance. Since it does not account for the obstacles in the environment, it provides a less satisfactory variance value of zero for all non-sensed cells.

2. At time step t_1, after inspecting both rooms for the first time, both methods estimate correctly the absence of ethanol in the environment as appreciated in their respective mean maps. Since e-nose observations are now available in both rooms, the GMRF provides low variances for the visited cells, while for those falling on obstacles or in non-visited rooms (as the isolated room at the top left of the map), still maintain high variance values, as desired.

3. Once the gas is released, we take snapshots at two time instants: t_2 after only one loop through the contaminated area, and t_3 after three loops, as displayed in Figure 5.5. Since the Kernel DM+V (as any other previous GDM method) does not consider the time at which observations are taken, the estimated mean maps are the result of averaging recent observations with older ones, that is, with observations gathered when no gas was still released. Notice how at such time instants, the kernel DM+V method fails to detect the high concentration values present in the ground truth maps. This, which represents one of the main limitations of existing GDM methods, is successfully overcome by the proposed approach.

   By increasing the uncertainty of observations as they become older, together with the fact that we consider the presence of obstacles in the environment, our approach is able to detect and correctly localize the high gas emissions with only one lap over the contaminated zone. The main difference between time-instants t_2 and t_3, is that at t_3 the variances at the right room start increasing as a consequence of the rise in the uncertainty of observations, which allows us to remove them from the set of observations.

4. Finally, at time steps t_4 and t_5, when the gas has been removed, the "average" effect of the Kernel DM+V method can still be appreciated as gas patches in the mean map. On the contrary, our approach adapts faster to changing gas concentrations, and so it correctly provides a zero mean map at the left room, even after only one lap (t_4).

For the purpose of providing a quantitative comparison of both GDM methods, we choose the negative log probability density (NLPD) of the ground truth state evaluated in the probabilistic map estimation. The resulting magnitude not only accounts for the mean estimations, but also considers the respective variance values. Figure 5.7

summarizes such comparison for the different time-steps (t_i) of the simulation, where low NLPD values are desired.

As can be appreciated, our approach performs much better than the kernel DM+V method for all the studied time-steps t_i. This notable improvement in the NLPD is the result of providing high variance values to those cells where none or little information is available, together with the fact that our mean estimation does not "average", but combine observations attending to their "age".

5.5.3 Real experiment

Following the same setup than in the simulated experiment, we show next the results of a real experience. In this case the gas observations are collected with a photo ionization detector (PID[2]), mounted on the gas sensitive robot *Rhodon* (see Appendix A). To generate the gas leak in the left room, we place an ethanol bottle in front of a fan to boost the gas dispersion. The ethanol bottle was timely opened or closed to match the setup described in Section 5.5.1. This configuration generates a gas plume towards the door heading to the corridor when the ethanol bottle is open, while it helps cleaning the room from gases when the bottle is closed.

For a fair comparison among the GDM methods, the experimental data is collected and saved to a log file for off-line processing. In this way, differences from

[2]Model *ppbRAE2000* from RAESystem, with a 10.6 eV UV lamp.

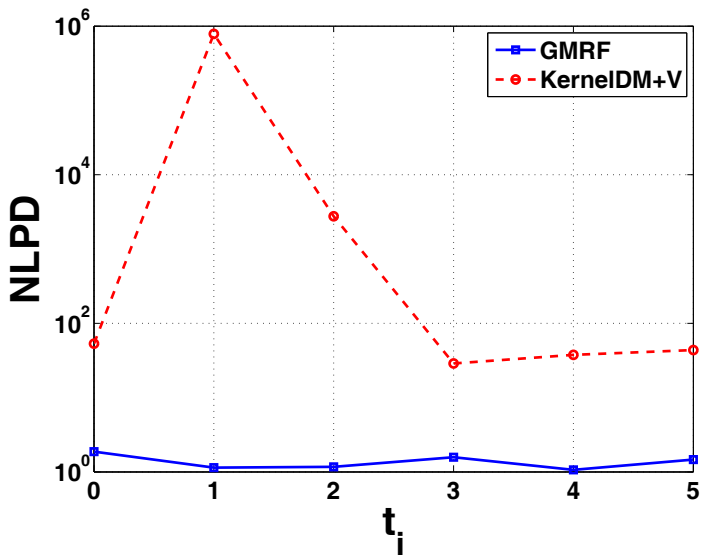

Figure 5.7: Negative log probability density (NLPD) at the different time steps (t_i) of the simulated experiment, for the GMRF and Kernel DM+V methods.

both estimations will be only attributable to differences in the methods and not to the data.

Figure 5.8 shows the mean and variance maps generated by the GMRF and the Kernel DM+V. As in the simulation experiment, snapshots of the maps at the different time steps (t_i) are depicted. Since the ground truth of the gas distribution is not available, we plot instead the obstacles map together with the robot localization.

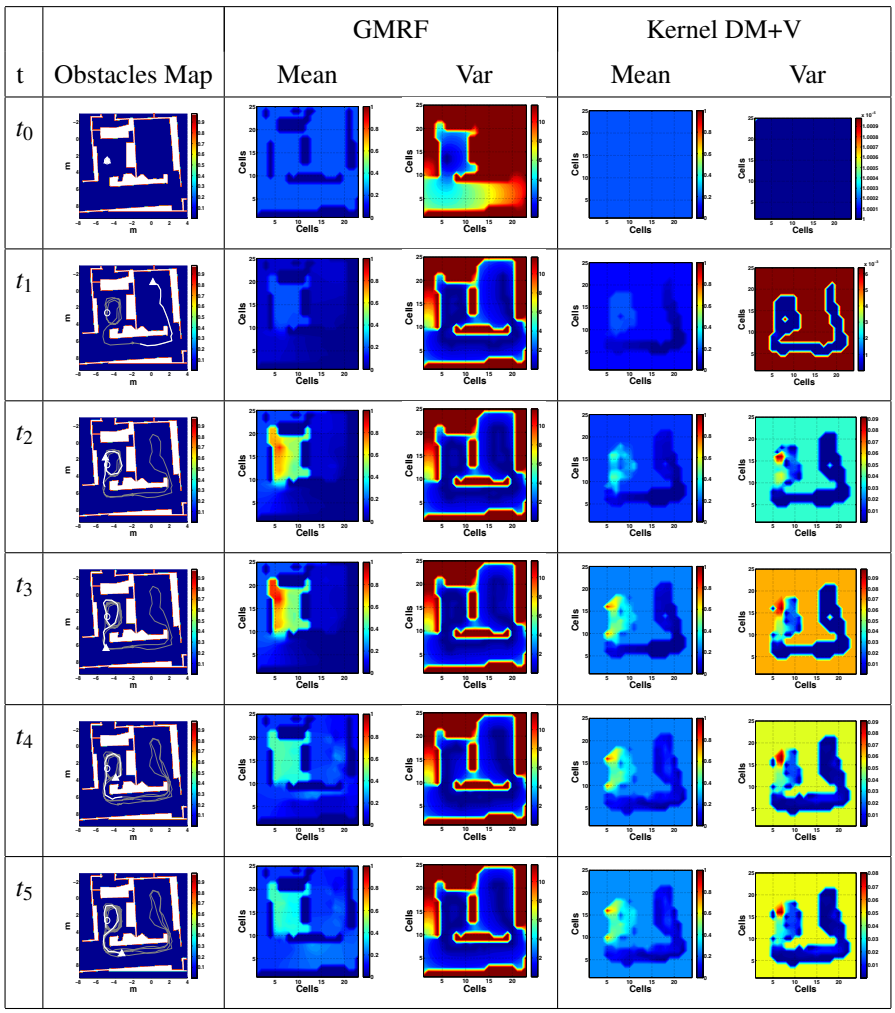

Figure 5.8: Mean and variance gas distribution maps estimated by two different methods (GMRF and Kernel DM+V) at the different time steps t_i of the real experiment.

As can be noticed, the results are quite similar to those obtained in simulation, which corroborates the already highlighted advantages of our proposal. The main difference resides in the presence of low concentration gases in the left room at time-steps t_4 and t_5, after the ethanol bottle is closed. This can be attributed to the fact that the time elapsed since the ethanol bottle was closed till the robot inspected back the room, was not enough to allow the gas to be completely "*removed*" from the environment.

5.6 Conclusions and outlook

In this chapter we have revised the problem of creating a map of the gas distribution and proposed a new approach that accounts for two important issues that had been ignored in previous works: the validity of the observations over time and the presence of obstacles in the environment. The former is achieved by introducing a time-decreasing weighting factor to each gas observation. Thus, when combining observations taken at close locations, their longevity is used to determine their respective relevances, avoiding the detrimental "average" effect of previous approaches. On the other hand, obstacles such as walls or furniture are now taken into account by modeling the correlation between the map cells they separate. We have addressed the problem in a probabilistic manner modeling the problem as a MAP estimation over a Gaussian Markov Random-Field (GMRF).

Our approach has been validated with simulated and real experiments, providing qualitative and quantitative comparison with classical methods and demonstrating the advantages in the estimation of the gas distribution.

Future work includes the consideration of additional sources of information such as vision or semantics to improve the way cell correlations are modeled, exploiting properties of the obstacles that influence the gas dispersion (shape, height, etc). Furthermore, efficient alternatives to obtain the GDM of non rectangular scenarios will be studied, for example, by only estimating the gas concentration at desired areas instead of the complete rectangular lattice.

Chapter 6
Conclusions

This thesis has addressed the problem of gas sensing with mobile robots, paying particular attention to the detection of gases with an array of metal oxide semiconductor (MOX) gas sensors, and to the tasks of gas distribution mapping in natural environments.

In robotics applications gas sensors are usually deployed in the so called open sampling systems (OSS). Here, sensors are directly exposed to the environment, which enables gathering valuable information about the dynamics of the interaction between the sensors and the gases to be analyzed. However, this also entails that gas readings are strongly influenced by the dispersal mechanisms of gases, especially in natural, real environments where turbulent advection produces the distribution of gases to be patchy and chaotic. Additionally, the limitations of the current state of the gas sensing technology make even more difficult the development of olfactory robots able to autonomously accomplish tasks like gas source localization or gas distribution mapping. This is specially important for MOX sensors due to their lack of selectivity, slow recovery and dependence of the environmental conditions. Two approaches to palliate one of their main drawbacks, the long recovery period, were proposed in Chapter 3. This disadvantage limits their suitability to applications where the sensors are exposed to rapid changes of the gas concentration. The long duration of the acquisition cycles (up to tens of seconds) is of especial concern for mobile robotics, since inaccurate readings are inevitable when measuring the gas concentration. It is particularly noticeable in the decaying phases, when the output of the sensor recovers to the baseline level. Consequently, steady state values are rarely reached, and therefore gas sensing based on MOX technology must deal with the transient information of the signals.

The first proposal to overcome the slow recovery of MOX sensors consisted of the introduction of a new electronic nose, the multi chamber electronic nose (MCE-nose). The key of the proposed design is the accommodation of a set of redundant sensors in different chambers, being alternatively activated to ignore the inaccurate

sensor output when a decay phase is detected. The output signal then results from the concatenation of the rise phases of a sequence of MOX sensors. Increasing the number of chambers also increase the possibility of having an array of sensors at the desired baseline level and, consequently, the possibility of sensing faster changes in the gas concentration. However, a higher cost, power consumption and complexity are also a direct consequence of having more chambers, thus the optimal number strongly depends on the characteristics of the target application. A prototype of the MCE-nose was built and integrated in a mobile robotic platform under the OpenMORA robotic architecture, being tested in real experiments to measure the improvement in the sensing of rapid changing gas concentrations. Through these experiments it was found that although small, the differences on the replicated sensors hosted in the different chambers, compromised the correct chamber-switching functionality of the MCE-nose. These differences, which are mostly due to the fabrication process, but also ageing or poisoning, require therefore a pre-calibration of the system to achieve proper results.

In contrast to the MCE-nose, the second approach suggested in Section 3.4 compensates the slow dynamic behavior of MOX sensors by forecasting the steady state values of the sensor resistance from a sequence of transient measurements. In general, MOX sensor models proposed in the literature seek to predict the sensor output (resistance) when exposed to a certain gas concentration profile. The suggested approach is inspired on reversing such a model: given a sequence of measurements from the transient response of the MOX sensor, the concentration profile of the exciting gas is forecasted by estimating the steady state values of the sensor resistance. The exploited model is based on two first order systems (rise and recovery) with time constants that depend on the sensor reading amplitude. Different experimental scenarios were presented for the validation of the suggested model-based approach. These experiments demonstrated how the "long tail" effect of the recovery phases could be effectively avoided, something that prevents the overlapping between the rise and recovery phases of the MOX sensor. On this basis, the profile of the gas concentration as estimated by the model, provides values more consistent with the reality than when using the raw sensor signal. Additionally, it was proved that both, the MCE-nose and the inverse MOX model, entail a considerable increase in the speed at which a mobile base carrying a MOX-based e-nose can inspect the environment. The latter directly implies an important reduction in the execution times of olfaction related tasks.

In mobile robotics, many of the olfaction related tasks require determining the gas concentration the sensors are being exposed to. This is a crucial step for realistic gas sensing applications since legal requirements and regulations are expressed in terms of absolute gas concentration, toxicity levels, etc. Additionally, and due to the chaotic nature that dominates gas dispersal in OSS, in most cases it is desirable to provide, together with an estimate of the mean concentration, an estimate of the uncertainty of the prediction. As a solution to this quantification problem, a supervised machine learning approach based on Gaussian Processes (GPs) was presented in Chapter 4. This issue has been addressed in a probabilistic manner, estimating a posterior distribution over the gas concentration given the response from an array of

MOX gas sensors. This has the advantage of enabling not only predictions of the expected gas concentration but also predictions of the uncertainty of this estimate. Different configurations of the array of gas sensors were studied, and automatic relevance determination (ARD) was considered to exclude those sensors not relevant for the estimation of the posterior gas distribution. Finally, two suggestions to account for the dynamics of the MOX response within the quantification problem were analyzed. Interestingly enough, Experiments showed that neither using additional samples from previous time steps, nor adding the signal derivatives, produced any significant improvement.

From a mobile robotics perspective, GDM addresses the problem of estimating the spatial distribution of volatile substances using a mobile robot equipped with an electronic nose (e-nose). The proposed algorithm described in Chapter 5 presents the novelty of accounting for two important issues that have been ignored in previous works: the validity of a gas observation as time goes by (ageing) and the presence of obstacles in the environment. Odors are ephemeral, thus, we can say that the information carried by a given gas measurement quickly vanishes as time goes by. A time-decreasing weighting factor was then suggested as a measure of the "age" of each gas measurement, modeling the fact that recent measurements more significantly represent the current gas distribution than older ones. On the other hand, obstacles such as walls or furniture are also taken into account in the GDM process by modeling the correlation between the map cells they separate. These two considerations produce maps of the gas distribution which are more compliant with the actual mechanisms of gas dispersion. The problem was addressed in a probabilistic manner, modeling it as an efficient MAP estimation over a Gaussian Markov Random-Field (GMRF). Results of time-variant GDM experiments were presented, providing a detailed comparison with existing mapping methods in both, simulated and real scenarios. Among the important advantages introduced with this new method, highlight the fact that the "averaging" effect between recent observations and older ones is avoided, which presents a faster adaptation of the map to changing gas concentrations.

Apart from the specific lines for future work included in the different chapters of this thesis, we envisage as a long-term path the integration of the gas sensing advances and algorithms presented in this thesis within the area of home-oriented robotics applications. Service robots are a potential market niche that is nowadays blooming because of the continuous advances in the field (SLAM, navigation, semantics, etc) and the greater availability of commercial robots. However, since artificial olfaction is not mature enough, most commercial robots do not yet incorporate gas sensing devices. In spite of this, we believe that olfaction can play a key role in making service robots more intelligent and valuable, by combining it with other sensing modalities like vision, range sensing or touch.

Appendices

Appendix A
Rhodon: a gas sensitive robot

This appendix is devoted to *Rhodon*, the gas-sensitive robot that has been used in most of the experiments presented in this thesis. *Rhodon* has been built with the funding supplied by the Andalucía Government and the European Regional Development Fund (ERDF) under project TEP-2008-4016. The author, among the rest of the MAPIR group, has devoted great efforts to the development of this robot, including both software and hardware implementation tasks.

Rhodon is a mobile robot intended to perform olfaction tasks. It has been constructed upon a Patrolbot commercial platform [106], on which a number of sensors and devices have been integrated. Additionally, a structure has been mounted on the platform in order to support the sensory system, which comprises, among other devices, two radial laser scanners (front and back), a ring of sonar sensors and a Kinect sensor. However, during the course of this thesis, both its structure and components have varied significatively to cope with the requirements of the new experiments and to adapt to the advances in the sensing technology. Next, we detail the main characteristics of the robot along its different stages.

A.1 First Stage: Experimental mobile platform

At the beginning of this thesis, Rhodon was composed only by the Patrolbot mobile platform on which a laptop, a SICK LMS200 and a custom designed e-nose were mounted. Due to the simplicity of the sensory system, no structure was necessary as shown in Figure A.1.

The SICK laser range scanner (embedded in an internal compartment designed to accommodate such sensor) was only intended to provide precise detection of obstacles, since localization was still not achieved at this early stage. The robot olfaction capability was provided by an e-nose which design was based in a mono-PCB with a central circular opening. The array of gas sensors composed by 6 TGS (Figaro) and 2 MiCS (e2V) gas sensors, was distributed around this aperture. A small fan placed over

(a) (b) (c)

Figure A.1: Images of Rhodon at its initial stage of development. Only a laptop and a custom designed e-nose were then mounted on the Patrolbot mobile platform.

the circular opening generated then a constant airflow through the array of sensors, allowing the sensing of different volatile substances. This simple setup was employed on the first olfaction experiments of this thesis, those related to improve the gas sensing process with a mobile robot as detailed in Chapter 3.

The onboard e-nose was then substituted by an improved sensor: the MCE-nose (see Section 3.3). Figure A.2 shows different snapshots of experiments carried out with this robotic setup. Due to the simplicity of the experiments at this stage where neither precise localization nor high computation was required, no additional sensors or electronic devices where integrated. Even the sonar ring (embedded in the Patrolbot platform) was not employed at this stage.

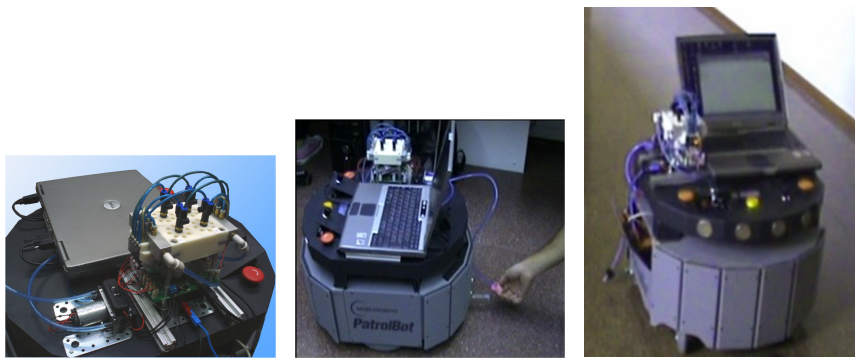

(a) (b) (c)

Figure A.2: Images of Rhodon during different experiments carrying the novel MCE-nose.

A.2 Second Stage: Autonomous mobile olfaction robot

Raised by the the need of computational power for executing the more and more sophisticated and complex algorithms, as well as to provide autonomy to the robot with tasks such as navigation, obstacle avoidance, path planning, etc., Rhodon was rebuilt completely.

At this stage, a structure to organize the different devices, sensors and electronics was built. To improve the range sensing capabilities for tasks as localization and mapping, another laser scanner (Hokuyo UTM-30LX) was integrated at the rear of Rhodon, providing together with the already mounted SICK a 360° field of vision. The laptop was then replaced by a compact and powerful computer to cope with the requirements of high computational capacity. Additionally, a touch screen was placed at the top of the structure to improve the interaction with the user. Due to the consequent increase in power consumption, a set of commercial $12V$ batteries was mounted to avoid an excessive discharge of the mobile platform batteries, improving the robot autonomy time. Finally, a panel of switches was placed at one side of Rhodon, allowing the control of which components are turned-on or off. Figure A.3 shows two pictures of the resulting setup, pinpointing the onboard components.

Naturally, a software architecture was then a necessary step towards an autonomous gas sensitive robot. The open mobile robot architecture (OpenMORA) [41] was then introduced. It is a MOOS [114] and MRPT [108] based distributed architecture for

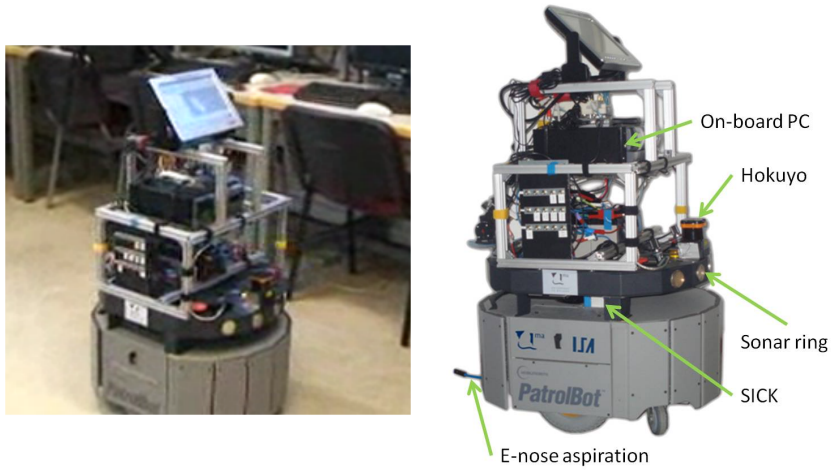

Figure A.3: (a) Picture of Rhodon during an olfaction experiment.(b) Description of the different elements onboard.

mobile robots, providing off-the-shelf modules for common robotics platforms and sensors, MonteCarlo localization, reactive navigation, simulation, etc. Despite its extensive variety of robotic modules, further development was necessary to incorporate the desired olfaction functionality, by means of different modules and the extension of the simulator as described in Appendix B.

A.3 Third Stage: Multi-purpose gas sensitive robot

A gas sensitive robot as the described above is of only use for carrying out specific olfaction-related tasks exclusively (see Section 2.4). That is, even if the main characteristic of a gas sensitive robot is to be able to detect and analyze volatile substances, it is desired for such a robot to be able to interact with the environment based on the sensed information. Additionally, for high level olfaction applications it is almost necessary to attend to other sensing disciplines rather than only olfaction, thus, the collaboration between sensing modalities should be allowed by the robot.

With this goal, the structure and components of Rhodon were once more customized. The main modification was the incorporation of a robotic arm, a research edition of the commercial JACO arm from Kinova [125]. This robotic arm is ideal to be used on a mobile robot, meeting the constraints of being lightweight and having a low power consumption. To increase the area of operation without interfering with other parts of the robot, the arm was placed at the top-center of the structure. To provide weight-compensation when the arm is fully extended and prevent the swaying of the robot, a new set of batteries were placed at its rear. Figure A.4 depicts the location of the different components onboard Rhodon.

A metal bar was then attached to the structure of the robot to allow vision cameras to obtain images from a wider perspective. A Kinect sensor and a thermal camera

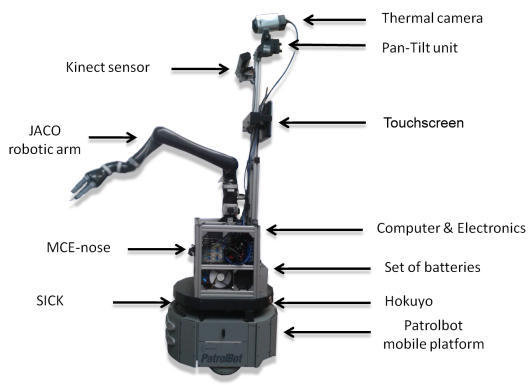

Figure A.4: Detailed view of the sensors and components of Rhodon at its current stage.

(a) (b)

Figure A.5: Images of Rhodon at its current stage. (a) Picture of an experiment where Rhodon was commanded to prepare a cocktail. (b) Detail of the attached e-nose aspiration to the robotic hand, and grasping of the plastic cup after the cocktail has been prepared.

were then positioned over the robotic arm, allowing the use of vision and range sensing for manipulating objets among other robotic tasks. Figure A.5 shows different pictures of Rhodon from an experiment where it was commanded to prepare a cocktail. In this experiment, almost all the capabilities of Rhodon were tested, ranging from localization and path planning, to olfaction and object manipulation. A video describing this experiment can be found at the Youtube channel of the MAPIR group at `http://mapir.isa.uma.es/` or directly at Youtube: `http://youtu.be/BsQMewX8yNQ`

Appendix B
An olfaction plugin for simulating mobile robotics olfaction

A noticeable characteristic of most mobile robotics olfaction experiments is their complex and time consuming development life-cycle. The main reason behind it lies in the impossibility of deriving a *ground truth* (GT) representation of the gas distribution in the environment, making difficult, almost impossible, to validate new algorithms or to compare different proposals aiming at the same objective. For this reason, real experiments usually employ complex setups with the intention of controlling, as much as possible, the dispersion of volatiles (generation of plumes with fans, shutting doors and windows to reduce undesired airflows, etc.), but even then, the results cannot be completely validated because of the lack of information about the real state of the gas distribution.

As in many other scientific fields, this drawback can be considerably reduced by making use of simulation environments. Among their many advantages, we can highlight:

- They provide a ground truth of the simulated environment, a fundamental requisite for validating new algorithms.
- Repeatability is easily achieved and so the comparison of different algorithms.
- Simulation allows testing an algorithm under different and complex environmental conditions (e.g. laminar flow, turbulent flow, plumes, etc.), and can be even more realistic than traditional experiments, as they allow the free configuration of the environmental parameters found in the operational application.
- Using simulated data is generally cheaper, safer (e.g. toxic gases) and less time-consuming than conducting laboratory experiments.

However, due to the still low presence of artificial olfaction in the area of mobile robotics, most currently-existing robotic simulators do not account yet for olfaction methods and models. In this appendix, an olfaction plugin developed to provide the necessary mechanisms for efficiently testing and validating algorithms related, but not limited to, gas distribution mapping (GDM), is introduced. This plugin, used in some of the olfaction experiments presented in this thesis, is designed to work with the robotic simulator included into the Open Mobile Robotics Architecture (OpenMORA) [52].

The organization of this appendix is as follows: after a review of the most spread mobile robotics simulators, an overview of the chosen OpenMORA's simulation environment is provided in Section B.2. Then, the proposed olfaction plugin is introduced in Section B.3, to finish with an illustrative example in Section B.4.

B.1 Related research

In robotics research, simulation plays a significant role as a tool for quickly and efficiently testing new concepts, strategies, and algorithms. Its importance can be appreciated in the number of simulation toolkits that have been proposed, including: Webots [105] a commercial development environment used to model, program and simulate mobile robots, with a large choice of simulated sensors and actuators, the Virtual robot experimentation platform (v-rep) [35], a distributed and modular robotic framework that concurrently simulates control, actuation, sensing and monitoring, or CARMEN [17] an open-source collection of software for mobile robot control from the university of Carnegie Mellon, which implements a moderately useful 2-D robotic simulator.

When simulating multiple mobile robots, Stage [46] is possibly the most commonly used 2D swarm robot simulator for research and university teaching nowadays, while Gazebo [75] extends the framework to 3D dynamic multi-robot environment, providing tools for recreating complex world scenarios.

Related to robotics olfaction, Cabrita et al. [16] proposed a Player/Stage based simulator for mobile robot olfactory experiments, focusing mainly on extending its functionality to support chemical plume simulation and plume tracking. This framework lacks, nevertheless, of a realistic model for the gas sensors and the inclusion of gas distribution mapping algorithms. In [93], the simulation of a single robot in a dynamic time-variable environment generated with the computational fluid dynamics software Fluent [4] is proposed to study the gradient and equilateral triangle search methods. However, this software package only considers a simple mobile object equipped with a gas sensor, thus it does not provide the variety of tools necessary to simulate a mobile robot (collision avoidance, path planning, etc.). Finally, a more complete olfaction simulator for mobile robotics is presented as a plugin for Webots [105]. Although it includes the simulation of odor dispersion based on the filament-based atmospheric dispersion model proposed by Farrell et al. [28], it still lacks a model of the gas sensors and the implementation of gas mapping methods.

B.2 Description of the OpenMORA's simulation environment

The robotic simulation environment used along this thesis is a C++ based-application consisting of a number of modules built using the Mobile Robot Programming Toolkit (MRPT) [108] and integrated into the Open Mobile Robotics Architecture (OpenMORA) [52]. The set of libraries which this simulation environment is based on, has been developed to keep flexibility, modularity and reusability as much as possible [32, 33]. These libraries are continuously being updated as a collaborative open project, maintained by researchers and developers worldwide within the area of mobile robotics[1]. In this section, the main components of this robotic simulator, corresponding to the grey shaded blocks in Figure B.1, are outlined.

B.2.1 Geometric map

Although the robotic simulator included in the OpenMORA architecture is a full 3D simulation environment, the specification of the scenario where the robot will be simulated in (dimensions, walls, obstacles, doors, free space, etc.) is only two dimensional. The simulator is configured to receive an external geometric map file. This map file can be either a simple *png* format image, or a *simplemap*[2] format file, sometimes called "view-based map", which is composed of a set of poses and their

[1]For a detailed list of collaborators, please visit www.mrpt.org/Authors and sourceforge.net/p/openmora/_members

[2]For detailed information about robotics file formats visit: www.mrpt.org/Robotics_file_formats

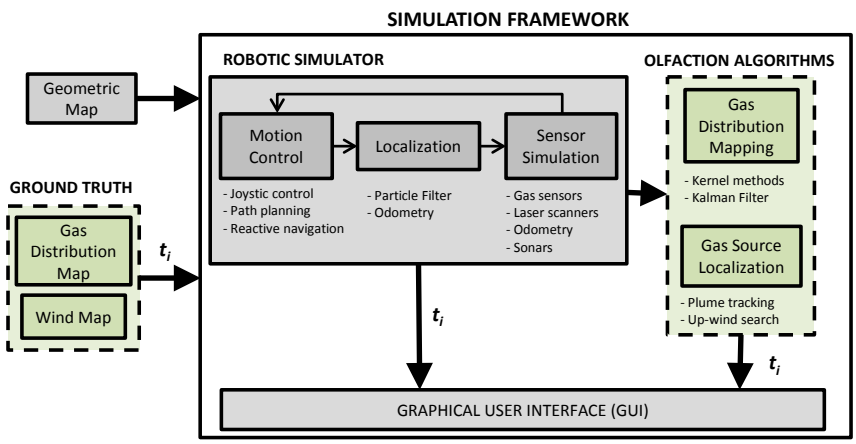

Figure B.1: Block diagram of the OpenMORA's simulation framework. The backbone robotic simulator (grey shaded/solid line blocks) is complemented with an olfaction plugin (green shaded/dashed line blocks) enabling the simulation of olfaction robotic tasks.

associated observations (generally laser scans). These *simplemap* files can be easily obtained by a generic ICP-based algorithm. Figure B.2 displays an example of both alternatives. Additionally, there exists also the possibility for arranging the map in a metric-topological structure as described in [8].

B.2.2 Robot simulation block

This block concerns three modules related to the simulation of the robot: the motion control, the localization of the robot in the environment, and the simulation of the sensors onboard.

Motion control: This module is in charge of simulating the movements of the robot. It implements different control alternatives:

1. Joystick mode, where the simulated robot moves following the commands provided by a user controlled joystick/keyboard.
2. Path planning with reactive navigation. In this mode, the robot automatically plans the best route to follow according to the distance to a user defined destination, while avoiding obstacles encountered along the path. The reactive navigation employs space transformations (PTGs) to drive the simulated robot using an internal simpler holonomic algorithm [10].
3. Predefined paths. In this mode the path is given as a sequence of nodes for the robot to go through and, as before, a reactive navigator takes care

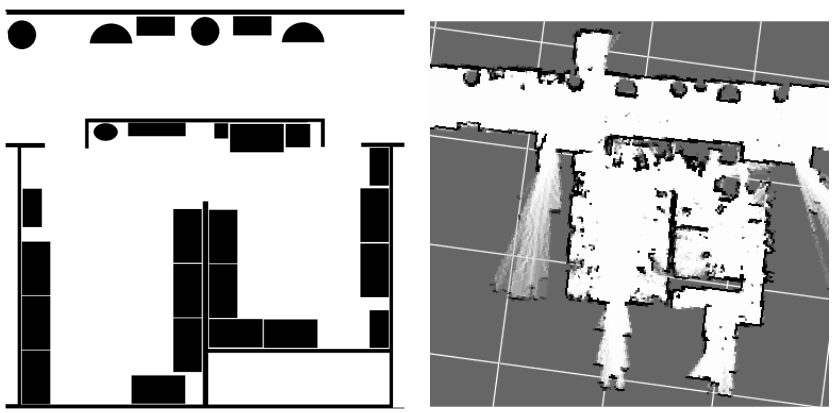

(a) (b)

Figure B.2: Two examples of geometric map input files: (a) a *png* image, and (b) a *simplemap* format file (created from odometry and laser scan observations of a real robot experiment). For both file formats, white colors represent free space while black occupied zones.

of the obstacle avoidance. This mode allows defining paths for inspecting complex scenarios, as well as for reproducing the same paths in the experiments.

Localization: The localization of the robot is handled by an efficient particle filter implementation [11], upon the information provided by the simulated robot odometry and laser scans, and a known map (see Section B.2.1). This method performs pose tracking of the robot using an adaptive number of particles.

Sensing: This module simulates the responses of different sensory systems that can be configured on the simulated robot. Implemented sensors include:

- Range scan lasers, as the SICK [140] or Hokuyo [58], which are useful for the correct localization of the robot, a fundamental requisite for the generation of GDM.
- The robot odometry, which is simulated by providing linear and angular speeds.
- Sonar sensors, which can be used to detect obstacles at low range.

Gas sensors as the MOX or PID (see Section 2.2) are also considered. Nonetheless, a detailed description of their implementation is left for Section B.3, where the specific modules related to olfaction are described.

B.2.3 Graphical user interface

To visualize and manage all the simulated information, a powerful 3D graphical user interface (GUI) is integrated in the environment. Figure B.3 depicts its main components:

1. The system bar, which hosts the navigation buttons related to the localization and motion control of the robot (e.g. re-localize, stop, move to, etc), as well as with some visualization options.
2. The navigation planning, which details the nodes visited by the robot along the planned path (see Section B.2.2), as well as the coordinates of the target node where the robot is heading to.
3. The panel view. It is composed of a checkbox list that allows the customization of the 3D view by selecting which sensors and layers are to be displayed (e.g. the robot shape, position of the particles used for localization, laser beams, mean and variance gas distribution maps, wind, etc).
4. The 3D view. An OpenGL based viewer to visualize the simulation.
5. The system messages panel. It displays valuable information about the simulation status (exceptions, warnings, flags, etc) of the different modules being executed.

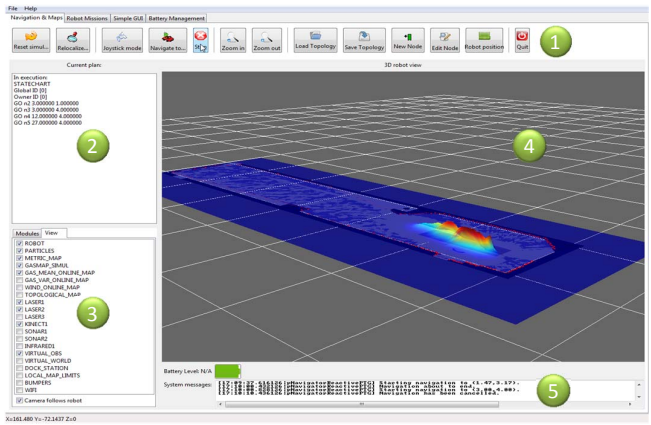

1. System bar
2. Current navigation planning
3. Panel view
4. 3D view
5. System messages

Figure B.3: Snapshot of the graphical user interface.

B.3 The proposed olfaction plugin

The olfaction plugin, which has been developed as part of this thesis, comprises three important parts as depicted in Figure B.1: a sensor model for MOX gas sensors, gas and wind ground truth maps, and several GDM algorithms.

B.3.1 Model of MOX gas sensors

A fundamental requirement for simulating olfaction tasks with mobile robots is to accurately describe the response of the gas sensors being simulated. This is even more important for the case of MOX gas sensors, since, as explained in Section 3.2, their long recovery times has a strong influence in the task performance. In the proposed plugin, the MOX sensor is modeled as low pass filters with different time constants for the rise and decay phases of the response (see Section 3.4 for a detailed description). Additionally, an additive Gaussian noise is considered to account for both the intrinsic sensor noise and the one in the measurement system.

The plugin allows the simulation of multiple MOX sensors simultaneously (array of sensors). Each simulated sensor is defined by selecting the target gas it is sensitive to and the characterization of its parameters: time constants and noise. This enables the simulation of complex tasks which require to detect multiple gases.

B.3.2 Ground-truth of gas concentration maps

As mentioned in above, having a ground truth (GT) of the gas distribution at each time step is fundamental for validating and comparing GDM approaches. The plugin provides two different types of ground truth maps (see Figure B.4):

- GT wind maps: By default, the plugin is only able to simulate constant wind fields, that is, homogeneous laminar flows. Although constant flows are not a very realistic model, they are very simple to configure and straightforward for comparison of different GDM approaches. To account for more realistic wind fields, including turbulence and obstacles, future development will consider the loading of computational fluid dynamic (CFD) simulations.

- GT gas maps: Given the complexity of the gas dispersal mechanisms and the many different configurations for the simulation (turbulence strength, single or multiple gas sources, source positions, obstacles, etc), the proposed plugin does not generate the GT gas maps itself, but rather it allows loading them from third party software. One interesting property is the fact that it can handle multiple GT gas maps simultaneously. This enables, for example, to account for the simulation of multiple gases in the environment (e.g. multiple and different gas sources). Additionally, GT gas maps can be static (e.g. being obtained from a simple grayscale image), or dynamic, being loaded from a sequence of successive gas maps at a configurable rate.

B.3.3 Gas distribution mapping algorithms

The plugin is specifically designed for developing and testing GDM algorithms, thus, it implements the most important alternatives proposed in literature.

- Kernel methods: Two different kernel methods for GDM are included. The first one, described in [85], convolves the sensor readings using a radially symmetric two-dimensional Gaussian functions to estimate the gas distribution, while

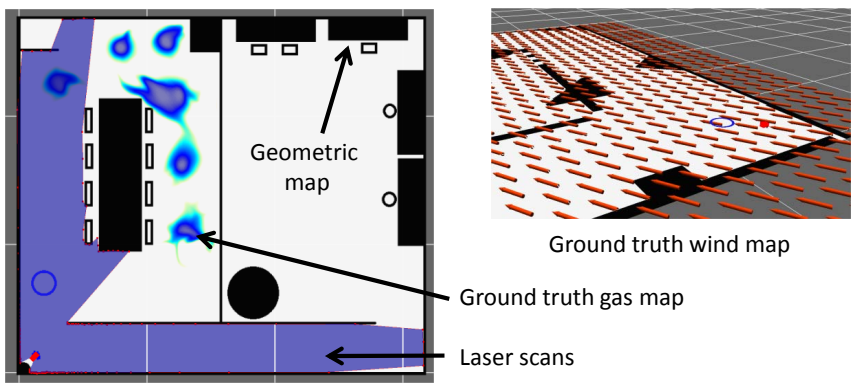

Figure B.4: Snapshot of a gas ground truth map (left) as displayed by the GUI, together with the geometric map and the simulated laser scans. (right) Sample of a wind ground truth map, displaying the vector grid that indicates the wind direction and strength (color scaled).

a later version called kernel DM+V [89], aims at probabilistic GDM by incorporating the variance in the estimation.

- Kalman filter based method: This approach to GDM is based on sequential Bayesian estimation given a lattice of 2D cells treated as hidden variables using a Kalman filter [9]. Two implementations of the algorithm are available: a standard version which employs the whole covariance matrix in the computation of the posterior, characteristic of Kalman filters, and an efficient and optimized version for real time applications. The implementation and advantages of the latter are described in detail in the cited work.

- Gaussian Markov Random-Field method: This is the novel approach to GDM presented in Chapter 5 which accounts for the aging of sensor observations as well as for the obstacles in the map.

B.4 An illustrative example

To illustrate the usefulness and potential of the proposed olfaction plugin, apart from the results presented in Section 5.5.2, in this section we estimate the GDM of a simulated environment, providing a detailed description for the different components, as well as the resulting maps for two different GDM methods.

The simulated scenario consists of two adjacent rooms that the robot inspects to check for dangerous levels of contaminant gases. The robot path is defined in order to explore both rooms sequentially, as depicted in Figure B.5(a). For this experiment, and considering that the tested GDM methods do not account for wind information, we set a constant wind field of zero strength, while the gas ground truth map (see Figure B.5(b)) is set to simulate a gas leak in the leftmost room, by generating different gas patches dispersed around the middle table.

Figure B.5(c-f) show the estimated maps using the kernel DM+V and the Kalman based method. The results can be analyzed in the GUI, or they can be saved to a file for a posterior analysis (e.g. using Matlab). A video of a similar simulation experiment can be found at: youtu.be/UKhtIem7u0g

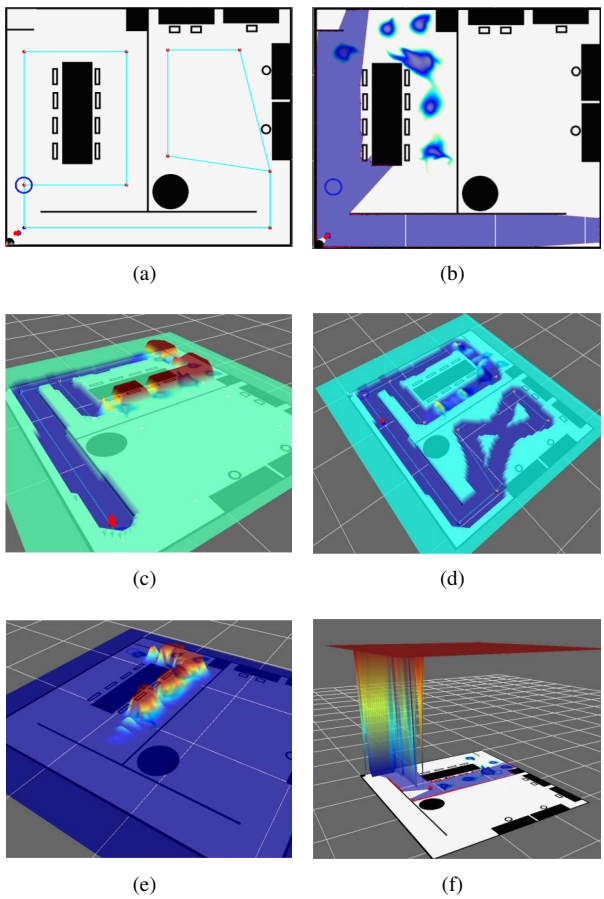

Figure B.5: Snapshots of the illustrative example: (a) Metric map with the predefined inspection path (red circles represent nodes to visit, while green lines represent allowed paths). (b) Gas ground truth map and detailed view of the laser beams emitted by the simulated SICK laser. (c-d) 3D snapshots of the mean and variance maps for the Kernel DM+V method, and (e-f) for the Kalman based method.

Appendix C
Data-Sets for odor classification

Odor classification by a robot equipped with an electronic nose (e-nose) is a challenging task for pattern recognition since volatiles have to be classified quickly and reliably even in the case of short measurement sequences, gathered under operation in the field. The discrimination of gases performed with this setup presents a number of additional challenges when compared to standard analyte identification applications, mostly due to the differences in the measurement conditions. While standard classification tasks usually host gas sensors inside a chamber with controlled humidity, temperature and airflow conditions; in robotics olfaction there is no control over the sensing conditions. This entails that the sensor signals to be processed are noisy and dominated by the signal transient behavior, as described in Section 2.3.

Although classification is not the main focus of this thesis, in this chapter different classification data-sets gathered with an array of MOX gas sensors in an open sampling system are presented. These data-sets, available online at `http://mrpt.org/Robotics_olfaction_dataset`, are the result of a collaboration with the University of Bielefeld (Germany) and the University of Groningen (The Netherlands), which aims at classifying different volatiles using only short data sequences of high dimensionality. Nevertheless, by the time this thesis was written no published results were available, thus only the data-sets gathered for the experiments are presented next.

C.1 Data-set 1: Controlled gas pulses

The first data-set is composed by samples gathered under restrained environmental conditions. The objective is to collect a set of e-nose signals from different volatile substances under very similar conditions. It can then be used for training and validation at the initial developing stages of new classification algorithms.

The odor samples are recorded from an e-nose containing an array of six different MOX gas sensors, whose readings are recorded to obtain an odor fingerprint of

the odor. These sensors are Figaro TGS-2600, TGS-2602, TGS-2611 and TGS-2620, and two e2v sensors: MICS-5135 and MICS-5521. Those sensors are hosted within a small-volume chamber where a continuous airflow is injected through a pneumatic circuit powered by a pump. The aspiration of that circuit is done through a thin tube, allowing an easy handling in the smelling process described in Section C.1.1. Figure C.1 shows a snapshot of this setup.

This data-set is comprised by 69 odor samples gathered by exposing the e-nose aspiration to gas pulses of seven different analytes (see Figure C.2): four of them are commercial alcoholic beverages (Negrita Rum, Larios and Gordon Gin, and Cointreau), and the other three are commercial polish remover based on acetone, standard ethanol and lighter's gas (usually butane, but sometimes mixed with propane). Acetone is given by 9 samples and the other classes by 10 samples each.

C.1.1 Setup

Odor samples have been collected following a fixed and systematic protocol for all the samples and test runs. In all cases, the bottle containing the target odor (in liquid or gas state) was used directly to smell from it. Each sample is collected according to the following three-phases procedure:

Seconds (0-30): For the initial 30s, the baseline level of each sensor in the array is estimated by measuring the sensor response in absence of the target gas. This is achieved by leaving the target volatile container closed and separated from the

Figure C.1: Snapshot during the collection of the first presented data-set. It can be seen the e-nose aspiration (blue thin tube) and containers of different target odors.

Negrita Rum | Larios Gin | Gordon's Gin | Cointreau | Polish remover (Acetone) | Ethanol | Butane-Propane

Figure C.2: The seven odor volatiles used on the first classification data-set.

e-nose aspiration. After the first 20s, the container is opened and left unattended for another 10 seconds, allowing the stabilization of the gas dispersion rate and avoiding measuring a burst of concentration due to the time the container was closed.

Seconds (30-90): For a duration of 60s the e-nose aspiration is placed next to the gas source (at a distance of 10cm approximately) exposing the sensor array to the target volatile.

Seconds (90-): Finally, the e-nose aspiration is taken away from the gas source allowing the sensor array to recover to its initial state (baseline). The volatile container is then closed to avoid an excessive contamination of the testing room. Due to the long recovery time of MOX sensors (see Chapter 3), the e-nose was left to recover its baseline level for almost 10 min. before starting a new run.

Figure C.3 shows two different samples of such data-set. Notice that although the gas exposure was "controlled" by time exposure and distance to the source, strong fluctuations in the sensor readings occur due to the chaotic nature of gas dispersal.

C.1.2 Data-set structure

Each target analyte is represented by a folder containing all the samples in a plain text file format ("*.txt"). Each of these txt files (recorded using a sampling frequency of approximately 7Hz) is organized in columns as follows:

Column 1 Time

Column 2 Temperature (Not available)

Column 3 MICS 5521 readings

Column 4 MICS 5135 readings

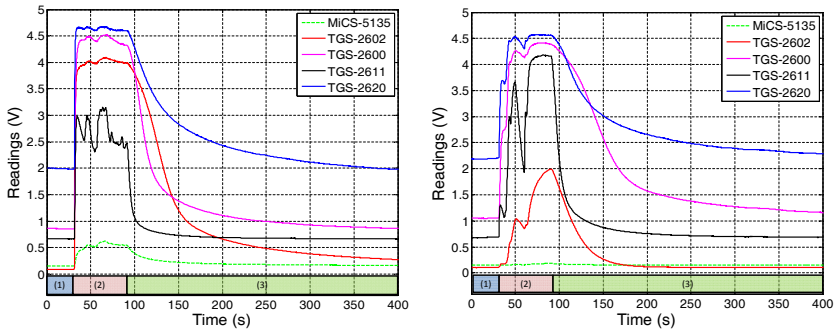

(a) (b)

Figure C.3: Two different samples of the first olfaction data-set. The three phases in which the samples can be decomposed are marked at the bottom of each figure as (1),(2) and (3). Readings of sensor MiCS-5521 have been omitted here because of the lack of sensitivity to the target gases.

Column 5 TGS 2602 readings

Column 6 TGS 2600 readings

Column 7 TGS 2611 readings

Column 8 TGS 2620 readings

Column 9 and following Not used

Additionally to the available data-set, a Matlab script file is provided to easily display the information graphically. This script asks for a *dataset.txt* file, and plots the contents of it according to the previous organization. This data-set can be found online at `http://mrpt.org/Robotics_olfaction_dataset`.

C.2 Data-set 2: Classification with different sensors

The objective of this second data-set is to provide a collection of olfaction samples that may be used to determine the strength of a classification algorithm under more realistic robotic scenarios, where environmental conditions cannot be controlled. More specifically, this data-set aims at studying the impact on the classification performance of different effects that disturb the behavior of gas sensors in real scenarios:

- **Ageing**: It is well known that most gas sensor technologies suffer from ageing, that is, the change of their chemical properties (size of the response to a certain amount of a given gas, selectivity, response speed, etc.) by the simple pass of time [43]. These changes give rise to *drift* in the sensor response.

- **Poisoning**: For real environments where a lot of different gases are present (several of them in very small amounts), little can be said beforehand regarding the chemical reactions and their reversibility. Therefore, some reactions will be irreversible, thereby blocking or creating reaction sites on the sensor surface and/or bulk of the sensing material. This will lead to a change in the sensitivity of the sensor towards other gases. [119]

- **Sensor replacement**: In the breakage of a sensor, it is replaced by an "identical" sensor, that is, a sensor from the same manufacturer and model than the initial one. Nevertheless, there are not two identical sensors, so differences in the sampled data will be appreciated after replacement.

To this extend, multiple sensors of the same model will be employed to sample the volatiles. Different sensors (even from the same manufacturer and model) will react differently to the same gas exposure, allowing the study of the above effects on the classification performance. An interesting point to be studied by the classification algorithm would be determining if a pre-calibration of all the sensors of the same model is required, or if the algorithm is capable of correctly classifying the volatiles despite the differences in the sensors responses.

C.2.1 Setup

The experimental setup consists of a robot that is commanded to sample the content of four recipients containing different substances (acetone, ethanol, butane [1] and gin). To address this, the robot is provided with a robotic arm [125] (see Figure C.4) which allows an easy displacement and correct positioning of the e-nose aspiration (attached

[1] Since butane is found at gas state at ambient temperature, the content of a lighter was released when the e-nose aspiration moved over the container.

Figure C.4: The JACO robotic arm mounted on a mobile platform, and a detailed view of the attached e-nose aspiration.

Figure C.5: Picture of the proposed setup. Each of the black plastic vessels contains a different substance.

to its hand) over the recipients containing the volatile samples. Figure C.5 shows a picture of the proposed setup.

To avoid waiting for the sensors to recover their baseline levels after each exposure (which would take more than a minute), this data-set employs the MCE-nose (see Section 3.3) which allows the measurement of fast changing gas concentrations. However, the consideration of the MCE-nose aims not only at increase the measurement rate of the different volatiles, but to enable the study of the effects of ageing, poisoning and sensor replacement on the classification performance, by using multiple arrays of "identical" MOX sensors, corresponding to the different chambers of the MCE-nose.

The robotic arm is then commanded to approximate to the containers following a predefined sequence. The exposition to each of the substances takes 20s, after which the arm moves to another container and the MCE-nose switches to a new, clean chamber. In particular, the MCE-nose employed to collect the samples for this data-set has four chambers, each of which hosts an array of 4 MOX gas sensors: TGS-2602, TGS-2600, TGS-2611 and TGS-2620. The volatile sequence and the gathered signals are depicted in Figure C.6.

C.2.2 Data-set structure

The samples of this data set have been conveniently organized in a *.mat* file containing two variables:

AC-index: Column array with the index of the MCE-nose *active chamber* (AC), that is, the chamber being exposed to the analyte. The index values are 0, 1, 2 or 3.

AC-readings: Matrix containing the readings of the gas sensors, as provided by the MCE-nose. Each column corresponds to:

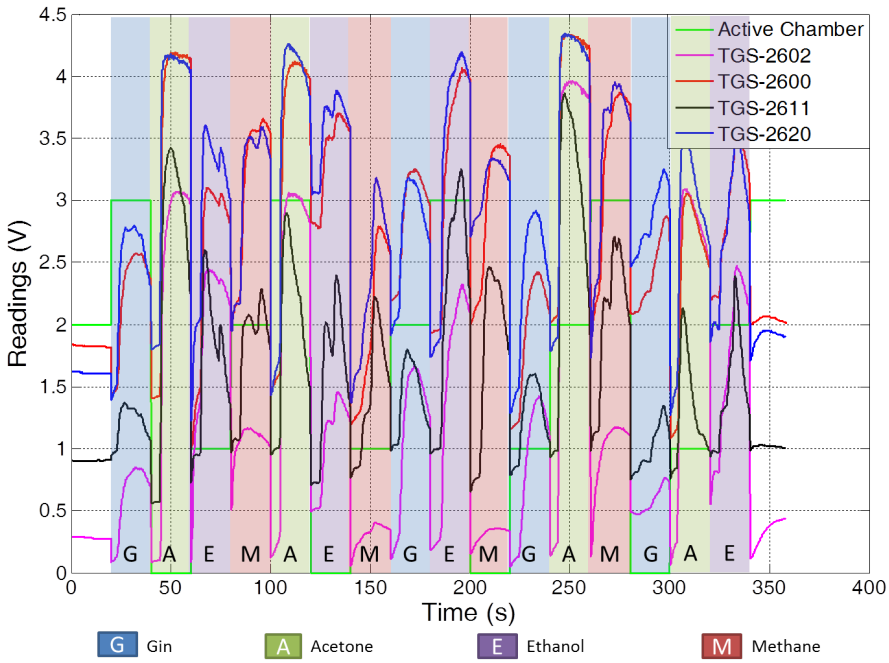

Figure C.6: Signals collected with the MCE-nose during the classification experiment with a robotic arm, and the "ground-truth" sequence of the inspected analytes. The active chamber [0,1,2,3] is switched every 20 seconds.

- Column 1: Timestamp (Elapsed time since the start of the experiment).
- Column 2: Readings of sensor Figaro TGS-2602
- Column 3: Readings of sensor Figaro TGS-2600
- Column 4: Readings of sensor Figaro TGS-2611
- Column 5: Readings of sensor Figaro TGS-2620

Additionally, with visualization purposes, the raw readings as gathered from the MCE-nose are provided in a *.txt* file in conjunction with a *Matlab* script designed to plot the readings of each chamber independently. This data-set can be downloaded at: `http://mrpt.org/Robotics_olfaction_dataset`.

C.3 Data-set 3: Classification of odors in turbulent environments

This third data-set has been built employing an open sampling system, that is, no control over the environmental conditions is performed during data acquisition. The

Figure C.7: Picture of the different analytes used in the generation of the third data-set.

objective of such setup is to obtain signals similar to those gathered by an e-nose carried by a mobile robot. Those signals are characterized by the continuous fluctuations (and consequently absence of steady state values) produced by the intrinsics mechanism of gas propagation: advection and turbulence (see Section 2.3).

Four different odor classes (acetone, ethanol, methane and gin) have been selected as target gases (see Figure C.7) for an e-nose composed by an array of 4 different MOX gas sensors: TGS-2600, TGS-2602, TGS-2611, TGS-2620. In order to enable the sensor surfaces to interact with the volatile molecules dispersed in the environment, a fan is used to generate a constant airflow through the array of sensors.

C.3.1 Cases of study

Four different cases have been considered in the generation of the presented data-set. In the first one, the e-nose is placed at a fixed distance from the gas source, while in the second case, the distance between the e-nose and the gas source varies recreating the movement of a mobile robot when passing over a gas source. Then, the third case of study exploits the exposition of the e-nose to different gas sources consecutively. Finally, the fourth case focusses in the mixture of different volatiles. Next, each of these cases of study is detailed:

First case

This part of the data-set deals with samples collected when the e-nose is placed at a fixed distance from the gas source. Initially, the e-nose is allowed to measure the baseline level for 20 seconds after which the volatile is released for a duration of another 60 seconds approximately. After that time, the volatile is removed to allow

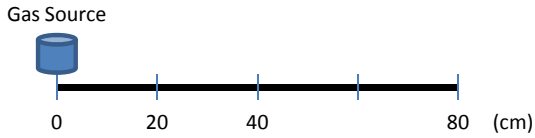

Figure C.8: Diagram of the setup employed to collect the samples of the first case of study. A fixed distance between the gas source and the e-nose is kept during the sampling process.

the sensors to slowly recover to their initial state. Three different distances between e-nose and gas source have been selected: 20cm, 40cm and 80cm (see Figure C.8). For each combination of volatile and distance, three repetitions are collected. A total of $3x4x3 = 36$ samples compose this section of the data-set.

Second case

This second case is conceived to recreate the aspects of the signals gathered by an e-nose carried by a mobile robot when passing near or over a gas source. Each sample contains the measurement of the baseline level for the initial 20 seconds, and then the sensor readings for two runs of the e-nose passing over the gas source (towards and back). Considering the generally slow movement speeds of robots carrying an e-nose, two different velocities of $5cm/s$ and $10cm/s$ are proposed here. As in the previous case, three repetitions of each speed are recorded, generating a total of 24 samples. Figure C.9 shows an example of the signals gathered with this setup.

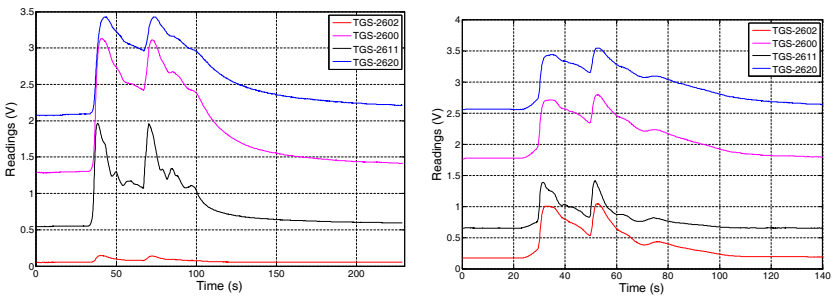

Figure C.9: Two different samples of the second case of study where an e-nose pass over the gas source at a constant speed: (a) Methane gas source at $5cm/s$ (b) Ethanol gas source at $10cm/s$.

Third case

This case of study exploits the exposition of the e-nose to different volatiles consecutively, which means that at each instant of time, only one volatile is presented to the e-nose; but the nature of the volatile changes with time. The objective of such configuration is to determine if a pattern recognition algorithm would be capable of determining the different analytes the e-nose has been exposed to and the respective time-slots, even when the sensors are not allowed to recover their baseline levels.

The samples collected have been categorized in three different groups attending to the order of exposure of the different volatiles, the distance between source and e-nose, and the exposure time:

Group 1: Samples of this group contains three out of the four volatiles, following the sequence: acetone - gin - ethanol. This sequence is repeated twice after the initial 20s assigned to baseline measurement. The distance between the e-nose and the gas source is kept constant at 30cm, and the different volatiles are presented to the e-nose each 20s.

Group 2: Again, only three out of the four volatiles available are used. In this case the sequence is based on acetone - ethanol - gin. As in the previous configurations, this sequence is repeated twice after the initial 20s assigned to baseline measurement. The distance to the source is increased to 65cm to scale up the effect of turbulence, and the exposure time is set to 40s.

Group 3: Finally, the third configuration employs all the four volatiles substances. The exposure sequence is set to: acetone - methane - ethanol - gin. As in the previous configuration, the distance to the source is set to 65cm and volatiles are presented each 40s.

To reduce the time necessary for the volatile molecules to reach the e-nose, in this scenario, a second fan was placed bellow the gas source container as depicted in Figure C.10. Additionally, an example of the signals gathered with this configuration is shown in Figure C.11.

Fourth case

To conclude this data-set, some experiments related to the mixture of multiple analytes are considered. Three repetitions of the volatile sequence detailed in Table C.1 compose this section of the data-set. The experimental setup is identical to the third case of study, setting a distance between e-nose and gas source of 65cm.

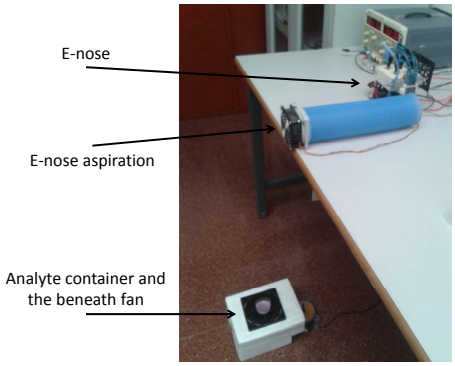

Figure C.10: Picture of the experimental setup employed in the third and fourth cases of study.

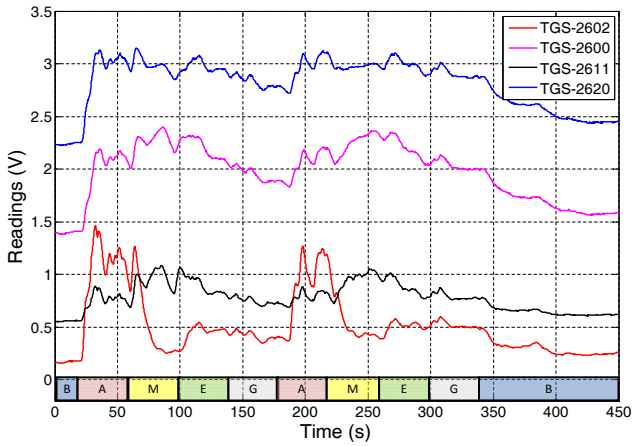

Figure C.11: Sample belonging to the third case of study (consecutive analyte exposition). The bottom rule depicts the time-slots of each analyte exposure, meaning: (B) baseline/recovery, (A) acetone, (M) methane, (E) ethanol and (G) gin.

C.3.2 Data-set structure

Samples of this third data-set have been conveniently organized in folders. Each sample (.txt file) is named using the following rule:

$$\mathbf{V}\_[at\_\mathbf{D}]\_[\mathbf{S}]\_[each\_\mathbf{P}]\_dataset\_\mathbf{DATE}.rawlog\_gasSensors.txt$$

where V is the ordered sequence of volatiles used in the sample, D is the distance between e-nose and gas source in cm, S is the movement speed of the e-nose measured in cm/s, P is the exposure time of each analyte in seconds and *DATE* is the time-stamp

Table C.1: Sequence of volatile exposure used for the fourth case of study.

Timeslot	**analytes**
0-20	Baseline
20-40	Methane
40-60	Methane + Acetone
60-80	Acetone
80-100	Acetone + Ethanol
100-120	Ethanol
120-140	Ethanol + Gin
140-160	Gin
160-180	Gin + Methane

in the format ($yyyy-mm-dd\_hhmmss$). Square brackets indicates that the value is optional depending on the case of study.

Samples of the two first cases of study have been separated by analyte, thus, each folder named as one of the four analytes contains the samples corresponding to these two cases of study. Two additional folders, each containing the samples of the third and fourth cases of study respectively, are also provided. As in the previous data-sets a *Matlab* script designed to visualize the readings of each sample is also attached. This data-set can be downloaded at: `http://mrpt.org/Robotics_olfaction_dataset`.

Appendix D
Least Squares

This appendix details the optimization of the least squares form of the cost function $F(\mathbf{m})$, defined as the following sum of quadratic forms:

$$F(\mathbf{m}) = \mathbf{r}^\top \Lambda \mathbf{r} \tag{D.1}$$

where $\mathbf{r}$ is the vector of errors or *residuals*, a measure of the mismatch between the prediction and the observation, and Λ stands for the information matrix, i.e. the inverse of the covariance matrix.

In general this process is iterative, providing for each iteration a small increment ($\Delta\mathbf{m}$) of the current state towards the optimal value ($\hat{\mathbf{m}}$). Nevertheless, in the case proposed in Chapter 5, since the residuals are linear functions, the optimal value can be found in only one iteration:

$$\hat{\mathbf{m}} \leftarrow \mathbf{m} + \Delta\mathbf{m}^* \tag{D.2}$$

This increment $\Delta\mathbf{m}^*$ can be shown to arise as a solution to the equation:

$$\Delta\mathbf{m}^* \leftarrow \left. \frac{\partial F(\mathbf{m}_k + \Delta\mathbf{m})}{\partial \Delta\mathbf{m}} \right|_{\Delta\mathbf{m}=0} = 0 \tag{D.3}$$

Approximating the cost function $F(\mathbf{m})$ by its second-order Taylor series expansion in the vicinity of its actual state $\mathbf{m}_k$:

$$\begin{aligned}
\mathrm{F}(\mathbf{m}_k+\Delta\mathbf{m}) &\approx \hat{F}_k(\mathbf{m}_k+\Delta\mathbf{m}) \\
&= \mathrm{F}(\mathbf{m}_k)+\underbrace{\left.\frac{\partial \mathrm{F}}{\partial \mathbf{m}}\right|_{\mathbf{m}=\mathbf{m}_k}}_{\nabla_{\mathbf{m}}\mathrm{F}(\mathbf{m}_k)}\Delta\mathbf{m}+\frac{1}{2}\Delta\mathbf{m}^T\underbrace{\left.\frac{\partial^2 \mathrm{F}}{\partial \mathbf{m}\partial \mathbf{m}^T}\right|_{\mathbf{m}=\mathbf{m}_k}}_{\nabla^2_{\mathbf{m}}\mathrm{F}(\mathbf{m}_k)}\Delta\mathbf{m} \\
&= \mathrm{F}(\mathbf{m}_k)+\underbrace{\nabla_{\mathbf{m}}\mathrm{F}(\mathbf{m}_k)}_{\mathbf{g}_k^T}\Delta\mathbf{m}+\frac{1}{2}\Delta\mathbf{m}^T\underbrace{\nabla^2_{\mathbf{m}}\mathrm{F}(\mathbf{m}_k)}_{\mathbf{H}_k}\Delta\mathbf{m} \qquad \text{(D.4)} \\
&= \mathrm{F}(\mathbf{m}_k)+\mathbf{g}_k^T\Delta\mathbf{m}+\frac{1}{2}\Delta\mathbf{m}^T\mathbf{H}_k\Delta\mathbf{m}
\end{aligned}$$

where we introduce the first and second-order derivatives of $F(\mathbf{m})$, namely the gradient vector $\mathbf{g_k}=\nabla_{\mathbf{m}}\mathrm{F}(\mathbf{m})^\top$ and the Hessian matrix $\mathbf{H_k}=\nabla^{\mathbf{2}}_{\mathbf{m}}\mathrm{F}(\mathbf{m_k})$.

Taking now derivatives with respect to an increment in the unknowns, we obtain:

$$\begin{aligned}
\frac{\partial \mathrm{F}(\mathbf{m}_k+\Delta\mathbf{m})}{\partial\Delta\mathbf{m}} &\approx \frac{\partial \hat{F}_k(\mathbf{m}_k+\Delta\mathbf{m})}{\partial\Delta\mathbf{m}} \qquad \text{(D.5)} \\
&= \underbrace{\frac{\partial}{\partial\Delta\mathbf{m}}\{\mathrm{F}(\mathbf{m}_k)\}}_{0}+\frac{\partial}{\partial\Delta\mathbf{m}}\left\{\mathbf{g}^T\Delta\mathbf{m}\right\}+\frac{\partial}{\partial\Delta\mathbf{m}}\left\{\frac{1}{2}\Delta\mathbf{m}^T\mathbf{H}\Delta\mathbf{m}\right\}
\end{aligned}$$

Provided that $\frac{\partial \mathbf{a^T M a}}{\partial \mathbf{a}}=\left(\mathbf{M}+\mathbf{M^T}\right)\mathbf{a}$ and $\frac{\partial \mathbf{a^T b}}{\partial \mathbf{b}}=\mathbf{a}$, and since the Hessian matrix is symmetric:

$$\frac{\partial \hat{F}_k(\mathbf{m}_k+\Delta\mathbf{m})}{\partial\Delta\mathbf{m}}=\mathbf{g}+\mathbf{H}\Delta\mathbf{m} \qquad \text{(D.6)}$$

The increment of the current state that leads to the optimal value of the gas distribution map is then determined by identifying Eq. (D.6) to zero:

$$\left.\frac{\partial \hat{F}_k(\mathbf{m}_k+\Delta\mathbf{m})}{\partial\Delta\mathbf{m}}\right|_{\Delta\mathbf{m}=0}=0 \quad\rightarrow\quad \mathbf{g}_k+\mathbf{H}_k\Delta\mathbf{m}_k^*=0$$

Therefore, $\Delta\mathbf{m}_k^*$ is computed by solving the linear system of the form $\mathbf{Ax}=\mathbf{b}$:

$$\mathbf{H_k}\Delta\mathbf{m_k^*}=-\mathbf{g_k} \qquad \text{(D.7)}$$

This linear system can be rewritten by taking the first order Taylor extension on the residuals $\mathbf{r}$, leading to:

$$\underbrace{(\mathbf{J}^\top\Lambda\mathbf{J})}_{\text{Hessian }\mathbf{H}}\Delta\mathbf{m}^*=-\underbrace{\mathbf{J}^\top\Lambda\mathbf{r}}_{\text{Gradient }\mathbf{g}} \qquad \text{(D.8)}$$

Bibliography

[1] W. C. Agosta. *Chemical Communication: The Language of Pheromones*. Library series. Scientific American Library, 1992.

[2] K. J. Albert, N. S. Lewis, C. L. Schauer, G. A. Sotzing, S. E. Stitzel, T. P. Vaid, and D. R. Walt. Cross-reactive chemical sensor arrays. *Chemical Reviews*, 100(7):2595–2626, 2000.

[3] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell*. Garland Science, 5th edition, nov 2008.

[4] Ansys. The ansys fluent project web page. `http://www.ansys.com`.

[5] S. Asadi, S. Pashami, A. Loutfi, and A. J. Lilienthal. TD Kernel DM+V: time-dependent statistical gas distribution modelling on simulated measurements. In *AIP Conference Proceedings Volume 1362: Olfaction and Electronic Nose - Proceedings of the 14th International Symposium on Olfaction and Electronic Nose (ISOEN)*, pages 281–283, 2011.

[6] C. Bishop. *Pattern recognition and machine learning*. Information science and statistics. Springer, 2006.

[7] A. Bjorck. *Numerical methods for least squares problems*. Number 51. Society for Industrial and Applied Mathematics, 1996.

[8] J.-L. Blanco, J.-A. Fernandez-Madrigal, and J. Gonzalez-Jimenez. A New Approach for Large-Scale Localization and Mapping: Hybrid Metric-Topological SLAM. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 2061–2067, 2007.

[9] J.-L. Blanco, J. G. Monroy, J. Gonzalez-Jimenez, and A. J. Lilienthal. A kalman filter based approach to probabilistic gas distribution mapping. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, SAC '13, pages 217–222. ACM, mar 2013.

[10] J.-L. Blanco, J. Gonzalez-Jimenez, and J.-A. Fernandez-Madrigal. Extending obstacle avoidance methods through multiple parameter-space transformations. *Autonomous Robots*, 24(1):29–48, 2008.

[11] J.-L. Blanco, J. Gonzalez-Jimenez, and J.-A. Fernandez-Madrigal. Optimal filtering for non-parametric observation models: Applications to localization and SLAM. *The International Journal of Robotics Research*, 29(14):1726–1742, 2010.

[12] M. Blanco, J. Coello, H. Iturriaga, S. Maspoch, and J. Pagés. Calibration in non-linear near infrared reflectance spectroscopy: A comparison of several methods. *Analytica Chimica Acta*, 384(2):207–214, 1999.

[13] K. Bradford and D. Desrochers. The use of scents to influence consumers: The sense of using scents to make cents. *Journal of Business Ethics*, 90(2):141–153, 2009.

[14] P. Brimblecombe. *Air Composition and Chemistry*. Cambridge Environmental Chemistry Series. Cambridge University Press, 1996.

[15] A. Bürkle, F. Segor, and M. Kollmann. Towards Autonomous Micro UAV Swarms. *Journal of Intelligent & Robotic Systems*, 61(1-4):339–353, 2011.

[16] G. Cabrita, P. Sousa, and L. Marques. Player/stage simulation of olfactory experiments. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1120–1125, 2010.

[17] CARMEN-Team. The Carnegie Mellon robot navigation toolkit. `http://carmen.sourceforge.net/home.html`.

[18] P. Clifford. Markov Random Fields in statistics. In G. Grimmett and D. Welsh, editors, *Disorder in Physical Systems: A Volume in Honour of John M. Hammersley*, pages 19–32. Oxford UniveOdor-source localization in clean roomrsity Press, Oxford, 1990.

[19] R. P. Cogdill and P. Dardenne. Least-squares support vector machines for chemometrics: An introduction and evaluation. *Journal of Near Infrared Spectroscopy*, 12(2):93–100, 2004.

[20] F. Dellaert and M. Kaess. Square root SAM: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research*, 25(12):1181–1203, 2006.

[21] J. Dennis and R. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 1983.

[22] G. E. DeVaull, J. A. King, R. J. Lantzy, and D. J. Fontaine. *Understanding Atmospheric Dispersion of Accidental Releases*. John Wiley & Sons, Inc., 2010.

[23] K. Domanský, D. L. Baldwin, J. W. Grate, T. B. Hall, J. Li, M. Josowicz, and J. Janata. Development and calibration of field-effect transistor-based sensor array for measurement of hydrogen and ammonia gas mixtures in humid air. *Analytical Chemistry*, 70(3):473–481, 1998.

[24] B. Ehret, K. Safenreiter, F. Lorenz, and J. Biermann. A new feature extraction method for odour classification. *Sensors and Actuators B: Chemical*, 158(1):75 – 88, 2011.

[25] J. Elkinton, R. Cardé, and C. Mason. Evaluation of time-average dispersion models for estimating pheromone concentration in a deciduous forest. *Journal of Chemical Ecology*, 10(7):1081–1108, 1984.

[26] J. Engelberger. *Robotics in Service*. MIT Press, 1989.

[27] A. M. Farah and T. Duckett. Reactive localisation of an odour source by a learning mobile robot. In *In Proceedings of the Second Swedish Workshop on Autonomous Robotics*, pages 29–38, 2002.

[28] J. Farrell, J. Murlis, X. Long, W. Li, and R. Cardé. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environmental Fluid Mechanics*, 2(1-2):143–169, 2002.

[29] J. Farrell, S. Pang, and W. Li. Chemical plume tracing via an autonomous underwater vehicle. *Oceanic Engineering, IEEE Journal of*, 30(2):428–442, 2005.

[30] J. Farrell, S. Pang, W. Li, and R. Arrieta. Chemical plume tracing experimental results with a REMUS AUV. In *OCEANS 2003. Proceedings*, volume 2, pages 962–968 Vol.2, 2003.

[31] J.-A. Fernandez-Madrigal and J.-L. Blanco. *Simultaneous Localization and Mapping for Mobile Robots: Introduction and Methods*. IGI Global, sep 2012.

[32] J.-A. Fernandez-Madrigal, C. Galindo, J. Gonzalez-Jimenez, E. Cruz-Martín, and A. Cruz-Martín. A software engineering approach for the development of heterogeneous robotic applications. *Robotics and Computer-Integrated Manufacturing*, 24(1):150 – 166, 2008.

[33] J.-A. Fernandez-Madrigal and J. Gonzalez-Jimenez. NEXUS: a flexible, efficient and robust framework for integrating software components of a robotic system. In *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, volume 1, pages 524–529 vol.1, 1998.

[34] G. Ferri, E. Caselli, V. Mattoli, A. Mondini, B. Mazzolai, and P. Dario. SPIRAL: A novel biologically-inspired algorithm for gas/odor source localization in an indoor environment with no strong airflow. *Robotics and Autonomous Systems*, 57(4):393 – 402, 2009.

[35] M. Freese, S. Singh, F. Ozaki, and N. Matsuhira. Virtual robot experimentation platform V-REP: A versatile 3d robot simulator. In N. Ando, S. Balakirsky, T. Hemker, M. Reggiani, and O. Stryk, editors, *Simulation, Modeling, and Programming for Autonomous Robots*, volume 6472 of *Lecture Notes in Computer Science*, pages 51–62. Springer Berlin Heidelberg, 2010.

[36] J. G. Monroy, J.-L. Blanco, and J. Gonzalez-Jimenez. An open source framework for simulating mobile robotics olfaction. In *15th International Symposium on Olfaction and Electronic Nose (ISOEN)*, jul 2013.

[37] J. G. Monroy, J.-L. Blanco, and J. Gonzalez-Jimenez. Time-Variant Gas Distribution Mapping with Obstacle Information. *Submitted*, 2013.

[38] J. G. Monroy, J. Gonzalez-Jimenez, and J.-L. Blanco. Overcoming the slow recovery of MOX gas sensors through a system modeling approach. *Sensors*, 12(10), oct 2012.

[39] J. G. Monroy, A. J. Lilienthal, J.-L. Blanco, J. Gonzalez-Jimenez, and M. Trincavelli. Calibration of MOX gas sensors in open sampling systems based on gaussian processes. In *Sensors, 2012 IEEE*, pages 1–4, 2012.

[40] J. G. Monroy, A. J. Lilienthal, J.-L. Blanco, J. Gonzalez-Jimenez, and M. Trincavelli. Probabilistic gas quantification with MOX sensors in Open Sampling Systems - A Gaussian Process approach. *Sensors and Actuators B: Chemical*, 188(0):298 – 312, 2013.

[41] C. Galindo, J. Gonzalez-Jimenez, and J.-A. Fernandez-Madrigal. A Control Architecture for Human-Robot Integration. Application to a Robotic Wheelchair. *IEEE Transactions on Systems, Man, and Cybernetics, part B*, 36(5):1053–1068, 2006.

[42] J. W. Gardner. A non-linear diffusion-reaction model of electrical conduction in semiconductor gas sensors. *Sensors and Actuators B: Chemical*, 1(1-6):166 – 170, 1990.

[43] J. W. Gardner and P. Bartlett. *Electronic noses: principles and applications*. Oxford science publications. Oxford University Press, 1999.

[44] J. W. Gardner and P. N. Bartlett. A brief history of electronic noses. *Sensors and Actuators B: Chemical*, 18(1-3):210 – 211, 1994.

[45] P. J. Gemperline, J. R. Long, and V. G. Gregoriou. Nonlinear multivariate calibration using principal components regression and artificial neural networks. *Analytical Chemistry*, 63(20):2313–2323, 1991.

[46] B. P. Gerkey, R. T. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *In Proceedings of the 11th International Conference on Advanced Robotics*, pages 317–323, 2003.

[47] M. Giurfa and J. Núñez. Honeybees mark with scent and reject recently visited flowers. *Oecologia*, 89(1):113–117, 1992.

[48] G. H. Golub and R. J. Plemmons. Large-scale geodetic least-squares adjustment by dissection and orthogonal decomposition. *Linear Algebra and its Applications*, 34(0):3 – 28, 1980.

[49] J. Gonzalez-Jimenez, J. G. Monroy, and J.-L. Blanco. The Multi-Chamber electronic nose - An improved olfaction sensor for mobile robotics. *Sensors*, 11(6):6145–6164, 2011.

[50] J. Gonzalez-Jimenez, J. G. Monroy, J.-L. Blanco, and F. Garcia. Electronic nose having a high sensing frequency and method for determining the quantitative and qualitative composition of a gas or mixture of gases using same, 2012. WO Patent 2012-049341.

[51] J. Gonzalez-Jimenez, J. G. Monroy, F. Garcia, and J.-L. Blanco. The Multi-Chamber Electronic Nose (MCE-nose). In *IEEE International Conference on Mechatronics*, apr 2011.

[52] M. group. The Open Mobile Robot Arquitecture (OpenMORA) website. `http://sourceforge.net/projects/openmora`.

[53] R. Gutierrez-Osuna. Pattern analysis for machine olfaction: a review. *Sensors Journal, IEEE*, 2(3):189–202, 2002.

[54] R. Gutierrez-Osuna, H. T. Nagle, and S. S. Schiffman. Transient response analysis of an electronic nose using multi-exponential models. *Sensors and Actuators B: Chemical*, 61(1-3):170 – 182, 1999.

[55] L. Hadjiiski, P. Geladi, and P. Hopke. A comparison of modeling nonlinear systems with artificial neural networks and partial least squares. *Chemometrics and Intelligent Laboratory Systems*, 49(1):91–103, 1999.

[56] A. Hayes, A. Martinoli, and R. Goodman. Distributed Odor Source Localization. *IEEE Sensors Journal, Special Issue on Electronic Nose Technologies*, 2(3):260–273, 2002.

[57] V. Hernandez Bennets, A. J. Lilienthal, P. Neumann, and M. Trincavelli. Mobile robots for localizing gas emission sources on landfill sites: is bio-inspiration the way to go? *Frontiers in Neuroengineering*, 4(0), 2012.

[58] Hokuyo. The Hokuyo automated co. ltd web page. `http://www.hokuyo-aut.jp`.

[59] M. Holmberg, F. Winquist, I. Lundstrom, F. Davide, C. DiNatale, and A. DAmico. Drift counteraction for an electronic nose. *Sensors and Actuators B-Chemical*, 36(1-3):528–535, OCT 1996.

[60] S. B. i Badia, U. Bernardet, A. Guanella, P. Pyk, and P. F. M. J. Verschure. A Biologically Based chemo-sensing UAV for Humanitarian Demining. *International Journal of Advanced Robotic Systems*, 4(2):187–198, 06/2007 2007.

[61] K. Ihokura and J. Watson. *The stannic oxide gas sensor: principles and applications*. CRC Press, 1994.

[62] H. Ishida. Blimp robot for three-dimensional gas distribution mapping in indoor environment. volume 1137, pages 61–64, 2009.

[63] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi. Odor-source localization in clean room by autonomous mobile sensing system. In *Solid-State Sensors and Actuators, 1995 and Eurosensors IX. Transducers '95. The 8th International Conference on*, volume 1, pages 783–786, 1995.

[64] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi. Odor-source localization in the clean room by an autonomous mobile sensing system. *Sensors and Actuators B: Chemical*, 33(1–3):115 – 121, 1996.

[65] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote sensing and localization of gas/odor source and distribution using mobile sensing system. In *Solid State Sensors and Actuators, 1997. TRANSDUCERS '97 Chicago., 1997 International Conference on*, volume 1, pages 559–562 vol.1, 1997.

[66] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote sensing of gas/odor source location and concentration distribution using mobile system. *Sensors and Actuators B: Chemical*, 49(1-2):52 – 57, 1998.

[67] H. Ishida, G. Nakayama, T. Nakamoto, and T. Moriizumi. Controlling a gas/odor plume-tracking robot based on transient responses of gas sensors. *Sensors Journal, IEEE*, 5(3):537–545, 2005.

[68] H. Ishida, K. Suetsugu, T. Nakamoto, and T. Moriizumi. Study of autonomous mobile sensing system for localization of odor source using gas sensors and anemometric sensors. *Sensors and Actuators A: Physical*, 45(2):153 – 157, 1994.

[69] H. Ishida, H. Tanaka, H. Taniguchi, and T. Moriizumi. Mobile robot navigation using vision and olfaction to search for a gas/odor source. *Autonomous Robots*, 20(3):231–238, 2006.

[70] H. Ishida, T. Ushiku, S. Toyama, H. Taniguchi, and T. Moriizumi. Mobile robot path planning using vision and olfaction to search for a gas source. In *Sensors, 2005 IEEE*, page 4 pp., 30 2005-nov. 3 2005.

[71] F. J. Urban air quality - their physical and chemical characteristics. *Atmospheric Environment*, 33(29):4877–4900, 1999.

[72] J. Janata. *Principles of Chemical Sensors*. Springer Publishing Company, 2nd ed. edition, 2009.

[73] C. Jones. On the structure of instantaneous plumes in the atmosphere. *Journal of Hazardous Materials*, 7(2):87 – 112, 1983.

[74] M. Kac. Random walk and the theory of Brownian Motion. *The American Mathematical Monthly*, 54(7):pp. 369–391, 1947.

[75] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 2149–2154, 2004.

[76] M. Kovacina, D. Palmer, G. Yang, and R. Vaidyanathan. Multi-agent control algorithms for chemical cloud detection and mapping using unmanned air vehicles. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 3, pages 2782–2788 vol.3, 2002.

[77] G. Kowadlo and R. A. Russell. Robot odor localization: A taxonomy and survey. *The International Journal of Robotics Research*, 27(8):869–894, 2008.

[78] Y. Kuwana, S. Nagasawa, I. Shimoyama, and R. Kanzaki. Synthesis of the pheromone-oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors. *Biosensors and Bioelectronics*, 14(2):195 – 202, 1999.

[79] Y. Kuwana, I. Shimoyama, and H. Miura. Steering control of a mobile robot using insect antennae. In *Intelligent Robots and Systems 95. 'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on*, volume 2, pages 530–535 vol.2, 1995.

[80] Y. Kuwana, I. Shimoyama, Y. Sayama, and H. Miura. Synthesis of pheromone-oriented emergent behavior of a silkworm moth. In *Intelligent Robots and Systems '96, IROS 96, Proceedings of the 1996 IEEE/RSJ International Conference on*, volume 3, pages 1722–1729 vol.3, 1996.

[81] A. J. Lilienthal. *Gas Distribution Mapping and Gas Source Localisation with a Mobile Robot*. PhD thesis, Wilhelm-Schickard Institute, University of Tübingen, December 2004.

[82] A. J. Lilienthal and T. Duckett. Approaches to gas source tracing and declaration by pure chemo-tropotaxis. In *Autonome Mobile Systeme (AMS), 18. Fachgespräch*, pages 161–171. GDI, December 4 – 5 2003.

[83] A. J. Lilienthal and T. Duckett. Creating gas concentration gridmaps with a mobile robot. In *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 1, pages 118–123 vol.1, 2003.

[84] A. J. Lilienthal and T. Duckett. Gas source localisation by constructing concentration gridmaps with a mobile robot. In *In Proceedings of the European Conference on Mobile Robots (ECMR 2003*, 2003.

[85] A. J. Lilienthal and T. Duckett. Building gas concentration gridmaps with a mobile robot. *Robotics and Autonomous Systems*, 48(1):3–16, August 31 2004.

[86] A. J. Lilienthal and T. Duckett. Experimental analysis of gas-sensitive braitenberg vehicles. *Advanced Robotics*, 18:817–834, 2004.

[87] A. J. Lilienthal, A. Loutfi, J.-L. Blanco, C. Galindo, and J. Gonzalez-Jimenez. A Rao-Blackwellisation approach to GDM-SLAM. Integrating SLAM and gas distribution mapping. In *Proceedings of the European Conference on Mobile Robots (ECMR)*, pages 126–131, September 19–21 2007.

[88] A. J. Lilienthal, A. Loutfi, and T. Duckett. Airborne chemical sensing with mobile robots. *Sensors*, 6(11):1616–1678, 2006.

[89] A. J. Lilienthal, M. Reggente, M. Trincavelli, J.-L. Blanco, and J. Gonzalez-Jimenez. A statistical approach to gas distribution modelling with mobile robots - the Kernel DM+V algorithm. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 570–576, October 11 – October 15 2009.

[90] A. J. Lilienthal, F. Streichert, and A. Zell. Model-based shape analysis of gas concentration gridmaps for improved gas source localisation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3575–3580, April 18 – 22 2005.

[91] A. J. Lilienthal, M. R. Wandel, U. Weimar, and A. Zell. Experiences using gas sensors on an autonomous mobile robot. In *Proceedings of EUROBOT 2001, 4th European Workshop on Advanced Mobile Robots*, pages 1–8. IEEE, IEEE Computer Press, September 19 - 21 2001.

[92] A. J. Lilienthal, A. Zell, M. Wandel, and U. Weimar. Sensing odour sources in indoor environments without a constant airflow by a mobile robot. In *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, volume 4, pages 4005–4010 vol.4, 2001.

[93] Q. Liu, C. shu Li, and X. tian Guan. Simulation study on robot active olfaction based on concentration and equilateral triangle search. In *Robotics and Biomimetics (ROBIO), 2010 IEEE International Conference on*, pages 625–628, 2010.

[94] E. Lobet. Dynamic pattern recognition methods and system identification. In T. Pearce, S. Schiffman, H. Nagle, and J. Gardner, editors, *Handbook of Machine Olfaction.*, chapter 12, pages 293–321. Wiley-VCH, 2000.

[95] A. Loutfi, S. Coradeschi, L. Karlsson, and M. Broxvall. Putting olfaction into action: using an electronic nose on a multi-sensing mobile robot. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 1, pages 337–342 vol.1, 2004.

[96] A. Loutfi, S. Coradeschi, A. J. Lilienthal, and J. Gonzalez-Jimenez. Gas distribution mapping of multiple odour sources using a mobile robot. *Robotica*, 27(2):311–319, June 4 2009.

[97] A. G. MacDiarmid. Synthetic metals: a novel role for organic polymers. *Current Applied Physics*, 1(4–5):269 – 279, 2001.

[98] D. J. C. MacKay. *Bayesian methods for back-propagation networks*, pages 211–254. Models of neural networks 3. Springer, 1994.

[99] A. Marjovi and L. Marques. Optimal spatial formation of swarm robotic gas sensors in odor plume finding. *Autonomous Robots*, pages 1–17, 2013.

[100] L. Marques, N. Almeida, and A. De Almeida. Olfactory sensory system for odour-plume tracking and localization. In *Sensors, 2003. Proceedings of IEEE*, volume 1, pages 418–423 Vol.1, 2003.

[101] L. Marques and A. de Almeida. Electronic nose-based odour source localization. In *Advanced Motion Control, 2000. Proceedings. 6th International Workshop on*, pages 36–40, 2000.

[102] D. Martinez and L. Perrinet. Cooperation between vision and olfaction in a koala robot. Technical report, Report on the 2002 Workshop on Neuromorphic Engineering, 2002.

[103] D. Martinez, O. Rochel, and E. Hugues. A biomimetic robot for tracking specific odors in turbulent plumes. *Autonomous Robots*, 20(3):185–195, 2006.

[104] D. McGurk, J. frost, E. Eisenbraun, K. Vick, W. Drew, and J. Young. Volatile compounds in ants: Identification of 4-methyl-3-heptanone from pogonomyrmex ants. *Insect Physiology*, (12):1433 – 1441, 1966.

[105] O. Michel. Webots: Professional Mobile Robot Simulation. *International Journal of Advanced Robotic Systems*, 1(1):39–42, 2004.

[106] MobileRobots Inc. Corporate website. `http://www.mobilerobots.com`.

[107] P. T. Moseley. Solid state gas sensors. *Measurement Science and Technology*, 8(3):223, 1997.

[108] MRPT-Team. The mobile robot programming toolkit (MRPT) website. `http://mrpt.org`.

[109] T. Nakamoto, H. Ishida, and T. Moriizumi. An odor compass for localizing an odor source. *Sensors and Actuators B: Chemical*, 35(1 - 3):32 – 36, 1996.

[110] T. Nakamoto and K. Yoshikawa. Movie with scents generated by olfactory display using solenoid valves. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, E89-A(11):3327–3332, Nov. 2006.

[111] C. D. Natale, S. Marco, F. Davide, and A. D'Amico. Sensor-array calibration time reduction by dynamic modelling. *Sensors and Actuators B: Chemical*, 25(1-3):578 – 583, 1995. Proceedings of the Fifth International Meeting on Chemical Sensors.

[112] R. Neal. *Bayesian learning for neural networks*. Lecture notes in statistics. Springer, 1996.

[113] P. P. Neumann. *Gas Source Localization and Gas Distribution Mapping with a Micro-Drone*. PhD thesis, Fachbereichs Mathematik und Informatik der Freien Universität Berlin, March 2013.

[114] P. Newman. Mission Oriented Operating Suite (MOOS). `http://www.robots.ox.ac.uk/~mobile/MOOS/wiki/pmwiki.php`, 2006.

[115] K. Ogata. *Modern Control Engineering*. Prentice-Hall electrical engineering series. Instrumentation and controls series. Prentice Hall, 2010.

[116] S. Pang and J. Farrell. Chemical plume source localization. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 36(5):1068–1080, 2006.

[117] A. Pardo, S. Marco, and J. Samitier. Nonlinear inverse dynamic models of gas sensing systems based on chemical sensor arrays for quantitative measurements. *Instrumentation and Measurement, IEEE Transactions on*, 47(3):644 –651, jun 1998.

[118] S. Pashami, S. Asadi, and A. J. Lilienthal. Integration of OpenFOAM flow simulation and filament-based gas propagation models for gas dispersion simulation. In *Proceedings of the Open Source CFD International Conference*, 2010.

[119] T. Pearce. *Handbook of machine olfaction: electronic nose technology*. Wiley-VCH, 2003.

[120] K. Persaud and G. H. Dodd. Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature*, 299:352–355, 1982.

[121] P. Pyk, S. Bermúdez i Badia, U. Bernardet, P. Knüsel, M. Carlsson, J. Gu, E. Chanie, B. Hansson, T. Pearce, and P. J. Verschure. An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. *Autonomous Robots*, 20(3):197–213, 2006.

[122] M. Rachkov, L. Marques, and A. T. Almeida. Multisensor demining robot. *Autonomous Robots*, 18(3):275–291, 2005.

[123] C. Rasmussen and C. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.

[124] M. Reggente and A. J. Lilienthal. Three-dimensional statistical gas distribution mapping in an uncontrolled indoor environment. In *AIP Conference Proceedings Volume 1137: Olfaction and Electronic Nose - Proceedings of the 13th International Symposium on Olfaction and Electronic Nose (ISOEN)*, pages 109–112, 2009.

[125] K. Robotics. The Kinova corporate web page. `http://www.kinovarobotics.com`.

[126] R. Rozas, J. Morales, and D. Vega. Artificial smell detection for robotic navigation. In *Advanced Robotics, 1991. 'Robots in Unstructured Environments', 91 ICAR., Fifth International Conference on*, pages 1730–1733 vol.2, 1991.

[127] A. Russell, D. Thiel, and A. Mackay-Sim. Sensing odour trails for mobile robot navigation. In *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*, pages 2672–2677 vol.3, 1994.

[128] R. Russell. *Odour Detection by Mobile Robots*. World Scientific Publishing Co, London, UK, 1999.

[129] R. A. Russell. Laying and sensing odor markings as a strategy for assisting mobile robot navigation tasks. *Robotics Automation Magazine, IEEE*, 2(3):3–9, 1995.

[130] R. A. Russell. Ant trails - an example for robots to follow? In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, volume 4, pages 2698–2703 vol.4, 1999.

[131] R. A. Russell. Survey of robotic applications for odor-sensing technology. *The International Journal of Robotics Research*, 20(2):144–162, 2001.

[132] R. A. Russell. Chemical source location and the RoboMole project. In *Proceedings of the Australasian Conference on Robotics and Automation*, 2003.

[133] J. Samitier, J. Lopez-Villegas, S. Marco, L. Cámara, A. Pardo, O. Ruiz, and J. Morante. A new method to analyse signal transients in chemical sensors. *Sensors and Actuators B: Chemical*, 18(1 - 3):308 – 312, 1994.

[134] G. Sandini, G. Lucarini, and M. Varoli. Gradient driven self-organizing systems. In *Intelligent Robots and Systems '93, IROS '93. Proceedings of the 1993 IEEE/RSJ International Conference on*, volume 1, pages 429–432 vol.1, 1993.

[135] G. Sberveglieri. Classical and novel techniques for the preparation of SnO2 thin-film gas sensors. *Sensors and Actuators B: Chemical*, 6(1 - 3):239 – 247, 1992.

[136] R. E. Shaffer, S. L. Rose-Pehrsson, and R. McGill. A comparison study of chemical sensor array pattern recognition algorithms. *Analytica Chimica Acta*, 384(3):305 – 317, 1999.

[137] H. Shen. *Environmental Fluid Mechanics: Theories and Applications: Committee Report*. Amer Society of Civil Engineers, 2002.

[138] A. Shmilovici, G. Bakir, S. Marco, and A. Perera. Finding the best calibration points for a gas sensor array with support vector regression. In *Intelligent Systems, 2004. Proceedings. 2004 2nd International IEEE Conference*, volume 1, pages 174 – 177 Vol.1, june 2004.

[139] B. I. Shraiman and E. D. Siggia. Scalar turbulence. *Nature*, 405:639–646, 2000.

[140] SICK. The SICK corportate web page. `http://www.sick.com`.

[141] S. Sklavounos and F. Rigas. Validation of turbulence models in heavy gas dispersion over obstacles. *Journal of Hazardous Materials*, 108(1–2):9 – 20, 2004.

[142] E. Stella, F. Musio, L. Vasanelli, and A. Distante. Goal-oriented mobile robot navigation using an odour sensor. In *Intelligent Vehicles Symposium., Proceedings of the*, pages 147–151, 1995.

[143] H. Sundgren, F. Winquist, I. Lukkari, and I. Lundstrom. Artificial neural networks and gas sensor arrays: Quantification of individual components in a gas mixture. *Measurement Science and Technology*, 2(5):464–469, 1991.

[144] S. Tauseef, D. Rashtchian, and S. Abbasi. CFD-based simulation of dense gas dispersion in presence of obstacles. *Journal of Loss Prevention in the Process Industries*, 24(4):371 – 376, 2011.

[145] M. Trincavelli. *Gas Discrimination for Mobile Robots*. PhD thesis, Orebro University, 2010.

[146] M. Trincavelli. Gas discrimination for mobile robots. *Künstliche Intelligenz*, 25(4):351 – 354, 2011.

[147] M. Trincavelli, S. Coradeschi, and A. Loutfi. Odour classification system for continuous monitoring applications. *Sensors and Actuators B: Chemical*, 58:265 – 273, 2009.

[148] M. Trincavelli, S. Coradeschi, and A. Loutfi. Online classification of gases for environmental exploration. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3311 – 3316, 2009.

[149] M. Trincavelli, M. Reggente, S. Coradeschi, H. Ishida, A. Loutfi, and A. J. Lilienthal. Towards environmental monitoring with mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2210 – 2215, 2008.

[150] W. Tsujita, A. Yoshino, H. Ishida, and T. Moriizumi. Gas sensor network for air-pollution monitoring. *Sensors and Actuators B: Chemical*, 110(2):304 – 311, 2005.

[151] S. D. Vito, E. Massera, M. Piga, L. Martinotto, and G. D. Francia. On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario. *Sensors and Actuators B: Chemical*, 129(2):750 – 757, 2008.

[152] M. R. Wandel, A. J. Lilienthal, T. Duckett, U. Weimar, and A. Zell. Gas distribution in unventilated indoor environments inspected by a mobile robot. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003.

[153] M. R. Wandel, U. Weimar, A. J. Lilienthal, and A. Zell. Leakage localisation with a mobile robot carrying chemical sensors. In *Electronics, Circuits and Systems, 2001. ICECS 2001. The 8th IEEE International Conference on*, volume 3, pages 1247–1250 vol.3, 2001.

[154] D. R. Webster, K. Y. Volyanskyy, and M. J. Weissburg. Bioinspired algorithm for autonomous sensor-driven guidance in turbulent chemical plumes. *Bioinspiration & Biomimetics*, 7(3):036023, 2012.

[155] M. Webster. *Webster's Seventh New Collegiate Dictionary*. Merriam-Webster Inc. (Springfield, Massachusetts), 1969.

[156] E. O. Wilson and M. Pavan. Glandular sources and specificity of some chemical releasers of social behavior in dolichoderine ants. *Psyche*, (66):70 – 76, 1959.

[157] G. Winkler. *Image analysis, random fields and Markov chain Monte Carlo methods: a mathematical introduction*, volume 27. Springer Verlag, 2003.

[158] T. Yamanaka, H. Ishida, T. Nakamoto, and T. Moriizumi. Analysis of gas sensor transient response by visualizing instantaneous gas concentration using smoke. *Sensors and Actuators A: Physical*, 69(1):77 – 81, 1998.