



Doctoral Dissertation

Probabilistic Methods for Robotic Gas Source Localization

José Ojeda Morala
2024

Dirigida por
Javier González Jiménez
Javier González Monroy

Tesis doctoral por compendio de publicaciones
Programa de Doctorado en Ingeniería Mecatrónica
Dpt. de Ingeniería de Sistemas y Automática
Universidad de Málaga

UNIVERSIDAD
DE MÁLAGA





UNIVERSIDAD
DE MÁLAGA

AUTOR: José Ojeda Morala

 <https://orcid.org/0000-0002-2326-5269>

EDITA: Publicaciones y Divulgación Científica. Universidad de Málaga



Esta obra está bajo una licencia de Creative Commons Reconocimiento-NoComercial-SinObraDerivada 4.0 Internacional:

<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>

Cualquier parte de esta obra se puede reproducir sin autorización pero con el reconocimiento y atribución de los autores.

No se puede hacer uso comercial de la obra y no se puede alterar, transformar o hacer obras derivadas.

Esta Tesis Doctoral está depositada en el Repositorio Institucional de la Universidad de Málaga (RIUMA): riuma.uma.es

UNIVERSIDAD
DE MÁLAGA





DECLARACIÓN DE AUTORÍA Y ORIGINALIDAD DE LA TESIS PRESENTADA PARA OBTENER EL TÍTULO DE DOCTOR

D./Dña JOSÉ OJEDA MORALA

Estudiante del programa de doctorado de INGENIERÍA MECATRÓNICA de la Universidad de Málaga, autor/a de la tesis, presentada para la obtención del título de doctor por la Universidad de Málaga, titulada: PROBABILISTIC METHODS FOR ROBOTIC GAS SOURCE LOCALIZATION

Realizada bajo la tutorización de JAVIER GONZÁLEZ JIMÉNEZ y dirección de JAVIER GONZÁLEZ JIMÉNEZ y JAVIER GONZÁLEZ MONROY (si tuviera varios directores deberá hacer constar el nombre de todos)

DECLARO QUE:

La tesis presentada es una obra original que no infringe los derechos de propiedad intelectual ni los derechos de propiedad industrial u otros, conforme al ordenamiento jurídico vigente (Real Decreto Legislativo 1/1996, de 12 de abril, por el que se aprueba el texto refundido de la Ley de Propiedad Intelectual, regularizando, aclarando y armonizando las disposiciones legales vigentes sobre la materia), modificado por la Ley 2/2019, de 1 de marzo.

Igualmente asumo, ante la Universidad de Málaga y ante cualquier otra instancia, la responsabilidad que pudiera derivarse en caso de plagio de contenidos en la tesis presentada, conforme al ordenamiento jurídico vigente.

En Málaga, a 16 de Enero de 2025

Fdo.: JOSÉ OJEDA MORALA Doctorando/a	Fdo.: JAVIER GONZÁLEZ JIMÉNEZ Tutor/a
Fdo.: JAVIER GONZÁLEZ JIMÉNEZ Director/es de tesis	



UNIVERSIDAD DE MÁLAGA
DEPARTAMENTO DE
INGENIERÍA DE SISTEMAS Y AUTOMÁTICA

El Dr. D. Javier González Jiménez y el Dr. D. Javier González Monroy, directores de la tesis titulada “Probabilistic Methods for Robotic Gas Source Localization” realizada por D. José Ojeda Morala, certifican su idoneidad para la obtención del título de Doctor en Ingeniería Mecatrónica.

Málaga, 16 de Enero de 2025

Dr. D. Javier González Jiménez

Dr. D. Javier González Monroy

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



Probabilistic Methods for Robotic Gas Source Localization

AUTHOR: José Ojeda Morala

SUPERVISOR: Javier González Jiménez
Javier González Monroy



All models are wrong, but some of them are useful
George E. P. Box

Prefiero caminar con una duda que con un mal axioma
Javier Krahe

Contents

Acknowledgments	iii
Abstract	vi
Resumen (Spanish)	vii
Introducción	viii
1 Introduction	1
1.1 Contributions	4
2 Background	7
2.1 Theoretical Framework and Mathematical Tools	7
2.2 Gas Dispersion Models and Airflow Estimation	15
2.3 Robotic Olfaction Sensors	19
3 Modeling the Relation between Measurements and Source Location	21
3.1 Introduction	21
3.2 Contribution	23
3.A Information-driven Gas Source Localization Exploiting Gas and Wind Local Measurements for Autonomous Mobile Robots . .	24
3.B On Gas Source Declaration Methods for Single-Robot Search .	25
4 Source Localization through Probabilistic Gas Mapping	26
4.1 Introduction	26
4.2 Contribution	28

4.A	Robotic Gas Source Localization with Probabilistic Mapping and Online Dispersion Simulation	29
5	Using Semantic Information to Leverage Additional Sensory Data	30
5.1	Introduction	30
5.2	Contribution	31
5.A	PSGSL: A Probabilistic Framework Integrating Semantic Scene Understanding and Gas Sensing for Gas Source Localization . .	33
6	Development of Robotic Olfaction Research Tools	34
6.1	Introduction	34
6.2	Contribution	35
6.A	VGR Dataset: A CFD-based Gas Dispersion Dataset for Mobile Robotic Olfaction	36
6.B	A Simulation Framework for the Integration of Artificial Olfaction into Multi-Sensor Mobile Robots	37
7	Conclusions and Future Work	38

Acknowledgments

Las palabras “tesis doctoral” suenan tan grandes que me cuesta trabajo reconciliarlas con lo que este documento es: un trabajo de clase especialmente largo, una colección de cosas que se me han ido ocurriendo, un intento de convencer a una imaginaria audiencia de que sé de lo que estoy hablando. Me cuesta, también, procesarlo como un gran logro. No puedo decir que acabar el doctorado haya sido un titánico esfuerzo, en el que la pasión por la ciencia me ha llevado a sobreponerme a grandes dificultades. Y sin embargo, eso es más o menos lo que se espera, ¿no? Una historia de vocación y superación personal.

A decir verdad, entré en el mundo de la investigación sin realmente buscarlo. Quedarme en la universidad era solo una manera de huir de la horripilante perspectiva de tener que buscar trabajo. Ni siquiera sabía que el tema sobre el que trata esta tesis existía hasta que empecé a trabajar en él. No puedo, por tanto, venderos que esto siempre fue mi sueño.

Tampoco ha sido un grave sufrimiento. No me he quedado sin dormir por las noches, no he caído en una depresión, no me han maltratado los jefes. Ni siquiera me han obligado a irme de estancia. Quizá es que esto del doctorado no es para tanto, entonces. Pueden dispersarse, falsa alarma, siento haberles hecho perder el tiempo.

Y, aún así.

Aún así, leyendo mi propia tesis, no puedo evitar una cierta incredulidad, un cierto orgullo, una voz que pregunta “¿de verdad hemos hecho todo esto?”. No un gran logro, monolítico e impresionante, sino una larga lista de logros más pequeños, más asequibles, reunidos. Presentados por compendio, podría decirse.

No podré contaros que la investigación siempre ha sido mi trabajo soñado, que el olfato era mi vocación. Pero sí que puedo decir, sin faltar a la verdad, que he llegado a emocionarme escribiendo sobre probabilidades. Que he descubierto en lo que puede parecer un campo técnico y frío una oportunidad para la creatividad. Que he encontrado una cierta belleza, guardada en secreto solo para los iniciados, en los conceptos más abstractos.

Y supongo que no, no ha sido trivial llegar aquí. Que el motivo por el que no siento que haya sido tan difícil es que no habido un desafío singular, sino una larga, larguísima lista de pequeños esfuerzos, de los cuales la mayoría ni siquiera son míos. Y si el esfuerzo no es todo mío, tampoco lo es el logro. Así pues, no me queda más remedio que compartir el momento con todos los que me han acompañado, y darles las gracias:

A mis compañeros de Mapir, pasados y presentes, por crear un ambiente tan bueno que hasta me gusta ir a trabajar.

A mis tutores, por estar siempre dispuestos a guiarme, y por escuchar pacientemente mis divagaciones incoherentes sobre variables aleatorias.

A mis amigos, por compartir conmigo la vida más allá del trabajo, y por fingir que les gusta mi comida incluso cuando me sale mal.

A todos los profesores que he tenido, porque no podría haber hecho nada de esto sin todo lo que me han enseñado.

Y a mi familia, por tantas cosas que no tiene sentido enumerarlas.

Pepe

The work for this PhD thesis was funded by the grant for pre-doctoral researchers from the Junta of Andalucía. Some of the works presented here were also funded by projects WISER (DPI2017-84827-R), ARPEGGIO (PID2020-117057GB-I00), and HOUNDBOT (P20_01302).

Abstract

The field of robotics is a complex one. Many things can go wrong when an autonomous machine interacts with the real world, from sensors producing unreliable measurements to actuators behaving in unpredictable manners. A robot operating with such unreliable tools must account for the uncertainty that they cause in its estimations of the state of the world.

Probabilistic robotics is an approach that embraces this uncertainty by modeling all the knowledge and estimation processes of the robot with the formal tools of probability theory. This creates a rigorous, robust way of handling uncertainty even as the complexity of the problems increases. Such an approach is beneficial in most cases, but becomes crucial when the information that the robot obtains can only be related to the state it tries to estimate through unreliable models.

One such problem is gas source localization (GSL), where a robot equipped with a gas sensor and an anemometer must solve a fluid dynamics problem to figure out which point in the environment a gas is being emitted from. Due to the complexity of the phenomenon of gas dispersion, GSL techniques must make assumptions and simplifications to be able to reason about it, which increases uncertainty.

In this thesis, we tackle the subject of source localization with an autonomous mobile robot from the perspective of probabilistic robotics. We explain the difficulties the problem presents, and propose solutions based on probabilistic modeling and Bayesian reasoning.

Resumen (Spanish)

La robótica puede resultar compleja. Hay muchas cosas que pueden salir mal cuando una máquina autónoma interactúa con el mundo real: sensores que producen medidas erróneas, actuadores que se comportan de manera impredecible, *etc.* Un robot que trabaja con herramientas tan poco fiables debe tener en cuenta la incertidumbre que estos errores causan en sus estimaciones.

La robótica probabilística es un enfoque que se centra en el manejo de esta incertidumbre, modelando la información que el robot posee, y sus procesos de estimación, mediante las herramientas formales de la teoría de probabilidades. Estas herramientas permiten gestionar la incertidumbre de una manera rigurosa y robusta, incluso cuando se trata con problemas de gran complejidad. Un enfoque de estas características resulta beneficioso en la mayoría de los casos, y se vuelve crucial cuando la información disponible al robot solo puede relacionarse con el estado que éste intenta estimar a través de modelos inexactos.

Un ejemplo de este tipo de problemas es la búsqueda de fuentes de gas (GSL por sus siglas en inglés), donde un robot equipado con un sensor de gas puntual y un anemómetro debe resolver un problema de dinámica de fluidos para determinar desde qué punto del entorno se está emitiendo un gas. Debido a la complejidad del fenómeno de la dispersión de gases, es necesario que las técnicas de GSL hagan suposiciones y simplificaciones sobre las condiciones ambientales para hacer el problema tratable, aumentando así la incertidumbre.

En esta tesis, estudiaremos el problema de la localización de fuentes de gas con un robot móvil autónomo desde la perspectiva de la robótica probabilística. Explicaremos las dificultades que el problema presenta, y propondremos soluciones basadas en el modelado probabilístico y el razonamiento bayesiano.

Introducción

La robótica es un campo de estudio apasionante. Los robots existen justo en la intersección entre los ordenadores, cuyas increíbles capacidades han figurativamente dado forma al mundo moderno, y la maquinaria física que es literalmente capaz de dar forma a cosas en el mundo real. Si la tecnología continúa avanzando, los robots pueden llegar a ser capaces de realizar todas las tareas peligrosas, repetitivas, o de otro modo problemáticas que los humanos nos vemos forzados a hacer a día de hoy. Por tanto, en la robótica existe un gran potencial para cambiar la sociedad en que vivimos. Cabe esperar que, en el futuro, la robótica sea no solo el tema de discusiones técnicas, sino también de debates éticos y políticos.

Sea cual sea su potencial, a día de hoy la robótica es también un campo de estudio dificultoso. El salto desde el estructurado mundo de las abstracciones matemáticas a la caótica realidad del nuestro no es sencillo para un ordenador. La interacción con el mundo real es impredecible, y la información sensorial de la que los robots disponen para entenderlo es muy limitada.

Es precisamente esta última cuestión, la información sensorial disponible, la que define gran parte de la investigación en robótica. Campos de estudio enteros surgen de la mera existencia de un tipo concreto de sensor, puesto que *utilizar* una forma de información sensorial puede ser tan difícil (o más) como diseñar el sensor que la proporciona. La *robótica olfativa*, que es el tema sobre el que trata esta tesis, es uno de estos casos.

Para poder entender correctamente el campo de la robótica olfativa, así como la investigación que esta tesis presenta, debemos primero introducir el tema del olfato artificial y los sensores de gas.

Cuando hablamos sobre *olfato artificial*, nos referimos a la habilidad de una máquina para detectar y cuantificar sustancias volátiles en el aire. Este tipo de tecnología está a día de hoy presente en numerosas situaciones: la monitorización de la polución en las grandes ciudades [1, 2], la detección de humo en el interior de edificios [3], el diagnóstico de enfermedades [4, 5, 6], el control de calidad en la elaboración de productos alimenticios [7, 8, 9], *etc.* A medida que la tecnología en sí mejora, más y más aplicaciones aparecen.

Es común que los investigadores nos refiramos a las sustancias detectadas por estos sensores con el término “gases” (y así lo haremos a lo largo de esta tesis, meramente por conveniencia), pero lo cierto es que no siempre lo son. Por ejemplo, la mayoría de los sensores de humo modernos no detectan ninguno de los gases que se producen en la combustión, sino las pequeñas partículas sólidas de material quemado que quedan suspendidas en el aire [10]. De manera rigurosa, por tanto, no son *sensores de gas*, pero aún así los consideraríamos sensores olfativos, puesto que detectan la presencia de una sustancia en el aire.

Utilizando este tipo de sensores se puede resolver un amplio abanico de problemas. Taxonómicamente hablando, estos problemas se pueden dividir en tres áreas principales:

Reconocimiento de olores, o clasificación de gases

El problema del reconocimiento de olores [11, 12] es el primero que viene a la mente de la mayoría de personas cuando oyen hablar del olfato artificial. La pregunta en la que se basa este tema es “¿Qué estoy oliendo?”. Como el lector podrá imaginar, este problema está estrechamente relacionado con varias de las aplicaciones que hemos listado anteriormente (por ejemplo, el control de calidad de productos). Ser capaz de detectar un gas específico dentro de una mezcla es de gran importancia en tales casos, y ese es precisamente el tipo de tareas en las que se centra la investigación de este campo.

El reconocimiento de olores requiere una comprensión profunda del funcionamiento de los propios sensores, y de cómo reaccionan ante distintas sustancias. A menudo no resulta posible encontrar *hardware* que detecte específicamente la sustancia en la que se está interesado, sin ser afectado por otros gases presentes. En estos casos es necesario analizar la respuesta de múltiples tipos de sensores para poder discernir qué es lo que se está detectando. Por tanto, gran parte de la investigación en este campo trata sobre el concepto de *nariz electrónica* (en inglés, *e-nose*) [13, 14]. Estos dispositivos contienen grupos de sensores que utilizan una variedad de tecnologías y están diseñados para detectar diferentes gases, así como otros sensores para medir magnitudes como la temperatura o la humedad. La información proporcionada por todos estos sensores puede procesarse en conjunto para producir una medida más detallada y completa que la que sería posible basándose sólo en la elección de *hardware*.

Creación de mapas de distribución de gases

Habitualmente, los gases no crean una distribución uniforme al dispersarse. La concentración de gas en un punto del entorno puede ser significativamente distinta a la que hay en otro punto, incluso si solo los separa una distancia pequeña. Estas variaciones pueden estar causadas por la presencia y posición de las fuentes de las que el gas emana, o por las condiciones del viento. La creación de mapas de distribución de gases (GDM por sus siglas en inglés) [15, 16, 17] consiste en registrar estas variaciones en la concentración del gas en el entorno.

Existen muchos casos en los que conocer qué áreas acumulan una mayor concentración puede ser importante. Por supuesto, están los casos en los que el gas estudiado es peligroso de por sí, pero también es posible que la acumulación del gas permita identificar otros problemas. Por ejemplo, durante la pandemia del COVID-19, los mapas de concentración de CO_2 fueron utilizados para identificar áreas con ventilación insuficiente, donde la probabilidad de contagio de enfermedades transmitidas por vía aérea es mayor [18].

La creación de mapas de concentración con un robot móvil es una tarea que requiere razonar sobre la manera en que los gases se dispersan y los mecanismos físicos subyacentes. A medida que el robot se mueve por el entorno y obtiene

un pequeño número de medidas sensoriales, las técnicas de GDM intentan rellenar los huecos en el mapa generado utilizando información adicional, como la geometría del entorno o la dirección del viento [19]. El problema de la creación de mapas de concentración se vuelve significativamente más complejo cuando se considera el hecho de que la distribución del gas no es constante en el tiempo [20], y que por tanto las medidas sensoriales tomadas en un instante concreto no constituyen una representación fiable del mapa de concentración existente un tiempo después.

Localización de fuentes de gas

De manera similar al problema de GDM, la localización de fuentes de gas (GSL en inglés) [21, 22, 23] trata el modo en que los gases se dispersan. Sin embargo, a diferencia de las técnicas de GDM, los algoritmos de GSL no intentan crear un mapa exhaustivo de la concentración de gas en el entorno, sino solamente identificar el punto desde el que el gas es emitido. La lista de aplicaciones prácticas del problema de GSL incluye la identificación de fugas en tuberías o depósitos, la detección de contrabando, la localización de explosivos, *etc.*

La localización de fuentes de gas es una tarea de gran dificultad, particularmente cuando se desarrolla en entornos complejos. Describir la dispersión del gas en un entorno requiere el uso de modelos de dinámica de fluidos [24], que son notoriamente complejos tanto a nivel conceptual como a nivel computacional. Complicando aún más el asunto, los algoritmos de GSL tratan de resolver el problema a la inversa: en lugar de comenzar con la posición de la fuente y utilizar el modelo para predecir cómo el gas se dispersa, intentan deducir dónde la fuente debe estar situada para haber podido producir las medidas sensoriales que se han obtenido.

Este problema, la búsqueda de fuentes de gas, es el principal tema de esta tesis. En los capítulos restantes exploraremos la complejidad del problema en sí, y propondremos soluciones novedosas basadas en los mecanismos y la filosofía de la robótica probabilística.

- - -

Si bien puede resultar conveniente considerar estos tres problemas por separado, es evidente que existe una profunda relación entre ellos. Tratar de estudiar esta conexión puede resultar imponente, ya que la complejidad de cada uno de los problemas aumenta si la tenemos en cuenta. Sin embargo, esta conexión nos permite diseñar maneras más sofisticadas de razonar sobre el problema, y nos proporciona información adicional que utilizar en estos razonamientos.

Esta idea, que es posible diseñar métodos mejores y más inteligentes al aumentar deliberadamente la complejidad del problema que consideramos, es una parte fundamental de la propuesta de esta tesis. Para evitar que esta

complejidad adicional nos supere, utilizamos las herramientas formales de la teoría de probabilidades y la poderosa maquinaria de la estimación bayesiana. Con estas herramientas somos capaces de definir las relaciones existentes entre conceptos abstractos, y de combinar múltiples fuentes de información de manera rigurosa, llegando así a las soluciones que presentamos.

Contribuciones

Nuestras contribuciones en esta tesis tratan sobre el modelado del problema de la localización de fuentes de gas en términos de teoría de probabilidades y razonamiento bayesiano.

Proponemos varios métodos nuevos de búsqueda de fuentes, procurando en cada caso no solo presentar un algoritmo concreto, sino también un análisis más profundo del problema y de cómo nuestros razonamientos intuitivos pueden formalizarse. El código que implementa nuestras distintas propuestas puede encontrarse en el siguiente repositorio público *online*:

<https://github.com/MAPIRlab/GasSourceLocalization>

La tesis incluye también algunos trabajos que no presentan algoritmos de GSL, sino que contribuyen a la investigación en este campo presentando nuevas herramientas de investigación que se hacen también públicas. Los enlaces a estas herramientas se encuentran en las secciones relevantes del capítulo 6.

Lista de publicaciones

Esta tesis se presenta como un compendio de artículos de investigación. A continuación está la lista de los artículos que la componen:

- **Information-driven Gas Source Localization Exploiting Gas and Wind Local Measurements for Autonomous Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *IEEE Robotics and Automation Letters (RAL)*, vol. 6, no. 2, pp. 1320-1326, 2021.
- **On Gas Source Declaration Methods for Single-Robot Search**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the ISOCS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*, Aveiro (Portugal), 2022.
- **Robotic Gas Source Localization with Probabilistic Mapping and Online Dispersion Simulation**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *IEEE Transactions on Robotics (TRO)*, vol. 40, pp. 3551-3564, 2024
- **PSGSL: A Probabilistic Framework Integrating Semantic Scene Understanding and Gas Sensing for Gas Source**

Localization, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Not yet published*

- **VGR Dataset: A CFD-based Gas Dispersion Dataset for Mobile Robotic Olfaction**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Journal of Intelligent & Robotic Systems*, 2023.
- **A Simulation Framework for the Integration of Artificial Olfaction into Multi-Sensor Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *MDPI Sensors*, vol. 21, no. 6, pp. 2041, 2021.

Los siguientes trabajos también fueron desarrollados durante la investigación en la que se basa la tesis, si bien no forman parte de este documento:

- **An Evaluation of Gas Source Localization Algorithms for Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the International Conference on Applications of Intelligent Systems (APPIS)*, Las Palmas de Gran Canaria (Spain), 2020.
- **Experimental Analysis of the Impact of Sensor Response Time on Robotic Gas Source Localization**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the ISOCs/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*, Aveiro (Portugal), 2022.
- **GadenTools: A Toolkit for Testing and Simulating Robotic Olfaction Tasks With Jupyter Notebook Support**, Pepe Ojeda, Jose-Raul Ruiz-Sarmiento, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the Fifth Iberian Robotics Conference (ROBOT2022)*, Zaragoza (Spain), 2022.
- **A Feasibility Study of a Leader-Follower Multi-Robot Formation for TDLAS Assisted Methane Detection in Open Spaces**, Javier Monroy, Pepe Ojeda and Javier Gonzalez-Jimenez, *Proceedings of the Iberian Robotics Conference (ROBOT)*, Coimbra (Portugal), 2023.

Estructura de la tesis

Los capítulos restantes están organizados de la siguiente manera:

El capítulo 2 explica los fundamentos teóricos y las herramientas matemáticas requeridas para entender los contenidos de la tesis y nuestras propuestas. Si bien estas herramientas serán explicadas cuando sea relevante

en el propio texto de los artículos presentados, la redacción de artículos científicos requiere asumir una cierta familiaridad del lector con el tema tratado. Por tanto, recomendamos la lectura de este capítulo para obtener el contexto necesario.

El capítulo 3 explora la formulación probabilística más sencilla del problema de búsqueda de fuentes, donde la posición de la fuente es estimada directamente a partir de las medidas sensoriales. Como parte de este capítulo, presentamos un nuevo algoritmo de GSL diseñado para entornos de interiores con geometría compleja. Así mismo, presentamos un segundo trabajo en el que se explora el sub-problema de la declaración de la fuente, utilizando el algoritmo previamente presentado para llevar a cabo la búsqueda.

El capítulo 4 introduce una versión más compleja del problema, que combina la búsqueda de fuentes con la creación de mapas de distribución. En este capítulo discutimos algunas de las limitaciones del método más simple, y cómo la introducción de los mapas de distribución nos permite superarlas.

El capítulo 5 trata sobre el concepto de la comprensión semántica del entorno, y sobre cómo ésta nos permite razonar sobre la posición de la fuente de gas incluyendo tipos distintos de información. Como parte de este capítulo presentaremos una formulación probabilística para el problema de la localización de fuentes con información semántica, y discutiremos la influencia del reconocimiento de olores en la utilización de esta nueva información.

El capítulo 6 presenta nuevas herramientas de investigación, desarrolladas durante la tesis. Estas herramientas incluyen un conjunto de datos de simulaciones de dispersión de gases de alta fidelidad, y la integración de simulaciones olfativas y visuales para el desarrollo de algoritmos tales como los presentados en el capítulo 5.

Por último, **el capítulo 7** presenta nuestras conclusiones y discute las principales líneas para el trabajo futuro.

Robotics is an exciting research field. It exists right at the intersection between computers, whose incredible capabilities have figuratively shaped the modern world, and the physical machinery that is literally capable of shaping things in the real world. As technology advances and more refined techniques are developed, robots may become able to perform all the dangerous, repetitive, or otherwise problematic tasks that humans must now do, drastically changing society. One should expect that, in the future, robotics will be the subject of not only technical discussions, but also political and ethical debates.

However promising it may be, robotics is also a challenging research field. Making the jump from the neatly organized world of mathematical abstraction to the messy reality of our own is not an easy thing for computers. Interaction with the real world is unpredictable, and very limited sensory information is available for the robots to try and make sense of it.

It is this last thing, the sensory information that is available, which defines much of the research on robotics. Entire fields of study stem from the mere existence of a certain type of sensor, as *making use* of some kind of sensory information can be as challenging (or more) as designing the sensor that provides it. One such case is *robotic olfaction*, which happens to be the subject of this thesis.

To understand the field of robotic olfaction, and the research this thesis presents, we must first introduce the broader topic of artificial olfaction and gas sensors.

When we talk about *artificial olfaction*, we refer to the ability of a machine to detect and measure volatile substances in the air. This type of technology is already in common use in everyday scenarios: monitoring pollution in cities [1, 2], detecting smoke in buildings [3], medical diagnosis [4, 5, 6], performing

quality control on food products [7, 8, 9], *etc.* As the technology itself advances, and more substances can be monitored, more applications appear.

We researchers commonly refer to these substances as *gases* (and we will do so throughout this thesis, for convenience), but in truth they not always are. For example, many smoke sensors do not detect any of the gases that are released by combustion, but instead the small solid particles of burnt material that get suspended in the air and carried by it [10]. Rigorously speaking, then, they are not *gas sensors*, but we would consider them olfactory sensors nonetheless, as they measure the presence of something in the air.

Using these kinds of sensors, we can try to solve a great variety of problems. Taxonomically speaking, there are three main subjects of interest in the application of artificial olfaction:

Smell recognition or gas classification

The subject of smell recognition [11, 12] is the first thing that comes to most people's mind when hearing about artificial olfaction. The question at the heart of this problem is "What am I smelling?". As the reader can imagine, this pertains to several of the previously discussed applications, like product quality controls. Being able to detect a specific gas in a mixture of multiple substances and measuring its concentration is of great importance for such goals, and that is the type of task that smell recognition is concerned with.

Smell recognition requires a deep understanding of the sensors themselves and how they react to different substances. It is often not possible to find hardware that will detect exactly the substance the application requires, without being affected by other volatiles. Instead, one must often rely on analyzing the response of multiple sensors of different types to discern what is being sensed. As a result, much of the research in this field has to do with the concept of *electronic noses*, or *e-noses* [13, 14]. These devices contain arrays of gas sensors using multiple technologies and targeting different gases, alongside other sensors like thermometers or hygrometers. The information from these sensor arrays can be processed to produce a more comprehensive and detailed form of sensing, compared to relying simply on the choice of hardware.

Gas distribution mapping

As gases disperse, they most often do not distribute uniformly. The concentration of a gas at one point of the environment might be significantly different from the concentration at another, even a small distance away. These variations can be caused by the presence and location of sources, or because of the characteristics of the airflow. When performing gas distribution mapping (GDM) [15, 16, 17], one attempts to create a map of the environment that registers these variations.

Knowledge about which areas accumulate more concentration of a given gas can be important in many cases. Of course, there are the situations where a specific gas being released is dangerous by itself, but sometimes the accumulation of gas can signal a different problem. For example, during the COVID-19 pandemic, concentration maps of CO_2 were used to identify areas of public buildings with insufficient ventilation, which increases the probability of contagion of airborne diseases [18].

GDM with mobile robots is a task that requires thinking about the way gas disperses and the underlying physical mechanisms. As a robot moves and gathers sparse measurements at different points of the environment, GDM algorithms attempt to fill in the gaps by using additional information, like the geometry of the environment or the direction of the airflow [19]. The problem of GDM is made significantly more complex when considering the fact that the distribution of a gas is not time-constant [20], and thus measurements taken at one point of the search will not correctly represent the concentration map at a later stage.

Gas source localization

Much like GDM, gas source localization (GSL) [21, 22, 23] is concerned with the idea of how gases disperse. Unlike GDM, however, the goal of GSL is not to produce an exhaustive map of the concentration of the gas, but instead to merely identify the specific point where the gas is being released. Practical applications for GSL include locating leaks in pipes or tanks, detecting contraband, locating explosives, *etc.*

GSL is a very challenging task, particularly when dealing with complex environments. Describing the dispersal of gases in an environment requires the usage of fluid dynamics models [24], which are notoriously complex both conceptually and computationally. Further complicating matters, GSL is trying to solve the problem backwards: instead of starting with the source position and modeling how the gas disperses, it tries to reconstruct where the gas must have been released from to cause a specific set of measurements.

This problem is the main focus of this thesis. In the remaining chapters, we will explore the complexity of the problem itself, and propose new solutions rooted in the mechanisms and philosophy of probabilistic robotics.

- - -

While it is convenient to consider these three problems as separate research subjects, they are of course profoundly connected. Tackling this connection can seem daunting, as the complexity of the problem increases when we consider it. However, in this connection we can also find more sophisticated ways to reason about the problems, and more information to base our reasoning on.

This idea, that we can design better, more intelligent methods by deliberately increasing the complexity of the considered problem, is an important part of the proposal of this thesis. To avoid being overwhelmed by this additional complexity, we rely on the formal tools of probability theory and the powerful machinery of Bayesian estimation. These tools allow us to describe the relations between abstract concepts and to combine multiple sources of information in a rigorous way, thus allowing us to arrive at the solutions we propose.

1.1 Contributions

Our contributions in this thesis all revolve around framing the problem of gas source localization in the terms of probability theory and Bayesian estimation.

We propose several new methods for performing GSL, attempting to contribute in each case not only a specific algorithm but a more fundamental analysis of the problem and of how our intuitions about it can be formalized. The code for the implementation of our proposed methods is made publicly available at an online repository:

<https://github.com/MAPIRlab/GasSourceLocalization>

The thesis also includes some works which do not present GSL methods, but which instead contribute to the field by presenting new research tools which are also made available to the research community. Links to these tools are available in the relevant chapters.

List of publications

The thesis is built as a compendium of research articles. Following is the list of the contributions that compose it:

- **Information-driven Gas Source Localization Exploiting Gas and Wind Local Measurements for Autonomous Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *IEEE Robotics and Automation Letters (RAL)*, vol. 6, no. 2, pp. 1320-1326, 2021.
- **On Gas Source Declaration Methods for Single-Robot Search**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the ISOCS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*, Aveiro (Portugal), 2022.
- **Robotic Gas Source Localization with Probabilistic Mapping and Online Dispersion Simulation**, Pepe Ojeda, Javier Monroy and

1.1. CONTRIBUTIONS

Javier Gonzalez-Jimenez, *IEEE Transactions on Robotics (TRO)*, vol. 40, pp. 3551-3564, 2024

- **PSGSL: A Probabilistic Framework Integrating Semantic Scene Understanding and Gas Sensing for Gas Source Localization**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Not yet published*
- **VGR Dataset: A CFD-based Gas Dispersion Dataset for Mobile Robotic Olfaction**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Journal of Intelligent & Robotic Systems*, 2023.
- **A Simulation Framework for the Integration of Artificial Olfaction into Multi-Sensor Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *MDPI Sensors*, vol. 21, no. 6, pp. 2041, 2021.

While not included in the thesis, the following works have also been published during the research that lead to it:

- **An Evaluation of Gas Source Localization Algorithms for Mobile Robots**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the International Conference on Applications of Intelligent Systems (APPIS)*, Las Palmas de Gran Canaria (Spain), no. 24, pp. 1-6, 2020.
- **Experimental Analysis of the Impact of Sensor Response Time on Robotic Gas Source Localization**, Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the ISOCS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*, Aveiro (Portugal), 2022.
- **GadenTools: A Toolkit for Testing and Simulating Robotic Olfaction Tasks With Jupyter Notebook Support**, Pepe Ojeda, Jose-Raul Ruiz-Sarmiento, Javier Monroy and Javier Gonzalez-Jimenez, *Proceedings of the Fifth Iberian Robotics Conference (ROBOT2022)*, Zaragoza (Spain), 2022.
- **A Feasibility Study of a Leader-Follower Multi-Robot Formation for TDLAS Assisted Methane Detection in Open Spaces**, Javier Monroy, Pepe Ojeda and Javier Gonzalez-Jimenez, *Proceedings of the Iberian Robotics Conference (ROBOT)*, Coimbra (Portugal), 2023.

Outline

The remaining chapters are organized as follows:

Chapter 2 discusses the theoretical background and mathematical tools required to understand our proposals and contributions. While specific tools will be explained when relevant in the presented works themselves, the process of redacting research articles requires assuming the reader a certain familiarity with the field. Thus, we recommend reading this chapter to obtain the necessary context.

Chapter 3 explores the simplest probabilistic formulation for the GSL problem, where the location of the source is estimated directly from the sensory measurements. As part of this chapter, we present a new probabilistic method for source localization in indoors environments featuring complex geometry. We also present a second work which explores the sub-problem of source declaration using the previously discussed method.

Chapter 4 introduces a more complex version of the problem, which combines source localization and gas distribution mapping. We discuss some of the limitations of the simpler approach, and how the introduction of mapping allows us to overcome them.

Chapter 5 deals with the subject of semantic scene understanding, and how it allows us to reason about the location of the gas source by including fundamentally different types of information. We present a formally derived probabilistic framework for semantics-based source localization, and discuss the relevance of smell recognition for making use of this new information.

Chapter 6 presents several newly developed research tools. These include a dataset of high-fidelity gas dispersion simulations and an integration of olfactory and visual simulation tools for the design and testing of algorithms like the ones discussed in Chapter 5.

Lastly, **Chapter 7** presents our conclusions and discusses future work.

This chapter does not present any new contributions. Instead, it discusses the mathematical tools and concepts necessary to understand the rest of the thesis. While a certain level of familiarity with the fields of probability theory and robotic olfaction will no doubt help the reader to follow what is discussed here, the intent is for this chapter to be accessible to anyone. Our main goal here is to make it so that the reader does not at any point need to put down the thesis to go read a probability theory textbook.

We will begin by discussing the most relevant mathematical tools, and giving a brief overview of the general theoretical framework we use to model the problem of source localization. More specific details might be discussed in later chapters as they become relevant. After this, we will discuss the subject of gas dispersion models and airflow estimation, which is of great importance for the problem of source localization. We will close the chapter by briefly discussing the sensing hardware used in robotic olfaction.

2.1 Theoretical Framework and Mathematical Tools

One of the most important schools of thought in the field of robotics since the 1990s is that of *probabilistic robotics* [25]. The main idea on which this approach

is cemented is that, when dealing with the real world (whether interacting with it or merely perceiving it), there is always a degree of uncertainty. Sensors have limited precision, and can be subject to random noise; control actions can lead to slightly different results than expected due to mechanical problems or imprecise models of the world.

We would like, of course, to reduce or eliminate the uncertainty by having more reliable sensors and better actuators. Indeed, this is a subject that much research effort is dedicated to, and the tools available to roboticists continue to improve. In some carefully controlled settings (*e.g.* assembly lines), robots with high-quality actuators can achieve great levels of precision and reliability.

Still, a robust framework that allows us to use imperfect information is of great importance. It not only allows us to work with the current (maybe transient) limitations of hardware, but it also matters due to a more fundamental issue, which is particularly relevant for the topic of this thesis: when basing our understanding of the world on sensory measurements, there are certain things that cannot be observed directly, and must instead be inferred from the observations through the use of models.

In our case, a robot equipped with a gas sensor cannot use it to directly observe the location of the source. Instead, all it can measure is the concentration of the target gas in its current location. Deriving from this the source location is a complex task that requires describing the behavior of the real-world phenomenon of gas dispersion with some model, and as the old adage goes, all models are wrong (even if some of them are useful). Whenever we must model the connection between *what we can observe* and *what we want to know*, the model itself is a source of uncertainty.

Probabilistic robotics embraces the idea that these imprecisions are inevitable, and the only thing that is certain is uncertainty itself. By modeling everything through probability distributions, we not only account for the possibility that we might be wrong, but also provide a robust, formal way to quantify the uncertainty of any given estimation. Uncertainty quantification is crucial, for there are many situations where being aware that you do not know something is almost as useful as actually knowing the thing.

Probability theory also provides us with the tools to update our estimations as new evidence is gathered. This incremental refinement allows us to be more accurate and more certain even if our tools remain just as unreliable. This is the main strength of probabilistic robotics, and with this idea in mind we can now talk about the main tools of the trade: the conditional probability, and the Bayesian Filter.

2.1.1 Random Variables and Conditional Probabilities

In probabilistic robotics we represent every concept or measure through a *random variable*, and describe its behavior and our estimations of it through *probability distributions*. To avoid excessive verbosity, we will use the term

2.1. THEORETICAL FRAMEWORK AND MATHEMATICAL TOOLS

probability distribution for both the discrete case (probability mass function) and the continuous case (probability density function) indistinctly. Similarly, we will talk about the probability of a variable “taking a certain value”, which in continuous cases should be interpreted as the value falling within a certain interval.

Even though probability theory itself does not make this distinction, it is useful for us to reason about these random variables as belonging to two separate groups: *state* and *observations*.

State is whatever we are trying to estimate. The position of a gas source, the pose of the robot, whether a door is open or closed. The intuition for why we model this through a random variable with a probability distribution is hopefully clear to the reader after the previous discussion: uncertainty. There is a whole space of possible hypotheses for the value of a state variable (*e.g.* all the possible locations for the source of the gas), and, accounting for uncertainty, all we can say is how probable each of those hypotheses are.

The concept of *observations* is self-explanatory: these are the values that the sensory measurements directly give us. Since these are evidence, not estimations, conceptualizing them as random variables does not encode uncertainty, but rather reliability. To understand this properly, we need to think about conditional probabilities.

We will use the notation $p(x)$ to denote $p(X = x)$ –that is, the probability that random variable X is equal to the value x . A probability expressed like this is a probability *a priori*, in lack of any relevant information that will modify the probability distribution. In general, we are not overly interested in these *a priori* (or prior) probabilities, as we have no way to calculate them and must often assume their distributions to be equiprobable over the space of possible values. Much more interesting to us is the concept of the conditional probability, which we will denote $p(x|y)$. This represents the probability of X taking the value x given that a separate variable Y takes the value y . If there is a relation between X and Y , one should expect that $p(x) \neq p(x|y)$, and thus knowing the value of Y gives you information about the value of X .

Armed with this concept, we can now understand the probabilistic nature of the observations. The most common model for an observation variable Z that is measuring state X is the following:

$$z = x + \epsilon, \tag{2.1}$$

where the value of ϵ represents random errors and will be different any time the sensor takes a measurement. Thus, the observation can be interpreted as a sample from a random variable whose distribution is defined by the value of the state and the distribution of the error. In more general terms, any distribution $p(z|x)$ that describes the probability of obtaining a certain observation given the state of the world is called the *sensor model* (see Figure 2.1 for an example).




Figure 2.1: An example of a sensor model slightly more complex than the one described by Eq. 2.1. This model describes a beam-based range sensor. The most probable sensor reading is the correct distance value, but sensor noise causes nearby values to also be likely. There is also a small increase in probability for the maximum range of the sensor, which corresponds to the case in which the measurement fails.

There is also the concept of the *inverse sensor model*, $p(x|z)$, which defines the probability of the state taking a certain value given a specific observation has been obtained. These two concepts are related through Bayes' Theorem:

$$p(x|z) = \frac{p(z|x) \cdot p(x)}{p(z)} \quad (2.2)$$

This expression already gives a basic version of the general form of probabilistic state estimation, where knowing the value of a variable (the observation) gives us information about a different variable we are interested in (the state). A more useful, generalized version of this idea that can include information from multiple observations to gradually improve our estimation and reduce uncertainty is the recursive Bayesian filter.

2.1.2 Bayesian Filtering

Bayesian filtering is a technique through which the conditional probability distribution of the state given the last acquired measurement, $p(x|z_t)$, is used to modify the pre-existing probability distribution of the state given all previous observations, $p(x|z_{1:t-1})$. This pre-existing distribution is often called the *belief* of the state. Note that, when receiving the first observation, the belief we update is just the a priori distribution, $p(x)$.

Because we use the belief from the previous step to generate the next belief, Bayesian filtering is a recursive technique. The generalized version of the Bayes Filter algorithm considers not only sensor measurements, but also control actions that modify the state, and thus considers that the state itself may change over time (See Figure 2.2A). Think, for example of the problem of robot localization: estimating the pose of the robot requires knowledge of how




Figure 2.2: (A) Bayesian network that represents the recursive state estimation in its most general version. The state x_t at a specific time step conditions the measurement z_t , but it is in turn affected by the previous state x_{t-1} and by the control action u_{t-1} . (B) Bayesian network that represents the simpler case where the state is constant over time.

the robot is trying to move, and the fact this movement exists means that the pose we estimate in this iteration of the algorithm will be different from the one we were estimating in the previous iteration.

For the purposes of this thesis, however, we can consider a slightly simplified version of the Bayes filter. Our state (the position of the gas source) can safely be assumed to remain constant for the duration of the search, and none of the actions that the robot takes can affect it in any way. Thus, we are left with the much simpler model represented by the Bayesian network in Figure 2.2B.

Using this model, we can derive the core algorithm of the Bayes Filter, starting with Bayes' Theorem itself:

$$p(x|z_{1:t}) = p(x|z_t, z_{1:t-1}) = \frac{p(z_t|x, z_{1:t-1}) \cdot p(x|z_{1:t-1})}{p(z_t|z_{1:t-1})} \quad (2.3)$$

Exploiting the relations of conditional independence defined by the Bayesian network in Figure 2.2B, we can simplify $p(z_t|x, z_{1:t-1})$ to $p(z_t|x)$. Also, because $p(z_t|z_{1:t-1})$ does not depend on the value of X at all, it is just a constant; it does not modify the probability distribution and serves merely as a normalization term. We will thus omit it from the formula and replace it with a generic normalization term η :

$$p(x|z_{1:t}) = \eta \cdot p(z_t|x) \cdot p(x|z_{1:t-1}) \quad (2.4)$$

Thus, we arrive at the canonical form of the Bayesian filter, where the current state belief is proportional to the sensor model multiplied by the previous belief. Readers who are familiar with this formulation will notice

that we do not consider distinct stages for the prediction and correction of the belief, as we assume the state to be constant in time.

A slightly modified version of the expression, which will be used indistinctly in later chapters, applies Bayes' Theorem again to use the *inverse sensor model* instead:

$$\begin{aligned} p(x|z_{1:t}) &= \eta \cdot \frac{p(x|z_t) \cdot p(z_t)}{p(x)} \cdot p(x|z_{1:t-1}) \\ &= \eta \cdot \frac{p(x|z_t)}{p(x)} \cdot p(x|z_{1:t-1}) \end{aligned} \quad (2.5)$$

This equation can be simplified even further if we assume that the prior of the state is uniform, and thus also part of the normalization term η :

$$p(x|z_{1:t}) = \eta \cdot p(x|z_t) \cdot p(x|z_{1:t-1}) \quad (2.6)$$

The Bayesian filter is the cornerstone of recursive state estimation, and will be used extensively throughout this thesis, but the underlying mechanism (applying Bayes' Theorem and exploiting conditional independence) can be used to fuse information from any arbitrary number of sources. For example, a very similar derivation to this one is used in chapter 5 to combine multiple types of information in a non-recursive manner with a more complex Bayesian network.

For the purposes of this thesis, which focuses on problems that are modeled through discrete random variables, equation 2.6 accurately describes not only the theoretical concept of the Bayesian filter but also its implementation. It should be noted, however, that in many cases (particularly when dealing with continuous variables) the implementation of the abstract Bayes filter needs to be particularized for the type of distribution being used –for example, the Kalman Filter and its extended version are particularizations for working with variables that follow Gaussian distributions.

2.1.3 Uncertainty Quantification

One of the main advantages of using a probabilistic framework is the ability to consider uncertainty when making decisions. This sometimes requires considering each individual possible state and how likely they are. Other times, it is enough to simply obtain a single metric (a number) that serves as a summary of *how much* uncertainty our belief carries. We will briefly present some of these metrics, which will be relevant in later chapters of the thesis. For the sake of brevity, we will only consider discrete random variables with finite support, as they are the only ones considered for the




Figure 2.3: Illustration of entropy and variance as measures of uncertainty in GSL. Both of these probability distributions (left and right) describe the location of the gas source within the limits of the environment. Both distributions have the same entropy, but the second case has a much higher variance. This is because variance also considers the similarity (i.e. distance) between each of the considered values.

contents of the thesis, but the concepts discussed here are all generalizable to any random variable.

One of the most commonly used metrics of uncertainty of a probability distribution is Shannon’s entropy (H) [26]:

$$H(X) = - \sum_x p(x) \cdot \log(p(x)), \quad (2.7)$$

where it must be noted that in the case in which $p(x) = 0$, it is conventional to treat $0 \cdot \log(0)$ as equal to 0, since

$$\lim_{y \rightarrow 0^+} y \cdot \log(y) = 0 \quad (2.8)$$

Entropy has multiple subtly different interpretations according to the field of study that the author favors. For example, an information theorist would talk about the size of the optimal encoding for a message, which would leave you and I confused. For our purposes, however, it can simply be understood as a measure of how “spread out” the probability of a distribution is over the set of possible outcomes. When $p(x) = 1$ for a specific case $X = x$, and all other possible states have 0 probability, the entropy of the distribution is 0. When all possible states are equiprobable, the entropy is maximized (although the exact value of the entropy will depend on the number of values that X can take).

Another common measure of uncertainty is the variance:

$$Var(X) = \sum_x p(x) \cdot (x - \mu)^2, \quad (2.9)$$

where μ is the expected value of X :

$$\mu = \sum_x p(x) \cdot x \quad (2.10)$$

Unlike entropy, which simply considers how many values a variable can take and how probable each of them are, variance also accounts for how *similar* the possible values are to each other (see Figure 2.3). While this additional information can be in many cases be advantageous, it also limits the applicability of variance. Since it requires defining the “distance” between the different values that the variable can take, variance is only applicable to numerical variables.

2.1.4 Information Gain

An important idea that will be mentioned multiple times throughout the thesis, and which is a fundamental part of exploiting uncertainty quantification, is the estimation of information gain. When deciding what to do next (where to move to take the next measurement), a promising idea is to always move to wherever the most information is expected to be gained. That way, we can reduce our uncertainty as quickly as possible and arrive at the correct estimation of the state.

Although this idea can be somewhat problematic, as it requires estimating what observations will be received in the future, it is a very commonly discussed idea in the field of GSL, popularized by the Infotaxis [27] algorithm. Many works (including those presented in this thesis) are published with a discussion of how the idea of information maximization can be applied to them, and proposing a related movement strategy. Even though a more extensive discussion of these concepts will be provided when relevant, we will briefly introduce here the most important metrics of information gain.

Conditional Entropy [26] builds on the now familiar concept of Shannon’s Entropy. Specifically, it measures the expected value of the entropy of the state variable given a value of the observation:

$$H(X|Z) = - \sum_z p(z) \cdot \sum_x p(x|z) \cdot \log(p(x|z)) \quad (2.11)$$

Note that this, being an expected value, does not assume a specific value for the observation, but instead relies on having a probability distribution for the values of the observation. Conditional entropy can be used to define the concept of Mutual Information [26], which is a measure of how much information one can expect to gain about a variable X by determining the value of another variable Z :

$$I(X; Z) = H(X) - H(X|Z) \quad (2.12)$$

Another related measure is the Kullback-Leibler Divergence [26] (often noted KLD or D_{KL}), which is also known as relative entropy:

$$KLD(p||q) = \sum_x p(x) \cdot \log\left(\frac{p(x)}{q(x)}\right) \quad (2.13)$$

Unlike mutual information, KLD does not relate two random variables, but instead two *distributions*. It is a measure of how different a probability distribution $p(x)$ is from another distribution $q(x)$. While the definition of KLD does not require that any specific relation exists between the two distributions, its application to estimating information gain most often involves p being the current state belief and q being the posterior that would result after applying an iteration of the Bayesian filter with a specific observation. Formally, this could be expressed:

$$KLD(p_X||p_{X|z}) = \sum_x p(x) \cdot \log\left(\frac{p(x)}{p(x|z)}\right) \quad (2.14)$$

It is important to highlight that, unlike mutual information, using the KLD like this does not provide an expected value of information gain. It considers only one specific value for the future observation. If we wanted to consider the probability distribution of the measurements, we could of course take the expected value of the KLD itself:

$$\mathbb{E}[KLD(p_X||p_{X|z})] = \sum_z p(z) \cdot \sum_x p(x) \cdot \log\left(\frac{p(x)}{p(x|z)}\right) \quad (2.15)$$

2.2 Gas Dispersion Models and Airflow Estimation

The problem of gas source localization is closely related to the field of fluid dynamics, if not always in terms of specific techniques, at least on a conceptual level. Many of the state-of-the-art computational fluid dynamics (CFD) techniques require an amount of computational effort that is simply unfeasible for an autonomous robot trying to solve the problem of source localization *on-line*. Still, if we frame the problem of source localization in the terms used in the previous section, the relation between *observations* (sensory readings) and *state* (the location and/or parameters of the source) is necessarily defined by a model describing the dispersion of the gas throughout a given environment.

A gas dispersion model, in its most abstract version, is simply a set of equations and algorithms that, given a source location, predict the distribution

of gas that will exist in a particular environment after the source has been active for a certain amount of time. This type of prediction inevitably requires more information about the source than just its position (parameters like the release rate are necessary as well).

It is often the case that, for the purposes of GSL, very simplified dispersion models are used (for example, the Gaussian Plume model [28]). These models only really offer any degree of accuracy under specific environmental conditions (constant, homogeneous wind; absence of obstacles) which makes them unfit for many prospective applications. The applications that we discuss in this thesis, in particular, prove rather challenging to simple plume models: indoors environments most often feature large obstacles and spaces that are divided into separate rooms, with the airflow drastically changing direction from one point of the environment to another.

It is important to note that the specific dispersion model that relates observation and state for a given source localization technique does not always need to be explicitly solved or used to simulate the dispersion. Any relation one establishes between the sensor reading at one specific point in the map and the probability of the possible source locations carries implicit assumptions about the dispersion model. It is reasonable to argue that even bio-inspired, reactive navigation methods like Surge-Cast or Gradient Climbing, which do not perform any sort of probabilistic estimation of the source location, rely on simple, implied dispersion models (such as the existence of unbroken downwind plumes, or a clear concentration gradient caused entirely by diffusion).

One must conclude from this reasoning that the choice of a dispersion model is an important step in designing a gas source localization algorithm, even when the model is not formally described and merely codifies a loose set of assumptions about the nature of the airflow and the presence of obstacles. There are two works in this field of dispersion modeling that have been of particular importance to the development of the contributions of this thesis: the filament dispersion model [29] and wind estimation through Gaussian Markov Random Fields (GMRF) [30].

2.2.1 Filament Model

The filament model was initially proposed as a simulation model to generate synthetic data with which to test robotic olfaction techniques. It is based on the idea that there are two main mechanisms for the dispersal of gas after it is released from the source: advection and diffusion. Advection is the phenomenon through which the airflow in the environment carries the target gas, moving it downwind. Diffusion is the tendency for the gas to spread out over time, going from covering a small area with high concentration to occupying a large area with low concentration. It should be pointed out that, even though diffusion occurs at the molecular level due to the Brownian

motion of the individual gas molecules, at the scales that we are concerned with it is mostly explained by small-scale turbulence pushing patches of gases apart from each other (turbulent diffusion) [31].

In the filament model, the gas release is represented by a sequence of discrete “puffs” being emitted from the source. Each of these puffs, which need not be released at a constant rate, is composed by a finite number of gas filaments. A filament represents a fixed quantity of gas, distributed in space. The filaments are then pushed by the airflow vector at its current location (which must be provided to the model), causing it to move through the environment due to advection.

Diffusion is simulated in two ways: the filaments that were initially released as part of the same puff tend to drift apart from each other by a small amount of random noise that is added to the airflow vectors, and each individual filament spreads out over time. This second part is most commonly implemented by modeling the filament as a Gaussian distribution of concentration and slowly increasing its variance as the simulation progresses, so that the total amount of gas represented by a single filament remains the same but it is more spread out.

The filament model has been utilized for its original purpose of creating simulations for testing in numerous works [32, 33, 34, 23], including Gaden [35] (see Figure 2.4A), which is a simulation software that was extended in one of the works presented in chapter 6. Its main strengths are its ability to deal with scenarios of arbitrary complexity and its relatively light computational cost. The main disadvantage it carries is the need for a complete map of the airflow to be provided. This is not an insurmountable issue when creating simulations for testing scenarios, as any CFD method could be used to generate a suitable airflow map, but poses an interesting challenge for *on-line* source localization. Because of this, airflow estimation methods like the one discussed in the following section become of great importance for our purposes.

2.2.2 GMRF Wind Estimation

The problem of estimating the map of the airflow vectors in a given environment is not a trivial one. Numerical simulation methods, while able to achieve highly accurate results, are computationally intractable in the context of *on-line* source localization. They also require knowledge about the boundary conditions of the environment which is often not available to the robotic agent.

A different approach is to estimate the airflow map by interpolating the sparse measurements the robot gathers with an anemometer as it explores the environment. This means that, at the start of the search process, no information about the airflow is available, but it removes the need to precisely know the boundary conditions. Of course, simple interpolation is not enough to provide a satisfactory estimation of the airflow maps. It would




Figure 2.4: (A) A screenshot of Gaden [35], which uses the filament model to simulate the dispersion of gases in complex environments. (B) Airflow map reconstructed with the GMRF technique [30].

require the measurements to encompass the entirety of the search area to produce a complete map, and it does not provide a way to handle obstacles and walls, which have a strong effect on the direction and speed of the airflow.

The method presented in [30] provides a way to tackle the problem of reconstructing the airflow map from sparse measurements in a more robust manner than simple interpolation. It uses a probabilistic framework, Gaussian Markov Random Fields (GMRF), to combine all obtained sensory observations and produce a most probable map given this evidence. The method relies on defining a set of *energy functions* for the airflow map which a minimization process is applied to in order to generate a maximum a posteriori estimation of the airflow. These energy functions impose certain constraints on the solution:

- The airflow vector estimated in a certain cell must be similar to the observations taken in that cell.
- The volume of air entering a cell should be equal to the volume exiting it (incompressibility).
- Airflow cannot go into or out of obstacles, only be tangential to them.
- Neighboring cells should have aligned wind vectors, with no sharp direction changes.

Since these energy functions are used in a numerical minimization algorithm, the method is always able to provide a solution that “best fits” the list of requirements, even if some of them are contradictory (for example, multiple wind measurements taken in the same cell with different directions).

Because this method establishes relations between all neighboring cells, it is able to not only interpolate but also extrapolate measurements to the rest of the environment. It also takes obstacles into account through the definition of the energy functions, avoiding interpolation issues and using the geometry of the environment as an additional source of information (see Figure 2.4B).

This method is used by the algorithms presented in chapters 3 and 4 of this thesis as a source of estimated airflow. We discuss it here because we believe it is important to address the issue of airflow estimation and how we have tackled it in our GSL implementations, but the specifics of this method are ultimately not directly relevant to any of our contributions. A different technique for airflow estimation could be used in its place, and the formulation for our proposed GSL algorithms would not be affected by the change.

2.3 Robotic Olfaction Sensors

While we will not concern ourselves too much with gas sensors themselves and how they work, it is worth briefly explaining some of the most common types. That way, readers unfamiliar with the field will later understand some of the discussion about their limitations and how they condition the details of robotic olfaction methods.

Most of the gas sensors in common use today measure the concentration of gas by contact. That is, they can only measure at one point, the current position of the sensor. While some technologies exist that allow for ranged gas sensing, like Tunable Diode Laser Spectroscopy (TDLAS) [36], the use of these sensors carries a complete set of separate challenges that are well outside the scope of this thesis, and we only mention them here for completion.

The point-like sensors can use multiple different technologies to sense the gas. Some of the most common include electrochemical sensors [37], optical sensors [38] and photoionization detectors (PID) [39].

Electrochemical sensors work by having the material of the sensor itself chemically react with the gas it comes into contact with. This causes a change in the properties of the material (usually, the electrical resistance), which can be measured. They are very commonly used due to being cheap to manufacture and very lightweight. Another great advantage of these sensors is that, by altering the chemical composition of the sensor, it is possible to significantly change the response of the sensor to specific gases. For these reasons, they are a common component in *e-noses* and widely used in smell recognition. The main disadvantage of these sensors is their response time. The chemical reactions are not instantaneous, and therefore it can take a long time for the sensor to produce an accurate reading after being exposed to a sudden change in gas

concentration. This is of course a big problem for mobile robotic applications, and the main reason we will generally prefer other technologies.

Optical sensors generally rely on having a small internal chamber with a light emitter and a receiver. The emitted light has a specific frequency that is part of the absorption spectrum of the target gas. By measuring how much of the light intensity is lost in the travel between emitter and receiver, the sensor can estimate the concentration of the target gas inside of the chamber. This mechanism allows for faster response times, which makes them more interesting for our applications.

The sensor we will most commonly rely on, however, is the photoionization detector (PID). A PID works in a similar way to an optical sensor, in that it relies on drawing the air into an internal chamber and emitting electromagnetic radiation of a specific frequency into it (in this case, in the ultraviolet range). Unlike optical sensors, however, the PID does not measure light absorption, but instead the electrical potential caused by the ionization of the gas inside of the chamber. This mechanism is also fast, and allows the device to detect many different gases, as the frequency of the radiation does not need to be tuned to a specific substance's absorption spectrum. This fact is also their main disadvantage, as their low specificity means they are not able to discern which type of gas is being sensed.

Beyond gas sensors, we will also mention that robotic olfaction techniques commonly employ anemometers to gain information about the direction and speed of the airflow. For the indoors scenarios that we consider, ultrasonic anemometers [40] are the most appropriate type. Curiously similar to some of the gas sensing technologies we just discussed, ultrasonic anemometers rely on pairs of emitters and receptors to measure the wind movement. In this case, they measure the time of flight of mechanical (ultrasonic) waves, which is altered by the movement of the air.

Modeling the Relation between Measurements and Source Location

This chapter presents two works that address the problem of source localization by directly modeling the connection between the source location and the sensory observations the robot perceives. The first work presents a new such model and an algorithm based on it, alongside an information-maximization movement strategy. The second work, based on the algorithm presented in the first work, analyzes the problem of source declaration: deciding when the search is finished and giving a final estimation for the source location.

3.1 Introduction

Designing a probabilistic gas source localization algorithm is mostly about finding a model that relates the sensory observations to the location of the source. Such a model will often be an explicitly described dispersion model; that is, a predictive model that generates the expected gas distribution from the location and characteristics of the source. Such is the case in much of the existing literature, using models like the Gaussian Plume to perform Source Term Estimation (STE) [41, 42].

If we want to frame this type of technique in the terms of Bayesian estimation, we would say that the dispersion model defines the calculation of the likelihood function, $p(z|x)$, with z being the observation (a concentration

reading), and x being a vector of the parameters of the dispersion model (source location, release rate, wind direction and speed, *e.t.c.*). With this value calculated, the Bayes Filter algorithm can be used to update the belief about the parameters of the source as described in section 2.1.2.

In this case, the definition of $p(z|x)$ is more than just a *sensor model*, like it would be in the traditional Bayesian formulation. It instead encodes both the sensor's response to being exposed to a concentration of gas, and the prediction of the dispersion model itself. It is not uncommon for STE works to make this distinction explicit, and calculate $p(z|x)$ in a way that accounts for both of these distinct models (sensor/dispersion). For example, in [43], the authors describe a particle-filter algorithm that models the sensor noise as a random variable that is separate from the dispersion model:

$$z = \begin{cases} C(x) + \epsilon & \text{if } D = 1 \\ \epsilon & \text{if } D = 0, \end{cases} \quad (3.1)$$

where $C(x)$ is the concentration predicted by the model, and the binary event D determines whether the sensor is actually responding to that concentration. The probability of D (which is given as an input parameter) then controls how much the dispersion model is trusted when calculating the value of $p(z|x)$.

This approach, of course, requires a suitable dispersion model, which is not trivial for the type of indoors scenarios we consider in this thesis. It also poses certain problems when it comes to generating probabilistic estimations from the dispersion model, as the model itself will usually just predict a deterministic concentration. Tackling this problem requires doing things like we saw in the previous example, where an externally defined parameter is used to control the probability that the model is accurate. Ultimately, this issue can be traced back to the fact that simple dispersion models predict time-averaged concentrations, while reality has unpredictable instantaneous variations due to turbulence [29].

For the works we introduce in this chapter, we considered a different approach. Instead of picking a specific dispersion model and describing the sensor response to calculate $p(z|x)$, we designed an inverse sensor model $p(x|z)$ that defines the probability of the source being in a specific position given an observation.

The definition of this model carries several implicit assumptions about the environmental conditions which we deem acceptable for indoor environments, discussed in more detail in the paper itself. We base the model in the idea that gas is mainly being dispersed through advection. With this assumption, we exploit knowing the geometry of the environment and the locally measured wind direction to reason about the path a specific patch of gas must have followed before reaching the robot.



3.2. CONTRIBUTION

With this model that lets us predict the source location from a measurement, we then apply the formulation of the discrete Bayes filter presented in Eq 2.6 to accumulate information recursively.

3.2 Contribution

In this chapter, we first present a novel method for probabilistic gas source localization in indoors environments based on defining an inverse sensor model that implicitly describes the gas dispersion process.

Unlike the algorithms based on plume models, this proposal is able to deal with environments that feature obstacles and multiple rooms, thus being a better fit for indoor environments. The formulation for the estimation of the source is based on a discretization of the environment into a grid of cells and the recursive updating of the probability distribution through Bayesian filtering.

The performance of the proposed method is evaluated under different environmental conditions, in both simulated and real environments. Two different movement strategies are used in this evaluation: a naive exploitation strategy that just moves in the direction towards the most likely source position, and an information maximization strategy. We recommend reading section 2.1.4 for context on the idea of information maximization as a navigational strategy.

On the second work included in this chapter, we present a study on the use of multiple methods to perform source declaration. Employing the previously presented algorithm for the search task, the focus is now on methods of uncertainty quantification (explained in 2.1.3) to declare this search task complete. This work uses the previously presented algorithm for the search, and discusses several of the methods of uncertainty quantification that were explained in 2.1.3. This article also presents experimental results to complement the theoretical analysis of the implication of choosing different methods and uncertainty metrics to perform source declaration.

3.A Information-driven Gas Source Localization Exploiting Gas and Wind Local Measurements for Autonomous Mobile Robots

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

Gas source localization (*GSL*) by an olfactory robot is a research field with a great potential for applications but also with numerous unsolved challenges, particularly when the search must take place in realistic, indoor environments that feature obstacles and turbulent airflows.

In this work, we present a new probabilistic *GSL* method for a terrestrial mobile robot that revolves around the propagation of local estimations throughout the environment. By exploiting the geometry of the environment as the basis for this propagation, we avoid relying on analytical dispersion models, eliminating the need to assume controlled environmental conditions.

Simulated and real experiments are presented in different indoor environments featuring multiple rooms and turbulent flows, demonstrating the suitability of our approach for locating the emitting gas source.

Published in
IEEE Robotics and Automation Letters (RAL)
2021

DOI: 10.1109/LRA.2021.3057290

3.B On Gas Source Declaration Methods for Single-Robot Search

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

Source declaration, along with plume finding and plume tracking, is one of the needed processes for gas source localization (GSL). It is a fundamental part of the search, since it is responsible to decide whether the gas source has been found, and also to pinpoint its location. Despite its importance, source declaration is often ignored in most of the GSL research, the criteria for termination being selected in a seemingly arbitrary manner, or even not being discussed at all. A clear example of this is the large number of experiments in the literature that are declared concluded whenever the robot manages to physically reach the source, without any declaration process. In this work, we seek to provide an overview of the most significant declaration methods that have been used in state-of-the-art research for single-robot GSL, analyzing their strengths and weaknesses. We also provide a preliminary experimental validation of these methods, focusing on how stable their performance is when their input parameters are modified.

Published in the proceedings of the
ISOCS/IEEE International Symposium on Olfaction and Electronic Nose
(ISOEN)

Aveiro, Portugal, 2022

DOI: 10.1109/ISOEN54820.2022.9789673

Source Localization through Probabilistic Gas Mapping

This chapter presents a work that combines the concepts of Gas Source Localization and Gas Distribution Mapping. We present a novel way to perform the gas distribution mapping by replacing the gas concentration with the concept of “gas hit probabilities”, which is designed specifically to be convenient for the purpose of GSL. These maps are then used to estimate the source’s probability distribution through real-time simulation with the filament model.

Beyond the specifics of the method itself, the underlying idea that we present is that a map of how the gas has dispersed through the environment offers a clearer connection to the candidate source positions than individual measurements. We also discuss certain implications that this change has on the probabilistic formulation and the recursive belief updating.

4.1 Introduction

In the previous chapter we discussed how defining an inverse sensor model, of the form $p(x|z)$, can implicitly model our assumptions about the gas dispersion. With such a model, one can estimate the probability distribution of the source

4.1. INTRODUCTION

location as new measurements are gathered. There are, however, some issues with this approach.

The first issue is best understood when described in an informal manner. Simply put, a single gas/wind measurement offers too little information. There are too many possible source configurations that could produce a distribution such that a robot in this position would receive this measurement. Since our model needs to use a single measurement to produce its estimation, there is a great degree of uncertainty associated to this prediction, specially in complex environments.

The second issue requires thinking about the probabilistic formulation itself. In the derivation of Bayes' Theorem, there is a step where it is assumed that one measurement is independent of another given the value of the state:

$$p(z_t|x, z_{1:t_1}) = p(z_t|x) \quad (4.1)$$

We need to consider this assumption carefully, as it has an important effect on the equation. Roboticists reading this text might be reminded of the Markov Property [44], which is a form of conditional independence that is often mentioned in such derivations.

A system satisfies the Markov Property if the probability of a certain event occurring (in this case, the observation) is only dependent on the current state, and not on any past history of the system. This is always true when the state is being observed directly.

As a trivial example (explained in more detail in [25]), consider a sensor that detects whether a door is open or closed. The probability that the sensor reports the value *open* depends only on whether the door actually is open or not, and the reliability of the sensor itself. If we know that the door is open now, then telling us that it was closed an hour ago does not modify the probability distribution for the sensor's response. Similarly, knowing *what the sensor reported* in the past offers no new information about its current response –provided the current state is known.

In our case, however, the state is not observed directly. Instead, as we have already mentioned, we connect observations and state through some model. This means that the question of whether our system is a Markov process is not a trivial one. In fact, it can be argued that this assumption is not really met, which can be seen in the following case.

Consider that we know the location of the source, x . This source location could create many different gas distributions, accounting for the airflow, the unknown release parameters and the inherent uncertainty of the dispersion model. Thus, we could define a probability distribution $p(z|x)$ that tells us how likely it is that we will find gas at a particular location, knowing where the source is located.

Now, let's assume that we take a measurement at one point of the map, and the sensor reports that there is a high concentration of gas there. This implies

that, out of the many possible gas distributions that the source could cause, we are dealing with one which covers this position. Should we now estimate the probability of finding gas here to be the same as before the first measurement? Formally:

$$p(z_t|x, z_{1:t_1}) \stackrel{?}{=} p(z_t|x) \quad (4.2)$$

It is intuitive that the answer is no. Our definition of state (the source location) is not complete, and thus we should not assume the Markov Property. Otherwise, we can imagine that taking many measurements in the same position would make us as certain about the source position as taking measurements all over the map (this is not a formal equality and would of course depend on the specifics of the observation model, but it illustrates the point).

It is this very problem which lead us to explore the idea of integrating mapping into the source localization process. By defining a “is there gas in this cell” intermediate variable, we can provide a more complete definition of state, such that different measurements become conditionally independent and there is a limit to how much information about the source can be derived from re-observing the same part of the map. The actual formulation itself is explained in detail in 4.A.

4.2 Contribution

Our main contribution in this chapter is an integration of Gas Distribution Mapping and Gas Source Localization. We describe some of the problems that derive from the simpler formulation introduced in the previous chapter, and explain the theoretical reason for these issues. We then propose an alternate formulation that provides a more complete definition of state to address the problem of conditional dependency.

The article that constitutes the remainder of this chapter applies these ideas to design a new source localization algorithm based on gas mapping. We propose a specific way to implement the ideas discussed so far through the new concept of *hit probabilities*, and then exploit the created maps to find the gas source. This is done by using a modified version of the filament model (see Sec. 2.2.1) which is used to carry out *on-line* simulations as the robot explores the environment. The article includes an extensive discussion of the computational considerations and the intelligent allocation of resources necessary to make this method feasible.

4.A Robotic Gas Source Localization with Probabilistic Mapping and Online Dispersion Simulation

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

Gas source localization (*GSL*) with an autonomous robot is a problem with many prospective applications, from finding pipe leaks to emergency-response scenarios. In this work, we present a new method to perform *GSL* in realistic indoor environments, featuring obstacles and turbulent flow. Given the highly complex relationship between the source position and the measurements available to the robot (the single-point gas concentration, and the wind vector) we propose an observation model that derives from contrasting the online, real-time simulation of the gas dispersion from any candidate source localization against a gas concentration map built from sensor readings. To account for a convenient and grounded integration of both into a probabilistic estimation framework, we introduce the concept of probabilistic *gas-hit* maps, which provide a higher level of abstraction to model the time-dependent nature of gas dispersion. Results from both simulated and real experiments show the capabilities of our current proposal to deal with source localization in complex indoor environments.

Published in
IEEE Transactions on Robotics (TRO)
2024
DOI: 10.1109/TRO.2024.3426368

Using Semantic Information to Leverage Additional Sensory Data

This equation-heavy chapter presents a probabilistic framework developed to integrate semantic knowledge into the source localization process. By formally defining a connection between types of objects and types of gas they may release, we are able to use a great amount of information to reason about the location of the gas source that we would otherwise be unable to exploit.

5.1 Introduction

A great deal of the problems and difficulties discussed in the previous chapters stem from the complexity of gas dispersion itself. Even with large amounts of information, solving the problem of gas dispersion backwards is an inherently challenging task. However, we are most definitely not helped by how little information is available to us. The vast majority of gas sensors offer single-point concentration readings, which may not even be accurate due to their long response times and their dependence on secondary variables like temperature and humidity.

Therefore, it is interesting to consider the possibility of obtaining information from additional sources to compensate for the limitations of the gas sensors. Moreover, depending on the type of information we include, it might be possible to establish its relation to the source location in a way that

5.2. CONTRIBUTION

does not only rely on our ability to exploit fluid dynamics and the mechanics of gas dispersion.

To explore this idea, this chapter deals with the subject of semantics [45], which is not trivial to define. In very abstract terms, semantics defines the relations between concepts. In the field of intelligent robotics, we use the term to refer to the ability to understand aspects of the environment and how they relate to each other. Problems like determining which type of room a robot is in (*a kitchen? a bathroom?*) [46, 47], or deciding whether two observations of an object actually correspond to the same object (the same *instance*) or to two objects of the same type [48] are good examples of what we mean when we talk about semantics.

For our purposes, the simplest version of semantic scene understanding that would be of relevance is a pre-existing map of all the objects in the environment and the category they belong to. With the semantic information (an *ontology* [49]) to tell our robot which types of objects emit the type of gas we are searching for, this map of objects could be used directly to estimate the source location without the need for a single gas measurement. Of course, if we *also* have gas measurements, we can combine both sources of information to produce a better estimation through our usual Bayesian techniques. This combined estimation can allow us to obtain greater precision in the final declared source position, by identifying the specific object that is emitting the gas.

As an example of a practical case where semantic knowledge could help locate the source of the gas, consider an accidental release of a particular chemical in a factory. Right from the start, knowledge about which areas/rooms the gas bottles are usually kept in allows us to focus our efforts into exploring those areas, since they are much more likely to contain the gas source than the rest of the factory. Once the robot is exploring, if it manages to recognize a gas bottle from an image captured by its vision system, that can allow it to quickly pinpoint the location of the source with higher precision than would be possible if relying entirely on olfactory data.

5.2 Contribution

The work we present in this chapter shows a formal way to combine semantic information and olfactory data through a unified probabilistic framework. The formulation itself is kept as abstract as possible, since the intent is not to propose a specific GSL algorithm, but instead a template through which any probabilistic GSL method may be extended with semantic information.

This formulation is then used to prove that it is formally sound to estimate the source probability from olfactory measurements and from semantic data

as two separate processes, and then combine the results. We also propose a specific method to probabilistically estimate the source location from semantic data using a dense object category map and an ontology.

Finally, we discuss the subject of expected information gain and how it can be calculated for an algorithm that follows the structure outlined by our formulation.

5.A PSGSL: A Probabilistic Framework Integrating Semantic Scene Understanding and Gas Sensing for Gas Source Localization

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

Semantic scene understanding allows a robotic agent to reason about problems in complex ways, using information from multiple and varied sensors to make deductions about a particular matter. As a result, this form of intelligent robotics is capable of performing more complex tasks and achieving more precise results than simpler approaches based on single data sources. However, these improved capabilities come at the cost of higher complexity, both computational and in terms of design. Due to the increased design complexity, formal approaches for exploiting semantic understanding become necessary.

We present here a probabilistic formulation for integrating semantic knowledge into the process of gas source localization (GSL). The problem of GSL poses many unsolved challenges, and proposed solutions need to contend with the constraining limitations of sensing hardware. By exploiting semantic scene understanding, we can leverage other sources of information, such as vision, to improve the estimation of the source location. We show how our formulation can be applied to pre-existing GSL algorithms and the effect that including semantic data has on the produced estimations of the location of the source.

Yet to be published

DOI: 10.48550/arXiv.2501.12812

Development of Robotic Olfaction Research Tools

This chapter consists of two distinct works which present new research tools. Specifically, we present a large dataset of realistic gas dispersion simulations, and an integration of visual and olfactory simulation tools. These tools, made publicly available to the research community, greatly simplify the process of implementing and testing new robotic olfaction techniques.

6.1 Introduction

The field of robotic olfaction in general, and source localization in particular, is (at the time of writing this thesis) a relatively small one. Not many people are doing research on this topic, which means that progress can be slow and difficult. Science is, after all, a collaborative effort, with each new piece of research building on top of everything that has come before.

Probably the clearest example of how a specific work can help the rest of the field move forwards is the development of research tools. Software and datasets that are designed specifically to facilitate the research process can be a valuable tool for all other researchers, simplifying the process of implementing and testing their ideas. It also makes the field more accessible and thus more appealing to new researchers, and allows their effort to be focused onto understanding the problem at hand and coming up with new

6.2. CONTRIBUTION

solutions, rather than on dealing with complex technologies and building their own tools.

The field of robotic olfaction is indeed an interesting case as far as this idea goes, due to the high complexity of the basic process of simulation. Simulating scenarios in which to test a new idea is often one of the first steps in research, and saves a great deal of time compared to jumping directly into the complexity of real-world experiments. However, CFD simulation with numerical models is notoriously difficult to set up, with not only many parameters that control the behavior of the model itself, but also strict requirements on the quality of the input data (mostly, the 3D meshes used to define the geometry of the environment).

6.2 Contribution

In this chapter, we present two works that describe novel research tools that are made publicly available to the robotic olfaction research community. Both of them are based on Gaden [35], our in-house gas dispersion simulator, which uses CFD airflow and the filament model to generate realistic gas dispersal simulations.

The first work is a large dataset of CFD and gas dispersion simulations taking place in realistic, complex environments. The environments themselves are extracted from the Robot@VirtualHome Dataset [50], and which contains 3D models of real houses.

The 3D meshes of these environments are preprocessed to conform to the standards of quality required by CFD simulation, and multiple completed simulations are made available to the users of the dataset. The dataset also includes code, configuration files, and intermediate data required to re-generate the simulations, in case the users want to make any modifications without having to set up the simulations from scratch.

The second work is an integration of Gaden and the Unity Engine [51], with the main intention of enabling research on multi-sensor techniques, such as the ones discussed in chapter 5. This integration includes a form of ray-marching based realistic volumetric renderization for working with plumes of visible substances, and a native re-implementation of part of the Gaden functionality inside of Unity to achieve the best possible performance.

6.A VGR Dataset: A CFD-based Gas Dispersion Dataset for Mobile Robotic Olfaction

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

There are many potential applications for an autonomous robotic agent capable of sensing gases in the environment, from locating leaks in pipes to monitoring air quality. However, the current state of the art in the field of robotic olfaction is not mature enough for most real-world applications. Due to the complexity of gas dispersion phenomena and the limitations of sensors, a great deal of research into the development of techniques and algorithms remains necessary. A very important part of this research is thorough experimentation, but carrying out robotic olfaction experiments is far from trivial. Real world experiments are usually limited to very simplified, wind-tunnel-like environments, as it is impossible to closely monitor or control the airflow in more complex scenarios. For this reason, simulation with CFD offers the most plausible alternative, allowing researchers to study the behavior of their algorithms in more challenging and complex situations. This work presents a CFD-based gas dispersion dataset composed of 120 cases generated under variable environmental conditions, taking place in 30 realistic and detailed models of real houses. All the data is made available in multiple formats, and is directly accessible through ROS, to permit easy integration with other robotic tools.

Published in
Journal of Intelligent & Robotic Systems
2023
DOI: 10.1007/s10846-023-02012-z

6.B A Simulation Framework for the Integration of Artificial Olfaction into Multi-Sensor Mobile Robots

Pepe Ojeda, Javier Monroy and Javier Gonzalez-Jimenez

The simulation of how a gas disperses in a environment is a necessary asset for the development of olfaction-based autonomous agents. A variety of simulators already exist for this purpose, but none of them allows for a sufficiently convenient integration with other types of sensing (such as vision), which hinders the development of advanced, multi-sensor olfactory robotics applications.

In this work, we present a framework for the simulation of gas dispersal and sensing alongside vision by integrating GADEN, a state-of-the-art Gas Dispersion Simulator, with the Unity 3D, a video game development engine that is used in many different areas of research and helps with the creation of visually realistic, complex environments. We discuss the motivation for the development of this tool, describe its characteristics, and present some potential use cases that are based on cutting-edge research in the field of olfactory robotics.

Published in
MDPI Sensors
2021
DOI: 10.3390/s21062041

Conclusions and Future Work

This thesis has addressed the problem of gas source localization with a mobile robot from the perspective of probabilistic robotics. We have discussed the reasons why this approach is particularly well suited for dealing with the uncertainty and imprecisions that are inherent to this problem: we not only have the issues of sensor reliability that are common to much of robotics, but we also must deal with simplified models to describe the phenomenon being studied.

We contributed three new probabilistic GSL techniques for indoors environments, which are based on an attempt to formalize our intuitions about the subject in the terms of Bayesian estimation:

- In the first method, we directly modeled the connection between the measurements and the source location, relying on information about the environment geometry to bridge the gap between single-point sensor readings and global estimations.
- In the second method, we attempted to correct some unintended effects of the simpler first formulation by introducing an intermediate variable that mediates the relation between source and measurement. This “gas presence” variable is directly related to the problem of gas distribution mapping, although we propose a new abstraction of “hit probabilities” instead of working with a map of gas concentration.

- The third technique is not a specific GSL algorithm, but instead a generalized probabilistic formulation that describes how semantic knowledge can be integrated into the source localization process. By introducing semantic understanding, we are able to leverage the information gathered by other, seemingly unrelated sensors like cameras. We also relate this approach to the problem of smell recognition, as estimating which type of gas is being sensed becomes an important part of utilizing the semantic information.

Besides these GSL techniques, we also presented other works. We discussed the under-studied problem of source declaration and different methods that can be used for it, highlighting their advantages and disadvantages. The findings from this study were applied in all later works. We also presented some newly developed research tools, including a dataset of gas dispersion simulations and a simulation framework for the development of complex, multi-sensor techniques such as the semantics approach we previously discussed.

As a final thought on the findings presented here, we conclude that the probabilistic approach is a solid way to tackle source localization. Considering more complex versions of the problem by combining it with distribution mapping or semantic knowledge can be challenging, but it also allows us to describe the relevant concepts and their relations in a more complete manner, giving us access to additional information.

Future Work

Of course, the field of robotic olfaction is still far from maturity. Our proposals here are only one small step in advancing the research, and much work remains to be done. We will now briefly discuss some promising avenues for future research, and some of the details that we did not study completely in the works presented here.

3D Source Localization

Almost all the work we have presented here ignores the existence of verticality. This fabled *third dimension* presents numerous unsolved challenges when it comes to airflow estimation and gas dispersion models, which currently limits us to working in 2D. Although there are other works which do consider 3D GSL with other approaches, it would be interesting to extend the techniques we presented here to three dimensions as well. The probabilistic formulations of our methods should be easily generalizable to this more complex case, but it is

to be expected that new challenges will appear in the process of implementing these methods.

Information Gain and Movement Strategies

While we have dedicated time and effort to tackling the subject of information gain optimization, it is a very complex topic, and much remains to be said about it.

Firstly, we have only considered greedy optimization being performed one step at a time, which is a reasonable approach given computational considerations, but is of course myopic. Studying strategies that consider more than a single movement is not a trivial task, but given the interest that the GSL research community dedicates to infotactic searching, one should expect future works to explore this problem.

There is also the subject of which metric to use to estimate the information gain. In chapter 4 we implemented an efficient method for estimating the information gain based on a modified version of variance. The problem with such ad-hoc methods is that, while they allow us to tell if one movement offers more information than another, they do not provide a reliable absolute quantification, with clear units. This makes our algorithms more difficult to integrate with other sensing modalities, which require their information gain to be calculated in a different way.

It is for this reason that the work we presented in chapter 5 proposes the use of Mutual Information, a formally defined metric with clear units. However, efficient calculation of such metrics can prove very challenging, specially as multiple interacting random variables are considered. The issue of information gain is thus as much about computational efficiency as it is about formal analysis of the related concepts.

Reducing Required Information

The olfaction methods in chapters 3 and 4 rely on pre-existing knowledge about the environment geometry. While there are some applications where this is not a problem, since the environment is known in advance (imagine a robot that patrols a factory monitoring for gas leaks), there are other cases where it is not a reasonable assumption (a search and rescue mission).

Removing such assumptions is not trivial, as the geometry information is a crucial part of the airflow estimation and the dispersion models used. Developing new GSL algorithms that are able to work with an SLAM method that dynamically updates the map and includes uncertainty would go a long way towards making these sort of techniques applicable in the real world.

Bibliography

- [1] Emily G Snyder, Timothy H Watkins, Paul A Solomon, Eben D Thoma, Ronald W Williams, Gayle SW Hagler, David Shelow, David A Hindin, Vasu J Kilaru and Peter W Preuss. The changing paradigm of air pollution monitoring. *Environmental science & technology*, 2013.
- [2] Víctor González, Javier Godoy, Patricia Arroyo, Félix Meléndez, Fernando Díaz, Ángel López, José Ignacio Suárez and Jesús Lozano. Development of a Smartwatch with Gas and Environmental Sensors for Air Quality Monitoring. *Sensors*, 2024.
- [3] Shin-Juh Chen, David C Hovde, Kristen A Peterson and André W Marshall. Fire detection using smoke and gas sensors. *Fire Safety Journal*, 2007.
- [4] Julian W Gardner, Hyun Woo Shin and Evor L Hines. An electronic nose system to diagnose illness. *Sensors and Actuators B: Chemical*, 2000.
- [5] Carlos Sánchez-Vicente, José Pedro Santos, Jesús Lozano, Isabel Sayago, José Luis Sanjurjo, Alfredo Azabal and Santiago Ruiz-Valdepeñas. Graphene-Doped Tin Oxide Nanofibers and Nanoribbons as Gas Sensors to Detect Biomarkers of Different Diseases through the Breath. *Sensors*, 2020.
- [6] Yuxin Xing and Julian W. Gardner. Classification of Urine Odour Using Machine Learning Methods. In *2022 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. 2022.
- [7] Rafa Aparicio, Sílvia Rocha, Ivonne Delgadillo and Maria Morales. Detection of Rancid Defect in Virgin Olive Oil by the Electronic Nose. *Journal of agricultural and food chemistry*, 2000.
- [8] Pablo Cano Marchal, Chiara Sanmartin, Silvia Satorres Martínez, Juan Gómez Ortega, Fabio Mencarelli and Javier Gámez García. Prediction of Fruity Aroma Intensity and Defect Presence in Virgin Olive Oil Using an Electronic Nose. *Sensors*, 2021.
- [9] Silvia Satorres Martínez, Sergio Illana Rico, Pablo Cano Marchal, Diego Manuel Martínez Gila and Juan Gómez Ortega. Zero Defect Manufacturing in the Food Industry: Virgin Olive Oil Production. *Applied Sciences*, 2022.
- [10] Hingman To and Naikong Fong. Investigation of the performance and improvement of optical smoke detectors. *Procedia engineering*, 2013.
- [11] Pengfei Jia, Xiaoyu Li, Min Xu and Lin Zhang. Classification techniques of electronic nose: a review. *International Journal of Bio-Inspired Computation*, 2024.

- [12] Javier Monroy and Javier Gonzalez-Jimenez. Real-Time Odor Classification Through Sequential Bayesian Filtering. *16th International Symposium on Olfaction and Electronic Nose (ISOEN)*, 2015.
- [13] JW Gardner. *Electronic Noses, Principles and Applications*. Oxford University Press, 1999.
- [14] Diclehan Karakaya, Oguzhan Ulucan and Mehmet Turkan. Electronic Nose and Its Applications: A Survey. *International Journal of Automation and Computing*, 2020.
- [15] Callum Rhodes, Cunjia Liu and Wen-Hua Chen. Structurally Aware 3D Gas Distribution Mapping Using Belief Propagation: A Real-Time Algorithm for Robotic Deployment. *IEEE Transactions on Automation Science and Engineering*, 2024.
- [16] Andres Gongora, Javier Monroy and Javier Gonzalez-Jimenez. Joint Estimation of Gas & Wind Maps for Fast-Response Applications. *Applied Mathematical Modelling*, 2020.
- [17] Andres Gongora, Javier Monroy, Faezeh Rahbar, Chiara Ercolani, Javier Gonzalez-Jimenez and Alcherio Martinoli. Information-Driven Gas Distribution Mapping for Autonomous Mobile Robots. *MDPI Sensors. Special Issue Robotics for Environment Sensing*, 2023.
- [18] Bingxu Li and Wenjian Cai. A novel CO₂-based demand-controlled ventilation strategy to limit the spread of COVID-19 in the indoor environment. *Building and Environment*, 2022.
- [19] Matteo Reggente and Achim J. Lilienthal. Using Local Wind Information for Gas Distribution Mapping in Outdoor Environments with a Mobile Robot. In *Proceedings of IEEE Sensors*. 2009.
- [20] Javier G. Monroy, Jose-Luis Blanco and Javier Gonzalez-Jimenez. Time-variant gas distribution mapping with obstacle information. *Autonomous Robots*, 2016.
- [21] Adam Francis, Shuai Li, Christian Griffiths and Johann Sienz. Gas source localization and mapping with mobile robots: A review. *Journal of Field Robotics*, 2022.
- [22] Xin-xing Chen and Jian Huang. Odor source localization algorithms on mobile robots: A review and future outlook. *Robotics and Autonomous Systems*, 2019.
- [23] Yaqub A. Prabowo, Bambang R. Trilaksono, Egi M. I. Hidayat and Brian Yulianto. Integration of Bayesian Inference and Anemotaxis for Robotics Gas Source Localization in a Large Cluttered Outdoor Environment. *IEEE Access*, 2023.

BIBLIOGRAPHY

- [24] Joel H Ferziger and Milovan Perić. *Computational methods for fluid dynamics*. Springer, 2002.
- [25] S. Thrun, W. Burgard and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
- [26] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 2012.
- [27] Massimo Vergassola, Emmanuel Villermaux and Boris I Shraiman. ‘Infotaxis’ as a strategy for searching without gradients. *Nature*, 2007.
- [28] Wm J Veigle and James H Head. Derivation of the Gaussian plume model. *Journal of the Air Pollution Control Association*, 1978.
- [29] Jay A Farrell, John Murlis, Xuezhong Long, WEI Li and Ring T Cardé. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environmental fluid mechanics*, 2002.
- [30] Javier G. Monroy, Mariano Jaimez and Javier Gonzalez-Jimenez. Online Estimation of 2D Wind Maps for Olfactory Robots. *ISOEN 2017 - ISOCS/IEEE International Symposium on Olfaction and Electronic Nose, Proceedings*, 2017.
- [31] Philip JW Roberts and Donald R Webster. Turbulent diffusion. *Environmental Fluid Mechanics-Theories and Application*, 2002.
- [32] Yatai Ji, Yong Zhao, Bin Chen, Zhengqiu Zhu, Yu Liu, Hai Zhu and Sihang Qiu. Source searching in unknown obstructed environments through source estimation, target determination, and path planning. *Building and Environment*, 2022.
- [33] Miao Wang, Bin Xin, Mengjie Jing and Yun Qu. An Exploration-Enhanced Search Algorithm for Robot Indoor Source Searching. *IEEE Transactions on Robotics*, 2024.
- [34] Xin-Kai Zhang, Ji-Gong Li and Jing Yang. Plume Finding with Frontier-Based Planner. In *2024 6th International Conference on Industrial Artificial Intelligence (IAI)*. 2024.
- [35] Javier Monroy, Victor Hernandez-Bennetts, Han Fan, Achim Lilienthal and Javier Gonzalez-Jimenez. GADEN: A 3D gas dispersion simulator for mobile robot olfaction in realistic environments. *Sensors*, 2017.
- [36] Muhammad Asif Arain. *Efficient Remote Gas Inspection with an Autonomous Mobile Robot*. Ph.D. thesis, Örebro University, School of Science and Technology, 2020.

- [37] Ananya Dey. Semiconductor metal oxide gas sensors: A review. *Materials science and Engineering: B*, 2018.
- [38] Jane Hodgkinson and Ralph P Tatam. Optical Gas Sensing: A Review. *Measurement Science and Technology*, 2013.
- [39] JE Lovelock. A photoionization detector for gases and vapours. *Nature*, 1960.
- [40] Vipin Prakash Yadav, Alakh Sinha and Arun Khosla. Design and implementation of ultrasonic anemometer. In *2017 4th International Conference on Power, Control & Embedded Systems (ICPCES)*. 2017.
- [41] Michael Hutchinson, Hyondong Oh and Wen-Hua Chen. A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors. *Information Fusion*, 2017.
- [42] Fei Xue, Hideki Kikumoto, Xiaofeng Li and Ryozo Ooka. Bayesian source term estimation of atmospheric releases in urban areas using LES approach. *Journal of hazardous materials*, 2018.
- [43] Michael Hutchinson, Cunjia Liu and Wen Hua Chen. Information-Based Search for an Atmospheric Release Using a Mobile Robot: Algorithm and Experiments. *IEEE Transactions on Control Systems Technology*, 2018.
- [44] Evgeniĭ Borisovich Dynkin and Evgenij Borisovič Dynkin. *Markov processes*. Springer, 1965.
- [45] Weiyu Liu, Angel Daruna, Maithili Patel, Kartik Ramachandrani and Sonia Chernova. A survey of semantic reasoning frameworks for robotic systems. *Robotics and Autonomous Systems*, 2023.
- [46] J.R. Ruiz-Sarmiento, C. Galindo and J. Gonzalez-Jimenez. Joint categorization of objects and rooms for mobile robots. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2015.
- [47] Shrihari Vasudevan and Roland Siegwart. Bayesian space conceptualization and place classification for semantic maps in mobile robotics. *Robotics and Autonomous Systems*, 2008.
- [48] Jose-Luis Matez-Bandera, David Fernandez-Chaves, Jose-Raul Ruiz-Sarmiento, Javier Monroy, Nicolai Petkov and Javier Gonzalez-Jimenez. LTC-Mapping, Enhancing Long-Term Consistency of Object-Oriented Semantic Maps in Robotics. *Sensors*, 2022.
- [49] David Paulius and Yu Sun. A survey of knowledge representation in service robotics. *Robotics and Autonomous Systems*, 2019.

BIBLIOGRAPHY

- [50] David Fernandez-Chaves, Jose-Raul Ruiz-Sarmiento, Alberto Jaenal, Nicolai Petkov and Javier Gonzalez-Jimenez. Robot@ VirtualHome, an ecosystem of virtual environments and tools for realistic indoor robotic simulation. *Expert Systems with Applications*, 2022.
- [51] Unity 3D Engine. Available Online: <https://www.unity.com> (accessed 15 November 2024).