

RESEARCH ARTICLE

Fast and automated biomarker detection in breath samples with machine learning

Angelika Skarysz^{1*}, Dahlia Salman², Michael Eddleston³, Martin Sykora⁴, Eugénie Hunsicker⁵, William H. Nailon⁶, Kareen Darnley⁷, Duncan B. McLaren⁶, C. L. Paul Thomas², Andrea Soltoggio^{1*}

1 Computer Science Department, School of Science, Loughborough University, Loughborough, United Kingdom, **2** Centre for Analytical Science, School of Science, Loughborough University, Loughborough, United Kingdom, **3** Pharmacology, Toxicology & Therapeutics Unit, University of Edinburgh, Edinburgh, United Kingdom, **4** Centre for Information Management, School of Business and Economics, Loughborough University, Loughborough, United Kingdom, **5** Mathematical Sciences Department, School of Science, Loughborough University, Loughborough, United Kingdom, **6** Edinburgh Cancer Centre, NHS Lothian, Edinburgh, United Kingdom, **7** Clinical Research Facility, Western General Hospital, NHS Lothian, Edinburgh, United Kingdom

* A.Skarysz@lboro.ac.uk (AS); A.Soltoggio@lboro.ac.uk (AS)

**OPEN ACCESS**

Citation: Skarysz A, Salman D, Eddleston M, Sykora M, Hunsicker E, Nailon WH, et al. (2022) Fast and automated biomarker detection in breath samples with machine learning. *PLoS ONE* 17(4): e0265399. <https://doi.org/10.1371/journal.pone.0265399>

Editor: Sarah Ostadabbas, Northeastern University, UNITED STATES

Received: June 20, 2021

Accepted: March 1, 2022

Published: April 12, 2022

Copyright: © 2022 Skarysz et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data files are available from the Dryad database (<https://doi.org/10.5061/dryad.51c59zwb2>).

Funding: The study was partially funded by the EU H2020 TOXI-Triage Project #653409. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Abstract

Volatile organic compounds (VOCs) in human breath can reveal a large spectrum of health conditions and can be used for fast, accurate and non-invasive diagnostics. Gas chromatography-mass spectrometry (GC-MS) is used to measure VOCs, but its application is limited by expert-driven data analysis that is time-consuming, subjective and may introduce errors. We propose a machine learning-based system to perform GC-MS data analysis that exploits deep learning pattern recognition ability to learn and automatically detect VOCs directly from raw data, thus bypassing expert-led processing. We evaluate this new approach on clinical samples and with four types of convolutional neural networks (CNNs): VGG16, VGG-like, densely connected and residual CNNs. The proposed machine learning methods showed to outperform the expert-led analysis by detecting a significantly higher number of VOCs in just a fraction of time while maintaining high specificity. These results suggest that the proposed novel approach can help the large-scale deployment of breath-based diagnosis by reducing time and cost, and increasing accuracy and consistency.

Introduction

A typical human breath sample is thought to contain thousands of volatile organic compounds (VOCs), which are the products of metabolic, catabolic and exogenous exposure processes occurring continuously in the human body [1]. This makes breath a particularly interesting medium for metabolomics, which describes the individuals' specific phenotype and health status by measuring the metabolites present in the biological sample and changes in their expressions [1, 2]. Breathomics [3] (i.e., breath metabolomics) can bring insight into all the metabolic processes in the body and thus provide comprehensive information about the organism's condition, additionally enabling non-invasive and rapid sample acquisition. Breath

analysis has the potential to expand the range of diagnosis platforms for fast and accurate detection of a disease at an early stage, or for metabolic phenotyping, and so contributing to the development of precision medicine and treatment optimisation [3]. Due to such benefits, breathomics is currently an extensively researched area. Over the past few years, studies have applied breathomics for biomarker discovery and presented the relationships among the changes in VOC patterns and different types of diseases, including chronic obstructive pulmonary disease [4], diabetes [5], as well as breast [6], colorectal [7] and lung cancer [8]. Most recently, a study by Ruszkiewicz et al. demonstrated that breath analysis may allow rapid diagnosis of Covid-19 [9].

Accurate detection of the VOCs present in breath, in particular biomarkers related to a specific metabolic change caused by a disease, is essential to obtain a reliable diagnosis. Gas chromatography-mass spectrometry (GC-MS) is a well-known analytical technology that, due to its low limits of detection, orthogonal data structure and molecular structural characteristics, is the gold standard for VOC measurement in breath samples [10]. The sample processed is passed through the GC capillary column, i.e., a very narrow tube lined with a particular chemical material. Each VOC carried in the sample elutes from the GC column after a specific amount of time, called retention time (*RT*), related to its chemical and physical properties but also dependent on the GC column features. That results in a separation of the compounds contained in the mixture of the sample. Subsequently, each eluted VOC is characterised in MS by a mass spectrum (mass-to-charge ratio, *m/z*) of its ion fragmentation. As different VOCs produce different ion fragmentation patterns, an ion pattern enables the identification of the VOC. GC-MS produces a two-dimensional data matrix, known as abundance matrix [11], such as the one shown in Fig 1c. Each column of the abundance matrix represents one ion channel *m/z*, whereas each row contains a mass spectrum delivered at a specific *RT*. GC-MS data is also often visualised compactly as total ion current (TIC) chromatogram (Fig 1b). For more details on GC-MS data, we refer readers to [12].

Deriving a comprehensive and reliable list of VOCs detected in GC-MS breath data is a difficult task. The VOC separation in GC may not fully occur, resulting in an overlapping of co-eluted compounds (i.e., peaks on a chromatogram, Fig 1b), or the ion patterns produced in MS by some VOCs may be similar and thus difficult to distinguish. Moreover, a GC column degrades over time changing instrumentation features and thus causes *RT* shifts, which significantly reduce the reliability of this parameter in VOC identification and requires its conversion into a system-independent constant (i.e., retention index) [13]. Additionally, the data produced by GC-MS are noisy and high dimensional: one single sample may contain over 9 million variables (in our study over 22500 *RT* points by 411 *m/z* channels).

For such complexity, the current breathomics workflow for GC-MS metabolic phenotyping employs various preprocessing steps performed by a trained and experienced chemical analyst. At the core of these operations is a spectral deconvolution, i.e., an extraction of the overlapping co-eluted VOCs along with their mass spectra prints [14]. The GC-MS breath data analysis (Fig 2, top) includes baseline correction, spectral deconvolution, peak detection and feature alignment. These steps collectively enable the clustering and identification of the VOCs [15] for further multivariate statistical analysis [1, 16]. These processes are subject to high variability. Usually, 350 to 500 VOCs are detected in the sample. The dynamic range of the variables may span 10^4 to 10^5 and some of the spectra for the lower abundance VOCs may be incomplete with minor ion fragments below the limits of detection. Variations in exogenous factors mean that different deconvolution approaches may need to be invoked from sample-to-sample. Current state-of-the-art deconvolution-based breathomics methods require analytical expertise and skilled analyst judgement to choose the techniques and parameters settings for processing the data from every sample. In particular, the optimisation of deconvolution needs

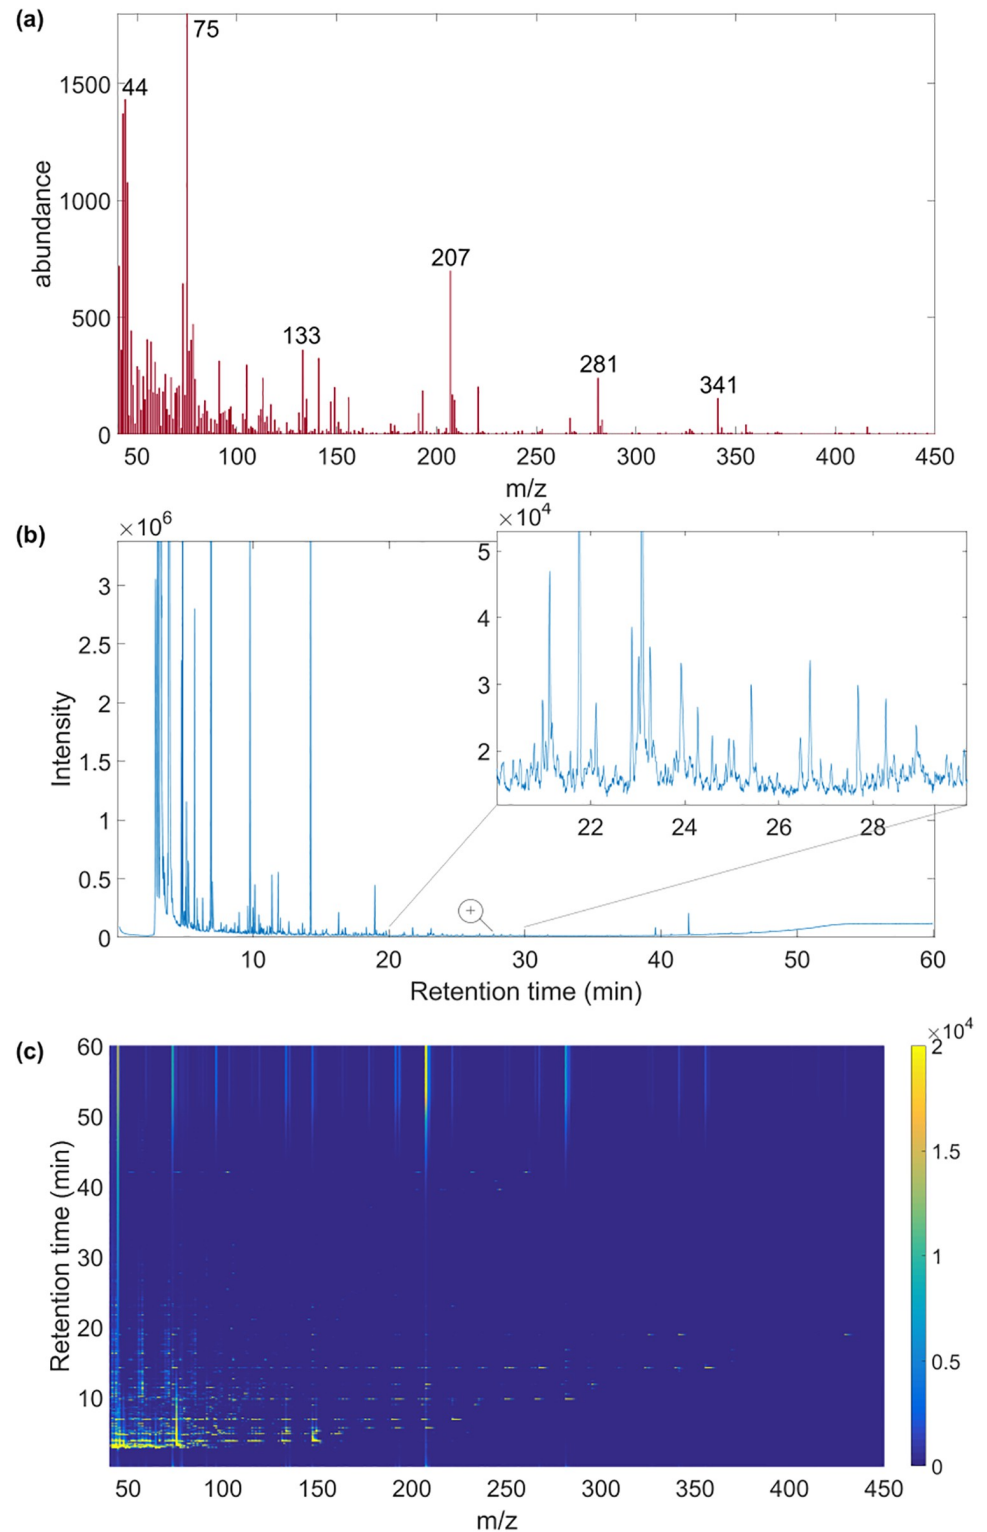


Fig 1. GC-MS breath data. (a) A full mass spectrum corresponds to each retention time (*RT*) point on the chromatogram. (b) The total ion current (TIC) chromatogram plots along the *RT* dimension the cumulative abundance of ions (intensity). Each peak generally represents one specific VOC, although superposition of peaks also occurs [1]. (c) GC-MS abundance matrix presented as a heat map, with the x-axis being the mass-to-charge ratio (*m/z*) and the y-axis the retention time.

<https://doi.org/10.1371/journal.pone.0265399.g001>

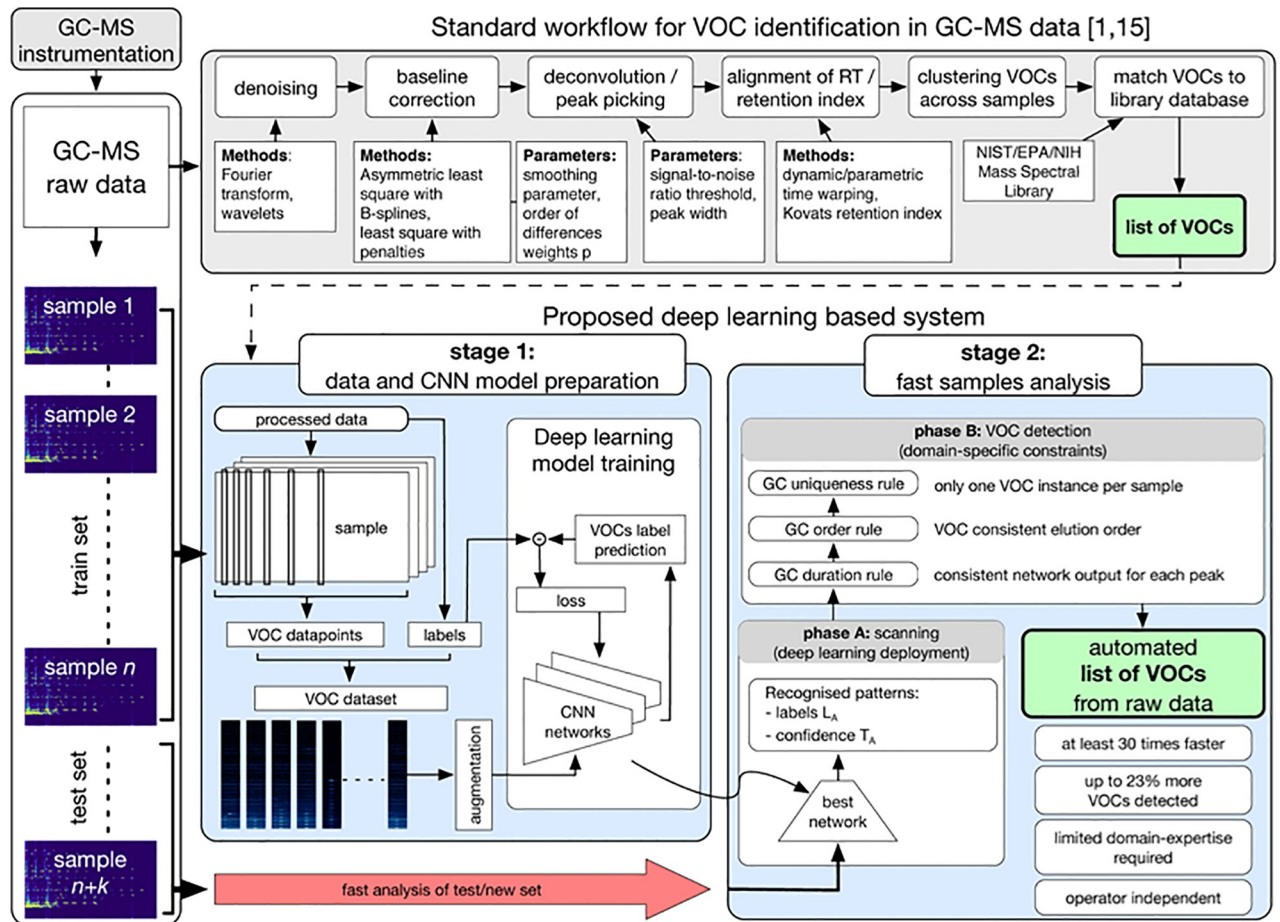


Fig 2. Graphical representation of the current GC-MS state-of-the-art analysis process (grey area) and the novel CNN-based method, stage 1 and stage 2. Both methods require GC-MS data storage (left column) and provide a list of VOCs as output.

<https://doi.org/10.1371/journal.pone.0265399.g002>

to strike the right balance between the inclusion of all the relevant VOCs and the omission of spectral noise and signal artefacts. The operator-subjective nature of the GC-MS data processing, combined with the data complexity, has the potential to introduce errors and variability to results on top of the observed biological variations from one subject to another, and thus the results may not be reproducible [17]. Additionally, expert-operated deconvolution is often labour-intensive time-consuming procedure [18]; the processing time of a single breath sample is estimated by experts as 60 to 120 minutes.

The limitations outlined above call for better algorithms for GC-MS data processing, now possible by exploiting recent advances in machine learning and deep learning. A number of machine learning and deep learning applications in biomedical studies have been reported in the literature [19, 20]. Several studies, such as [4, 7, 21], successfully applied machine learning in the area of breathomics. These reported applications, however, do not process GC-MS breath data directly but analyse a list of selected VOCs, provided by expert-led preprocessing, to classify patients and control group. Consequently, they allow for high-level GC-MS data classification by black-box like decisions without justifications, rather than detection of particular VOCs of interest. Thus, these methods provide limited information about individual's metabolomics condition. Moreover, such data processing (e.g., [4, 7, 21]) involves treating an

entire GC-MS sample as a single data point; in the case of clinical samples, obtaining large-scale datasets desired for machine learning applications may be a challenging and demanding task. On the other hand, to the best of the authors' knowledge, there is a lack of studies applying machine learning directly on raw GC-MS breath data to detect VOCs of interest.

This study introduces a new approach to VOC detection in GC-MS samples: we propose the application of convolutional neural networks (CNNs) [22] to learn to detect VOCs in breath sample automatically and directly from raw GC-MS data, thereby bypassing the current labour-intensive, time-consuming and operator-subjective data preprocessing steps. CNNs [23] are a popular type of deep learning algorithms that is particularly effective in image analysis (e.g., [24–26]). CNNs can autonomously learn useful features directly from low-level data, e.g., pixels [27], and construct high-level features without human intervention. CNNs can also exploit geometrical properties of the data and thus adapt well to image-based tasks [28]. Presenting GC-MS data as abundance matrix enables one to see GC-MS data as an image (Fig 1c). Therefore, we propose a new approach, exploiting recent advancement in CNN applications for image analysis and pattern recognition, to learn to recognise ion patterns directly from raw abundance matrix. Ion patterns derived from specific compounds, although noisy, present unique features that distinguish them. A recognised ion pattern is effectively a recognised VOC, which in turn could be a biomarker of a given physical condition [29].

The promise of such an approach was first demonstrated in [30]. In that study, CNNs were shown to have considerably better performance than support vector machines [31] and shallow neural networks [32]. However, the study reported a high number of false positives and targeted only 8 VOCs, thus leaving open the question of reliability in detection and scalability. Nevertheless, the findings from [30] justify the application of CNN over shallow methods in the considered problem.

The CNN-based approach proposed here has two main stages (Fig 2, bottom); stage 1: data and CNN model preparation, and stage 2: raw GC-MS samples analysis. In stage 1, expert knowledge is exploited to create a dataset of target VOCs and their corresponding ion patterns on the raw GC-MS data. A CNN architecture is then trained on such dataset to learn to recognise ion patterns specific to the target VOCs. In stage 2, new raw breath samples are analysed to detect the targeted VOCs quickly and automatically. Firstly (in phase A, Fig 2), an entire GC-MS sample is scanned by the trained CNN network to provide a list of recognised patterns. Secondly (in phase B), domain-specific constraints are used to derive a final list of detected VOCs.

The robustness and scalability of the proposed CNN-based method were investigated on a dataset of 120 GC-MS samples and the set of 30 target VOCs. A wide range of target compounds generates various challenges: several target VOCs produce similar ion fragmentation patterns, some others have close *RT* positions and thus overlap, which may provide an obstacle to their accurate discrimination. In this study we tested different types of CNNs: VGG16 and VGG-like networks [33], residual neural networks [34] and densely connected convolutional networks [35] with different configurations, to compare their performance and select the most efficient strategy for GC-MS data processing.

The novel CNN-based approach proposed here provides a comprehensive system for fast and automated detection of any set of VOCs in raw GC-MS data, outperforming current expert-led deconvolution-based methods. The analysis of raw GC-MS breath data may reduce human-related errors and has the potential to detect compounds of very low abundances. Consequently, the proposed novel approach showed the ability to correctly detect VOCs missed by the current methods, while improving specificity and significantly reducing processing time. The system may support experts to put much more accurate hypotheses on the VOCs related to the specific health conditions. Moreover, by the significant acceleration of the VOC

detection process, the CNN-based method allows for much quicker hypotheses validation on new GC-MS samples. To the best of the authors' knowledge, this is the first study that proposes a comprehensive system to reveal VOC ion patterns directly from raw GC-MS samples, with high sensitivity and specificity.

Results

GC-MS sample dataset

Breath samples were obtained in a clinical study from 25 participants with different types of cancer receiving radiotherapy treatment. Four breath samples (prior and 1, 3, and 6 hours post radiation) were collected from each participant along with one environmental sample.

Each clinical sample was processed with GC-MS and stored in a data file containing both metadata and the abundance matrix $A \in \mathbb{R}^{R \times 411}$, where 411 is the number of measured m/z channels and R is the number of the mass spectrum measurements performed over retention time. As the instrumental scanning rate was approximately 6.25 Hz, for about one-hour processing it gave approximately 22500 RT points of measurements. The first dimension of the abundance matrix A is called here the RT dimension, whereas the second one is called m/z dimension. Subsequently, all GC-MS data were processed with the current expert-led methods to identify VOCs contained in the sample and their RT positions. This process was completed for 120 clinical samples (see [Methods](#)).

The GC-MS sample dataset was divided into training and testing sets in the proportion 82/38: the training set contains 65 breath samples and 17 environmental samples associated with 17 randomly chosen participants, the testing set contains 30 breath samples and 8 environmental samples associated with remaining 8 participants. Both raw and expert-processed GC-MS samples from the training set were used to generate a VOC dataset (described later). The VOC dataset was applied in stage 1 ([Fig 2](#)) for the CNN model training, selected through cross-validation (CV) [36]. The testing GC-MS sample set was used for the assessment of the proposed CNN-based system performance. Precisely, the 38 raw samples from the testing set were used as inputs for automated VOC detection in stage 2, whereas their expert-processed equivalents made a ground truth for the system evaluation. [Fig 3](#) shows the sample and data flow through the study.

Target VOCs

A set of 30 target VOCs ([Table 1](#)) was designed to contain compounds commonly found in the breath, including alkanes, aldehydes, ketones, furans and siloxanes and sulfur-containing compounds [15]. The confidence in the expert-led identification process for these VOCs varies according to potential confounding factors. In particular, confounding factors include low concentrations of a VOC in a sample, which results in a low signal-to-noise ratio in GC-MS data. Propionic acid is an example of a target VOC reporting relatively low concentration in the clinical samples from the dataset. Other confounding factors are mass spectra overlapping caused by the co-elution of VOCs from the GC column, and the similar mass spectra among VOCs. For example, Octane and Hexanal frequently overlap, additionally sharing three of the five top ions. What is more, octane produces highly similar mass spectrum profiles as another target VOC—2,4-dimethylheptane. These VOCs also elute at a relatively close RT locations and thus may be difficult to differentiate with the current expert-led processing [15]. As a consequence of the abovementioned challenges, the process of expert-led VOC identification cannot be guaranteed to be error-free, resulting in a possibly noisy-labelled VOC dataset (stage 1) and ground truth (stage 2). The S1 Table in [S1 File](#) provides, for each target VOC, the details

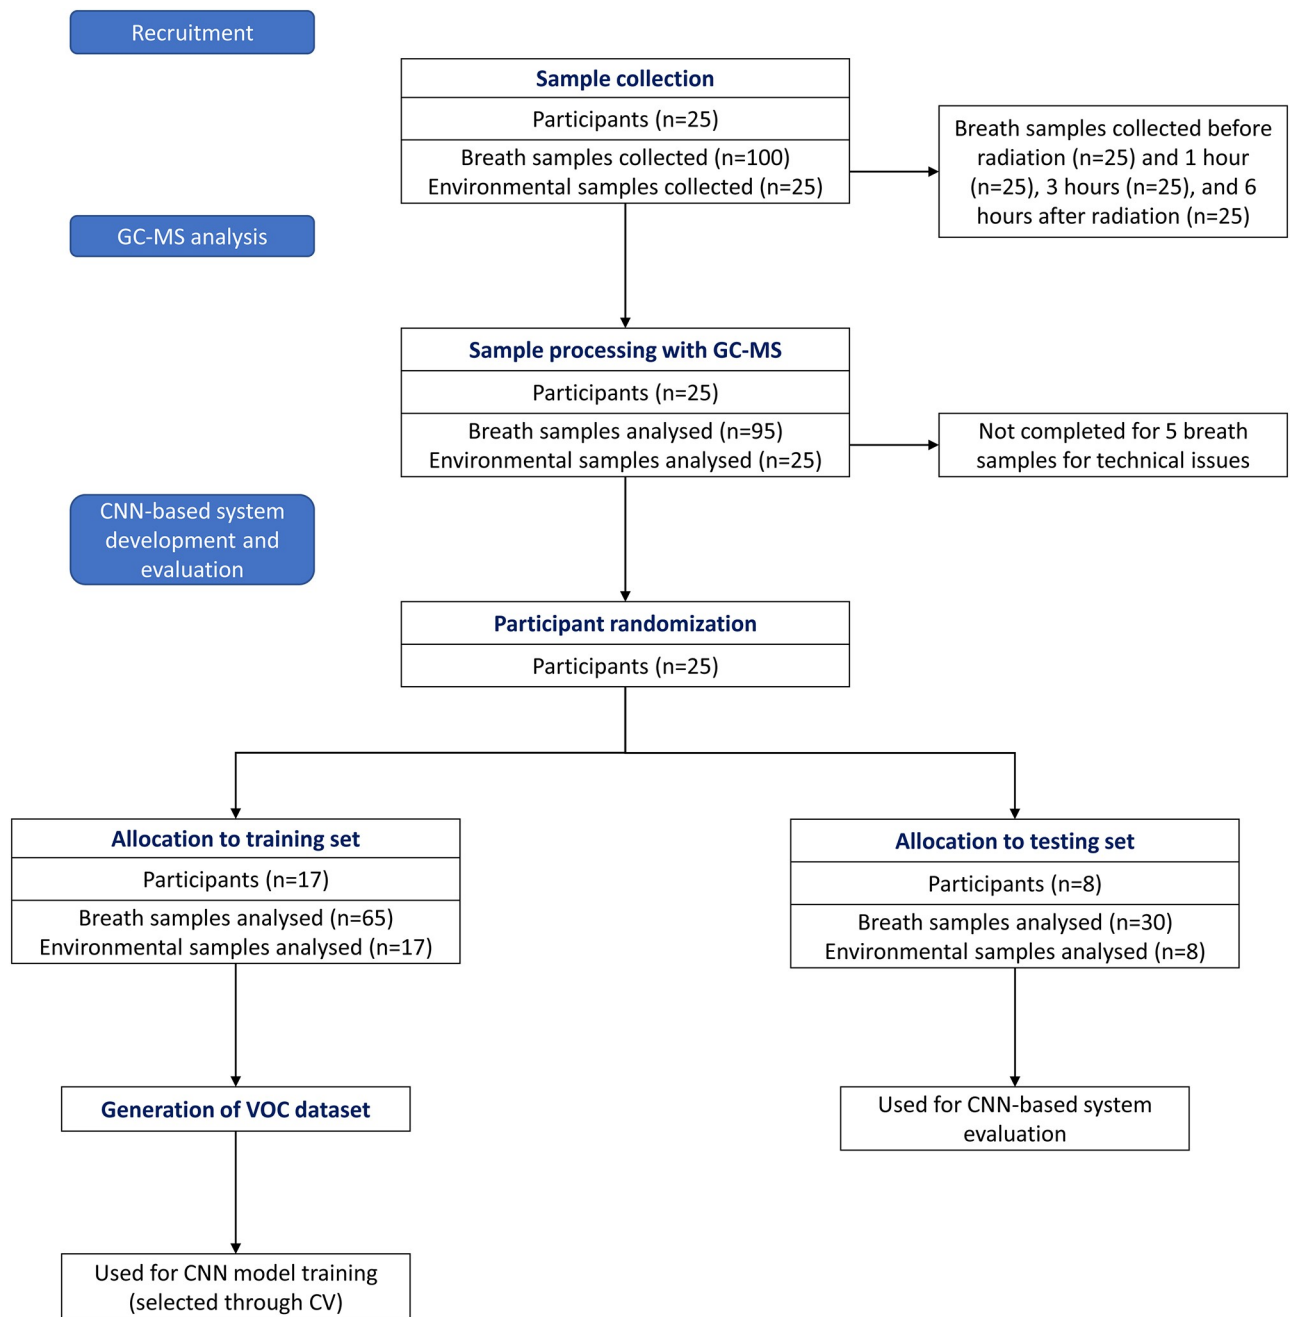


Fig 3. Sample and data flow through the study.

<https://doi.org/10.1371/journal.pone.0265399.g003>

of its ion pattern, mean *RT* location and mean concentration in GC-MS sample dataset, and a number of VOC instances reported respectively in breath and environmental samples.

VOC dataset

The proposed CNN-based approach relies on the construction of a dataset of target VOC patterns, derived from raw GC-MS samples for the network training. Each VOC appears on the TIC chromatogram (Fig 1b) as one peak (sometimes overlapped) over a small segment of *RT*

Table 1. List of the target VOCs with class labels in the elution order.

Label	Target compound	Label	Target compound
0	Negative class	16	Heptanal*
1	Ethanol	17	Benzaldehyde
2	Dimethyl sulfide	18	Benzonitrile
3	2-Methylfuran	19	Octanal*
4	Trichloromethane-d	20	Limonene
5	Benzene	21	2-Ethylhexanol
6	Propionic acid*	22	Nonanal
7	Dimethyl disulfide	23	Dodecane
8	Toluene-D8	24	Decanal*
9	Toluene	25	2-Phenoxyethanol
10	3-Methylthiophene	26	Phthalic acid**
11	Octane**	27	Tetradecane
12	Hexanal**	28	1,4-Diacetylbenzene
13	2,4-Dimethylheptane	29	3,3,6,6-Tetraphenyl-1,2,4,5-tetroxane*
14	2,4-Dimethyl-1-heptene	30	2,5-Diphenyl-1,4-benzoquinone
15	3-Heptanone		

The frames indicate pairs of overlapping compounds. Compounds of relatively low concentrations are marked with * and **: Mean EIC-Area* $\leq 10^3 \leq$ Mean EIC-Area** $\leq Q_1$ (first quartile); see S1 Table in S1 File.

<https://doi.org/10.1371/journal.pone.0265399.t001>

axis, corresponding to a specific range of retention times when the VOC was eluting from the GC column. A sub-matrix of abundance matrix A , encompassing such an RT range, contains the ion pattern for that specific VOC. Thus, we defined a VOC data point as a matrix $\hat{s} \in \mathbb{R}^{\delta \times 411}$, where the retention time window size was selected as $\delta = 80$ (see Methods).

Data points corresponding to all target VOC instances, previously identified by expert-led processing in the training GC-MS samples, were extracted to form the labelled VOC dataset for network training. In addition to data points representing target VOCs, a negative class of data points was created from randomly chosen sub-matrices \hat{s} of A that did not contain target VOCs. The size of the negative class was selected as equal to the total number of the data points representing the target VOCs, i.e., as half of the generated dataset. It was motivated by the compromise between significant unbalance in the VOC dataset (whether all the negative examples appearing in the training GC-MS samples would be included) and coverage of a wide variety of ion patterns that the negative class represents. In total, from the 82 training GC-MS samples, 3,736 VOC data points were extracted with a minimum of 22 and a maximum of 82 data points in each class of target compounds (S1 Table in S1 File). Fig 4 shows examples of VOC data points extracted from raw GC-MS data.

Data augmentation and normalisation

Data augmentation, i.e., methods for enlargement of the dataset by insertion of unobserved data examples [37], is known to benefit machine learning models where data points are scarce [37]. Often these new examples are constructed from the observed ones, by the introduction of some variations to data points, which do not change their underlying distribution (here, VOC ion patterns) and classes.

Data augmentation was applied to the VOC dataset to increase the robustness of the training. We introduced two methods for VOC dataset augmentation: *translation along RT* and *intensity variation* (see Methods). Combined collectively, these two methods generated 100

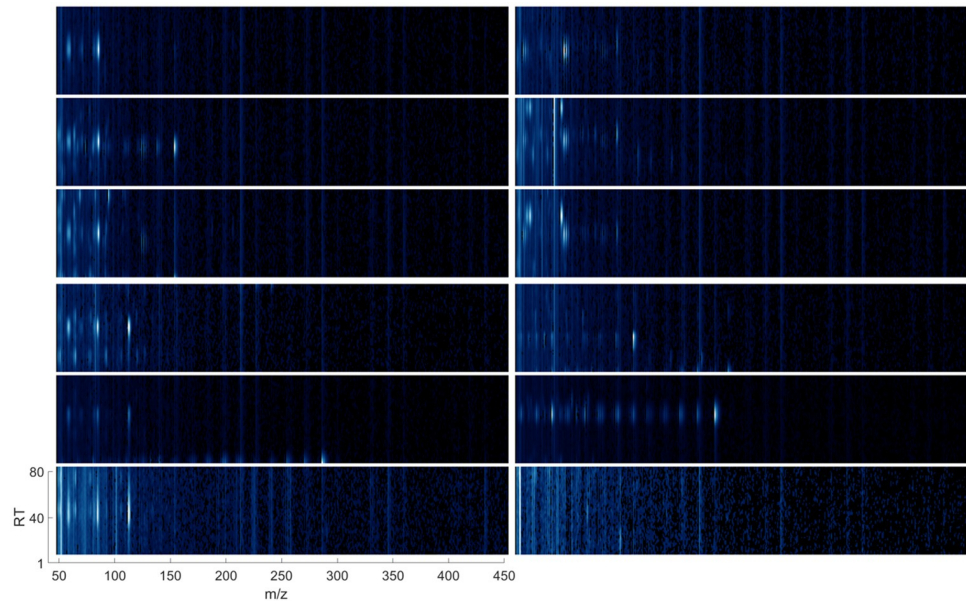


Fig 4. Examples of the VOC data points extracted from the raw GC-MS abundance matrices. Left, from top: three examples of Benzene, three examples of Benzaldehyde; right, from top: three examples of Trichloromethane-d and three example data points from the negative class. For better visualisation, the segments were mapped to RGB format.

<https://doi.org/10.1371/journal.pone.0265399.g004>

variations for each VOC data point, resulting in the fully-augmented dataset of 373,600 VOC data points. To test the impact of the proposed augmentation methods and select the most efficient approach for network training, we compared the performance achieved on the fully-augmented dataset, partially-augmented dataset (after translation along *RT* only) and original VOC dataset. The intensity values in each data point were normalised in the range [0, 1].

Deep learning models

We assessed the learning and inference capabilities of the following recently developed deep CNN architectures, which achieved state-of-the-art performance in several image classification challenges: VGG16 [33], residual CNNs (ResNets) [34] and densely connected CNNs (DenseNets) [35]. We also tested smaller VGG-like networks with 4, 6 and 8 layers.

Along the *m/z* axis the values of a GC-MS abundance matrix are spatially only weakly correlated, as opposed to the data typically used in image-based tasks. Therefore, we tested the effectiveness of 1D filters (see [Methods](#)) and compared them with typically used 2D filters to investigate the suitability of such a network variation to capture the specific nature of correlation in GC-MS data. The details on tested CNN architectures, their implementations as well as parameters settings are given in S3.1-S3.11 Tables in [S1 File](#).

CNN model training

To evaluate the CNN models' ability to learn ion patterns from the VOC dataset, and to select the best-performing configurations of the networks, five-fold cross-validation (CV) was applied. CV is a common technique used for model assessment and hyperparameter selection [36]. Each split of the VOC dataset for CV was performed at the level of participant.

Table 2. Five-fold cross-validation (CV) performance achieved by tested CNN models on the VOC dataset. (Model VGG16 with 1D filters has not been tested due to its extensive demand for memory resources).

Model	Depth	Filters	CV accuracy: mean (\pm sth) %		
			Fully-aug.	Partially-aug.	Original data
VGG-like	4	1D	97.16 (± 0.82)	96.89 (± 0.65)	94.60 (± 1.02)
VGG-like	4	2D	97.13 (± 0.92)	96.48 (± 0.89)	93.89 (± 1.06)
VGG-like	6	1D	97.60 (± 1.08)	97.43 (± 0.78)	94.75 (± 1.11)
VGG-like	6	2D	97.37 (± 0.95)	96.69 (± 0.79)	93.73 (± 1.12)
VGG-like	8	1D	98.06 (± 0.38)	97.83 (± 0.38)	94.51 (± 1.43)
VGG-like	8	2D	97.34 (± 0.94)	96.89 (± 1.03)	88.11 (± 1.62)
VGG16	16	2D	97.89 (± 0.47)	97.68 (± 0.40)	96.14 (± 1.63)
DenseNet	40	1D	97.34 (± 1.18)	96.51 (± 1.67)	94.55 (± 2.69)
DenseNet	40	2D	96.73 (± 0.82)	95.92 (± 1.98)	69.98 (± 1.96)
ResNet	34	1D	97.71 (± 0.68)	97.64 (± 0.45)	94.26 (± 2.50)
ResNet	34	2D	97.25 (± 0.88)	96.98 (± 0.67)	92.84 (± 1.53)

<https://doi.org/10.1371/journal.pone.0265399.t002>

Table 2 presents the classification performance of the models in the CV stage. All tested models achieved high accuracy (i.e., a proportion of data points for which the network label matched the ground truth), which confirms that ion patterns can be learned by CNNs from VOC data points derived directly from raw GC-MS data. The results also validate the hypothesis that 1D filters are more effective than standard 2D filters. Among the tested augmentation strategies, the best models' performances were obtained with the fully-augmented dataset. Shallower architectures, i.e., VGG-like, performed no worse than other deeper networks. The VGG-like network with 8 layers and 1D filters (VGG-8-1D) reported the highest accuracy.

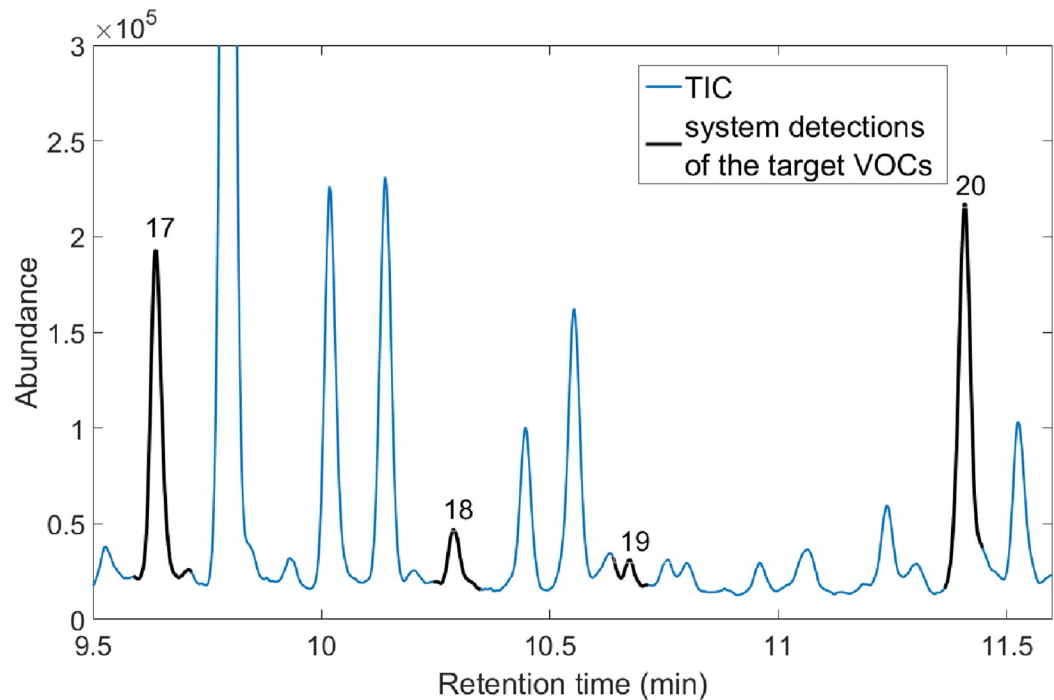
The best performing hyperparameter configuration for each investigated CNN architecture type, i.e., VGG-8-1D, DenseNet-40-1D and ResNet-34-1D, was used to create the system for automated VOC detection in raw GC-MS samples (Fig 2, stage 2). Before the deployment to stage 2, the models were retrained on the entire, fully-augmented VOC dataset. S8 Fig in S1 File shows a learning curve for VGG-8-1D model.

Automated samples analysis and VOC detection

The trained networks were employed in stage 2 (Fig 2) to analyse the 38 raw GC-MS samples from the testing set. The proposed stage 2 analysis was composed of two phases: (A) scanning of each clinical sample with the network; (B) identification of VOC detections from the scan results.

Phase A—Scanning of a breath sample. The abundance matrix A of a clinical sample was scanned along the RT dimension. Precisely, the network was fed consecutive normalised sub-matrices $s_i \in \mathbb{R}^{\delta \times 411}$ of A , starting at retention point i (for each i in the RT dimension for that sample). Phase A classifies each s_i into one of the 31 possible classes. Thus, scanning a single GC-MS sample required approximately 22500 queries of the network. The output of the process are two sequences: a sequence L_A that contained approximately 22500 class labels, and a sequence T_A that contained the classification confidence given by the network for each respective label (see Methods).

Phase B—VOC detection. To obtain a list of target VOCs detected in each clinical GC-MS sample, the pair of phase A output vectors for that sample was analysed. The following general properties of the GC-MS samples are considered in the phase B analysis:



VOC Label	Start detection sRT	End detection eRT	Confidence \mathcal{T}
17	9.588	9.712	1.0000
18	10.243	10.359	1.0000
19	10.638	10.715	0.9974
20	11.364	11.447	1.0000

Fig 5. Example output from clinical sample scanning for VOC detection (stage 2) with a CNN-based system, range between 9.5 and 11.5 minutes; sample Test-01-BS01, model VGG-8-1D (S4.2.1 Table in S1 File). Top: VOCs detected by the system in this RT range, visualised on the TIC chromatogram. Bottom: Table of detected VOCs along with their RT positions and detection confidences.

<https://doi.org/10.1371/journal.pone.0265399.g005>

- (i). Along the RT dimension, a VOC is measured multiple times consecutively for the duration of the VOC elution (typically about 6 seconds).
- (ii). The order of elution of VOCs from the GC column is usually constant across different samples, as indicated in Table 1.
- (iii). Each VOC elutes from the GC column only once in a GC-MS process.

According to (i), a detection $d_j \in \mathcal{D}_A$ was defined as a consecutive sub-sequence of L_A of length at least γ with constant label values j (*duration rule*, see Methods). Subsequently, according to (ii), the system ignored detections that were far from the expected order (*order rule*), giving a set $\tilde{\mathcal{D}}_A \subset \mathcal{D}_A$. Finally, according to (iii), if multiple detections of one VOC occurred in any sample, the system selected the one with the highest *detection confidence*, derived from T_A , expressing the confidence of the model in such detection (*uniqueness rule*); a set $\approx \mathcal{D}_A \subset \tilde{\mathcal{D}}_A$.

As a result, for each abundance matrix A , stage 2 delivered a list of detected VOCs along with their location along the RT axis and the detection confidence. A graphical representation of the output is provided in Fig 5.

Evaluation of the automated sample analysis and VOC detection

The performance of the system was assessed with two approaches: (1) including both the VOC localisation (along *RT* dimension) and its classification; and additionally, to report the system specificity, (2) including only the VOC presence in a sample.

(1)—VOC localisation and classification. VOC detection involves both localisation of the compound in the GC-MS sample and its classification. Therefore, each VOC detection in a sample was identified as a true positive (TP) if there was matching in both label and *RT* position with the ground truth for that sample, and as a false positive otherwise. Similarly, if a specific target VOC was not detected at the *RT* position reported by the ground truth, it was identified as a false negative. Note that a retention time window of any size in which the system and the expert both did not identify any target VOC could be considered as a true negative: therefore true negatives were not measured here.

System-derived ground truth corrections. To gain further insight on the results, we examined all the detections $\hat{d}_j \in \mathcal{D}_A$ (i.e., before applying order and uniqueness rules) that were not reported in the ground truth or had mismatching *RT* positions. The range of retention times in the ground truth was calculated for each VOC (see [Methods](#)). We observed that in most cases, the detection \hat{d}_j was reported by the system in the expected (compatible) narrow *RT* range, characteristic for that VOC *j*. The upper bound probability of a random false positive detection occurring within a precise and restricted *RT* range was calculated as 4% (see [Methods](#)). Accordingly, it is highly likely that such false positive is actually a true positive and reveals an error in the (noisy-labelled) ground truth, i.e., a VOC occurrence missed in the expert-led analysis. False positives with a compatible *RT* value (i.e., characteristic for that specific VOC detected) were named *tentative true positives* (TTP), whereas false positives that are not TTP are called certain false positives (FP); see [Fig 6](#).

In some cases, the VOCs detected by the system and identified as TTP were, in fact, reported by the ground truth, but at different *RT* positions ([Fig 6](#)). Accepting the high chance of tentative true positives to be true positives and the constraints that a VOC appears only once in a sample, the corresponding false negatives at the *RT* positions reported by the expert were identified as *tentative true negatives* (TTN). Note that only FN with respective TTP were verified; in fact, more examples identified as FN may be correctly not detected by the system. False negatives that were not TTN were called semi-certain false negatives (FN). Further details on the concepts of tentative true and false positives, and certain and semi-certain false positives are provided in the [Methods](#).

[Table 3](#) reports the performance achieved on the testing set by the system with different CNN models. Additionally, the result intersection, i.e. detections consistent (in terms of label and *RT*) among all the models, was evaluated. Results are presented in two forms: according to the expert-derived ground truth and according to the system-derived corrections. The total quantities of TP, TTP, FP (certain), FN (semi-certain) and TTN across all samples are given. We report the system sensitivity and mean average precision score (mAP), which is an evaluation metric commonly used in the object detection domain [38] (see [Method](#)). [Fig 6](#) shows a graphical example, in which TP, TTP, FN and TTN were identified.

(2)—VOC presence. To compute the system's specificity, we considered only the question of the presence of each VOC in each sample and ignore its *RT* position. Since each VOC may appear in a sample at most once, such an approach leads to a binary classification problem. True negatives (TN) were thus well defined here as VOCs not detected in a sample by both the system and the expert-led processing. False positives were defined as VOCs detected by the system and not reported by the ground truth. In particular, this definition is different than

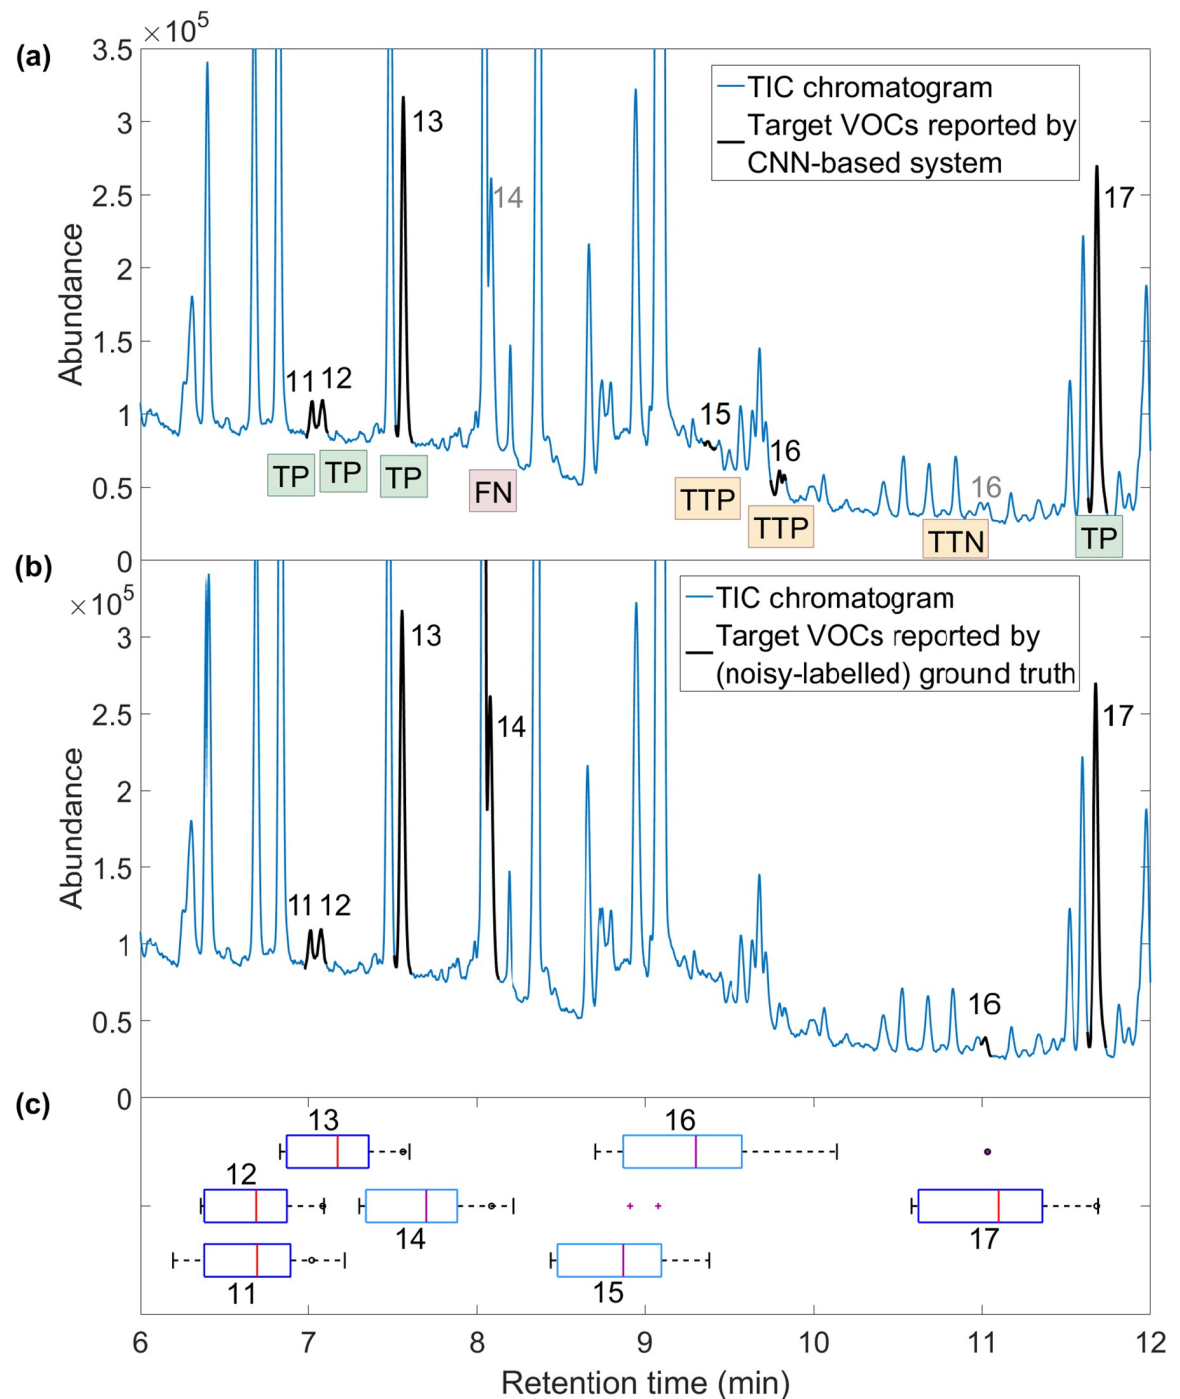


Fig 6. An evaluation of the results of a GC-MS sample scanning for VOC detection (stage 2) with the CNN-based system, range between 6 and 12 minutes; sample Test-04-BS01, model VGG-8-1D (S4.2.16 Table in S1 File). (a) VOCs detected by the system presented on total ion current (TIC) chromatogram and marked accordingly to the outcome of the evaluation. VOCs 11, 12, 13 and 17 detected alike by the system and the ground truth, true positives. VOC 14 reported by the ground truth and not detected by the system, false negative. VOC 15 not reported by the ground truth, but detected by the system in compatible RT range, tentative true positive. VOC 16 not detected by the system on the position reported by the ground truth, but detected on a different position within its RT range, tentative true negative and tentative true positive. (b) VOCs reported by the (noisy-labelled) ground truth presented on TIC chromatogram. (c) RT ranges specific for the VOCs, derived from the ground truth. Black dots indicate the exact RT positions of VOCs in the ground truth. Note that VOC 16 is reported by the ground truth at RT position being an outlier for the RT range of this compound.

<https://doi.org/10.1371/journal.pone.0265399.g006>

Table 3. Evaluation (1) of the results of stage 2 analysis of testing GC-MS samples.

Model	Classification					Benchmark	Metrics	
	TP	TTP	FP	TTN	FN		sensitivity	mAP
VGG-8-1D	816	226	2	18	11	expert	0.9657	0.9019
						system-derived correction	0.9896	0.9894
DenseNet-40-1D	808	181	6	15	22	expert	0.9562	0.93
						system-derived correction	0.9782	0.9994
ResNet-34-1D	804	187	9	14	27	expert	0.9515	0.9282
						system-derived correction	0.9735	0.998
Intersection	787	138	1	8	49	expert	0.9325	0.9398
						system-derived correction	0.9497	1

Benchmarks: expert-derived ground truth (yellow) [tentative true positives (TTP) are considered FP, tentative true negatives (TTN) are considered FN]; system-derived correction (green) [tentative true positives (TTP) are considered TP, tentative true negatives (TTN) are considered TN].

<https://doi.org/10.1371/journal.pone.0265399.t003>

above, i.e., in (1) false positives covered also VOCs reported by both the system and the expert but at different *RT* positions.

System-derived ground truth corrections. Similarly as above, false positives reported in their respective *RT ranges* were identified as tentative true positives (TTP*), false positives that were not TTP were called certain false positives (FP*).

Table 4 reports the system specificity in the analysis of the 38 clinical samples from the testing set. Results are presented according to the expert-derived ground truth and according to the system-derived corrections. The total quantities of TN, FP* (certain), TTP* across all samples are given along with the system specificity. The full tables of detections for each CNN model and their intersection, presented per each target VOC and per each testing sample, are reported in S4.1–7.2 Tables in [S1 File](#).

Discussion

The results show that all tested system configurations, with various CNN models employed, achieved high performance in the detection of target VOCs in the clinical samples, reporting high sensitivity and specificity (when admitting the noisy-labelled ground truth). The results demonstrate that ion patterns can be effectively learnt directly from the raw GC-MS data. The

Table 4. Evaluation (2) of the results of stage 2 analysis of testing GC-MS samples.

Model	Classification			Benchmark	Metric
	TN	FP*	TTP*		specificity
VGG-8-1D	86	1	208	expert	0.2915
				system-derived correction	0.9885
DenseNet-40-1D	125	4	166	expert	0.4237
				system-derived correction	0.969
ResNet-34-1D	116	6	173	expert	0.3932
				system-derived correction	0.9508
Intersection	165	0	130	expert	0.5593
				system-derived correction	1

Benchmarks: expert-derived ground truth (yellow) [tentative true positives (TTP*) are considered FP*]; system-derived correction (green) [tentative true positives (TTP*) are considered TP].

<https://doi.org/10.1371/journal.pone.0265399.t004>

problem of VOC detection includes various challenges such as detecting VOCs of low intensities and distinguishing overlapping VOCs and those with similar ion patterns (Table 1). Nevertheless, the system performed comparably well for all 30 VOCs from the dataset (S4.1, S5.1, S6.1 Tables in S1 File). Consequently, we claim that the CNN-based system proposed here allows bypassing time-consuming and labour-intensive expert-led data processing in real-world GC-MS data targeted analysis.

We found strong evidence that the proposed CNN-based system may outperform human experts. In our tests, the system detected 17% to 23% more occurrences of the VOCs than expert-led deconvolution-based method (TTP, Table 3). As much as 138 TTP were reported by each of the tested models. Additionally, the system did reveal possible errors, i.e., incorrectly labelled VOCs (TTN). There are two possible reasons to explain the CNN-based system out-performance over the current expert-led processing. The current method involves complex preprocessing steps, based on the operator-subjective spectral deconvolution [17]. The proposed system instead analyses raw GC-MS data and thus bypasses multiple steps and possibly suboptimal choices of parameters. Another reason why the proposed workflow delivers more VOC detections may be its potential to detect compounds of low intensities. Additional analysis of the results (S1, S3.1, S4.1, S5.1 Tables in S1 File) showed indeed that over 50% of TTP (for VGG-8-1D; over 40% for rest of the models) were reported by 25% of target VOCs with the lowest average concentration (Table 1). Most TTP, 27, were reported for Propanoic acid (VOC 6), a compound of low concentration. Interestingly, before augmentation, the VOC dataset had only 22 training examples of VOC 6. Despite that, the CNN-based system was able to detect this compound correctly: none FP and FN reported, all 27 TTP reported by all the models. In our tests, the proposed system improves on the state-of-the-art performance of the current deconvolution-based process.

The proposed CNN-based system requires expert knowledge to be trained (stage 1), but consequently, it can detect VOCs autonomously and significantly faster than a human-driven procedure. The training stage in the proposed approach, requiring about 23 to 80 hours (depending on the CNN model), has to be performed only once to be able to scan new samples at a rate of as low as just around 2 minutes per sample (see Table 5 in Methods). Interestingly, the fastest among the tested architectures, i.e., VGG-8-1D, is also the one with the highest sensitivity. Consequently, the proposed system may support experts in much faster validation of hypotheses regarding compounds related to specific health conditions on new GC-MS samples. What is more, those hypotheses may be more accurate due to the system's ability to provide a more comprehensive list of VOCs than the deconvolution-based methods.

Very deep networks are often thought to have more discrimination capabilities than shallower ones, but in this application, the VGG-like networks with reduced depth performed no worse than other very deep models (Table 2). This suggests that the detection and classification of VOC ion patterns can be effectively performed with modest-depth networks, which are also less resource-intensive. All tested networks achieved higher performance when implemented with 1D filters, adapted to the nature of GC-MS data.

The proposed CNN-based approach provides specific information on individual VOCs in breath for further analysis and diagnosis, rather than a high-level sample classification (e.g., [4, 7, 21]). As opposed to black-box and end-to-end machine learning diagnostic systems, the proposed system quickly produces accurate lists of VOCs, thus enabling a transparent and explainable pipeline for breathomics-based diagnosis. This approach avoids the common problem of having limited datasets of clinical samples: each of many target VOC occurrences in a sample makes a data point; in this study, 82 clinical samples and 30 target VOCs resulted in a dataset with 3,736 unique data points and 373,600 augmented data points.

The proposed system does not use *RT* values, which are a major source of variation among measurements with different instruments (or even the same instrument but on different days). Instead, the proposed system analyses the patterns of the VOCs, which remain substantially the same when are processed with similar instruments, in terms of type and properties such as resolution [29]. Therefore, the CNN-based system may be transferable to data obtained from another GC-MS instrument than the one from which the training dataset originates. Further tests are required to assess such proprietary precisely.

The proposed approach exploits GC-MS general properties and thus potentially extends beyond breath data. Further studies may extend tests to GC-MS data from a large variety of domains, e.g.: detection of CBRN (chemical, biological, radiological, nuclear) biomarkers in breath, saliva and skin for casualty triage [39]; tracking organic pollutants in water for environment monitoring [40]; detection of accelerants in fire debris for criminal forensics [41]; detecting drug ingredients in urine samples for law enforcement or sports anti-doping analysis [42, 43]; analysis of chemical composition of the planets' atmosphere in astrochemistry [44]; as well as in chemical engineering [45], food, beverage and perfume analysis [46–48] and medicine [8].

The proposed new approach to GC-MS data analysis, exploiting the application of deep learning, has the potential for extensive development in the future. Increasing breath analysis as a diagnostic technology will also increase the number of available GC-MS datasets: a larger number of VOC patterns, reflecting more of the possible variations in the data points, may benefit the accuracy of deep neural network training. The use of GPU computing and dedicated hardware can help process the large amount of data collected through GC-MS; additionally, its rapid development, seen in recent years, may reduce the processing time even further.

Future studies can extend the proposed CNN-based approach to also measure VOC intensities. The proposed system is currently limited to detecting the presence of VOCs of interest, but not their abundances. However, in real life scenarios, additional analysis by experts to determine peaks intensities and then concentrations of particular VOCs, e.g., biomarkers related to a specific disease, may become necessary only if the system reveals their presence in a sample. In fact, as a result, the proposed workflow delivers along with a list of detected VOCs their *RT* positions in the sample, which can significantly facilitate and accelerate the quantification of the compounds.

In summary, the proposed CNN-based system delivers a faster, more accurate and scalable method for automated targeted analysis of raw GC-MS data than the current state-of-the-art expert-led processing. The proposed approach has a significant potential to contribute to the development of breath analysis as a diagnostic platform to detect various diseases quickly, efficiently, and reliably.

Methods

Ethical approval

The study was approved by the South East Scotland Research Ethics Committee 01 (16/SS/0059) and all clinical staff were trained, and proficiency tested for breath analysis prior to the start of patient recruitment. Informed and written consent was given by all participants.

Breath sample collection

Breath samples were collected from 25 participants before and after radiotherapy at 1, 3, and 6 hr. A Respiration Collector for In Vitro Analysis ReCIVA™ device (Owlstone Medical, Cambridge, UK) was used for the breath samples collection. Clean air was provided from room-air filtered with an activated-carbon scrubber and HEPA filter, at a flow of 35 dm³ min⁻¹. (The

air-supply unit was built and tested by the Centre for Analytical Science, Loughborough, UK). The total sample volume of breath was set to 1000 cm³ with a sampling duration cap of 900 s. Tenax[®]/Carbotrap 1TD hydrophobic adsorbent tube (Markes International Ltd, Llantrisant, UK) was used. All materials were conditioned and sterilised before use to reduce exogenous VOC artefacts. Clinical staff were trained and proficiency-tested prior to clinical breath sampling. Environmental and air-supply samples were collected with each set of breath samples. Samples were sealed and stored at ca. 4°C immediately after collection and transported to Loughborough Centre for Analytical Science within 48 hr [49]. Samples were dry-purged on receipt with a 120 cm³ of purified nitrogen at a flow rate of 60 cm³ min⁻¹. Toluene-D8 (0.069 ng) and trichloromethane-d (0.28 ng) internal standards were spiked into the sample during the dry-purge process using a six-port valve. All samples were then sealed and stored at 4°C prior to analysis.

GC-MS processing of clinical samples

Thermal desorption (Unity-2, Markes International) interfaced to a GC (Agilent, 7890A) coupled to a quadrupole mass spectrometer (Agilent, MS 5977A) was used for the analysis of all clinical samples, see S9 Table in [S1 File](#) for operating details. The ion channels from 40 *m/z* to 450 *m/z* were measured with unit resolution. The instrumental scanning rate was approximately 6.25 Hz, which for about one-hour processing gave approximately 22500 *RT* points. The size of the derived abundance matrix of each sample is $R \times 411$, where $R \approx 22500$.

The samples were analysed over a year period (Sep 2016—Sep 2017) as part of a wider multi-centre clinical study campaign. Instruments were serviced whenever statistical process controls indicated a *z*-score >3 or more than 3 consecutive *z*-scores >2. The frequency of the service interventions was determined by the quality and contamination levels of the samples returned from the clinics. The column was replaced once during this phase of the campaign, in March 2017.

Abundance matrix

For each sample, GC-MS processing produces an abundance matrix $A \in \mathbb{R}^{R \times 411}$ (see S2 Fig in [S1 File](#) for an example). Let $z_i \in \mathbb{R}^{411}$, $i = 1, \dots, R$ be mass spectrum (i.e., intensities across consecutive ion channels derived by MS) at particular retention time points r_i , i.e.:

$$A = \begin{bmatrix} z_1 \\ \vdots \\ z_R \end{bmatrix}.$$

For each ion channel z_i its corresponding *RT* point r_i is specified by the function \mathcal{RT} :

$$\mathcal{RT} : z_i \rightarrow r_i \in \mathbb{R}^+.$$

Expert-led GC-MS data processing

GC-MS data denoising and baseline correction and feature deconvolution were carried out (AnalyzerPro Spectral Works, UK) by a highly-qualified and experienced chemical analyst. As a result, 350 to 500 VOC features per sample were recorded. Features were aligned to correct for retention time variation, and the VOCcluster algorithm [15] was used to cluster all features into groups. Each feature was assigned an identifier in the format: BRI—*m/z*₁ *m/z*₂ . . . *m/z*_{*n*}. BRI indicated the retention index for the VOC breath feature and *m/z*₁ . . . *m/z*_{*n*} are the nominal

masses of the ion-fragments in decreasing order of abundance needed to uniquely define the feature. The processing time of a single breath sample is estimated as 60 to 120 minutes.

To produce the ground truth for this study, this process delivered a list of VOCs with Level 2 chemical identification [50] and their corresponding positions in the matrix of raw data, this is $startRT$, $peakRT$ (denoted also as RT) and $endRT$ —the indexes along the retention time where the compound was measured to start, peak and end the release from the GC column.

The process was not completed for 5 breath samples for technical issues, resulting in 120 expert-processed GC-MS samples.

VOC dataset structure

The VOC dataset was extracted from the abundance matrices of the 82 samples from the training set, using the ground truth for the 30 target VOCs. A VOC data point \hat{s} has a size 411 along the m/z dimension, which ensures inclusion of the entire mass spectrum. The size along RT dimension, δ , was computed with the aim to capture the entire elution process for each target VOC instance. Precisely, the maximum duration of compound elution, $\max(endRT - startRT)$, was measured across 1868 occurrences of target VOCs reported by expert led-processing in training samples. This was computed as ~ 9.8 seconds, corresponding to ~ 61 time steps on the RT axis. To allow for translation-based data augmentation, the size δ was increased by 19, resulting in a final $\delta = 80$. Each data point $\hat{s} \in \mathbb{R}^{\delta \times 411}$ contains mass spectra of a VOC peak centred over RT dimension, i.e.:

$$\hat{s} = \left[z_k, \dots, z_{k+\frac{\delta}{2}-1}, \dots, z_{k+\frac{\delta}{2}}, \dots, z_{k+\frac{\delta}{2}+1}, \dots, z_{k+\delta} \right]^T$$

where

$$\mathcal{RT}\left(z_{k+\frac{\delta}{2}-1}\right) = startRT,$$

$$\mathcal{RT}\left(z_{k+\frac{\delta}{2}+1}\right) = endRT,$$

$$\mathcal{RT}\left(z_{k+\frac{\delta}{2}}\right) = \mu,$$

for $\mu = \frac{startRT+endRT}{2}$ —a middle point of the VOC peak shape. Note that μ does not necessarily coincide with $peakRT$ value as VOC peaks can be not symmetric but skewed. l value depends on the VOC instance.

The VOC dataset for CNN training was unbalanced: the groups representing target VOCs were unequal since each of the considered VOCs did not necessarily occur in each of the GC-MS samples from the training set. What is more, as during the scanning of entire raw breath samples the negative examples appear much often than target VOCs and cover a broad variety of ion patterns, the negative class was created as half of the VOC dataset. This aligns with common practice in object detection domain, where the negative class usually dominates with the ratio up to 3:1 [51, 52]. The size of each class is reported in S1 Table in [S1 File](#).

Data augmentation

In the proposed approach, we devised and tested two augmentation methods that maintain the underlying structure of a VOC pattern.

Translation along RT. The VOC data point \hat{s} is centred on the VOC peak so that the middle point of the VOC peak shape, $\mu = \frac{startRT+endRT}{2}$, is located at $\delta/2$. Twenty VOC data points were created from \hat{s} by shifting the extraction point from the abundance matrix A from -9

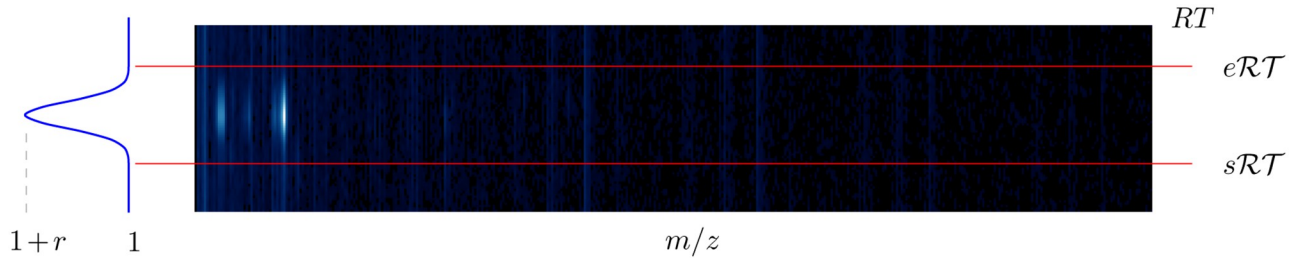


Fig 7. Scheme of the data augmentation by the intensity variation in VOC the data point. Along RT points containing the VOC pattern (i.e., between sRT and eRT), the intensities of each m/z channel were multiplied by a Gaussian-shaped function to simulate variations of the VOC concentration.

<https://doi.org/10.1371/journal.pone.0265399.g007>

to +10 steps along the RT axis (0 indicates \hat{s}), i.e.:

$$\hat{s}_n^* = \left[z_{k+n}, \dots, z_{k+\frac{\delta}{2}}, \dots, z_{k+n+\frac{\delta}{2}}, \dots, z_{k+n+\delta} \right]^T$$

Such data augmentation presents the VOC pattern at different positions in the data point: top, centre, bottom (translation along RT) and represents variations in the pattern’s location in the subsequent sub-matrices of A seen by the network, while scanning entire raw GC-MS breath samples.

Intensity variation. In the VOC data point, the intensities along RT interval corresponding to the VOC pattern were varied to simulate variations of the VOC concentration (see Fig 7). Background (i.e., values along RT not containing the pattern) remains unchanged. In this process, it was important to maintain the VOC ion pattern, i.e., relative ratios along the m/z dimension while increasing peak intensity along the RT axis. This was achieved by multiplication of each mass spectrum $z_{startRT}, \dots, z_{endRT}$ (corresponding to VOC location) by a particular value of a Gaussian-shaped function along RT interval. Precisely, for a data point \hat{s} :

$$\hat{s} = \begin{bmatrix} z_k \\ \vdots \\ z_{k+\delta} \end{bmatrix}.$$

the augmented data point \hat{s}^* can be computed as

$$\hat{s}^* = \begin{bmatrix} z_k \cdot G_{RT(z_k)} \\ \vdots \\ z_{k+\delta} \cdot G_{RT(z_{k+\delta})} \end{bmatrix},$$

where

$$G_x = \begin{cases} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \cdot r + 1, & startRT < x < endRT \\ 1, & otherwise, \end{cases}$$

and r is a random value in the range (0, 0.1). This augmentation step was repeated to obtain 4 additional data points for each data point derived with translation along RT . Therefore, the

VOC dataset for CNN training was augmented 100 times (i.e., 20×5 times). The intensity values in each data point from the augmented dataset were normalised in the range [0, 1].

The original VOC dataset had 3,736 data points. Following *translation along RT* the dataset for CNN training consisted of 74,720 segments (partially-augmented dataset). When *intensity variation* was also applied, the dataset contained 373,600 data points (fully-augmented dataset). Results presented in the [Table 2](#) indicate that both developed augmentation methods bring benefits to the system: the best performance was achieved on the fully-augmented dataset.

CNN filter adaptation

To adjust the deep learning models specifically to VOC detection, we adapt the CNN filters to the GC-MS data characteristics. In CNNs, filters specify the local receptive fields, i.e., the regions of the image (or feature maps) visible by convolutional and pooling layers of the network at a time; this is an essential concept that enables CNNs to capture local geometric spatial correlations in the data [22]. With GC-MS data, such a local correlation occurs only in the retention time dimension. Along this dimension, the abundance of different m/z increases and decreases depicting peaks as the VOC exit the GC column. On the other hand, the abundance values across different m/z channels also correlate as the particular ions make up the VOC pattern ([Fig 1a and 1b](#)). However the values along this dimension represent independent ion channels and locally they are only weakly correlated, thus their correlation cannot be captured by small local filters.

One of the hypotheses in this study is that the convolutional and pooling layers in the network do not need to be two-dimensional as it is usual for image classification. Thus, two types of filters are tested: a traditional two-dimensional filter and a specific one-dimensional filter along the *RT* axis to cover only this dimension. The filters sizes in the particular network layers are given, along with the detailed architecture of each network, in S3.1-S3.11 Tables in [S1 File](#).

Phase A output

The last layer in all the tested networks is a fully connected layer with softmax activation. The softmax function $\sigma : \mathbb{R}^k \rightarrow \mathbb{R}^k$, defined as

$$\sigma(\mathbf{x})_i = \frac{e^{x_i}}{\sum_{j=0}^k e^{x_j}}$$

for $i = 1, \dots, k$, $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{R}^k$, takes as input a vector of k real numbers and normalises it into a probability distribution consisting of k probabilities; all the components are mapped to the interval (0, 1) such that their sum is 1. Therefore, the network returns a probability distribution over output classes, i.e the probabilities of the allocation of each particular normalised data point $s \in \mathbb{R}^{\delta \times 411}$ to each of 31 VOC groups:

$$s \rightarrow (P_0(s), \dots, P_{30}(s) \mid P_i(s) = P(s \in C_i))$$

for $i = 0, \dots, 30$, C_i —class representing the VOC labelled with i . (Note that $P_i(s) = \sigma(\mathbf{x})_i$, where \mathbf{x} is a vector of 31 entries generated for data point s by the last fully connected layer of the network.)

The VOC class that gives the highest probability for a data point s is selected as its classification c :

$$c(s) = \arg \max_i \{P_i(s) \mid i = 0, \dots, 30\}$$

with confidence t :

$$t(s) = \max_i \{P_i(s) \mid i = 0, \dots, 30\}.$$

As a result, for each analysed raw GC-MS sample with abundance matrix $A \in \mathbb{R}^{R \times 411}$, the CNN scanning in phase A produces two sequences L_A and T_A . L_A values indicate classification labels of the subsequent sub-matrices s of A :

$$L_A = (c(s_i) \mid s_i \subset A, i = 1, \dots, N);$$

T_A values indicate classification confidence given by the network for each respective label:

$$T_A = (t(s_i) \mid s_i \subset A, i = 1, \dots, N),$$

where $N = RT - \delta + 1 \approx 22500$ is a number of data points in A .

Duration rule: VOC detection

Each maximal sub-sequence of L_A with constant label values j and length equal or greater than a constant γ indicates one detection d_j of the VOC j , i.e.:

$$d_j = (s_l, \dots, s_{l+n-1}),$$

where

$$c(s_i) = j, i \in [l, l+n-1], \quad c(s_{l-1}) \neq j, \quad c(s_{l+n}) \neq j, \quad n \geq \gamma, \quad j > 0.$$

The value γ was derived from the width setting of matrix s : for the augmentation, width of s was enlarged by 19 pixels in respect to the maximum peak width (see Data augmentation). Hence, during the scanning of the abundance matrix A , the entire shape of a target VOC peak is seen by the model at least $\gamma = 20$ times. All the detections, of any class, from the abundance matrix A (i.e., sample) constitutes the set \mathcal{D}_A .

RT of detection

Let the retention time value corresponding to the specific data point $s = [z_k, \dots, z_{k+\delta}]$ be defined as the retention time value of its middle ion channel, i.e.:

$$\mathcal{RT}(s) = \mathcal{RT}\left(z_{k+\frac{\delta}{2}}\right).$$

Retention time values corresponding to the first and last data points s_i of detection d_j are called respectively the *start detection*, $s\mathcal{RT}(d_j)$, and *end detection*, $e\mathcal{RT}(d_j)$:

$$s\mathcal{RT}(d_j) = \min\{RT(s_i) \mid s_i \in d_j\},$$

$$e\mathcal{RT}(d_j) = \max\{RT(s_i) \mid s_i \in d_j\}.$$

We denote the resulting detection interval by

$$DI(d_j) = [s\mathcal{RT}(d_j), e\mathcal{RT}(d_j)].$$

Order rule

The elution of particular VOCs from the GC column can be expected in a certain order (small variations may occur) related to compounds' chemical properties. Therefore, to derive a reliable list of detected VOCs, the detections that fall outside the expected order are removed.

Let \leq be a linear order on a set \mathcal{D}_A of detections, such that

$$\forall d_i, d_j \in \mathcal{D}_A : d_i \leq d_j \Leftrightarrow s\mathcal{RT}(d_i) \leq s\mathcal{RT}(d_j).$$

Then:

$$\begin{aligned} \tilde{\mathcal{D}}_A = \mathcal{D}_A \setminus \{d_f \mid \exists i, j, k < f : d_f < d_i < d_j < d_k \\ \vee \exists i, j, k > f : d_i < d_j < d_k < d_f\}. \end{aligned}$$

Detection confidence

For each detection d_j , the detection confidence $\mathcal{T}(d_j)$ was calculated as the maximum value of the moving average, with size $\gamma = 20$, of the classification confidence values $t(s)$ for the consecutive data points s from d_j , i.e.:

$$\mathcal{T}(d_j) = \max_{(s_i, \dots, s_{i+\gamma-1}) \subset d_j} \left\{ \frac{t(s_i) + \dots + t(s_{i+\gamma-1})}{\gamma} \right\}.$$

The value $\gamma = 20$ was derived as explained for the duration rule above.

Uniqueness rule

Since each VOC may be present in a GC-MS sample at most once, if multiple detections of one VOC occur, $d_j^1, \dots, d_j^{K_j}$, the one with the highest detection confidence value is kept in the set of detections:

$$d_j(A) = \arg \max_{d_j^k \in \mathcal{D}_A} \{\mathcal{T}(d_j^k) \mid k = 1, \dots, K_j\},$$

$$\tilde{\mathcal{D}}_A = \mathcal{D}_A \setminus \{d_j^k \neq d_j(A) \mid k = 1, \dots, K_j, j = 1, \dots, 30\}.$$

Phase B output

As an output, phase B produces a set \mathcal{L} of detections $d_j \in \tilde{\mathcal{D}}_A$ as 4-tuples of label, start detection, end detection and detection confidence:

$$\begin{aligned} \mathcal{L} = \{[j, \quad s\mathcal{RT}(d_j), \quad e\mathcal{RT}(d_j), \quad \mathcal{T}(d_j)] : \\ d_j \in \tilde{\mathcal{D}}_A, j = 1, \dots, 30\}. \end{aligned}$$

RT range of a VOC

Each target VOC in the GC-MS dataset is narrowly distributed over a specific range of the RT dimension. From the processed data files, the RT range of occurrences was extracted for each target VOC j across all its instances in the GC-MS dataset reported by expert-led

processing, i.e.:

$$RTrange_j = [\min\{peakRT_j\}, \max\{peakRT_j\}].$$

What is important, such *RT* ranges depend on the GC column. During the course of our study, the GC column was changed and thus we extracted two *RT* ranges for each VOC, $RTrange_j^1$ and $RTrange_j^2$, valid respectively before and after column change. S1 Table in [S1 File](#) presents *RT* ranges for each target VOC.

Tentative true positives

Let \hat{d}_j be the detections in the set \mathcal{D}_A that were not reported in the ground truth at the *RT* position reported by the system for that sample *A*. Note that we consider the set \mathcal{D}_A , i.e., before reducing the number of detections for the order and uniqueness rules, as their application changes the detection distribution. We observed that most (60% to 74% depending on the network) detections $\hat{d}_j \in \bigcup_A \mathcal{D}_A$ were reported by the system in the *RT* range of the *j* VOC (accordingly before or after column change, $i = 1$ or 2), i.e.:

$$DI(\hat{d}_j) \cap RTrange_j^i \neq \emptyset,$$

Let *I* denote an interval of length

$$|I| \leq \max_{\hat{d}_j \in \bigcup_A \mathcal{D}_A} |DI(\hat{d}_j)|,$$

chosen uniformly at random from within the *RT* dimension (of length *R*). Then the probability of the intersection of the interval *I* with a given $RTrange_j^i$ can be computed as:

$$P(I \cap RTrange_j^i \neq \emptyset) = \frac{|RTrange_j^i| + |I|}{R - |I|} \leq \frac{\max_{i,j} |RTrange_j^i| + \max_{\hat{d}_j} |DI(\hat{d}_j)|}{R - \max_{\hat{d}_j} |DI(\hat{d}_j)|} =: P_{\max},$$

where the $\max_{\hat{d}_j \in \bigcup_A \mathcal{D}_A}$ depends on the network used. P_{\max} is the upper bound probability of a random (false) detection $\hat{d}_j \in \mathcal{D}_A$ of a target VOC *j* at specific *RT* range in the sample *A*. We computed the following bounds:

Network	P_{\max}
VGG-8-1D	0.0431
DenseNet-40-1D	0.0439
ResNet-38-2D	0.0441

The probabilities P_{\max} are all very low in comparison to the actual proportion of such considered detections (i.e., 60% to 74% depending on the network). Therefore, it is highly likely that the detections \hat{d}_j within a compatible *RT* range are not random false detections, but are correctly detected VOCs. Such detections are named *tentative true positives* (TTP).

Tentative true negatives

Several VOCs, not detected at the specific *RT* points reported by the ground truth (preliminary marked as false negatives), were detected as tentative true positives on the different *RT*

positions within their compatible RT range. Because of the high chance of tentative true positives to be true positives, the noisy-labelled character of the ground truth and the constraints that a VOC appears only once in a sample, we conclude that such particular examples may indicate errors in the expert-led processing. Such examples are called *tentative true negatives*.

Average precision score

The system performance of each network was assessed by an average precision score (AP) for each target VOC and mean average precision (mAP). For each target VOC j , a list of all system detections $d_j \in \bigcup_A \tilde{\mathcal{D}}_A$ of the VOC in all clinical samples from the testing set was derived from stage 2 output and sorted in descending order associated with the detection confidence scores $T(d_j)$ (for Intersection of the models, it was sorted by mean value of detection confidence scores for tested models). For the first n elements of this list, the precision function $Precision(n)$ was defined as the proportion of TP. The recall function $Recall(n)$ was defined as the proportion of all detections in the ground truth that appear in the first n elements of the system detection list. As usual, the AP score for each target VOC was calculated as the integral under the graph of precision against recall (S4.1-S7.30 Figs in [S1 File](#)). mAP was calculated as the average of AP values across all target VOCs. AP values for each target VOC for each tested model, along with the respective graphs of precision vs recall functions, are given in the [S1 File](#).

Sensitivity and specificity

The system's sensitivity was computed for the *expert* benchmark (i.e., TTP are considered FP, TTN are considered FN) as

$$sensitivity = \frac{TP}{TP + FN + TTN},$$

and for the *system-derived correction* benchmark (i.e., TTP are considered TP, TTN are considered TN) as

$$sensitivity = \frac{TP + TTP}{TP + TTP + FN}.$$

The system specificity was computed for the *expert* benchmark (i.e., TTP are considered FP) as

$$specificity = \frac{TN}{TN + FP^* + TTP^*},$$

and for the *system-derived correction* benchmark (i.e., TTP are considered TP) as

$$specificity = \frac{TN}{TN + FP^*}.$$

Resources

All experiments were run on a server running Linux Ubuntu with 20 cores, 128GB RAM and NVIDIA Tesla K80 GPU cards. [Table 5](#) compares the training time (stage 1, [Fig 2](#)) and average scanning time (stage 2) of each of the tested network architectures. Memory requirements are given in S3.1-S3.11 Tables in [S1 File](#).

Table 5. Time resources: Tcime of a network training on the fully-augmented VOC dataset and scanning time of a single GC-MS sample.

Network	Training time	Scan time
VGG-8-1D	23 h	2 min
DenseNet-40-1D	37 h	10 min
ResNet-38-2D	80 h	12.5 min

<https://doi.org/10.1371/journal.pone.0265399.t005>

Supporting information

S1 File. The combined file with the supplementary tables and figures.
(PDF)

Acknowledgments

We are thankful to Iain Phillips and Yaser Al Khalifah for sharing their work on VOC clustering, Yang Hu for discussions on deep neural networks, Joanna Turner for discussions on the analogy with object detection.

Author Contributions

Conceptualization: Angelika Skarysz, Dahlia Salman, Martin Sykora, C. L. Paul Thomas, Andrea Soltoggio.

Data curation: Dahlia Salman.

Formal analysis: Angelika Skarysz, Dahlia Salman, Martin Sykora, Eugénie Hunsicker, Andrea Soltoggio.

Funding acquisition: C. L. Paul Thomas.

Investigation: Angelika Skarysz, Andrea Soltoggio.

Methodology: Angelika Skarysz, Dahlia Salman, Martin Sykora, C. L. Paul Thomas, Andrea Soltoggio.

Project administration: C. L. Paul Thomas.

Resources: Michael Eddleston, William H. Nailon, Kareen Darnley, Duncan B. McLaren.

Software: Angelika Skarysz.

Supervision: Martin Sykora, C. L. Paul Thomas, Andrea Soltoggio.

Validation: Dahlia Salman, Eugénie Hunsicker, Andrea Soltoggio.

Visualization: Angelika Skarysz.

Writing – original draft: Angelika Skarysz, Andrea Soltoggio.

Writing – review & editing: Dahlia Salman, Michael Eddleston, Martin Sykora, Eugénie Hunsicker, William H. Nailon, Kareen Darnley, Duncan B. McLaren, C. L. Paul Thomas.

References

1. Smolinska A., Hauschild A., Fijten R., Dallinga J., Baumbach J. & Schooten F. Current breathomics—a review on data pre-processing techniques and machine learning in metabolomics breath analysis. *Journal Of Breath Research*. 8, 27105 (2014), <http://www.ncbi.nlm.nih.gov/pubmed/24713999> PMID: 24713999

2. Hollywood K., Brison D. & Goodacre R. Metabolomics: Current technologies and future trends. *Proteomics*. 6 pp. 4716–4723 (2006) <https://doi.org/10.1002/pmic.200600106> PMID: 16888765
3. Rattray N., Hamrang Z., Trivedi D., Goodacre R. & Fowler S. Taking your breath away: Metabolomics breathes life in to personalized medicine. *Trends In Biotechnology*. 32 pp. 538–548 (2014) <https://doi.org/10.1016/j.tibtech.2014.08.003> PMID: 25179940
4. Van Berkel J., Dallinga J., Möller G., Godschalk R., Moonen E., Wouters E., et al. A profile of volatile organic compounds in breath discriminates COPD patients from controls. *Respiratory Medicine*. 104, 557–563 (2010) <https://doi.org/10.1016/j.rmed.2009.10.018> PMID: 19906520
5. Li W., Liu Y., Liu Y., Cheng S. & Duan Y. Exhaled isopropanol: new potential biomarker in diabetic breathomics and its metabolic correlations with acetone. *RSC Advances*. 7, 17480–17488 (2017) <https://doi.org/10.1039/C7RA00815E>
6. Fuchs P., Loeseken C., Schubert J. & Miekisch W. Breath gas aldehydes as biomarkers of lung cancer. *International Journal Of Cancer*. 126, 2663–2670 (2010) PMID: 19839051
7. Altomare D., Di Lena M., Porcelli F., Trizio L., Travaglio E., Tutino M., et al. Exhaled volatile organic compounds identify patients with colorectal cancer. *British Journal Of Surgery*. 100 pp. 144–150 (2012) <https://doi.org/10.1002/bjs.8942>
8. Phillips M., Cataneo R., Saunders C., Hope P., Schmitt P. & Wai J. Volatile biomarkers in the breath of women with breast cancer. *Journal Of Breath Research*. 4, 026003 (2010) <https://doi.org/10.1088/1752-7155/4/2/026003> PMID: 21383471
9. Ruszkiewicz, D., Sanders, D., O'Brien, R., Hempel, F., Reed, M., Riepe, A., et al. Diagnosis of COVID-19 by Analysis of Breath with Gas Chromatography-Ion Mobility Spectrometry—A Feasibility Study.. *SSRN Electronic Journal*. (2020)
10. Watson, J. & Sparkman, O. Introduction to Mass Spectrometry: Instrumentation, Applications and Strategies for Data Interpretation: Fourth Edition. *Introduction To Mass Spectrometry: Instrumentation, Applications And Strategies For Data Interpretation: Fourth Edition*. pp. 1–819 (2008)
11. Stein S. An integrated method for spectrum extraction and compound identification from gas chromatography/mass spectrometry data. *Journal Of The American Society For Mass Spectrometry*. 10, 770–781 (1999) [https://doi.org/10.1016/S1044-0305\(99\)00047-1](https://doi.org/10.1016/S1044-0305(99)00047-1)
12. Hübschmann, H. Handbook of GC-MS. *Handbook Of GC-MS*. (2015)
13. The Kováts Retention Index System *Analytical Chemistry* (2012)
14. Colby B. Spectral deconvolution for overlapping GC/MS components. *Journal Of The American Society For Mass Spectrometry*. 3, 558–562 (1992) [https://doi.org/10.1016/1044-0305\(92\)85033-G](https://doi.org/10.1016/1044-0305(92)85033-G) PMID: 24234499
15. Alkhalifah Y., Phillips I., Soltoggio A., Darnley K., Nailon W., McLaren D., et al. VOCCluster: Untargeted metabolomics feature clustering approach for clinical breath gas chromatography/mass spectrometry data. *Analytical Chemistry*. 92, 2937–2945 (2019) <https://doi.org/10.1021/acs.analchem.9b03084>
16. Ren S., Hinzman A., Kang E., Szczesniak R. & Lu L. Computational and statistical analysis of metabolomics data. *Metabolomics*. 11 pp. 1492–1513 (2015) <https://doi.org/10.1007/s11306-015-0823-6>
17. Coombes K., Baggerly K. & Morris J. Pre-processing mass spectrometry data. *Fundamentals Of Data Mining In Genomics And Proteomics*. pp. 79–102 (2007) https://doi.org/10.1007/978-0-387-47509-7_4
18. Likić V. Extraction of pure components from overlapped signals in gas chromatography-mass spectrometry (GC-MS). *BioData Mining*. 2 (2009) <https://doi.org/10.1186/1756-0381-2-6> PMID: 19818154
19. Sajda Paul Machine learning for detection and diagnosis of disease.. *Annual Review Of Biomedical Engineering*. 8, 8.1–8.29 (2006), <http://liinc.bme.columbia.edu/publications/SajdaAnnRevBME.pdf> PMID: 16834566
20. Mamoshina P., Vieira A., Putin E. & Zhavoronkov A. Applications of Deep Learning in Biomedicine. *Molecular Pharmaceutics*. 13 pp. 1445–1454 (2016) <https://doi.org/10.1021/acs.molpharmaceut.5b00982> PMID: 27007977
21. Baranska A., Tigchelaar E., Smolinska A., Dallinga J., Moonen E., Dekens J., et al. Profile of volatile organic compounds in exhaled breath changes as a result of gluten-free diet. *Journal Of Breath Research*. 7 (2013) <https://doi.org/10.1088/1752-7155/7/3/037104> PMID: 23774130
22. Le Cun Y., Boser B., Denker J., Henderson D., Howard R., Hubbard W., et al. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*. 1 pp. 541–551 (1989) <https://doi.org/10.1162/neco.1989.1.4.541>
23. Rawat W. & Wang Z. Deep convolutional neural networks for image classification: A comprehensive review. *Neural Computation*. 29 pp. 2352–2449 (2017) https://doi.org/10.1162/neco_a_00990 PMID: 28599112

24. LeCun, Y., Huang, F. & Bottou, L. Learning Methods for Generic Object Recognition with Invariance to Pose and Lighting. *Computer Vision And Pattern Recognition, 2004. CVPR 2004. Proceedings Of The 2004 IEEE Computer Society Conference On.* 2 pp. II-97–104 (2004)
25. Cireřan D., Meier U. & Schmidhuber J. Multi-column Deep Neural Networks for Image Classification. *International Conference Of Pattern Recognition.*, 3642–3649 (2012)
26. Krizhevsky A., Sutskever I. & Hinton G. ImageNet Classification with Deep Convolutional Neural Networks. *Nips.* pp. 1–9 (2012)
27. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R. & LeCun, Y. OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks. *ArXiv Preprint ArXiv.* pp. 1312.6229 (2013), <http://arxiv.org/abs/1312.6229>
28. Nielsen M. Neural Networks and Deep Learning. *Machine Learning.* pp. 875–936 (2015), <http://neuralnetworksanddeeplearning.com/chap4.html>
29. Garcia A. & Barbas C. Gas chromatography-mass spectrometry (GC-MS)-based metabolomics.. *Methods In Molecular Biology (Clifton, N.J.).* 708 pp. 191–204 (2011) https://doi.org/10.1007/978-1-61737-985-7_11 PMID: 21207291
30. Skarysz, A., Alkhalifah, Y., Darnley, K., Eddleston, M., Hu, Y., McLaren, D., et al. Convolutional neural networks for automated targeted analysis of raw gas chromatography-mass spectrometry data. *Proceedings Of The International Joint Conference On Neural Networks.* 2018–July (2018)
31. Cortes C. & Vapnik V. Support-Vector Networks. *Machine Learning.* 20, 273–297 (1995) <https://doi.org/10.1007/BF00994018>
32. Zhang G. Neural networks for classification: a survey. *IEEE Transactions On Systems, Man And Cybernetics, Part C (Applications And Reviews).* 30, 451–462 (2000), <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=897072>
33. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *ArXiv Preprint ArXiv:1409.1556.* pp. 1–13 (2014), <http://arxiv.org/abs/1409.1556>
34. He K., Zhang X., Ren S. & Sun Jian. Deep Residual Learning for Image Recognition. *Multimedia Tools And Applications.* pp. 1–17 (2016)
35. Huang, G., Liu, Z., Weinberger, K. & Maaten, L. Densely connected convolutional networks. *Proceedings Of The IEEE Conference On Computer Vision And Pattern Recognition.* 1, 3 (2017)
36. Hastie, T., Tibshirani, R. & Friedman, J. The Elements of Statistical Learning (Second Edition, 10th print). *Springer.* 1 pp. 337–387 (2009), <http://www.springerlink.com/index/10.1007/b94608>
37. Dyk D. & Meng X. The art of data augmentation. *Journal Of Computational And Graphical Statistics.* 10, 1–50 (2001) <https://doi.org/10.1198/10618600152418584>
38. Everingham M., Van Gool L., Williams C., Winn J. & Zisserman A. The pascal visual object classes (VOC) challenge. *International Journal Of Computer Vision.* 88, 303–338 (2010) <https://doi.org/10.1007/s11263-009-0275-4>
39. TOXI-triage project, <http://toxi-triage.eu>
40. Moore R. & Karasek F. GC/MS identification of organic pollutants in the caroni river, trinidad. *International Journal Of Environmental Analytical Chemistry.* 17, 203–221 (1984) <https://doi.org/10.1080/03067318408076973>
41. Keto R. & Wineman P. Detection of Petroleum-Based Accelerants in Fire Debris by Target Compound Gas Chromatography/Mass Spectrometry. *Analytical Chemistry.* 63, 1964–1971 (1991) <https://doi.org/10.1021/ac00018a013>
42. Lee J., Park J., Go A., Moon H., Kim S., Jung S., et al. Urine Multi-drug Screening with GC-MS or LC-MS-MS Using SALLE-hybrid PPT/SPE. *Journal Of Analytical Toxicology.* 42, 617–624 (2018) <https://doi.org/10.1093/jat/bky032> PMID: 29762685
43. Tsivou M., Kioukia-Fougia N., Lyris E., Aggelis Y., Fragkaki A., Kiouisi X., et al. An overview of the doping control analysis during the Olympic Games of 2004 in Athens, Greece. *Analytica Chimica Acta.* 555, 1–13 (2006) <https://doi.org/10.1016/j.aca.2005.08.068>
44. Krasnopolsky V. & Parshev V. Chemical composition of the atmosphere of Venus. *Nature.* 292, 610–613 (1981) <https://doi.org/10.1038/292610a0>
45. Tekin K., Karagöz S. & Bektaş S. A review of hydrothermal biomass processing. *Renewable And Sustainable Energy Reviews.* 40 pp. 673–687 (2014) <https://doi.org/10.1016/j.rser.2014.07.216>
46. Bianchi F., Careri M., Musci M. & Mangia A. Fish and food safety: Determination of formaldehyde in 12 fish species by SPME extraction and GC-MS analysis. *Food Chemistry.* 100, 1049–1053 (2007) <https://doi.org/10.1016/j.foodchem.2005.09.089>

47. Garruti D., Franco M., Da Silva M., Janzantti N. & Alves G. Assessment of aroma impact compounds in a cashew apple-based alcoholic beverage by GC-MS and GC-olfactometry. *LWT—Food Science And Technology*. 39, 373–378 (2006)
48. Van Asten A. The importance of GC and GC-MS in perfume analysis. *TrAC—Trends In Analytical Chemistry*. 21, 698–708 (2002) [https://doi.org/10.1016/S0165-9936\(02\)00807-5](https://doi.org/10.1016/S0165-9936(02)00807-5)
49. Thomas, C. D3.1 Prototype sampling system for reproducible non-invasive clinical sampling protocol. . (2016)
50. Salek R., Steinbeck C., Viant M., Goodacre R. & Dunn W. The role of reporting standards for metabolite annotation and identification in metabolomic studies. *GigaScience*. 2 (2013) <https://doi.org/10.1186/2047-217X-2-13> PMID: 24131531
51. Shrivastava, A., Gupta, A. & Girshick, R. Training region-based object detectors with online hard example mining. *Proceedings Of The IEEE Computer Society Conference On Computer Vision And Pattern Recognition*. 2016-December pp. 761–769 (2016)
52. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C., et al. SSD: Single shot multibox detector. *Lecture Notes In Computer Science (including Subseries Lecture Notes In Artificial Intelligence And Lecture Notes In Bioinformatics)*. 9905 LNCS pp. 21–37 (2016)