

Accelerating Plasmonic Hydrogen Sensors for Inert Gas Environments by Transformer-Based Deep Learning

Viktor Martvall,¹ Henrik Klein Moberg,¹ Athanasios Theodoridis,¹ David Tomeček,¹ Pernilla Ekborg-Tanner,¹ Sara Nilsson,¹ Giovanni Volpe,² Paul Erhart,^{1,*} and Christoph Langhammer^{1,†}

¹*Department of Physics, Chalmers University of Technology, SE-41296, Göteborg, Sweden*

²*Department of Physics, University of Gothenburg, 412 96 Göteborg, Sweden*

The ability to rapidly detect hydrogen gas upon occurrence of a leak is critical for the safe large-scale implementation of hydrogen (energy) technologies. However, to date, no technically viable sensor solution exists that meets the corresponding response time targets set by stakeholders at technically relevant conditions. Here, we demonstrate how a tailored Long Short-term Transformer Ensemble Model for Accelerated Sensing (LEMAS) accelerates the response of a state-of-the-art optical plasmonic hydrogen sensor by up to a factor of 40 in an oxygen-free inert gas environment, by accurately predicting its response value to a hydrogen concentration change before it is physically reached by the sensor hardware. Furthermore, it eliminates the pressure dependence of the response intrinsic to metal hydride-based sensors, while leveraging their ability to operate in oxygen-starved environments that are proposed to be used for inert gas encapsulation systems of hydrogen installations. Moreover LEMAS provides a measure for the uncertainty of the predictions that is pivotal for safety-critical sensor applications. Our results thus advertise the use of deep learning for the acceleration of sensor response, also beyond the realm of plasmonic hydrogen detection.

INTRODUCTION

The ability to detect, quantify and distinguish chemical species accurately and rapidly is crucial for technologies requiring swift data capture to support well informed decision-making, automation and process-monitoring. Such technologies span a wide range of applications, including environmental monitoring [1], biosensing for real-time disease diagnostics [2], chemical process control [3] and food quality evaluation [4]. They all have in common that they critically rely on the development of sensors that are not only precise, sensitive and selective but also respond rapidly to their target substance and are able to deliver an accurate quantitative measure of the concentration of that target.

A domain that is rapidly expanding and where sensing will play a pivotal role in facilitating safe large scale implementation is hydrogen-based technologies, including fuel cells for heavy transport, shipping and aviation, energy storage solutions and green steel production. They all have in common that they promise substantial reductions of greenhouse gas emissions. However, this prospect also generates new demands for active process monitoring and control, and introduces safety concerns owing to the high flammability of H₂-air mixtures. All of these issues can be effectively addressed by the development of fast and accurate H₂ sensors.

From a sensing environment perspective, two distinct settings exist, where ambient conditions characterized by an abundance of oxygen constitutes the most obvious one. The second setting, which is of significant technological relevance but much less discussed in the scientific literature to date, is so-called “inert” or “oxygen starved” environments. They are established to encapsulate/enclose large scale H₂ installations, such as entire en-

gine rooms on fuel-cell powered ships, or fuel pipes on H₂-powered airplanes, to avoid the formation of flammable air-H₂ mixtures. The rapid detection of even the tiniest H₂ leaks inside these inert gas encapsulation infrastructures is critical to provide enough time for the implementation of appropriate measures to eliminate, as well as to spatially localize, the leak by placing a large number of sensors at strategic locations inside the system. Specifically, in such installations the system is continuously flushed by an inert gas, such as N₂ or Ar, to eliminate the presence of molecular oxygen. Importantly, we note that the inert gas used in such systems will be of low quality from a purity perspective with respect to species such as H₂O, CO or SO_x, for cost reasons. This combination of lack of O₂ and presence of sizable amounts of “poisoning” molecules that bind strongly to many sensor surfaces poses a significant challenge because (i) established H₂ sensors of the catalytic and thermal type require O₂ to work and (ii) because the strong molecular bonds either block/poison surface sites required for H₂ dissociation and/or detection, or facilitate surface reactions that consume hydrogen species and thus prevent them from being detected [5].

To steer the development of next-generation H₂ sensors that meet the upcoming demands of widely implemented H₂ technologies outlined above, agencies and stakeholders have defined performance targets. The most well-know ones are defined by the U.S. Department of Energy (DOE), which identify sensor speed at ambient conditions as one of the key unresolved metrics [6]. To this end, a small number of studies exist where H₂ sensors with response times just below 1 s for a 1 mbar H₂ pulse have been demonstrated experimentally [7–9]. However, while indeed important breakthroughs, these demonstrations were made in an idealized pure H₂-vacuum environment

that constitutes a severe simplification. As the main reason for this simplification we identify the aforementioned challenge of “poisoning” molecular species in technologically relevant sensing environments due to their impact on the surface chemistry of a sensor. Hence, even though these demonstrations of H_2 detection with sub-second response in idealized vacuum/pure H_2 conditions exist, it is clear that further advances in this field are necessary [5].

Traditionally, such advances are attempted by developing new sensing materials, by nanostructuring the sensing materials and/or signal transducers, and by the refinement or modification of fundamental physical sensing mechanisms [7, 9–14]. Interestingly, however, only very limited attention has been directed towards harnessing the potential of tailoring the treatment of output data of existing sensor platforms with the aim to improve sensor response, e.g., by machine learning techniques. Accordingly, only few recent studies have leveraged the potential of machine learning to enhance the *accuracy* or *sensitivity* of different kinds of gas sensors [15, 16], leaving the potential to enhance *sensor response times* still largely unexplored.

In response to and motivated by the high demand for faster sensors in general and H_2 sensors for inert gas environments in particular, here, we develop an approach for accelerating H_2 -sensing that combines optical nanoplasmonic sensors based on hydride-forming metal nanoparticles, such as Pd and its alloys with coinage metals [5, 17, 18], that enable operation in oxygen-starved environments, with deep learning. We show that this combination enables predicting the H_2 concentration in an inert gas (far) more quickly than the conventional approach, which is limited by the need to reach full thermodynamic equilibration of the sensor after a change in H_2 concentration and hampered by sensor deactivation effects due to the presence of molecular contaminants.

To analyze the output data of plasmonic hydrogen sensors, which typically consists of time series of scattering or extinction spectra in the visible light spectral range [5, 18], the current state of the art (SotA) analysis widely applied in the field collapses each such measured spectrum to a single spectral descriptor, such as the spectral peak position, the full-width half maximum or the centroid position [19]. In other words, the (potentially) vast amount of information contained in the full spectrum is lost in this analysis.

To harness this information with the aim to accelerate plasmonic H_2 sensor response in inert gas environment, here, we introduce LEMAS, short for Long Short-term Transformer Ensemble Model for Accelerated Sensing, which improves the sensor speed by learning the relationship between the time dependence of the full spectrum and the H_2 concentration, while simultaneously assessing uncertainty in the model predictions through model ensembles. The long short-term transformer (LSTR) ar-

chitecture consists of a long and short-term memory and has been demonstrated to be well suited for modeling long time sequences [20]. We demonstrate that LEMAS reduces the response time of a $Pd_{70}Au_{30}$ alloy plasmonic H_2 sensor by up to 40 times when exposed to a distinct H_2 pulse down to 0.06 vol.% H_2 in an inert gas environment at atmospheric pressure, in a scenario simulating a sudden large leak. Furthermore, we illustrate the ability of LEMAS to rapidly discern and quantify slow gradual changes in H_2 concentration from mere noise in a simulated scenario of detecting a small leak in an enclosed inert gas environment. This ability is critical for detecting H_2 at as low concentrations as possible as quickly as possible, allowing sufficient time to apply safety measures, such as system shutdown, before a safety-critical H_2 concentration is reached. Finally, as an ensemble model LEMAS enables one to obtain uncertainty estimates, which is of fundamental importance for safety-critical applications, including but not limited to H_2 sensing.

RESULTS AND DISCUSSION

$Pd_{70}Au_{30}$ alloy plasmonic H_2 sensors. As the plasmonic H_2 sensor platform of choice, we selected the well-established $Pd_{70}Au_{30}$ alloy system, which we have investigated in detail earlier [7, 21–24]. This material system is especially suited for inert gas sensing environments since its sensing mechanism, the interstitial sorption of hydrogen into the lattice of the metal host, does not require O_2 to be present. The Au alloyant serves the purpose of eliminating the intrinsic hysteresis characteristic for pure Pd by lowering the critical temperature of the system [25–28]. At 30% Au the best compromise between completely eliminating hysteresis, establishing linear optical response to H_2 and maximizing optical contrast per unit sorbed H_2 is reached. Therefore, we nanofabricated quasi-random arrays of $Pd_{70}Au_{30}$ alloy nanodisks with a mean diameter of 210 nm and 25 nm height onto fused silica substrates using hole-mask colloidal lithography (Fig. 1a-b), following the procedures described in detail in our earlier work [29] and in Sect. in Sect. .

The working principle of plasmonic H_2 sensors is based on the localized surface plasmon resonance (LSPR) phenomenon, characteristic for metal nanoparticles irradiated by visible light. In an optical transmission, scattering or extinction spectrum, the LSPR manifests itself as a distinct peak with a maximum at a specific wavelength. The spectral position of this peak maximum, as well as related peak descriptors such as width and intensity, exhibits a linear dependence on the H_2 partial pressure surrounding the particles and on the amount of hydrogen species absorbed into interstitial lattice sites of the Pd or Pd alloy host (Fig. 1c-d) [23]. Since the ab- and desorption of hydrogen into and from these interstitial

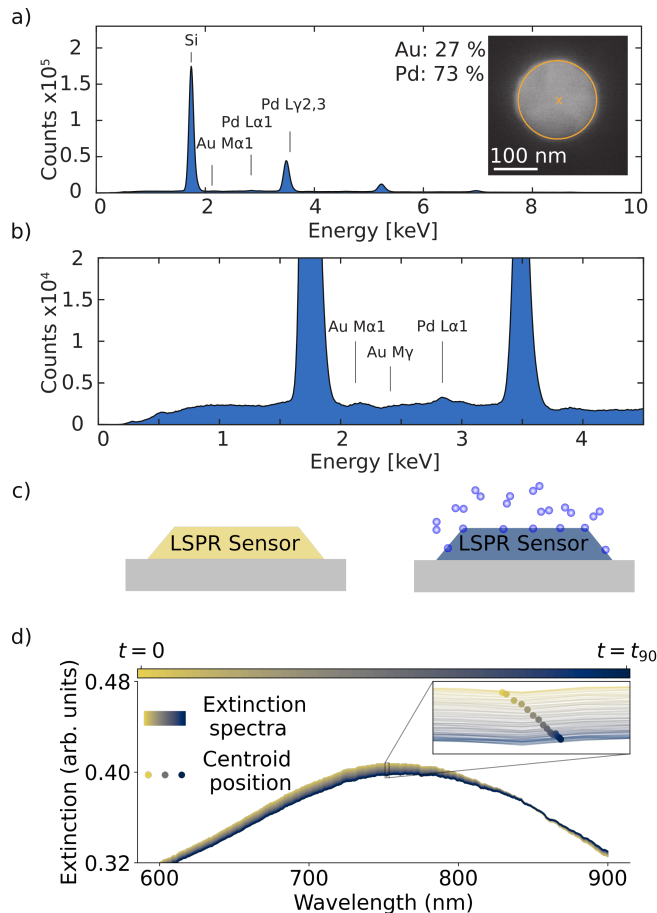


FIG. 1. $\text{Pd}_{70}\text{Au}_{30}$ alloy nanoparticle plasmonic sensor characterization and operating principle. (a) Energy-dispersive X-ray (EDX) spectrum collected from a single $\text{Pd}_{70}\text{Au}_{30}$ alloy nanodisk in the quasi-random array of such disks that constitutes the active sensor surface. (b) Zoom-in of the EDX-spectrum in (a) up to 4.5 keV to focus on the characteristic Pd and Au peaks. (c) Schematic illustration of the plasmonic H_2 sensing principle, where the sorption of hydrogen into hydride-forming metal nanoparticles induces a change in their localized surface plasmon resonance frequency, which leads to a color change that is resolved in a spectroscopic measurement in the visible light spectral range. (d) Example of the spectral response of the $\text{Pd}_{70}\text{Au}_{30}$ alloy plasmonic sensor used in this work, resolved as a gradual shift in the extinction spectrum as hydrogen is absorbed the crystal lattice. Inset: Temporal evolution of the peak centroid position, one of the spectral descriptors that can be tracked to enable real time H_2 detection.

lattice positions, respectively, occurs spontaneously and reversibly at ambient conditions, and also in oxygen-free environments, tracking of the spectral position (as well as other peak descriptors) of the LSPR peak as function of H_2 partial pressure enables real time H_2 detection (inset in Fig. 1d). In this work, for what we refer to as the SotA analysis, we use the centroid position as a spectral descriptor which we relate to the H_2 concentration by a

calibration function (see Sect. in Sect. for details).

Deep learning model selection. We base our choice of a LSTR model for accelerating the plasmonic H_2 sensor response on several key characteristics of the output data generated by this type of sensor (see Fig. 2 and Sect. in Sect. for details). The first important characteristic to take into account is that the measured extinction spectra that constitute the raw sensor response over time, exhibit intrinsic noise (due to intensity fluctuations of the halogen light source and detection noise of the spectrometer used) that is comparable to the magnitude of changes induced in the spectra by small variations in H_2 concentration. Consequently, a crucial criterion for selecting the deep learning model is its ability to accurately model long temporal sequences, enabling the differentiation between relevant temporal trends in the extinction spectrum and the inherent noise.

The second critical aspect influencing the performance of the LSTR model, based on the characteristics of the sensor data, is the pre-processing of the measured extinction spectra. Such pre-processing is needed due to drift in the sensor response over time (mainly due to long-term variations of light source intensity), as well as small variations in the extinction spectra obtained in different measurements using the same sensor, due to slightly different placement of the sensor in the measurement chamber for each independent experiment. Here, we found that using several pre-processing methods is beneficial for the performance of the LSTR model. Therefore, we used four different pre-processing techniques (see Eq. and Note S2 for details) and concatenated them into a single array [30]. As a result, the input data for the deep learning models was a time series, where each element in the sequences consisted of the concatenation of the different pre-processing techniques (see Fig. 2a).

Another modeling choice that we make is to employ an *ensemble* of LSTR models. This choice is motivated by the safety-critical nature of the hydrogen sensor application and yields a more robust prediction, as well as a measure of uncertainty by aggregating the predictions of several LSTR models to compute the mean and the standard deviation (see Sect. in Sect. for details). Combining these modeling choices, we arrive at LEMAS, characterized by an ensemble of LSTR models that can both rapidly predict the H_2 concentration and provide a measure of uncertainty from a time series of pre-processed spectra.

LEMAS model training and testing. Having introduced the architecture of the LEMAS model, we discuss the training and testing data used for optimizing the sensor response in (i) a large and fast leak scenario and (ii) a slow gradual leak scenario. These data were generated by measuring tailored time series of optical extinction spectra of the sensor localized in a custom-made measurement chamber with small volume to enable rapid gas exchange at atmospheric pressure to expose the sensor to




FIG. 2. **Long Short-term Transformer Ensemble Model for Accelerated Sensing.** Illustration of the deep learning model used in this work based on the LSTR architecture. (a) The input data to the model consists of a time sequence of the past evolution of the spectral response of the sensor. In this figure the time sequence consist of 600 time steps corresponding to 198 s. (b) The time sequence is first split into a long and short-term memory and pre-processed using four different methods, including wavelength dependent min-max normalization, standard normal variate standardization, global min-max normalization, and level scaling, before the concatenation of the pre-processed data is being fed to the LSTR. (c) The LSTR firstly compresses the long-term memory to a fixed length latent representation in the LSTR encoder. Secondly, the LSTR decoder extracts relevant temporal features in the short term memory while also querying the compressed long-term memory. The extracted temporal features are then passed through a stack of MLP layers to obtain a prediction of the current H_2 concentration.

varying H_2 concentrations in inert gas environment (see Sect. in Sect. and Fig. S1 for details). Specifically, we used three different H_2 profiles for generating the training data: (i) step-wise increase/decrease of H_2 from 0.00 vol.% H_2 to 0.06–1.97 vol.% H_2 in inert Ar environment (Fig. S4), (ii) linear increase/decrease of H_2 from 0.06 vol.% H_2 to 0.09–1.97 vol.% H_2 in inert Ar environment (Fig. S5a) and (iii) exponential increase/decrease of H_2 from 0.06 vol.% H_2 to 0.09–1.97 vol.% H_2 in inert Ar environment (Fig. S5b), see Note S1 for details.

For the first case of a large simulated leak characterized by a rapid step-wise increase of H_2 concentration in the sensor surroundings, we trained LEMAS using two independent measurements of step-wise H_2 concentration increase/decrease and subsequently tested the trained LEMAS model on a third measurement not used for training. For the second case of a simulated small slow leak, we trained LEMAS on one measurement of linear H_2 concentration increase/decreases and tested the performance of the trained model on one measurement with exponential H_2 increase/decreases (see Sect. for details).

The models trained for optimizing the sensor response in a large and fast leak scenario used a total input sequence length corresponding to the past 3 min, whereas the models trained for optimizing the sensor response in a slow gradual leak scenario used an input sequence length corresponding to the past 22 min of the sensor history. These choices were made based on an analysis of the change in centroid position in the training data and an estimation of the length of the time sequence needed to differentiate the slowest occurring process in the sensor output data from the noise in the measurement (see

Fig. S10 and Fig. S11 for details).

Accelerating sensor response to a simulated large leak in inert Ar environment. To assess the ability of LEMAS to accelerate the response of a plasmonic H_2 sensor, we first consider a scenario where a 0.06% H_2 pulse in inert Ar gas is applied to our device at 30 °C (Fig. 3a). For this analysis, we define the response time t_{90} as the first point in time where the sensor response has reached 90% of its new steady state value. Applying first the SotA analysis that tracks the centroid position, reveals that it takes on the order of 85 s to reach t_{90} . Deploying the LEMAS analysis on the same data, shows that it is able to predict the saturated H_2 level after only 3.6 s and thus long before the response of the system has saturated, leading to a more than 20-fold reduction of the response time. This result is corroborated when comparing t_{90} values obtained by SotA and LEMAS analysis across a range of H_2 pressure pulses from 0.06 vol.% H_2 to 1.97 vol.% H_2 (Fig. 3b and Fig. S15).

Remarkably, LEMAS also achieves a response time that is practically independent of H_2 concentration whereas the t_{90} from the SotA analysis quickly increases with decreasing H_2 concentration. We attribute this behavior to the fact that LEMAS only requires a certain number of data points to make its prediction, the availability of which is dictated by the read-out frequency of the spectrometer rather than the H_2 pressure. By contrast, the SotA analysis is limited by the intrinsic kinetics of the material platform, causing a strong dependence on the H_2 pressure. As a result, LEMAS yields the largest boost in acceleration in the application critical range of lower H_2 pressures and LEMAS, overcoming one of the




FIG. 3. **Accelerating sensor response to a simulated large leak in inert gas environment.** (a) Comparison of the prediction of LEMAS and the SotA centroid analysis for a pulse of 0.06 vol.% H₂ in inert Ar environment. By utilizing the full time-dependent spectrum of the measured sensor response LEMAS is able to accurately predict the final value H₂ concentration before the sensor physically reaches its new state in equilibrium with the new H₂ concentration level. (b) Comparison of response times obtained by LEMAS and the SotA centroid analysis as a function of H₂ concentration in inert Ar environment. Note the significant acceleration by LEMAS, in particular at the lowest H₂ concentrations, and the elimination of the concentration dependence of the response.

most important intrinsic limiting factors of hydride-based H₂ sensors.

Specifically, for the smallest concentrations considered in our experiment, at 0.1 vol.% H₂ and below, the response times range between 1.6 s – 3.6 s for the LEMAS analysis compared to 50 s – 85 s for the SotA analysis. This corresponds to an 21 – 40 fold improvement compared to the SotA. At the same time, we also note that

even the accelerated response obtained by LEMAS in the present inert gas conditions, is slower than the state-of-the-art in vacuum/H₂ environment without acceleration [7–9]. As the main reasons, we identify the following points: (i) The traces of poisoning species such as H₂O, CO, etc. present in the Ar inert gas used, significantly decelerate the sensor, as expected[7] (see Fig. S2 and Fig. S3 for quantitative mass spectrometric analysis of the background molecular species present in the Ar inert gas used). (ii) We have used relatively large nanoparticles, and it is known that reducing size increases sensor speed due to reduced hydrogen diffusion path lengths[7]. (iii) We have not applied any polymer coatings, which are known to accelerate sensor response, as well as protect them from the poisoning molecular species present in the inert gas [7].

Furthermore, we highlight that the amount of response time acceleration that LEMAS can produce not only depends on the obvious intrinsic response speed of the active sensor material (in our case the PdAu alloy nanoparticles) but also on the sampling rate of the sensor hardware, where a higher sampling rate enables a larger degree of acceleration. For our experiments discussed so far, we have used a sampling frequency of 3 Hz, which was the highest rate enabled by the used spectrometer. Consequently, in this specific implementation, LEMAS has only three data points available to identify a change in the H₂ concentration in less than 1 s. Crucially, the acceleration observed in the present case is thus not limited by LEMAS but the underlying materials and read-out of the used light sampling device.

Having established the overall LEMAS concept and demonstrated its ability to substantially accelerate sensor speed in inert sensing environments, in particular in the low concentration regime, it is interesting to evaluate the performance of LEMAS in more detail. To do so, we select three different H₂ concentration pulses, i.e., pulses to 0.06, 0.08, and 0.10 vol.% H₂, and plot the sensor response predicted by LEMAS as a function of time, with the standard deviation of the prediction at each time point indicated in the corresponding graphs (Fig. 4a and Fig. S16 for all pulses). We also define the sensor settling time, t_s^{LEMAS} , as the first time point where the predicted H₂ concentration lies within $\pm 10\%$ of the target H₂ concentration and the relative standard deviation is smaller than 10%. This metric complements the response time by also considering cases where LEMAS either underestimates or overestimates the H₂ concentration after t_{90} .

First, we note that at the onset of each pulse, there is a brief interval where LEMAS predicts 0 vol.% H₂, while the actual H₂ concentration has already increased. This behavior occurs since the change in the extinction spectrum induced by the presence of H₂ is not yet distinguishable from the noise level in the measurement. This initial phase is followed by an interval where a clear change in the extinction spectrum is detected but where both er-

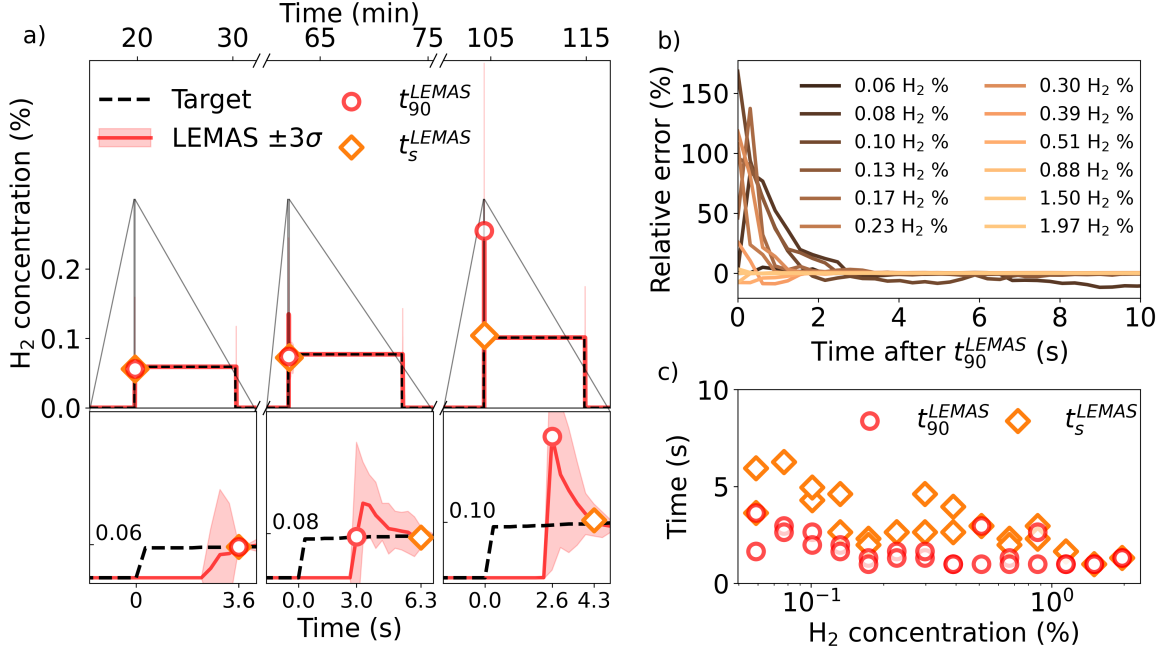


FIG. 4. **LEMAS prediction accuracy to H₂ pulses of different concentration.** (a) The prediction and standard deviation from LEMAS for three selected H₂ concentration pulses in the test set. The lower panels display zoom-ins on the initial response to the pulse. (b) The relative error of the LEMAS predictions for the entire range of H₂ concentration pulses, starting at the response time, t_{90}^{LEMAS} , and forward. (c) LEMAS settling and response times as a function of the H₂ concentration.

ror (Fig. 4b) and uncertainty are still rather large (red shaded areas in Fig. 4a). In the final phase, the LEMAS-predicted H₂ concentration settles at the correct value once the change in the extinction spectrum is sufficiently distinct, such that all models in the ensemble predict a similar H₂ concentration, and the uncertainty becomes very small.

Finally, we note that for some pulses (illustrated by 0.08 and 0.10 vol.% H₂ in Fig. 4a) the predicted H₂ concentration is overestimated for a brief interval past t_{90}^{LEMAS} , before the mean prediction settles around the target value. Consequently, t_s^{LEMAS} is larger than t_{90}^{LEMAS} . Conversely, in other cases (illustrated by the pulse to 0.06 vol.% H₂) the H₂ concentration is underestimated until t_{90}^{LEMAS} , at which point the uncertainty has also been reduced. As a result, t_s^{LEMAS} equals t_{90}^{LEMAS} in this (and similar) cases. To further examine the overestimations we analyze the relative error of the predicted response, starting from t_{90}^{LEMAS} (Fig. 4b). Overall the relative error tends to be larger for lower H₂ concentration pulses since there is a transient overestimation in the LEMAS prediction. Consequently, t_s^{LEMAS} is generally larger for lower H₂ concentrations (Fig. 4c). This is likely the consequence of the early predictions being more affected by measurement noise, since lower H₂ concentrations are associated with slower absorption kinetics and smaller changes in the extinction spectrum. An impor-

tant implication of these initial over-estimations is that the accuracy of the sensor initial response can be compromised if one relies on a single model, providing further evidence for the benefit of using an ensemble model, as we do with LEMAS.

Improving sensor response to a simulated small and slow leak in inert Ar environment. In a practical application in inert gas environment, H₂ sensors are not only required for the rapid detection of large leaks with fast and essentially instantaneous increase of H₂ concentration, but also in scenarios where a small leak will lead to a slow increase in H₂ concentration in an enclosed environment over time. Technically, this translates into the challenge of being able to as quickly as possible discern a tiny sensor signal from noise. To address this scenario in the LEMAS framework, we define the limit of detection (LOD) of a sensor as the minimal amount of H₂ required for the mean H₂ prediction to change by more than three times the standard deviation of the H₂ prediction at a baseline where the H₂ concentration is kept constant. In other words, the smallest H₂ required to discern (but not quantify) the presence of hydrogen gas with a confidence of 3σ . Furthermore, we define the limit of quantification (LOQ) as the minimal amount of H₂ for which the mean relative error of the H₂ prediction is less than 5 %. The mean absolute relative error $\frac{1}{N} \sum_{t=1}^N \left| \frac{P_t - T_t}{T_t} \right|$ is calculated over a time window of 2 min, where N is the

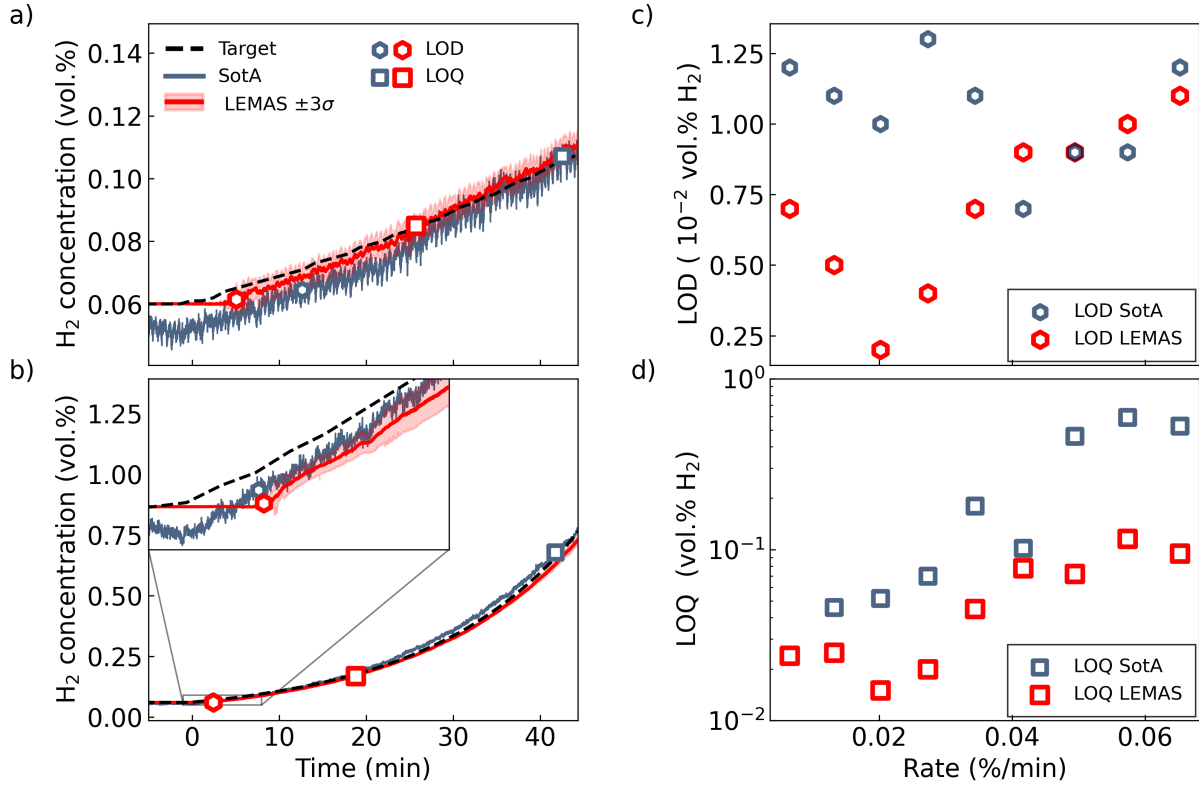


FIG. 5. **Small leak detection and quantification.** (a-b) Time evolution of the H₂ concentration obtained using the SotA centroid analysis and LEMAS, respectively, for exponentially increasing H₂ concentration in inert Ar environment for leak rates of (a) $1.32 \times 10^{-3} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ and (b) $5.73 \times 10^{-2} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$. The correspondingly obtained limit of detection (LOD) and limit of quantification (LOQ) are also indicated. (c-d) LOD and LOQ as a function of leak rate. Note that at the smallest leak rate the LOD is not reached for the SotA centroid analysis.

total number of time steps across the window, P_t is the predicted H₂ value for time step t and T_t is the true H₂ value for time-step t .

We then assess the sensor response to a first scenario with a very small exponential leak rate of $1.32 \times 10^{-3} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ using both the centroid-shift based SotA analysis and LEMAS (Fig. 5a). This analysis reveals that the LEMAS-predicted H₂ concentrations contain considerably less noise compared to the SotA analysis. In the SotA analysis the noise in the H₂ signal is initially comparable to the change in H₂ concentration. As a result, the LOD is reached faster, i.e., at lower H₂ concentrations, for LEMAS analysis due to its ability to faster discern changes in the H₂ concentration. We also note that the LEMAS analysis generally delivers a more accurate response, where the predicted H₂ concentration values are closer to the target values, resulting in the LOD being reached earlier for LEMAS.

We also perform a similar analysis of a second scenario with a higher exponential leak of rate $5.73 \times$

$10^{-2} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ (Fig. 5b; see also Fig. S17 for the SotA and LEMAS-predicted H₂ concentrations for all exponential leak rates). Consistent with the previous analysis, LEMAS demonstrates a significantly lower LOQ, attributed to its overall higher accuracy. However, despite LEMAS exhibiting less noise than the SotA analysis the LOD of the two approaches are very similar. This can be understood by noting that, in contrast to the previous case, here the fluctuations in the prediction of the SotA analysis are much smaller than the change in H₂ due to the higher exponential leak rate.

These results are further corroborated when comparing the LOD and LOQ obtained by the SotA and LEMAS analysis methods, respectively, across a range of exponential leak rates from $6.52 \times 10^{-3} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ to $6.52 \times 10^{-2} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ (Fig. 5c-d). From Fig. 5c we identify that for exponential leak rates at $3.44 \times 10^{-2} \text{ vol.}\% \text{ H}_2 \text{ min}^{-1}$ and below LEMAS has a significantly lower LOD than SotA. At larger exponential leak rates, the LOD for the LEMAS and SotA analysis be-

comes approximately equal. This occurs because, at larger rates, the change in H_2 is sufficiently large, such that the initial change in the sensor signal is much larger than the intrinsic noise. Consequently, the ability of LEMAS to discern small signals from noise does not significantly contribute to decreasing the LOD. At lower leak rates, however, LEMAS indeed makes it possible to extract a discernible signal earlier, at lower leaked concentrations, thereby significantly increasing the time window from triggered sensor response to the leak having reached the flammability limit of 4 vol.% H_2 . Finally, in Fig. 5d, we see that LEMAS has a lower LOD for all investigated leak rates, which is a consequence of the higher accuracy obtained through LEMAS (see Note S3). In summary, these results underscore on one hand the effectiveness of LEMAS in detecting small and slow leaks earlier, as it consistently achieves a lower LOD than the SotA analysis at small rates. On the other hand, they demonstrate that LEMAS consistently outperforms the SotA analysis in terms of leak quantification, as reflected in its lower LOQ across all rates.

CONCLUSIONS

In this work, we have leveraged plasmonic H_2 sensors with deep learning to address a crucial challenge in H_2 sensing: the need for faster H_2 detection in technically relevant conditions. We have focused on the so far relatively unexplored but technically important application area of inert gas environments planned to be used to encapsulate/enclose large scale H_2 installations, such as fuel pipes on H_2 -powered airplanes or entire engine rooms on fuel-cell powered ships, to avoid the formation of flammable air- H_2 mixtures. For this application area of H_2 detection, hydride-forming plasmonic sensors are particularly well-suited, since they do not require molecular oxygen for their operation — in contrast to, e.g., catalytic or thermal H_2 sensors, which are most commonly used today.

From the deep learning perspective, we have developed LEMAS, short for Long Short-term Transformer Ensemble Model for Accelerated Sensing, which accelerates sensor response by learning the relationship between the time dependence of the full spectral response of the plasmonic sensor and the H_2 concentration to predict the final sensor response before it is reached physically, while simultaneously assessing uncertainty in the model predictions through model ensembles. To obtain accurate models for the ensemble and mitigate artifacts from measurement noise, drift, and variations between different measurements, we found it crucial to use a sufficiently long time series and combine the result of several different pre-processing approaches, such as wavelength dependent min-max normalization, standard normal variate standardization, global min-max normalization, and

level scaling.

In summary, our results demonstrate the ability of deep learning concepts to significantly accelerate sensor response times and to enhance quantification of H_2 leaks by enabling their detection earlier and therefore at lower leaked concentrations. This, in turn, is important from a practical application perspective, since it provides a longer time window for the implementation of appropriate measures for handling the leak. Specifically, we have demonstrated (i) that LEMAS is able to accelerate the response time of a plasmonic H_2 sensor by a factor of 20 to 40 in an inert gas environment for H_2 concentrations of 0.1 vol.% H_2 and below, and (ii) that it effectively eliminates the intrinsic H_2 concentration dependence of metal hydride-based sensors. This is an important result because the characteristic increase of sensor response time for decreasing H_2 concentrations constitutes one of the long-standing unresolved limitations in H_2 sensor technology.

Taken together, our findings underscore the significant potential of deep learning for overcoming current limitations in H_2 sensor performance, such as slow response in technologically relevant sensing environments like the inert gas environment investigated here. Furthermore, while our investigation primarily focused on plasmonic H_2 sensors, it is important to emphasize that the insights gained from this study hold broader implications for sensor technology in general as they illustrate a relatively unexplored generic concept for accelerating sensors by means of deep learning.

Methods

Hydrogen sensing experiments. The measurements were conducted in a custom-built reactor chamber that is comprised of a customized DN 16 CF spacer flange (Pfeifer Vacuum), equipped with a gas in- and outlet, and two fused-silica viewports (1.33" CF Flange, Accu-Glass). The effective chamber volume is ca. 1.5 mL. The gas flow rates were controlled by mass flow controllers (El-Flow Select series, Bronkhorst High-Tech) (Fig. S1). The sample inside the chamber was illuminated using an unpolarized halogen white light source (AvaLight-HAL, Avantes) and an optical fiber equipped with a collimating lens. The transmitted light was collected and analyzed using a fiber-coupled fixed-grating spectrometer (SensLine AvaSpec-HS1024TEC, Avantes). The temperature was controlled with a heating coil wrapped around the chamber and a temperature controller (Eurotherm 3216) in a feedback loop manner, where the sample surface temperature inside the chamber was continuously used as input.

All measurements were performed at 30 °C in Argon background, with a constant gas flow of 300 mL/min. The hydrogen concentration in all of the following

measurements was in the range of 0.06 vol.% H_2 – 1.97 vol.% H_2 (detailed description of the different pulse schemes as found in Note S1). The sampling frequency of the spectrometer was set to 3 Hz.

SotA analysis. In this work, we used the centroid position as a spectral descriptor. The centroid position is defined as $\lambda_c = \sum_{\lambda} \lambda I(\lambda) / \sum_{\lambda} I(\lambda)$, where λ is the wavelength in nm and $I(\lambda)$ is the intensity at wavelength λ . To enable comparison between the SotA analysis and LEMAS on the test measurements we fit a calibration function, using the measured H_2 concentration in the training measurements, as

$$H_2(\Delta\lambda_c) = a\Delta\lambda_c^b, \quad (1)$$

where $\Delta\lambda_c$ is the change centroid position, taken from the smallest centroid position in each measurement. The values of the parameters a and b are determined by minimizing the mean absolute percentage error between the measured H_2 and $H_2(\lambda_c)$. We fit two different calibration functions, one for the data consisting of of step-wise increase/decrease of H_2 and one for the data consisting of of linear/exponential increase/decrease of H_2 (see Note S3 for details).

Data pre-processing. Before the data was fed into the deep learning model it was pre-processed using four different methods, and the concatenation of these methods was fed as input to the deep learning model. Each measurement was pre-processed individually by using the initial sequence of 5 pulses of 1.97 vol.% H_2 , for each measurement, to estimate the minimum/maximum/and mean intensity. The scaling methods were (i) wavelength dependent min-max normalization: for each wavelength subtracting the estimated minimum intensity at the corresponding wavelength and dividing by the difference between the estimated maximum intensity of all wavelengths and the estimated minimum intensity measured at the specific wavelength, (ii) standard normal variate standardization: scaling each spectrum using its mean and standard deviation, (iii) global min-max normalization: subtracting the estimated minimum intensity and dividing by the difference between the estimated maximum and minimum intensity in and (iv) level scaling: subtracting and dividing each spectrum in the measurement by the estimated mean intensity.

Deep learning model. The deep learning architecture that was used in this work was a long short-term transformer (LSTR) [20] which operates as illustrated in Fig. S14. Each temporal feature consisting of the concatenation is linearly mapped to a vector of size $d_{model} = 256$. Subsequently, positional encoding is added and the data is split into a short-term memory and long-term memory. Here, we down-sample the long-term memory using a stride of 4. Firstly, the long-term memory undergoes a two stage memory compression through the LSTR encoder, using a set of learnable token embeddings of dimensions $d_{model} \times n_1$ and $d_{model} \times n_0$. Here,

we used $n_0 = 8$ and $n_1 = 4$ and the encoder consisted of 4 transformer decoder units. Secondly, the LSTR decoder extracts relevant temporal features in the short-term memory, while also querying the encoded long-term memory to retrieve useful information from the history of the sensor. Here, the decoder consisted of 8 transformer decoder units. The extracted temporal features are then passed through $n_{mlp} = 8$ MLP layers of dimension $d_{mlp} = 512$ to obtain H_2 concentration predictions. Here all the transformer decoder units performed multi head attention as in with $h = 8$ heads and the $d_k = d_q = d_v = d_{model}/h = 32$, and the dimension of the MLP inside the transformer decoder units was $d_{ff} = 512$. Furthermore, in the LSTR encoder, masked multi-head attention was performed such that during training the H_2 concentration corresponding to each time step in the short term memory could be used for supervision during training.

Ensembles The constructed ensembles comprised ten models, each varying in the lengths of short-term and long-term memory. This variation was designed to induce diversity in the predictive capabilities of the models within the ensemble. Specifically, two models were designated for each combination of long and short-term memory lengths, while ensuring a consistent total input sequence length across all models.

For the ensemble tailored for leak detection, the preceding 4000 time steps constituted the input. By contrast, the ensemble model, which was aimed at minimizing response time, utilized the previous 600 time steps as input. For both ensembles, the selected lengths for short-term memory were 20, 40, 60, 80, and 100 time steps, respectively. This approach was adopted to enhance the generalization capabilities the ensemble. It is important to note that apart from the variation in memory lengths, all models shared identical hyperparameters (see Table S1).

To make the prediction of the ensemble more robust to potential outliers, we only included predictions that fall between the first and third percentile to compute the ensemble prediction and uncertainty.

Deep learning training The models were implemented using TENSORFLOW [31] and were trained for 100 epochs on Nvidia A100 graphical processing units using the AdamW [32] optimizer with weight decay 5×10^{-5} , a batch size of 128, and mean-absolute-error loss. The learning rate was increased linearly from zero to 5×10^{-5} during the first 15 epochs then decaying to zero following a cosine curve. To analyze the impact of different pre-processing methods we used the first half of the data from measurement Fig. S4a and Fig. S4c as training data and the other half as validation data (see Note S2). For the first case of a large simulated leak characterized by a rapid step-wise increase of H_2 concentration in the sensor surroundings, we trained LEMAS using the data from measurement Fig. S4a and Fig. S4c as training data and

the data from measurement Fig. S4b as test data. For the second case of a simulated small slow leak, we trained LEMAS the data from measurement Fig. S5a as training data and data from measurement Fig. S5b as test data. Furthermore, during the training phase, each model in the ensemble was exposed to a distinct subset of the training data, comprising a random 90% of the total dataset.

Sample fabrication. Quasi-random PdAu alloy (nominal 70:30 at. %) nanodisk arrays with 210 nm average disk diameter and 25 nm height, were fabricated using Hole-Mask Colloidal Lithography (HCL) [33]. The metals were deposited layer-by-layer via electron beam evaporation, onto 1 cm \times 1 cm fused silica substrates (Siebert Wafer GmbH). Subsequent annealing was performed at 500 °C for 18 h under a flow of 4 vol.% H₂ in Ar to induce alloy formation. A more detailed description of the nanofabrication procedure can be found in our earlier work [34].

Data availability

AUTHOR CONTRIBUTIONS

COMPETING INTERESTS

C. L. is co-founder and scientific advisor at Inspilorion AB who markets plasmonic hydrogen sensors.

ACKNOWLEDGMENTS

This work was funded by the Vinnova project 2021-02760, the Swedish Research Council (grant numbers 2018-06482, 2020-04935, 2021-05072), the Swedish Energy Agency (grant No. 45410-1), the Area of Advanced Nano at Chalmers, and the Competence Centre TechForH2. The Competence Centre TechForH2 is hosted by Chalmers University of Technology and is financially supported by the Swedish Energy Agency (P2021-90268) and the member companies Volvo, Scania, Siemens Energy, GKN Aerospace, PowerCell, Oxeon, RISE, Stena Rederier AB, Johnsson Matthey and Inspilorion. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at C3SE partially funded by the Swedish Research Council through grant agreement no. 2022-06725. This work was performed in part at Myfab Chalmers and the Chalmers Materials Analysis Laboratory (CMAL).

[†] clangham@chalmers.se

- [1] M. Zhang, Z. Yuan, J. Song, and C. Zheng, *Sensors and Actuators B: Chemical* **148**, 87 (2010).
- [2] F. Ribet, G. Stemme, and N. Roxhed, *Biomedical microdevices* **20**, 1 (2018).
- [3] L. Rolinger, M. Rüdts, and J. Hubbuch, *Analytical and Bioanalytical Chemistry* **412**, 2047 (2020).
- [4] A. Poghosian, H. Geissler, and M. J. Schöning, *Biosensors and Bioelectronics* **140**, 111272 (2019).
- [5] I. Darmadi, F. A. A. Nugroho, and C. Langhammer, *ACS Sensors* **5**, 3306 (2020).
- [6] U.S. department of energy; 2015 safety, codes and standards section (2015).
- [7] F. A. Nugroho, I. Darmadi, L. Cusinato, A. Susarrey-Arce, H. Schreuders, L. J. Bannenberg, A. B. da Silva Fanta, S. Kadkhodazadeh, J. B. Wagner, T. J. Antosiewicz, *et al.*, *Nature Materials* **18**, 489 (2019).
- [8] H. M. Luong, M. T. Pham, T. Guin, R. P. Madhogaria, M.-H. Phan, G. K. Larsen, and T. D. Nguyen, *Nature communications* **12**, 2414 (2021).
- [9] L. Bannenberg, H. Schreuders, and B. Dam, *Advanced Functional Materials* **31**, 2010483 (2021).
- [10] I. Darmadi, F. A. A. Nugroho, S. Kadkhodazadeh, J. B. Wagner, and C. Langhammer, *ACS Sensors* **4**, 1424 (2019).
- [11] E. K. Alenezy, Y. M. Sabri, A. E. Kandjani, D. Korcoban, S. S. A. Abdul Haroon Rashid, S. J. Ippolito, and S. K. Bhargava, *ACS Sensors* **5**, 3902 (2020).
- [12] S. Kabcum, D. Channei, A. Tuantranont, A. Wisitsoraat, C. Liewhiran, and S. Phanichphant, *Sensors and Actuators B: Chemical* **226**, 76 (2016).
- [13] K. Hassan, A. Iftekhar Uddin, and G.-S. Chung, *Sensors and Actuators B: Chemical* **234**, 435 (2016).
- [14] M. Losurdo, Y. Gutiérrez, A. Suvorova, M. M. Giangregorio, S. Rubanov, A. S. Brown, and F. Moreno, *Advanced Materials* **33**, 2100500 (2021).
- [15] S. Huang, A. Croy, B. Ibarlucea, and G. Cuniberti, *Machine learning-driven gas identification in gas sensors*, in *Machine Learning for Advanced Functional Materials*, edited by N. Joshi, V. Kushvaha, and P. Madhusri (Springer Nature Singapore, Singapore, 2023) pp. 21–41.
- [16] M. Kang, I. Cho, J. Park, J. Jeong, K. Lee, B. Lee, D. Del Orbe Henriquez, K. Yoon, and I. Park, *ACS Sensors* **7**, 430 (2022).
- [17] B. Ai, Y. Sun, and Y. Zhao, *Small* **18**, 2107882 (2022).
- [18] C. Wadell, S. Syrenova, and C. Langhammer, *ACS Nano* **8**, 11925 (2014).
- [19] A. B. Dahlin, J. O. Tegenfeldt, and F. Höök, *Analytical Chemistry* **78**, 4416 (2006).
- [20] M. Xu, Y. Xiong, H. Chen, X. Li, W. Xia, Z. Tu, and S. Soatto, in *NeurIPS 2021* (2021).
- [21] C. Wadell, F. A. A. Nugroho, E. Lidström, B. Iandolo, J. B. Wagner, and C. Langhammer, *Nano Letters* **15**, 3563 (2015).
- [22] L. J. Bannenberg, F. A. A. Nugroho, H. Schreuders, B. Norder, T. T. Trinh, N.-J. Steinke, A. A. van Well, C. Langhammer, and B. Dam, *ACS Applied Materials & Interfaces* **11**, 15489 (2019).
- [23] F. A. A. Nugroho, I. Darmadi, V. P. Zhdanov, and C. Langhammer, *ACS Nano* **12**, 9903 (2018).
- [24] P. Ekborg-Tanner, J. M. Rahm, V. Rosendal, M. Bancerek, T. P. Rossi, T. J. Antosiewicz, and P. Erhart, *ACS Applied Nano Materials* **5**, 10225 (2022).

* erhart@chalmers.se

- [25] S.-M. Lee, H. Noh, T. B. Flanagan, and S. Luo, *Journal of Physics: Condensed Matter* **19**, 326222 (2007).
- [26] S. Luo, D. Wang, and T. B. Flanagan, *The Journal of Physical Chemistry B* **114**, 6117 (2010).
- [27] M. Mamatkulov and V. P. Zhdanov, *Phys. Rev. E* **101**, 042130 (2020).
- [28] J. M. Rahm, J. Löfgren, E. Fransson, and P. Erhart, *Acta Materialia* **211**, 116893 (2021).
- [29] F. A. A. Nugroho, B. Iandolo, J. B. Wagner, and C. Langhammer, *ACS Nano* **10**, 2871 (2016).
- [30] P. Mishra, D. Passos, F. Marini, J. Xu, J. M. Amigo, A. A. Gowen, J. J. Jansen, A. Biancolillo, J. M. Roger, D. N. Rutledge, and A. Nordon, *TrAC Trends in Analytical Chemistry* **157**, 116804 (2022).
- [31] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, *TensorFlow: Large-scale machine learning on heterogeneous systems* (2015), software available from tensorflow.org.
- [32] I. Loshchilov and F. Hutter, *Decoupled weight decay regularization* (2019), arXiv:1711.05101 [cs.LG].
- [33] H. Fredriksson, Y. Alaverdyan, A. Dmitriev, C. Langhammer, D. S. Sutherland, M. Zäch, and B. Kasemo, *Advanced Materials* **19**, 4297 (2007).
- [34] F. A. A. Nugroho, B. Iandolo, J. B. Wagner, and C. Langhammer, *ACS Nano* **10**, 2871 (2016).

Supporting Information

Accelerating Plasmonic Hydrogen Sensors in Inert Gas Environments by Transformer-Based Deep Learning

Viktor Martvall¹, Henrik Klein Moberg¹, Athanasios Theodoridis¹, David Tomeček¹,
Pernilla Ekborg-Tanner¹, Sara Nilsson¹, Giovanni Volpe², Paul Erhart^{1,*}, and
Christoph Langhammer^{1,†}

¹ *Department of Physics, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden*

² *Department of Physics, University of Gothenburg, 412 96 Gothenburg, Sweden*

^{*} *erhart@chalmers.se*

[†] *clangham@chalmers.se*

Contents

Supplementary Notes	2
S1. Hydrogenation protocols	2
S2. Data pre-processing	2
S3. SotA analysis	2
Supplementary Figures	3
S1. Experimental setup	3
S2. Gas background quadrupole mass spectrometer measurements	4
S3. Gas background quadrupole mass spectrometer measurements	5
S4. Step wise hydrogenation measurements	6
S5. Linear and exponential hydrogenation measurements	6
S6. Calibration function for step-wise increases/decreases	7
S7. Calibration function applied to the test data	7
S8. Calibration function for linear/exponential increases/decreases	8
S9. Calibration function applied to the test data	8
S10. Time series length for models optimized for accelerating sensor response	8
S11. Time series length for models optimized for leak detection	9
S12. Different pre-processing methods	9
S13. Impact of using different pre-processing methods	10
S14. Schematic illustration of the deep learning architecture used in this work	10
S15. Comparison of LEMAS and SotA analysis for step-wise increases in test data	11
S16. Initial prediction of mean and standard deviation by LEMAS for step-wise increases in test data	12
S17. Comparison of LEMAS and SotA analysis for exponential leak rates in test data	13
Supplementary Tables	14
S1. Hyperparameters used for all models	14
Supplementary References	14

Supplementary Notes

Supplementary Note S1: Hydrogenation protocols

The hydrogenation protocols were as follows: 5 pulses of 1.97 vol.% H_2 concentration were introduced at the beginning of each measurement, in order to activate the nanodisks. For the step wise hydrogen profile (Fig. S4), 3 different measurements were performed which consisted of a logarithmic distribution of H_2 pulses (14 steps increase/decrease, 0.06 vol.% H_2 minimum and 1.97 vol.% H_2 maximum concentration), where the pulse durations were set to 5 min, 10 min and 15 min. The time between pulses was 30 min for all measurements.

For the linear hydrogen profile (Fig. S5a), the H_2 concentration was slowly increased linearly to a maximal concentration. For each of the ramp up/down cycles, the maximum concentrations were: 0.09, 0.13, 0.19, 0.28, 0.41, 0.61, 0.90, 1.33, 1.97 vol.% H_2 . The ramp-up/down duration was 1 hour. The maximal concentrations were increasing after each ramping up/down and therefore the hydrogen increase rate was changing as well. For the exponential hydrogen profile (Fig. S5b) the H_2 concentration was increased exponentially to a maximal concentration instead of linearly. The duration of the ramp-up/down was 1 hour and the same maximum concentrations were used.

Supplementary Note S2: Data pre-processing

Due to variations in the spectra from different measurements (see Fig. S12), it was necessary to pre-process the data before training a model (see Data pre-processing of the main paper). Initially, we only considered one pre-processing method (wavelength dependent min-max normalization), however, as seen by the blue curve in Fig. S13 this caused errors in the model predictions, which are likely due to the pre-processed spectra for similar H_2 concentration being relatively similar. Sequentially, we therefore added standard normal variate standardization (orange curve in Fig. S13), global min-max normalization (green curve in Fig. S13) and level scaling (red curve in Fig. S13). By using all four pre-processing methods the model accuracy increases substantially.

Supplementary Note S3: SotA analysis

To obtain the calibration curve for the step-wise increases we used the measurement illustrated in Fig. S4a. For each pulse, we extracted the steady-state centroid shift as well as the H_2 concentration (Fig. S6). The motivation behind only using the steady-state centroid shift is that a unique mapping between the centroid shift and the H_2 concentration does not exist during the transient state. Then we minimized the mean absolute relative error and obtained $a = 0.334$ and $b = 2.91$ (see Eq. (1)). In Fig. S6 we illustrate the use of the obtained calibration function for predicting the H_2 concentration on the measurement used for testing LSTR Ensemble Model for Accelerated Sensing (LEMAS) (Fig. S4a). The mean absolute relative error of the state-of-the-art (SotA) analysis for this measurement is 14.73×10^{-3} vol.% H_2 , which is higher than the mean absolute relative error of LEMAS being 1.048×10^{-3} vol.% H_2 , showing that LEMAS has a higher accuracy.

To obtain the calibration curve for the linear/exponential increases we used the measurement consisting of linear increases/decreases (Fig. S5a). Here, we used the entire measurement to obtain a calibration curve, as the H_2 concentration changes slower than the response time of the system (Fig. S8). Then we minimized the mean absolute relative error and obtained $a = 0.40$ and $b = 2.51$. In Fig. S9 we illustrate the use of the obtained calibration function for predicting the H_2 concentration on the measurement used for testing LEMAS (Fig. S5b). The mean absolute relative error of the SotA analysis for this measurement is 13.43×10^{-3} vol.% H_2 , which is higher than the mean absolute relative error of LEMAS being 7.4×10^{-3} vol.% H_2 , confirming again that LEMAS has a higher accuracy.

Supplementary Figures




Figure S1: **Experimental setup.** Schematic illustration of the (a) reaction chamber and (b) gas mixing system. Adapted with permission from Ref. 1.




Figure S2: **Gas background quadrupole mass spectrometer measurements.** The gas flows of 100 % Ar (blue, left y-axis) followed by 100 % H₂ (red, left y-axis) and temperature (black, right y-axis) during the quadrupole mass spectrometer (QMS) measurements of the concentrations of possible trace gases in the mixture. The temperature was first increased to 400 °C in 100 % Ar flow, to desorb possible absorbed gases in the flow reactor. This was followed by one 15 minutes 100 % H₂ pulse, and again one pulse of 100 % Ar. The QMS measurement was started at time 0, marked by the dashed line.




Figure S3: **Gas background quadrupole mass spectrometer measurements.** The concentrations of Ar and the trace gases during the Ar and H₂ pulses in Fig. S2. The concentrations of the trace gases at the end of the Ar pulse at 8000 s were 99.7 % Ar, 0.26 % H₂O, 0.027 % CO or N₂, 31 ppm O₂, 18 ppm NO, and 6.6 ppm NO₂.

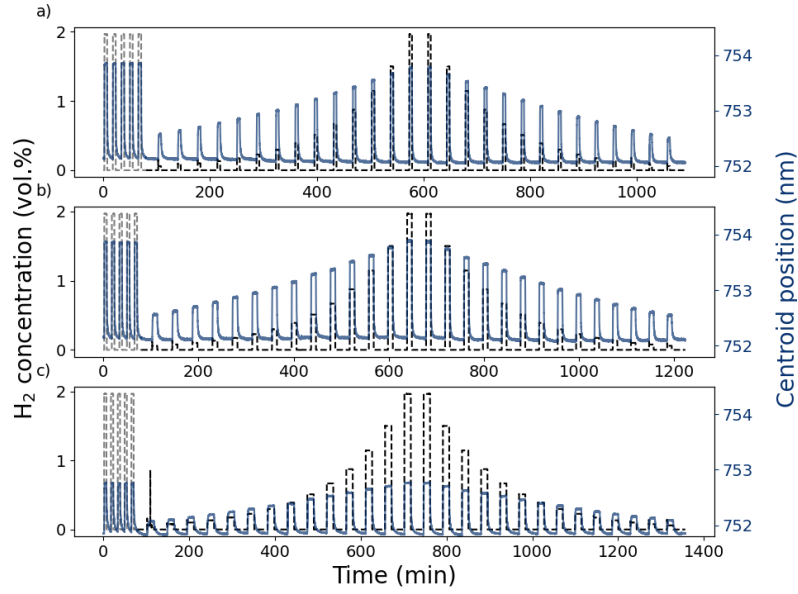


Figure S4: **Step wise hydrogenation measurements.** Hydrogenation protocol for the step wise increase/decrease, alongside the measured centroid shift. (a) 5 min, (b) 10 min, (c) 15 min of pulse duration followed by 30 minutes between each pulse for all measurements.




Figure S5: **Linear and exponential hydrogenation measurements.** Hydrogenation protocol for the (a) linear and (b) exponential increase/decrease alongside the measured centroid shift. The ramp-up/down duration for each cycle was 1 h.




Figure S6: **Calibration function for step-wise increases/decreases.** Experimental data used from measurement Fig. S4a to fit the calibration function for the step-wise increases and the obtained calibration function.




Figure S7: **Calibration function applied to the test data.** Calibration function obtained from Fig. S6 used on the test data, measurement Fig. S4b.




Figure S8: **Calibration function for linear/exponential increases/decreases** . Experimental data used from measurement Fig. S5a to fit the calibration function for the linear/exponential increases and the obtained calibration function.




Figure S9: **Calibration function applied to the test data.** Calibration function obtained from Fig. S8 used on the test data, measurement Fig. S5b.

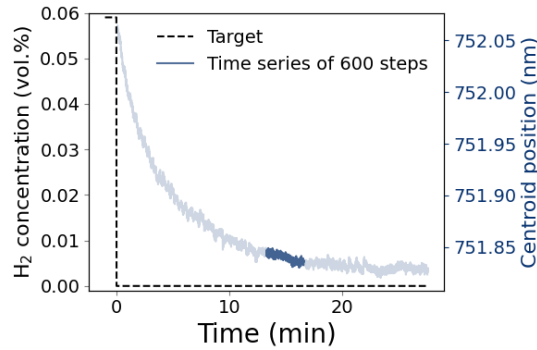


Figure S10: **Time series length for models optimized for accelerating sensor response.** Change in centroid position during the slowest occurring process in the training data, namely the desorption from 0.06 vol.% H₂ in measurement Fig. S4c. A time series comprising 600 time steps is chosen here to allow the model to estimate the rate of change in the spectrum, which is a prerequisite for accurately describing this process.

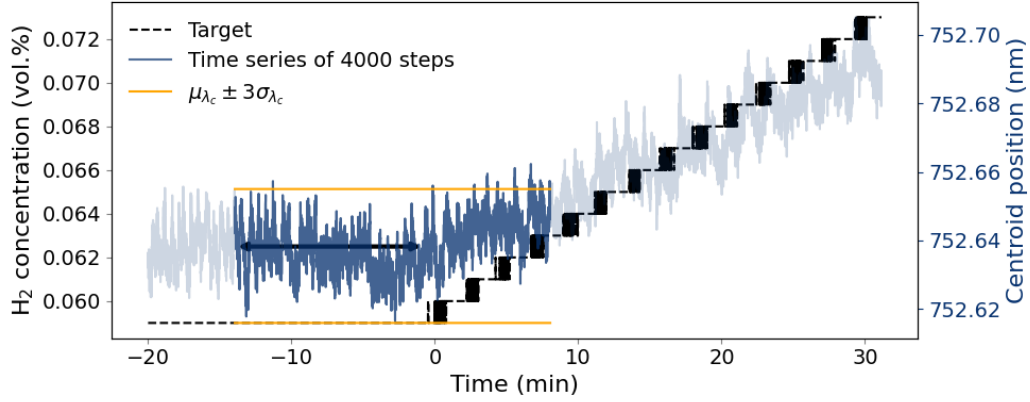


Figure S11: **Time series length for models optimized for leak detection.** The change in centroid position in the beginning of the leak with the smallest leak rate in Fig. S5a, compared to the mean and standard deviation of the centroid position before the leak has started (indicated by the double-headed arrow). Through using a time series comprising 4000 time steps the model should, in the worst case, be able to differentiate the leak occurring with the smallest slope from the noise in the sensor output data approximately 10 min after the leak has started.

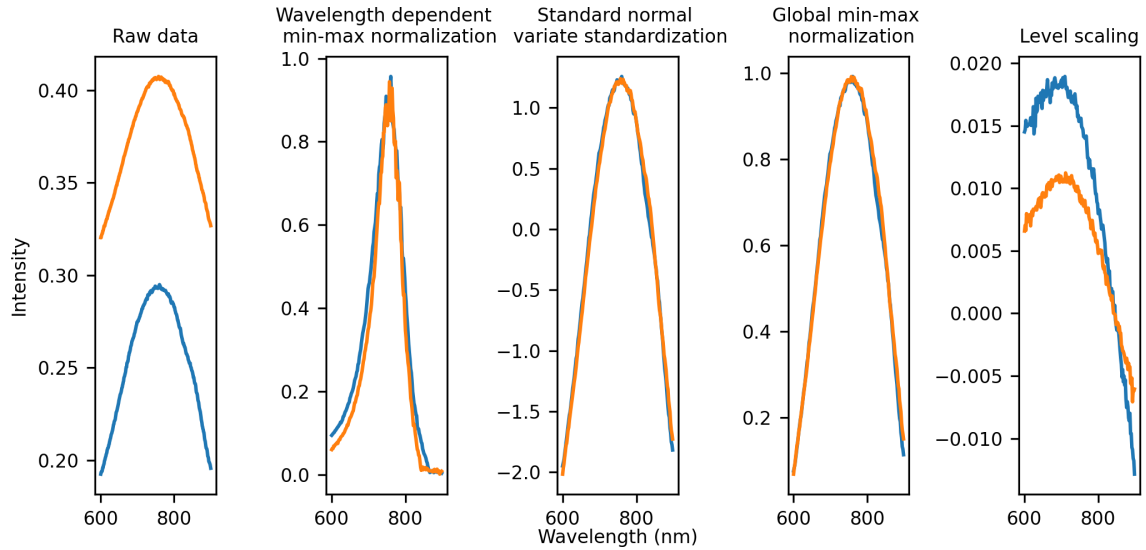


Figure S12: **Different pre-processing methods.** Illustration of the effect of different pre-processing methods on model prediction as well as raw data for spectra from measurement Fig. S4a (blue) and Fig. S4c (orange) at 0 vol.% H₂.




Figure S13: **Impact of using different pre-processing methods.** Predicted H_2 concentration on the two last step-wise increases in measurement Fig. S4c for models trained on the first half of measurement Fig. S4a and Fig. S4c, using different combinations of pre-processing methods.

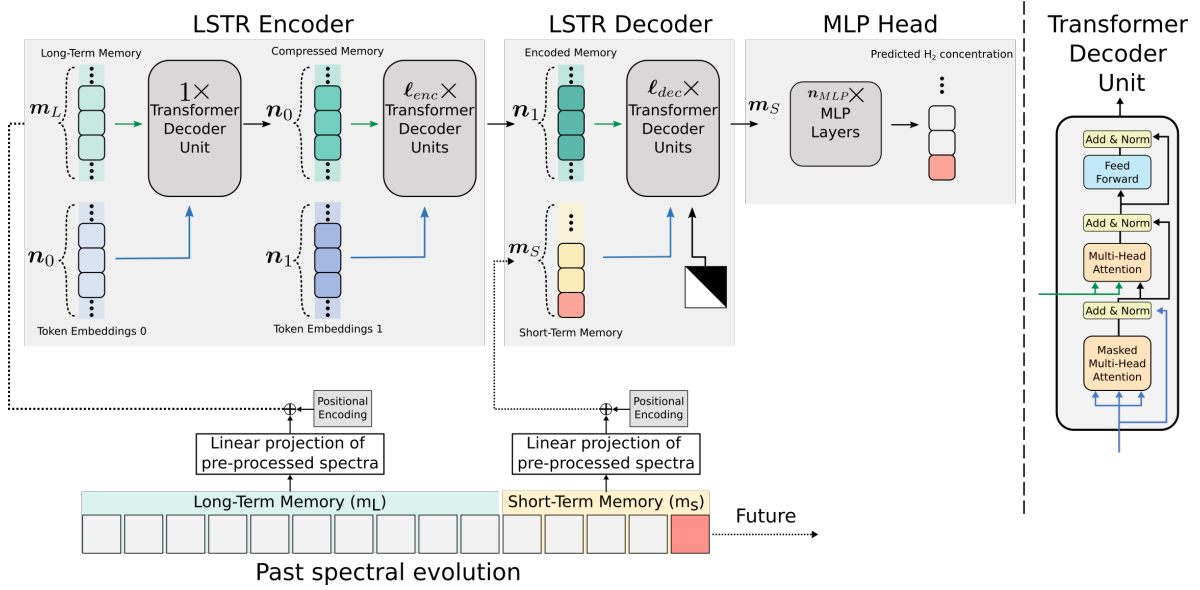


Figure S14: **Schematic illustration of the deep learning architecture used in this work.** Adapted from Ref. 2.

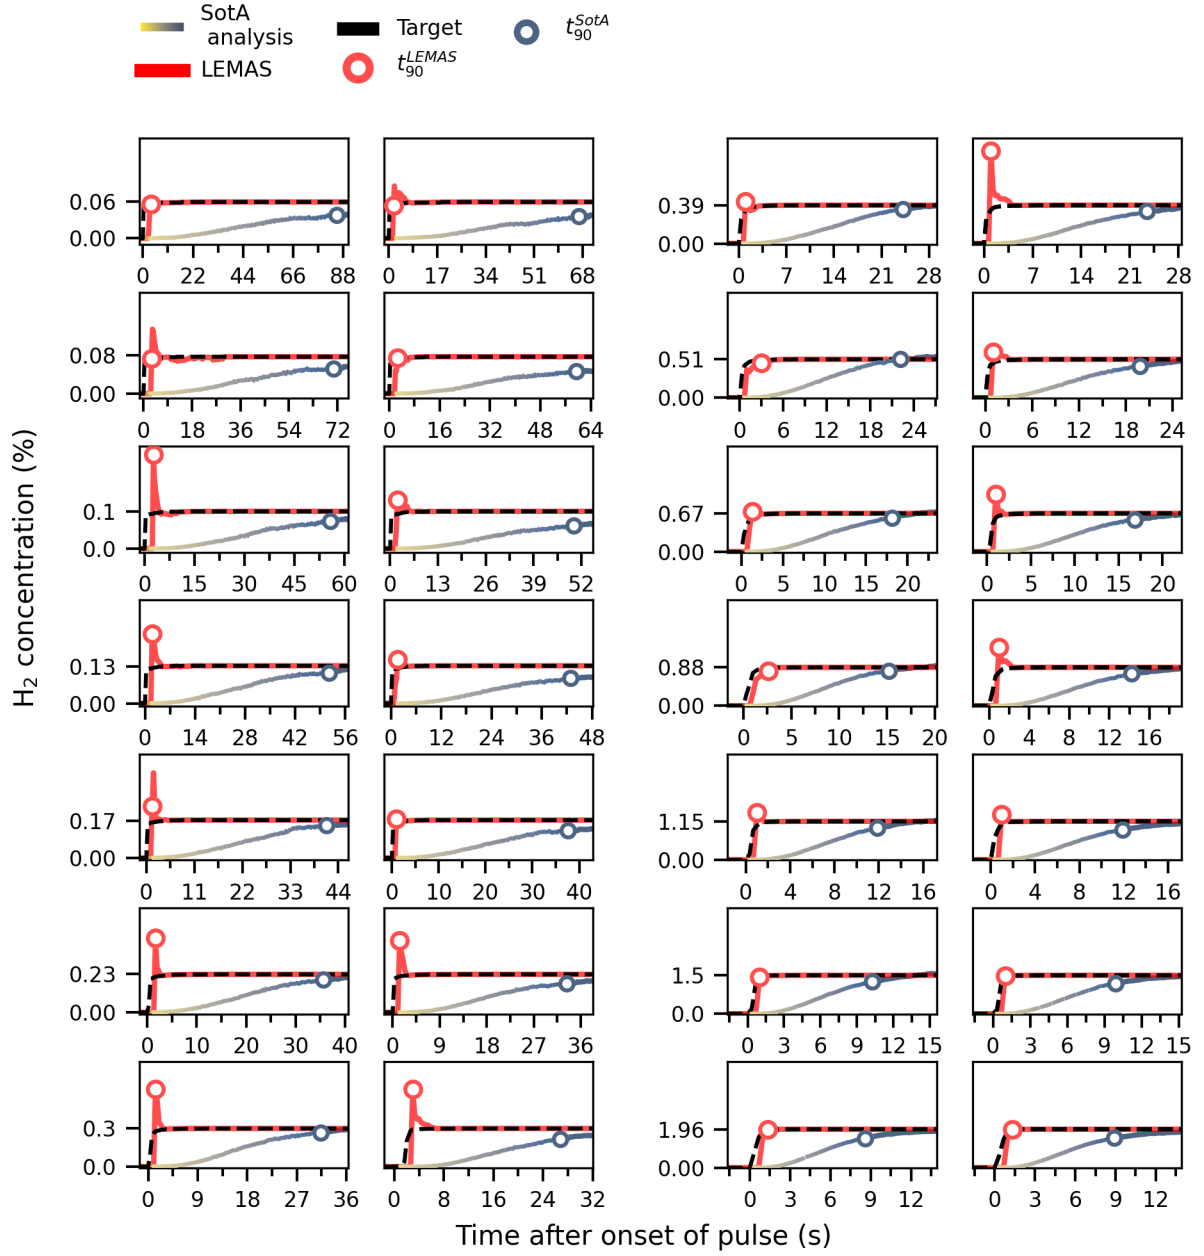


Figure S15: Comparison of LEMAS and SotA analysis for step-wise increases in test data.

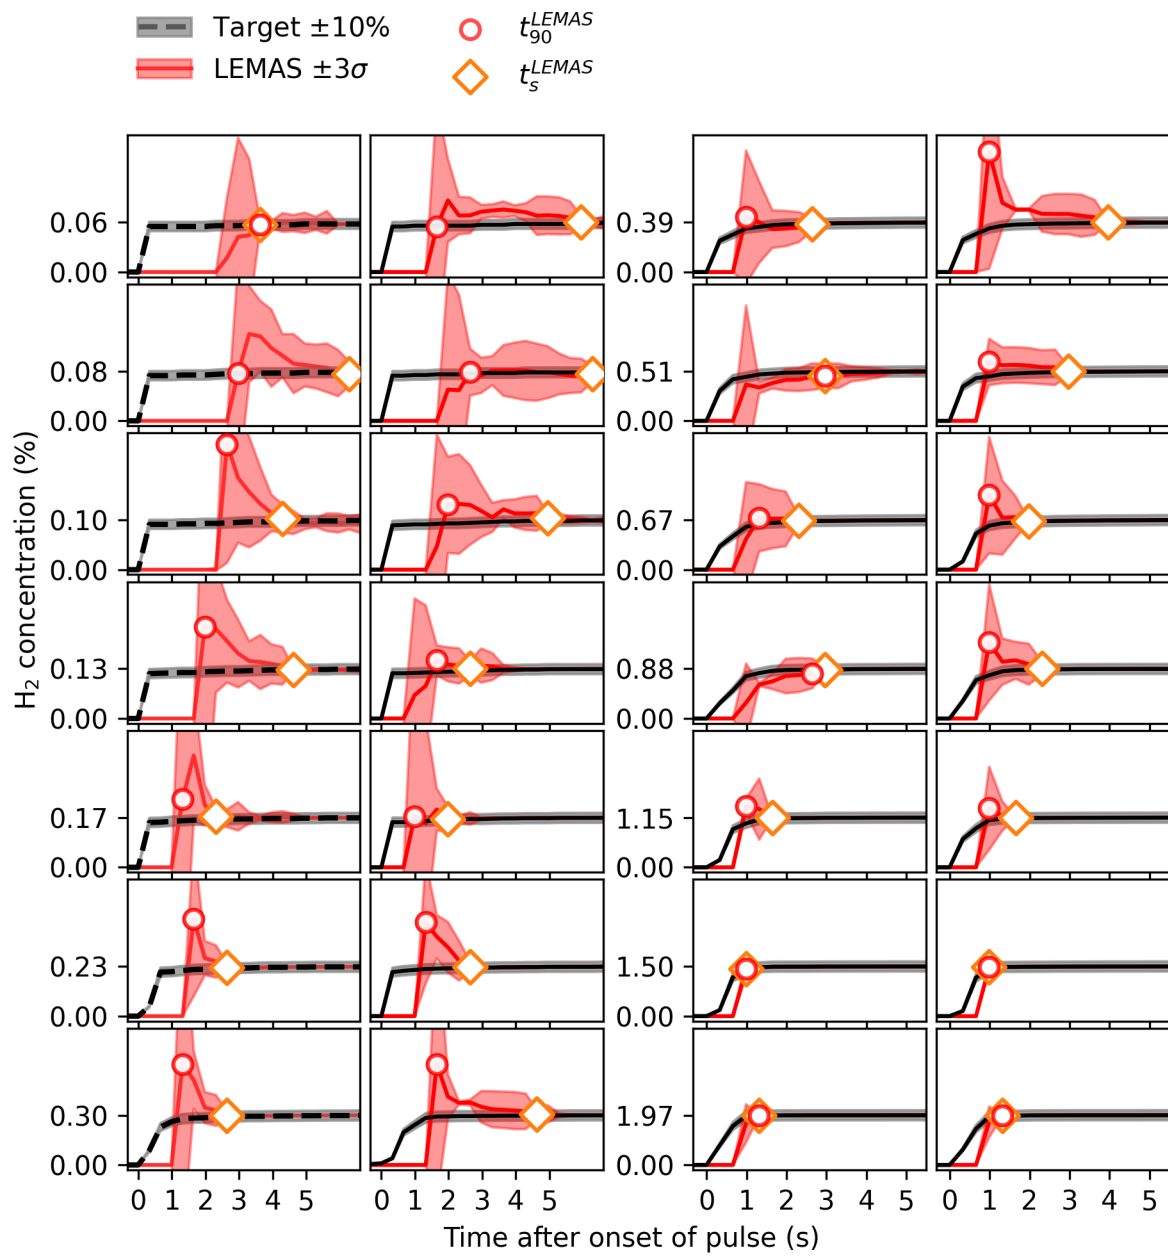


Figure S16: Initial prediction of mean and standard deviation by LEMAS for step-wise increases in test data.

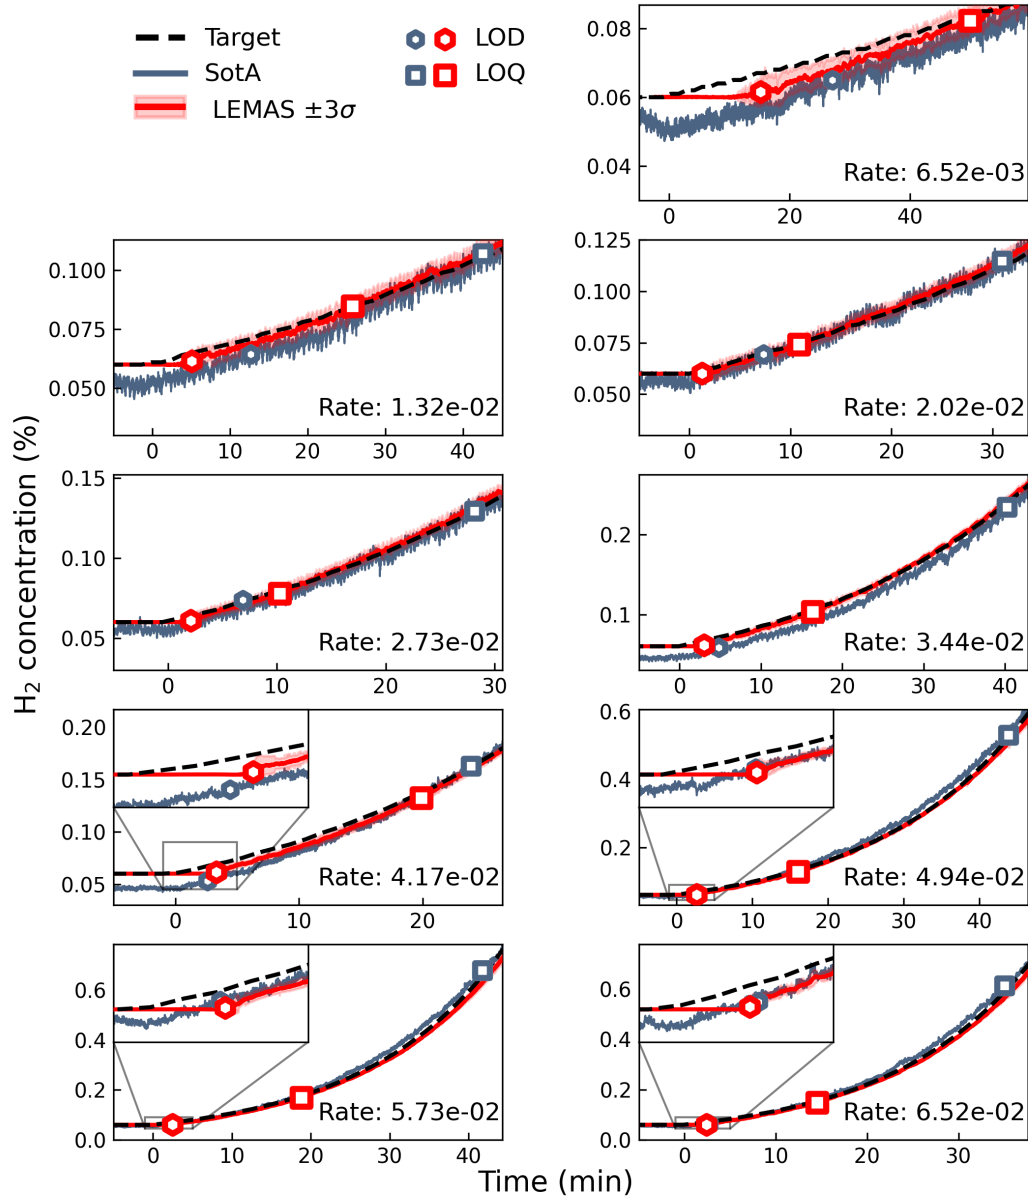


Figure S17: Comparison of LEMAS and SotA analysis for exponential leak rates in test data.

Supplementary Tables

Table S1: **Hyperparameters used for all models.**

Parameter	Value
n_{enc}	8
n_{dec}	4
n_0	8
n_1	4
d_{model}	256
d_{ff}	512
num heads	8
head size	32
Temporal stride on long term memory	4
n_{mlp}	8
d_{mlp}	512
Dropout rate	0.1

Supplementary References

- [1] Iwan Darmadi. *Polymer-Nanoparticle Hybrid Materials for Plasmonic Hydrogen Detection*. PhD thesis, Chalmers University of Technology, Gothenburg, Sweden, January 2021. Available at <https://research.chalmers.se/en/publication/521833>.
- [2] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. In *NeurIPS 2021*, 2021. URL <https://www.amazon.science/publications/long-short-term-transformer-for-online-action-detection>.