

Machine Learning for Chemistry Reduction in N₂–H₂ Low-Temperature Plasmas

Diogo R. Ferreira¹(✉), Alexandre Lança¹, and Luís Lemos Alves²

¹ Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal

² Instituto de Plasmas e Fusão Nuclear (IPFN), Instituto Superior Técnico,
Universidade de Lisboa, 1049-001 Lisboa, Portugal
`diogo.ferreira@tecnico.ulisboa.pt`

Abstract. Low-temperature plasmas are partially ionized gases, where ions and neutrals coexist in a highly reactive environment. This creates a rich chemistry, which is often difficult to understand in its full complexity. In this work, we develop a machine learning model to identify the most important reactions in a given chemical scheme. The training data are an initial distribution of species and a final distribution of species, which can be obtained from either experiments or simulations. The model is trained to provide a set of reaction weights, which become the basis for reducing the chemical scheme. The approach is applied to N₂–H₂ plasmas, created by an electric discharge at low pressure, where the main goal is to produce NH₃. The interplay of multiple species, as well as of volume and surface reactions, make this chemistry especially challenging to understand. Reducing the chemical scheme via the proposed model helps identify the main chemical pathways.

Keywords: Plasma Physics · Plasma Chemistry · Machine Learning.

1 Introduction

The production of ammonia (NH₃) is critically important to a wide range of economic activities in the agricultural and industrial sectors. For example, ammonia is a key ingredient in the manufacturing of fertilizers, pharmaceuticals, plastics, textiles, and other chemicals. For more than a hundred years, the main industrial procedure for the production of ammonia has been the Haber-Bosch process [6]. However, this procedure requires temperatures in the range of 400 to 500 °C and pressures in the range of 150 to 300 atmospheres, which translate into considerable requirements in terms of energy consumption. These conditions are necessary in order to overcome the triple bond of nitrogen molecules (N₂), enabling them to react with hydrogen (H₂), via $\text{N}_2 + 3\text{H}_2 \rightarrow 2\text{NH}_3$.

Low-temperature plasmas have emerged as a more energy efficient and environmentally friendly process for ammonia synthesis [2]. In this case, the process begins by exciting and ionizing a gas mixture of N₂ and H₂. Several ionization methods are available, but one of the simplest and most direct is to apply an electrical discharge [3]. This brings the gas mixture into a partially ionized state,

containing electrons, ions, radicals, and neutral particles. In this plasma state, the triple bonds of N_2 and the covalent bonds of H_2 can be broken more effectively than in thermal processes, leading to the formation of highly reactive nitrogen and hydrogen species. A vast and complicated series of reactions then unfolds, eventually producing NH_3 . One of the main goals of this work is to contribute to the understanding of this chain of reactions.

Studying the chemistry of ammonia production is important for additional reasons. Ammonia is known to be present in the atmosphere of the gas giants in our solar system, especially in Jupiter, Saturn, and Saturn’s moon Titan [5]. Ammonia has also been detected in interstellar space and in exoplanets [10]. The complex chemistry that takes place in a planetary atmosphere has prompted scientists to carry out laboratory experiments, including plasma discharges, to study the formation of organic molecules and even prebiotic molecules required for the origin of life [4]. Understanding the production of ammonia in those environments, through lab experiments and computer-based simulations, is an integral part of those endeavors. Using machine learning, our goal is to identify the most important reactions in such complex chemical processes.

The paper is organized as follows. Section 2 provides a brief introduction to the chemistry of N_2 – H_2 plasmas. Section 3 explains how a chemical scheme containing many reactions can be represented as a Petri net. This representation can be translated into matrix form, and we use this matrix form to develop a machine learning model in Section 4. In Section 5 we describe where the training data comes from, and how the model is trained. Finally, in Section 6 we analyze the results and identify the main pathways for NH_3 production.

2 Chemistry of N_2 – H_2 Plasmas

The basis of the chemistry of N_2 – H_2 low-temperature plasmas is fairly well documented in the literature [11]. It consists of over 200 reactions, including vibrational processes, electron impact processes, as well as several reactions involving heavy (neutral and charged) species. Since it is impractical (and unnecessary) to reproduce the chemical scheme here in full, we provide only a few examples of the type of reactions that it includes.

For vibrational processes, there are reactions such as:

- Vibrational excitation or de-excitation, e.g. $e + N_2(X,v) \rightleftharpoons e + N_2(X,w)$, where an electron collides with a nitrogen molecule in a vibrational level v , and excites (or de-excites) it to a higher (or lower) vibrational level w . Similar reactions exist for hydrogen molecules.
- Vibrational energy transfer, e.g. $N_2(X,v) + H_2(X,w) \rightleftharpoons N_2(X,v+1) + H_2(X,w-1)$, where nitrogen and hydrogen molecules collide, and energy is transferred from one to the other through their vibrational levels.

For other electron impact processes, there are reactions such as:

- Electronic excitation and de-excitation, e.g. $e + N_2(X) \rightleftharpoons e + N_2(A)$, where the collision of an electron with a nitrogen molecule causes its excitation from the ground state (X) to a higher energy state (A), or vice-versa.

- Dissociation, e.g. $e + \text{H}_2(\text{X}) \rightarrow e + 2\text{H}$, where the collision of an electron with a hydrogen molecule breaks its molecular bond and dissociates it into two hydrogen atoms.
- Ionization, e.g. $e + \text{N}_2(\text{X}) \rightarrow \text{N}_2^+ + 2e$, where the collision of an electron with a nitrogen molecule ejects an electron from the molecule, thereby ionizing it.
- Attachment and detachment, e.g. $e + \text{H}_2 \rightarrow \text{H} + \text{H}^-$, where the collision of an electron with a hydrogen molecule causes the dissociation of the molecule, and the attachment of the electron to one of the resulting hydrogen atoms, forming a hydride ion (H^-).
- Neutralization, e.g. $e + \text{N}_2^+ \rightarrow \text{N} + \text{N}$, where the incoming electron recombines with the positively charged nitrogen ion, thereby neutralizing it, and, in addition, the excess energy of the newly formed molecular state leads to the breaking of the molecular bond between the two nitrogen atoms.

For heavy-species collisions, there are reactions such as:

- Collisional excitation and de-excitation, possibly with dissociation, e.g. $\text{N}_2(\text{A}) + \text{H}_2 \rightarrow \text{N}_2(\text{X}) + 2\text{H}$, where a nitrogen molecule in an excited state transfers energy to a hydrogen molecule. The nitrogen molecule returns to its ground state, and the energy transferred breaks the molecular bond between the two hydrogen atoms.
- Ionization, e.g. $\text{N}_2(\text{A}) + \text{N}_2(\text{a}') \rightarrow \text{N}_2(\text{X}) + \text{N}_2^+ + e$, where there is an energy transfer between two nitrogen molecules in different excited states. One of them returns to the ground state, while the other receives sufficient energy to eject an electron and ionize.
- Neutral-neutral reactions, e.g. $\text{N}(\text{D}) + \text{H}_2(\text{X}) \rightarrow \text{H} + \text{NH}$, where an excited nitrogen atom transfers energy to a hydrogen molecule, enough to break its molecular bond. The nitrogen atom then forms a new bond with one of the free hydrogen atoms.
- Ion-molecule reactions, e.g. $\text{N}_2^+ + \text{H}_2 \rightarrow \text{N}_2\text{H}^+ + \text{H}$, where the nitrogen ion breaks the bond of a hydrogen molecule, and a new bond is formed between the nitrogen ion and one of the hydrogen atoms.
- Ion-ion recombination, e.g. $\text{H}^- + \text{H}_2^+ \rightarrow \text{H}_2 + \text{H}$, where a hydride ion (H^-) transfers an excess electron to neutralize the charge of the dihydrogen cation (H_2^+), resulting in a hydrogen molecule and a hydrogen atom.

The reactions above illustrate the types of mechanisms at work, and these mechanisms apply not only to N₂ and H₂ but also to atoms, radicals and ions resulting from the various chemical reactions. The chemistry is therefore rich enough to create a multitude of species (namely H, H⁺, H₂, H₂⁺, H₃⁺, N, N⁺, N₂, N₂⁺, N₃⁺, N₄⁺, NH, NH⁺, NH₂, NH₂⁺, NH₂⁻, NH₃, NH₃⁺, NH₄⁺, N₂H⁺, N₂H₂, N₂H₃, N₂H₄) and these species can be in different excited states, particularly electronic and vibrational. For simplicity, we will omit the details about the electronic/vibrational configuration of those species.

Besides the reactions that take place in the bulk of the plasma – often referred to as *volume* processes [7] – there is another set of reactions that take place at the boundaries of the experiment (at the interface between the gas and its

surrounding walls), which are referred to as *surface* processes [8]. The plasma-surface interaction stimulates the reaction of some species (namely H, N, NH, NH₂) with the surface, by some physical or chemical means. A species may stick to the surface (physically) or react with the surface (chemically). Subsequently, there may be reactions between species that have previously adhered to the wall (involving only surface species), and there may be reactions between species that have previously adhered to the wall and species that come from the bulk of the plasma (involving both surface and volume species). In some cases, the products of these reactions can return to the volume.

To distinguish between volume and surface species, we use the “wall” prefix for the latter. For example, NH₂ refers to the volume species, but wall_NH₂ refers to the surface species. In addition, we represent vacant sites where species can physically attach to the surface as wall_F_v, and vacant sites where species can chemically bind to the surface as wall_S_v.

The surface chemistry can then be formulated as, for example:

- $\text{NH(X)} + \text{wall_S}_v \rightarrow \text{wall_NH(S)}$, where NH adheres chemically to the wall.
- $\text{NH(X)} + \text{wall_H(S)} \rightarrow \text{wall_NH}_2(\text{S})$, where an incoming NH molecule reacts with a hydrogen atom that is chemically bound to the wall.
- $\text{wall_H(F)} + \text{wall_NH(S)} \rightarrow \text{wall_NH}_2(\text{S}) + \text{wall_F}_v$, where a hydrogen atom that was physically attached to the wall reacts with a NH molecule that is chemically bound to the wall. The resulting NH₂ is bound to the wall and the physical site that released the hydrogen atom becomes vacant.
- $\text{H}_2(\text{X}) + \text{wall_NH(S)} \rightarrow \text{NH}_3(\text{X}) + \text{wall_S}_v$, where an incoming hydrogen molecule reacts with a NH molecule that was chemically bound to the wall, releasing NH₃ and leaving the chemical site vacant.

These examples illustrate the surface chemistry that needs to be taken into account, simultaneously with the reactions that take place in the plasma volume.

3 Petri Net Modeling

One way to model chemical reactions is through the use of Petri nets [14]. In essence, a Petri net is a graphical model with two types of nodes: *places* and *transitions*. Places connect to transitions, and transitions connect to places, through directed arcs. In addition, places have *tokens*. When a transition *fires*, it removes tokens from its input places, and adds tokens to its output places. By default, it removes one token from each of its input places, and adds one token to each of its output places. However, the number of tokens to be removed from (or added to) each place can be configured by labeling the arc connecting the place to (or from) the transition. For a transition to be *enabled*, each of its input places must contain a sufficient number of tokens; otherwise, it cannot fire.

Figure 1 shows an example of how a chemical reaction can be represented as a Petri net transition, with the reactants as input places, and the products as output places. For illustrative purposes, we use the Haber-Bosch reaction $\text{N}_2 + 3\text{H}_2 \rightarrow 2\text{NH}_3$. In Fig. 1(a), the transition t is enabled by having a sufficient




Fig. 1. A chemical reaction modeled as a Petri net transition t .

number of tokens in each of its input places; this is called the initial *marking*. When t fires, one token is subtracted from place N_2 , three tokens are subtracted from place H_2 , and two tokens are added to place NH_3 . The resulting marking is shown in Fig. 1(b). In this case, t is no longer enabled but, in a general case, the presence of tokens in NH_3 could enable subsequent transitions.

The same principles can be applied to model chemical reaction networks [13]. The idea is perhaps better explained with an example. For this purpose, we consider a set of reactions that is by no means a complete description of the $\text{N}_2\text{-H}_2$ plasma chemistry, but nevertheless includes some of the mechanisms that are present in such chemistry, with the mere purpose of illustrating an over-simplified ammonia production process:

- $\text{N}_2 \rightleftharpoons 2\text{N}$, which can be regarded as two separate (forward and backward) reactions, to be represented as transitions t_1 and t_2 .
- $\text{N}_2 + \text{H}_2 \rightarrow \text{N}_2 + 2\text{H}$, to be represented as transition t_3 .
- $2\text{H} \rightarrow \text{H}_2$, to be represented as transition t_4 .
- $\text{N} + \text{H} \rightarrow \text{NH}$, to be represented as transition t_5 .
- $\text{NH} + \text{H} \rightarrow \text{NH}_2$, to be represented as transition t_6 .
- $\text{NH}_2 + \text{H} \rightarrow \text{NH}_3$, to be represented as transition t_7 .
- $\text{NH}_3 + \text{N} \rightarrow \text{NH}_2 + \text{NH}$, to be represented as transition t_8 .

(In this example, the list of reactions ignores the presence of electrons, ions and wall species, and accordingly does not consider ionization reactions.)

Figure 2 shows the corresponding Petri net. One can imagine that by placing a generous amount of tokens in places N_2 and H_2 , such initial marking could enable a series of transition firings until, eventually, some amounts of NH , NH_2 and NH_3 are produced. Due to the cyclic chaining of some of these reactions, such firings could go on indefinitely until the experiment is interrupted, or until all the hydrogen and/or nitrogen tokens are consumed.

Mathematically, the structure of a Petri net can be represented as a transition matrix, where each element specifies the number of tokens that are subtracted from (or added to) a given place by a given transition. For the example in Fig. 2, the transition matrix is:




Fig. 2. A chemical reaction network modeled as a Petri net.

$$A = \begin{array}{cccccccc|l} & t_1 & t_2 & t_3 & t_4 & t_5 & t_6 & t_7 & t_8 & \\ \hline 0 & 0 & 2 & -2 & -1 & -1 & -1 & 0 & 0 & \text{H} \\ 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & \text{H}_2 \\ 2 & -2 & 0 & 0 & -1 & 0 & 0 & 0 & -1 & \text{N} \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \text{N}_2 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 1 & 0 & \text{NH} \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & 1 & 0 & \text{NH}_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & \text{NH}_3 \end{array} \quad (1)$$

Let us suppose that, initially, there are two tokens in N_2 and two tokens in H_2 . This initial marking of the Petri net can be represented as a vector:

$$b = \begin{bmatrix} 0 \\ 2 \\ 0 \\ 2 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \begin{array}{l} \text{H} \\ \text{H}_2 \\ \text{N} \\ \text{N}_2 \\ \text{NH} \\ \text{NH}_2 \\ \text{NH}_3 \end{array} \quad (2)$$

From this initial marking, let us suppose that t_1 fires once, t_3 fires twice, and t_5 , t_6 , t_7 each fire once. This can be represented as a firing vector:

$$x = \begin{bmatrix} 1 \\ 0 \\ 2 \\ 0 \\ 1 \\ 1 \\ 1 \\ 0 \end{bmatrix} \begin{array}{l} t_1 \\ t_2 \\ t_3 \\ t_4 \\ t_5 \\ t_6 \\ t_7 \\ t_8 \end{array} \quad (3)$$

Then the final marking of the Petri net can be calculated as $Ax + b$:

$$y = Ax + b = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} \begin{matrix} \text{H} \\ \text{H}_2 \\ \text{N} \\ \text{N}_2 \\ \text{NH} \\ \text{NH}_2 \\ \text{NH}_3 \end{matrix} \quad (4)$$

(The interested reader might want to check that this is indeed the case by imagining the movement of tokens in the Petri net of Fig. 2.)

In this formulation, the most important problem for this work is how to find x given A , b and y . In general, the chemical scheme (i.e. the set of possible reactions), represented as matrix A , is known. The initial amounts of hydrogen and nitrogen that are provided at the beginning of the experiment or simulation, represented as vector b , are also known. The final amounts of species at the end of the experiment or simulation, represented as vector y , can be measured or calculated. Therefore, what is left to know is the vector x , which is proxy for the reaction rates, i.e. the number of times those reactions occur per unit time and unit volume (or surface). In this work, we will interpret x as a measure of weight or importance of each reaction in the chemical scheme.

4 Machine Learning Model

As we have seen in the previous section, A is an $n \times m$ matrix, where n is the number of species, and m is the number of reactions. In practice, depending on the chemical scheme, the number of reactions m can be (much) larger than the number of species n . This means that $Ax + b = y$ is an under-determined system. While there might be many solutions for x , a possible approach would be to find an approximation $\tilde{x}$ (and its corresponding prediction $\tilde{y} = A\tilde{x} + b$) that minimizes the residual $\|\tilde{y} - y\|$ and has the smallest norm $\|\tilde{x}\|$.

In our case, there are additional restrictions to be imposed on the model:

- The first restriction is that every element of the solution (or approximation) $\tilde{x}$ should be non-negative, i.e. $\forall_i : \tilde{x}_i \geq 0$. This means that negative transition firings (or, in a chemical sense, negative reaction rates) are disallowed.
- The second restriction has to do with how the experimental/calculated data about the initial and final amounts (i.e. the density) of species are typically provided. While absolute values for the density of species (in m⁻³ for volume species, and in m⁻² for surface species) are usually available, it is quite common to describe an experiment as comprising certain percentages of molecular hydrogen and nitrogen as input (e.g. 5% of H₂ and 95% of N₂) and producing a certain *distribution* of species as output (also expressed in percentages). As a result, this means that the model should provide a normalized output $\tilde{y}$. However, the distributions of volume species and surface

species should be normalized separately, hence part of $\tilde{y}$ (the partition $\tilde{y}_{i \in \mathcal{V}}$ that corresponds to volume species) is normalized to 1, and another part of $\tilde{y}$ (the partition $\tilde{y}_{i \in \mathcal{S}}$ that corresponds to surface species) is also normalized to 1. In other words:

$$\begin{cases} \sum_{i \in \mathcal{V}} \tilde{y}_i = 1 & \text{for volume species in } \mathcal{V} = \{\text{H}, \text{H}_2, \text{N}, \text{N}_2, \dots\} \\ \sum_{i \in \mathcal{S}} \tilde{y}_i = 1 & \text{for surface species in } \mathcal{S} = \{\text{wall_H}, \text{wall_N}, \dots\} \end{cases} \quad (5)$$

The non-negativity restriction on $\tilde{x}$ and the normalization requirements on $\tilde{y}$ can be implemented via standard machine learning constructs, one of them being the *rectified linear unit* (ReLU) and the other being the *softmax* function. Both functions are extensively used in the field of machine learning [1].

In particular, ReLU is an element-wise function that can be defined as:

$$\text{ReLU}(\tilde{x}) = [\dots, \text{ReLU}(\tilde{x}_i), \dots]^\top = [\dots, \max(0, \tilde{x}_i), \dots]^\top \quad (6)$$

On the other hand, softmax is a normalization function defined as:

$$\text{softmax}(\tilde{y}) = [\dots, \text{softmax}(\tilde{y}_i), \dots]^\top = \left[\dots, \frac{e^{\tilde{y}_i}}{\sum_j e^{\tilde{y}_j}}, \dots \right]^\top \quad (7)$$

The purpose of using these functions is to allow $\tilde{x}$ and $\tilde{y}$ to remain unconstrained, while enforcing the non-negativity of $\tilde{x}$ and the normalization of $\tilde{y}$ by applying ReLU and softmax to them, respectively.

Since softmax is applied separately on volume and surface species, the model can be expressed as:

$$\begin{cases} \text{softmax}((A \cdot \text{ReLU}(\tilde{x}) + b)_{i \in \mathcal{V}}) = \tilde{y}_{i \in \mathcal{V}} \\ \text{softmax}((A \cdot \text{ReLU}(\tilde{x}) + b)_{i \in \mathcal{S}}) = \tilde{y}_{i \in \mathcal{S}} \end{cases} \quad (8)$$

In the model above, a guess for $\tilde{x}$ will produce predictions for $\tilde{y}_{i \in \mathcal{V}}$ and $\tilde{y}_{i \in \mathcal{S}}$. The goal is to find the best guess $\tilde{x}$, i.e. the guess $\tilde{x}$ for which the predictions $\tilde{y}_{i \in \mathcal{V}}$ and $\tilde{y}_{i \in \mathcal{S}}$ best approximate, in some sense, the true $y_{i \in \mathcal{V}}$ and $y_{i \in \mathcal{S}}$.

Now, if the true $y_{i \in \mathcal{V}}$ and the prediction $\tilde{y}_{i \in \mathcal{V}}$ are both normalized (the same applies to $y_{i \in \mathcal{S}}$ and $\tilde{y}_{i \in \mathcal{S}}$), they can be reinterpreted as probability distributions. This provides the opportunity for using a loss function based on statistical measures, such as the Kullback–Leibler (KL) divergence, to quantify the difference between a true probability distribution $y_{i \in \mathcal{V}}$ and a predicted probability distribution $\tilde{y}_{i \in \mathcal{V}}$ (and, similarly, between $y_{i \in \mathcal{S}}$ and $\tilde{y}_{i \in \mathcal{S}}$).

Specifically, the KL divergence measures how much information is lost when a true distribution p is approximated by a predicted distribution q . In this case, the KL divergence is defined as:

$$D_{\text{KL}}(p \parallel q) = \sum_i p_i \log \left(\frac{p_i}{q_i} \right) \quad (9)$$

where p and q are vectors of probabilities over the same outcomes.

In our case, the loss function is defined as:

$$\begin{aligned} L(\tilde{x}) &= D_{\text{KL}}(y_{i \in \mathcal{V}} \parallel \tilde{y}_{i \in \mathcal{V}}) + D_{\text{KL}}(y_{i \in \mathcal{S}} \parallel \tilde{y}_{i \in \mathcal{S}}) \\ &= \sum_{i \in \mathcal{V}} y_i \log \left(\frac{y_i}{\tilde{y}_i} \right) + \sum_{i \in \mathcal{S}} y_i \log \left(\frac{y_i}{\tilde{y}_i} \right) \end{aligned} \quad (10)$$

The goal is to find $\tilde{x}$ that minimizes $L(\tilde{x})$. For this purpose, the gradient of $L(\tilde{x})$, denoted as $\nabla L(\tilde{x})$, provides the direction in which the loss increases. Since the objective is to minimize the loss, we update $\tilde{x}$ iteratively via:

$$\tilde{x}^{(k+1)} = \tilde{x}^{(k)} - \eta \nabla L(\tilde{x}^{(k)}) \quad (11)$$

where η is a positive learning rate. In other words, the model is trained by *gradient descent*, or variants thereof [15].

At the end of the training process, $\tilde{x}$ will provide the weights assigned to the reactions in the chemical scheme. Any negative weight found in $\tilde{x}$, by force of the ReLU function, will be turned into zero, and therefore the corresponding reaction does not contribute to the result $\tilde{y}$. In fact, all those reactions for which $\tilde{x}_i \leq 0$ can be considered as candidates for removal from the chemical scheme, and this becomes the basis for chemistry reduction in this work.

5 Chemical Scheme Reduction

In [3], the authors provide an experimental study of N₂-H₂ discharge plasmas, based on a low-pressure (50–500 Pa), high-voltage (1–3 kV), direct-current (10–40 mA) setup. The gas mixture, containing 0–5% H₂, is pumped into a 23 cm-long tube with a 2 cm diameter, and mass spectrometry is used to measure the formation of neutral species (especially NH₃) and many ion species, such as H⁺, H₂⁺, H₃⁺, N⁺, N₂⁺, N₃⁺, N₄⁺, NH⁺, NH₂⁺, NH₃⁺, NH₄⁺, and N₂H⁺. The formation of NH₃ (and of positive ions in the form NH_x⁺) was found to be correlated with current and pressure.

An alternative way to study the chemistry of low-temperature plasmas is through numerical modeling, using computer-based simulation tools, such as the LisbOn KInetics (LoKI) code [16,9]. LoKI comprises two modules:

- The first module is a solver for the electron Boltzmann equation that provides the statistical distribution of electrons according to their kinetic energy, i.e. the so-called *electron energy distribution function* (often abbreviated as *eedf*). Knowing the *eedf* is crucial to estimate the rate coefficients at which electrons excite, dissociate and ionize atoms and molecules in the gas, in order to study the overall dynamics of the system.
- In a second module, LoKI handles the plasma chemistry by determining the reaction rates for a provided chemical scheme. The chemical scheme is specified in a configuration file that contains hundreds of reactions collected from the literature. A comprehensive list is provided in [11]. The rate coefficients of some reactions are specified in the chemical scheme, while others

are determined based on the *eedf* calculated by the first module. The reaction rates are calculated by multiplying the corresponding rate coefficients by the densities of the reactants intervening in each reaction. By balancing the creation and destruction rates for the different reactions in the chemical scheme, LoKI computes the final densities for all the species in the plasma.

LoKI receives the initial densities of species as input fractions in a configuration file (e.g. 5% of H_2 expressed as 0.05, and 95% of N_2 expressed as 0.95, for a pure $\text{N}_2\text{-H}_2$ mixture similar to experimental conditions). Then it computes and saves the final densities in an output file. These final densities can be converted to fractions, so that input and output fractions can be directly compared.

Table 1 shows an example of the input and output fractions for a LoKI simulation. For simplicity, we abstract from vibrational levels and electronic states (e.g. N_2 represents all $\text{N}_2(A,v)$ species regardless of vibrational level v and

Table 1. Input and output fractions for a LoKI simulation.

Species	Input fraction	Output fraction
H	0.045	$4.076\,311 \times 10^{-2}$
H^+	0	$1.521\,387 \times 10^{-11}$
H_2	0.005	$2.821\,790 \times 10^{-2}$
H_2^+	0	$1.232\,711 \times 10^{-12}$
H_3^+	0	$1.350\,015 \times 10^{-10}$
N	0	$1.385\,849 \times 10^{-2}$
N^+	0	$1.149\,678 \times 10^{-12}$
N_2	0.95	$9.171\,353 \times 10^{-1}$
N_2^+	0	$1.998\,340 \times 10^{-11}$
N_2H^+	0	$2.301\,581 \times 10^{-7}$
N_2H_2	0	$1.186\,624 \times 10^{-9}$
N_2H_3	0	$1.913\,678 \times 10^{-11}$
N_2H_4	0	$2.236\,752 \times 10^{-15}$
N_3^+	0	$8.794\,668 \times 10^{-11}$
N_4^+	0	$1.844\,171 \times 10^{-13}$
NH	0	$1.477\,812 \times 10^{-6}$
NH^+	0	$1.616\,889 \times 10^{-12}$
NH_2	0	$1.188\,158 \times 10^{-8}$
NH_2^+	0	$1.287\,649 \times 10^{-10}$
NH_2^-	0	$6.742\,121 \times 10^{-9}$
NH_3	0	$2.347\,209 \times 10^{-5}$
NH_3^+	0	$1.444\,042 \times 10^{-9}$
NH_4^+	0	$2.753\,958 \times 10^{-8}$
wall_ F_v	0.9976	$9.974\,761 \times 10^{-1}$
wall_ H	0.002	$2.313\,191 \times 10^{-3}$
wall_ N	0.0002	$1.184\,274 \times 10^{-4}$
wall_ NH	0.0002	$9.229\,275 \times 10^{-5}$
wall_ NH_2	0	$8.429\,000 \times 10^{-9}$
wall_ S_v	0	$1.202\,643 \times 10^{-8}$

electronic state A). Therefore, the input/output fractions of each species include all the electronic and vibrational populations, if applicable. More importantly, the table shows volume species first, and surface species at the end. As can be seen, the input fractions are normalized separately, i.e. the fractions for volume species sum up to 1, and the fractions for surface species also sum up to 1; the same applies to the output fractions.³

With the previous information, it is possible to apply the model described in Section 4, which includes matrix A and vectors b and y :

- Matrix A is built from the reactions in the chemical scheme, by following the same approach that led to Eq. (1). Basically, matrix A contains the stoichiometric coefficients of each reaction, with reactants being represented as negative values, and products being represented as positive values. Converting the chemical scheme to A results in a matrix of size 29×160 , where there are 160 reactions involving the 29 species in Table 1.
- As for b and y , these vectors correspond, respectively, to the input and output fractions presented in Table 1. The conversion of those fractions to b and y results in two vectors of size 29×1 .

In the framework of $Ax + b = y$, the goal would be to find the solution x that explains how much each reaction contributes to transform the input fractions b into the output fractions y , according to the chemical scheme A . In this context, x will be a vector of size 160×1 with a weight for each reaction.

However, in Section 4 we have seen that such system is under-determined, as shown here by the fact that we have 29 equations for 160 unknowns. Besides, there are certain restrictions to be imposed, namely the non-negativity of any approximation $\tilde{x}$, and the normalization of any prediction $\tilde{y}$.

This has led us to an approach in the form $\text{softmax}(A \cdot \text{ReLU}(\tilde{x}) + b) = \tilde{y}$, where the objective is to find an approximate solution $\tilde{x}$ via gradient descent, based on the model in Eq. (8), the loss function in Eq. (10), and a learning rule similar to Eq. (11). In this work, we use the popular Adam optimizer [12] with a small learning rate, i.e. $\eta = 10^{-4}$.

Figure 3 shows the evolution of the loss values across training. The loss (i.e. the KL divergence) decreases sharply during the first few thousand iterations, and then continues to improve, although at a much slower pace, until it eventually converges to a value on the order of 10^{-6} .

At the end of the training process, we extract $\tilde{x}$ and analyze the weights assigned to reactions. Reactions with a positive weight will be kept in the chemical scheme, as they are found to play a role in transforming the input fractions into the output fractions. On the other hand, reactions with zero or negative weight are considered as candidates for removal from the chemical scheme since, by force of the ReLU function, their contribution will be zero.

³ The input fractions may be different for a simulation and for a lab experiment. For example, simulations usually define not only the fractions of molecular species (e.g. H₂ and N₂) but also of atomic species (e.g. H and N) and wall species (e.g. wall_H, wall_N, etc.). On the other hand, in a lab experiment, the input for the gas mixture is limited to the actual fraction of gases available in commercial bottles.




Fig. 3. Loss curve obtained during training.




Fig. 4. Petri net representation of the chemical scheme, with reactions to be kept (in blue), and reactions to be removed (in orange) [colors available in the online version].

Figure 4 shows a Petri net representation of the chemical scheme, where the reactions with positive weight (these ones that will be kept) are shown in blue, and the reactions with zero or negative weight (the ones that will be considered for removal) are shown in orange. From the 160 reactions in the original chemical scheme, only about 60 reactions (less than 40%) are guaranteed to survive, with the rest becoming candidates for removal. This highlights the extent to which the chemical scheme can be potentially reduced.

6 Analysis of the Results

An analysis of the reduced chemical scheme, based on the machine learning model described in Section 4, reveals that there are several possible mechanisms for the production of NH₃:

- One of those mechanisms is: $\text{wall_H} + \text{wall_NH}_2 \rightarrow \text{NH}_3$ (simplified form). In this case, a hydrogen atom attached to the wall reacts with a NH₂ molecule bound to the wall, releasing NH₃ into the plasma volume. This requires the availability of wall_NH₂, which can be produced either by the surface chemistry ($\text{wall_H} + \text{wall_NH} \rightarrow \text{wall_NH}_2$) or may come from the volume chemistry ($\text{NH}_2 \rightarrow \text{wall_NH}_2$). The former depends on the availability of wall_NH, which can be produced via similar pathways (either from wall_H and wall_N, or from the volume species NH binding to the wall); the latter depends on the volume species NH₂, which can be produced in a number of ways, usually involving neutralization of NH₂⁺ or NH₂⁻. These ion species can be traced back to a multitude of volume reactions.
- Another mechanism is $\text{H} + \text{wall_NH}_2 \rightarrow \text{NH}_3$ (simplified form), where a hydrogen atom reacts with a NH₂ molecule bound to the surface, releasing NH₃ into the plasma volume. Much of the same considerations apply here as well, and this serves to illustrate how the surface chemistry can play an important role in the production of NH₃. Once H and N bind to the wall, they can react chemically to produce wall_NH, wall_NH₂, and eventually NH₃. Another chemical pathway is for the volume species NH and NH₂ to bind to the wall midway through this process.
- In the partially ionized environment of the plasma volume, there are multiple ion species. Some of these ion species are highly reactive, and can produce NH₃ in a number of different ways. For example, when H₃⁺ encounters NH₂⁻, they recombine to produce NH₃ (via $\text{H}_3^+ + \text{NH}_2^- \rightarrow \text{H}_2 + \text{NH}_3$). If NH₂⁻ encounters NH₄⁺, they can also produce NH₃ (via $\text{NH}_2^- + \text{NH}_4^+ \rightarrow \text{NH}_2 + \text{NH}_3 + \text{H}$). In this case, a curious fact is that NH₄⁺ can only exist if it has been produced from NH₃ (e.g. via $\text{NH}^+ + \text{NH}_3 \rightarrow \text{N} + \text{NH}_4^+$). Therefore, there is an interplay between ions and neutral species, where they produce each other recursively with the help of other species. Another example is when NH₃ is produced by neutralization of NH₃⁺, but NH₃⁺ itself is produced from NH₃ and other ions (e.g. via $\text{H}_2^+ + \text{NH}_3 \rightarrow \text{H}_2 + \text{NH}_3^+$). All of these mechanisms contribute to the co-existence and continuous interplay between ions and neutrals in the plasma volume.




Fig. 5. Graph representation of the chemical scheme, with reactions to be kept (solid lines), and reactions to be removed (dashed lines) [colors available in the online version].

- Besides ion reactions, NH₃ can also be produced from neutrals, especially through reactions that involve NH₂ (e.g. $\text{NH} + \text{NH}_2 \rightarrow \text{N} + \text{NH}_3$; or $2\text{NH}_2 \rightarrow \text{NH} + \text{NH}_3$). The heavier species N₂H₂, N₂H₃ and N₂H₄, if present, can also react with NH₂ to produce NH₃ (e.g. $\text{N}_2\text{H}_2 + \text{NH}_2 \rightarrow \text{H} + \text{N}_2 + \text{NH}_3$). However, the origin of NH₂ is difficult to disentangle, since it appears that NH₂ arises from the neutralization of NH₂⁺, but NH₂⁺ itself arises from the ionization of NH₂ when it collides with ions such as N⁺, H₂⁺ or NH⁺. Another pathway to NH₃ is through neutralizing reactions involving NH₂⁻, but this also unclear, since NH₂⁻ itself originates from electron impact on NH₃. In conclusion, these species (i.e. NH₂, NH₂⁺, NH₂⁻, NH₃) seem to coexist and continuously recreate each other with the help of other species.

Figure 5 shows a graph representation of the chemical scheme that may shed some light on these intricate mechanisms. Here, each arrow represents a reactant-product relationship (reactant→product) between two species, so a single reaction (such as $\text{H}_3^+ + \text{NH}_2^- \rightarrow \text{H}_2 + \text{NH}_3$) may be drawn as multiple arrows (e.g. $\text{H}_3^+ \rightarrow \text{H}_2$; $\text{H}_3^+ \rightarrow \text{NH}_3$; $\text{NH}_2^- \rightarrow \text{H}_2$; $\text{NH}_2^- \rightarrow \text{NH}_3$). The solid arrows represent reactions that are to be kept in the reduced chemical scheme, while the dashed arrows represent candidates for removal.

One of the main features that can be observed in this graph is the clear separation between volume and surface chemistry. It is clear that H and N will attach to the wall and, from there, there is a pathway towards the production of NH₃ through wall_NH₂. On the other hand, the volume chemistry shows a multitude of pathways towards NH₃, as well as the reciprocal relationships that exist among NH₃, NH₂, NH₂⁻, NH₂⁺, and even NH₄⁺. This visual portrayal is in agreement with our analysis above.

7 Conclusion

In this work, we developed a machine learning model inspired by a Petri net representation of the chemical scheme. When such representation is converted to matrix form, it becomes possible to use standard machine learning constructs. Here, we defined a learning model based on matrix-vector multiplication, and on the ReLU and softmax functions. Such model is trained by minimizing the Kullback-Leibler divergence in order to find a set of reaction weights. These weights become the basis for reducing the chemical scheme.

The model was applied to the chemistry of low-temperature N₂-H₂ plasmas. Despite being a relatively simple setup, this gas discharge exhibits a rich chemistry, where both ions and neutrals, as well as volume and surface species, have an interconnected role. With the proposed model, it was possible to identify the main reactions and chemical pathways involved in the production of NH₃. In future work, we plan to further develop the approach by comparing the simulation results obtained with the full and the reduced chemical schemes, and also validate them against experimental results obtained in the lab.

Acknowledgments. L. L. Alves would like to acknowledge the support of FCT (Fundação para a Ciência e a Tecnologia, I.P.) through project reference 2022.04128.PTDC. IPFN activities were supported by FCT through project references UIDB/50010/2020, UIDP/50010/2020 and LA/P/0061/202.

References

1. Asadi, B., Jiang, H.: On approximation capabilities of ReLU activation and softmax output layer in neural networks. arXiv 2002.04060 (2020)
2. Carreon, M.L.: Plasma catalytic ammonia synthesis: state of the art and future directions. *Journal of Physics D: Applied Physics* **52**(48), 483001 (2019)
3. Chatain, A., Morillo-Candas, A.S., Vettier, L., Carrasco, N., Cernogora, G., Guaitella, O.: Characterization of a DC glow discharge in N_2 - H_2 with electrical measurements and neutral and ion mass spectrometry. *Plasma Sources Science and Technology* **32**(3), 035002 (2023)
4. Chudják, S., Kozáková, Z., Krčma, F.: Study of chemical processes initiated by electrical discharge in Titan-related atmosphere at laboratory temperature and pressure. *ACS Earth and Space Chemistry* **5**(3), 535–543 (2021)
5. Encrenaz, T., Owen, T., Woodman, J.H.: The abundance of ammonia on Jupiter, Saturn and Titan. *Astronomy and Astrophysics* **37**(1), 49–55 (1974)
6. Erisman, J.W., Sutton, M.A., Galloway, J., Klimont, Z., Winiwarter, W.: How a century of ammonia synthesis changed the world. *Nature Geoscience* **1**(10), 636–639 (2008)
7. Gordiets, B., Ferreira, C.M., Pinheiro, M.J., Ricard, A.: Self-consistent kinetic model of low-pressure N_2 - H_2 flowing discharges: I. Volume processes. *Plasma Sources Science and Technology* **7**(3), 363 (1998)
8. Gordiets, B., Ferreira, C.M., Pinheiro, M.J., Ricard, A.: Self-consistent kinetic model of low-pressure N_2 - H_2 flowing discharges: II. Surface processes and densities of N, H, NH_3 species. *Plasma Sources Science and Technology* **7**(3), 379 (1998)
9. Guerra, V., Tejero-del-Caz, A., Pintassilgo, C.D., Alves, L.L.: Modelling N_2 - O_2 plasmas: volume and surface kinetics. *Plasma Sources Science and Technology* **28**(7), 073001 (2019)
10. Huang, J., Seager, S., Petkowski, J.J., Ranjan, S., Zhan, Z.: Assessment of ammonia as a biosignature gas in exoplanet atmospheres. *Astrobiology* **22**(2), 171–191 (2022)
11. Jiménez-Redondo, M., Chatain, A., Guaitella, O., Cernogora, G., Carrasco, N., Alves, L.L., Marques, L.: N_2 - H_2 capacitively coupled radio-frequency discharges at low pressure: II. Modeling results: the relevance of plasma-surface interaction. *Plasma Sources Science and Technology* **29**(8), 085023 (2020)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv:1412.6980 (2014)
13. Koch, I.: Petri nets – A mathematical formalism to analyze chemical reaction networks. *Molecular Informatics* **29**(12), 838–843 (2010)
14. Marwan, W., Wagler, A., Weismantel, R.: Petri nets as a framework for the reconstruction and analysis of signal transduction pathways and regulatory networks. *Natural Computing* **10**(2), 639–654 (2011)
15. Ruder, S.: An overview of gradient descent optimization algorithms. arXiv 1609.04747 (2017)
16. Tejero-del-Caz, A., Guerra, V., Gonçalves, D., da Silva, M.L., Marques, L., Pinhão, N., Pintassilgo, C.D., Alves, L.L.: The LisbOn KInetics Boltzmann solver. *Plasma Sources Science and Technology* **28**(4), 043001 (2019)