



HAL
open science

Stochastic Chemical Reaction Networks for MAP Detection in Cellular Receivers

Bastian Heinlein, Lukas Brand, Malcolm Egan, Maximilian Schäfer, Robert Schober, Sebastian Lotter

► **To cite this version:**

Bastian Heinlein, Lukas Brand, Malcolm Egan, Maximilian Schäfer, Robert Schober, et al.. Stochastic Chemical Reaction Networks for MAP Detection in Cellular Receivers. NANOCOM '23: The 10th Annual ACM International Conference on Nanoscale Computing and Communication, Sep 2023, Coventry, United Kingdom. pp.65-71, 10.1145/3576781.3608709 . hal-04361281

HAL Id: hal-04361281

<https://hal.science/hal-04361281>

Submitted on 22 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stochastic Chemical Reaction Networks for MAP Detection in Cellular Receivers

Bastian Heinlein

Friedrich-Alexander-Universität
Erlangen-Nürnberg,
Erlangen, Germany

Lukas Brand

Friedrich-Alexander-Universität
Erlangen-Nürnberg,
Erlangen, Germany

Malcolm Egan

Univ. Lyon, Inria, INSA Lyon,
Villeurbanne, France

Maximilian Schäfer

Friedrich-Alexander-Universität
Erlangen-Nürnberg,
Erlangen, Germany

Robert Schober

Friedrich-Alexander-Universität
Erlangen-Nürnberg,
Erlangen, Germany

Sebastian Lotter

Friedrich-Alexander-Universität
Erlangen-Nürnberg,
Erlangen, Germany

ABSTRACT

In order to fully exploit the potential of molecular communication (MC) for intra-body communication, practically implementable cellular receivers are an important long-term goal. A variety of receiver architectures based on chemical reaction networks (CRNs) and gene-regulatory networks (GRNs) has been introduced in the literature, because cells use these concepts to perform computations in nature. However, practical feasibility is still limited by stochastic fluctuations of chemical reactions and long computation times in GRNs. Therefore, in this paper, we propose two receiver designs based on stochastic CRNs, i.e., CRNs that perform computations by exploiting the intrinsic fluctuations of chemical reactions with very low molecule counts. The first CRN builds on a recent result from chemistry that showed how Boltzmann machines (BMs), a commonly used machine learning model, can be implemented with CRNs. We show that BMs with optimal parameter values and their CRN implementations can act as maximum-a-posteriori (MAP) detectors. Furthermore, we show that BMs can be efficiently trained from simulation data to achieve close-to-MAP performance. While this approach yields a fixed CRN once deployed, our second approach based on a manually designed CRN can be trained with pilot symbols even within the cell and thus adapt to changing channel conditions. We extend the literature by showing that practical robust detectors can achieve close-to-MAP performance even without explicit channel knowledge.

CCS CONCEPTS

• **Applied computing** → **Telecommunications**; • **Hardware** → **Biology-related information processing**; • **Mathematics of computing** → *Probabilistic inference problems*.

KEYWORDS

Molecular communication, Boltzmann machine, chemical reaction network, maximum-a-posteriori detection, machine learning

1 INTRODUCTION

Molecular Communication (MC) is a new paradigm for information exchange in conditions that are unfavorable for traditional wireless communication, e.g., at nano-scale or inside the human body. MC has great potential, for example in the context of the Internet of Bio-Nano Things which will enable groundbreaking improvements in healthcare by employing a network of connected nano-sensors

within the body to diagnose and treat diseases [1].

For both theoretical work [2] and experimental testbeds [3] progress has been made to realize this ambitious vision. A further leap forward could be possible if practical implementations of receivers for cell-to-cell communication were found.

In contrast to electromagnetic wave-based mobile communication, there are no general-purpose processors available to implement receivers for MC on the scale of individual cells. Yet, in nature, cells do communicate with each other using signaling molecules. Often, cells have ligand-binding receptors on the surface and use chemical reaction networks (CRNs), i.e., networks of interacting chemical species, to perform computations. For example, bacteria use CRNs to determine whether to tumble or to move forward in the context of chemotaxis. Also, receptor states can cause long-term changes by interacting with the cells' gene-regulatory networks (GRNs), i.e., by influencing which genes are expressed depending on the sensed environmental conditions.

Therefore, it is a natural choice to formulate signal detection problems in such a way that CRNs and GRNs can be used to perform the computations required for artificial MC.

In [4, 5, 6], CRNs were designed manually for a given channel to implement receiver components. However, these approaches either do not account for the stochastic fluctuations of CRNs, require many molecules or are suboptimal compared to a maximum-a-posteriori (MAP) detector.

The authors of [7] used CRNs to implement feed-forward neural networks. This idea is especially interesting for MC if analytical channel models are not available or their parameters are unknown. However, the resulting CRNs are very complex and the impact of stochastic fluctuations has not been considered in [7].

The use of GRNs to perform computations in general has been reviewed in [8] and proposed for receiver design in [9]. While this approach is especially interesting because it can exploit natural mechanisms already present in cells and it has been shown that logical functions can be implemented via GRNs, very long time scales are needed to perform complex computations. A recent review highlighted that using transcriptional elements is only feasible for very simple receiver implementations in the near future [10]. Finally, several concepts for implementing molecular machine learning based on CRNs, bacterial multi-species communication, and Calcium signaling were proposed in [11]. Potentially, these approaches could be also applied for receiver design. However, it is

unclear how complex the resulting architectures would be and how long the computations would take.

A common issue for all previously mentioned approaches is the diversity of transmission channels. Each channel involves different parameters and possibly time-variant parameter values. Yet, most model-based approaches for receiver design in the literature are not concerned with estimating correct parameters and adapting to changing channel conditions. While the existing learning-based approaches could in principle resolve this issue, it remains unclear whether their training would be fast enough to adapt to changing conditions.

In this paper, we consider the practical implementation of MAP detectors for cellular receivers with ligand-binding receptors using CRNs. While the stochastic fluctuations of molecule counts are usually considered as noise source in CRNs, our designs exploit this randomness to perform computations. To emphasize this, we call the proposed CRNs "stochastic".

For the first CRN, we start by approximating the joint distribution of the transmitted symbol and the receptor states using a Boltzmann machine (BM). BMs are commonly used graphical models in machine learning to approximate arbitrary joint probability distributions of binary random variables (RVs) [12]. Then, we exploit the methods reported in [13] to represent BMs in stochastic CRNs. In order to obtain posterior estimates for the transmitted symbol, i.e., to perform detection, we then condition the BM, or rather its CRN representation, on the receptor states.

The second CRN proposed in this paper is not derived from BMs, but is designed directly for our purposes. It has a lower complexity and can adapt to changing channel conditions over time by learning efficiently from pilot symbols. We demonstrate the performance of the proposed approach for a MC channel subject to time-varying background noise levels.

The remainder of this paper is organized as follows. In Section 2, we review the fundamentals of BMs, CRNs, and how BMs can be represented in CRNs. In Section 3, we introduce the considered system model and the corresponding MAP detector. In Sections 4 and 5, we introduce the proposed receiver designs which are then evaluated and compared to the optimal MAP detector in Section 6. Finally, we summarize our main findings and potential future work in Section 7.

2 BOLTZMANN MACHINES AND CHEMICAL REACTION NETWORKS

2.1 Boltzmann Machines

In order to approximate a joint distribution of N binary RVs with probability mass function $q_{\mathbf{Z}}(\mathbf{z})$, a BM with at least $M \geq N$ nodes is required. N of these nodes are each identified with a binary RV of the original distribution. The additional $M - N$ nodes are "hidden", i.e., they do not correspond to a specific RV but simply allow for additional degrees of freedom in the BM in order to achieve a higher approximation quality. However, as we will show in Section 4, no hidden nodes are required for the eventual task of MAP detection. Therefore, we restrict ourselves to the simpler case of a BM without hidden nodes, which we call fully visible Boltzmann machine (FVBM).

A FVBM consists of N nodes, each of which corresponds to a

binary RV $Z_i, i \in \{1, \dots, N\}$, which we collect in a vector $\mathbf{Z} = [Z_1, \dots, Z_N]^T$, where $[\cdot]^T$ denotes the transpose operator. The probability mass function of the FVBM nodes is then given by

$$p_{\mathbf{Z}}(\mathbf{z}) = \frac{1}{\mathcal{Z}} \exp\left(\frac{1}{2} \mathbf{z}^T \mathbf{W} \mathbf{z} + \mathbf{z}^T \boldsymbol{\theta}\right). \quad (1)$$

Here, $\boldsymbol{\theta} \in \mathbb{R}^{N \times 1}$ and $\mathbf{W} \in \mathbb{R}^{N \times N}$ denote respectively the vector of biases and a symmetric weight matrix with all-zero diagonal, which captures the correlations between the nodes. \mathcal{Z} is a normalization constant that ensures that $p_{\mathbf{Z}}(\mathbf{z})$ is a probability distribution.

The state of node i depends on its associated bias θ_i and the current state of all other nodes $j \neq i$ with non-zero correlations to node i , captured by weight matrix entries $W_{i,j}$. Formally, the probability that $Z_i = 1$ is given by

$$p_{Z_i | \mathbf{Z}_{-i}}(Z_i = 1 | \mathbf{Z}_{-i} = \mathbf{z}_{-i}) = \sigma\left(\theta_i + \sum_{j \neq i} W_{i,j} z_j\right), \quad (2)$$

where \mathbf{Z}_{-i} is the vector of all RVs except the one associated with node i , \mathbf{z}_{-i} contains their observed realization, and $\sigma(x) = \frac{1}{1+e^{-x}}$. Eq. (2) is for example well known from the Gibbs sampling algorithm [12].


When using a BM to approximate $q_{\mathbf{Z}}(\mathbf{z})$, \mathbf{W} and $\boldsymbol{\theta}$ can be learned from the first- and second-order moments $\mathbb{E}_q\{\mathbf{z}\}$ and $\mathbb{E}_q\{\mathbf{z}\mathbf{z}^T\}$, respectively, where $\mathbb{E}_q\{\cdot\}$ denotes the expectation operator for a probability distribution $q_{\mathbf{Z}}(\mathbf{z})$. For an introduction to learning weights and biases for BMs, we refer to [12].

2.2 Chemical Reaction Networks

As previously mentioned, BMs can be realized as CRNs [13]. Formally, a CRN $C = (\mathcal{S}, \mathcal{R}, \mathbf{k})$ consists of a set of species \mathcal{S} , a set of reactions \mathcal{R} defined over \mathcal{S} , and a vector of reaction rate constants \mathbf{k} .

A chemical reaction $r \in \mathcal{R}$ converts molecules into other molecules. For mass-action kinetics, the propensity of reaction r , i.e., how likely it occurs per unit time, is proportional to the reaction rate constant $k_r > 0$ and the number of available reactants in the considered reaction network.

Consider for example the reaction



Here, the propensity of reaction r is given by the product of the reaction rate constant, the number of W^{ON} molecules, and the number of X^{ON} molecules. Thus, if there is no X^{ON} or no W^{ON} molecule, the reaction cannot happen at all.

In the remainder of this paper, we often write that a molecule can be in one of two states, namely ON or OFF. This is to represent binary states, e.g., of a RV, by molecules. Chemically speaking, the ON-version of a molecule might for example contain a phosphor group that is not contained in the complementary OFF species. Otherwise, the molecules are identical. We also say that the ON species is "active" whereas the OFF species is "inactive". For example, reaction r in (3) deactivates the X molecule, i.e., it converts X^{ON} to X^{OFF} .

To implement a BM using a CRN, we exploit the fact that a CRN can be described by a continuous-time Markov Chain (CTMC) [13]. We assume a CRN with species set $\{Z_1^{\text{ON}}, Z_1^{\text{OFF}}, \dots, Z_N^{\text{ON}}, Z_N^{\text{OFF}}\}$ and demand that there is exactly either one Z_i^{ON} or one Z_i^{OFF} molecule

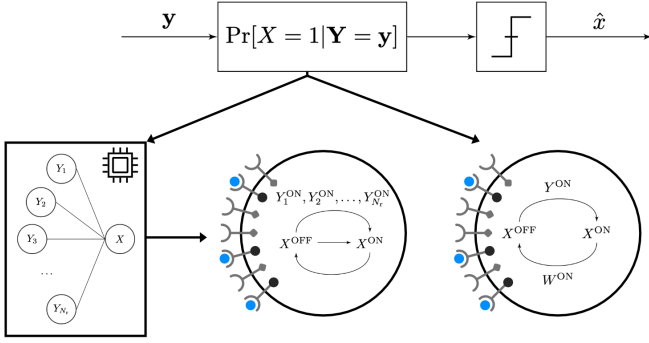


Figure 1: Proposed receiver model. The probabilistic model can be implemented in silico using a BM (bottom left). Its in vivo implementation could be a CRN based on the BM (bottom center) or a directly designed CRN (bottom right).

for $\forall i$ at any point in time. Then we can say that the molecule Z_i is ON if the i -th node of the corresponding BM has the value 1 whereas the Z_i molecule is OFF if the node has value 0 at a given point in time. By defining appropriate reactions that activate or deactivate molecules, one can ensure that the CTMC describing the CRN eventually reaches a stationary distribution and that this stationary distribution is the same as the one of the corresponding BM. In this case, the CRN is said to implement the BM.

This idea has been introduced and formalised in [13]. In fact, in [13] three representations of BMs using CRNs are provided. Two of them are exact but require a number of reactions $|\mathcal{R}|$ that scales exponentially with the number of nodes of the BM, where $|\mathcal{S}|$ denotes the cardinality of a set \mathcal{S} . The third one is an approximation of the BM, which requires much fewer reactions at the cost of a mismatch between the stationary distributions of the CRN and the BM.

3 SYSTEM MODEL

As our focus is the detection process, we keep the modulation scheme and channel model simple. Specifically, we assume the transmission of a binary source symbol $x \in \{0, 1\}$ via binary concentration shift keying (BCSK) and the absence of inter-symbol interference (ISI). The latter holds for sufficiently long symbol intervals or enzymatic ligand degradation [14]. Considering that existing receiver implementations may require extremely long decoding times of up to several hours per symbol [9, 10], choosing long intervals to avoid ISI would not be a performance bottleneck compared to existing approaches.

A cellular receiver, as depicted in Fig. 1, senses its environment through N_r cell surface receptors. Depending on whether receptor i is bound to a ligand, its intra-cellular domain might have different chemical properties. Thus, we represent receptor i 's intra-cellular domain by different chemical species depending on its state, namely by Y_i^{ON} if the receptor is bound to a ligand and by Y_i^{OFF} otherwise. For our purposes, we assume that the receptor states are sampled once and then stored in a vector $\mathbf{y} \in \{0, 1\}^{N_r}$ where $y_i = 1$ if we observe Y_i^{ON} and $y_i = 0$ for Y_i^{OFF} . Mathematically, the considered MC channel is characterized by the joint distribution $q_{\mathbf{Y}, X}(\mathbf{y}, x)$ where

X denotes the binary RV corresponding to the transmitted symbol and \mathbf{Y} contains the binary RVs Y_i corresponding to the different receptors.

A MAP detector for BCSK would simply compute the estimated symbol \hat{x} as

$$\hat{x}^{\text{MAP}} = \begin{cases} 1 & , \text{ if } \Pr[X = 1 | \mathbf{Y} = \mathbf{y}] \geq \frac{1}{2} \\ 0 & , \text{ otherwise} \end{cases} \quad (4)$$

In the remainder of this paper, we assume that all receptors have identical chemical properties, such that $\Pr[X = 1 | \mathbf{Y} = \mathbf{y}]$ depends only on the number of bound receptors $N_{r,b} = \sum_{i=0}^{N_r} y_i$. We further assume that the likelihoods $l_x(n) = \Pr[N_{r,b} = n | X = x]$ have exactly one local maximum $n_{m,x}$ and are monotonically increasing for $n \leq n_{m,x}$ and monotonically decreasing for $n \geq n_{m,x}$. For these quite general assumptions, the maximum-likelihood (ML) detector is a simple threshold detector, i.e.,

$$\hat{x}^{\text{ML}} = \begin{cases} 1 & , \text{ if } N_{r,b} \geq \nu \\ 0 & , \text{ otherwise} \end{cases} \quad (5)$$

for some threshold $\nu \in \mathbb{N}_0$. Moreover, we assume equiprobable symbols x , such that the ML and the MAP detector coincide.

Now, we observe that it is actually not necessary to know $\Pr[X = x | \mathbf{Y} = \mathbf{y}]$ perfectly to make MAP decisions. Instead, it is sufficient if we find a function $f_1(N_{r,b}) \in [0, 1]$ that fulfills $f_1(N_{r,b}) \geq \frac{1}{2}$ if and only if $\Pr[X = 1 | N_{r,b}] \geq \frac{1}{2}$. We say that such a function has the *MAP property*. Thus, even if a probabilistic model does not output the exact values of the posterior distribution, it can still serve as MAP detector as long as it has the *MAP property*.

4 BOLTZMANN MACHINE-INSPIRED MAP DETECTORS

4.1 MAP-Capability of Boltzmann Machines

In order to use a BM for BCSK detection, we set $\mathbf{Z} = [X \ \mathbf{Y}]^T$.

THEOREM 1. *For a known threshold ν and properly chosen parameters, BMs have the MAP property.*

PROOF. Setting $Z_i = X$, $Z_{-i} = \mathbf{Y}$ in (2) and using equal weights $w_{x,y}$ for all $w_{i,j}$, $j \neq i$, yields

$$p_{X|Y}(X = 1 | \mathbf{Y} = \mathbf{y}) = \sigma \left(\theta_i + \sum_{i=1}^{N_r} w_{x,y} y_i \right) \quad (6)$$

$$= \sigma \left(\theta_i + N_{r,b} w_{x,y} \right). \quad (7)$$

By choosing $\theta_i = -(\nu - \frac{1}{2})w_{x,y}$, we ensure $p_{X|Y}(X = 1 | \mathbf{Y} = \mathbf{y}) > 0.5$ if $N_{r,b} \geq \nu$ and $p_{X|Y}(X = 1 | \mathbf{Y} = \mathbf{y}) < 0.5$ if $N_{r,b} < \nu$ and thus can realize a MAP detector for a known optimal threshold ν . \square


4.2 Representation via Chemical Reaction Networks

Formally, it would be possible to define a CRN C_{EM} based on the *Edge Species Mapping* proposed in [13] that would implement a trainable BM. However, even for a few dozen receptors, this is not feasible in practice because $(N_r + 1) \cdot 2^{N_r+1}$ reactions, each with up to $2N_r + 1$ reactants, would be required [13].

Therefore, we resort to the *Taylor Mapping* from [13] instead. The

resulting CRN C_{TM} approximates a given BM by $2N_r^2 + 4N_r + 2$ reactions, each with at most two reactants. However, to perform inference, only a subset of these reactions is required, namely only those that activate or deactivate the molecule identified with the X -node in the BM.

The resulting CRN $C_{TM,X}$ then requires only the reactions



Here, θ_x is the bias associated with the X -node, $W_{y_i,x}$ is the weight between the Y_i - and the X -node of the BM, and k is an arbitrary scaling factor for the reaction rate constants. X^{ON} and X^{OFF} are the species associated with node X . At any point in time, there is either a single X^{ON} molecule or a single X^{OFF} molecule. In the absence of any ligands all Y_i are OFF and X^{ON} switches to X^{OFF} more likely than vice versa due to the bias θ_x . On the other hand, the more receptors are bound to a ligand, the larger the tendency to switch from X^{OFF} to X^{ON} .

THEOREM 2. $C_{TM,X}$ preserves the MAP property of the BM.

PROOF. First, we set $\theta_x = -(\nu - \frac{1}{2})w_{x,y}$ and $W_{y_i,x} = w_{x,y}$ according Theorem 1 and assume $N_{r,b}$ to be constant. Then, we want to compute the steady-state-probability to observe X^{ON} at any given point in time. Therefore, we first sum up all the rate constants for all possible reactions converting X^{OFF} to X^{ON} in (8). This yields $k(1 + N_{r,b}w_{x,y})$. On the other hand, X^{ON} is converted to X^{OFF} with rate constant $k(1 + |\theta_x|)$.

Now, we can use the detailed balanced approach, i.e., we set

$$\Pr[X^{\text{ON}}|N_{r,b}] \cdot k(1 + |\theta_x|) = \Pr[X^{\text{OFF}}|N_{r,b}] \cdot k(1 + N_{r,b}w_{x,y}), \quad (9)$$

where $\Pr[X^{\text{ON}}|N_{r,b}]$ is the probability to observe X^{ON} at any given point in time if $N_{r,b}$ receptors are bound. Using $\Pr[X^{\text{OFF}}|N_{r,b}] = 1 - \Pr[X^{\text{ON}}|N_{r,b}]$, we find $\Pr[X^{\text{ON}}|N_{r,b}]$ to be

$$\Pr[X^{\text{ON}}|N_{r,b}] = \frac{1 + N_{r,b}w_{x,y}}{2 + N_{r,b}w_{x,y} + |\theta_x|}. \quad (10)$$

Solving $\Pr[X^{\text{ON}}|N_{r,b}] \geq \frac{1}{2}$ for $N_{r,b}w_{x,y}$ yields

$$N_{r,b}w_{x,y} \geq |\theta_x|. \quad (11)$$

Since $\theta_x = -(\nu - \frac{1}{2})w_{x,y}$, $\Pr[X^{\text{ON}}|N_{r,b}] \geq 0.5$ holds if and only if $N_{r,b} \geq \nu$ receptors are bound. \square

Finally, it should be noted that $\Pr[X^{\text{ON}}|N_{r,b}]$ can be approximated by a time-average

$$\Pr[X^{\text{ON}}|N_{r,b}] \approx \frac{1}{t_{\text{obs}}} \int_0^{t_{\text{obs}}} \mathbb{1}_t\{X^{\text{ON}}\} dt, \quad (12)$$


where $\mathbb{1}_t\{\cdot\}$ is 1 if X^{ON} is present at time t and 0 otherwise, and t_{obs} is the duration of the observation interval. As t_{obs} increases, the time average becomes an increasingly exact estimate for $\Pr[X^{\text{ON}}|Y = y]$.

5 TRAINABLE LOW-COMPLEXITY MAP DETECTORS

5.1 A MAP-Capable CRN


For the architecture introduced in the previous section, the number of reactions scales linearly with the number of receptors. This is due to the fact that the mappings proposed in [13] do not exploit the fact that some weights are identical. Furthermore, there is no known way to train the CRNs from Section 4.

Therefore, we propose a low-complexity CRN C_{LC} that embraces identical receptor weights and can be trained with pilot symbols. As before, we assume N_r identical receptors that activate the decision molecule if bound to a ligand. Formally, this is given by



with reaction rate constant k_{on} .

However, in contrast to the BM-inspired CRN, we do not define an explicit reaction to switch from X^{ON} to X^{OFF} due to a fixed bias θ . Instead, we assume that there are N_w^{ON} active weight molecules W^{ON} present in the receiver cell. These can deactivate X^{ON} with reaction rate constant k_{off} , i.e.,



THEOREM 3. For an appropriate value of N_w^{ON} , C_{LC} , described by (13) and (14), has the MAP property.

PROOF. We take the same detailed balance-based approach to compute the steady state probability of observing X^{ON} as in the proof of Theorem 2. However, now the rate with which X^{OFF} is converted to X^{ON} is given by $N_{r,b} \cdot k_{\text{on}}$. X^{ON} is converted to X^{OFF} with rate $N_w^{\text{ON}} \cdot k_{\text{off}}$.

For these rates, we obtain

$$\Pr[X^{\text{ON}}|N_{r,b}, N_w^{\text{ON}}] = \frac{N_{r,b}}{N_{r,b} + \frac{k_{\text{off}}}{k_{\text{on}}} N_w^{\text{ON}}}. \quad (15)$$

Solving $\Pr[X^{\text{ON}}|N_{r,b}, N_w^{\text{ON}}] \geq 0.5$ for $N_{r,b}$ yields $N_{r,b} \geq \frac{k_{\text{off}}}{k_{\text{on}}} N_w^{\text{ON}}$.

For example, when setting $\frac{k_{\text{off}}}{k_{\text{on}}} = 1$, MAP performance is achieved for $N_w^{\text{ON}} = \nu$. \square

Practically, one can estimate $\Pr[X^{\text{ON}}|N_{r,b}, N_w^{\text{ON}}]$ using a time-average like in (12).


5.2 Pilot Symbol-based Learning Rule

In order to learn the optimal value of N_w^{ON} online, we use a simple learning rule based on pilot symbols. For each pilot symbol, the detector output \hat{x} is compared to the actually transmitted symbol x_{pilot} . If an estimate is wrong, N_w^{ON} is slightly adapted such that we eventually converge to the value of N_w^{ON} that would achieve MAP performance.


Formally, the number of active weight molecules after the k -th pilot symbol is given by

$$N_w^{\text{ON}}[k+1] = \begin{cases} N_w^{\text{ON}}[k] & , \text{ if } \hat{x} = x_{\text{pilot}} \\ N_w^{\text{ON}}[k] + 1 & , \text{ if } \hat{x} = 1 \text{ and } x_{\text{pilot}} = 0 \\ N_w^{\text{ON}}[k] - 1 & , \text{ if } \hat{x} = 0 \text{ and } x_{\text{pilot}} = 1 \end{cases} \quad (16)$$

To implement this chemically, we introduce a reservoir of N_w^{OFF} inactive weight molecules W^{OFF} that can be converted into W^{ON} molecules. At any point in time, $N_w^{\text{ON}} + N_w^{\text{OFF}} = N_w$, where $N_w \in \mathbb{N}$ remains constant over time. Clearly, one should choose N_w large enough such that C_{LC} has the MAP property for all possible values of v . Because $v \leq N_r$, one could simply choose $N_w \geq \left\lceil \frac{k_{\text{off}}}{k_{\text{on}}} N_r \right\rceil$. Naively, one could try to implement (16) chemically by the following reactions



where $k_{u,1}$ and $k_{u,2}$ are the respective reaction rate constants. We assume that a $X_{\text{pilot}}^{\text{OFF}}$ molecule is present in the receiver if a 0 has been transmitted and a $X_{\text{pilot}}^{\text{ON}}$ molecule if a 1 has been transmitted. The estimated symbol \hat{x} is encoded in the presence of \hat{X}^{ON} and \hat{X}^{OFF} for $\hat{x} = 1$ and $\hat{x} = 0$, respectively. Thus, if both $X_{\text{pilot}}^{\text{OFF}}$ and \hat{X}^{ON} are present, this indicates that $\hat{x} = 1$ whereas $x_{\text{pilot}} = 0$ and thus N_w^{ON} should be increased, i.e. a W^{OFF} molecule should be activated. However, for mass-action kinetics, the propensities, and thus how often each reaction would occur in a given time frame, actually depend on the number of W^{OFF} and W^{ON} molecules. Therefore, (17) actually does not implement (16) exactly. A more accurate chemical implementation of (16) is instead given by



Here, we introduce an additional H molecule that is consumed by the reaction. If we ensure that only a single H molecule is present at the beginning of each pilot symbol interval, N_w^{ON} cannot change by more than 1. By further specifying that $k_{u,1}$ and $k_{u,2}$ are chosen such that both reactions occur with probability close to 1 for all possible values of N_w^{ON} , N_w^{ON} changes then exactly by 1. Therefore the reactions in (18) implement (16).

6 PERFORMANCE EVALUATION

6.1 Evaluation Setup

We evaluate our proposed detectors using a similar channel model as in [4]. Namely, the ligand concentration for a transmitted symbol x around the receiver is given by

$$c_x = c^n + \Delta_c \cdot x, \quad (19)$$

where c^n and Δ_c denote the concentration due to the expected background noise and the concentration increase due to the release of molecules if $x = 1$ is transmitted, respectively.

The receiver cell uses $N_r \in \{30, 50\}$ independent and identically distributed (i.i.d.) receptors to estimate the ligand concentration and thus the transmitted symbol x . From [4], we obtain the binding probability of receptor i for x as

$$\Pr[Y_i = 1|X = x] = \frac{c_x}{c_x + \frac{k_-}{k_+}}. \quad (20)$$

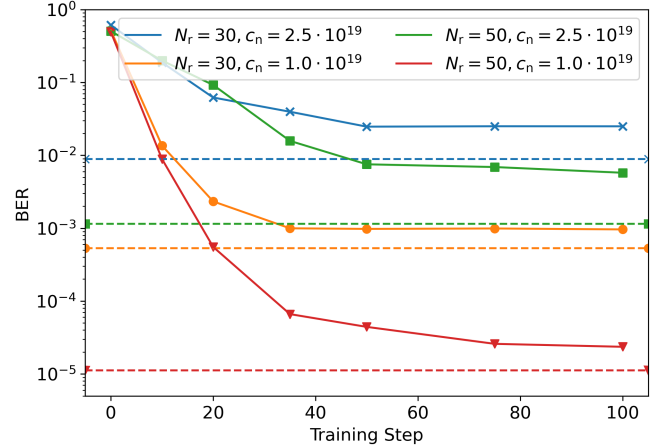


Figure 2: Average bit error rates (BERs) as a function of the training step for the four considered scenarios (solid curves). After a fast initial decay, the BERs approach the optimal MAP performance (dashed curves) slowly.

Here, k_+ and k_- are the binding and unbinding rate constants between ligand and receptor, respectively.

For equiprobable symbols, the joint distribution $q_{\mathbf{Z}}(\mathbf{z}) = q_{Y,X}(\mathbf{y}, x)$ is then given by

$$q_{Y,X}(\mathbf{y}, x) = \frac{1}{2} \prod_{i=1}^{N_r} \Pr[Y_i = y_i|X = x]. \quad (21)$$

We choose $k_+ = 2 \cdot 10^{-19} \text{m}^3 \text{s}^{-1}$ and $k_- = 20 \text{s}^{-1}$ as in [4]. We also consider the system model for diffusive ligand propagation and instantaneous molecule release from [4], given by

$$c(\tau) = \frac{\gamma}{(4\pi D\tau)^{3/2}} \exp\left(-\frac{d^2}{4D\tau}\right). \quad (22)$$

Here, d , D , γ and τ denote the distance between receiver and transmitter, the diffusion coefficient, the number of released molecules, and the time since the molecule release, respectively. For $\gamma = 10^3$ released molecules, a diffusion coefficient $D = 10^{-10} \frac{\text{m}^2}{\text{s}}$ and a distance of $0.75 \mu\text{m}$, $c(t)$ has a peak value in the order of $10^{20} \frac{\text{molecules}}{\text{m}^3}$. Thus, we choose $c_1 = 1.5 \cdot 10^{20} \frac{\text{molecules}}{\text{m}^3}$. For the noise levels, we assume two scenarios. The first one with $c^n = 2.5 \cdot 10^{19} \frac{\text{molecules}}{\text{m}^3}$ and the second one with $c^n = 1.0 \cdot 10^{19} \frac{\text{molecules}}{\text{m}^3}$.

6.2 Training Boltzmann Machines

We train the BMs with estimated expectations $\mathbb{E}_{q_{Y,X}}\{\mathbf{z}\}$ and $\mathbb{E}_{q_{Y,X}}\{\mathbf{z}\mathbf{z}^T\}$ based on 10^4 samples from our system model defined by (21).

In each training step, we also need to estimate the first- and second-order moments of the BM for the current biases and weights. Therefore, we generate 10^4 Gibbs samples using the Gibbs sampling algorithm [12]. This algorithm exploits that (2) can be easily evaluated and returns samples distributed according to the Boltzmann distribution $p_{X,Y}(x, y)$. From the obtained Gibbs samples, we then compute the expectations $\mathbb{E}_{p_{X,Y}}\{\mathbf{z}\}$ and $\mathbb{E}_{p_{X,Y}}\{\mathbf{z}\mathbf{z}^T\}$.

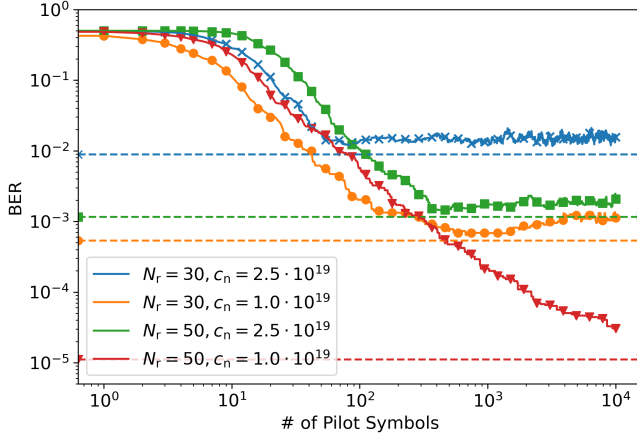


Figure 3: Average BERs over time for the four considered scenarios (solid curves). All systems eventually approach MAP performance (dashed curves).

Similar to deep learning [15], we use a step-wise reduced learning rate $\eta[k]$. In training steps $k \in \{0, \dots, 19\}$, $\eta[k] = 1.0$, for $k \in \{20, \dots, 49\}$ $\eta[k] = 0.5$ and for $k \in \{50, \dots, 99\}$, we use $\eta[k] = 0.1$.

Initially, we define the matrix $\mathbf{W}' = \frac{1}{2}(\mathbf{V} + \mathbf{V}^T) \in \mathbb{R}^{(N_r+1) \times (N_r+1)}$, where the entries of a random matrix \mathbf{V} are i.i.d. normally distributed with zero mean and variance $\frac{1}{N_r+1}$. We then obtain the initial weight matrix \mathbf{W}_0 from \mathbf{W}' by setting the diagonal entries to zero. To reduce the number of parameters, we also set all entries that capture only correlations among receptors to zero, as well, and we also do not update them during training.

We train 20 BMs for each scenario using this approach. In Fig. 2, the obtained BERs are shown as a function of the training step. For each BM, we compute the BER by comparing \hat{x} to x until 100 errors are made. Then, we average the BERs of all BMs. For reference, we also show the BER obtained from the corresponding MAP detectors (dashed lines).

Fig. 2 confirms that BMs are indeed able to learn close-to-MAP performance from training data. If the expectations were estimated from more samples, the detectors would come even closer to MAP performance for more training steps.

6.3 Convergence of CRN Online Learning

We investigate the convergence over time for our online learning system. Therefore, we use (16) as the learning rule with $N_w^{\text{ON}} = 0$ as initial value. We compute \hat{x} as

$$\hat{x} = \begin{cases} 1 & , \text{ if } \Pr[X^{\text{ON}} | N_{r,b}, N_w^{\text{ON}}] \geq \frac{1}{2} \\ 0 & , \text{ otherwise} \end{cases} \quad (23)$$

by computing $\Pr[X^{\text{ON}} | N_{r,b}, N_w^{\text{ON}}]$ analytically using (15) for $k_{\text{on}} = k_{\text{off}} = 1$. This corresponds to $t_{\text{obs}} \rightarrow \infty$.

After every pilot symbol, we compute the BER that would be obtained with the current value of N_w^{ON} . We repeat this procedure 20 times for each scenario and then average the obtained BERs. This way, we can observe how many pilot symbols are necessary to

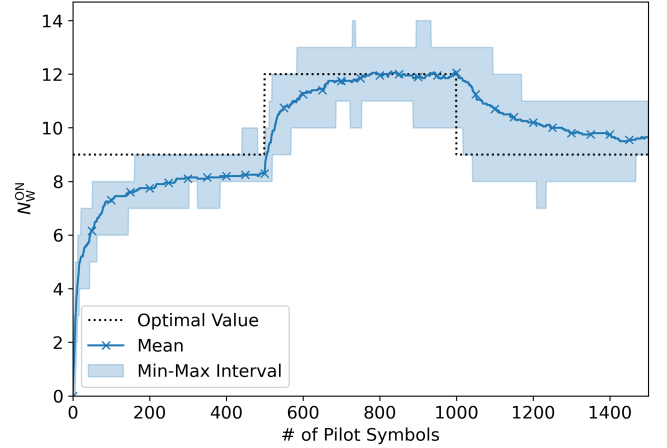


Figure 4: Average number of N_w^{ON} over time (solid line) and interval with minimum and maximum value of N_w^{ON} (shaded area) for time-variant background noise levels. For reference, the optimal values are shown by the dotted black line.

achieve a certain performance. The resulting BERs are shown in Fig. 3. Clearly, all proposed detectors eventually achieve close-to-MAP performance.

To understand the BER curves better, we first observe that errors can be due to a mismatch between N_w^{ON} and the MAP threshold ν , or due to the intrinsic randomness due to the stochastic binding process of the receptors. The BER due to the mismatch between N_w^{ON} and ν is lower if the mismatch is lower. Therefore, as N_w^{ON} is adapted over time to come closer to the MAP value, the BER goes down, and consequently the adaption of N_w^{ON} slows down. Once the MAP threshold is reached, there are no mismatch-errors. However, errors can still be introduced due to the intrinsic channel randomness. Then, N_w^{ON} is updated and a mismatch is introduced again, that will be corrected once a new mismatch error occurs. Therefore, only close-to-MAP performance can be expected with this update rule. For the yellow curve, one can observe a temporary BER minimum after about 10^3 pilot symbols. At this point, there are no mismatch errors anymore and randomness-caused errors have not yet occurred. Thus, the BER is temporarily lower than in the steady state.

6.4 Time-Variant Background Noise Levels

In real-world applications, channel parameters might be time-variant, e.g., background noise levels could depend on the activity of other users in the channel. To explore the impact of this on the performance of our online-learning rule, we consider a receiver with $N_r = 30$ receptors and assume $\Delta_c = (1.5 \cdot 10^{20} - 1.0 \cdot 10^{19}) \frac{\text{molecules}}{\text{m}^3}$. Initially, we set $c_1^n = 1.0 \cdot 10^{19} \frac{\text{molecules}}{\text{m}^3}$ as noise level, change it after 500 pilot symbols to $c_2^n = 2.5 \cdot 10^{19} \frac{\text{molecules}}{\text{m}^3}$ and after 500 more pilot symbols back to the original value.

The resulting average number of active weight molecules is shown in Fig. 4 over time. The shaded area shows the lowest and highest

values of N_w^{ON} for our 20 model runs. Clearly, N_w^{ON} follows the optimal values relatively closely over time. Still, the current learning rule requires a considerable amount of pilot symbols which indicates that the proposed detector is best suited for slowly changing channels and relatively high BERs.

7 CONCLUSION

In this paper, we introduced two CRNs that can be used to realize MAP detection for appropriately chosen reaction rate constants and molecule counts. In contrast to existing approaches, both detectors exploit the intrinsic fluctuations of chemical reactions.

The first detector exploits that BMs can be represented using CRNs. It can be trained offline using simulations, for example when no analytical channel model exists. We showed for a simple system model that even with relatively few training data close-to-MAP performance can be achieved.

The second detector is based on a manually designed CRN that can be trained with pilot symbols even after deployment in a cellular receiver. Therefore, this detector can adapt to changes in the channel. We showed that this detector achieves close-to-MAP performance, too.

The combination of low molecule counts, trainability, and their embrace of stochastic fluctuations make our low-complexity detectors promising candidates on the way towards practical receiver implementations for intra-body communication. In future work, we thus plan to extend them to M-ary transmission schemes and improve the training process further.

REFERENCES

- [1] W. Haselmayr et al. 2019. Integration of molecular communications into future generation wireless networks. In *Proc. 6G Wireless Summit*. Levi, Finland.
- [2] N. Farsad, H. B. Yilmaz, A. Eckford, C.-B. Chae, and W. Guo. 2016. A comprehensive survey of recent advancements in molecular communication. *IEEE Commun. Surv. & Tuts.*, 18, 3, 1887–1919.
- [3] S. Lotter et al. 2023. Experimental research in synthetic molecular communications – part I: Overview and short-range systems. *IEEE Nanotechnol. Mag.*, (Apr. 2023).
- [4] M. Kuscu and O. B. Akan. 2018. Maximum likelihood detection with ligand receptors for diffusion-based molecular communications in Internet of Bio-nano Things. *IEEE Trans. Nanobiosci.*, 17, 1, (Jan. 2018), 44–54.
- [5] C. T. Chou. 2019. Designing molecular circuits for approximate maximum a posteriori demodulation of concentration modulated signals. *IEEE Trans. Commun.*, 67, 8, (Aug. 2019), 5458–5473.
- [6] M. Egan, T. Q. Duong, and M. Di Renzo. 2019. Biological circuits for detection in MoSK-based molecular communication. *IEEE Access*, 7, (Feb. 2019), 21094–21102.
- [7] D. F. Anderson, Joshi Badal, and Deshpande Abhishek. 2021. On reaction network implementations of neural networks. *J. R. Soc. Interface.*, 18, 177, (Apr. 2021).
- [8] B. Saltepe et al. 2018. Cellular biosensors with engineered genetic circuits. *ACS Sensors*, 3, 1, (Jan. 2018), 13–26.
- [9] B. D. Unluturk, A. O. Bicen, and I. F. Akyildiz. 2015. Genetically engineered bacteria-based biotransceivers for molecular communication. *IEEE Trans. Commun.*, 63, 4, (Apr. 2015), 1271–1281.
- [10] M. Femminella and G. Reali. 2022. Implementation issues of diffusion-based molecular communications receivers based on transcriptional elements. *Digit. Signal Process.*, 124, (May 2022), 103160.
- [11] S. Balasubramaniam et al. 2022. Realizing molecular machine learning through communications for biological AI: Future directions and challenges. eprint: arXiv:2212.11910.
- [12] D. J. C. MacKay. 2002. *Information Theory, Inference & Learning Algorithms*. Cambridge Univ. Press, Cambridge, U.K.
- [13] W. Poole et al. 2017. Chemical boltzmann machines. In *Proc. Int. Conf. DNA-Based Comput.* Austin, TX, USA.
- [14] A. Noel, K. C. Cheung, and R. Schober. 2014. Improving receiver performance of diffusive molecular communication with enzymes. *IEEE Trans. Nanobiosci.*, 13, 1, (Mar. 2014), 31–43.
- [15] I. Goodfellow, Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press, Cambridge, MA, USA.