SSM

Università di Napoli Federico II – Scuola Superiore Meridionale

Ph.D. program in Genomic and Experimental Medicine (GEM)

ENGINEERING GENE EXPRESSION IN MAMMALIAN CELLS WITH DEEP LEARNING MODELS

A dissertation presented by

ADELE DE HOFFER

Matricola n.

DR995935

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Genomic and Experimental Medicine (GEM)- 37º Cycle

Supervisor:
Prof. Diego di Bernardo, PhD,
Full Professor in BioEngineering
Department of Chemical, Materials and Industrial Engineering
Università di Napoli Federico II
Principal Investigator- Telethon Institute of Genetics and Medicine (TIGEM)

GEM program coordinator:
Prof. Brunella Franco, MD.,
Professor of Medical Genetics
Department of Translational Medicine
Università di Napoli Federico II
Principal Investigator- Telethon Institute of Genetics and Medicine (TIGEM)

To all the amazing people
who walk with me
through this beautiful chaos called life.

Thank you for being
my greatest strength and deepest inspiration.

Abstract

Gene expression regulation in mammalian cells is a complex process that can be modulated at multiple levels. This complexity makes it challenging to achieve precise spatiotemporal control, requiring both dynamic regulatory systems and cell-type–specific regulatory elements. This thesis explores two deep learning–based approaches to controlling gene expression in synthetic biology contexts. The first project develops a data-driven model predictive controller (MPC) that employs a neural network to model and dynamically regulate the expression of an exogenous gene controlled by a synthetic, inducible Tet-Off promoter. This approach enables precise temporal modulation of gene expression in response to external stimuli. This part is subject to embargo. While the two projects focus on distinct aspects of gene regulation—dynamic temporal control versus sequence-driven cellular specificity—both demonstrate the utility of neural networks as versatile tools for modeling and controlling gene expression. Together, these studies advance synthetic biology by providing data-driven strategies for both temporal modulation and cell-type–specific design of gene regulatory elements.

Contents

Chapter 1

Aim and scope of the thesis

1.1 Gene Expression Overview

Gene expression is the process by which the information stored in DNA is converted into functional elements, such as proteins and RNA molecules, that carry out essential cellular functions [1]. The regulatory architecture governing gene expression spans multiple hierarchical levels, from chromatin organization to protein modification [2]. At the chromatin level, histone modifications and chromatin remodeling determine the accessibility of DNA to the transcriptional machinery. Transcriptional regulation is primarily mediated by promoter and enhancer elements that recruit transcription factors and RNA polymerase complexes, thereby controlling the rate of mRNA synthesis. Once transcribed, mRNA molecules undergo post-transcriptional regulation through alternative splicing, microRNA targeting, and RNA-binding proteins, which modulate their processing, stability, and localization. Finally, post-translational modifications such as phosphorylation and ubiquitination fine-tune protein activity, localization, and stability. Together (Fig 1.1), these interconnected regulatory layers define the precise control of gene expression and offer multiple points of intervention for designing controllable gene expression systems in biotechnological applications.

1.1.1 Biotechnological Applications

The precise control of gene expression is fundamental to cellular function and organismal development. For example in regenerative medicine, precise temporal control of gene expression enables the directed differentiation of stem cells into specific cell types [3]. For instance, coordinated expression of transcription factors such as OCT4, SOX2, and NANOG must be dynamically regulated to guide pluripotent stem cells through intermediate progenitor states toward desired mature cell lineages, such as neurons or cardiomyocytes for tissue repair. Moreover, in gene therapy, cell-type-specific promoters are critical for targeting therapeutic gene expression exclusively to diseased tissues while avoiding off-target effects in healthy cells [4]. A notable example is the use of liver-specific promoters to

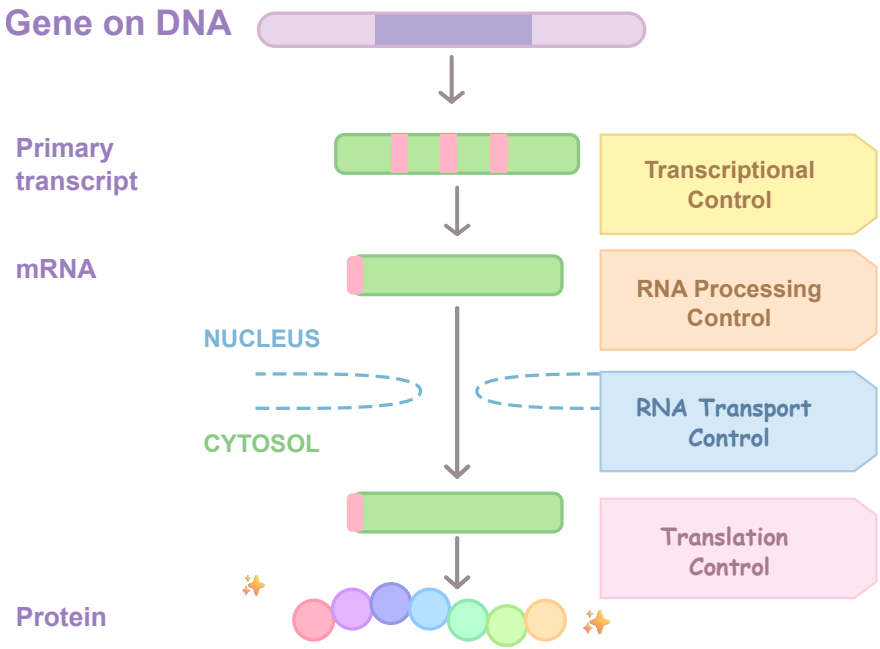

Figure 1.1: Overview of gene expression regulation: Gene expression is controlled at multiple levels, including transcriptional control of the primary transcript, RNA processing, RNA transport from the nucleus to the cytosol, and translational control, ultimately determining the final protein output.

drive the expression of clotting factors in hemophilia patients, or neuron-specific promoters to deliver therapeutic genes for neurodegenerative diseases, ensuring both efficacy and safety of the treatment.

1.2 AI-Driven Approaches to Model Cellular Behaviour

The convergence of artificial intelligence and biology has opened new paradigms for understanding, modeling, and controlling cellular systems. Neural networks, in particular, are capable of capturing complex, multidimensional, and nonlinear aspects of biology—relationships that traditional linear methods cannot model effectively. Building on this, machine learning has enabled the development of predictive models that can decode the intricate regulatory logic governing gene expression [5]. These AI-driven approaches provide powerful tools for the rational analysis and control of cellular behavior [6], allowing us to simulate and screen a vast number of scenarios rapidly and efficiently [7]. This dramatically reduces reliance on traditional trial-and-error experimentation and enables modeling of phenomena that have never been observed experimentally. Together, these strategies offer unprecedented opportunities to bridge the gap between our understanding of molecular mechanisms and our ability to model biological systems with precision.

1.3 Aim and scope of the thesis

This thesis aims to explore two strategies for controlling gene expression, demonstrating how deep learning models can be leveraged to engineer biological systems at multiple regulatory levels.

The first approach focuses on temporal control through a data-driven controller that dynamically regulates protein expression over time. Specifically, I implement a neural network-based controller to modulate the expression of an exogenous gene placed upstream of a tetracycline-inducible promoter driving GFP expression, achieving precise temporal trajectories of gene expression in CHO cells in response to varying input conditions.

This part is subject to embargo.

Together, these two approaches demonstrate the versatility of deep learning methods in addressing distinct challenges in gene expression control: temporal dynamics and cellular specificity. By combining predictive modeling with sequence design, this thesis highlights how neural networks can be harnessed to rationally engineer cellular behavior with unprecedented precision.

Chapter 2

Gene Expression and Automatic Control

Precise temporal control of gene expression is essential for numerous applications, from optimizing metabolic production to engineering therapeutic cells. Traditional approaches rely on open-loop systems where gene expression is set at a fixed level, unable to respond to changing cellular conditions or environmental perturbations [8]. Achieving dynamic, closed-loop control requires continuously measuring gene expression and adjusting regulatory inputs accordingly—a task naturally suited to automatic control systems.

However, the underlying dynamics underneath biological systems are often nonlinear, stochastic, and incompletely characterized. Mathematical models, even when available, typically capture only a simplified view of the true system complexity [9]. This motivates the use of data-driven control approaches that can learn control policies directly from experimental observations without requiring explicit mathematical models.

2.1 Challenges of Automatic Control in Gene Expression

Traditional control design relies on detailed mathematical models derived from first principles, such as ordinary differential equation (ODE) systems [10]. However, these approaches face fundamental limitations when applied to complex mammalian gene regulatory networks, due to the inherent non-linearity and complexity of biological systems, the presence of multiple feedback loops and stochastic effects. Moreover, significant cell-to-cell variability further prevents model generalization across different conditions. These limitations motivate data-driven approaches that learn control strategies directly from experimental observations, achieving precise gene expression regulation without requiring complete mechanistic models [11].

Moreover, to implement control in living cells, I need to introduce synthetic tools that allow us to monitor and manipulate the system [12]. Achieving this

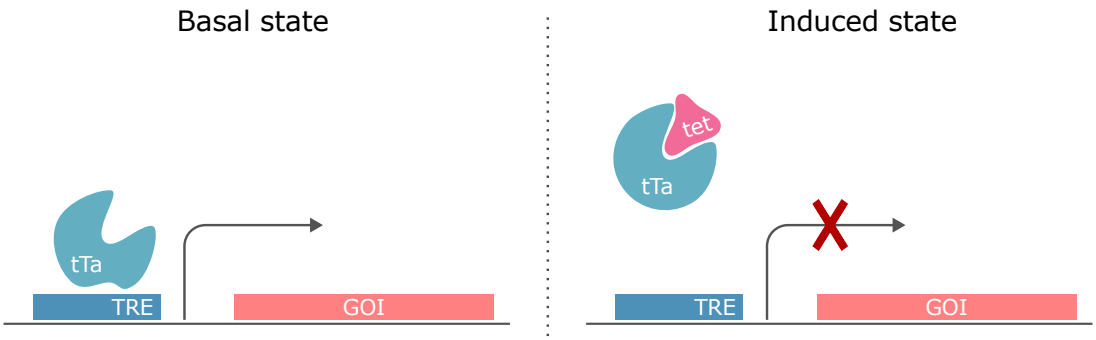

Figure 2.1: Schematic representation of the Tet-Off expression system. Basal state: In the absence of tetracycline, tTA binds to the TRE and activates expression of the gene of interest. Induced state: In the presence of tetracycline, tTA cannot bind to the TRE, thereby preventing expression of the gene.

requires leveraging approaches from synthetic biology to embed the necessary control components directly into the cells. In particular, synthetic biology aims to apply engineering principles to biological system. At its core, it treats cells as programmable substrates, where gene expression can be rationally engineered to implement desired computational and regulatory functions.

To achieve this, synthetic biology [13, 14] designs and introduces genetic circuits into cells, using engineered DNA sequences to control gene expression in a predictable manner [15]. These circuits can include sensors, logic gates, and feedback controllers that allow the cell to respond to internal and external signals in a programmable way. Essentially, it provides the tools to monitor, manipulate, and control cellular behavior, enabling the implementation of complex regulatory functions that would be difficult or impossible with natural cellular networks alone [16, 17].

2.2 The Tet-OFF expression system

The Tet-Off system represents a paradigmatic example of synthetic biology implementation for controllable gene expression [18]. This artificial regulatory mechanism operates through tetracycline-mediated transcriptional repression, where the expression of a gene of interest (GOI) is controlled by the administration of tetracycline.

2.2.1 Molecular Mechanism

The Tet-Off system consists of two key components:

- **tTA Protein**: A fusion protein composed of the Tet repressor (TetR) DNA-binding domain and the VP16 activation domain
- **TRE (TetR-responsive element)**: A promoter containing tetO operator sequences upstream of the gene of interest

The system constitutively express the artificial transcriptional activator **tTA** (tetracycline-responsive transcriptional activator) through synthetic biology engineering. In the absence of tetracycline (Fig 2.1), the system expresses the GOI at its maximum level, since the tTA protein specifically binds to the tetracycline response element (TRE) located in the promoter region upstream of the target gene. This synthetic protein-DNA interaction results in the recruitment of the endogenous transcriptional machinery, leading to active transcription of the GOI. Consequently, in the absence of tetracycline, tTA freely binds to the TRE sequence, allowing transcription of the GOI to remain constitutively active, which causes cells to exhibit measurable green fluorescence. When tetracycline is present (Fig 2.1), the antibiotic molecules bind with high affinity to the tTA protein, forming a stable tTA-tetracycline complex. This conformational change renders the tTA protein incapable of binding to the TRE sequence, effectively blocking transcriptional activation. In this sense the Tet-off system operates as an ON-OFF production system.

This synthetic gene circuit has been engineered into Chinese Hamster Ovary (CHO) Fig 2.2 cells through genetic modification techniques. To measure the level of expression, I coupled the Tet-Off system with GFP. The experimental setup employs destabilized variants that accelerate the dynamic response. These modified molecular components contain destabilization sequences that are not present in their natural counterparts, accelerating both the decay and recovery phases of the fluorescence signal [19].

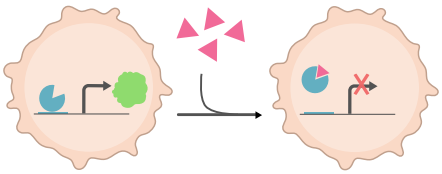

Figure 2.2: Schematic representation of the Tet-Off system integrated in CHO cells. In the absence of tetracycline (pink), the cells express the GPT. When tetracycline is added, expression is inhibited, and production stops.

2.2.2 Mathematical Model of the system

This biological processes can be captured through a system of ordinary differential equations (ODEs) that describe the temporal evolution of key molecular species. Through systematic experimentation Postiglione et al. [20] were able to characterize the system dynamics and derive a linear state-space model that accurately captures the essential features of the Tet-Off circuit behavior. The model was obtained through careful parameter identification based on time-series fluorescence measurements under controlled tetracycline perturbations.

The identified linear model takes the form of a continuous-time state-space system:

$$\dot{x} = \begin{pmatrix} -0.0256 & 0 & 0 \\ 0.0501 & -0.0257 & 0 \\ 0 & 0.0209 & -0.0045 \end{pmatrix} x + \begin{pmatrix} 0.0029 \\ 0 \\ 0 \end{pmatrix} u \tag{2.1}$$

$$y = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix} x \tag{2.2}$$

where the state vector $x \in \mathbb{R}^3$ represents the concentrations of key molecular species, the scalar input $u \in \mathbb{R}$ represents the tetracycline concentration, and the scalar output $y \in \mathbb{R}$ represents the measured fluorescence intensity.

This system presents two equilibrium points (Fig 2.3): one at maximum expression (absence of tetracycline) and one at minimum expression (presence of tetracycline). The second equilibrium point is reached after some descent dynamics, exhibiting the characteristic bistable behavior where the system transitions between stable states through well-defined temporal trajectories. Maintaining expression at an intermediate level requires active dynamic control of the system.

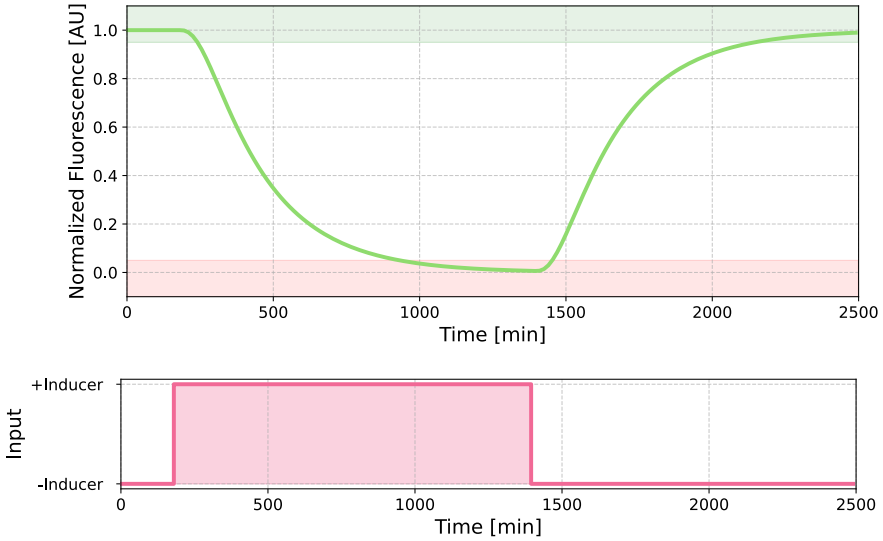

Figure 2.3: Dynamics of the Tet-Off system showing normalized fluorescence response over time. The green trace shows GFP expression levels, while the red shaded regions indicate periods of tetracycline administration (inducer present). The system demonstrates bistable behavior with rapid fluorescence decay upon tetracycline addition (200 min) and slower recovery upon removal (1300-2000 min), illustrating the asymmetric dynamics characteristic of the Tet-Off regulatory circuit.

2.3 Data-Driven Control Approaches

When accurate mathematical models are unavailable or impractical, data-driven control offers an alternative by learning control policies directly from observed

system behavior. Such approaches have proven successful in robotics, autonomous systems, and complex industrial processes where first-principles modeling is intractable [21, 22].

Reinforcement learning (RL) has emerged as a dominant paradigm, where an agent learns optimal control policies through trial-and-error interaction with the system. Modern deep RL methods combine neural networks with RL algorithms—such as Deep Q-Networks, Policy Gradient methods, and Actor-Critic architectures [23]—enabling learning in high-dimensional state spaces [22, 24, 25]. Although powerful, these methods are often sample inefficient, a significant limitation in biological systems where experimental data are costly and time-consuming to generate. Model Predictive Control (MPC) provides a complementary framework by optimizing control actions based on predicted future system behavior, with data-driven variants learning predictive models directly from experimental observations [25, 26]. Deep MPC integrates neural networks with MPC, allowing control over complex dynamics without explicit first-principles models [26].

Data-driven control has been successfully demonstrated in bacterial systems, where short experimental timescales and relatively simple dynamics make trial-and-error feasible [27, 28, 29, 30]. For example, Lugagne et al. [31] implemented in-vivo control using deepMPC, combining automated microscopy for real-time fluorescence measurement with optogenetic inducers to maintain bacterial cells at target expression levels. These studies illustrate the feasibility of closed-loop control in living organisms [32].

Extending these approaches to mammalian cells is substantially more challenging. Longer experimental timescales make iterative learning slow and costly, while complex gene regulatory networks result in high-dimensional, nonlinear dynamics. Pronounced cell-to-cell variability and stochastic gene expression further complicate control, especially in small experimental populations. Finally, limited data availability restricts learning efficiency, making data-efficient control strategies essential for practical implementation in mammalian systems.

Pioneer Studies in Mammalian Control

Despite these challenges, pioneering studies have begun demonstrating classical control strategies in mammalian cells. Fracassi et al. [33] and Postiglione et al. [34] implemented feedback control in mammalian cells using microfluidic platforms, demonstrating that closed-loop regulation is feasible despite the experimental challenges. These studies employed model-based control approaches (PID controllers, MPC) that require less data than learning-based methods, establishing proof-of-concept for temporal control in mammalian systems.

However, these classical approaches rely on simplified linear models or heuristic tuning, limiting their applicability to systems where such approximations are valid. The gap remains for data-driven methods that can learn complex control policies while operating within the constraints of limited mammalian cell data.

2.4 LSTM for Gene Expression Dynamics

A critical challenge in data-driven control of gene expression is the slow dynamics of mammalian gene expression. Combined with limited sampling frequency, this means that control decisions must often be made based on incomplete temporal information. This problem can naturally be formulated as a time series forecasting task.

Long Short-Term Memory networks (LSTMs) have demonstrated remarkable success in capturing temporal dependencies in biological time series data [35, 36, 37]. Their ability to learn from irregularly sampled time series while maintaining memory of relevant historical information makes them particularly well-suited for the sparse measurement scenarios typical of biological experiments. While feedforward neural networks process each input independently, without retaining information about previous observations, recurrent neural networks (RNNs) can theoretically maintain temporal context through recurrent connections. However, vanilla RNNs suffer from vanishing and exploding gradient problems [38], which make learning long-term dependencies—precisely the timescale at which gene regulatory processes operate—very difficult.

LSTMs address these limitations through a carefully designed memory cell architecture that can selectively retain, update, or forget information over long time horizons. This allows the network to capture the slow dynamics and temporal dependencies characteristic of gene expression. By training on historical time-series data from the biological system, the LSTM learns an implicit model of gene regulatory dynamics without requiring explicit knowledge of the underlying biology, providing a data-efficient approach to modeling complex biological processes. Once trained, the model can be integrated into a control platform to implement closed-loop regulation.

2.4.1 Overview of the Data-Driven Controller

Negative feedback loops are fundamental regulatory motifs that provide stability. The output inhibits its own production, creating self-regulation within a desired range. When protein abundance increases beyond the setpoint, the controller decreases induction to reduce production. Conversely, when levels drop, induction is increased to restore expression. This continuous adjustment maintains stable expression despite perturbations. A data-driven negative feedback control system operates as follows (Fig 2.4):

1. **Measurement**: Monitor current protein expression level
2. **Error calculation**: Compare measured level with desired setpoint
3. **Model prediction**: Use LSTM to predict system response to potential actions
4. **Optimization**: Select action that minimizes deviation from setpoint

5. **Actuation**: Apply selected control input to the system

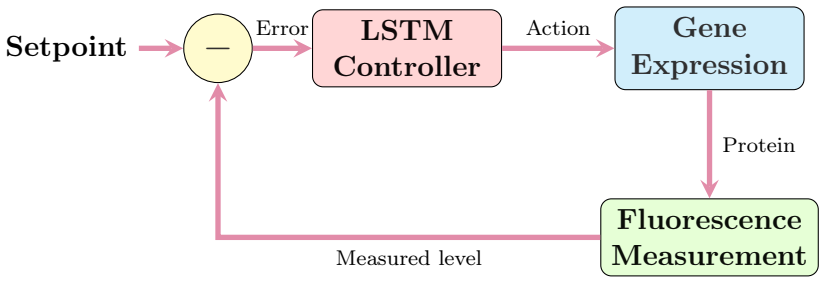

Figure 2.4: Block diagram of the data-driven controller. The selected setpoint is compared with the measured protein level to compute the error, which is fed into the LSTM-based controller. The controller outputs an action that modulates gene expression, altering the level of the fluorescent protein. The updated measurement is fed back into the controller, closing the loop.

Chapter 3

Dynamic Control of Gene Expression via Neural Model Predictive Control (MPC)

The goal is to model in a data-driven fashion the tet-off expression system described in Section 2.2.2 and use this model to control the expression level of the gene of interest (GOI). To achieve this, I first explore an *in-silico* approach, simulating mathematically how the biological system may respond. Subsequently, I aim to validate the approach *in vitro* using a microfluidic platform.

3.1 In-Silico Control

As a preliminary step, I evaluate the regulation of the *Tet-Off* expression system using a data-driven approach in a fully simulated environment. This allows me to design and test control strategies computationally before performing experimental validation.

Specifically, I first generated synthetic time-series data using a mathematical model based on ordinary differential equations (ODEs) that describe the dynamics of the Tet-Off system. A Neural Network (NN) is then trained on this simulated dataset to learn the system's behavior. The trained NN is subsequently integrated into a Model Predictive Control (MPC) framework, which determines the optimal sequence of tetracycline (Tet) inputs to regulate gene expression.

Finally, the MPC is tested *in silico* by simulating the control process, where the ODE-based model acts as the environment. This setup enables me to validate the effectiveness of the data-driven control strategy before moving to real biological experiments.

3.1.1 Synthetic experiments generation

To train the neural network model, I generated synthetic datasets using the mathematical model described in Section 2.2.2. The simulations were designed to closely mimic real experimental conditions, adopting a sampling interval of

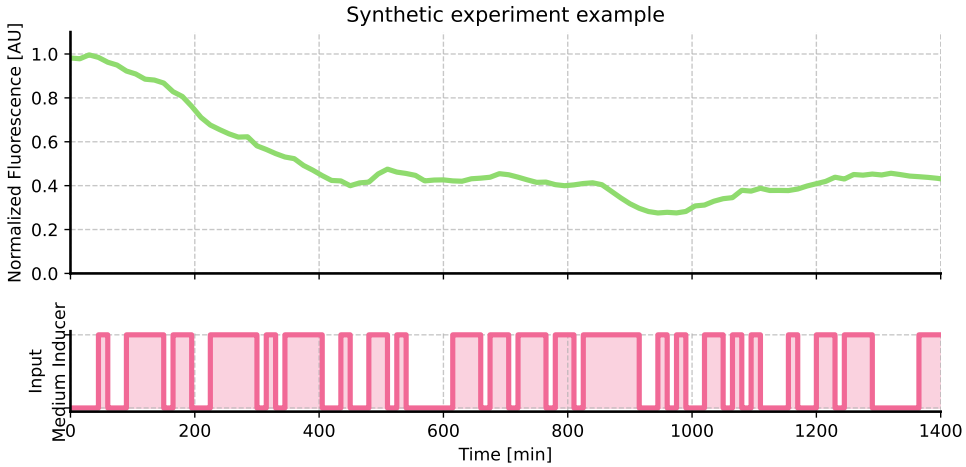

Figure 3.1: Simulated experiment example. The x-axis represents time [min]. The green line shows the normalized fluorescence level, while the pink trace indicates the presence or absence of the inducer over time.

15 minutes and a total duration of 1400 minutes—consistent with the *in vitro* studies.

I simulated six independent experiments, each driven by a sequence of random input signals representing tetracycline induction levels (An example is shown in Figure 3.1). These inputs modulate the system's dynamics, producing corresponding fluorescence trajectories of GFP expression.To account for experimental variability, I added Gaussian noise $\eta(t)$ with zero mean and standard deviation 0.05 to the simulated data, defined as

$$\eta(t) \sim \mathcal{N}(0, 0.05),$$

so that the noisy GFP signal $G_{\text{noisy}}(t)$ is given by

$$G_{\text{noisy}}(t) = G(t) + \eta(t),$$

where $G(t)$ is the noise-free simulated GFP expression. Each fluorescence trajectory, together with the administered input (*tet* or *no tet*), can be represented as a sequence of observations

$$\{(x_1, u_1), (x_2, u_2), \ldots, (x_t, u_t)\},$$

where x_t represents measured biological variables (e.g., fluorescence) and u_t represents control actions (e.g., inducer concentrations) at time t.

3.1.2 Time-Series Forecasting with LSTM

I modeled the system as a time series forecasting problem to capture its temporal dynamics and predict future fluorescence levels. Each experiment can be described as a trajectory formed by the state at time t, consisting of fluorescence

and induction levels: $\text{system}_t = [s_t, a_t]$, where s_t represents the fluorescence intensity and a_t represents the tetracycline concentration. the goal is to use this data to predict the system behavior at future time points. The model takes as input a sequence of the past n time steps containing both fluorescence measurements and tetracycline concentrations to predict the fluorescence value at the next time step (Fig 3.3):

$$\begin{aligned} s_{t+1} &= f_\theta(x_{t-n+1}, x_{t-n+2}, \ldots, x_t) \\ \text{where} \quad x_t &= [s_t, a_t] \end{aligned} \tag{3.1}$$

where s_t represents the fluorescence intensity at time t, a_t denotes the tetracycline concentration, and f_θ represents the LSTM-based mapping parameterized by weights θ.

3.1.3 LSTM Architecture

The LSTM architecture enables the model to learn a representation of the hidden system states by processing sequential information, making it particularly suitable for partially observable systems where direct state measurements are unavailable. To leverage this capability, I structured the data using a sliding window approach with a window size of 12 time points and a step size of 1. This windowing strategy provides the LSTM with sufficient historical context about both fluorescence and tetracycline concentrations to predict the next fluorescence level while creating multiple overlapping training samples from each experiment.

Given that each experiment spans 1400 minutes with a sampling time of 15 minutes, I obtained 93 total samples per experiment. Using 12 overlapping windows creates 82 training samples per experiment, yielding a total of approximately 492 samples across all experiments. I divided the data as follows: 4 experiments for training (328 samples), 1 for validation (82 samples), and 1 for testing (82 samples) to simulate real-world data distribution.

Given the limited dataset size, I employed a compact neural network (Fig 3.2) architecture to prevent overfitting . The model comprises an input layer of size 12×2, a single LSTM layer with h hidden units, and a fully connected output layer, resulting in approximately 600 trainable parameters (Table 3.1.3).

Layer (type)	**Output Shape**	**Param #**
InputLayer	(None, 12, 2)	0
LSTM	(None, 20)	20,440
Dense	(None, 20)	4,970
Dense	(None, 1)	71

Table 3.1: Summary of the model layers, including the number of trainable parameters for each layer.

The hyperparameters, including the window size n and the number of hidden

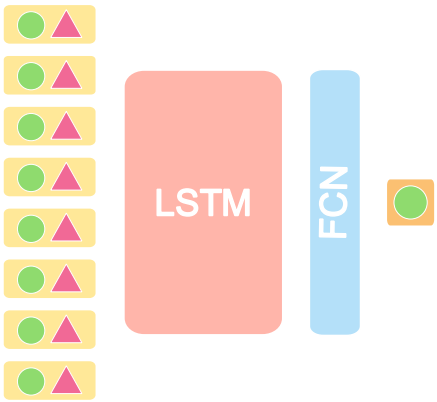

Figure 3.2: LSTM architecture: The neural network takes an input of size $n_{\text{psat}} \times 2$, where the two features correspond to fluorescence (green) and the control action (pink) applied to the system. The input is processed by an LSTM layer, followed by a fully connected layer that maps the output to the predicted fluorescence.

units h, were optimized through systematic grid search based on validation set performance.

3.1.4 Training Procedure

The model was trained for 300 epochs using the Adam optimizer [39] with a learning rate of $\alpha = 0.001$. I employed mean squared error (MSE) as the loss function:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2 \tag{3.2}$$

where y_i represents the target expression level and $\hat{y}_i$ the predicted output for sample i.

To prevent overfitting, I implemented early stopping with a patience of 50 epochs, monitoring validation loss for convergence. Gradient clipping with a maximum norm of 1.0 was applied to ensure training stability and prevent exploding gradients during backpropagation.

As shown in Figure 3.4, both training and validation losses decreased monotonically until convergence, indicating successful model optimization without signs of overfitting. The convergence behavior demonstrates that the model effectively learned the relationship between control actions and gene expression dynamics.

3.1.5 Model Performance Evaluation

I assessed the model's predictive accuracy using mean squared error (MSE) on the held-out test set for single-step forecasting. The results demonstrate strong

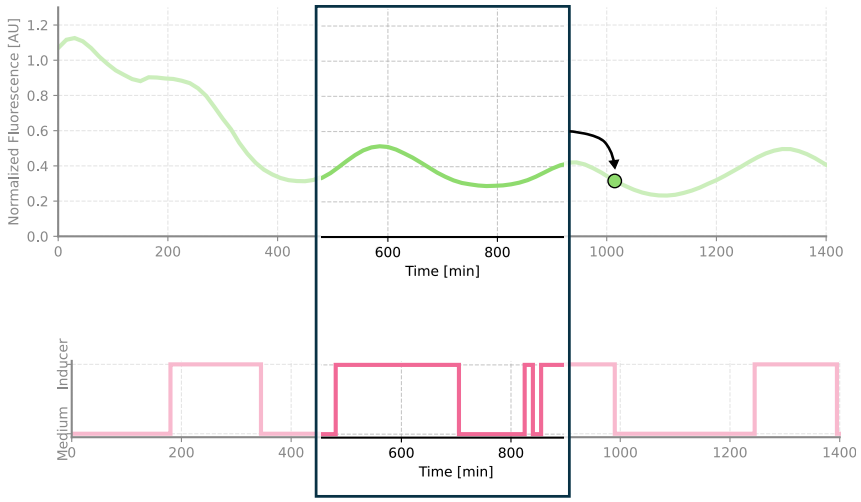

Figure 3.3: Schematic of the prediction process. Inputs consist of past fluorescence and experimental actions using sliding windows, while the output is the predicted fluorescence at the next time step.

predictive performance, achieving an MSE of less than 7.5% on the test data. As shown in Figure 3.4, the predicted trajectory closely follows the ground truth measurements, appearing nearly linear over short time intervals due to the slow dynamics of the system relative to the 15-minute sampling period. This behavior is expected given that the characteristic time constants of gene expression and protein degradation occur on timescales significantly longer than the observation interval, resulting in smooth, gradually evolving trajectories that the LSTM model accurately captures.

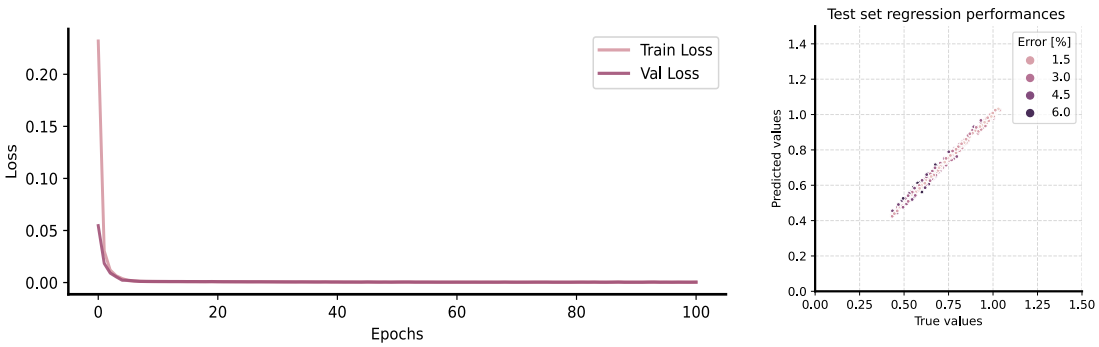

Figure 3.4: Training and validation loss convergence without overfitting. **Left:** Model loss during training, showing convergence. **Right:** Regression accuracy on the test set, demonstrating a clear correlation between true and predicted values.

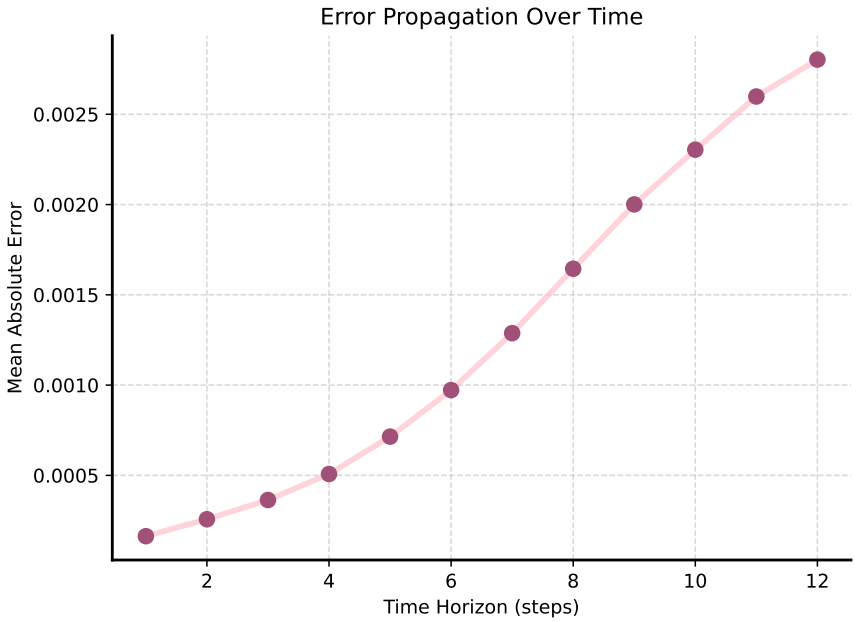

Figure 3.5: Model error propagation through time steps, calculated as mean of the MSE among all the time points

3.1.6 Recursive Multi-Step Prediction

To evaluate the model's ability to simulate long-term system behavior, I assessed its performance in a recursive forecasting setting (Fig 3.6). In this mode, the model operates in an autoregressive manner: the predicted fluorescence at time $t+1$ is fed back as input for predicting the state at time $t+2$, and this process continues iteratively for extended time horizons. This recursive prediction strategy is critical for evaluating whether the model has learned the true underlying dynamics rather than merely fitting short-term correlations in the training data.

I extended the prediction horizon from single-step (15 minutes) to multi-step forecasting up to 120 minutes ahead (8 time steps). At each prediction step, the model's output is incorporated into the sliding window input for the subsequent prediction, creating a fully autonomous simulation of the system trajectory. Despite the accumulation of prediction errors inherent in recursive forecasting, the model maintains accurate predictions that remain consistent with the true system dynamics over the extended time window. This result demonstrates that the LSTM has successfully captured the essential temporal dependencies governing the tet-OFF system, enabling reliable long-term forecasting without direct access to ground truth measurements as shown in Figure 3.5.

3.1.7 Model Predictive Control Design

Having validated the LSTM model's predictive capabilities, I decided to employ it as the foundation for a model predictive control (MPC) framework as shown in Fig3.7. To evaluate the controller's performance in a realistic scenario, I per-

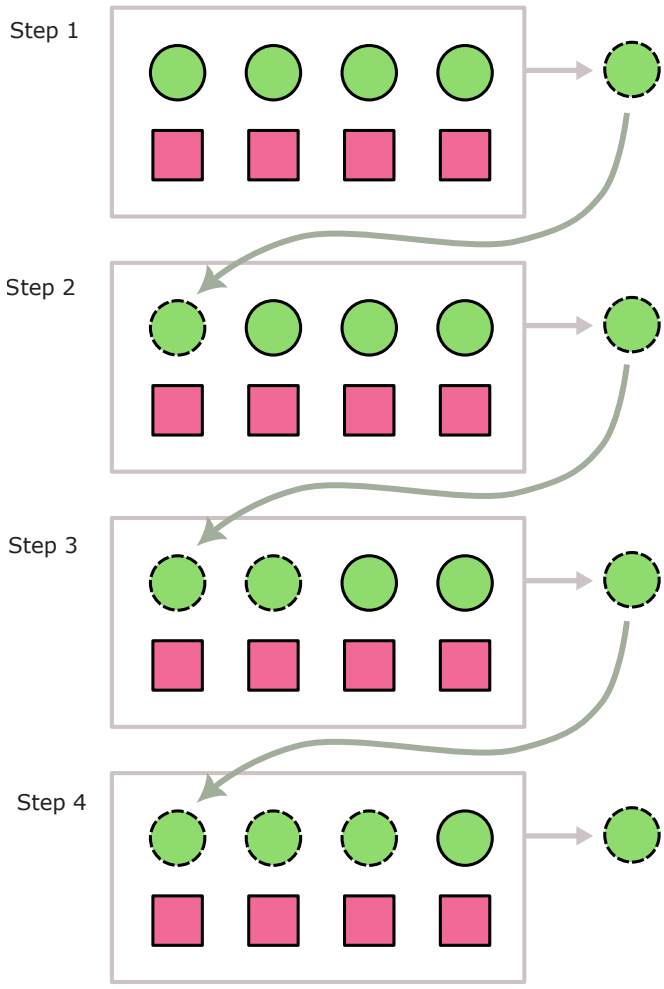

Figure 3.6: Recurrent multi-step prediction: at each iteration, the predicted fluorescence is fed back as input to the model to simulate longer time horizons. In Step 1, only measured fluorescence data are used to predict the next fluorescence level. In Step 2, the prediction from Step 1 is incorporated into the input, resulting in a combination of predicted and measured data that is used to predict the fluorescence level for the next time step. The process is iterated for all the desired steps

formed simulated control experiments where the LSTM-based controller interacts with a modified version of the mathematical model described in Section 2.2.2. Specifically, I introduce parametric uncertainty by randomly perturbing the system matrices by up to 30% to simulate model mismatch and biological variability. Each control experiment begins with the system at maximum fluorescence (absence of tetracycline) and tasks the controller with driving and maintaining the fluorescence at a specified intermediate reference level.

The MPC strategy operates through receding horizon optimization at each time step t. The controller evaluates potential control sequences over a prediction horizon of 180 minutes (12 time steps) and selects the action that minimizes the predicted deviation from the target fluorescence. The optimization procedure consists of the following steps:

1. **Control sequence generation**: Generate N candidate control sequences $\{u_t, u_{t+1}, \ldots, u_{t+H}\}$, where each $u_i \in \{0, 1\}$ represents the absence or presence of tetracycline, and H denotes the prediction horizon.

2. **Trajectory simulation**: For each candidate control sequence, use the trained LSTM model to recursively simulate the resulting fluorescence trajectory $\{\hat{y}_{t+1}, \hat{y}_{t+2}, \ldots, \hat{y}_{t+H}\}$.

3. **Cost evaluation**: Compute the mean squared error between the predicted trajectory and the constant reference target y_{ref}:

$$J_i = \frac{1}{H} \sum_{k=1}^{H} (\hat{y}_{t+k} - y_{\text{ref}})^2 \tag{3.3}$$

4. **Action selection**: Select the control sequence with minimum cost:

$$u^* = \arg\min_{u \in \mathcal{U}} J_i \tag{3.4}$$

   where $\mathcal{U}$ represents the set of all candidate control sequences, and apply only the first action u_t^* of the optimal sequence to the system at time t.

This process repeats at each time step in a receding horizon fashion, continuously re-planning based on the most recent system state. By implementing only the first control action and then re-optimizing, the controller adapts to actual system behavior and compensates for modeling errors and disturbances while leveraging the LSTM's learned dynamics to anticipate future evolution.

3.1.8 Control Performance and Robustness Analysis

I evaluate the controller's performance across a range of reference setpoints 0.25, 0.50 and 0.75 arbitrary units (AU) of normalized fluorescence. For each reference value, I conduct 10 independent control experiments and reported the averaged results to assess consistency and robustness.

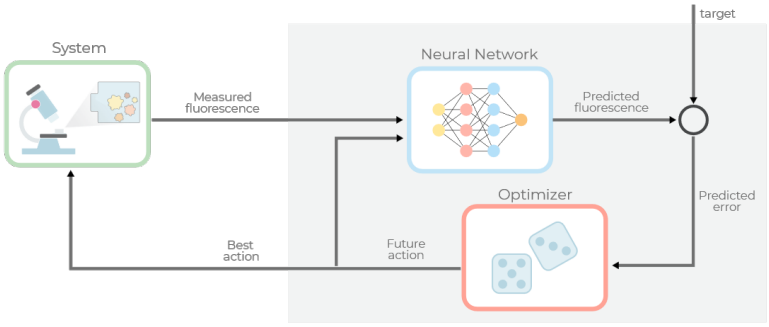

Figure 3.7: Block diagram of the neural model predictive control system integrated into the negative feedback loop. Fluorescence at time t is measured and fed into the neural network along with the input history. The algorithm then computes the optimal control action with respect to the target value, which is subsequently applied to the system.

Fig 3.8 illustrates a representative control experiment with a reference setpoint of 0.5 AU. The system initially operates at maximum fluorescence in the absence of tetracycline. At $t = 180$ minutes, the MPC controller is activated and begins administering tetracycline to reduce the fluorescence level. The controller successfully drives the system toward the desired setpoint and, once reached, maintains stable regulation around the target value for the remainder of the experiment. This behavior demonstrates the controller's ability to both achieve accurate setpoint and provide sustained regulation despite model uncertainty and process noise.

To quantify control performance, I computed the mean squared error (MSE) between the actual fluorescence and the reference value during the second half of each experiment (after initial transient dynamics):

$$\text{MSE} = \frac{1}{T_{\text{control}}} \sum_{t=T_{\text{start}}}^{T_{\text{end}}} (y_t - y_{\text{ref}})^2 \tag{3.5}$$

where T_{start} represents the time when the system first reaches the vicinity of the setpoint, and T_{end} marks the conclusion of the experiment.

Table 3.2 summarizes the control performance across all tested reference values. The results demonstrate consistently low tracking error across the entire operational range. These results compare favorably with classical model-based control approaches, while offering the advantage of learning directly from experimental data without requiring explicit mechanistic models.

The consistent performance across the full range of setpoints confirms the robustness of the LSTM-based MPC framework. The controller successfully handles the inherent challenges of partial observability, slow system dynamics, and model mismatch introduced by the 30% parameter perturbations. This demonstrates that the learned neural network model captures sufficient information

Table 3.2: Control performance across different reference setpoints. Results are averaged over 10 independent experiments for each condition.

Reference (AU)	MSE (AU^2)	Std. Dev. (AU^2)
0.25	0.0032	0.0008
0.5	0.0035	0.0007
0.75	0.0025	0.0004

about the system's dynamics to enable effective predictive control, offering a promising data-driven alternative to traditional mechanistic modeling approaches for synthetic gene circuit regulation.

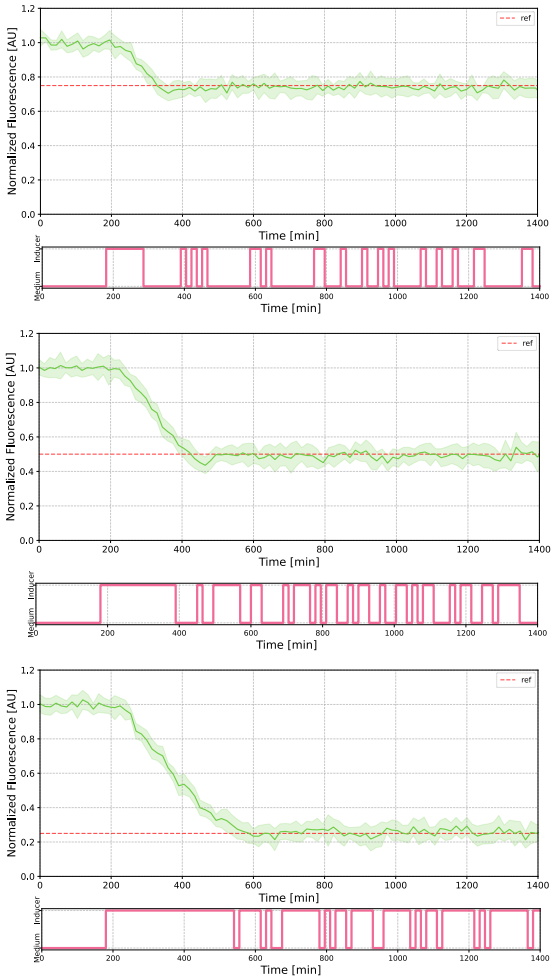

Figure 3.8: In-silico control results for different reference values (0.75, 0.5, and 0.25). For all targets, the controller is able to reach and maintain the desired fluorescence levels. The experiment was performed 10 times; the solid line represents the mean, and the shaded area indicates the standard deviation. Below, the control action over the system from a single experiment is shown.

3.2 In Vitro Control Implementation

To validate the data-driven MPC in a real biological system, I integrated the LSTM-based MPC controller into an automated experimental platform. The control system consists of three main components: (i) a Nikon fluorescence microscope for real-time monitoring of cellular GFP expression, (ii) a temperature-controlled incubator maintaining physiological conditions at 37°C and 5% CO_2, and (iii) a computer-controlled microfluidic perfusion system equipped with two motorized syringes that automatically deliver either tetracycline-free medium or tetracycline-containing medium to the cell culture, as shown in Figure 3.9.

The closed-loop control architecture operates as follows: at each 15-minute sampling interval, the microscope acquires fluorescence images of the cell population, the image processing pipeline quantifies the mean GFP intensity, and this measurement is fed to the LSTM controller running on the computer. The controller then executes the MPC optimization procedure described in Section 3.1.7 to determine the optimal medium selection, and the corresponding motorized syringe automatically perfuses the appropriate medium into the microfluidic chamber. This automated feedback loop enables continuous, real-time regulation of gene expression without manual intervention.

3.2.1 Microfluidic Platform and Experimental Setup

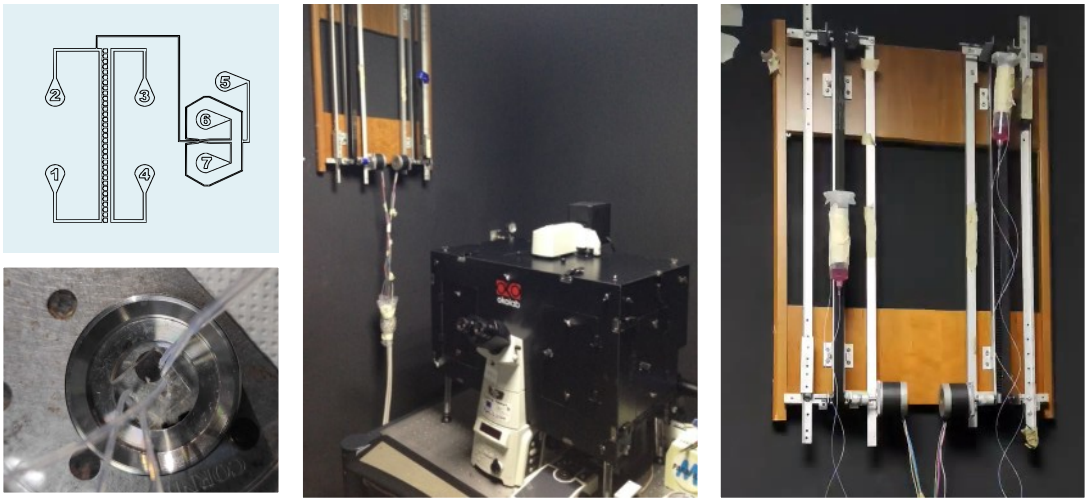

Figure 3.9: Microfluidic platform setup integrating automated microscopy, environmental control, and medium perfusion for closed-loop gene expression regulation. The setup includes a chip for culturing cells, microfluidic channels for nutrient delivery and waste removal, a microscope with an incubator, and motorized syringes for precise fluid control.

To maintain cell viability during extended experiments and enable precise control over the cellular microenvironment, genetically modified CHO cells are

cultured within a specialized microfluidic platform. The experimental configuration integrates real-time microscopy, automated medium delivery, and computational image analysis to implement the closed-loop control system described above.

Microfluidic Chip Design and Cell Confinement

Cells are loaded into a polydimethylsiloxane (PDMS) [40] microfluidic chip specifically engineered to confine individual cells within microscale traps (Fig 3.9). This design immobilizes cells at fixed positions throughout the experiment, enabling continuous tracking and longitudinal analysis of the same cell population over extended periods. The chip architecture provides well-defined microenvironments for each cell while allowing uniform medium distribution across all cellular compartments through a network of microchannels.

Environmental Control and Medium Delivery

The microfluidic chip is mounted on the stage of a Nikon Ti inverted microscope and positioned within a temperature-controlled incubator maintaining physiological conditions at 37°C and 5% CO_2 (Fig 3.9). The platform incorporates a dual-syringe perfusion system where each syringe contains a different medium composition:

- **Control medium**: Standard growth medium without tetracycline (baseline condition)
- **Induction medium**: Growth medium supplemented with tetracycline (repression condition)

A computer-controlled coupled with two motorized syringes, enables automated selection and delivery of the desired medium to the microfluidic chip.

Image Acquisition and Analysis Pipeline

The Nikon Ti microscope acquires live-cell images at 15-minute intervals in both brightfield and fluorescence channels. This temporal resolution provides sufficient sampling frequency for control purposes while minimizing phototoxicity and photobleaching effects that could compromise cell viability during long-term experiments.

The image analysis pipeline processes each acquired image through the following automated steps:

1. **Cell segmentation**: Brightfield images are processed using Cellpose, a deep learning-based segmentation algorithm that accurately identifies individual cell boundaries even in densely packed populations, as shown in (Figure 3.10)

2. **Population analysis**: Mean fluorescence intensity is computed across the entire segmented cell population to generate a single scalar feedback signal representing the population-level expression state. For each frame, I calculate the background-corrected mean fluorescence μ_{pop} and standard deviation σ_{pop} across the cell population:

$$\mu_{\text{pop}} = \frac{1}{N_{\text{cells}}} \sum_{j=1}^{N_{\text{cells}}} \left(\frac{\sum_{i \in M_j} (I_{\text{green}}(i) - I_{\text{bg}})}{|M_j|} \right) \tag{3.6}$$

where N_{cells} is the number of segmented cells, M_j is the binary mask for cell j, $|M_j|$ denotes the number of pixels in mask j, $I_{\text{green}}(i)$ is the green channel intensity at pixel i, and I_{bg} is the background intensity. The resulting μ_{pop} serves as the scalar feedback signal for the control system.

3. **Controller input**: The population mean fluorescence value is normalized and fed to the LSTM-based MPC controller to determine the optimal medium selection for the next time step

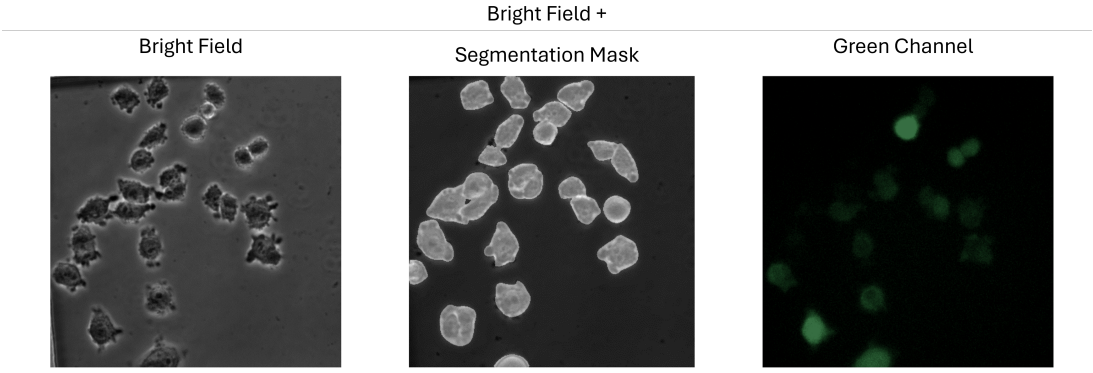

Figure 3.10: Image analysis pipeline showing (left) raw brightfield image, (center) Cellpose segmentation masks identifying individual cells, and (right) corresponding fluorescence image with quantified GFP expression overlaid on segmented regions.

This integrated platform enables fully automated, closed-loop control experiments spanning multiple days without manual intervention, providing a robust testbed for validating data-driven control strategies in living biological systems.

3.2.2 Dataset and Experimental Results

Figure 3.11 illustrates a representative microfluidic experiment. The microfluidic platform contains multiple independent cell chambers, each housing a distinct population of CHO cells. While the controller makes decisions based on the mean fluorescence measured in a single designated target chamber, the dual-syringe system simultaneously delivers the selected medium to all chambers in

parallel. This configuration provides an inherent source of biological replication: each chamber experiences the identical sequence of tetracycline exposures, yet the fluorescence dynamics may vary due to chamber-to-chamber differences in cell density, initial conditions, and stochastic gene expression.

As shown in Fig 3.11, I observe fthe replicate trajectories evolving under the same control input sequence, each representing a different chamber within the same experiment. The left panel displays the temporal evolution of mean fluorescence intensity for each cells chamber (individual colored traces) along with the population-averaged trajectory across all chambers (bold line) and the associated standard deviation (shaded region). The right panel shows a microscopy image of the microfluidic chip, highlighting the spatial arrangement of the cell chambers. This multi-chamber design enables me to assess both control performance in the target chamber and the reproducibility of system response across biological replicates.

For model training and validation, I collected six independent experiments with varied tetracycline administration patterns of 1400 minutes each. The dataset is partitioned as follows: fthe experiments for training the LSTM model, one for hyperparameter validation, and one held-out experiment for final testing. Each experiment contributes multiple replicate trajectories from its constituent chambers, enriching the dataset while maintaining independence at the experiment level to prevent data leakage between training and testing conditions. In this way I acquired 1819 data points for train and 213 for testing.

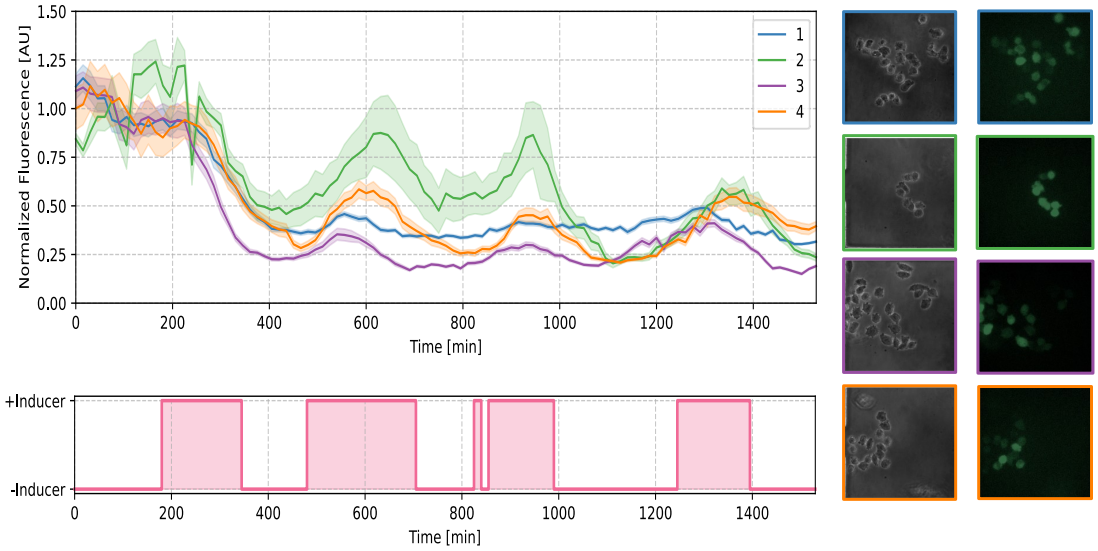

Figure 3.11: Example of *in vitro* experimental data showing fluorescence dynamics. Each line represents a different cells chamber, with the solid line indicating the mean and the shaded area representing the standard deviation. The cells were subjected to the input train shown below. While the dynamics are similar across chambers, the absolute values differ. On the left, the fluorescence levels of the cells at time point $t = 0$ are shown.

3.2.3 Invitro training

The model was trained for 300 epochs using the Adam optimizer [39] with a learning rate of $\alpha = 0.001$. I employed mean squared error (MSE) as the loss function:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{3.7}$$

where y_i represents the target expression level and $\hat{y}_i$ the predicted output for sample i.

To prevent overfitting, I implemented early stopping with a patience of 50 epochs, monitoring validation loss for convergence. Gradient clipping with a maximum norm of 1.0 was applied to ensure training stability and prevent exploding gradients during backpropagation. As shown in Figure 3.12, both training and validation losses decreased monotonically until convergence, indicating successful model optimization without signs of overfitting. The convergence behavior demonstrates that the model effectively learned the relationship between control actions and gene expression dynamics.

To evaluate the model's ability to simulate long-term system behavior, I assessed its performance in a recursive forecasting setting (Fig 3.6)

I extended the prediction horizon from single-step (15 minutes) to multi-step forecasting up to 180 minutes ahead 128 time steps). At each prediction step, the model's output is incorporated into the sliding window input for the subsequent prediction, creating a fully autonomous simulation of the system trajectory. Despite the accumulation of prediction errors inherent in recursive forecasting, the model maintains accurate predictions that remain consistent with the true system dynamics over the extended time window. This result demonstrates that the LSTM has successfully captured the essential temporal dependencies governing the tet-OFF system, enabling reliable long-term forecasting without direct access to ground truth measurements as shown in Figure 3.13.

3.2.4 In Vitro Control Protocol

The closed-loop control experiment proceeds through the following protocol: at each 15-minute sampling interval, the microscope acquires brightfield and fluorescence images of the target chamber. The image analysis pipeline performs automated cell segmentation and quantifies the mean population fluorescence, which is transmitted to the control computer. During the initial 180-minute period, the system operates in open-loop mode without control intervention, allowing cells to reach steady-state maximum expression in the absence of tetracycline. This baseline period serves two purposes: it ensures cells have adapted to the microfluidic environment, and it establishes the maximum fluorescence level used for normalization of all subsequent measurements.

After the initialization period, closed-loop control is activated. At each time step, the LSTM-based MPC controller receives the normalized fluorescence mea-

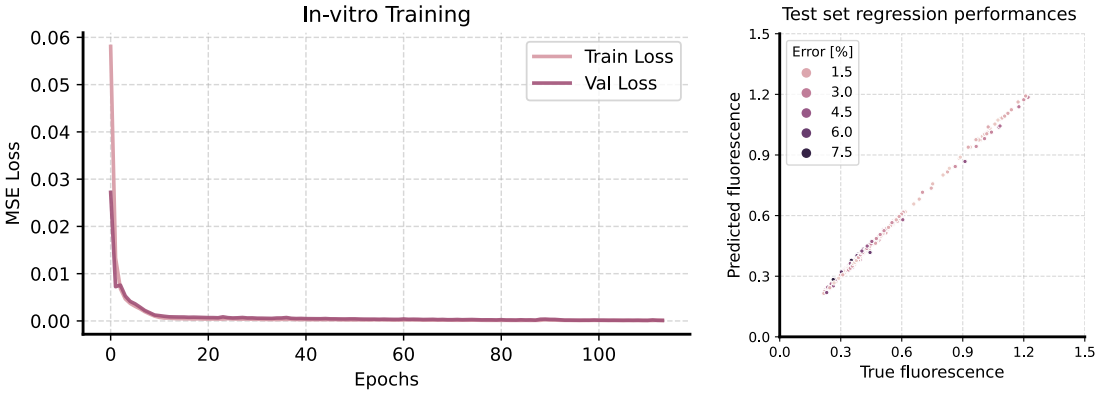

Figure 3.12: Training and validation loss convergence without overfitting. **Left:** Model loss during training, showing convergence. **Right:** Regression accuracy on the test set, demonstrating a clear correlation between true and predicted values.

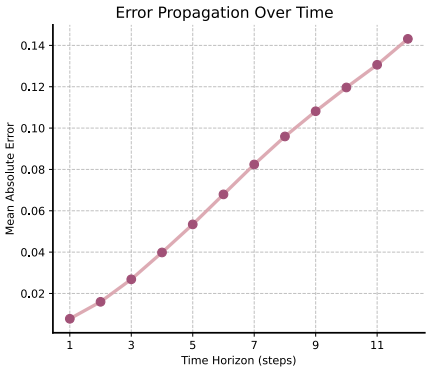

Figure 3.13: MSE error propagation over time. Following the recursive prediction, the mean MSE is estimated over the entire predictive horizon, demonstrating that even at longer time points the error remains below 0.15.

surement, executes the optimization procedure described in Section 3.1.7 to determine the optimal medium selection, and commands the actuator to switch between control and induction media as needed. The system then waits 15 minutes for the cells to respond to the medium change before acquiring the next measurement, completing one control cycle. This process repeats continuously for the duration of the experiment.

3.2.5 In Vitro Control Performance

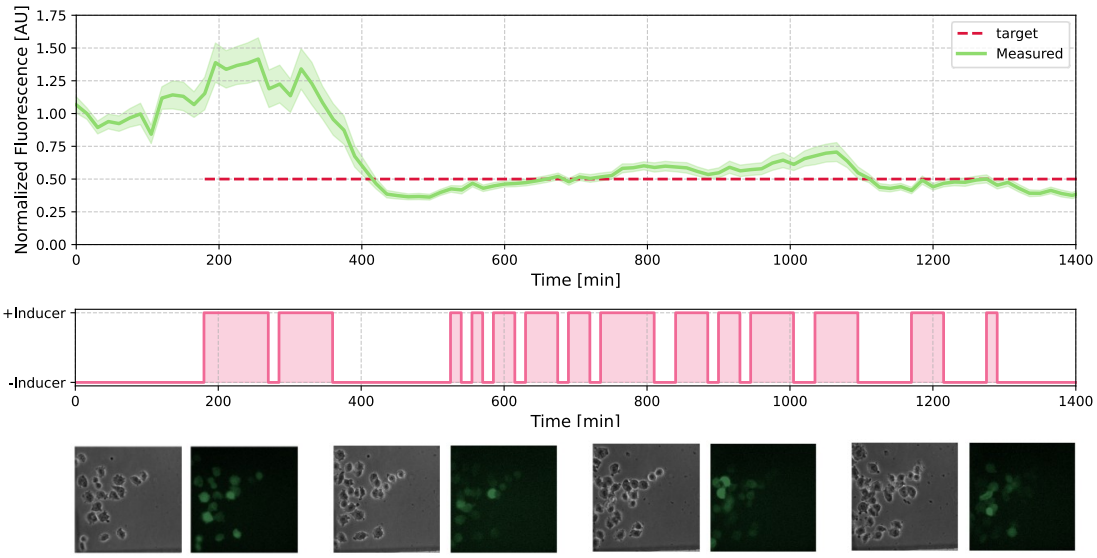

Figure 3.14: In vitro control results showing performance across multiple cell chambers. The dashed line represents the reference set-point of 0.5 a.u. The top panel shows normalized fluorescence measurements over time, with the shaded region indicating measurement variance. The middle panel displays the inducer addition schedule (+Inducer indicates addition, -Inducer indicates no addition). The bottom panel shows representative microscopy images at different time points, with brightfield (left) and fluorescence (right) images demonstrating cellular response to control inputs.

Figure 3.14 demonstrates successful closed-loop regulation of gene expression in living cells. Following the 180-minute initialization period, all fluorescence measurements are normalized by the maximum observed intensity during this baseline phase. I then set the target reference at 0.5 AU (arbitrary units of normalized fluorescence), representing an intermediate expression level between fully repressed and fully induced states.

Upon activation of the controller at $t = 180$ minutes, the system begins administering tetracycline-containing medium to reduce the initially high fluorescence. The controller successfully drives the population mean fluorescence toward the target setpoint and maintains stable regulation around 0.5 AU for the remainder of the 12-hthe experiment. The upper panel of Figure 3.14 shows

the controlled fluorescence trajectory (green line) tracking the reference setpoint (dashed line), while the lower panel displays representative time-lapse microscopy images illustrating the temporal evolution of cellular fluorescence intensity. The experiment achieves a mean squared error of 0.008 AU^2 during the control phase. While the robustness analysis is limited to steady-state MSE, the consistent convergence to the target setpoint and the absence of oscillations or divergence indicate that the controller maintains stable performance under the tested conditions, providing an initial indication of its robustness.

Critically, these results validate that the LSTM model trained on real-world data can successfully generalize to control the real biological system. This demonstrates the controller's ability to regulate previously unseen dynamics and confirms the practical applicability of the data-driven MPC framework.

Chapter 4

Controlling Gene Expression through Proximal Promoter Engineering

The content of this chapter is subject to embargo.

Chapter 5

Design of Cell-Type Specific Synthetic Promoters with PromoterGPT

The content of this chapter is subject to embargo

Chapter 6

Conclusions

The central aim of this thesis has been to demonstrate how deep learning models can be leveraged to engineer biological systems at multiple regulatory levels, addressing two complementary challenges in gene expression control: temporal dynamics and cellular specificity. First, I achieved data-driven control of temporal gene expression without requiring mechanistic mathematical models. Through 1,400-minute experiments in CHO cells, I successfully maintained desired expression levels using an LSTM-based controller trained directly from experimental data. This part is subject to embargo.

6.1 Key Contributions

I successfully implemented an LSTM-based controller for dynamic regulation of gene expression in living cells. the approach consisted of three phases.

In the first phase, I trained and validated the neural network controller in simulation to ensure feasibility. Using the Tet-Off system (where the gene of interest is expressed in the absence of tetracycline), I performed six 1,400-minute experimental conditions to capture the full system dynamics. In the second phase, I integrated real-time microscopy feedback with the predictive neural network model, enabling dynamic modulation of the tetracycline-responsive promoter. Finally, I validated controller performance through trajectory tracking experiments. The controller maintained GFP expression at desired setpoints and followed complex temporal patterns. These results validate that neural network-based control can function in living biological systems without deriving first-principles mathematical models, even with limited training data.

This part is subject to embargo.

6.2 Scientific Context and Implications

The temporal control results demonstrate that effective controllers can be designed without prior knowledge of underlying biological mechanisms by training neural networks to capture system dynamics directly from data. This approach

is particularly valuable for mammalian systems where mechanistic models are difficult to derive and experimental iteration is costly.

This part is subject to embargo.

Together, these contributions represent a paradigm shift from mechanistic to data-driven design, opening new possibilities for rational engineering of biological systems.

6.3 Limitations

While these results are promising, several important limitations must be acknowledged. The neural network controller has been validated exclusively on the tetracycline-inducible GFP system, and its performance with endogenous proteins or therapeutically relevant genes remains to be demonstrated. The biological response dynamics may differ substantially when controlling proteins that themselves participate in cellular regulatory networks or exhibit toxic effects at high concentrations. The experimental platform relies on microfluidic chambers with limited cell populations and chemically inducible systems, which may not fully capture the complexity of in-vitro environments. Furthermore, the field is increasingly moving toward optogenetic control systems that offer faster response kinetics and reduced perturbation of cellular metabolism. The temporal resolution and reversibility of light-based induction may ultimately prove more suitable for precise dynamic control applications.

This part is subject to embargo.

6.4 Future Perspectives

The approaches developed in this thesis provide a foundation for several promising directions that could significantly enhance both the capabilities and applicability of neural network-based gene expression control.

The current controller relies solely on fluorescence intensity as feedback signal. Future implementations could incorporate additional morphological and population-level features to enable more sophisticated control strategies. Integration of cell morphology features such as cell shape, diameter, area, and aspect ratio as additional state variables could enable detection of cellular stress or differentiation states. Rather than relying on extracted features, end-to-end learning from raw microscopy images using convolutional architectures could capture subtle phenotypic changes not captured by scalar metrics. Additionally, the current system administers input in a binarized form (inducer present or absent). Future work should explore continuous input spaces, enabling fine-tuned control at intermediate inducer concentrations (e.g., 20% of maximum concentration). This would provide smoother transitions and potentially reduce cellular stress associated with abrupt environmental changes. With larger datasets and greater computational resources, more sophisticated neural architectures could improve

controller performance. Transformer-based controllers with attention mechanisms could better capture long-range temporal dependencies and enable more precise trajectory tracking, together with enhanced interpretability and visualization capabilities. Furthermore, transfer learning approaches could drastically reduce the number of experiments required. By using a network pre-trained on diverse cellular systems and then adapting it to new control tasks with limited data, the experimental burden for each new application can be minimized. Specifically, a controller trained on one genetic circuit could be fine-tuned for a different circuit by updating only the final layers while keeping the earlier feature extraction layers frozen. This approach reduces data requirements, potentially lowering the number of experiments needed to fit the dynamics of the biochemical reaction. Moreover, transfer learning could address the challenge of changing cellular conditions over time. As cell populations age, undergo metabolic shifts, or experience passage-dependent changes, system dynamics may drift from the original training regime. A pre-trained controller could be continuously adapted through online learning, periodically updating model parameters based on recent observations. This would require quantifying how many data points are necessary to detect distribution shifts. Such adaptive systems could maintain control performance despite long-term changes in cellular behavior. Moving beyond reporter genes to control of functionally important pathways represents a critical next step. One particularly promising application is the autonomous differentiation of stem cells, where neural network controllers could dynamically modulate key transcription factors to guide differentiation trajectories toward desired cell fates. This could enable automated production of specific cell types for regenerative medicine with unprecedented precision and reproducibility. Another important direction is threshold maintenance, where the controller must keep expression of potentially toxic genes below toxicity thresholds while maintaining sufficient activity for desired function.

6.4.1 Cell-Type-Specific Promoter Applications

This part is subject to embargo.

6.5 Concluding Remarks

This thesis demonstrates that machine learning approaches can address fundamental challenges in synthetic biology that have traditionally required extensive mechanistic understanding or exhaustive experimental screening. By training neural networks directly on experimental data for temporal control and leveraging transformer architectures for sequence design, I have shown that data-driven methods can achieve precise, predictable regulation of gene expression.

The temporal control framework provides a template for autonomous regulation of therapeutic genes, while the promoter design methodology enables targeted delivery to specific cell populations. Together, these complementary

approaches—controlling when and where genes are expressed—establish a foundation for the next generation of engineered biological systems.

As these methods mature and expand to encompass broader biological contexts, they promise to accelerate the development of sophisticated gene therapies, enable precise control of cellular behavior, and ultimately bridge the gap between computational prediction and experimental reality in synthetic biology.

Bibliography

[1] Bruce Alberts, Alexander Johnson, Julian Lewis, et al. *Molecular Biology of the Cell*. Garland Science, New York, 4th edition, 2002. Available from: `https://www.ncbi.nlm.nih.gov/books/NBK21054/`.

[2] C. David Allis and Thomas Jenuwein. The molecular hallmarks of epigenetic control. *Nature Reviews Genetics*, 17:487–500, 2016.

[3] Thomas Graf and Matthias Stadtfeld. Heterogeneity of embryonic and adult stem cells. *Cell Stem Cell*, 3(5):480–483, 2008.

[4] Suzanne L. Ginn, Angelica K. Amaya, Ian E. Alexander, Michael Edelstein, and Mahshid R. Abedi. Gene therapy clinical trials worldwide to 2017: An update. *Journal of Gene Medicine*, 20(5):e3015, May 2018. Epub 2018 Apr 19. Erratum in: J Gene Med. 2019 Sep;21(9):e3124. doi: 10.1002/jgm.3124.

[5] Jianzhu Ma, Michael K. Yu, Stephen Fong, Kosuke Ono, Eric Sage, Barry Demchak, Roded Sharan, and Trey Ideker. Using deep learning to model the hierarchical structure and function of a cell. *Nature Methods*, 15(4):290–298, April 2018. Epub 2018 Mar 5.

[6] E. Seo, Y. N. Choi, Y. R. Shin, D. Kim, and J. W. Lee. Design of synthetic promoters for cyanobacteria with generative deep-learning model. *Nucleic Acids Research*, 51(13):7071–7082, July 2023.

[7] James Zou, Maryam Huss, Aakrosh Abid, Pooya Mohammadi, Ali Torkamani, and Amal Telenti. A primer on deep learning in genomics. *Nature Genetics*, 51(1):12–18, January 2019. Epub 2018 Nov 26.

[8] Wen-Chi W. Chen, Leonid Gaidukov, Yu-Ting Lai, et al. A synthetic transcription platform for programmable gene expression in mammalian cells. *Nature Communications*, 13:6167, 2022.

[9] J. Uhlendorf, A. Miermont, T. Delaveau, G. Charvin, F. Fages, S. Bottani, G. Batt, and P. Hersen. Long-term model predictive control of gene expression at the population and single-cell levels. *Proceedings of the National Academy of Sciences of the United States of America*, 109(35):14271–14276, 2012.

[10] Tianyu Lu, Hua Liang, Haihong Li, and Hulin Wu. High dimensional ODEs coupled with mixed-effects modeling techniques for dynamic gene regulatory network identification. *Journal of the American Statistical Association*, 106(496):1242–1258, 2011.

[11] Andreas Milias-Argeitis, Marc Rullan, S. K. Aoki, Peter Buchmann, and Mustafa Khammash. Automated optogenetic feedback control for precise and robust regulation of gene expression and cell growth. *Nature Communications*, 7:12546, August 2016.

[12] Domitilla Del Vecchio, Aaron J. Dy, and Yili Qian. Control theory meets synthetic biology. *Journal of The Royal Society Interface*, 13(120):20160380, 2016.

[13] Steven A. Benner and Amy M. Sismour. Synthetic biology. *Nature Reviews Genetics*, 6:533–543, 2005.

[14] D. Ewen Cameron, Caleb J. Bashor, and James J. Collins. A brief history of synthetic biology. *Nature Reviews Microbiology*, 12(5):381–390, 2014.

[15] D. Chandran, W. B. Copeland, S. C. Sleight, and H. M. Sauro. Mathematical modeling and synthetic biology. *Drug Discovery Today: Disease Models*, 5(4):299–309, 2008.

[16] Alec A. K. Nielsen et al. Genetic circuit design automation. *Science*, 352:aac7341, 2016.

[17] Jennifer A. Brophy and Christopher A. Voigt. Principles of genetic circuit design. *Nature Methods*, 11(5):508–520, May 2014.

[18] M. Gossen and H. Bujard. Tight control of gene expression in mammalian cells by tetracycline-responsive promoters. *Proceedings of the National Academy of Sciences of the United States of America*, 89(12):5547–5551, June 1992.

[19] M. Santorelli et al. Reconstitution of an ultradian oscillator in mammalian cells by a synthetic biology approach. *ACS Synthetic Biology*, 7:1447–1455, 2018.

[20] L. Postiglione, S. Napolitano, E. Pedone, D. L. Di Rocca, F. Aulicino, M. Santorelli, B. Tumaini, L. Marucci, and D. di Bernardo. Regulation of gene expression and signaling pathway activity in mammalian cells by automated microfluidics feedback control. *ACS Synthetic Biology*, 7(11):2558–2565, November 2018. Epub 2018 Oct 22.

[21] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

[22] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, et al. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.

[23] Vijay Konda and John Tsitsiklis. Actor-critic algorithms. In S. Solla, T. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.

[24] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[25] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[26] Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J. Zico Kolter. Differentiable mpc for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, 2018.

[27] Neythen J. Treloar, Alex J. H. Fedorec, Brian Ingalls, and Chris P. Barnes. Deep reinforcement learning for the control of microbial co-cultures in bioreactors. *PLOS Computational Biology*, 16:1–18, 04 2020.

[28] Jong Woo Kim, Byung Jun Park, Tae Hoon Oh, and Jong Min Lee. Model-based reinforcement learning and predictive control for two-stage optimal control of fed-batch bioreactor. *Computers & Chemical Engineering*, 154:107465, 2021.

[29] B. Jaganatha Pandian and Mathew Mithra Noel. Control of a bioreactor using a new partially supervised reinforcement learning algorithm. *Journal of Process Control*, 69:16–29, 2018.

[30] P. Petsagkourakis, I.O. Sandoval, E. Bradford, D. Zhang, and E.A. del Rio-Chanona. Reinforcement learning for batch bioprocess optimization. *Computers & Chemical Engineering*, 133:106649, 2020.

[31] Jean-Baptiste Lugagne, Caroline M. Blassick, and Mary J. Dunlop. Deep model predictive control of gene expression in thousands of single cells. *bioRxiv*, 2022.

[32] R. Chait, J. Ruess, T. Bergmiller, et al. Shaping bacterial population behavior through computer-interfaced control of individual cells. *Nature Communications*, 8:1535, 2017.

[33] Chiara Fracassi, Lorena Postiglione, Gianfranco Fiore, and Diego di Bernardo. Automatic control of gene expression in mammalian cells. *ACS Synthetic Biology*, 5(4):296–302, 2016. PMID: 26414746.

[34] Lorena Postiglione, Sara Napolitano, Elisa Pedone, Daniel L. Rocca, Francesco Aulicino, Marco Santorelli, Barbara Tumaini, Lucia Marucci, and

Diego di Bernardo. Regulation of gene expression and signaling pathway activity in mammalian cells by automated microfluidics feedback control. *ACS Synthetic Biology*, 7(11):2558–2565, 2018.

[35] Zachary Chase Lipton. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*, 2015.

[36] Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps, 2017.

[37] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, et al. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8:6085, 2018.

[38] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1310–1318. JMLR.org, 2013.

[39] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015.

[40] Martin Kolnik, Lev S. Tsimring, and Jeff Hasty. Vacuum-assisted cell loading enables shear-free mammalian microfluidic culture. *Lab on a Chip*, 12(22):4732–4737, 2012.