

---

# Differentiable Physics-based Greenhouse Simulation

---

**Nhat M. Nguyen**  
Koidra.ai  
nhat.nguyen@koidra.ai

**Hieu T. Tran**  
Koidra.ai  
hieu.tran@koidra.ai

**Minh V. Duong**  
Koidra.ai  
minh.duong@koidra.ai

**Hanh Bui**  
Koidra.ai  
hanh.bui@koidra.ai

**Kenneth Tran**  
Koidra.ai  
ken@koidra.ai

## Abstract

We present a differentiable greenhouse simulation model based on physical processes whose parameters can be obtained by training from real data. The physics-based simulation model is fully interpretable and is able to do state prediction for both climate and crop dynamics in the greenhouse over very a long time horizon. The model works by constructing a system of linear differential equations and solving them to obtain the next state. We propose a procedure to solve the differential equations, handle the problem of missing unobservable states in the data, and train the model efficiently. Our experiment shows that the procedure is effective. The model improves significantly after training and can simulate a greenhouse that grows cucumbers accurately.

## 1 Introduction

There has been a growing interest in applying machine learning in various agricultural domains and such attempts have achieved multiple successes to allow for more efficient and better farming ([7], [14]). With the world population still growing, food security remains to be one of the most important issues for many countries and machine learning applications in agriculture and horticulture can help alleviate this issue. One of such applications is to use machine learning to improve optimal control in greenhouses. For tuning and learning of the optimal controllers in the greenhouses, there is a need to use simulators, as it is prohibitively expensive to do so in the real systems. The most prominent and successful class of models to simulate the greenhouses are based on physical processes. Physics-based greenhouse models are well-studied ([12], [6], [13]) and have been used for optimal greenhouse control, for example, in [1]. However, those models are usually fixed, and their parameters have to be hand-tuned for each greenhouse. This limits the rapid deployment and adaptation of such models in commercial settings. On the other hand, differentiable physics ([10], [11]) is an emerging interdisciplinary paradigm where automatic differentiation [8] is applied in the numerical simulation of physical systems and has the potential to facilitate automatic learning of physical systems parameters. Furthermore, many optimal control algorithms, for example, the Direct Multiple Shooting method [2], require not only an accurate model of the dynamics but also the Jacobian and Hessian of the dynamic with respect to the state vector. Thus, there is a need to develop an accurate and differentiable physics-based greenhouse simulation implementation that can be trained end-to-end in order to address those challenges and allow better adoption of machine learning in greenhouse optimal control. This work presents such a greenhouse model implementation and proposes an efficient training procedure to address various problems when training the model.

The contributions of this paper are as follow: **(1)** We present a fully interpretable and differentiable physics-based greenhouse model implementation that can be trained end-to-end and can be used for accurate simulation of the greenhouse and for optimal greenhouse control. **(2)** We propose an

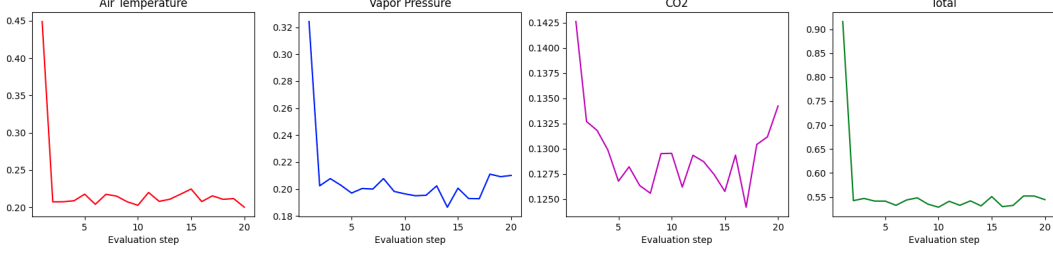


Figure 1: Individual loss components and total loss of the climate states on Croperator dataset.

efficient and numerically stable way to solve the linear system of differential equations for greenhouse dynamics, which can be used for similar simulation systems that use differential equations. **(3)** We propose an efficient and effective training procedure to handle difficulties in training long sequential greenhouse data. **(4)** As a proof of concept, we show some good preliminary results in predicting the cumulative weight and number of fruits harvested in a greenhouse that grows cucumbers.

## 2 Physics-based Greenhouse Simulation Model

### 2.1 Greenhouse simulation model formulation

Let  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{u}_i) | i = 0..N\}$  be the sequential trajectory of a greenhouse with length  $N$  in a single growing season.  $\mathbf{x}_i$  is the  $D_x$ -dimension state vector and  $\mathbf{u}_i$  is the  $D_u$ -dimension control vector at time step  $t$ . The time steps are divided equally in time, with  $\Delta t$  being the interval between time steps.

The state vector  $\mathbf{x}_i$  includes both climate states and crop states. Climate states are states that are related or contribute to the climate dynamic in the greenhouse. For example: air temperature, air humidity, how much CO2 is in the air, *etc.* Crop states are states that are related to the plants. For example: average plant temperature in the last 24 hours, carbohydrate stored in the leaf, fruits, *etc.* However, many states are unobservable, either due to the lack of sensors or due to the impossibility of measuring them. For example, a greenhouse may not have sensors to measure the temperature of the plant canopy, and it is considered an unobservable state. Another example is that it is not possible to quantify the amount of carbohydrate in the plant stem without causing harm to the plant and this also an unobservable state. The unobservable states are not available in the data collection process but may contribute significantly to the dynamics of a greenhouse. To accurately simulate the dynamics of a greenhouse, it is important to handle the missing states. A process for handling missing states during training is described in the experiment section 3.2. The control vector  $\mathbf{u}_i$  also includes control for both climate and crop. For example, setpoints of the heating pipes, setpoints of the fogging system, or harvesting action of the fruits.

A greenhouse simulator  $\mathcal{S}_\theta$ , parameterized by the parameter vector  $\theta$ , calculates a prediction of  $T$  next states  $\mathbf{x}_{i+1}, \dots, \mathbf{x}_{i+T}$  given the starting state  $\mathbf{x}_i$  and the control vectors  $\mathbf{u}_i, \dots, \mathbf{u}_{i+T-1}$  using the recursive formulation:  $\mathbf{x}_{k+1} = \mathcal{S}_\theta(\mathbf{x}_k, \mathbf{u}_k)$  for  $k = i \rightarrow i + T - 1$ . The physics-based greenhouse simulation model predicts the next state  $\mathbf{x}_{k+1}$  by solving a system of linear differential equations:

$$\dot{\mathbf{x}} = A_\theta(\mathbf{x}, \mathbf{u})\mathbf{x} + b_\theta(\mathbf{x}, \mathbf{u}) \quad (1)$$

where  $\dot{\mathbf{x}}$  is the derivative of the state vector  $\mathbf{x}$  over time.  $A_\theta(\mathbf{x}, \mathbf{u})$  is a  $D_x \times D_x$  matrix and  $b_\theta(\mathbf{x}, \mathbf{u})$  is a  $D_x$  dimensional vector. Both terms are functions of the state vector  $\mathbf{x}$  and the control vector  $\mathbf{u}$ , and are parameterized by  $\theta$ . The parameter vector  $\theta$  consists of parameters that affect the physical and biological processes in the greenhouse. Such parameterization scheme is fully interpretable while also allows human experts to contribute to the process of tuning the model. Examples of physical processes parameters include the height of the greenhouse, the normalized area of the ventilation, parameters of greenhouse screens, the capacity of the fogging system, *etc.* The biological processes parameters depend on the type of crop being grown. For vine crops (*e.g.* cucumber, tomato, *etc.*), examples are maximum leaf area index, light absorption coefficient of the canopy, *etc.*

The next state  $\mathbf{x}_{k+1}$  is obtained by substituting  $\mathbf{x}_k$  into (1) and solve it for time step  $\Delta t$ . Detail how to solve (1) is outlined in 2.2. The outstanding task of the simulator is to find  $A_\theta(\mathbf{x}, \mathbf{u})$  and

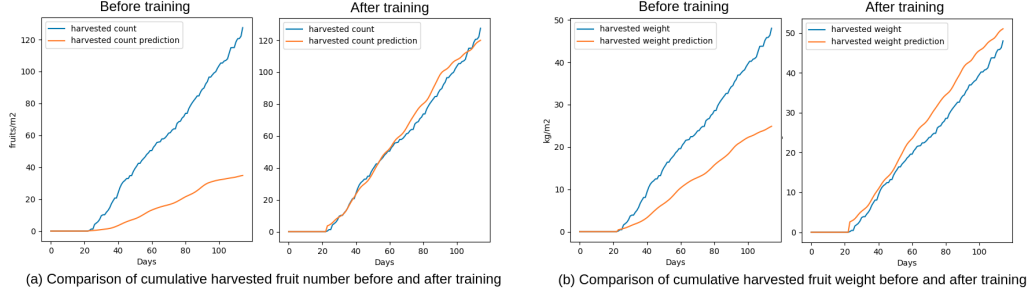


Figure 2: Crop prediction performance before and after training on the Croperators dataset. Blue: Ground truth. Orange: Our physics-based simulator prediction. Left: Cumulative harvested fruit number. Right: Cumulative harvested fruit weight.

$b_\theta(\mathbf{x}, \mathbf{u})$ . To construct them, we follow the physical and biological processes of vine crops described in [12]. The parameters in  $\theta$  are also obtained from [12]. To make our simulator differentiable, we implemented both the construction of the terms in (1) as well as the solving of (1) in Pytorch [9].

The training of our simulator reduces to finding the best values for the parameters in  $\theta$ . In practice, those parameters can only take values in physically feasible ranges. Therefore, we further formulate the parameters in  $\theta$  using the sigmoid function  $\sigma$ :  $\theta_j = \theta_j^{min} + \sigma(\theta_j^*) (\theta_j^{max} - \theta_j^{min})$ . This constrains  $\theta_j \in (\theta_j^{min}, \theta_j^{max})$ . In this paper, we consider the greenhouses where cucumber is grown.  $(\theta_j^{min}, \theta_j^{max})$  are chosen to be from 0.5 to 2.0 times the value of each parameter in [12] for parameters of the climate processes and in [6] for parameters of the cucumber crop model.

## 2.2 Solving non-homogeneous linear system of differential equations

Suppose that  $\mathbf{x}$  follows the dynamics  $\dot{\mathbf{x}} = A\mathbf{x} + b$  where  $A$  and  $b$  are constants. Denote  $\mathbf{x}(t)$  as the state vector at time  $t$ . We want to find  $\mathbf{x}(t+h)$  given  $\mathbf{x}(t)$ ,  $h$  is the time step. The non-homogeneous linear system of differential equations has a closed form solution:

$$\mathbf{x}(t+h) = e^{hA}\mathbf{x}(t) + \left( \int_0^h e^{(h-\tau)A} d\tau \right) b = e^{hA}\mathbf{x}(t) + \left( \int_0^h e^{\tau A} d\tau \right) b \quad (2)$$

To obtain the next state  $\mathbf{x}_{k+1}$  from  $\mathbf{x}_k$ , we substitute  $\mathbf{x}_k$  for  $\mathbf{x}(t)$ ,  $\mathbf{x}_{k+1}$  for  $\mathbf{x}(t+h)$  and set  $h = \Delta t$ . In our simulation model,  $A$  and  $b$  are functions of the state  $\mathbf{x}$  and the control  $\mathbf{u}$ , which change over time. When  $h$  is small enough,  $A$  and  $b$  can be treated as constants and we can use (2) for our simulation model. The first matrix exponential term in (2) can be approximated by Taylor expansion:

$$e^X = \sum_{j=0}^{\infty} \frac{(X)^j}{j!} = \sum_{j=0}^{\infty} C_j \text{ where } C_j = \frac{1}{j} \times C_{j-1} \times X \text{ for } j > 0 \text{ and } C_0 = 1 \quad (3)$$

For computational purposes, we follow this recursive formulation until the change is less than  $10^{-6}$  with a maximum of ten Taylor expansion terms.

The integral  $Q = \int_0^h e^{\tau A} d\tau = h \int_0^1 e^{\tau(hA)} d\tau$  in the second term of (2) can be computed exactly in a numerically stable fashion as follow. Let  $Z = \begin{bmatrix} hA & I \\ 0 & 0 \end{bmatrix}$  where  $I$  is the identity matrix and 0 are zero matrices so that  $Z$ 's dimension is twice that of  $A$ . The matrix exponential of  $Z$  can be proved, by Taylor expansion, to be  $e^Z = \begin{bmatrix} e^{hA} & \int_0^1 e^{\tau(hA)} d\tau \\ 0 & I \end{bmatrix}$ . Hence, we can obtain  $Q/h$  by computing the matrix exponential of  $Z$  using the same procedure in (3) and slice it.

### 3 Experiment

#### 3.1 Dataset

We used the public dataset <sup>1</sup> from the Autonomous Greenhouse Challenge 2018 [3] to train our simulator’s parameters and test the results. The original dataset contains data on greenhouse climate, irrigation, outdoor weather, control setpoints, crop management, and resource consumption. The dataset has a time step of  $\Delta t$  equal to five minutes. The collected data were from 4-month Hi-Power cucumber crops from six compartments of five participants (AiCU, Croperators, Deep Greens, iGrow, Sonoma) and a reference team, resulting in tens of thousands of data points total for each team. In this work, we trained on the AiCU dataset and validated on the Croperators dataset.

#### 3.2 Training details

A sketch of the general training procedure is available in Algorithm 1. We considered the tasks of predicting the following quantities in the state vector: **(1)** Climate states: air temperature, air vapor pressure, CO2 in the air. **(2)** Crop states: cumulative harvested fruit weight and cumulative harvested fruit number everyday. Denote  $\mathbf{x}_k^{Temp}$ ,  $\mathbf{x}_k^{VP}$ ,  $\mathbf{x}_k^{CO2}$ ,  $\mathbf{x}_k^{HW}$ ,  $\mathbf{x}_k^{HC}$  as the aforementioned quantities of a state  $\mathbf{x}_k$  in that order. Furthermore, denote  $\hat{\cdot}$  as the prediction of a state vector by the simulator. We consider the multiple-step prediction problem where the formula  $\mathbf{x}_{k+1} = \mathcal{S}_\theta(\mathbf{x}_k, \mathbf{u}_k)$  is applied repeatedly to obtain the multistep prediction  $\hat{\mathbf{x}}_{i+1}, \hat{\mathbf{x}}_{i+2}, \dots, \hat{\mathbf{x}}_{i+T}$  from the starting state  $\mathbf{x}_i$  and the known control vectors  $\mathbf{u}_i, \mathbf{u}_{i+1}, \dots, \mathbf{u}_{i+T-1}$ . In practice, however, we don’t have  $\mathbf{x}_i$  because many states are unobservable. To solve this issue, we replace the starting  $\mathbf{x}_i$  with  $\tilde{\mathbf{x}}_i$  which is a combination of available data of the true states  $\mathbf{x}_i$  and some approximation of the missing states computed by a reference model  $\mathcal{S}_{ref}$ . Exactly how to compute  $\tilde{\mathbf{x}}_i$  is described in section 3.3.

The simulation model was trained using gradient descent with backpropagation [4]. The Adam optimizer [5] with a learning rate of  $10^{-1}$  and betas equal to (0.9, 0.999) was used. The absolute values of the gradients were clipped to a maximum of 1.0. We separated the training processes of the climate states and crop states losses due to their differences in sampling scheme to compute the losses. For the climate states loss, we randomly sampled the starting state  $\tilde{\mathbf{x}}_i$ , used the simulator to compute the next  $T$  states prediction, then computed the mean squared error loss for the prediction trajectory:

$$\mathcal{L}_1(\theta, i) = \frac{1}{T} \sum_{k=i+1}^{i+T} \left[ \omega^{Temp} (\hat{\mathbf{x}}_k^{Temp} - \mathbf{x}_k^{Temp})^2 + \omega^{VP} (\hat{\mathbf{x}}_k^{VP} - \mathbf{x}_k^{VP})^2 + \omega^{CO2} (\hat{\mathbf{x}}_k^{CO2} - \mathbf{x}_k^{CO2})^2 \right] \quad (4)$$

where  $\omega^{Temp} = 10^{-1}$ ,  $\omega^{VP} = 5 * 10^{-6}$ ,  $\omega^{CO2} = 2 * 10^{-5}$ . We sampled multiple starting states at the same time for minibatch training of size 32. The prediction trajectory length  $T$  was 24 (2 hours).

The crop states training procedure is similar to that of the climate states, with a different loss:

$$\mathcal{L}_2(\theta, i) = \sum_k \left[ \omega^{HW} (\hat{\mathbf{x}}_k^{HW} - \mathbf{x}_k^{HW})^2 + \omega^{HC} (\hat{\mathbf{x}}_k^{HC} - \mathbf{x}_k^{HC})^2 \right] \quad (5)$$

where  $\omega^{HW} = \omega^{HC} = 1$ . Because the harvested fruit weight and harvested fruit numbers data are only available daily, the index  $i$  of the starting state was chosen to be start-of-day, the prediction length  $T$  was chosen so that the last timestep is end-of-day and the indices  $k$  were all the end-of-day timesteps in the prediction horizon. We used  $T = 2016$  (equal 7 days). The mini-batch size was 32.

#### 3.3 Compute the state approximation $\tilde{\mathbf{x}}_i$

At the start of the season, the missing states are initialized using other available state that is closest to it. For example, temperature of the attic air will be initialized using values of the general air temperature. We then use the reference greenhouse simulation model  $\mathcal{S}_{ref}$  to simulate the whole growing season using the known control vectors  $\mathbf{u}$ . This gives predictions of the missing states. The predictions of the missing states are then combined with the states available in the data to make  $\tilde{\mathbf{x}}_i$ . The intuition is that over a long time, the missing states should depend mostly on the control and not on their initialization. Because this computation is very expensive, we only do this occasionally during the training process.  $\mathcal{S}_{ref}$  is chosen to be the physics-based greenhouse simulation model but with an old parameters vector and is updated periodically to reflect the updated parameters.

<sup>1</sup>Available at this Github repository.

---

**Algorithm 1: General Training Procedure**

---

```
1 Initialize parameter vector  $\theta$  // Randomly in physically feasible ranges
2 for  $E$  iterations do
  //  $E=10$  in the experiment
  3 Set  $\mathcal{S}_{ref}$  to be the latest version of  $S_\theta$ 
  4 Simulate the whole growing season with the reference model  $\mathcal{S}_{ref}$  to compute all  $\tilde{x}_i$ 
  5 for  $L_{climate}$  iterations do
    //  $L_{climate}=50$  in the experiment
    6 Sample a random mini-batch of initial states  $\tilde{x}_i$  and control vectors  $u$ 
    7 Simulate using  $\mathcal{S}_{ref}, \tilde{x}_i, u$ , to compute the prediction trajectories
    8 Compute the loss function in 4
    9 Update  $\theta$  using backpropagation
  10 Set  $\mathcal{S}_{ref}$  to be the latest version of  $S_\theta$ 
  11 Simulate the whole growing season with the reference model  $\mathcal{S}_{ref}$  to compute  $\tilde{x}_i$ 
  12 for  $L_{crop}$  iterations do
    //  $L_{crop}=25$  in the experiment
    13 Sample a mini-batch of initial states  $\tilde{x}_i$  and control vectors  $u$  using the sampling scheme
    associated with the loss 5
    14 Simulate using  $\mathcal{S}_{ref}, \tilde{x}_i, u$ , to compute the prediction trajectories
    15 Compute the loss function in 5
    16 Update  $\theta$  using backpropagation
```

---

### 3.4 Result

We show some preliminary evaluation of our model in this paper. Figure 1 shows the component of the loss (4) on the Croperator validation dataset. All the individual climate states losses and the total loss trended down over time. Figure 2 shows the comparison between the ground truth and the prediction by our simulation model before and after training on the Croperator dataset. The loss terms in (5) decreased from over 2500 to 400 for harvested fruit numbers and from 155 to just over 20 for harvested fruit weight. Furthermore, for the prediction performance on the last day of the season (which matters the most for economic purpose), the relative errors decreased from 48% to 6% for predicting harvested fruit weight and from 73% to 6% for harvested fruit numbers. Hence, we can conclude that the training significantly improved the simulation model to have good performance.

## 4 Conclusion

We presented an interpretable and differentiable physics-based greenhouse simulation model and proposed an efficient inference and training procedure for it. The work has the potential to improve optimal control for greenhouses. Future directions are to experiment with more crop types and to exploit the fact that the crop dynamics changes slower than the climate dynamics to improve speed.

**Limitations of the work:** The work is still in progress and lacks rigorous experiments and further experiments on more crop types, such as tomato and lettuce. **Potential negative societal impacts:** Automation of greenhouse simulation and control may take away jobs from human experts. However, we argue that this will free up human labor capital. Moreover, experts are still needed to better tune the simulator models and the optimal controllers.

### Broader Impacts

Our work has the potential to facilitate improvement and allow faster deployment of optimal control algorithms for greenhouses. It is also a proof of concept that a physics-based greenhouse simulation model can be trained end-to-end to achieve good performance. Better automation of greenhouse control can help address the food security issue that many countries currently face. The procedure to solve the differential system of differential equations and the training procedure in this work can also be applied to similar systems.

## References

- [1] Wei-Han Chen, Neil S Mattson, and Fengqi You. Intelligent control and energy optimization in controlled environment agriculture via nonlinear model predictive control of semi-closed greenhouse. *Applied Energy*, 320:119334, 2022.
- [2] Moritz Diehl, Hans Georg Bock, Holger Diedam, and P-B Wieber. Fast direct multiple shooting algorithms for optimal robot control. In *Fast motions in biomechanics and robotics*, pages 65–93. Springer, 2006.
- [3] S. (Silke) Hemming, H.F. (Feije) de Zwart, A. (Anne) Elings, Isabella Righini, and A. (Anna) Petropoulou. Autonomous Greenhouse Challenge, First Edition (2018). 3 2019.
- [4] Henry J Kelley. Gradient theory of optimal flight paths. *Ars Journal*, 30(10):947–954, 1960.
- [5] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [6] LFM Marcelis. A simulation model for dry matter partitioning in cucumber. *Annals of botany*, 74(1):43–52, 1994.
- [7] Vishal Meshram, Kailas Patil, Vidula Meshram, Dinesh Hanchate, and SD Ramkteke. Machine learning in agriculture domain: A state-of-art survey. *Artificial Intelligence in the Life Sciences*, 1:100010, 2021.
- [8] Richard D Neidinger. Introduction to automatic differentiation and matlab object-oriented programming. *SIAM review*, 52(3):545–563, 2010.
- [9] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc., 2019.
- [10] Bharath Ramsundar, Dilip Krishnamurthy, and Venkatasubramanian Viswanathan. Differentiable physics: A position piece. *arXiv preprint arXiv:2109.07573*, 2021.
- [11] Nils Thuerey, Philipp Holl, Maximilian Mueller, Patrick Schnell, Felix Trost, and Kiwon Um. *Physics-based Deep Learning*. WWW, 2021.
- [12] Bram HE Vanthoor. *A model-based greenhouse design method*. Wageningen University and Research, 2011.
- [13] Colin Mark Wells. *Modelling the greenhouse environment and the growth of cucumbers (Cucumis sativus L.): a thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Agricultural Engineering at Massey University*. PhD thesis, Massey University, 1992.
- [14] Biyun Yang and Yong Xu. Applications of deep-learning approaches in horticultural research: a review. *Horticulture Research*, 8, 2021.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
  - (b) Did you describe the limitations of your work? **[Yes]** See section 4
  - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** See section 4
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**

2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
  - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The code is proprietary.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A] No GPUs was used. The model is very small and training can be done on CPU in a couple of hours.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes]
  - (b) Did you mention the license of the assets? [Yes]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]