
DEEMD: DRUG EFFICACY ESTIMATION AGAINST SARS-CoV-2 BASED ON CELL MORPHOLOGY WITH DEEP MULTIPLE INSTANCE LEARNING

A PREPRINT

M. Sadegh Saberian
School of Computing Science
Simon Fraser University
Burnaby, BC V5A 1S6
msaberia@sfu.ca

Kathleen P. Moriarty
School of Computing Science
Simon Fraser University
Burnaby, BC V5A 1S6
kmoriart@sfu.ca

Andrea D. Olmstead
Department of Microbiology and Immunology
Life Sciences Institute
University of British Columbia
Vancouver, BC V6T 1Z3
andrea.olmstead@ubc.ca

Ivan R. Nabi
Department of Cellular and Physiological Sciences
School of Biomedical Engineering
Life Sciences Institute
University of British Columbia
Vancouver, BC V6T 1Z3
irnabi@mail.ubc.ca

François Jean
Department of Microbiology and Immunology
Life Sciences Institute
University of British Columbia
Vancouver, BC V6T 1Z3
fjean@mail.ubc.ca

Maxwell W. Libbrecht
School of Computing Science
Simon Fraser University
Burnaby, BC V5A 1S6
maxwl@sfu.ca

Ghassan Hamarneh*
School of Computing Science
Simon Fraser University
Burnaby, BC V5A 1S6
hamarneh@sfu.ca

May 13, 2021

ABSTRACT

Drug repurposing can accelerate the identification of effective compounds for clinical use against SARS-CoV-2, with the advantage of pre-existing clinical safety data and an established supply chain. RNA viruses such as SARS-CoV-2 manipulate cellular pathways and induce reorganization of subcellular structures to support their life cycle. These morphological changes can be quantified using bioimaging techniques. In this work, we developed DEEMD: a computational pipeline using deep neural network models within a multiple instance learning (MIL) framework, to identify putative treatments effective against SARS-CoV-2 based on morphological analysis of the publicly available RxRx19a dataset. This dataset consists of fluorescence microscopy images of SARS-CoV-2 non-infected cells and infected cells, with and without drug treatment. DEEMD first extracts discriminative morphological features to generate cell morphological profiles from the non-infected and infected cells. These morphological profiles are then used in a statistical model to estimate the applied treatment efficacy on infected cells based on similarities to non-infected cells. DEEMD is capable of localizing infected cells via weak supervision without any expensive pixel-level annotations. DEEMD identifies known SARS-CoV-2 inhibitors, such as *Remdesivir* and *Aloxistatin*, supporting the validity of our approach. DEEMD is scalable to process and screen thousands of treatments in parallel and can be explored for other emerging viruses and datasets to rapidly identify candidate antiviral treatments in the future. Our implementation is available online¹.

¹github.com/Sadegh-Saberian/MIL-COVID19-RxRx19a

Keywords Drug Repurposing · Deep Multiple Instance Learning · Morphological Analysis · SARS-CoV-2

1 Introduction

The COVID-19 global pandemic has urged the research community to focus their resources towards studying SARS-CoV-2 and discovering or identifying potential therapeutics. To date, despite intense efforts, very few treatment options are available for those suffering from COVID-19 [1]. Drug repurposing is an attempt to identify existing clinically approved treatments with established pharmacological and safety profiles that could be rapidly redirected towards clinical treatment of novel diseases such as COVID-19 [2, 3] (and references within). The antiviral activity of candidate compounds can be tested using cell-based systems of viral infection. The detection of the viral infection is achieved using molecular tools such as antibodies directed at virus-encoded proteins. In the case of newly emerging viruses such as SARS-CoV-2, access to such molecular tools may represent an important limiting step to rapidly developing cell-based assays to discover novel antiviral agents. Alternatively, since human pathogenic viruses such as SARS-CoV-2 manipulate cellular pathways to reorganize the host cell morphology to support their life cycle [4, 5], developing a computational method to perform quantitative analysis of virus-induced cell morphology provides a unique approach to discover candidate antiviral molecules without the use of viral biomarkers [6]. Virus infected cells can be treated with thousands of compounds at different concentrations followed by staining of cellular structures with fluorescent dyes that can be imaged using high-content screening microscopes [7, 8]. Morphological features can be extracted from images of infected and non-infected cells and then applied to images of infected, drug-treated cells to predict antiviral efficacy based on cellular morphology. Thus, quantitative morphological analysis of cells as a computational method for drug repurposing allows for accelerated parallel screening of multiple therapeutics [9].

This paper is structured as follows: In Section 2 we review the literature on drug repurposing, high-throughput cell imaging, and computational multiple instance learning framework. In Section 3 we present DEEMD and its components and discuss them in detail. In Section 4 we describe the RxRx19a dataset and DEEMD implementation details. In Section 5 we discuss our findings along with DEEMD’s limitations. Finally, in Section 6 we conclude the paper.

2 Background

2.1 Drug Repurposing and High-throughput Imaging

Drug repurposing is an active research area in the pharmaceutical industry in which existing drugs are used in alternative applications for which they were originally designed and it has been shown to be an effective approach in many cases [2]. Various approaches are used to identify candidate drugs for repurposing applications including computational methods, i.e. genetic association or molecular docking, and experimental methods, i.e. phenotypic screening or binding assays. An additional approach is to apply high-throughput imaging assays for morphological cell profiling and to predict antiviral effectiveness. Morphological cell profiling combined with recent advancements in computer vision models can be leveraged to extract the morphological changes induced by each compound in a cell population. These morphological features can also provide information about how compounds may be interacting with the host cell, their molecular targets, and the affected cellular pathways. Image-based morphological analysis has been successfully utilized in small molecule profiling [6, 10], identifying drug mechanism of action [11, 12, 13], and in drug repurposing [14, 9].

Recently, many drug repurposing studies have focused on COVID-19 to find FDA-approved and clinically evaluated therapeutic compounds that exhibit antiviral activity against SARS-CoV-2 [15, 16] (and references within). Among multiple methodologies, high-throughput screens have enabled image-based morphological analysis of cells infected with SARS-CoV-2 [9, 17, 18, 19]. Mirabelli *et al.* [9] used Huh-7 cells infected with SARS-CoV-2 and treated with a library of 1425 FDA-approved compounds at different concentrations to identify compounds with antiviral activity. The samples were then probed for SARS-CoV-2 nucleocapsid protein and for dyes specific to cell organelles for imaging. CellProfiler software [20] was used to extract morphological features from the infected cells, using SARS-CoV-2 nucleocapsid as an indicator for regions of interest, and trained a random forest model to predict efficacy scores for each treatment. These scores were used to select efficacious compounds for follow-up experimental triplicate dose-response confirmation. They identified 17 effective compounds including *Remdesivir*, of which 10 are novel *in vitro* identifications.

Heiser *et al.* [17] used a morphological analysis pipeline on the RxRx19a dataset, a fluorescence microscopy dataset of Human Renal Cortical Epithelial (HRCE) cells, to identify potential treatments for SARS-CoV-2. They used a proprietary deep convolutional neural network to embed the sample images for calculating on-disease projection and off-disease rejection scores per dose for each treatment. These scores were then aggregated and normalized using a proprietary algorithm to compute the hit-scores. They reported that *Remdesivir*, *GS-441524*, *Aloxistatin*,

Silmitasertib, and *Almitrine* showed moderate to strong effectiveness against SARS-CoV-2 in their model, whereas, neither *chloroquine* nor *hydroxychloroquine* demonstrated any significant antiviral effectiveness. Similarly, Cuccarese *et al.* [18] used an *in vitro* deep-learning-driven analysis of cellular morphology on Human Umbilical Vein Endothelial Cells (HUVEC) treated with protein cocktails that mirror those from severe COVID-19 to identify potential drug repurposing candidates against COVID-19-associated cytokine storm.

2.2 Multiple Instance Learning

Multiple instance learning (MIL), as a form of weakly supervised learning, has been under the spotlight of research communities recently due to its ability to leverage weak supervision for tasks that are conventionally considered heavily dependent on laborious human annotations [21]. MIL was originally introduced for drug activity prediction [22] and recently it has been applied in many different domains such as computer vision [23], medical imaging and diagnosis [24], and in biology for applications such as mechanism of action classification using microscopy images [11], identifying antigen binding peptides [25], and predicting specific functional binding sites in microRNA targets [26]. Deep multiple instance learning uses different MIL approaches combined with a deep neural network model as a learner in the MIL framework, and it has shown performance competitive with state-of-the-art in recent studies [27, 28, 11, 29, 30].

Contrary to conventional supervised learning methods in which every instance in the training data is associated with a label, in the MIL the learner is provided with a training set of labeled bags, where each bag is a set of instances. The learner is tasked to predict a label for an unseen bag given its instances. Instance-space algorithms are a category of the MIL which assume that the discriminative information lies within instances and is local. Hence, they predict the bag label by training an instance classifier and then aggregating the instance-level predictions into the bag label [31]. The most commonly used assumption in these methods –the standard multiple instance assumption– states that in a binary bag classification set up, each bag labeled with the positive label contains at least one instance representing the positive class, whereas the negative class bags do not contain any instance of positive class.

Virus-specific markers are not always included in high-throughput imaging studies that investigate the effects of viral infection and drug treatments on host cell morphology; in these instances, the extent of viral infection is unknown. This scenario may also occur during the very early stages of an outbreak or pandemic when limited virus-specific tools are available, such as when COVID-19 first emerged. When no pixel-level annotations are available to specifically identify infected regions of the imaged cell population, MIL may be used as an alternative strategy to localize viral infection and the corresponding effects on host cell morphology. The imaged cell populations can be broken down into smaller instances, single cells or patches, and the sample cell population label is utilized in the MIL formulation. An instance-space approach would very well suit this problem for viral infections such as SARS-CoV-2, that induce local morphological changes to cells [5]. Based on the standard assumption, we assume that if a sample cell population is infected, then there should be at least one instance that contains infected cells. And a non-infected sample cell population does not have any infected cells.

2.3 Contribution

In this work, we present DEEMD: a computational pipeline that estimates treatment efficacy using morphological features of cells extracted using a deep learning model. DEEMD is scalable to process and screen multiple cell lines and thousands of treatments and can accelerate the identification of clinically evaluated and FDA-approved compounds with antiviral activity. We applied DEEMD on a public microscopy image dataset of SARS-CoV-2 infected cells, and it identifies treatment compounds that have been suggested in the literature to be effective against COVID-19, further supporting the validity of the pipeline. The specific design of the deep learning model allows us to localize infected cells via a weakly supervised training procedure without expensive annotations. Localizing the infected cells allows us to better explain the model predictions, as well as providing an opportunity to generate hypotheses about the biology behind the morphological changes in cells induced by SARS-CoV-2.

To the best of our knowledge, this is the first work to approach the drug repurposing problem via leveraging weak supervision through multiple instance learning. Some morphological analysis pipelines utilize a viral-specific marker to localize infected regions of interest in sample images for studying morphological changes induced by the virus, such as [9]. Another approach is to rely on computer vision models to extract viral-related morphological features through massive datasets without providing any annotations on infected regions of interest, such as [17]. DEEMD is an attempt to integrate both these approaches through identifying infected regions in sample images without requiring any viral-specific markers or annotations. DEEMD can be applied to quantify the effectiveness of any available compound and identify lead candidates that show therapeutic potential against viral infection. The identified candidates can then be

evaluated using other molecular tools once they become available, such as fluorescent-tagged viruses, nucleic acid-based methods (PCR), or animal models.

3 Methodology

We first describe a high-level overview of the DEEMD before delving into the details. DEEMD consists of two components: First component is a deep learning model which is trained on microscopy images of untreated samples to extract discriminative morphological features in MIL framework to distinguish between uninfected and SARS-CoV-2-infected cell population images, further described in Section 3.1. Following infection of cells treated with a drug, the trained deep learning model is used to extract a morphological profile for each sample. These profiles are then aggregated and used in a statistical model, the second component, to estimate the treatment efficacy per concentration, Section 3.4. Finally, all dose-dependent efficacy scores are aggregated for each treatment to form the final identified set of potentially effective treatments. In this context, we define treatment effectiveness as how close a treated sample’s morphological profile is to uninfected cell profiles. We discuss each component detail in the following sections. Table 1 includes all notations used in Section 3 for brevity.

3.1 Classification with deep MIL

In this work, we follow closely the training procedure described in [27, 29] for training an instance classifier based on a relaxed version of the standard multiple instance assumption. The training set \mathcal{D} consists of n data points (\mathbf{X}, \mathbf{Y}) in form of (x_i, y_i) pairs where x_i is a fluorescence microscopy image along with its associated sample-level label y_i with two classes: SARS-CoV-2 infected (c^+) and non-infected (c^-). Each sample image x_i can be split into N patches each referred to as $x_i^{p_j}$. We assume that each $x_i^{p_j}$ is associated with a patch-level label $y_i^{p_j}$ which is unknown and not included in the training set \mathcal{D} . Let $\mathcal{M}(\theta)$ be a deep neural network model responsible for predicting a patch-level label $\mu_i^{p_j}$ for every patch $x_i^{p_j}$ of each sample image x_i . In the MIL terminology, model $\mathcal{M}(\theta)$ is an instance classifier since it classifies each instance into SARS-CoV-2 infected and non-infected classes. We define $\mu_i^{p_j}$ to be the model $\mathcal{M}(\theta)$ estimation of the unknown true patch-level label $y_i^{p_j}$ given the dataset \mathcal{D} . This can be expressed mathematically as:

$$\mu_i^{p_j} \triangleq \mathbb{P}\left[y_i^{p_j} = c^+ \mid \mathcal{D}; \mathcal{M}(\theta)\right], \quad (1)$$

where $\mathbb{P}[A]$ shows the probability of the event A .

Each iteration within a training epoch starts with an exhaustive inference using the model from previous iteration over all patches in the training set and all corresponding $\mu_i^{p_j}$ are estimated. Then for each sample image x_i , all of its patch-level estimated labels $\mu_i^{p_j}$, denoted as M_i , are sorted and the set of k patches with the highest probability, denoted as \mathcal{K}_i^k , are selected for model training.

$$t_r \triangleq \left\{ j \mid \mu_i^{p_j} = m(M_i, r) \right\}, \quad \mathcal{K}_i^k \triangleq \bigcup_{r=1}^k t_r, \quad (2)$$

where we define $m(S, r)$ to be the r -th greatest element in the set S .

The sample-level label y_i is assigned to each set of patches in the set \mathcal{K}_i^k , and the model $\mathcal{M}(\theta)$ is trained to minimize the binary cross entropy loss, $\mathcal{L}(\cdot)$, between y_i and $\mu_i^{p_j}$ for all patches in the set \mathcal{K}_i^k . More specifically, we can break this process down into multiple steps:

Table 1: List of Notations

\mathcal{D}	Training set in form of (\mathbf{X}, \mathbf{Y})
n	Number of data points in \mathcal{D}
x_i	Fluorescence microscopy sample image
y_i	Sample-level label of x_i
c^+	SARS-CoV-2 infected class
c^-	Non-infected class
N	Number of patches in each x_i
$x_i^{p_j}$	j -th patch of x_i
$y_i^{p_j}$	Patch-level label of $x_i^{p_j}$
$\mu_i^{p_j}$	Predicted patch-level label of $x_i^{p_j}$
$\mathbb{P}[A]$	Probability of the event A
M_i	Set of all predicted patch-level labels for x_i
\mathcal{K}_i^k	Set of k -top patches with highest $\mu_i^{p_j}$ for x_i
$m(S, r)$	r -th greatest element in the set S
$\mathcal{L}(\cdot)$	Binary cross entropy loss
w^+, w^-	Class weights
\hat{y}_i	Predicted sample-level label for x_i
η	Cut-off threshold for sample classification
$\mathbb{I}(\cdot)$	Indicator function
A_i	Calculated infection map for sample x_i
$\mathcal{O}_{x_i}^{(l,m)}$	Set of patches overlapping at pixel $x_i^{(l,m)}$
z_i	Infection probability of x_i
t_i	i -th treatment
c_j	j -th concentration
$\mathcal{T}_{t_i}^{c_j}$	Set of all replicates treated with t_i at c_j
$\mathcal{C}_{t_i}^{c_j}$	Set of all concentrations of t_i
$e_{t_i}^{c_j}$	Dose-dependent efficacy score for t_i at c_j
$\beta_{t_i}^{c_j}$	Median point estimator for z_i ’s of $\mathcal{T}_{t_i}^{c_j}$
$\text{CI}(\cdot)$	Confidence interval an estimation
\sup	Supremum (least upper bound)
$T(\cdot)$	Descriptive statistic
ζ	Cut-off threshold for effective treatment
$\Omega(\omega)$	Model Ω with learned weights ω
$\mathcal{E}_{\Omega(\omega)}$	Set of effective identified treatments by $\Omega(\omega)$
$\mathcal{W}(\phi)$	Whole-image based model
$\mathcal{V}(\psi)$	Patch-based model
$\mathcal{M}(\theta)$	MIL deep neural network model

1. Given a patch $x_i^{p_j}$ and the label y_i associated with its corresponding sample image x_i , we can calculate binary cross entropy loss as follows:

$$\mathcal{L}(x_i^{p_j} | \mathcal{D}; \mathcal{M}(\theta)) = -y_i \log(\mu_i^{p_j}) - (1 - y_i) \log(1 - \mu_i^{p_j}). \quad (3)$$

2. We define the loss value for the sample image x_i to be only inclusive of informative patches, e.g. patches that we have confidence in their predicted infection probability. In other words, we only include patches which are in the set \mathcal{K}_i^k and take an average over them. Thus the value of the loss function for each sample image is defined as follows:

$$\mathcal{L}(x_i | \mathcal{D}; \mathcal{M}(\theta), w^+, w^-) = -\frac{1}{k} \sum_{j \in \mathcal{K}_i^k} \left[w^+ y_i \log(\mu_i^{p_j}) + w^- (1 - y_i) \log(1 - \mu_i^{p_j}) \right], \quad (4)$$

where w^+ and w^- are incorporated in the loss function $\mathcal{L}(\cdot | \mathcal{D}; \mathcal{M}(\theta))$ to account for class imbalance [32].

3. Finally the value of loss function over the whole dataset in each epoch is calculated as follows:

$$\mathcal{L}(\mathbf{X} | \mathcal{D}; \mathcal{M}(\theta), w^+, w^-) = -\frac{1}{n} \sum_{i=1}^n \mathcal{L}(x_i | \mathcal{D}; \mathcal{M}(\theta), w^+, w^-). \quad (5)$$

We will provide more details on how to tune hyper-parameter k in Section 5.4. The reader is encouraged to refer to Section 1 in the Supplementary Material for a detailed discussion on how hyper-parameter k affects the training dynamics for a MIL learner.

3.2 Inference with deep MIL

Since the classification labels are only available for the sample images, aggregation of the patch-level infection probabilities is required for performance evaluation and downstream analysis. For inference, each sample image is split into patches and a forward pass through the model. This provides us with the estimated infection probabilities, which are aggregated to form the sample-level label \hat{y}_i :

$$\hat{y}_i = \mathbb{I} \left(\min_{p_j \in \mathcal{K}_i^k} \{ \mu_i^{p_j} \} \geq \eta \right) = \mathbb{I} \left(m(M_i, k) \geq \eta \right), \quad (6)$$

where η is the cut-off threshold, selected based on the validation set and $\mathbb{I}(\cdot)$ is an indicator function.

For the MIL model $\mathcal{M}(\theta)$, hyper-parameter k is incorporated into the standard multiple instance assumption so that the model can generalize to a more relaxed constraint: a positive (SARS-CoV-2 infected) sample image should contain at least k positive patches to be considered positive and a sample image would be considered negative if less than k patches are predicted to be infected. This generalized assumption is mathematically expressed in equation (6).

3.3 MIL infection localization

Localization of the infected region is achieved using the estimated $\mu_i^{p_j}$. To form the infection map A_i all patches need to be aggregated. Based on the set of all patches that overlap at the pixel $x_i^{(l,m)}$, denoted as set $\mathcal{O}_{x_i}^{(l,m)}$, the value of infection map $A_i^{(l,m)}$ is calculated by the weighted average of its overlapping patches infection probabilities, $\{ \mu_i^{p_j} \}_{p_j \in \mathcal{O}_{x_i}^{(l,m)}}$. We opt to set sample weights based on the sample's values, similar to the quadratic mean, but using $(\mu_i^{p_j})^\alpha$, where $\alpha < 1$, instead of $\mu_i^{p_j}$. The infection map A_i is calculated as:

$$A_i^{(l,m)} = \frac{\sum_{p_j \in \mathcal{O}_{x_i}^{(l,m)}} (\mu_i^{p_j})^{1+\alpha}}{\sum_{p_j \in \mathcal{O}_{x_i}^{(l,m)}} (\mu_i^{p_j})^\alpha}. \quad (7)$$

By using this averaging method, with $\alpha < 1$ when all $\mu_i^{p_j}$'s are smaller than 1, the average is more sensitive to higher values, and smaller values have less diminishing power on the result compared to simple averaging. This property is well-suited for infection localization since patches with high infection probabilities do not get diminished by the adjacent overlapping low probability ones. Finally, prior to rendering the infection map, a low-pass Gaussian filter, $\sigma = 60$ pixels, is applied to make the infection map smoother. For generating the infection map A_i based on the MIL model, $\mathcal{M}(\theta)$ we set α to 0.2.

3.4 Treatment Efficacy Estimation

DEEMD estimates the treatment efficacy based on its trained MIL model that compares the morphological profile of drug-treated cells with those of untreated infected and uninfected cells. We assume that an effective treatment would prevent drastic infection-induced morphology changes in the cells, hence, the treated cells morphological profile would be similar to that of uninfected ones. A statistical model takes in the infection probabilities predicted by the MIL model and estimates the probability that a treatment is effective based on its morphological similarities to uninfected cells. For each sample image x_i , the infection probability z_i is calculated as:

$$z_i = \text{Median}\{\mu_i^{p_j} | p_j \in \mathcal{K}_i^k\}. \quad (8)$$

We define the set $\mathcal{T}_{t_i}^{c_j}$ such that it consists of infection probabilities for every sample image of a given treatment t_i at concentration c_j in the treated test set. We have to aggregate the infection probabilities, z_i , over the set $\mathcal{T}_{t_i}^{c_j}$ to minimize the inevitable variations that have occurred during the sample preparation, treatment administration and image acquisition. We define $e_{t_i}^{c_j}$ to be the dose-dependent estimated efficacy score of the treatment t_i at concentration c_j based on any model $\Omega(\omega)$ capable of assigning an infection probability to a sample image based its morphology. Since we do not have any assumptions on the distribution of infection probabilities z_i in the set $\mathcal{T}_{t_i}^{c_j}$, we opt to use non-parametric statistics related to the median of this distribution, $\beta_{t_i}^{c_j}$. We observed that the distribution infection probabilities z_i in the set $\mathcal{T}_{t_i}^{c_j}$ for any t_i is heavily skewed and asymmetric, thus using Wilcoxon test is misleading since it assumes a symmetric distribution for the data[33]. However, the sign test is still valid[34]. Given a confidence level a , we can calculate the exact confidence interval for the estimated median, denoted as $\text{CI}(\beta_{t_i}^{c_j} | a, \Omega(\omega))$. To be more conservative about the false positive rate, instead of using the point estimate $\beta_{t_i}^{c_j}$, we use the least upper bound, supremum, of $\text{CI}(\beta_{t_i}^{c_j} | a, \Omega(\omega))$. The mathematical expression for the dose-dependent estimated efficacy score of the treatment t_i at concentration c_j is:

$$e_{t_i}^{c_j} \triangleq T(\mathcal{T}_{t_i}^{c_j} | \Omega(\omega)) = 1 - \sup\{\text{CI}(\beta_{t_i}^{c_j} | a, \Omega(\omega))\}, \quad (9)$$

where $T(\cdot)$ is a descriptive statistic. The dose-dependent efficacy score $e_{t_i}^{c_j}$ reflects the model $\Omega(\omega)$ belief on the morphological similarity between the sample images in the set \mathcal{T}_{t_i, c_j} and the uninfected ones in the training set. Close to 1 values for $e_{t_i}^{c_j}$ imply that the treatment t_i at concentration c_j is effective against SARS-CoV-2. In this context, we assume that the morphological profile of cells treated with an effective treatment is more similar to the uninfected cells rather than infected ones.

To summarise all estimated dose-dependant efficacy scores of a given treatment t_i for ranking all treatment compounds and identifying the effective ones, the estimated efficacy score e_{t_i} is calculated:

$$\tilde{e}_{t_i}^{c_j} = \begin{cases} \{e_{t_i}^{c_j} | e_{t_i}^{c_j} \geq \zeta\}_{c_j \in \mathcal{C}_{t_i}} & \exists c_j : e_{t_i}^{c_j} \geq \zeta \\ \{e_{t_i}^{c_j}\}_{c_j \in \mathcal{C}_{t_i}} & \text{otherwise} \end{cases}, \quad e_{t_i} = \text{Median}\{\tilde{e}_{t_i}^{c_j}\}, \quad (10)$$

where ζ is a cut-off threshold and \mathcal{C}_{t_i} refers to all available concentrations of treatment t_i in the treated test set. The treatment compounds that DEEMD identifies as effective against SARS-CoV-2 are those for which at least one dose-dependent efficacy score is higher than ζ or equivalently: $\mathcal{E}_{\mathcal{M}(\theta)} = \{t_i | e_{t_i} \geq \zeta\}$.

4 Experimental Design and Implementation Details

4.1 Dataset

We are using the publicly available RxRx19a dataset, which is the first morphological imaging dataset of cells infected with SARS-CoV-2 [17]. The dataset consists of HRCE and African green monkey kidney epithelial (Vero) cell lines subjected to three different conditions: 1) mock non-infected control, 2) infection with ultraviolet (UV)-light inactivated SARS-CoV-2 (irradiated for 20 minutes), or 3) infection with SARS-CoV-2 at multiplicity of infection (MOI) of 0.4. All samples were incubated for 96 hours. A library of 1672 small molecules and FDA-approved treatments were applied to a subset of SARS-CoV-2 infected samples in 6+ half log concentrations with six replicates per dose for each compound. All cells were stained with five fluorescent dyes detecting various subcellular structures, each imaged in a separate channel on a fluorescent confocal high-content imaging microscope. The dyes include Hoechst (nucleus), Syto14 (nucleoli and cytoplasmic RNA), phalloidin (actin cytoskeleton), Concanavalin A (ConA; endoplasmic reticulum) and Wheat Germ Agglutinin (WGA; Golgi and plasma membrane). The dataset consists of more than 300K 5-channel labeled images of size 1024×1024 and is publicly available through the *Recursion* website: www.rxxx.ai/rxxx19a.

For this work, we only used the HRCE cell line and merged the mock control and the UV inactivated control into a single non-infected class, as no cytopathic effects were observed in either condition [17]. We also verified this by using

each sample image’s corresponding deep learning embeddings included in the dataset. We decided to merge mock control and UV-inactivated into a single class as they were indistinguishable by a trained classifier.

4.2 Cell nucleus count in the sample images

RxRx19a dataset includes Hoechst stain which binds double stranded DNA in the nucleus. We used this stain to identify and count the cells in each sample image for both explanatory analysis and preprocessing the data. We first stitch images of 4 adjacent sites of each well on each plate to reconstruct the image of whole sample cell population in each well. This step is necessary to avoid repeat counting of nuclei in multiple images. We localize and count the nuclei in each stitched image using a segmenting pipeline based on Otsu thresholding and watershed algorithm [35]. We used this pipeline to exclude any sample images which did not contain any detectable cells from the dataset. The distribution of the nuclei count in the sample images for infected and non-infected samples are shown in Fig 1. The cell nucleus count is significantly lower in the infected samples (using two-sided Mann-Whitney test with $p = 0.0$).

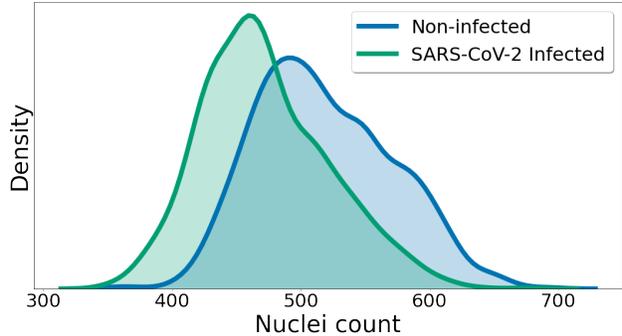


Figure 1: Distribution of cell nucleus count in the stitched well sample images from non-infected and SARS-CoV-2 infected HRCE cells. The non-infected class consists of samples from both Mock and UV Inactivated SARS-CoV-2 classes.

4.3 Baseline Classification Models

To compare the MIL model $\mathcal{M}(\theta)$ performance, two other deep learning models are trained: **1)** A model that is trained with the conventional training procedure for training convolutional neural networks by minimizing a cross-entropy-based loss function, similar to equation (4), with full resolution sample images as input. Denoted as $\mathcal{W}(\phi)$, this model will be referred to as the whole-image based model $\mathcal{W}(\phi)$ throughout the paper. Unlike most deep neural network classifiers, the input sample image is not down-sampled, and the full resolution version was used to make the comparison fair. The whole-image based model $\mathcal{W}(\phi)$ takes in the sample image x_i and calculates infection probability for that cell population.

2) Instead of limiting the model to only use k instances with the highest probabilities within each bag for training, we can have a reasonable alternative approach of using all instances in the bag and assigning them with the bag-level label. There are studies in histopathology such as [36] that used this alternative approach and their proposed model outperformed other models at their respective tasks. We refer to this model as the patch-based model $\mathcal{V}(\psi)$ in the text. Similar to MIL instance classifier, $\mathcal{V}(\psi)$ calculates a patch infection probability for all patches in the input sample image x_i . Intuitively, in this approach the model would potentially be able to eliminate the effect of noisy labels and estimate the true distribution of the instance labels given large enough training data. This approach is an extreme case of the MIL training procedure in which $k = N$ is where N is the total number of instances within a bag.

4.4 Implementation Details

The HRCE sample images in the RxRx19a dataset are split into 4 non-overlapping sets: **1)** training, **2)** validation, **3)** untreated, and **4)** treated test sets consisting of 20K, 5K, 6K, and 247K sample images respectively. The training set is used for updating the weights and training the neural networks using backpropagation, whereas the validation set is used to tune hyper-parameters. The performance of the tuned model is measured on the held-out untreated test set, while the treated test set is solely used for estimating the efficacy of the treatments. All sets are balanced in terms of class labels except for the treated test set which only consists of SARS-CoV-2 infected sample images.

All models are based on ResNet34 architecture, pretrained on ImageNet with a modified input layer to accommodate 5-channel inputs [37, 38]. Each model was trained for 150 epochs, monitored for early stopping, with Adam optimizer at a learning rate $\eta = 10^{-4}$, $\beta_1 = 0.9$ and $\beta_2 = 0.999$, [39], with a batch size of 128. All hyper-parameters were tuned with respect to the model’s performance on the validation set. Sample images are channel normalized using the empirical means and standard deviations calculated on the training set. We used *PyTorch* framework [40] for model training and evaluation and treatment efficacy estimation was coded using *R* [41].

For evaluating the performance of the models on both validation and untreated test set, we used $\eta = 0.5$. For the treatment efficacy estimation, ζ was set to 0.5, inspired by the concept of effective concentration 50 (EC_{50}) in dose-

response curves, and a confidence level of 0.95% was used in the sign test for estimating the confidence interval for the point estimated median used in the efficacy score estimations.

For both the MIL model $\mathcal{M}(\theta)$ and the patch-based model $\mathcal{V}(\psi)$, we used a uniform grid of 256×256 patches with 50% overlap, resulting in 49 patches per sample image. Since we wanted DEEMD to be applicable to other microscopy imaging datasets with minor modifications, we choose a uniform grid of patches over a cell-based segmentation map. Although a cell-based segmentation map could potentially provide a single cell resolution infection map, it is heavily dependent on the stains used in the imaging procedure, which in turn reduces the pipeline generalization. Based on the validation set, we set the hyper-parameter k for the MIL model $\mathcal{M}(\theta)$ to 2. The whole-image based model $\mathcal{W}(\phi)$ was trained on full-resolution (1024×1024) sample images without any input downsampling.

We trained the models on a single Nvidia TITAN V GPU. Notice that the effective training set for the MIL model $\mathcal{M}(\theta)$ on each training iteration is k/N of the whole training samples, thus its training is N/k times faster compared to the patch-based model $\mathcal{V}(\psi)$. The patch-based model $\mathcal{V}(\psi)$ even has longer training time than the whole-image based model $\mathcal{W}(\phi)$ since the patches are overlapping. Thus, the MIL model $\mathcal{M}(\theta)$ is computationally more efficient compared to the other two.

In the early stages of the training a deep learning model within MIL framework, the model has not learned the concepts of positive and negative instances from the training samples. Thus, the selected set \mathcal{K}_i^k , populated based on the the initial exhaustive inference at the beginning of each iteration, consists of randomly selected patches from each image. To prevent this random selection process from heavily affecting the training procedure and avoiding any drastic random changes in the weights, it is important to control the learning rate of the model. To this end, we used an one-cycle policy based on cosine annealing, [42], for the learning rate during the training to make sure that the network gradually learns to distinguish between the positive and negative samples. Using this approach, resulted in a stable training procedure with consistently comparable trained models.

4.5 Selecting the optimal value for hyper-parameter k

Choosing an optimal value for hyper-parameter k is crucial to the performance of MIL models since it effectively controls the amount of noise in the training labels. The hyper-parameter k is the number of top high probability instances in each bag used for loss calculation in each training iteration, equation (4). Tuning this hyper-parameter can be either based on problem-specific domain knowledge or a performance metric on the validation set. For the problem of estimating the treatment efficacy against SARS-CoV-2, the choice of optimal k was found based on incorporation of domain knowledge and common model performance metrics.

In a simple yet reasonable modeling, let’s assume that the virions enter the cells independent of each other and X is a random variable that counts the number of virions entering a cell. Since we are dealing with a counting process, it is reasonable to assume that random variable X comes from a Poisson distribution with parameter m , which is actually the multiplicity of infection or MOI $X \sim \text{Poisson}(m)$. The probability of a cell being infected is calculated as follows:

$$\mathbb{P}[\text{A cell is infected}] = 1 - \mathbb{P}[X = 0] = 1 - e^{-m}. \quad (11)$$

The probability of a cell being infected is shown in Fig 2-(a) as a function of MOI. Based on the data provided in [17], in the preparation process of the the RxRx19a dataset, the samples were infected with a MOI of 0.4. Based on our modeling of the infection process, this value would indicate that around 33% of the cells in each sample are infected on average.

We incorporate this domain knowledge about the expected number of infected cells within each sample image in selecting the optimal value for hyper parameter k . Since we are using a uniform grid of overlapping patches for the sample images, the number of the selected patches in the set \mathcal{K}_i^k can not be directly translated into the fraction of the infected cells within each sample image. The area of the infected regions based on the infection maps can be a good approximation for the faction of the infected imaged cells. We trained multiple MIL models with different values for the hyper-parameter $k \in \{1, 2, 3, 5, 10, 15, 25, 49\}$. Next, we calculated the fraction of pixels that were higher than the cutoff threshold η in the generated infection map for each sample images in the validation set for all of the MIL models, as shown in Fig 2-(c). We can see that when k equals to 1 or 2 or 3, the average of the distribution is close to the theoretical calculated value using the MOI. Hence, these 3 options are candidates for the optimal value for k . Finally, to select the optimal value between these 3 candidates, we turn to the average precision metric. As shown in Fig 2-(b), we can see that the models have a very close AUC, but when K is 2 the performance is slightly better and thus we set k to be 2.

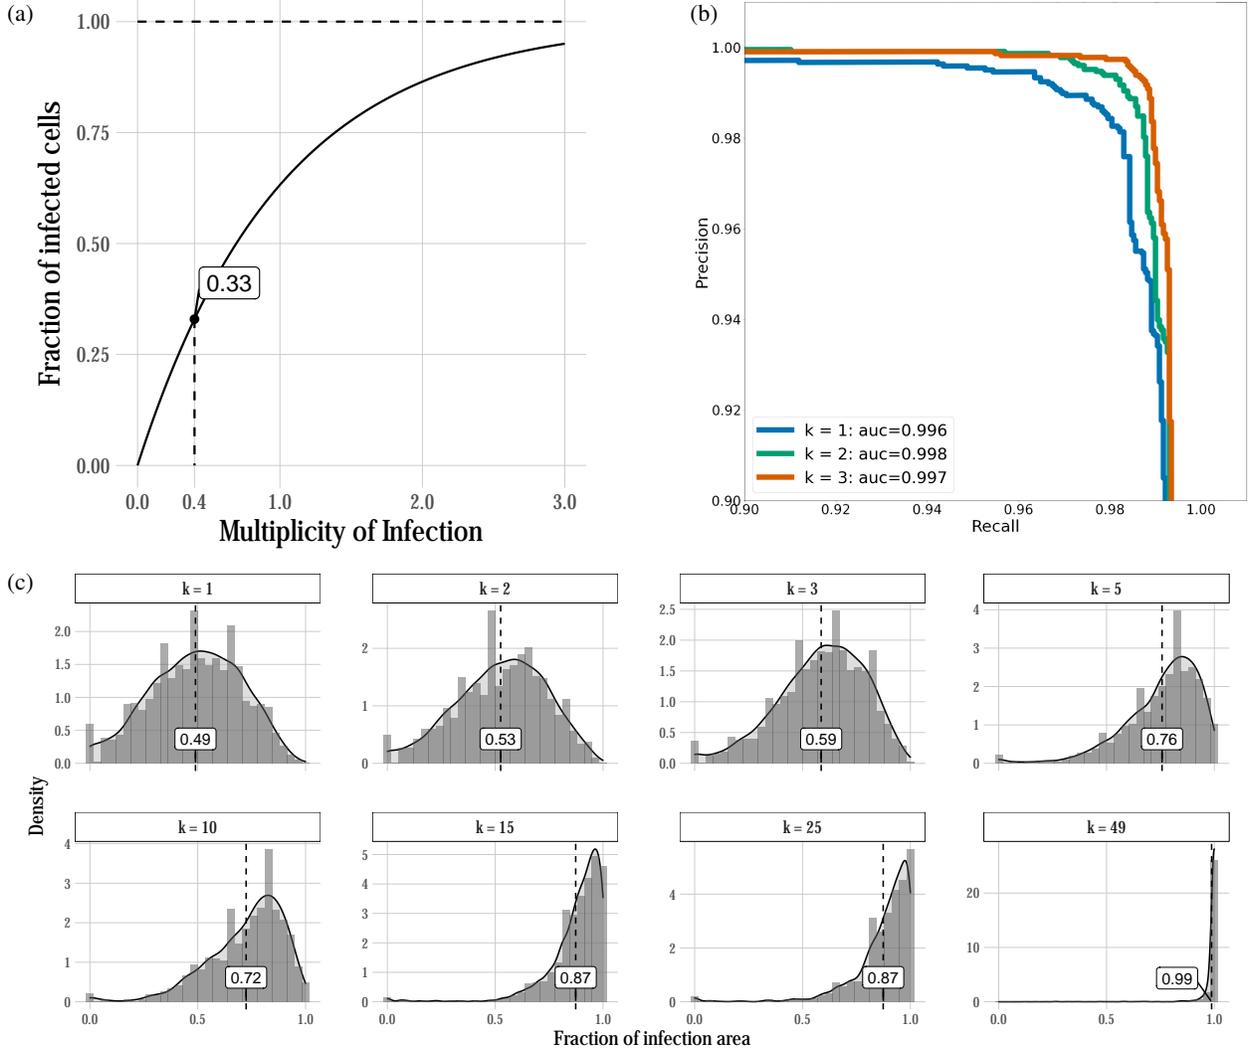


Figure 2: Selecting optimal value for hyper-parameter k . (a) Cell infection probability as a function of MOI in a simple modeling. (b) Zoomed Precision-Recall curve for candidate optimal values of k on the validation set. (c) The distributions of the fraction of infected pixels in each sample image in the validation set for multiple values of k , along with the mean of the distribution.

5 Results and Discussion

5.1 The MIL model can accurately predict SARS-CoV-2 infection

We evaluate the performance of the MIL model $\mathcal{M}(\theta)$, along with two baseline models, on the untreated test set in terms of the area under the curve for precision-recall curves, or average precision, as shown in Fig 3. The precision-recall curve effectively represents the trade-off between precision and recall for all possible cut-off values based on the model predictions. All three models are capable of accurately classifying the sample images into non-infected and SARS-CoV-2 infected classes, with an average precision of ≈ 0.99 . The performance of the models is independent of the learner architecture; similar results were observed with models based on VGG16 architecture [43]. High average precision implies that these models have learned morphological features that can be generalized to the untreated test set for accurate classification, however, the learned feature spaces have drastically different characteristics due to their input and training procedures. The MIL model $\mathcal{M}(\theta)$ has learned to extract highly discriminative features from micro-populations or single cells whereas the whole-image based model $\mathcal{W}(\phi)$ focuses on the macro-population of cells and learns morphological features of the population, which are not as detailed as in the MIL model $\mathcal{M}(\theta)$. The patch-based model $\mathcal{V}(\psi)$ lies in between those two models in the spectrum; on the one hand it has been trained using

patches with micro-population features, but on the other hand, the labels it was provided for training were noisy. Thus, this model is not focused on the details as the MIL model $\mathcal{M}(\theta)$. These shifts in the learned feature spaces are further transferred to the downstream analysis of treatment efficacy, as discussed in Section 5.2.

5.2 DEEMD dose-dependent efficacy scores are well-structured

We assume that an efficacious treatment is able to effectively stop viral infection and prevent major infection-induced morphological changes in the cell population. Using morphological analysis we can estimate treatment efficacy by profiling the treated and infected cell morphology and quantifying its similarity to non-infected and infected morphology. We applied DEEMD to the infected and treated cell images from the RxRx19a dataset, using $\zeta = 0.5$ inspired by concept of EC_{50} in dose-response curves, which resulted in a ranked list of 18 potential efficacious treatments against SARS-CoV-2, $\mathcal{E}_{\mathcal{M}(\theta)}$ (Supplementary Table 1). The dose-dependent efficacy scores for top 6 ranked effective compounds identified by DEEMD are shown in Fig 4-(a). As is expected for effective antiviral compounds, the dose-dependent efficacy score of the identified treatments increases with increasing concentration, similar to the fitted logistic curve. We expect to observe higher effectiveness with increasing concentration of a treatment with antiviral potency against a specific target, up to the point that it does not cause toxicity or drastically alter cell morphology. Importantly, many of the identified compounds were previously demonstrated to have antiviral activity against SARS-CoV-2 by other drug repurposing studies for COVID-19 based on morphological profiling [9, 17, 18], supporting DEEMD’s methodology. On the other hand, if we compare the same plots for the predictions based on the whole-image based model $\mathcal{W}(\phi)$ to those of the DEEMD, we notice that the estimated efficacy scores are scattered randomly among all treatments and lack any form of structure or pattern, Fig 4-(c). As mentioned in Section 5.1, the patch-based model $\mathcal{V}(\psi)$ is a hybrid of the other two models in terms of its training. This hybridization also manifests itself in Fig 4-(b). We can see that the estimations are slightly structured, not as much as DEEMD’s. Implying that the model was not able to extract the informative features and information completely due to its noisy training environment.

5.3 DEEMD identified treatments are reoccurring in the literature

DEEMD has identified treatments that are reported in the literature to possess therapeutic activity against SARS-CoV-2 or COVID-19 [9, 17, 18, 3]. In Fig 5 we compare DEEMD efficacy scores to hit-scores reported by Heiser *et al.* [17], which similarly used the RxRx19a dataset. We can see that reoccurring treatments with high DEEMD efficacy scores, namely *Remdesivir*, *GS-441524*, and *Aloxistatin* were also assigned a high hit-score by Heiser *et al.* ($\rho = 0.56, p = 0.02$). In the following, we review DEEMD top ranked identified treatments and briefly discuss their potential mechanisms of action against SARS-CoV-2 infection. A complete list of treatments identified to be effective by DEEMD along with the studies that have reported them to be effective against SARS-CoV-2 or COVID-19 is presented in the Supplementary Table 1. We also compare the DEEMD dose-dependent efficacy score for these compounds to previously reported EC_{50} values as a measure of validity whenever such data is available. However it should be noted that due to differences in experimental design, such as cell line, timing of treatment and duration of infection, EC_{50} values from cell-based studies can vary widely.

The most well-known treatment in $\mathcal{E}_{\mathcal{M}(\theta)}$ is *Remdesivir* and its metabolite, *GS-441524*. Previously studied for their antiviral effectiveness against Ebola virus, these compounds target the virus-encoded RNA-dependent RNA polymerase complex needed by RNA viruses to replicate their genome [44]. Multiple studies and clinical trials have found that both compounds are effective against SARS-CoV-2, which led to its emergency approval by the FDA [3, 17, 18, 9] (and references within). DEEMD dose-dependent efficacy scores for *Remdesivir* are shown in Fig 4-(a). We can see that the DEEMD efficacy score for *Remdesivir* is persistently close to 1 for concentrations of $0.3 - 10 \mu\text{M}$ which is

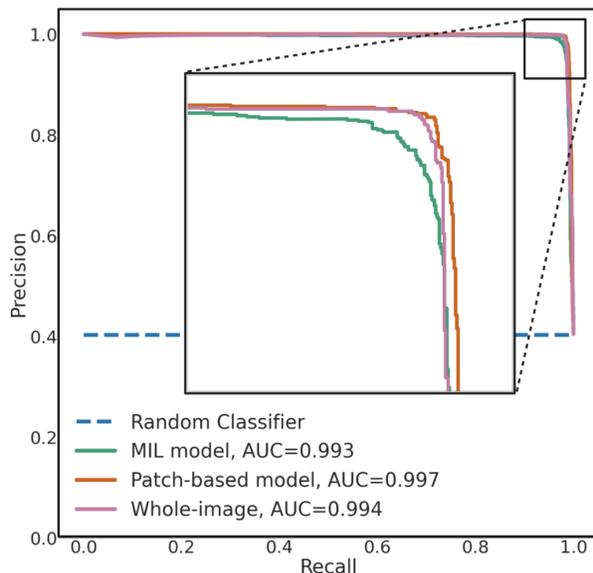


Figure 3: Comparison of precision-recall curves. All three models are capable of accurately classifying the sample images as infected or non-infected in the untreated test set. A random classifier is included as a reference.

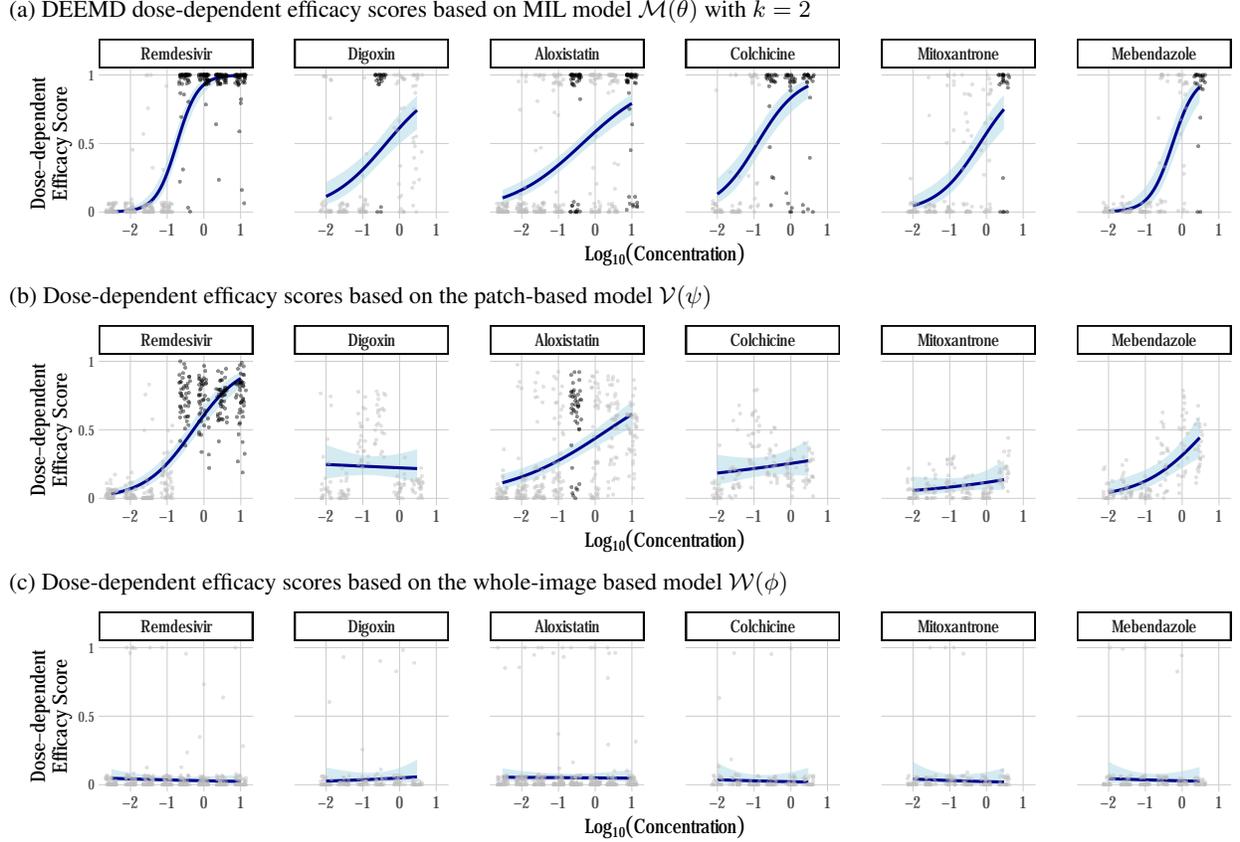


Figure 4: Estimated efficacy scores for top ranked identified treatments in $\mathcal{E}_{\mathcal{M}(\theta)}$. **(a)** Estimated Dose-dependent efficacy scores for DEEMD top-ranked treatments with $k = 2$. The x and y axes show $\text{Log}_{10}(\text{Concentration})$ and dose-dependent efficacy score for each compound. Each data point corresponds to a replicate sample for $t_i^{c_j}$. The points for a concentration are black if DEEMD identified that concentration to be effective, i.e. $e_{t_i}^{c_j} \geq \zeta$. For better visualization and to avoid overlapping points, a small noise on both x and y axis are added to the data points. We used all of the data points, \mathcal{T}_{t_i} , to fit a logistic regression to better visualize the trend. **(c)** and **(b)** are similar plots for predictions based on the whole-image based model $\mathcal{W}(\phi)$ and the patch-based model $\mathcal{V}(\psi)$ respectively.

consistent with EC_{50} values reported by a number of other cell-based studies ranging between 0.003 to 27 μM [45]. Similarly, *GS-441524* was estimated to be efficacious at concentrations of 3 – 10 μM by DEEMD, which is comparable to previously reported EC_{50} values ranging from 0.5 to 8.2 μM [45].

Digoxin is a treatment used for heart disease with a well-established safety profile. Multiple *in vitro* drug repurposing studies for COVID-19 reported its ability to inhibit SARS-CoV-2 infection [45]. The exact mechanism of action for viral inhibition is not identified yet, however, Cho *et al.* [46] hypothesised that *Digoxin* inhibition occurs at the step of viral RNA synthesis. They used multiple FDA-approved treatments, including *Digoxin*, on SARS-CoV-2 infected Vero cells. To understand how the drugs might inhibit SARS-CoV-2, they were administered at three different time points: 1) prior to infection (prophylactic), 2) at the time of infection (entry), and 3) after the infection (therapeutic). They reported that *Digoxin* showed high efficacy following prophylactic and therapeutic administration but failed to effectively inhibit the virus when administered at the time of infection. However, it should be noted that *Digoxin* exhibited cytotoxicity, as reported by Mirabelli *et al* [9]. Fig 4-(a) shows the DEEMD estimated dose-dependent efficacy scores for *Digoxin*; it passed the threshold at 0.3 μM which is aligned with multiple studies that found an EC_{50} between 0.04 and 0.2 μM [45]. We hypothesize that the DEEMD dose-dependent efficacy scores are not conclusive for higher concentrations because of *Digoxin* reported cytotoxicity.

Aloxistatin (E-64d) is another potential candidate treatment reported in the literature as an effective agent against SARS-CoV-2. *Aloxistatin* is a membrane-permeable irreversible cysteine-protease inhibitor of calpains and cathepsins. Recent studies have shown that SARS-CoV-2 requires cathepsin L to enter some cell types. *Aloxistatin* can significantly

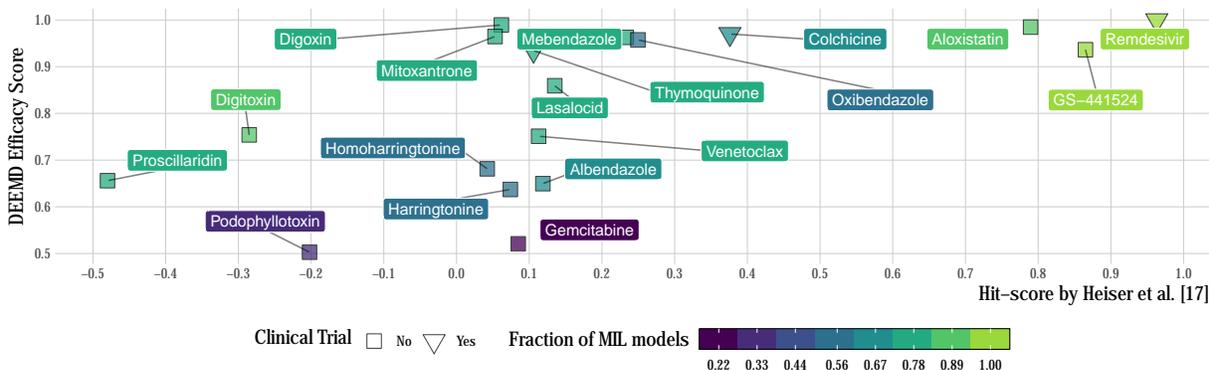


Figure 5: Comparison of DEEMD efficacy scores and Heiser *et al.* hit-scores [17] for treatments in $\mathcal{E}_{\mathcal{M}(\theta)}$. Each data point is color-coded based on the fraction of the MIL models that identify it as effective and its shape indicates whether the treatment has gone into clinical trial against COVID-19. The y-axis represents the DEEMD efficacy scores and hit-scores are on the x-axis.

reduce entry of SARS-CoV-2 pseudovirions by inhibiting cathepsin L [47, 3, 48] (and references within). Fig 4-(a) shows DEEMD estimated dose-dependent efficacy scores for *Aloxistatin*. As we can see, the estimated efficacy is higher than the cutoff threshold ζ at 0.3 and 10 μM . Olaley *et al.* used Vero cells to investigate the antiviral activity of several compounds including *Aloxistatin* for which they reported an EC_{50} of 22 μM [49]. Two additional studies [17] and [18] using morphology-based approaches found that *Aloxistatin* shows strong efficacy without inducing morphological changes to cells.

DEEMD also identified *Mitoxantrone* to be effective against SARS-CoV-2, as shown in Fig 4-(a). The cell surface heparan sulfate (HS) is a molecule commonly found on the membrane and on extracellular proteins of cells that assists the endocytosis of many cargos, including SARS-CoV-2 spike. Recently, Zhang *et al.* reported that HS facilitates spike-dependent viral entry of SARS-CoV-2 [50]. They experimented with inhibitor drugs that target the HS-dependent cell entry pathway, and observed that *Mitoxantrone* inhibited viral entry by directly binding to cell surface HS. DEEMD identified *Mitoxantrone* to be effective at 3.0 μM , while Zhang *et al.* reported an EC_{50} value of 0.03 μM .

As shown in Fig 4-(a), DEEMD also identified *Colchicine* as an effective compound against SARS-CoV-2 at concentrations of 0.3 – 3.0 μM . *Colchicine* is an anti-inflammatory compound used for multiple indications including gout and heart disease. The anti-inflammatory activity is predicted to result from the inhibitory effect on tubulin polymerization and microtubule assembly [3]. Due to the anti-inflammatory properties, and the known link between inflammation and severe COVID-19 outcomes, *Colchicine* has been evaluated as a potential COVID-19 therapeutic with conflicting results [51]. Encouragingly, a meta-analysis of multiple clinical trials suggests *Colchicine* is associated with improved outcomes for individuals with COVID-19 [52]. To date, a direct antiviral effect of *Colchicine* on SARS-CoV-2 has not been demonstrated, but tubulin has been implicated in coronavirus entry [53]. Our findings suggest that the antiviral activity of *Colchicine* against SARS-CoV-2 warrants further investigation.

DEEMD also identified three compounds, *Mebendazole*, *Oxibendazole* and *Albendazole* with effective concentrations ranging from 1.0 – 3.0 μM Fig 4-(a). These compounds belong to a large chemical family of *Benzimidazoles* that are used to treat nematode and trematode infections [54]. These three compounds specifically have not yet been demonstrated to exhibit therapeutic activity against SARS-CoV-2 or COVID-19. However, antiviral activity against other viruses has been shown using other *Benzimidazole* derivatives with EC_{50} values ranging from 0.02 – 90.0 μM [55]. Thus based on our analysis, these should be further examined for activity and possible mechanism of action against SARS-CoV-2.

5.4 DEEMD identified treatments are robust across multiple selections of k

When comparing the set of identified treatments $\mathcal{E}_{\mathcal{M}(\theta)}$ for different values of k , multiple reoccurring treatments were observed. Moreover, the fitted logistic regression curved on the estimated dose-dependent efficacy scores retained its shape and the inflection points are close together; suggesting a consensus, almost independent of choice of k , between the MIL models on the degree of effectiveness for various treatments (See Fig 2 in the Supplementary Material). As shown in Fig 5, most of the treatments in $\mathcal{E}_{\mathcal{M}(\theta)}$ have a high recurrence rate in the MIL models, implying their distinguishable ability to suppress viral-induced morphological changes robust across multiple MIL models trained with different values for k . We noticed that based on the patch-based model $\mathcal{V}(\psi)$ only three treatments were identified

to be effective: *Remdesivir*, *GS-441524*, and *Aloxistatin*, which are among the most commonly identified treatments both in COVID-19 drug repurposing studies and multiple MIL models trained with different k , as shown in Fig 5. This observation suggests that for high values of k the amount of noisy labels drastically changes the dynamics of the training. The learned feature space lacks the required sensitivity for downstream analysis of treatment identification, while the model is still capable of accurately classifying the sample images into infected and non-infected classes.

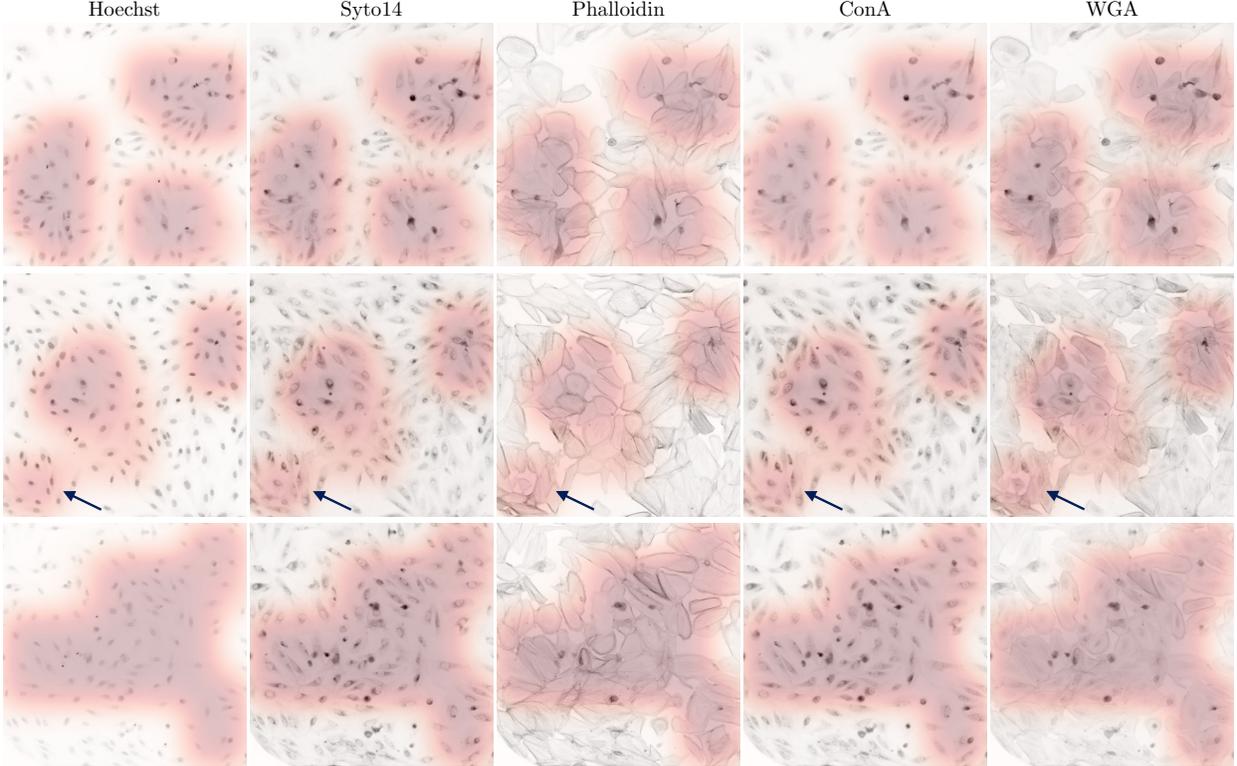


Figure 6: Generated infection maps for a representative set of infected samples from the untreated test set. Sample images are overlaid with their corresponding infection map on each of the indicated stains, displayed separately. Each row represents a SARS-CoV-2 infected sample image from the untreated test set and each column corresponds to stains used in the dataset. The intensity of the red color shows the probability of the infection on each position based on the MIL model $\mathcal{M}(\theta)$ predictions. On the middle row the arrows point to the same region, which is assigned a high infection probability by the MIL model $\mathcal{M}(\theta)$, despite it lacks visible signs of CPE compared to other high infection probability regions. For better visualization, each dye image is color-inverted.

5.5 The MIL infection maps and cytopathic effects

Localization of infected regions in a sample image can be done by aggregating all estimated patch-level labels $\mu_i^{p_i}$ to generate an infection map, as described in Section 3.3. These maps can be used to better understand and explain the MIL model $\mathcal{M}(\theta)$ classification predictions as well as providing an annotation for patches that contain SARS-CoV-2 infected cells. Notice that in the calculation of these maps, no form of annotation was used; instead, the MIL model uses sample-level labels to localize the infection. Figure 6 shows three representative examples from the untreated test set, stain-separated overlaid with infection maps predicted by the MIL model $\mathcal{M}(\theta)$. Visual inspection of the infection maps and regions that are predicted with high infection probability, annotated as red, suggests that the patches that trigger the MIL model $\mathcal{M}(\theta)$ include but not limited to high intensity stained areas, which we predict to reflect cell death or cytopathic effects (CPE), detectable in all channels, except for Hoechst. However, there are patches without the visible signs of CPE that are assigned a high infection probability by the model, e.g. notice the identified infected region on the bottom left corner of samples in the the middle row in Fig 6. CPE refers to changes in host cell structure as a result of viral infection and SARS-CoV-2 is known to be a cytopathogenic agent [56]; for example SARS-CoV-2 is known to cause cell death and to induce syncytia formation (fusion of adjacent cell membranes) [57]. CPE can be measured indirectly by using luminescent cell viability assays [58], however, the RxRx19a dataset does not include a specific viability marker, hence CPE can not be quantified on these samples. Nonetheless, it appears that CPE is

being detected in the highlighted regions of the infection maps. These regions often contain very brightly stained rounded cells or nuclei, and we hypothesize that the stronger fluorescent signals are associated with morphological changes related to cell death such as loss of membrane integrity, cell shrinkage, and nuclear fragmentation [59]. We also observed that the cell nucleus count is significantly lower in the infected samples, Fig 1, which is also indicative of cell death. This suggests that the MIL model $\mathcal{M}(\theta)$ has incorporated biologically relevant morphological features into the infection map, as is expected.

5.6 DEEMD is limited by drug toxicity

Drug toxicity refers to a compound’s negative side effects on a living cell, as the compound can disrupt crucial cellular functions and pathways to the extent of causing cell death. Taking drug toxicity into consideration is essential for optimizing the concentration of the compound needed for optimal efficacy [60]. It is noteworthy to clarify that a low efficacy score $e_{t_i}^{c_j}$ estimated by DEEMD does not necessarily indicate dose or treatment ineffectiveness against SARS-CoV-2. If the treatment has toxic effects or induces other changes in cell morphology, the sample morphology may no longer resemble either the uninfected or SARS-CoV-2 infected class; subsequently, the model’s prediction would not be conclusive. Currently, DEEMD only relies on the assumption that if the drug is toxic, the treated cell morphology would not be similar to uninfected cells, thus the model estimates a low efficacy for a toxic compound. A more complex model capable of integrating drug toxicity with cell morphology is required to properly capture the dynamics of the treatment compound, toxicity and effects on cellular morphology, including cell death.

The RxRx19a public dataset does not include images from healthy drug treated cells, which restricts the morphological feature space learned by the model during training. The model can only learn the regions corresponding to the uninfected and SARS-CoV-2 infected cell morphology, and it is unaware of the space structure outside of these regions. Including additional classes into the training dataset would allow the model to learn a wider range of morphological variations and enable it to differentiate between treatments that are ineffective from those that are impacting cell morphology through toxicity or any other mechanisms; thus a better control over the false negative rate would be in place. We noticed clear cases of drug toxicity with multiple compounds in the dataset, where high concentrations clearly disrupted cell morphology. These sample images lack a detectable signal for the different cell structures and thus we decided to exclude them from the treated test set. Training the model to identify cellular morphology associated with drug toxicity would support the identification of the compound and corresponding concentrations with optimal efficacy and minimal toxicity [60].

6 Conclusion

In this work, we present DEEMD: a pipeline capable of estimating the treatment efficacy of compounds based on morphological analysis of fluorescent-labelled cells. It includes a deep learning model trained within a MIL framework to extract morphological features corresponding to the predicted SARS-CoV-2 infection versus no infection in micro-populations, as well as generating an infection map in a weakly supervised fashion. We compared the performance of the MIL model $\mathcal{M}(\theta)$ to the conventionally trained whole-image based model $\mathcal{W}(\phi)$. As discussed in Section 5, both models are capable of accurately distinguishing between images from uninfected and SARS-CoV-2 infected sample images. By integrating a statistical test into the pipeline, DEEMD identifies efficacious compounds that have been reported to have antiviral effectiveness against SARS-CoV-2 using other methods, supporting the performance of the proposed pipeline, whereas the whole-image based model $\mathcal{W}(\phi)$ fails to estimate meaningful predictions.

We only applied DEEMD to the RxRx19a dataset [17] in this work. With the current configuration, DEEMD is not transferable to other datasets, cell lines, or viruses because it was only trained on the HRCE cells with SARS-CoV-2 infection. A version of DEEMD that is capable of identifying infection-induced morphological changes across multiple viruses and cell lines can be achieved by leveraging a compiled dataset of multiple cell lines and viruses. The training procedure would be the same except for the addition of an auxiliary input that conditions the model based on the cell line and the virus used in the generation of each sample input image.

DEEMD is designed to be adopted and applied to other datasets and problems with few adjustments and modifications. With proper assumptions and data preprocessing, this pipeline can be applied to any fluorescence microscopy datasets, regardless of the stains and image size, for treatment efficacy estimation.

In the future, we plan to generate and apply DEEMD to more comprehensive datasets that include drug-treated, uninfected cells along with specific markers of viral infection that can be used to properly address the shortcomings and limitations of the current version of DEEMD. Having treated uninfected samples would have a great impact on morphological profiling and would bring machine learning-based drug repurposing one step closer to being widely applied to therapeutics.

7 Acknowledgments and Disclosure of Funding

This study was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) Alliance COVID-19 grant: ALLRP 553515-20 (to GH, FJ, and IRN) and Canadian Institutes of Health Research (CIHR) / Operating Grant: COVID-19 May 2020 Rapid Research Funding Opportunity: VR3-172639 (to FJ).

8 Code and Data Availability

The RxRx19a dataset is publicly available through the *Recursion* website: rxrx.ai/rxrx19a. The code for reproducing the result of this paper is available on Github: github.com/Sadegh-Saberian/MIL-COVID19-RxRx19a.

References

- [1] National Institutes of Health (NIH). *Therapeutic Management of Patients with COVID-19*, (accessed October 27, 2020).
- [2] Sudeep Pushpakom, Francesco Iorio, Patrick A Eyers, K Jane Escott, Shirley Hopper, Andrew Wells, Andrew Doig, Tim Williams, Joanna Latimer, Christine McNamee, et al. Drug repurposing: progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.
- [3] Hassan Yousefi, Ladan Mashouri, Samuel C Okpechi, Nikhilesh Alahari, and Suresh K Alahari. Repurposing existing drugs for the treatment of covid-19/sars-cov-2 infection: A review describing drug mechanisms of action. *Biochemical pharmacology*, page 114296, 2020.
- [4] Rory KM Long, Kathleen P Moriarty, Ben Cardoen, Guang Gao, A Wayne Vogl, Francois Jean, Ghassan Hamarneh, and Ivan R Nabi. Super resolution microscopy and deep learning identify zika virus reorganization of the endoplasmic reticulum. *Scientific Reports*, 10(1):1–18, 2020.
- [5] Mirko Cortese, Ji-Young Lee, Berati Cerikan, Christopher J Neufeldt, Viola MJ Oorschot, Sebastian Köhrer, Julian Hennies, Nicole L Schieber, Paolo Ronchi, Giulia Mizzon, et al. Integrative imaging reveals sars-cov-2-induced reshaping of subcellular morphologies. *Cell host & microbe*, 28(6):853–866, 2020.
- [6] Mark-Anthony Bray, Shantanu Singh, Han Han, Chadwick T Davis, Blake Borgeson, Cathy Hartland, Maria Kost-Alimova, Sigrun M Gustafsdottir, Christopher C Gibson, and Anne E Carpenter. Cell painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes. *Nature protocols*, 11(9):1757, 2016.
- [7] Rainer Pepperkok and Jan Ellenberg. High-throughput fluorescence microscopy for systems biology. *Nature reviews Molecular cell biology*, 7(9):690–696, 2006.
- [8] V Starkuviene and R Pepperkok. The potential of high-content high-throughput microscopy in drug discovery. *British journal of pharmacology*, 152(1):62–71, 2007.
- [9] Carmen Mirabelli, Jesse W Wotring, Charles J Zhang, Sean M McCarty, Reid Fursmidt, Namrata S Kadambi, Anya T Amin, Teresa R O’Meara, Carla D Pretto-Kernahan, Jason R Spence, et al. Morphological cell profiling of sars-cov-2 infection identifies drug repurposing candidates for covid-19. *bioRxiv*, 2020.
- [10] Mathias J Wawer, Kejie Li, Sigrun M Gustafsdottir, Vebjorn Ljosa, Nicole E Bodycombe, Melissa A Marton, Katherine L Sokolnicki, Mark-Anthony Bray, Melissa M Kemp, Ellen Winchester, et al. Toward performance-diverse small-molecule libraries for cell-based phenotypic screening using multiplexed high-dimensional profiling. *Proceedings of the National Academy of Sciences*, 111(30):10911–10916, 2014.
- [11] Oren Z Kraus, Jimmy Lei Ba, and Brendan J Frey. Classifying and segmenting microscopy images with deep multiple instance learning. *Bioinformatics*, 32(12):i52–i59, 2016.
- [12] Michael J Cox, Steffen Jaensch, Jelle Van de Waeter, Laure Cougnaud, Daan Seynaeve, Soulayman Benalla, Seong Joo Koo, Ilse Van Den Wyngaert, Jean-Marc Neefs, Dmitry Malkov, et al. Tales of 1,008 small molecules: phenomic profiling through live-cell imaging in a panel of reporter cell lines. *Scientific reports*, 10(1):1–14, 2020.
- [13] Aravind Subramanian, Rajiv Narayan, Steven M Corsello, David D Peck, Ted E Natoli, Xiaodong Lu, Joshua Gould, John F Davis, Andrew A Tubelli, Jacob K Asiedu, et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. *Cell*, 171(6):1437–1452, 2017.
- [14] Jaak Simm, Günter Klambauer, Adam Arany, Marvin Steijaert, Jörg Kurt Wegner, Emmanuel Gustin, Vladimir Chupakhin, Yolanda T Chong, Jorge Vialard, Peter Buijnsters, et al. Repurposing high-throughput image assays enables biological activity prediction for drug discovery. *Cell chemical biology*, 25(5):611–618, 2018.

- [15] Yadi Zhou, Fei Wang, Jian Tang, Ruth Nussinov, and Feixiong Cheng. Artificial intelligence in covid-19 drug repurposing. *The Lancet Digital Health*, 2020.
- [16] Maxim Kuleshov, Daniel JB Clarke, Eryk Kropiwnicki, Kathleen Jagodnik, Alon Bartal, John Erol Evangelista, Abigail Zhou, Laura B Ferguson, Alexander Lachmann, and Avi Ma’ayan. The covid-19 gene and drug set library. 2020.
- [17] Katie Heiser, Peter F McLean, Chadwick T Davis, Ben Fogelson, Hannah B Gordon, Pamela Jacobson, Brett L Hurst, Ben J Miller, Ronald W Alfa, Berton A Earnshaw, et al. Identification of potential treatments for covid-19 through artificial intelligence-enabled phenomic analysis of human cells infected with sars-cov-2. *bioRxiv*, 2020.
- [18] Michael F Cuccarese, Berton A Earnshaw, Katie Heiser, Ben Fogelson, Chadwick T Davis, Peter F McLean, Hannah B Gordon, Kathleen-Rose Skelly, Fiona L Weathersby, Vlad Rodic, et al. Functional immune mapping with deep-learning enabled phenomics applied to immunomodulatory and covid-19 drug discovery. *bioRxiv*, 2020.
- [19] Thai-Hoang Pham, Yue Qiu, Jucheng Zeng, Lei Xie, and Ping Zhang. A deep learning framework for high-throughput mechanism-driven phenotype compound screening and its application to covid-19 drug repurposing. *Nature Machine Intelligence*, pages 1–11, 2021.
- [20] Claire McQuin, Allen Goodman, Vasiliy Chernyshev, Lee Kamentsky, Beth A Cimini, Kyle W Karhohs, Minh Doan, Liya Ding, Susanne M Rafelski, Derek Thirstrup, et al. Cellprofiler 3.0: Next-generation image processing for biology. *PLoS biology*, 16(7):e2005970, 2018.
- [21] Marc-André Carbonneau, Veronika Cheplygina, Eric Granger, and Ghyslain Gagnon. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition*, 77:329–353, 2018.
- [22] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997.
- [23] Fang Wan, Chang Liu, Wei Ke, Xiangyang Ji, Jianbin Jiao, and Qixiang Ye. C-mil: Continuation multiple instance learning for weakly supervised object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2199–2208, 2019.
- [24] Jiawen Yao, Xinliang Zhu, Jitendra Jonnagaddala, Nicholas Hawkins, and Junzhou Huang. Whole slide images based cancer survival prediction using attention guided deep multiple instance learning networks. *Medical Image Analysis*, 65:101789, 2020.
- [25] EL-Manzalawy Yasser, Drena Dobbs, and Vasant Honavar. Predicting mhc-ii binding affinity using multiple instance regression. *Ieee/acm transactions on computational biology and bioinformatics*, 8(4):1067–1079, 2010.
- [26] Sanghamitra Bandyopadhyay, Dip Ghosh, Ramkrishna Mitra, and Zhongming Zhao. Mbstar: multiple instance learning for predicting specific functional binding sites in microrna targets. *Scientific reports*, 5(1):1–12, 2015.
- [27] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Mirafior, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019.
- [28] Maximilian Ilse, Jakub M Tomczak, and Max Welling. Attention-based deep multiple instance learning. *arXiv preprint arXiv:1802.04712*, 2018.
- [29] Yongluan Yan, Xinggang Wang, Xiaojie Guo, Jiemin Fang, Wenyu Liu, and Junzhou Huang. Deep multi-instance learning with dynamic pooling. In *Asian Conference on Machine Learning*, pages 662–677, 2018.
- [30] Yu Zhao, Fan Yang, Yuqi Fang, Hailing Liu, Niyun Zhou, Jun Zhang, Jiarui Sun, Sen Yang, Bjoern Menze, Xinjuan Fan, et al. Predicting lymph node metastasis using histopathological images based on multiple instance learning with deep graph convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4837–4846, 2020.
- [31] Jaume Amores. Multiple instance classification: Review, taxonomy and comparative study. *Artificial intelligence*, 201:81–105, 2013.
- [32] Gary King and Langche Zeng. Logistic regression in rare events data. *Political analysis*, 9(2):137–163, 2001.
- [33] Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pages 196–202. Springer, 1992.
- [34] Wilfrid J Dixon and Alexander M Mood. The statistical sign test. *Journal of the American Statistical Association*, 41(236):557–566, 1946.
- [35] Luc Vincent and Pierre Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Computer Architecture Letters*, 13(06):583–598, 1991.

- [36] Nicolas Coudray, Paolo Santiago Ocampo, Theodore Sakellaropoulos, Navneet Narula, Matija Snuderl, David Fenyő, Andre L Moreira, Narges Razavian, and Aristotelis Tsirigos. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature medicine*, 24(10):1559–1567, 2018.
- [37] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [38] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [39] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [41] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013.
- [42] Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [44] Calvin J Gordon, Egor P Tchesnokov, Joy Y Feng, Danielle P Porter, and Matthias Götte. The antiviral compound remdesivir potently inhibits rna-dependent rna polymerase from middle east respiratory syndrome coronavirus. *Journal of Biological Chemistry*, 295(15):4773–4779, 2020.
- [45] Philip L Tzou, Kaiming Tao, Janin Nouhin, Soo-Yon Rhee, Benjamin D Hu, Shruti Pai, Neil Parkin, and Robert W Shafer. Coronavirus antiviral research database (cov-rdb): an online database designed to facilitate comparisons between candidate anti-coronavirus compounds. *Viruses*, 12(9):1006, 2020.
- [46] Junhyung Cho, Young Jae Lee, Je Hyoung Kim, Sang il Kim, Sung Soon Kim, Byeong-Sun Choi, and Jang-Hoon Choi. Antiviral activity of digoxin and ouabain against sars-cov-2 infection and its implication for covid-19. *Scientific reports*, 10(1):1–8, 2020.
- [47] Xiuyuan Ou, Yan Liu, Xiaobo Lei, Pei Li, Dan Mi, Lili Ren, Li Guo, Ruixuan Guo, Ting Chen, Jiixin Hu, et al. Characterization of spike glycoprotein of sars-cov-2 on virus entry and its immune cross-reactivity with sars-cov. *Nature communications*, 11(1):1–12, 2020.
- [48] Jian Shang, Yushun Wan, Chuming Luo, Gang Ye, Qibin Geng, Ashley Auerbach, and Fang Li. Cell entry mechanisms of sars-cov-2. *Proceedings of the National Academy of Sciences*, 117(21):11727–11734, 2020.
- [49] Omonike A Olaleye, Manvir Kaur, Collins C Onyenaka, and Tolulope O Adebusuyi. Discovery of clioquinol and analogues as novel inhibitors of severe acute respiratory syndrome coronavirus 2 infection, ace2 and ace2-spike protein interaction in vitro. *bioRxiv*, 2020.
- [50] Qi Zhang, Catherine Zhengzheng Chen, Manju Swaroop, Miao Xu, Lihui Wang, Juhung Lee, Amy Qiu Wang, Manisha Pradhan, Natalie Hagen, Lu Chen, Min Shen, Zhiji Luo, Xin Xu, Yue Xu, Wenwei Huang, Wei Zheng, and Yihong Ye. Heparan sulfate assists SARS-CoV-2 in cell entry and can be targeted by approved drugs in vitro. *Cell Discovery*, 6(1), November 2020.
- [51] Medine Cumhuri Cure, Adem Kucuk, and Erkan Cure. Colchicine may not be effective in COVID-19 infection; it may even be harmful? *Clinical Rheumatology*, 39(7):2101–2102, May 2020.
- [52] Timotius Ivan Hariyanto, Devina Adella Halim, Claudia Jodhinata, Theo Audi Yanto, and Andree Kurniawan. Colchicine treatment can improve outcomes of coronavirus disease 2019 (covid-19): A systematic review and meta-analysis. *Clinical and Experimental Pharmacology and Physiology*, 2021.
- [53] Naomi Schlesinger, Bonnie L Firestein, and Luigi Brunetti. Colchicine in covid-19: an old drug, new use. *Current Pharmacology Reports*, 6(4):137–145, 2020.
- [54] Julie R. Harris and Peter J. Hotez. 276 - intestinal nematodes. In Sarah S. Long, Charles G. Prober, and Marc Fischer, editors, *Principles and Practice of Pediatric Infectious Diseases (Fifth Edition)*, pages 1373–1381.e3. Elsevier, fifth edition edition, 2018.
- [55] Michele Tonelli, Matteo Simone, Bruno Tasso, Federica Novelli, Vito Boido, Fabio Sparatore, Giuseppe Paglietti, Sabrina Pricl, Gabriele Giliberti, Sylvain Blois, et al. Antiviral activity of benzimidazole derivatives. ii. antiviral activity of 2-phenylbenzimidazole derivatives. *Bioorganic & medicinal chemistry*, 18(8):2937–2953, 2010.

- [56] Na Zhu, Wenling Wang, Zhidong Liu, Chaoyang Liang, Wen Wang, Fei Ye, Baoying Huang, Li Zhao, Huijuan Wang, Weimin Zhou, et al. Morphogenesis and cytopathic effect of sars-cov-2 infection in human airway epithelial cells. *Nature communications*, 11(1):1–8, 2020.
- [57] Julian Buchrieser, Jeremy Duffoo, Mathieu Hubert, Blandine Monel, Delphine Planas, Maaran Michael Rajah, Cyril Planchais, Françoise Porrot, Florence Guivel-Benhassine, Sylvie Van der Werf, et al. Syncytia formation by sars-cov-2 infected cells. *bioRxiv*, 2020.
- [58] Laura Riva, Shuofeng Yuan, Xin Yin, Laura Martin-Sancho, Naoko Matsunaga, Sebastian Burgstaller, Lars Pache, Paul De Jesus, Mitchell V Hull, Max Chang, et al. A large-scale drug repositioning survey for sars-cov-2 antivirals. *BioRxiv*, 2020.
- [59] Brian S Cummings and Rick G Schnellmann. Measurement of cell death in mammalian cells. *Current protocols in pharmacology*, 25(1):12–8, 2004.
- [60] Jeremy K Nicholson, John Connelly, John C Lindon, and Elaine Holmes. Metabonomics: a platform for studying drug toxicity and gene function. *Nature reviews Drug discovery*, 1(2):153–161, 2002.

Supplementary Material of DEEMD: Drug Efficacy Estimation against SARS-CoV-2 based on cell Morphology with Deep multiple instance learning

1 Noisy label analysis in MIL

As mentioned in subsection IV-B, the patch-based model $\mathcal{V}(\psi)$ trains the model in the presence of noisy labels. To better understand this model and compare it to the MIL model $\mathcal{M}(\theta)$, it would be insightful to measure how much noise is in the training labels for both models. Each image is split into N patches and there are N_p positive samples and N_n negative samples in the dataset. We define the label noise in the dataset to be the ratio of the samples correctly labeled in the dataset to the total number of samples. Without loss of generality, assume that the fraction of the patches in a sample image that are truly infected has an expected value of λ . Therefore, the expected noise ratio (NR) for the MIL model $\mathcal{M}(\theta)$ and the patch-based model $\mathcal{V}(\psi)$ can be expressed as follows:

$$\text{NR}_{\mathcal{M}(\theta)} = \frac{|\frac{k}{N} - \lambda|N_p}{N_p + N_n}, \quad \text{NR}_{\mathcal{V}(\psi)} = \frac{(1 - \lambda)N_p}{N_p + N_n}. \quad (1)$$

By defining $r(\lambda, \frac{k}{N})$ to be the log ratio of $\text{NR}_{\mathcal{V}(\psi)}$ to $\text{NR}_{\mathcal{M}(\theta)}$, we can quantitatively analyze these two models behaviour in different configurations for λ and k .

$$r(\lambda, \frac{k}{N}) \triangleq \ln\left(\frac{\text{NR}_{\mathcal{V}(\psi)}}{\text{NR}_{\mathcal{M}(\theta)}}\right) = \ln\left(\frac{1 - \lambda}{|\frac{k}{N} - \lambda|}\right). \quad (2)$$

When $r(\lambda, \frac{k}{N}) \geq 0$ the MIL model $\mathcal{M}(\theta)$ has less noisy labels compared to the patch-based model $\mathcal{V}(\psi)$. The landscape of $r(\lambda, \frac{k}{N})$ is visualized in Fig 1 for closer inspection. To ensure the numerical stability of $r(\lambda, \frac{k}{N})$ and keep it bounded on the $y = x$ line, a small value ϵ was added to $\text{NR}_{\mathcal{M}(\theta)}$.

One can notice that in almost three fourth of the cases, $r(\lambda, \frac{k}{N}) \geq 0$, meaning that the MIL model $\mathcal{M}(\theta)$ has less noisy labels in its training procedure on average, hence, the model converges faster to an optimum point. Moreover, the MIL model $\mathcal{M}(\theta)$ training procedure results in a more efficient training both in terms of computation footprint and extracted features. The model is using only k/N of the training set for updating the weights. This is N/k times faster than using the whole dataset in the patch-based model $\mathcal{V}(\psi)$. And more

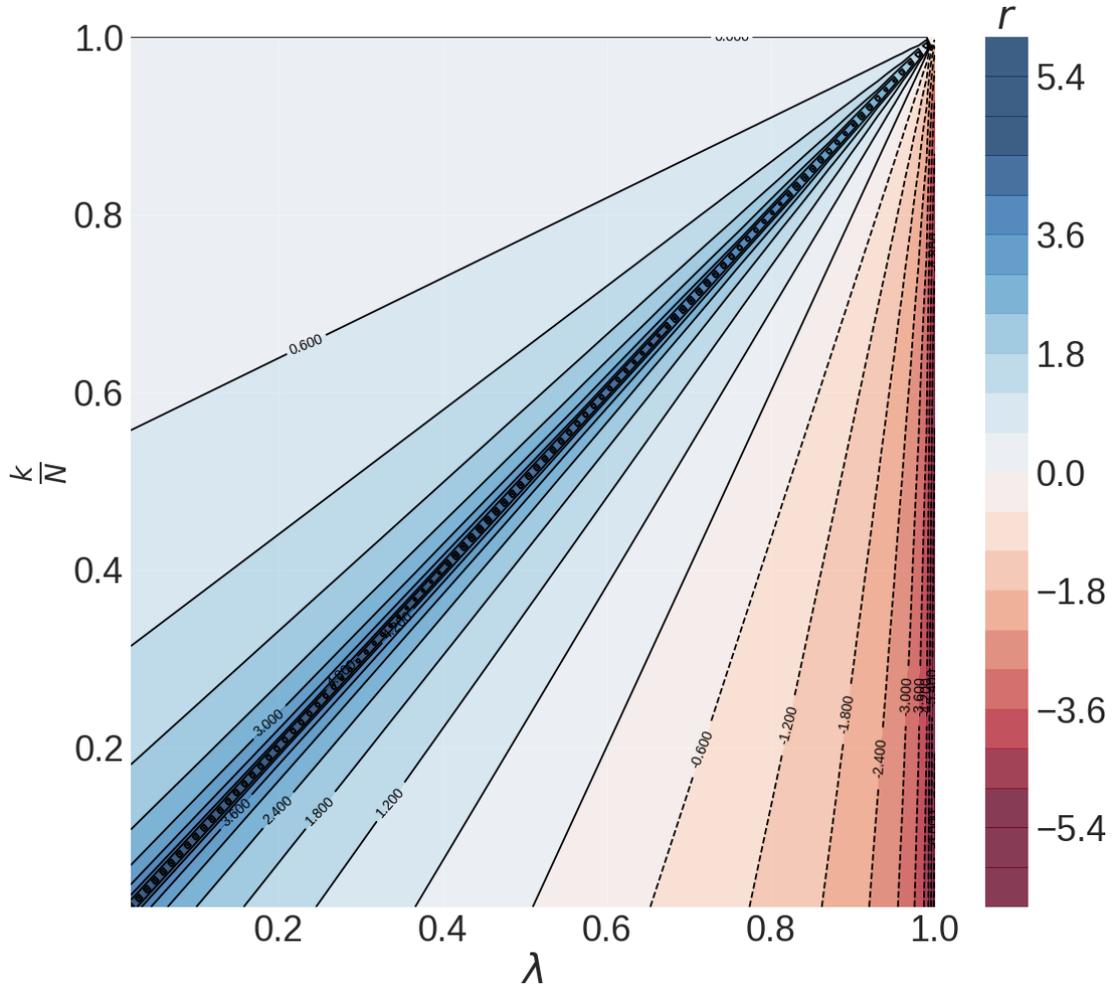


Figure 1: Contour plot for $r(\lambda, \frac{k}{N})$, all possible values for λ and $\frac{k}{N}$ are shown on the x and y axis respectively.

importantly, the patches that are selected for training the MIL model $\mathcal{M}(\theta)$, would contain more discriminative features since they were top-ranked among all patches in their samples for their informativeness. Finally, the MIL model $\mathcal{M}(\theta)$ has the capability of incorporating domain knowledge about the problem into the training by the choice of hyper-parameter k . By choosing k based on prior or domain knowledge, the MIL model $\mathcal{M}(\theta)$ can iteratively refine its training dataset towards for less noisy labels.

Rank	Compound	Studies
1	Remdesivir	[4, 5, 12, 19, 20, 24, 31, 33, 35–37]
2	Digoxin	[3, 31]
3	Aloxistatin	[5, 12, 21, 22, 26, 36]
4	Colchicine	[6, 11, 25, 30, 36]
5	Mitoxantrone	[38]
6	Mebendazole	[8, 18, 32]
7	Oxibendazole	[18]
8	GS-441524	[19, 24, 27, 33, 35, 37]
9	Thymoquinone	[1, 7, 28, 34]
10	Lasalocid	[29]
11	Digitoxin	[16, 17, 23]
12	Venetoclax	[9]
13	Homoharringtonine	[2, 4, 14, 15, 35, 36]
14	Proscillaridin	[10]
15	Albendazole	[18]
16	Harringtonine	-
17	Gemcitabine	[14, 35, 39]
18	Podophyllotoxin	[13]

Table 1: Ranked list of DEEMD identified treatments, $\mathcal{E}_{\mathcal{M}(\theta)}$, along with drug repurposing or clinical studies that reported to show effectiveness against SARS-CoV-2 or COVID-19.

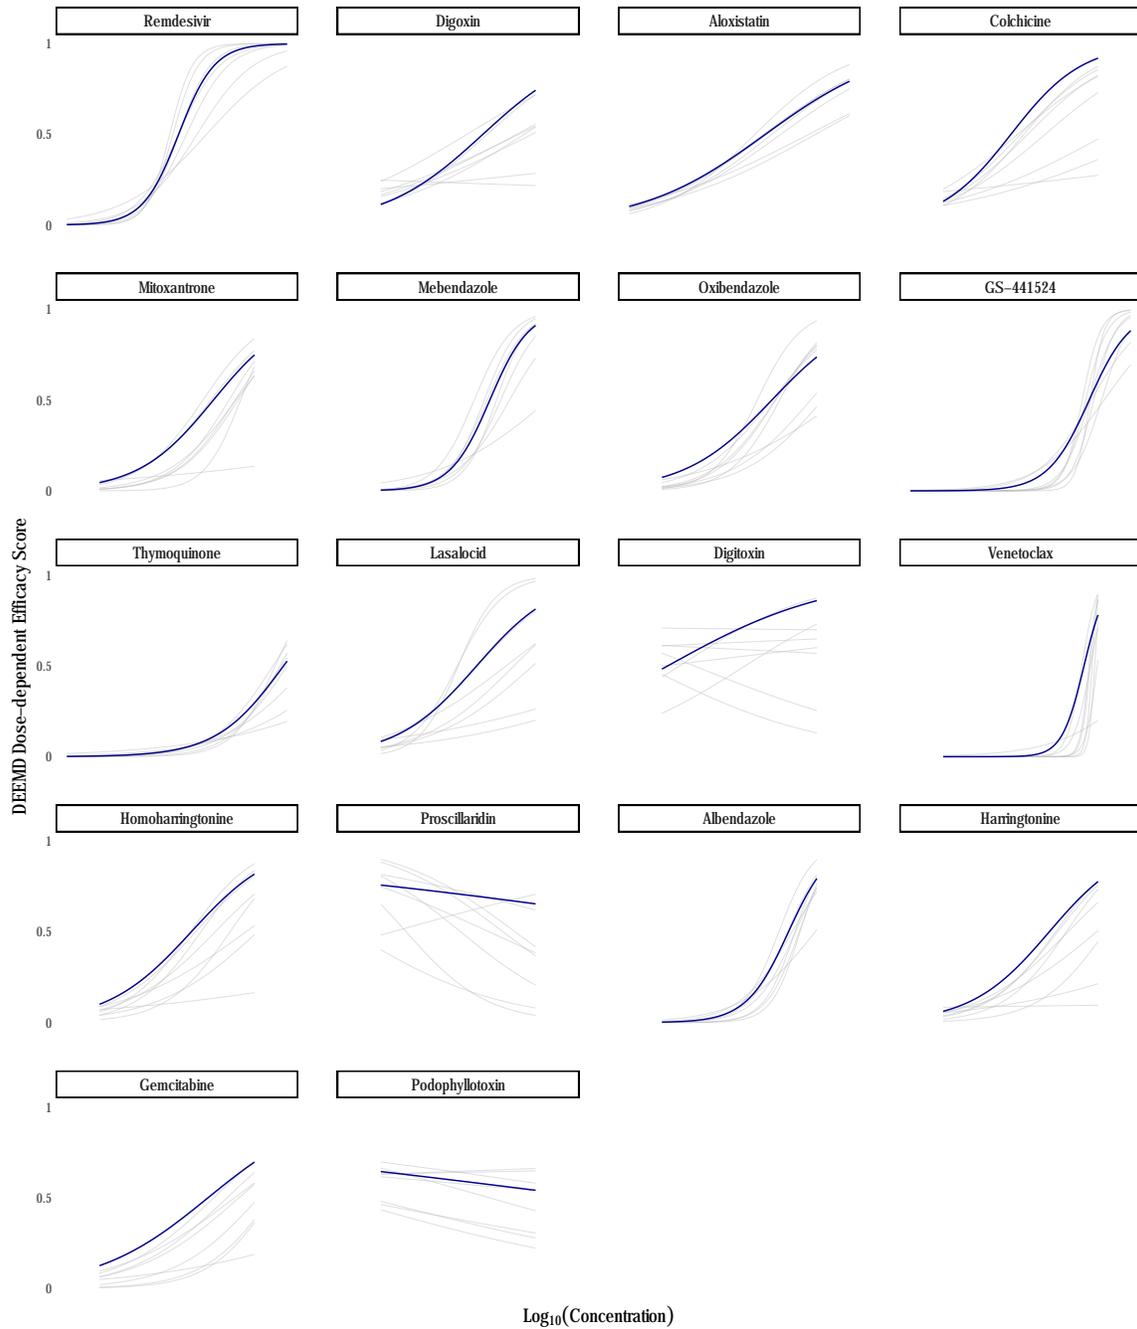


Figure 2: DEEMD estimated efficacy scores for identified treatments for different k values. The logistic curves are fitted based on each value for k . The blue curves represents the MIL model with $k = 2$ which showed the best classification performance on the validation set.

References

- [1] Ajaz Ahmad, Muneeb U Rehman, Parvaiz Ahmad, and Khalid M Alkharfy. Covid-19 and thymoquinone: Connecting the dots. *Phytotherapy Research*, 2020.
- [2] Yanwen Chen, Travis Lear, John Evankovich, Mads Larsen, Bo Lin, Irene Alfaras, Jason Kennerdell, Laura Salminen, Daniel Camarco, Karina Lockwood, et al. A high throughput screen for tmprss2 expression identifies fda-approved and clinically advanced compounds that can limit sars-cov-2 entry. 2020.
- [3] Junhyung Cho, Young Jae Lee, Je Hyoung Kim, Sang il Kim, Sung Soon Kim, Byeong-Sun Choi, and Jang-Hoon Choi. Antiviral activity of digoxin and ouabain against sars-cov-2 infection and its implication for covid-19. *Scientific reports*, 10(1):1–8, 2020.
- [4] Ka-Tim Choy, Alvina Yin-Lam Wong, Prathanporn Kaewpreedee, Sin-Fun Sia, Dongdong Chen, Kenrie Pui Yan Hui, Daniel Ka Wing Chu, Michael Chi Wai Chan, Peter Pak-Hang Cheung, Xuhui Huang, et al. Remdesivir, lopinavir, emetine, and homoharringtonine inhibit sars-cov-2 replication in vitro. *Antiviral research*, page 104786, 2020.
- [5] Michael F Cuccarese, Berton A Earnshaw, Katie Heiser, Ben Fogelson, Chadwick T Davis, Peter F McLean, Hannah B Gordon, Kathleen-Rose Skelly, Fiona L Weathersby, Vlad Rodic, et al. Functional immune mapping with deep-learning enabled phenomics applied to immunomodulatory and covid-19 drug discovery. *bioRxiv*, 2020.
- [6] Medine Cumhuri Cure, Adem Kucuk, and Erkan Cure. Colchicine may not be effective in COVID-19 infection; it may even be harmful? *Clinical Rheumatology*, 39(7):2101–2102, May 2020.
- [7] Sawsan Elgohary, Aya A Elkhodiry, Nada S Amin, Ulrike Stein, and Hend M El Tayebi. Thymoquinone: a tie-breaker in sars-cov2-infected cancer patients? *Cells*, 10(2):302, 2021.
- [8] Ayman Farag, Ping Wang, Mahmoud Ahmed, and Hesham Sadek. Identification of fda approved drugs targeting covid-19 virus by structure-based drug repositioning. *ChemRxiv*, doi: doi, 10, 2020.
- [9] Moritz Fürstenau, Petra Langerbeins, Nisha De Silva, Anna Maria Fink, Sandra Robrecht, Julia von Tresckow, Florian Simon, Karin Hohloch, Jolanda Droogendijk, Marjolein van der Klift, et al. Covid-19 among fit patients with cll treated with venetoclax-based combinations. *Leukemia*, 34(8):2225–2229, 2020.
- [10] Kaifu Gao, Duc Duy Nguyen, Jiahui Chen, Rui Wang, and Guo-Wei Wei. Repositioning of 8565 existing drugs for covid-19. *arXiv preprint arXiv:2005.10028*, 2020.
- [11] Timotius Ivan Hariyanto, Devina Adella Halim, Claudia Jodhinata, Theo Audi Yanto, and Andree Kurniawan. Colchicine treatment can improve outcomes of coronavirus disease 2019 (covid-19): A systematic review and meta-analysis. *Clinical and Experimental Pharmacology and Physiology*, 2021.

- [12] Katie Heiser, Peter F McLean, Chadwick T Davis, Ben Fogelson, Hannah B Gordon, Pamela Jacobson, Brett L Hurst, Ben J Miller, Ronald W Alfa, Berton A Earnshaw, et al. Identification of potential treatments for covid-19 through artificial intelligence-enabled phenomic analysis of human cells infected with sars-cov-2. *bioRxiv*, 2020.
- [13] Andreas Hensel, Rudolf Bauer, Michael Heinrich, Verena Spiegler, Oliver Kayser, Georg Hempel, and Karin Kraft. Challenges at the time of covid-19: Opportunities and innovations in antivirals from nature. *Planta medica*, 86(10):659, 2020.
- [14] Aleksandr Ianevski, Rouan Yao, Svetlana Biza, Eva Zusinaite, Andres Mannik, Gaily Kivi, Anu Planken, Kristiina Kurg, Eva-Maria Tombak, Mart Ustav, et al. Identification of novel antiviral drug combinations in vitro and tracking their development. *bioRxiv*, 2020.
- [15] Aleksandr Ianevski, Rouan Yao, Mona Høysæter Fenstad, Svetlana Biza, Eva Zusinaite, Tuuli Reisberg, Hilde Lysvand, Kirsti Løseth, Veslemøy Malm Landsem, Janne Fossum Malmring, et al. Potential antiviral options against sars-cov-2 infection. *Viruses*, 12(6):642, 2020.
- [16] Sangeun Jeon, Meehyun Ko, Jihye Lee, Inhee Choi, Soo Young Byun, Soonju Park, David Shum, and Seungtaek Kim. Identification of antiviral drug candidates against sars-cov-2 from fda-approved drugs. *Antimicrobial Agents and Chemotherapy*, 2020.
- [17] Meehyun Ko, Sangeun Jeon, Wang-Shick Ryu, and Seungtaek Kim. Comparative analysis of antiviral efficacy of fda-approved drugs against sars-cov-2 in human lung cells. *Journal of medical virology*, 93(3):1403–1408, 2021.
- [18] Jeffrey N Law, Nure Tasnina, Meghana Kshirsagar, Judith Klein-Seetharaman, Mark Crovella, Padmavathy Rajagopalan, Simon Kasif, and TM Murali. Identifying human interactors of sars-cov-2 proteins and drug targets for covid-19 using network-based label propagation. *arXiv preprint arXiv:2006.01968*, 2020.
- [19] Yingjun Li, Liu Cao, Ge Li, Feng Cong, Yunfeng Li, Jing Sun, Yinzhu Luo, Guijiang Chen, Guanguan Li, Ping Wang, et al. Remdesivir metabolite gs-441524 effectively inhibits sars-cov-2 infection in mouse models. *Journal of medicinal chemistry*, 2021.
- [20] Carmen Mirabelli, Jesse W Wotring, Charles J Zhang, Sean M McCarty, Reid Fursmidt, Namrata S Kadambi, Anya T Amin, Teresa R O’Meara, Carla D Pretto-Kernahan, Jason R Spence, et al. Morphological cell profiling of sars-cov-2 infection identifies drug repurposing candidates for covid-19. *bioRxiv*, 2020.
- [21] Omonike A Olaleye, Manvir Kaur, Collins C Onyenaka, and Tolulope O Adebusuyi. Discovery of clioquinol and analogues as novel inhibitors of severe acute respiratory syndrome coronavirus 2 infection, ace2 and ace2-spike protein interaction in vitro. *bioRxiv*, 2020.
- [22] Xiuyuan Ou, Yan Liu, Xiaobo Lei, Pei Li, Dan Mi, Lili Ren, Li Guo, Ruixuan Guo, Ting Chen, Jiixin Hu, et al. Characterization of spike glycoprotein of sars-cov-2 on

- virus entry and its immune cross-reactivity with sars-cov. *Nature communications*, 11(1):1–12, 2020.
- [23] Bette S Pollard, Jorge C Blanco, and John R Pollard. Classical drug digitoxin inhibits influenza cytokine storm, with implications for covid-19 therapy. *in vivo*, 34(6):3723–3730, 2020.
- [24] Andrea J Pruijssers, Amelia S George, Alexandra Schäfer, Sarah R Leist, Lisa E Gralinski, Kenneth H Dinnon III, Boyd L Yount, Maria L Agostini, Laura J Stevens, James D Chappell, et al. Remdesivir inhibits sars-cov-2 in human lung cells and chimeric sars-cov expressing the sars-cov-2 rna polymerase in mice. *Cell reports*, 32(3):107940, 2020.
- [25] Naomi Schlesinger, Bonnie L Firestein, and Luigi Brunetti. Colchicine in covid-19: an old drug, new use. *Current Pharmacology Reports*, 6(4):137–145, 2020.
- [26] Jian Shang, Yushun Wan, Chuming Luo, Gang Ye, Qibin Geng, Ashley Auerbach, and Fang Li. Cell entry mechanisms of sars-cov-2. *Proceedings of the National Academy of Sciences*, 117(21):11727–11734, 2020.
- [27] Yuejun Shi, Lei Shuai, Zhiyuan Wen, Chong Wang, Yan Yuanyuan, Zhe Jiao, Fenglin Guo, Zhen F Fu, Chen Huanchun, Zhigao Bu, et al. The preclinical inhibitor gs441524 in combination with gc376 efficaciously inhibited the proliferation of sars-cov-2 in the mouse respiratory tract. *Emerging Microbes & Infections*, pages 1–43, 2021.
- [28] Andrei P Sommer, Horst-Dieter Försterling, and Kurt G Naber. Thymoquinone: shield and sword against sars-cov-2. *Precis Nanomed*, 3(3):541–548, 2020.
- [29] Esben B Svenningsen, Jacob Thyrssted, Julia Blay-Cadanet, Han Liu, Shaoquan Lin, Jaime Moyano-Villameriel, David Olagnier, Manja Idorn, Søren R Paludan, Christian K Holm, et al. Ionophore antibiotic x-206 is a potent inhibitor of sars-cov-2 infection in vitro. *Antiviral research*, 185:104988, 2021.
- [30] Jean-Claude Tardif, Nadia Bouabdallaoui, Philippe L L’Allier, Daniel Gaudet, Binita Shah, Michael H Pillinger, Jose Lopez-Sendon, Protasio da Luz, Lucie Verret, Sylvia Audet, et al. Efficacy of colchicine in non-hospitalized patients with covid-19. *Medrxiv*, 2021.
- [31] Philip L Tzou, Kaiming Tao, Janin Nouhin, Soo-Yon Rhee, Benjamin D Hu, Shruti Pai, Neil Parkin, and Robert W Shafer. Coronavirus antiviral research database (cov-rdb): an online database designed to facilitate comparisons between candidate anti-coronavirus compounds. *Viruses*, 12(9):1006, 2020.
- [32] Zhihan Wang, Kai Guo, Pan Gao, Qinqin Pu, Min Wu, Changlong Li, and Junguk Hur. Identification of repurposable drugs and adverse drug reactions for various courses of covid-19 based on single-cell rna sequencing data. *ArXiv*, 2020.

- [33] Xuping Xie, Antonio E Muruato, Xianwen Zhang, Kumari G Lokugamage, Camila R Fontes-Garfias, Jing Zou, Jianying Liu, Ping Ren, Mini Balakrishnan, Tomas Cihlar, et al. A nanoluciferase sars-cov-2 for rapid neutralization testing and screening of anti-infective drugs for covid-19. *Nature communications*, 11(1):1–11, 2020.
- [34] Huan Xu, Bing Liu, Zhen Xiao, Meiling Zhou, Lin Ge, Fan Jia, Yanling Liu, Hongshan Jin, Xiuliang Zhu, Jian Gao, et al. Computational and experimental studies reveal that thymoquinone blocks the entry of coronaviruses into in vitro cells. *Infectious diseases and therapy*, pages 1–12, 2021.
- [35] Cheng-Wei Yang, Tzu-Ting Peng, Hsing-Yu Hsu, Yue-Zhi Lee, Szu-Huei Wu, Wen-Hsing Lin, Yi-Yu Ke, Tsu-An Hsu, Teng-Kuang Yeh, Wen-Zheng Huang, et al. Repurposing old drugs as antiviral agents for coronaviruses. *Biomedical journal*, 43(4):368–374, 2020.
- [36] Hassan Yousefi, Ladan Mashouri, Samuel C Okpechi, Nikhilesh Alahari, and Suresh K Alahari. Repurposing existing drugs for the treatment of covid-19/sars-cov-2 infection: A review describing drug mechanisms of action. *Biochemical pharmacology*, page 114296, 2020.
- [37] Keivan Zandi, Franck Amblard, Katie Musall, Jessica Downs-Bowen, Ruby Kleinbard, Adrian Oo, Dongdong Cao, Bo Liang, Olivia O Russell, Tamara McBrayer, et al. Repurposing nucleoside analogs for human coronaviruses. *Antimicrobial Agents and Chemotherapy*, 65(1), 2020.
- [38] Qi Zhang, Catherine Zhengzheng Chen, Manju Swaroop, Miao Xu, Lihui Wang, Juhung Lee, Amy Qiu Wang, Manisha Pradhan, Natalie Hagen, Lu Chen, Min Shen, Zhiji Luo, Xin Xu, Yue Xu, Wenwei Huang, Wei Zheng, and Yihong Ye. Heparan sulfate assists SARS-CoV-2 in cell entry and can be targeted by approved drugs in vitro. *Cell Discovery*, 6(1), November 2020.
- [39] Ya-Nan Zhang, Qiu-Yan Zhang, Xiao-Dan Li, Jin Xiong, Shu-Qi Xiao, Zhen Wang, Zhe-Rui Zhang, Cheng-Lin Deng, Xing-Lou Yang, Hong-Ping Wei, et al. Gemcitabine, lycorine and oxysophoridine inhibit novel coronavirus (sars-cov-2) in cell culture. *Emerging microbes & infections*, 9(1):1170–1173, 2020.