Dario Cardamone

Analysis and modelling of innovative therapeutic approaches

Human monoclonal antibodies

PhD Programme in Complex Systems for Life Sciences

Department:
Mathematics and Natural Sciences

XXXV Ciclo
Candidate: Dario Cardamone
Supervisors:
Prof.ssa Elisa Ficarra
Dott. Duccio Medini
Accademic years: 2019 - 2022

UNIVERSITÀ
DI TORINO

2023

Series of dissertations submitted to the department of
Mathematics and Natural Sciences, University of Turin

To Camilla, Sandro and Enrico

«L'attività scientifica è materia
per grandissima parte di sforzo fantastico;
chi è incapace di costruire ipotesi
non sarà mai scienziato»
A. Gramsci

Acknowledgements

As I write this thesis, I am reminded of the exceptional work that can be accomplished when individuals from different fields and background come together with an open mind and intelligence. The research we conducted at the crossroads of biology and mathematics was a testament to this, and I am proud to have worked alongside such outstanding colleagues and mentors. We were able to achieve remarkable results, and I am honoured to have the opportunity to present in this work some of the results we achieved.

Thank you, Fabiola, Nicola, Duccio, Elisa, Claudia, Alfredo, Chiara, Matteo, Concetta, Marco, Samuele, Santa, Emanuele, Soraya, Alessio, Francesco, Piero, Giorgio, Anna, Noemi, Vittoria, Sara, Giuseppe, Pardis, Ida, Giampiero, Giulio, Valentina and Rino.

This was not easy, but that was fun.

Dario Cardamone
Turin, June 2023

Abstract

Monoclonal antibodies are promising biomolecules for the prevention and treatment of infectious diseases. The COVID-19 pandemic poses unique challenges to the health systems worldwide, and monoclonal antibodies were the therapeutic product developed by many pharmaceutical companies to support the licensed vaccines in reducing hospitalisation and control the infections. We have described the potency, functionality and genetic characteristics of a panel of human monoclonal antibodies isolated from COVID-19 convalescent patients. We have shown that the most potent neutralising antibodies recognise the spike protein receptor-binding domain. Only 1.4% of them neutralise the authentic virus with a potency of $1 - 10\,ng/mL$. Prophylactic and therapeutic efficacy in the hamster model was observed at 0.25 and $4\,mg/kg$ respectively in absence of antibody Fc functions. Our results were used to select a candidate human monoclonal antibody evaluated in a Phase 2-3 clinical trial aimed to assess the antibody ability to treat SARS-CoV-2 infection and reduce the hospitalisation time.

Beside conventional microfluidics-based assay, we investigate the potential of using high-throughput confocal microscopy as a platform to develop functional assay for the identification of compounds, e.g. monoclonal antibodies or small-molecules, able to inhibit the SARS-CoV-2 in vitro. We used an image dataset released during the pandemic to show that Wasserstein Generative Adversarial Networks enable high-throughput compound screening based on raw images. We demonstrate this by classifying active and inactive compounds tested for the inhibition of SARS-CoV-2 infection in two different cell models: the primary human renal cortical epithelial cells (HRCE) and the African green monkey kidney epithelial cells (VERO). In contrast to previous methods, our deep learning-based approach does not require any annotation, and can also be used to solve subtle tasks it was not specifically trained on, in a self-supervised manner.

Our results motivated the effort for the development of a staining protocol and extend the result obtained on the SARS-CoV-2 to the study of bacteria opsonization mediated by monoclonal antibodies. In particular, our approach utilises high-content microscopy and image analysis techniques to assess how effective monoclonal antibodies are at promoting the process of phagocytosis in macrophages infected with bacteria. Through this method, we were able to measure the extent to which the monoclonal antibodies facilitated the engulfment of the bacteria by the macrophages.

List of Papers

Extremely potent human monoclonal antibodies from COVID-19 convalescent patients

Andreano, Emanuele and Nicastri, Emanuele and Paciello, Ida and Pileri, Piero and Manganaro, Noemi and Piccini, Giulia and Manenti, Alessandro and Pantano, Elisa and Kabanova, Anna and Troisi, Marco and Vacca, Fabiola and Cardamone, Dario and others.
In: *Cell*. Vol. 184, no. 7 (2021), pp. 1821–1835. DOI: 10.1016/j.cell.2021.02.035.
Chapter 2.

Exploiting generative self-supervised learning for the assessment of biological images with lack of annotations

Mascolini, Alessio and Cardamone, Dario and Ponzio, Francesco and Di Cataldo, Santa and Ficarra, Elisa In: *BMC bioinformatics*.
Vol. 23, no. 1 (2022), pp. 295–312. DOI: 10.1186/s12859-022-04845-1.
Chapter 3.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Throughout our biological history, humans have been faced with infectious diseases. While scientific advancements have allowed for better control and eradication of diseases like plague, polio, and smallpox through the development of antibiotics and vaccines, infectious diseases still pose a persistent and global threat. Tuberculosis, caused by a bacterial infection, remains the leading cause of infectious disease in modern times and is responsible for approximately one million deaths each year [1]. Additionally, viral pandemics such as the flu, AIDS, and COVID-19 have collectively claimed many lives [2], [3]. The biological complexity of pathogens like *Neisseria gonorrhoeae* and their ability to evolve and evade immune responses make infectious disease a difficult challenge. However, recent technological advancements in biology and computational power have the potential to expand our understanding of infectious disease and lead to new interventions that reduce human suffering. Human monoclonal antibodies (mAbs) are an attractive class of biologic drug that can significantly help in controlling infectious deseases caused by viral and bacterial pathogens. There are many reasons to study mAbs beyond the obvious goal of drug development for prophylaxis or therapy. One of the main benefits of working with mAbs is gaining a comprehensive understanding of the genetic, molecular, and cellular basis of the naturally occurring immunity to viruses and bacteria. Moreover, analysing mAbs isolated from humans after natural infection or vaccination provides fundamental insights into the functioning of the human immune system.

1.1 Objective

This thesis presents the results of our study, which aims to make novel contributions to the characterisation of heavy and light chain sequences of monoclonal antibodies (mAbs) capable of neutralising SARS-CoV-2. Additionally, we explore the utilisation of biological images in infectious disease research, as well as their applications in the broader fields of computation and biology. The objectives of this study are as follows:

Characterise SARS-CoV-2 Neutralising mAbs Sequences Our primary objective is to analyse and characterise the heavy and light chain sequences of monoclonal antibodies that demonstrate effective neutralisation of SARS-CoV-2. Through this analysis, we aim to identify key features and patterns within these sequences that contribute to their neutralising capabilities.

Biological Image Utilisation We explore the utilisation of biological images in the context of infectious disease research. Our objective is to assess the

value and potential of using fluorescent confocal microscopy to perform high-throughput characterisation of the activity of therapeutic approaches as monoclonal antibodies, or small molecules, and enhance our understanding of viral infections and aid in the development of targeted therapies.

Interdisciplinary Applications Furthermore, we aim to demonstrate the broader applications of biological images in the fields of computation and biology. This includes exploring how advanced computational methods, such as generative deep learning, can be integrated with biological image data to facilitate more accurate and efficient analysis, diagnosis, and treatment of infectious diseases.

By achieving these objectives, we intend to contribute to the existing knowledge of mAbs' neutralisation potential against SARS-CoV-2 and provide valuable insights into the utilisation of biological images for infectious disease research and the interdisciplinary fields of computational biology.

1.2 Literature Review

During the early phase of the COVID-19 pandemic, mAbs have been developed as one of the most rapid public health tool to prevent and treat viral infections[4]–[7]. Moreover, given the crucial role of adaptive immunity in safeguarding against viral infection and disease, extensive efforts have been undertaken by the scientific community to investigate mAbs mediated immunity. These endeavors aim to identify correlates of protection, develop immune-based therapies, and optimise vaccine design and delivery [8]–[12].

In the event of an outbreak caused by an infectious agent such as the SARS-CoV-2 virus or a bacterium with antimicrobial resistance, an exceptionally successful approach for the development of mAb-based therapies involves the identification of naturally occurring human monoclonal antibodies from individuals who have successfully recovered from the specific infection. This strategy expedites the discovery process, mitigates the risks of adverse effects, and enhances the accessibility of the therapy [13]–[15]. Key to this strategy is an effective high-throughput screening (HTS) technology, to quickly identify those few, most potent natural monoclonal antibodies among the many thousands potentially isolated from each patient. The screening should identify monoclonals for their ability to confer disease protection, and fluorescent microscopy and biological images offers a robust platform to develop those HTS screening assay [16]–[20].

Thanks to their inherent capability to discover hidden data structures and extract powerful features representation, Convolutional Neural Network (CNNs) have become the fundamental building blocks in many computer vision applications. Nevertheless, much of their recent success lies in the existence of large labelled datasets: CNNs are data-hungry supervised algorithms, and thus supposed to be fed with a large amount of high quality annotated training samples ([21]; [22]).

However, associating labels to a massive number of images to effectively train a CNN may be extremely problematic in a number of real-world applications. Significant examples are the medical and computational biology domains, where image annotation is an especially cumbersome and time-consuming task that requires solid domain expertise and, more often than not, necessitates consensus strategies to aggregate annotations from several experts to solve class variability problems ([23]; [24]). Moreover, biological systems are affected by multiple sources of variability that make the definition of a supervised task impractical, as they require discovering new effects that were not observed during the generation of the training set. Seeking answer to such limitations, a considerable amount of literature focuses on machine learning systems, especially CNNs, able to adapt to new conditions without needing a large amount of high-cost data annotations. This effort includes advances on transfer learning ([25]), domain adaptation ([26]), semi-supervised learning ([27]; [28]) and self-supervised representation learning ([21]; [29]; [30]). The self-supervised representation learning (SSRL) paradigm has especially received increasing attention in the research community. Yann LeCun, invited speaker at AAAI 2020 conference ([31]), has defined the SSRL as "the ability of a machine to predict any parts of its input from any observed part". In other words, SSRL can be realised by contextualising a supervised learning task in a peculiar form (known as pretext task) to predict only a subset of the information and using the rest to drive the decision process. Although the pretext task guides the learning by means of a supervised loss function, the performance of the model on the pretext is irrelevant, as the actual objective of SSRL is to learn an intermediate representation capable of solving a variety of practical downstream tasks that are completely different from the pretext one. Popular SSRL pretext tasks are rotation, jigsaw, instance discrimination and auto-encoders methods (colorization, denoising, inpainting) ([21]; [32]). There is a twofold explanation behind SSRL's recent success: on one hand it can make use of the tremendous amounts of unlabeled data, heritage of the big-data era; on the other hand it is able to dispose of the human supervision and turn back to the data's self-supervision ([21]; [32]).

Current literature has primarily exploited SSRL on general category object classification tasks (e.g. ImageNet classification) ([21]; [32]). Surprisingly few studies have investigated how to extend SSRL methodologies to other important domains like computational biology or medicine, even though they are among the ones that are most affected by the lack of labelled training data ([32]). In this regard, a longitudinal investigation by Wallace et al. ([32]) recently showed how traditional SSRL feature embedding fails in several biological downstream tasks. The authors suggest that the absence of canonical orientation, coupled with the textural nature of the problems, prevents classical SSRL methods from learning a pertinent representation space. They conclude that finding an optimal SSRL feature embedding for fine-grained, textural and biological domains is still an open question.

In an attempt to solve this problem, one the first works exploring image generation as a SSRL pretext task with biological images was undertaken by Goldsborough et al. ([33]), for the morphological profiling of human cultured

cells with fluorescence microscopy. While they speculated the superiority of adversarially learned representations over auto-encoders ones, the authors found their generative approach not competitive yet with traditional transfer learning-based methodologies ([33]). More recently, a number of studies have further investigated and improved generative-based SSRL methods for biological applications, with special focus on histopathological images ([34]; [35]) and, more recently, cancer cell cultures ([36]). Nonetheless, existing works typically make use of side information in the construction of their generative pretext task, that is tailored to the specific application ([37]).

1.3 Outline

The rest of the thesis is organised as follows:

Chapter 2 presents the result obtained from the identification and study of human monoclonal antibodies isolated from COVID-19 convalescent patients, posing a particular focus on their genetic characterisation.

Chapter 3 is dedicated to biological image analysis. Here we show that Wasserstein Generative Adversarial Networks [38] enable high-throughput compound screening based on raw fluorescent microscopy images. The proposed deep learning-based approach does not require any annotation, and can also be used to solve subtle tasks it was not specifically trained on, in a self-supervised manner.

Chapter 4 describes an opsono-phagocytosis assay, that relies on confocal microscopy and image analysis based on methods developed in Chapter 3, to quantify the impact of mAbs on phagocytosis of bacterium *Neisseria gonorrhoeae* by macrophages.

Chapter 2

Extremely potent human monoclonal antibodies from COVID-19 convalescent patients

Human monoclonal antibodies are safe, preventive, and therapeutic tools that have the potential to help restore the massive health and economic disruption caused by the coronavirus disease 2019 (COVID-19) pandemic. In a work we published in 2021 [39], by single-cell sorting 4,277 SARS-CoV-2 spike protein-specific memory B cells from fourteen COVID-19 convalescent patients, we identified 453 neutralising antibodies. The most potent neutralising antibodies recognised the spike protein receptor-binding domain, followed in potency by antibodies that recognise the S1 domain, the spike protein trimer, and the S2 sub-unit. Only 1.4% of them neutralised the authentic virus with a potency of 1–10 ng/mL. The most potent monoclonal antibody, engineered to reduce the risk of antibody-dependent enhancement and prolong half-life, neutralised the authentic wild-type virus and emerging variants containing D614G, E484K, and N501Y substitutions. During the first months of the COVID-19 pandemic, many research groups have been active in isolating and characterising human monoclonal antibodies from COVID-19 convalescent patients or from humanised mice. Notably, in November 2020 two of them have received the Emergency Use Authorization (EUA) from the FDA ([40] ; [41]). The vast majority of the isolated human monoclonal antibodies neutralise SARS-CoV-2 infection by binding to the spike glycoprotein (S protein), a trimeric class I viral fusion protein that mediates virus entry into host cells by engaging with the human angiotensin-converting enzyme 2 (hACE2) and cellular heparan sulfate as receptors ([42]). The S protein exists in a meta-stable pre-fusion conformation and in a stable post-fusion form ([43]; [44]). Each S protein monomer is composed of two distinct regions, the S1 and S2 subunits. The S1 sub-unit contains the receptor-binding domain (RBD), which is responsible for the interaction with hACE2 and heparan sulfate on host cell membranes triggering the destabilisation of the pre-fusion state of the S protein and consequent transition into the post-fusion conformation. This event results in the entry of the virus particle into the host cell and the onset of infection ([44]; [45]).

To retrieve human monoclonal antibodies specific for SARS-CoV-2 S protein, peripheral blood mononuclear cells (PBMCs) were collected and stained with fluorescently labelled S protein trimer to identify antigen-specific memory B cells (MBCs). A panel of 1,731 human monoclonal antibodies specific for the SARS-CoV-2 S protein were identified and characterised. The 1,731 S protein-specific mAbs, were screened in vitro for their ability to neutralise authentic SARS-CoV-

Subject ID	S-protein positive mAbs	SARS-CoV-2 Neutralizing mAbs
PT-004	158	13
PT-005	33	4
PT-006	34	4
PT-008	148	48
PT-009	61	12
PT-010	39	3
PT-012	132	14
PT-014	165	53
PT-041	165	77
PT-100	230	59
PT-101	219	66
PT-102	145	42
PT-103	129	54
PT-188	72	4
TOTAL	1.731	453

Table 2.1: **COVID-19 mAbs summary**, the table summarises the number of S-protein specific mAbs, and SARS-CoV-2 neutralising antibodies.

2 virus by in vitro micro-neutralisation assay. A panel of 453 (26.2%) mAbs neutralised the authentic virus and prevented infection of Vero E6 cells (Table 2.1). The majority of neutralising antibodies (nAbs) were able to specifically recognise the S protein S1 domain (57.5%; n = 244), while 7.3% (n = 53) of nAbs were specific for the S2 domain, and 35.2% (n = 156) did not recognise single domains but only the S protein in its trimeric conformation.

2.1 Genetic characterisation

From the panel of 453 nAbs, we recovered the heavy chain (HC) and light chain (LC) variable regions of 220 nAbs, which were expressed as full-length immunoglobulin G1 (IgG1) using the transcriptionally active PCR (TAP) [46] approach to characterise their neutralisation potency against the live virus at 100 TCID50. The vast majority of nAbs identified (65.9%; n = 145) had a low neutralising potency and required more than 500 ng/mL to achieve 100% inhibitory concentration (IC100). A smaller fraction of the antibodies had an intermediate neutralising potency (23.6%; n = 52) requiring between 100 and 500 ng/mL to achieve the IC100, while 9.1% (n = 20) required between 10 and 100 ng/mL. Finally, only 1.4% (n = 3) of the expressed nAbs were classified as extremely potent nAbs, showing an IC100 lower than 10 ng/mL.

Based on the first round of screening, 14 nAbs were selected for further genetic characterisation. The genes encoding the fragment antigen-binding region of

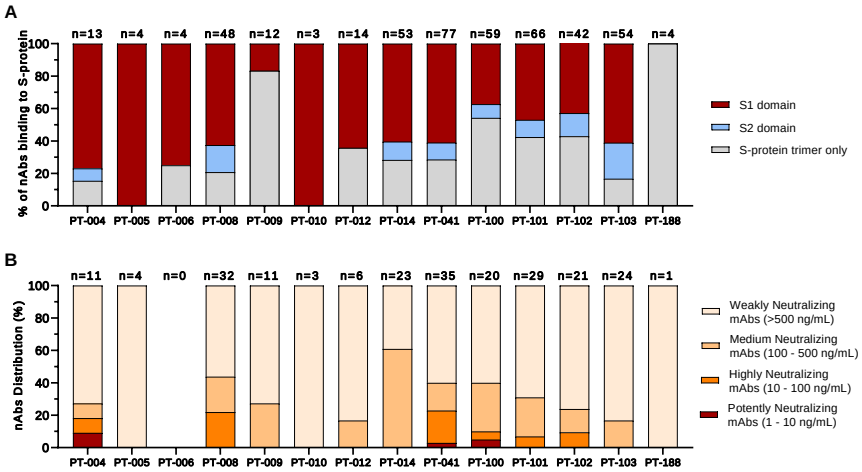

Figure 2.1: **Characterisation and distribution of SARS-CoV-2 S protein-specific nAbs**, (A) The bar graph shows the distribution of nAbs binding to different S-protein domains. In dark red, light blue and gray are shown antibodies binding to the S1-domain, S2-domain and S-protein trimer respectively. The total number (n) of antibodies tested per individual is shown on top of each bar. (B) The bar graph shows the distribution of nAbs with different neutralization potencies. nAbs were classified as weakly neutralizing (> 500 ng/mL; pale orange), medium neutralizing (100 – 500 ng/mL; orange), highly neutralizing (10 – 100 ng/mL; dark orange) and extremely neutralizing (1 – 10 ng/mL; dark red). The total number (n) of antibodies tested per individual is shown on top of each bar.

the Heavy and Light chains of the 14 selected nAbs were sequenced, and their IGHV and IGKV genes were compared with publicly available SARS-CoV-2 neutralising antibody sequences collected in the CoV-AbDab: the Coronavirus Antibody Database ([47]). Four nAbs used one of the most predominant IGHV genes for SARS-CoV-2 nAbs (IGHV1-69), while three nAbs used one of the least representative IGHV genes (IGHV1-24). Two other nAbs employed the most common germline observed for SARS-CoV-2 nAbs, which is IGHV3-53 (Figure 2.2) [48]. Interestingly, while IGHV1-69 and IGHV1-24 accommodate IGHJ diversity, nAbs belonging to the IGHV3-53 gene family only showed recombination with the IGHJ6 gene (Table 2.2). The IGHV genes somatic hypermutation level and complementary determining region 3 (H-CDR3) length were also evaluated. Our selected nAbs displayed a low level of somatic mutations when compared to the inferred germline with sequence identities ranging from 95.6% to 99.3% (Figure 2.2 **(b)**; Table 2.2). The H-CDR3 length spanned from 7 to 21 amino acids (aa) with the majority of the antibodies (n = 6; 42.0%) having a length of 14 to 16 aa that is slightly longer than previously observed

mAb ID	IGHV gene	IGHJ gene	H-CDR3 Length	IGHV Germline identity	IGLV gene	IGLJ gene	L CDR3 Length	IGLV Germline identity
J08	IGHV1-69	IGHJ4	18	96.9	IGKV3-11	IGKJ4	5	98.9
I14	IGHV1-58	IGHJ3	16	96.9	IGKV3-20	IGKJ1	9	98.6
F05	IGHV3-53	IGHJ6	14	97.6	IGKV1-17	IGKJ1	9	94.3
I15	IGHV1-24	IGHJ4	20	96.9	IGKV1-9	IGKJ2	9	98.9
G12	IGHV1-69	IGHJ4	14	95.6	IGKV3-15	IGKJ4	8	97.5
I21	IGHV3-30	IGHJ4	16	98.3	IGKV1-9	IGKJ4	10	98.6
F10	IGHV1-24	IGHJ6	21	96.6	IGKV2-24	IGKJ2	9	98.7
L19	IGHV3-11	IGHJ5	11	96.3	IGKV1-33	IGKJ5	9	98.2
H20	IGHV1-69	IGHJ2	15	97.3	IGKV3-11	IGKJ1	10	98.6
F20	IGHV1-24	IGHJ5	15	96.9	IGKV3-15	IGKJ2	10	98.2
J13	IGHV1-69	IGHJ6	17	99.3	IGKV3-11	IGKJ3	10	98.9
C14	IGHV3-53	IGHJ6	11	96.6	IGKV1-9	IGKJ5	8	98.2
D14	IGHV3-30	IGHJ4	13	96.6	IGKV1-39	IGKJ1	10	97.9
B07	IGHV1-46	IGHJ4	7	97.6	IGKV1-16	IGKJ5	9	98.6

Table 2.2: **Genetic description of fourteen selected SARS-CoV-2 nAbs**, The table describes the heavy and light chain V-J gene usage, heavy and light complementary determining region 3 (H-CDR3) length and percentage of nucleotide germline identity for all the fourteen antibodies characterised in this study.

(Figure 2.2 **(c)**; Table 2.2). All of our nAbs used the Kappa light chain, and the majority of them used the common genes IGKV1-9 and IGKV3-11 (n = 6; 42.0%) (Figure 2.3; Table 2.2).

The level of IGKV somatic hypermutation was extremely low for LCs showing a percentage of sequence identities ranging from 94.3% to 98.9% (Figure 2.3 **(b)**; Table 2.2). The LC CDR3 (L-CDR3) lengths were ranging from 5 to 10 aa, which is in line with what was previously observed for SARS-CoV-2 nAbs [49]–[51] (Figure 2.3 **(b)**; Table 2.2). When paired HC and LC gene analysis was performed, IGHV1-69-derived nAbs were found to rearrange exclusively with IGKV3 gene family, whereas IGHV1-24-derived nAbs accommodate LC diversity (Table 2.2). Of note, some of our candidates showed unique HC and LC pairing when compared to the public SARS-CoV-2 nAb repertoire. Particularly, five different HC and LC rearrangements not previously described for nAbs against SARS-CoV-2 were identified. These included the IGHV1-24;IGKV1-9, IGHV1-24;IGKV3-15, IGHV1-46;IGKV1-16, IGHV3-30;IGKV1-9, and IGHV3-53;IGKV1-17 (Figure 2.4).

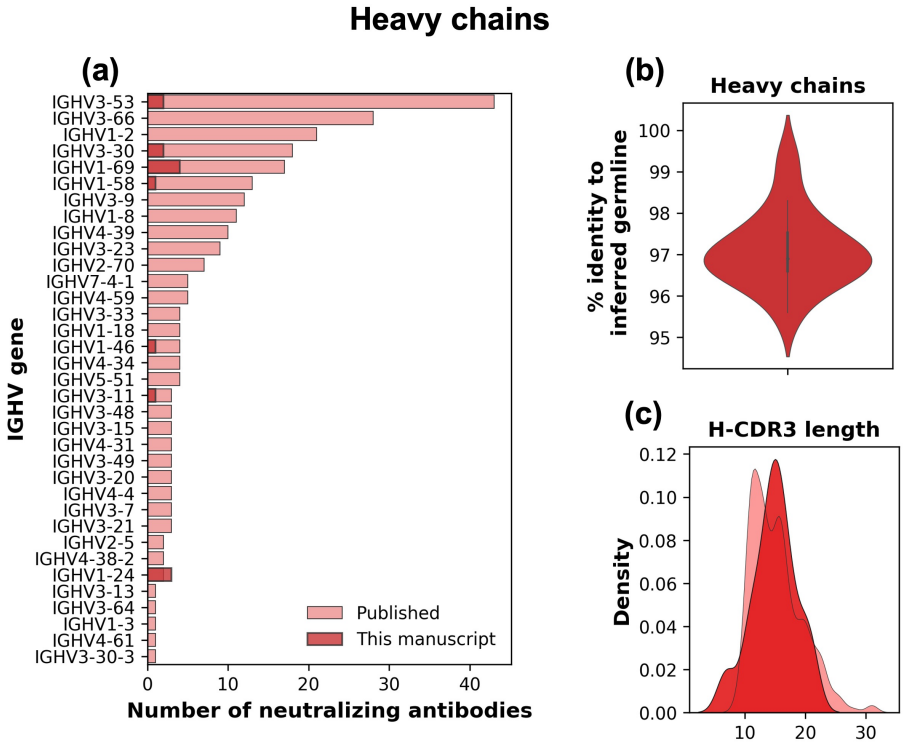

Figure 2.2: **Genetic characterisation of SARS-CoV-2 S selected monoclonal antibodies IGH genes**, **(a)** Bar graph show the heavy chains usage for neutralising antibodies against SARS-CoV-2 in the public repertoire compared to the antibodies identified in this study. Our and public antibodies are shown in dark and light colours, respectively. **(b-c)** The heavy chain percentage of identity to the inferred germline and amino acidic CDR3 length are shown as violin and distribution plot, respectively.

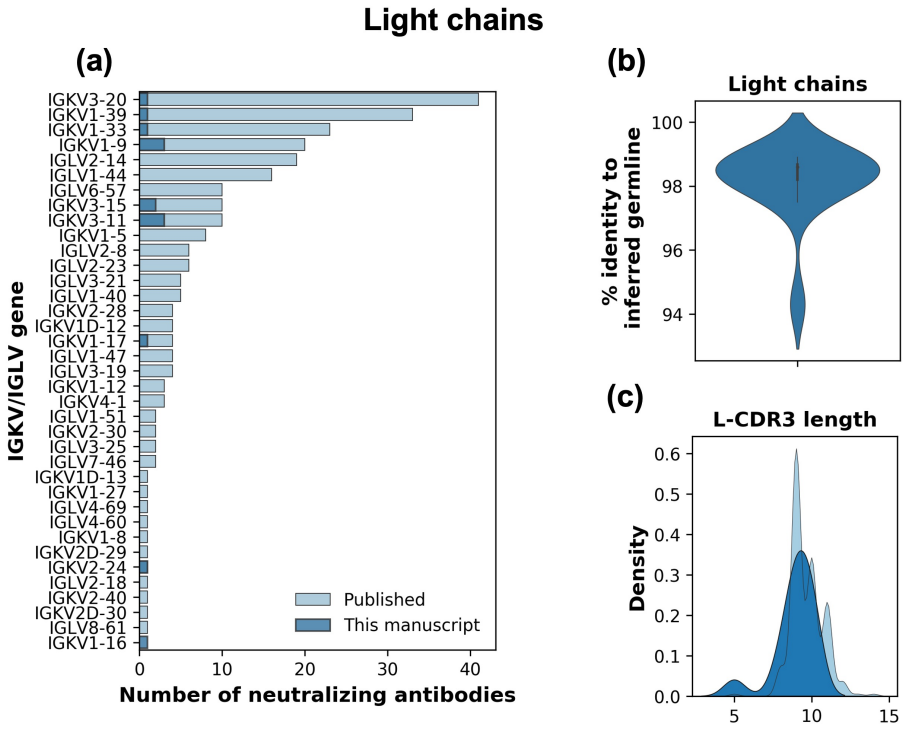

Figure 2.3: **Genetic characterisation of SARS-CoV-2 S selected monoclonal antibodies IGK/IGL genes**, **(a)** Bar graphs show the light chains usage for neutralising antibodies against SARS-CoV-2 in the public repertoire compared to the antibodies identified in this study. Our and public antibodies are shown in dark and light colours, respectively. **(b-c)** The light chain percentage of identity to the inferred germline and amino acidic CDR3 length are shown as violin and distribution plot, respectively.

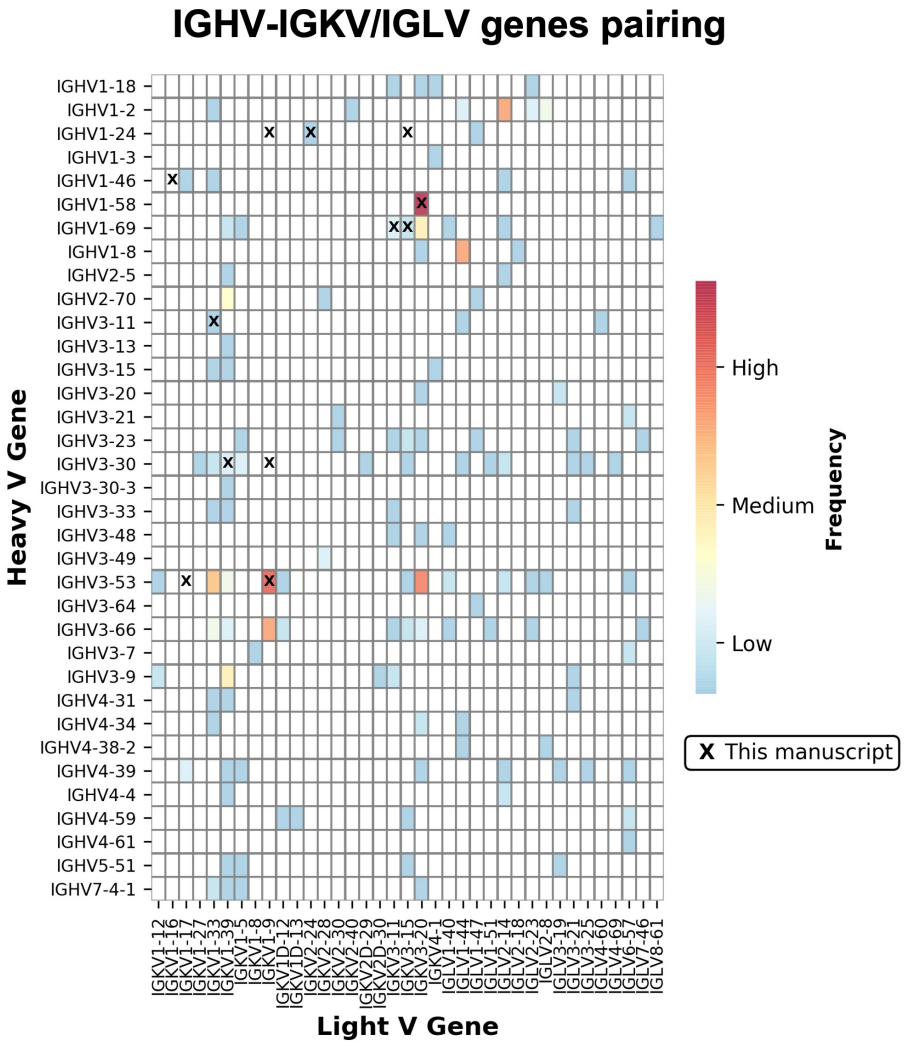

Figure 2.4: **Heavy and light chain pairing for SARS-CoV-2 neutralising human monoclonal antibodies**, the heat-map shows the frequency of heavy and light chain pairing for SARS-CoV-2 neutralising human monoclonal antibodies already published. The number within the heat-map cells represent the amount of nAbs described in this manuscript showing already published (coloured cells) or novel heavy and light chain rearrangements (blank cells).

2.2 Methods for sequence analysis and mAbs annotation

For the genetic Analyses of SARS-CoV-2 S-protein specific nAbs, a custom pipeline was developed for the analyses of antibody sequences and the characterisation and annotation of immunoglobulin genes. Raw sequences were stored as .ab1 files and transformed into .fastq using Biopython[52]. The reads were then quality checked using FastQC (https://www.bioinformatics.babraham.ac.uk/projects/fastqc/) and a report was generated using MultiQC [53]. The antibody leader sequence and the terminal part of the constant region were removed by trimming using Trimmomatic [54]. This latter program was also used to scan and remove low-quality reads using a sliding-window parameter. Once sequences were recovered, germline gene assignment and annotation were performed using the MiXCR suite [55]. The single-read alignment parameters was used, and a CSV-formatted output was generated.

2.3 Discussion

We have described a systematic genetic characterisation of memory B cells from SARS-CoV-2 convalescent patients and the identification of human monoclonal antibodies against SARS-CoV-2. We demonstrated that antibodies with relatively high percentage of identity to the inferred germline are able to neutralise the authentic virus with a potency of 1–10 ng/mL, and the antibody that showed the highest neutralisation potency is able to neutralise the authentic WT SARS-CoV-2 virus at pico molar concentration in vitro. To advance machine-learning-based approaches for designing and modifying antibodies, it is essential to develop experimental datasets comprising sequence data, binding affinity, neutralization potency, and structure conformation. This can save the time-consuming process of discovering new therapeutic antibodies. An example of this strategy is demonstrated in [56], where authors introduced a few mutations to a regulatory-approved antibody product that had escaped the virus and evaluated its enhanced binding efficacy against SARS-CoV-2 Omicron subvariants. The approach utilizes high-performance computing resources, simulation, and machine learning to co-optimize binding efficacy against multiple antigen targets, including RBDs from various SARS-CoV-2 strains, and other desirable qualities, such as thermostability.

In the search for potent antibodies, we found that approximately 10% of the total B cells against the S protein isolated produce neutralizing antibodies, and these can be divided into four different groups recognizing the S1 RBD, S1 domain, S2 domain, and the S protein trimer. Most potently neutralizing antibodies are extremely rare and recognize the RBD, followed in potency by antibodies recognizing the S1 domain, the trimeric structure and the S2 subunit. From these data we can conclude that in COVID-19 convalescent patients, most of the observed neutralization titers are likely mediated by antibodies with medium-high neutralizing potency. Indeed, the extremely potent antibodies

and the antibodies against the S2 subunit are unlikely to contribute to the overall neutralizing titers because they are respectively too rare and too poor neutralizers to be able to make a difference. We and others found that the antibody repertoire of convalescent patients is mostly germline-like. This may be a consequence of the loss of Bcl-6-expressing follicular helper T cells and the loss of germinal centers in COVID-19 patients, which may limit and constrain the B cell affinity maturation [57]. It will be therefore important to perform similar studies following vaccination as it is likely that the repertoire of neutralizing antibodies induced by vaccination may be different from the one described here.

A potential issue associated with the use of human mAbs against viral pathogens is the potential selection of escape mutants. This is usually addressed by using a combination of antibodies directed against non-overlapping epitopes. While this is an ultimate clear solution, it increases the complexity of development, costs of production, drug availability, and affordability. In our case, we believe that selection of escape mutants upon treatment with a single mAb may be quite difficult as the SARS-CoV-2 RNA-dependent polymerase possesses a proofreading machinery, and the epitope recognized by the antibodies herein described overlaps with the region necessary to bind the hACE2 receptor. In this regard, it took more than 70 days of continuous co-culture of the virus in presence of the antibodies before we were able to detect the first emergence of escape mutants of the WT SARS-CoV-2 (data not shown).

Nucleotide and amino acidic sequences of all SARS-CoV-2-neutralising antibodies were deposited in the Italian patent applications n. 102020000015754 filed on June 30th 2020 and 102020000018955 filed on August 3rd 2020. The accession number for the nucleotide sequences of all SARS-CoV-2-neutralising antibodies reported in this paper is GenBank: MW_598287 - MW_598314.

2.4 Limitations of the study

While we believe that our antibodies are extremely potent when compared to most of those described in literature, we acknowledge that in most cases, direct comparison was not performed, and we rely on published data.

Chapter 3

Generative self-supervised learning for the assessment of biological images with lack of annotations

A powerful, and high-throughput approach to characterise the activity of therapeutic approaches as monoclonal antibodies, or small molecules, is the analysis of microscopy images from eukariotic cell lines co-incubated with the treatment under investigation.

However, computer-aided analysis of biological images typically requires extensive training on large-scale annotated datasets, which is costly to produce in many practical applications. In this chapter, we present Generative Adversarial Network Discriminator Learner (GAN-DL), a self-supervised learning paradigm based on the StyleGAN2 ([58]) architecture, which we employ for self-supervised image representation learning in the case of fluorescent biological images.

In this chapter, we show that Wasserstein Generative Adversarial Networks enable high-throughput compound screening based on raw fluorescent microscopy images. We demonstrate this by classifying active and inactive compounds tested for the inhibition of SARS-CoV-2 infection in two different cell models: the primary human renal cortical epithelial cells (HRCE) and the African green monkey kidney epithelial cells (VERO). In contrast to previous methods, our deep learning-based approach does not require any annotation, and can also be used to solve subtle tasks it was not specifically trained on, in a self-supervised manner. For example, it can effectively derive a dose-response curve for the tested treatments.

3.1 Introduction and motivation

Motivated by the the preliminary results obtained by generative-based SSRL methods with different types of microscopy images, in this study we propose GAN Discriminator Learner (GAN-DL), a SSRL framework that exploits the discriminator of a Generative Adversarial Network (GAN) for feature extraction, using a state-of-the-art StyleGAN2 architecture as the backbone ([58]). In our framework, the adversarial training of the StyleGAN2 is exploited as a pretext task, and the trained features of the discriminator provide a new representation space to solve different downstream tasks. By doing so, GAN-DL does not make use of any annotations on the training images, nor of any side information about the specific tasks.

The idea of leveraging GAN's discriminator as feature extractor was first introduced by Radford et al. ([61]), but its employment has been mainly

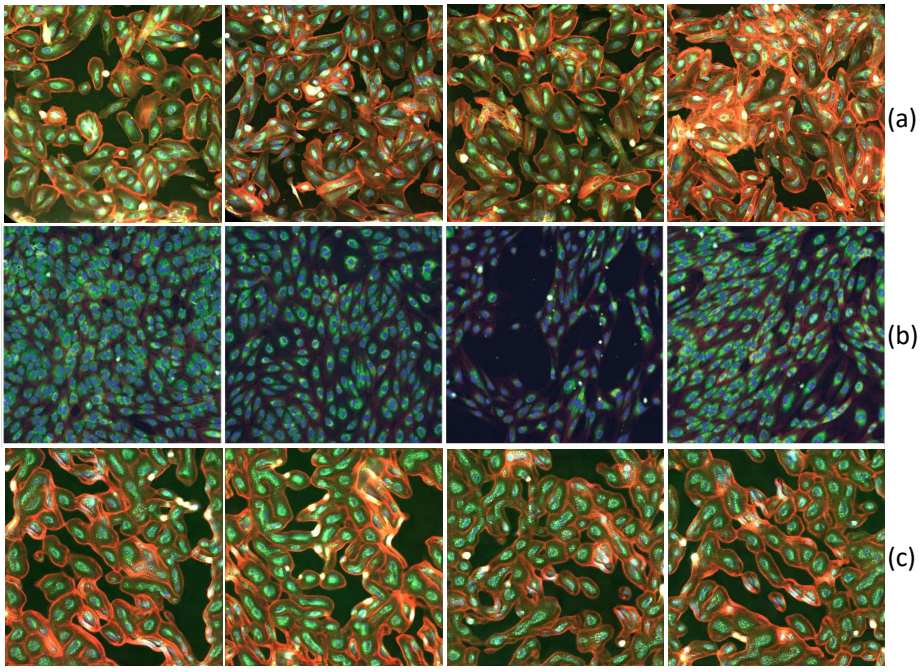

Figure 3.1: The first two rows of the figure show illustrative examples of RxRx19a [59](a) and RxRx1 [60] datasets (b). The third row (c) presents representative examples of Style-GAN generated images for the RxRx19a [59] dataset.

confined to non-biological applications ([62]; [63]). In 2020, Mao et al. ([64]) showed that the effectiveness and robustness of discriminator features strongly depend on avoiding mode collapse in the network. The Wasserstein GANs family StyleGAN2 belongs to are known to be particularly resistant to the mode collapse phenomenon ([63]; [38]). This motivated our choice of using StyleGAN2 ([58]) as the backbone of our method.

To characterise our framework, we focus on a relevant biological case-study, that is COVID-19 drug discovery, exploiting two recently released fluorescence microscopy datasets: (1) RxRx19a ([59]), a morphological imaging dataset that is specific of COVID-19; (2) RxRx1 ([60]; [65]), a non-COVID related collection of fluorescent microscopy images (a more detailed description will follow). In Fig. 1 we show a representative collection of images from RxRx19a ([59]) (a) and RxRx1 ([60]) (b) datasets, depicting different cell models stained with multiple fluorescent dyes. The reported datasets perfectly embody those features (absence of canonical orientation, fine-grained content, textural nature) that make the classical SSRL pretext tasks, as described in the work by Wallace and colleagues ([32]), difficult, if not unsolvable.

Besides the imaging data, a transfer learning-based image embedding for

the RxRx19a benchmark is also accessible online ([66]), which does not exploit any annotation of the target dataset. Such embedding is taken as baseline comparison to prove the goodness of our approach, and referred to as baseline in the rest of the manuscript.

To the best of our knowledge, the only other works addressing the problem of COVID-19 drug discovery with RxRx19a exploit labels of the target dataset in the construction of their embedding ([67]; [68]).

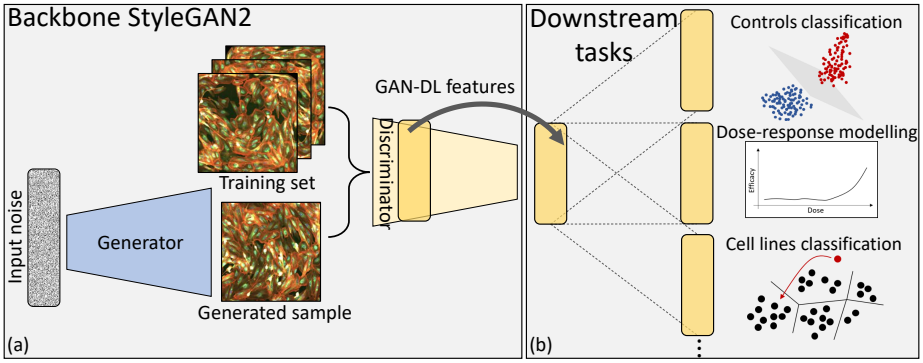

Figure 3.2: Overview of GAN-DL self-supervised representation learning framework, whose pretext task consists in the adversarial game between the generator and the discriminator of the backbone StyleGAN2 (a). The discriminator's features are exploited to several downstream tasks (b): (i) Controls classification - classification of active and inactive compounds against SARS-CoV2 in two different cell models; (ii) Dose-response modelling - disease-associated profiling from raw microscopy images; (iii) Cell models classification - zero-shot representation learning classification task consisting in categorising four different cell types.

The main contributions of the proposed method are the following:

- We propose GAN-DL, a fully SSRL-based approach to characterize relevant biological case studies. We specifically employ generative SSRL in a challenging, real-world biological application of microscopy imaging tailored to COVID-19 drug discovery. We show that GAN-DL, leveraging the pretext of creating diverse and realistic images, is capable not only of proficiently managing downstream classification tasks, but also of separating multiple unrelated features at once along different axis of the latent space.

- GAN-DL significantly deviates from the baseline featurization method proposed by Cuccarese et al. [66] and released together with the RxRx19a [59] benchmark. As a matter of fact, the authors proposed a classic transfer-learning approach featuring a deep network trained from scratch on the RxRx1 [60] and on an additional proprietary images, a very

large dataset that is similar in terms of imaging technology and content to their final application, the RxRx19a [59] dataset. The necessity of a pre-training phase leveraging about 300GB of annotated microscopy images puts serious limitations to the applicability of such method in other contexts affected by scarcity of labelled data. Conversely, as above-mentioned, GAN-DL is trained solely on the unlabelled RxRx19a [59].

- To assess GAN-DL's ability to solve different downstream tasks, we evaluate our method on the classification of active and inactive compounds against SARS-CoV2 in two different cell lines (see Figure 3.2(b)). We show that GAN-DL: (i) outperforms the classical transfer learning approach consisting of a CNN pre-trained on ImageNet; (ii) is comparable to the baseline method in terms of accuracy, even though it was not purposely trained for the downstream tasks; (iii) is able to model disease-associated profiles from raw microscopy images, without the use of any purposely labelled data during the training.
- Finally, to assess the generalization capability of our method, we exploit the GAN-DL embedding learnt on RxRx19a in a zero-shot representation learning task consisting in categorizing the four different cell types of the RxRx1 [60] benchmark: human liver cancer cells (HEPG2), human umbilical vein endothelial cells (HUVEC), retinal pigment epithelium cells (RPE) and human bone osteosarcoma epithelial cells (U2OS).

3.2 Dataset

The data used in this work are part of the RxRx datasets collections, that are available online [69]. More specifically, in our experiments we exploit:

1. The RxRx19a [59], which gathers several experiments aimed at investigating therapeutic potential treatments for COVID-19 from a library of FDA-approved and EMA-approved drugs or compounds in late-stage clinical trials [59]. After 24 hours post-seeding, the cells have been infected with SARS-CoV-2 and then incubated for 96 hours before fixation, staining and imaging. Images were produced using five channels to highlight the cell membrane and different cellular compartments, leveraging a specific fluorescent staining protocol, as described in the work by Cuccarese and colleagues [66]. The compounds were screened by treating cells in six half-log doses with six replicates per dose for each compound approximately two hours after cell seeding. Further details about the assays protocol can be found at the official dataset website [60]. The resulting dataset is made up of 305,520 fluorescent microscopy images of size equal to $1024 \times 1024 \times 5$. To assess the specificity of the tested compounds, two suitable control groups have been designed. The first one consists in conditioned media preparations generated from uninfected cells (Mock), the second one is made up of cells infected in vitro by active SARS-CoV-2 virus and not treated with any compounds.

2. The RxRx1 [60], a dataset consisting of 296 GB of 16-bit fluorescent microscopy images, created under controlled conditions to provide the appropriate data for discerning biological variation in the common context of changing experimental conditions. The RxRx1 [60] has been specifically created to push innovative machine learning and deep learning pipeline on large biological datasets, aimed at drug discovery and development [60].

We leverage the whole RxRx19a [59] to train our GAN-DL on the pretext task of creating diverse and realistic images. Notably, such task does not require any specific annotation. Experiments on downstream tasks were conducted by using 75% of the control images for training and 25% for testing (randomly split by well), with all images outside of the control group used for dose-response evaluation. Images from the same wells were put in the same partition and class imbalances were corrected by automatically adjusting weights inversely proportional to class frequencies in the input data.

No images outside of the training subset of the control group were used in the training of the downstream tasks. For both the RxRx19a and the RxRx1 we performed standard post-processing of the embedded images as described in [66], including normalisation to remove inter-plate variance.

3.3 GAN-DL's backbone: the StyleGAN2 model

The recent literature about GANs is focused on methodologies to improve their training and counteract the well known difficulties and limitations of this phase [70]. More specifically, Wasserstein Generative Adversarial Networks (W-GANs) [38] have been introduced to prevent two common problems of training GANs. First, mode collapse is a form of GAN failure in which the network learns to generate only a subset of the data, eventually a single image or a discrete set of images representing the modes the distribution has collapsed to. The discriminator ends up trapped into a local minimum and the generator easily presents the same examples over and over to convince the discriminator. This results in a model that is heavily over-fitted on this particular subset. Second, lack of convergence due to either the generator or the discriminator, which are improving at a faster pace than the other network. This prevents the mutual improvement that is necessary for convergence.

W-GANs have proved to be an efficient solution to overcome both those limitation at once, by replacing the classical discriminator model with a critic that scores the realness of a given image by means of the so-called Wasserstein distance [38]. For our GAN-DL we employed the Nvidia's StyleGAN2 architecture [58], that is an instance of W-GAN with residual connections in both the generator and the discriminator.

The original StyleGAN2 model has been scaled down to allow training on more reasonable hardware and time-frames. To reduce the number of parameters, we simplified the fully connected mapping network to be 3 layers deep instead of the original 8. The latent space we employ corresponds to the style vector, the sizing of which is 512 in accordance with the original paper. The network

analyses each sample as a 5-channel image, with each channel containing 1 stain. To do so, the only adaptation needed over the original StyleGAN2 model is to increase the filter size of the convolutional layer closest to the image of both the generator and the discriminator to 5. Refer to Supplementary Materials for the experimental setup.

3.4 Counterpart embedding

In our experiments, GAN-DL embedding is compared against several different counterparts:

- The RxRx19a [59] embedding released by Cuccarese et al. together with the imaging data and referred to as *baseline* in this manuscript [66]. It consists of 1024-dimensional vectors (one vector per image) obtained using a DenseNet CNN architecture specifically pre-trained for identifying the different 1,108 genetic perturbations across the four human cell types gathered in the RxRx1 dataset [59], [60]. Such dataset, which collects 125,514 high-resolution fluorescence microscopy images with corresponding labels, is a source annotated dataset with very similar imaging characteristics to the target one (the RxRx19a [59]). The author adapted their DenseNet-based network by firstly changing the initial convolutional layer to accept image input of size $512 \times 512 \times 5$. Like the original DenseNet model, they used Global Average Pooling to contract the final feature maps to a vector of length 2,208. Then, instead of following immediately with a classification layer, the authors added a fully connected layer of dimension 1,024 used as the embedding of the image. The weights of the network were learned by adding two separate classification layers to the embedding layer, one using softmax activation and the other using ArcFace activation [71], which were simultaneously optimised by training the network to recognise perturbations in the public dataset RxRx1 [59] and in a proprietary dataset of immune stimuli in various cell types, unfortunately not released by the authors. Due to operational constraints, a modified assay protocol, lacking one image channel, was used for the live-virus experiments of the RxRx19a dataset [59]. To accommodate this change, the network was trained on only the five first input channels of the RxRx1 images [59]. The proprietary model is not publicly released by the authors.

- The embedding of a DenseNet CNN pre-trained on a source dataset with completely different imaging characteristics and contents (ImageNet). For a fair comparison, the backbone of this methodology is a DenseNet, same as for the baseline solution.

  Pre-training a neural network with ImageNet data involves interpreting images in terms of RGB channels, while cellular images acquired by a fluorescent staining procedures, as for the generation of RxRx19a [59] and RxRx1 [60] datasets, are potentially represented by a variable number of

channels. The staining procedures adopted for the RxRx datasets collection produced images of 5 channels (RxRx19a [59]) and 6 channels (RxRx1 [60]). To account for this difference, we adopted two different strategies:

  - The *ImageNet-collapsed* strategy, where we introduce a trainable convolutional layer with a kernel size of 1 at the beginning of the RGB pre-trained networks, so that the fluorescent images are converted to 3 channels pseudo-RGB images where each channel is replicated three times. The weights of such input trainable layer were learnt via fine-tuning on the given downstream task, leveraging Adam optimiser with learning rate equal to 0.001. The training lasted only for a few epochs, since the number of trainable weights is low. We picked as final model the one giving best accuracy value during training. The ImageNet-collapsed strategy features the same dimension as the embedding of Cuccarese et al. [66] since it is based on the same architecture.
  - The *ImageNet-concatenated* strategy, where each channel is processed independently and then all the resulting features are concatenated. This strategy does not require any fine-tuning and produces an embedding of size 5120 (1024×5).

- The embedding of a convolutional auto-encoder (referred to as *ConvAE*) trained on the target dataset RxRx19a. For this purpose, we implemented the method presented by Wallace et al [32], that was demonstrated to be superior in term of classification accuracy to jigsaw, rotation and instance discrimination based self-supervised methods on biological images. To allow the auto-encoder to converge on the higher resolution images we are evaluating, we modify the original architecture by adding the same residual connection scheme used in the generator of StyleGAN2 and GAN-DL and a perceptual loss function obtained using an Imagenet pre-trained ResNet50 [72]. The embedding, extracted from the last layer of the encoder, features a size of 1024, same as the baseline and ImageNet-collapsed pre-trained method.

Note that the embedding size varies across the different counterparts. This is constrained by the specific architecture the given featurization strategy leverages. In our GAN-DL, as mentioned in the previous subsection, the latent space we employ corresponds to the style vector, which has a size of 512.

3.5 Experimental setup

ConvAE

The embedding of the ConvAE was obtained by training the model on the target dataset RxRx19a [59], employing an NVIDIA TITAN Xp GPU. To ensure convergence on the high resolution RxRx19a images, we modified the original auto-encoder architecture presented by Wallace et al. [32] in the following way:

we employed the same residual connection scheme used in the generator of StyleGAN2 and GAN-DL and a perceptual loss function obtained using an Imagenet pre-trained ResNet50 [72]. The learning rate was set to 10^{-4} and the training leverages Adam optimiser. Figure 3.3 provides some representative examples of ConvAE reconstructed images (b), alongside the original inputs (a), for the RxRx19a [59] dataset. Lastly, Figure 3.3(c) provides the training and validation trends of the ConvAE's loss with respect to training epochs.

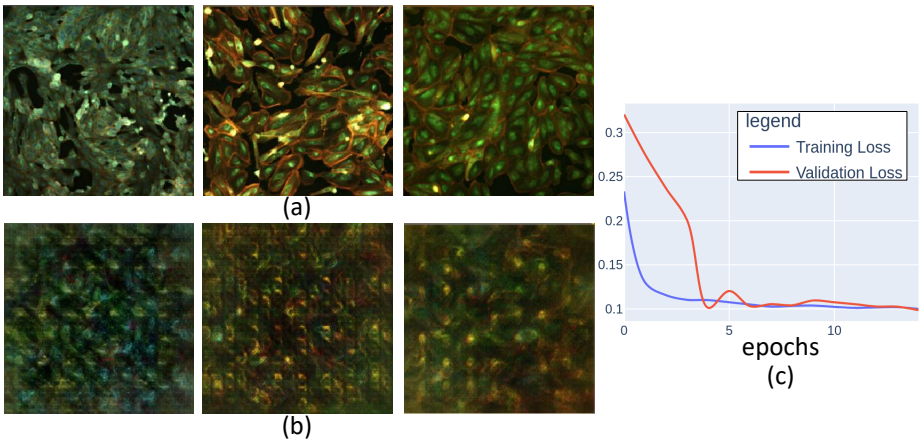

Figure 3.3: Illustrative examples of RxRx19a [59](a) and the corresponding images reconstructed by the ConvAE approach (b). ConvAE's training and validation curves (c).

GAN-DL

The StyleGAN2 backbone was trained on the RxRx19a [60] dataset using Adam optimiser with a learning rate of 10^{-4}, with the same loss as the one described in the StyleGAN2 paper [58]. No hyper-parameter optimisation was performed. Conversely, we employed two regularisation terms:

- Generator: Jacobian Regularization, also referred to as Perceptual Path Length regularisation in [58], Exponential Moving Average of the weights [73]
- Discriminator: Lipschitz L1 penalty [74], R1 regularisation [70]

For training we employed one TPU v3-8 node with 16GiB of RAM per core. The original StyleGAN2 took 9 days on 8 Tesla V100 GPUs to train on the FFHQ dataset [58], while our slimmed and repurposed version required 24 hours on a TPU v3-8 node or 48 hours on a single Tesla V100 GPU to obtain the results shown in this chapter. Most of the difference in training time can be attributed to the lower amount of parameters as well as the vastly different dataset used.

Our experiments specifically seek an answer to two research questions: (i) is it possible to learn an accurate and reliable image featurization, able to encode and describe biological relevant information, leveraging a self-supervised pretext task?; (ii) up to which extent the learned biological information can be transferred to a different dataset? To answer such questions, we have put into effect the properties of GAN-DL featurization in the following experiments.

3.6 Visualizing GAN-DL's representation learning capability

To characterize the representation capability of the proposed SSRL featurization methodology, we evaluate GAN-DL on the RxRx19a [59] dataset. We summarize the screening control samples into two sets of conditions, C^+ and C^-. C^+ represents uninfected samples treated with culture medium or a solvent, and C^- represents samples infected with wild-type SARS-CoV-2 virus. For simplicity and with abuse of notation, we refer to C^+ as *positive controls* and to C^- as *negative controls*.

In the RxRx19a [59] compound screening setting, only the images that correspond to positive and negative sets of conditions can be associated with either *live* or *dead* labels, where those labels refer to the viability of the cellular model imaged in that specific condition. The cellular model viability is unknown for the remaining part of the samples. In this regard the vast majority of the dataset is unlabelled. The large amount of unlabelled data, coupled with the textural and fine-grained aspect of the images, makes RxRx19a [59] a challenging case-study and a perfect candidate to assess our proposed SSRL methodology.

As Figure 3.2 suggests, GAN-DL embedding consists of a 512-dimensional features vector. To assess and interpret its inherent capability of learning a genuine representation, we need to define a projection space able to allow some degrees of visualization of the data structure. Hence, using the control samples, we promote the explainability of the projection procedure by defining:

1. The *effectiveness-space* $\mathbf{E}^2$, a two-dimensional space that represents the treatment effectiveness of the tested compounds on two axes. The *On-perturbation* axis captures the difference between uninfected samples C^+ and infected samples C^-. Intuitively, the screened compounds able to inhibit SARS-CoV-2 infection should have an *On-perturbation* value similar to the C^+ set of conditions. The *Off-perturbation* axis represents the remaining variability in the samples that cannot be unambiguously associated to the compound effectiveness.

2. The *cell models-space* $\mathbf{C}^2$, a two-dimensional space that captures the morphological properties of the two cell models into two dimensions. The *On-perturbation* axis projects the differences of the two cell models onto one direction. Intuitively, all the samples in which VERO cells were used should have a similar *On-perturbation* value. The same goes for the samples

in which HRCE cells were used. The *Off-perturbation* axis represents the remaining variability that cannot be associated to the cell model differences.

Projecting the data along *On-perturbation* and *Off-perturbation* allows us to visually represent the high-dimensional image embedding obtained by GAN-DL into two-dimensional plots. To obtain such directions, we leverage a linear Support Vector Machine (SVM) trained to classify C^+ versus C^- ($\mathbf{E}^2$ space) or HRCE versus VERO control cells ($\mathbf{C}^2$ space). In both the cases, the separation hyper-plane fitted by the SVM and its normal respectively represent the *Off-perturbation* and the *On-perturbation* axis. As shown later in this section, the scalar projection of the 512 GAN-DL features on such spaces are exploited on one hand to provide an effective visual representation of the high-dimensional data structure through point cloud scatter plots, on the other hand to derive dose-response curves for the tested compounds. For better readability, the *On-perturbation* axis is scaled so that C^+ are centered around $+1$ and C^- around -1 and the *Off-perturbation* axis is zero-centered.

The plots in the first row of Figure 3.4 compare our GAN-DL's embedding (a) with the baseline embedding [66](b) in the $\mathbf{E}^2$ projection space, where we expect a degree of separation between C^+ and C^-, since such space was spanned by the SVM trained on the embeddings of the negative and positive controls. The analysis is performed considering the two sets of conditions grouped by cell model. Hence, different colors identify C^+ and C^- for the two distinct cell models: blue and orange for the C^+ of HRCE and VERO cell model, respectively, green and red for the corresponding C^- conditions. As it can be gathered from the degree of separation between C^+ and C^- on the $\mathbf{E}^2$ projection space, both the embeddings behave coherently in separating mock-treated samples from those where the virus was active. A quantitative comparison in terms of degree of separation between C^+ and C^- is presented in the following subsection.

The second row of Figure 3.4 shows GAN-DL featurization (c) and the baseline featurization (d) projected onto the $\mathbf{C}^2$ space, where we expect a certain degree of separation between distinct cell types, irrespective of whether C^+ or C^- are considered. Same as in the previous experiment, results are reported separately for the two cell models. Here HRCE are represented with blue (C^+) and green (C^-) colors, while VERO with orange (C^+) and red (C^-), respectively. Even in this case, the plots demonstrate that GAN-DL is able to caught the inherent variability of the two cell models, in a comparable way to the transfer-learning baseline.

3.7 Assessing controls linear separability

To assess the goodness of our embedding, we try to demonstrate that it is able to establish a good linear separability of samples on two different downstream tasks it was not specifically trained for: (i) the categorization of C^+ versus C^- and (ii) the classification of HRCE and VERO cells.

For both the tasks, the linear separability is verified by exploiting soft margin linear SVMs for classification. More specifically, we compare the classification

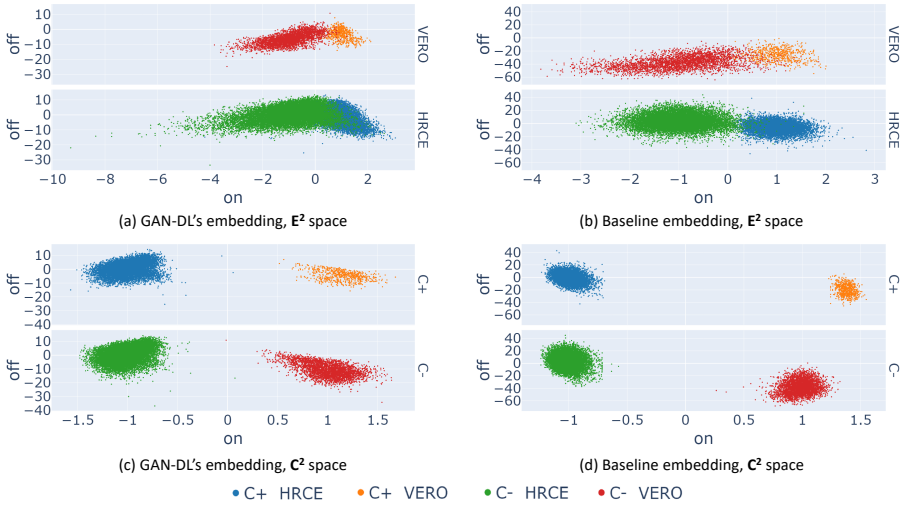

Figure 3.4: The left column of the figure shows the scatter plots of GAN-DL's embedding of the RxRx19a [59] dataset projected onto the E^2(a) and C^2(c) axes. The right column shows the baseline embeddings of the RxRx19a [59] dataset projected onto the E^2(b) and C^2(d) axes.

accuracy of a linear SVM built on top of our embedding with the ones obtained by the other counterpart embeddings: i) the baseline [66], ii) a DenseNet CNN model pre-trained on ImageNet, respectively with *collapsed* and *concatenated* strategy, iii) a convolutional auto-encoder [32] trained on RxRx19a. All results presented are obtained using data that was not included in the SVM training process.

The first two lines of Table 3.1 report the classification accuracy values of the two classification tasks (for the first one, C^+ versus C^-, the two cellular models are merged into the same dataset). From the reported values we can observe that GAN-DL provides informative features for both C^+ versus C^- categorization (91.4% accuracy) and cell models recognition (100% accuracy). The baseline, that leverages the RxRx1 [60] dataset as transfer learning source domain, outperforms GAN-DL by just 5% in terms of C^+ versus C^- classification accuracy, and is equivalently 100% accurate in the other task. This is a remarkable result for GAN-DL, given that its embedding is trained on a completely different pre-text task, which does not require any kind of image annotation. Lastly, GAN-DL outperforms by a large margin the traditional solutions based on ImageNet pre-training (respectively, by 26% and 14% for the two tasks with respect to ImageNet-collapsed solution, and by 11% for the first task with respect to ImageNet-concatenated solution). Finally, our GAN-based approach outperforms by a good margin the other SSRL method based on convolutional auto-encoder, especially in the C^+ versus C^- task.

The last two lines of Table 3.1 report the accuracy of the C^+ versus C^- categorization task, this time separated by the cellular models HRCE and VERO. For all the considered embeddings, we can observe that the accuracy is generally higher when the cell models are separated. Nonetheless, this variation is quite contained for the SSRL solutions. More specifically, GAN-DL shows an accuracy of 92.44% and 99.93% for respectively HRCE and VERO, against the 91.4% obtained with the two models considered together. The baseline, on the other hand, shows an accuracy of 99.28% and 100% for respectively HRCE and VERO, against the 95.81% for the two merged cell models. We can again observe that the ImageNet pre-trained solutions reported a much higher accuracy difference: 84.09% and 84.53% against 65.31% for the ImageNet-collapsed solution, and 90.24% and 99.8% against 79.61% for the ImageNet-concatenated strategy. Finally, even in this configuration, the embedding based on a convolutional auto-encoder obtained the lowest accuracy values.

Table 3.1: Classification accuracy on the downstream tasks.

	GAN-DL	**Baseline [66]**	**ImageNet-collapsed**	**ImageNet-concatenated**	**ConvAE**
C^+ vs C^-	91.4 %	95.81 %	65.31%	79.61%	64.50%
HRCE vs VERO	100.0 %	100.0 %	85.52%	100.0%	99.80%
C+ vs C- (HRCE only)	92.44 %	99.28 %	84.09 %	90.24%	68.41%
C+ vs C- (VERO only)	99.93 %	100 %	84.53 %	99.8%	82.89%

3.8 Automatically deriving dose-response curves from image data

In this section, we exploit the GAN-DL's featurization projected onto the *On-perturbation* axis of the $\mathbf{E}^2$ space, defined in section *Visualizing GAN-DL's representation learning capability*, to automatically derive the dose-response of all the 1,672 screened compounds in RxRx19a [59] dataset. Even in this case, the featurization is the one obtained from the image generation pre-text, which did not exploit any task-specific annotation.

As the figures of merit we propose: (i) the embedding distributions, in the form of a scatter plot at varying concentrations, of *Remdesivir* and *GS-441524*, two compounds proved to be effective on SARS-CoV-2 in vitro in both the cell models, and of *Polydatin*, a compound that is known to be ineffective [66], [75] (see Figure 3.5). These compounds are shown as representative examples for both our embedding (a) and the baseline embedding (b); (ii) the dose-response curves of a number of other compounds, obtained by reporting the corresponding mean *efficacy score* at each concentration (see Supplementary Material).

From Figure 3.5, we can draw a number of considerations. For the effective compounds *Remdesivir* and *GS-441524*, it is possible to see that progressively higher drug concentrations (corresponding to progressively darker red points in the scatter plots) tend to cluster towards positive values of the *On-perturbation* axis in the $\mathbf{E}^2$ space, closer to the region associated to the C^+group: the higher the concentration, the higher the *On-perturbation* value. This is generally true for both the GAN-DL and the baseline embedding (see sections (a) and (b) of

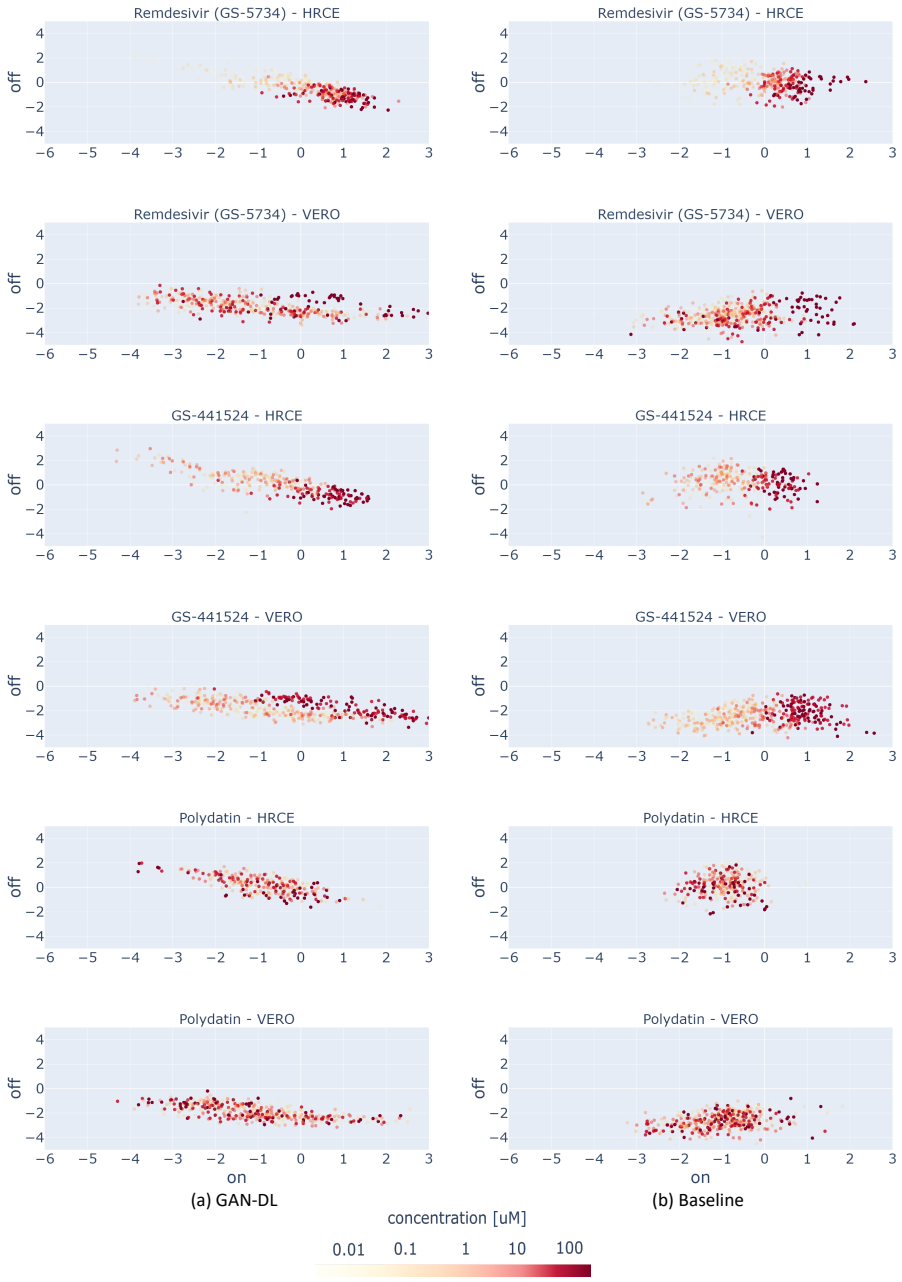

Figure 3.5: Drug effectiveness as a function of concentration, obtained using our GAN-DL (a) and the baseline embedding (b).

the figure, respectively), meaning that GAN-DL is equally able to represent the concentration-dependent ability of an active compound to preserve cell viability and inhibit SARS-CoV-2 infection.

Differently from the effective compounds, the ineffective ones should reasonably behave the same in terms of SARS-CoV-2 inactivation, independently of their concentration. When looking at the plot of *Polydatin* (a compound with no known effect on the virus in vitro), the values cluster towards the left side of the on perturbation axis where C^- samples are located and do not show any specific color-pattern at increasing values of dose concentration. This demonstrates that, same as for the baseline, with GAN-DL embedding the ineffective compounds do not show any specific dose-dependent behaviour. Accordingly, very few values of the ineffective compounds are located in the positive *On-perturbation* space (slightly greater then zero), suggesting no inactivation effect for SARS-CoV-2 infection in vitro.

3.9 Zero-shot representation learning

In this section we try to assess the generalization capabilities of the proposed embedding technique in a zero-shot representation learning experiment, that consists in a categorization problem where a classifier observes samples described by a featurization learnt not only on a different pretext task, but even on a different dataset.

For this purpose, we exploit the RxRx1 [60] image collection, a non-SARS-CoV2 related dataset consisting in 125,510 fluorescent microscopy images featuring human liver cancer cells (HEPG2), human umbilical vein endothelial cells (HUVEC), retinal pigment epithelium cells (RPE) and human bone osteosarcoma epithelial cells (U2OS). For the sake of channels compatibility, to perform a zero-shot inference on the RxRx1 [60] dataset we removed the channel corresponding to the MitoTracker, a dye that stains mitochondria, that is not present in the five-staining protocol of RxRx19a [59].

We exploit a soft margin linear SVM built on top of our GAN-DL embedding to categorize the four different cell models included in the RxRx1 [60] benchmark. We show the corresponding results in the form of a confusion matrix in Figure 3.6(a). From this matrix we can see that, despite the fact that the RxRx1 [60] cell models are totally new for GAN-DL, they can be linearly separated in the feature space with a mean accuracy of 92.68%.

For comparison, we show the results obtained by: (i) a DenseNet CNN model pre-trained on ImageNet, respectively with collapsed and concatenated strategy; (ii) a convolutional auto-encoder [32] trained on RxRx19a. As shown in the confusion matrices of Figure 3.6, both the DenseNet-based classifiers (ImageNet-collapsed and ImageNet-concatenated) and the convolutional auto-encoder (ConvAE) obtained an accuracy at least 5% lower than our GAN-DL.

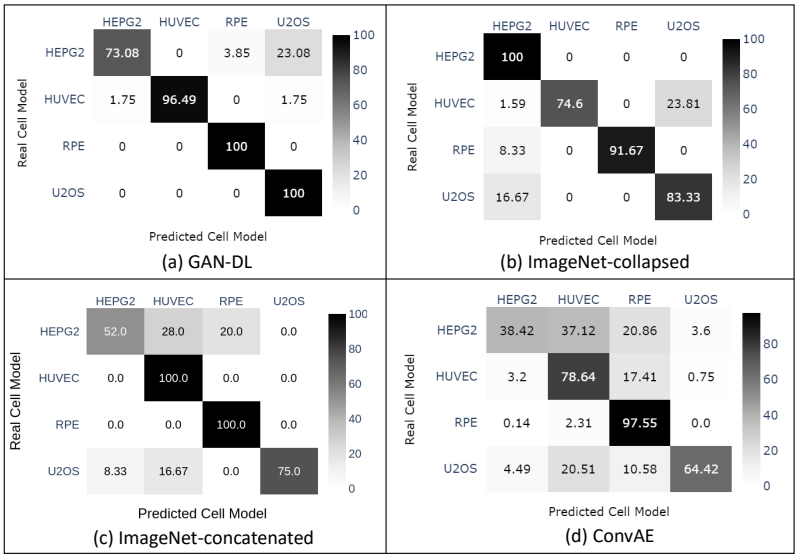

Figure 3.6: Confusion matrix of the cell classification task on the RxRx1 [60] dataset.

3.10 Discussion

In contexts where dataset annotation is costly, like medical and computational biology domains, the current standard, for the application of deep learning models on image data, involves the use of a ImageNet-pretrained CNN model [22], [37]. If, in general, SSRL seems a promising solution for those scenarios suffering a paucity of labelled data, the recent work by Wallace et al.[32] has shown how traditional SSRL featurization methodologies fail in several biological downstream tasks. This is mainly imputed on the difficulty in defining a pretext task which can be exploited by traditional contrastive SSRL.

On top of these considerations, in this chapter we propose GAN-DL, a fully SSRL method leveraging the representation learning acquired by the discriminator of a StyleGAN2 model [58]. Our GAN-DL does not require any task-specific label to obtain the image embedding, as the StyleGAN2 backbone is trained on a generative task based on the competition of a generator and of a discriminator, that is completely independent on the downstream task. By doing so, we address the problem of lack of annotated data, that is instead necessary for conventional CNN-based transfer learning methods. We demonstrated the goodness of our featurization methodology in two downstream supervised tasks: the classification of different cellular models (HRCE versus VERO cells) and the categorisation of positive versus negative control groups in the RxRx19a [59] benchmark. For this purpose, we trained a simple linear SVM on top of the self-supervised GAN-DL embedding, which does not require a large number of annotated data. Furthermore, we compared our solution with a baseline

state-of-the-art DenseNet model, pre-trained on the RxRx1 dataset [60] (the corresponding embedding is released together with the imaging data by [66]).

On the one hand, the baseline embedding is generally more accurate than GAN-DL in the downstream classification tasks, even though by a small margin. On the other hand, the baseline is pre-trained on a very large annotated dataset (RxRx1[60] dataset, consisting of 296 GB of fluorescent microscopy images), while training GAN-DL does not require any task-specific image annotations. This is indeed a major advantage for the re-usability of our method in different contexts where annotated data from a similar domain are few or even not available at all, which is a frequent challenge of many biological applications [22], [32].

We furthermore compare our GAN-DL with ImageNet-pretrained models, traditionally exploited as fixed feature extractor for biological images [22], as well as with an other SSRL method based on convolutional auto-encoder. We found our GAN-DL superior to both the ImageNet-based strategies and to the convolutional auto-encoder, the latter resulting less accurate with respect to ImageNet-based strategies by a narrow margin. We believe that the quality of representations extracted by the convolutional auto-encoder approach is less competitive than the ImageNet-based methods due to the limited capability of the auto-encoder in generating high-quality images (see some illustrative examples in the supplementary materials). The goodness of the results obtained by our GAN-DL, whose backbone is StyleGAN2 [58], state-of-the-art technology in image generation, corroborates the insight that the SSRL adversarial pretext task of learning to generate high-quality synthesised images allows to extract an SSRL representation featuring inherent relations that are not captured by previous techniques.

We speculate that our GAN-DL embedding, leveraging as pre-text task the generation of plausible and high resolution images through the adversarial game between the generator and the discriminator, proficiently learns an unbiased image featurization able to describe the fine-grained patterns that are typical of biological applications. This leads to an improved capability of separating multiple unrelated features along different axis of the latent space, which should be ultimately helpful to address any downstream tasks requiring knowledge of the salient attributes of the data [76]. To demonstrate our claim, we put this capability of GAN-DL into effect in a number of different applications: (i) the classification of active and inactive compounds against SARS-CoV-2 infection in two different cell lines; (ii) the generation of dose-response curves for the large scale molecule screening of RxRx19a [59], without the need of any training on purposely labelled data; (iii) the zero-shot representation learning of four different cell lines included in the RxRx1 [60] dataset.

In conclusion, the satisfactory results obtained in all the presented scenarios on the one hand demonstrate the goodness and generalisation capability of our approach, on the other hand legitimise the future exploitation of generative SSRL even in other biological applications, where the collection of annotated images is typically a cumbersome task.

Chapter 4

Development of a visual opsono-phagocytosis screening assay for monoclonal antibodies against *Neisseria gonorrhoeae*

Note: this Chapter is part of a manuscript in preparation.

Monoclonal antibodies (mAbs) can fight infectious diseases caused by pathogenic bacteria by exerting different functions, such as bactericidal activity, enhancement of phagocytosis or inhibition of adhesion [77]. In this chapter we present a high-throughput screening assay that identifies functional monoclonal antibodies (mAbs) against bacterial pathogens using fluorescent microscopy images. In particular we identify mAbs that promote phagocytosis of bacterium *Neisseria gonorrhoeae* by macrophages in single-point dilution experiments. In the preceding chapter, we illustrated the feasibility of conducting high-throughput compound screening through the use of fluorescent microscopy images and generative deep learning, enabling the derivation of a score that quantifies eukaryotic cell viability. In this chapter, we extend and refine this technology to create an opsono-phagocytosis assay that utilizes confocal microscopy and image analysis to measure the influence of monoclonal antibodies on the phagocytosis of *Neisseria gonorrhoeae* bacteria by macrophages. In particular, we fine-tuned a Dense Convolutional Network (DenseNet)[78] to classify images of positive and negative controls, and, as described in Chapter 3, used a simple linear Support Vector Machine (SVM) to quantify mAbs phagocytosis promoting activity by the Phagocytic score value. We named the proposed approach, presented in this chapter, visual opsono-phagocytosis assay (vOPA). Comparing the vOPA with the conventional protocol based on CFU counting, we observed, in a concentration-response experiments, that the results for the assays positive an negative controls were in agreement and that the variability we measured in EC_{50} estimation was one order of magnitude lower for the proposed assay. We have then used the vOPA as high-througput single point dilution "hits" selection assay to screen 96 human monoclonal antibodies that target *N. gonorrhoeae* and rank the best candidate mAbs based on their Phagocytic score. Furthermore, the flexibility of the staining protocol and the advantage of the deep-learning approach in utilising weakly labels derived from the experimental protocol, indicate the potential for expanding the assay to other bacterial species and cell lines to investigate mAbs against additional bacterial pathogens.

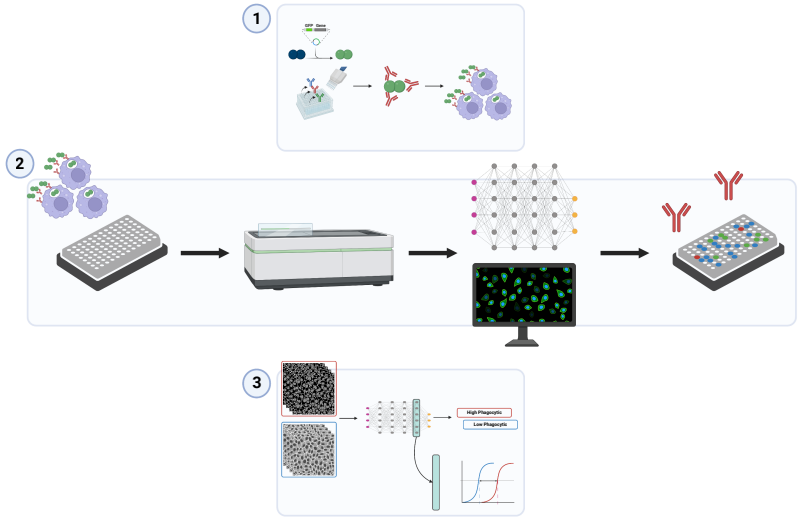

Figure 4.1: **The vOPA schematic representation**. (1) *N. gonorrhoeae* is engineered to express GFP and mAbs are expressed. Each mAb is then incubated with the fluorescent *N. gonorrhoeae* and used to infect differentiated THP-1. (2) The mAb/bacteria mixture is transferred into a 96-well plate for fluorescent imaging and used to infect pre-seeded THP-1 cells. Following staining and fixation, images are acquired with the Opera Phenix microscope and (3) analysed fine-tuning a CNN model to classify experimental positive and negative controls. The obtained score is used to quantify the phagocytic activity of cells in presence or absence of mAbs. Ultimately, the most prominent candidates are selected. *Image created with BioRender.com*

4.1 Introduction and motivation

Antimicrobial resistance (AMR) refers to the ability of microorganisms, primarily bacteria, viruses, fungi, and parasites, to survive exposure to drugs developed to kill them or limit their growth. As a result of AMR, standard treatments become less effective, and infections may become harder or impossible to control and more expensive to treat, thus increasing the risk of spread of disease and death. AMR is estimated to cause 33,000 deaths annually in the European Union and European Economic Area, at least 23,000 deaths annually in the United States [79]. Unfortunately, estimates suggest that AMR could cause 28 million people to fall into poverty worldwide due to an increase in healthcare costs of up to US$ 1 trillion globally by 2050 [80].

Addressing AMR involves various approaches. Some potential solutions include: effective implementation of administration principles and reduction of unnecessary or excessive uses of antibiotic, employing advanced diagnostic

tools, and promoting proper sanitation and hygiene, particularly in low-to-middle income areas, developing new antibiotics and alternative treatments like monoclonal antibodies. While these methods hold promise, tackling antimicrobial resistance requires an integrated approach since one solution on its own is unlikely to adequately mitigate this increasingly pressing public health concern [81]. Scientists have primarily focused on small-molecule drugs in their attempts to tackle AMR. However, recently, attention has shifted towards exploring biological approaches. A particularly effective strategy to fight against AMR is the identification of natural, human mAbs from patients that survived a specific infection, accelerating the discovery, reducing the risks of side-effects and increasing access to the therapy [13]–[15]. Key to this strategy is an effective high-throughput screening technology, to quickly identify those few, most potent natural monoclonal antibodies among the many thousands potentially isolated from each patient. The screening should identify monoclonals for their ability to confer disease protection. mAbs exert different modes of action when facing bacterial infections: they can trigger the complement cascade, induce agglutination of pathogens or, in the presence of specialised immune cells like macrophages or neutrophils, can enhance their intrinsic opsono-phagocytic activity [77]. Phagocytosys is defined as cellular uptake of particulates (>0.5 μm) within a plasma-membrane envelope. In particular, antibody mediated phagocytosis is the process by which a pathogen is marked for ingestion and eliminated by a phagocyte. The Fab region of the antibody binds to the antigen, whereas the Fc region of the antibody binds to an Fc receptor on the phagocyte, facilitating phagocytosis [82].

Unfortunately, currently available phagocytosis assays have limitations; they are low-throughput, requiring two-days procedures to screen mAbs [83], suffer from significant variability as most cellular assays, and require a unique setup specific to each pathogen species, which cannot be generalised across different microbes. Here we demonstrate that high-throughput fluorescent microscopy and deep-learning based image analysis enable the screening of a large number of candidates mAbs for opsonophagocytosis, reducing the variability of cell-based assay and potentially allowing the proposed set-up to be reused for other cell lines and other pathogens. To exemplify its effectiveness we focus on *N. gonorrhoeae*, the causative agent of Gonorrhea, the sexually transmitted infection that affects more than 100 million people annually, has remained a major global public health and AMR concern [84]. In fact, no vaccine against *N. gonorrhoeae* is currently available, and vaccine development has proven complicated in the past few decades [85]. Moreover, *N. gonorrhoeae*, has seen an alarmingly high number of resistant cases being reported [86], [87]. To limit the rising threat posed by N. gonorrhoea as an AMR pathogen, in spite of the lack of licensed vaccine for Gonorrhea prevention, we developed a functional high-throughput screening assay for the evaluation of antibody activity in host-pathogen interaction setting. In fact, the opsono-phagocytosis-promoting activity is considered as a major predictor of antibody protective efficacy [88] and an important feature of a candidate mAb for therapeutic use. Recently, the interactions of *N. gonorrhoeae* with two human monocytic cell lines, THP-1 and U937, differentiated into

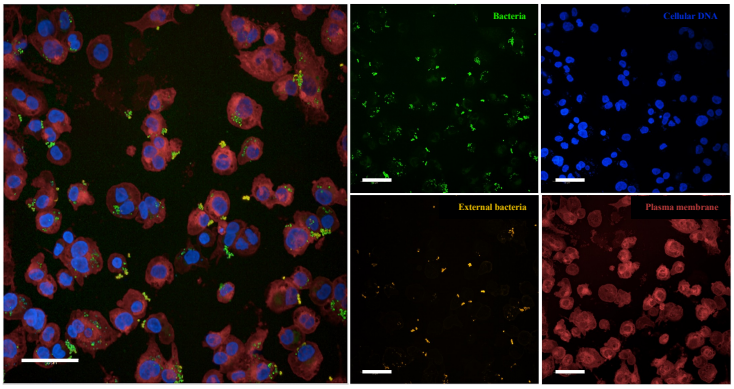

Figure 4.2: **The vOPA assay in THP-1 cells infected with *N.gonorrhoeae*.** The small pictures display the four channels imaged in the vOPA assay protocol as imaged using the Opera Phanix microscope: sfGFP, expressed by *N.gonorrhoeae* (green), DAPI, to stain the nuclei and bacterial DNA, (blue), Alexa 568 conjugated secondary antibody for immunostaining (orange) and CellMask Deep Red, to stain cell membranes, (red). Scale bar is 50 μm.

macrophages, were reported in [89]. Authors described the bacterial survival assays they used to quantify the number of macrophage internalised bacteria, released after cell lysis. The described approach, based on colony forming units (CFU) counts, was used to describe the interaction of the bacterium and macrophages alone, without the evaluation of compounds able to modulate that interaction.

In this chapter we adapt the methodology developed in [89] to quantify the phagocytosis promoting activity of mAbs against *Neisseria gonorrhoeae*. The adapted methodology relies on confocal microscopy and image analysis to overcome the long incubation time otherwise necessary using the CFU counting assay.

4.2 Methods

4.2.1 High throughput mAbs expression using TAP

To produce the anti-*Neisseria gonorrhoeae* monoclonal antibody, expression vectors containing the genes for its heavy and light chains were utilised in Transcriptionally Active Polymerase Chain Reaction (TAP)[46] to generate linear DNA fragments. These fragments were subsequently employed to conduct transient transfections of Expi293F cells, ensuring that there was a 1:2 ratio between the heavy and light chains. This procedure continued for six days, undergoing incubation at 37°C with 8% CO2, and consistent with the manufacturer's guidelines(Thermo Fisher Scientific, US). Following completion of the transfection period, the cultures were subject to centrifugation at 4,500 x g

Table 4.1: **vOPA infection condition set-up**. Median and standard deviation for the number of internalised bacteria per infected cells for the 2C7 and Unrelated mAbs are reported.

MOI	Infection time	mAb	Median	Std
20	30 min	2C7	2.84	0.22
		Unrelated	1.0	0.18
20	1 h	2C7	4.04	0.52
		Unrelated	1.85	0.18
40	**30 min**	2C7	4.7	0.44
		Unrelated	1.38	0.19
40	1 h	2C7	7.36	0.72
		Unrelated	2.83	0.23

and kept static for 15 minutes at 4°C, allowing them to settle. Post-centrifugation, only the cellular supernatants were retained and saved for later use.

4.2.2 Conventional mAbs expression and purification procedure

To achieve conventional medium-level expression and purification of the anti-*Neisseria gonorrhoeae* monoclonal antibody, the necessary heavy and light chain genetic sequences were inserted into expression vectors which were then used in transient transfection of Expi293F cells within a final volume of 60 ml. The subsequent antibody molecules were purified by affinity chromatography on protein G columns performed on an AKTA-Go system supplied by GE Healthcare Life Sciences.

4.2.3 vOPA experimental protocol

For our protocol, we engineered the *N.gonorrhoeae* strain FA1090 to constitutively express fluorescent protein GFP (FA1090::sfGFP). Then the THP-1 cells were seeded and differentiated into 96-well microplates for confocal imaging. After 5 days of differentiation, dTHP-1 cells were infected with FA1090::sfGFP grown to mid- logarithmic phase and pre-incubated with mAb supernatants diluted 1:5 in RPMI medium. After 30 minutes of pre-incubation, the mixture composed of mAbs and bacteria was added onto dTHP-1. To study in-vitro the antibody mediated enhancement of phagocytosis two crucial factors must be meticulously calibrated: the number of bacteria used to infect the target cell line and the time of infection. The number of bacteria is conventionally quantified by the multiplicity of infection (MOI) as the ratio of the number of bacteria to the number of target cells present in the well. To select the best conditions and observe antibody mediated phagocytosis in our set-up, we tested different MOIs, 20 and 40, and time of infections, 30 minutes and 1 hour.

According to the data shown in Figure 4.3 and Table 4.1, we found that MOI of 40 and duration of infection of 30 minutes yielded the optimal conditions

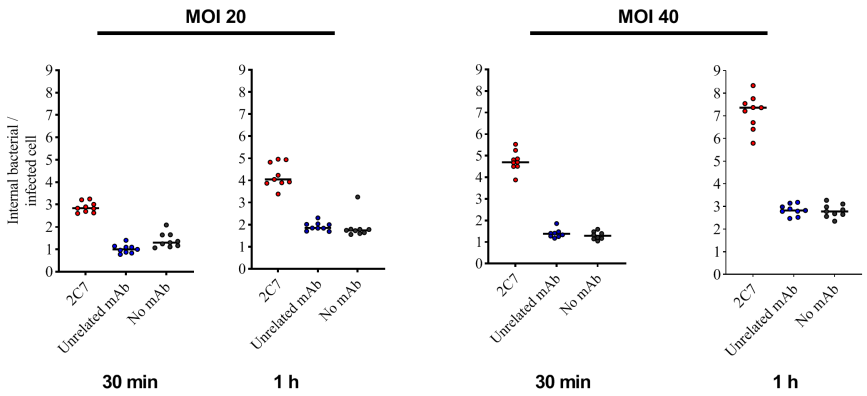

Figure 4.3: **Evaluation of two mAbs activity at different MOIs and times of infections**. The graphs reports the number of internal bacteria per infected cell for the conditions MOI 20 and MOI 40, and at two different infection time, 30 minutes and 1-hour. The 2C7 and the unrelated monoclonal were purified and tested at a concentration of $10\mu g/ml$. For each condition 9 technical replicates were performed.

needed for this assay setup. At these settings, we could effectively amplify assay range and keep variability levels low enough without disrupting normal biological variations in uninfected test samples.

4.2.4 Gentamicin protection assay (classical OPA assay)

THP-1 cells were seeded in 96-well tissue culture plates at 40,000 cells per well and subjected to the differentiation and infection protocol as described above. For determining the amount of pathogens internalised by the cells, extracellular bacteria were washed away three times with RPMI medium and gentamicin (100 μg/ml) was added for 30 minutes to kill adherent bacteria. Subsequent to the elimination of surface bound bacteria, the cells were lysed with 0.5% saponin for 5 minutes, permitting release of intra-cellular microbes. Dilutions from this suspended mixture were dispersed onto solid media (GC agar) to facilitate visualisation and counting of colony forming units, with incubation at optimal temperature ranges for subsequent 24 hours - 48 hours prior to enumerating counts.

4.2.5 vOPA stainings

In Figure 4.2 we show a representative example of the images we generated for the vOPA assay. In particular we used four fluorescent stainings to identify the whole bacteria population, the internalised bacteria, cellular DNA and plasma membrane. Extracellular FA1090::sfGFP bacteria were stained with primary

antibody 2C7, at final concentration of 3 mg/ml , for 1h at RT, followed by secondary goat anti- Human IgG Alexa Fluor 568 (Thermo Fisher, A-21090) diluted 1:2,000 at RT for 30 minutes. CellMaskTM Deep Red stain (Invitrogen) was used to stain the cell membrane, providing a means to delineate the cell boundary, and DAPI to stain cell nuclei and bacterial DNA.

4.2.6 Confocal microscopy image acquisition

96-well plates were imaged with the microscope Opera Phenix High-Content Screening System (PerkinElmer) using the 40x water objective (numerical aperture 1.1). In each well, 16 images were acquired in different locations to account for minor variability in the sample. Moreover in each location, 13 images on the vertical dimension were acquired to form a z-stack.

4.2.7 vOPA image dataset

The vOPA image dataset consists of 9.048 4-channel images generated in 13 different experiments performed in 96-well imaging plates. A single dataset's image corresponds to the maximum projection performed over the 13 images acquired on the vertical dimension. The raw images, with full resolution and original colour depth (one TIFF file per channel, 16 bit grayscale, lossless compression) and shape of 2160x2160 pixels (px), were transformed into 8 bit images of size 512x512 px to be processed with the Convolutional neural network (CNN) and derive the corresponding Phagocytic score, as described in the next section.

4.2.8 Deep Learning model and phagocytic score

We downloaded and fine-tuned a DenseNet161[78] model from the pytorch[90] hub repository. The fine-tuning was performed as a binary classification task on the images acquired at the Opera Phenix, where the two classes were represented by cropped images from positive and negative controls respectively. To process the 4-channels vOPA images, we added a convolutional layer at the top of the pre-trained model. As optimizer we used SGD with 0.001 learning-rate and 0.9 momentum. We trained the model for 15 epochs. No hyperparameter optimization was performed. For the training we used 1 NVIDIA A100 40GB GPU. To increase the amount of training data for our DenseNet model, we employed a data augmentation technique that involves cropping the input image. Specifically, we used a patch size of one half and a stride of one fourth of the input image respectively, generating with this process nine cropped images. These augmented images were then used to fine-tune the model. After training we exploited a simple soft margin linear SVM on top of the DenseNet last layer embeddings to derive a two-dimensional space to capture the differences between mAbs promoting phagocytosis and unrelated mAbs. The direction orthogonal with respect to the separating hyperplane fitted by the linear SVM was named Phagocytic score (Figure 4.4).

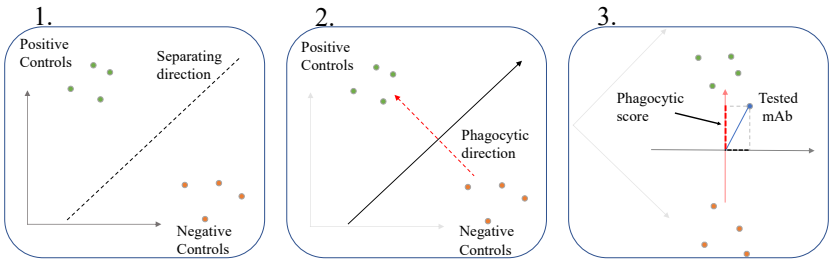

Figure 4.4: **Phagocytic score using linear SVM on networks embedding**. 1. Derive the separating direction (dashed line) of a binary classifier tasked to recognise positive and negative controls. 2. Compute the orthogonal direction with respect to the separating direction. This results in being the Phagocytic direction (red dashed line). 3. Testing a new monoclonal, the blue dot, we obtain the Phagocytic score measuring the orthogonal projection on the phagocytic direction.

4.2.9 Harmony software image analysis

We used the Harmony High-Content Imaging and Analysis Software version 4.9 to quantify the number of internalised bacteria per infected cells used to set-up the vOPA infection conditions. In a first phase, dTHP-1 cells were detected by analysing the combined signal of DAPI (nuclei) and CellMask (membrane) In a second phase, bacteria were localised and segmented by combining the DAPI and GFP signals. Bacteria that overlapped with the image regions segmented as cells were identified as infecting bacteria and further distinguished as internalised if they were negative for the immunostaining and adherent if they were positive for the immunostaining used in the vOPA staining protocol. Ultimately, cells that contain infecting bacteria were considered as infected cells.

4.3 Results

4.3.1 The phagocitic score separates positive and negative controls with TAP-expressed mAbs, achieving acceptable Z' assay value

In order to assess the ability of the vOPA phagocitic score defined in Section 4.2.8 to distinguish positive and negative control samples, we tested 18 replicates for three conditions, the two controls (2C7 and unrelated mAb) and the condition without any mAb (no-mAb), to account for possible variability of the mAbs in TAP expression procedure. Figure 4.5 reports on the y axis the value of the Phagocytic score that clearly separetes the positive 2C7, in red, from the unrelated mAb and the no mAb conditions, in blue and gray respectively. The corresponding statistical difference was quantified using the Student's t-test,

the p-values significance are reported on the corresponding boxes, where $****$ represent a p-value < 0.001 and *ns* a non-statistically significant difference. As no statistical difference was observed for the Unrelated and No mAb conditions, the vOPA resulted to be a specific assay to quantify antibody mediated *N. gonorrhoeae* phagocytosis in THP-1 cell line.

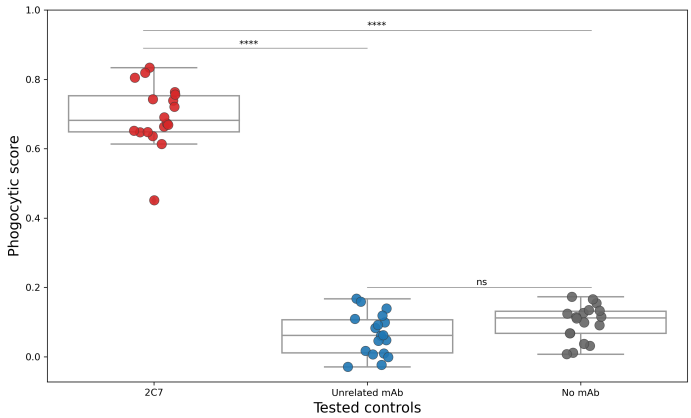

Figure 4.5: **Control separability for non-purified mAbs**. 2C7, Unrelated mAbs and No mAb condition were tested in 18 technical replicates. The phagocytic score is reported on the x-axis. On the graph the unpaired non-parametric t-test (Mann-Whitney test) p-value significance are reported.

Furthermore, we quantified the validity of vOPA as single point screening assay, by computing the Signal window and Z' performance measures. In Table 4.2, we report measures results and corresponding acceptance ranges, as described in [91]. We observed that, the measured values for the assay Signal window and Z' were in the recommended and acceptable range respectively, demonstrating that vOPA can be used to screen monoclonal antibodies for their phagocytosis-promoting activity in single point dilution experiments.

4.3.2 The vOPA Phagocytic score is linearly correlated with the cOPA CFU counts

To further explore a biological interpretation of the Phagocytic score, in Figure 4.6, we show that the score has a linear correlation with the phagocytosis assay based on CFU counting in dose-dependent experiment. In particular, we quantify the 2C7 phagocytosis-promoting activity in two dose-dependent experiments. In one case we quantify the mAbs phagocytic effect by the vOPA and in the other case by the classic OPA (cOPA), described in Method section 4.2.4 The 2C7 mAb was expressed, purified and starting from a concentration of $50\mu g/ml$ was diluted 1:10 for 5 steps. We verified that the vOPA Phagocytic score is linearly

Table 4.2: **Controls separability performance measures**. Phagocytic score separation measured between positive (2C7) and negative (Unrelated mAb) controls.

Assay performance measures	**Value**	**Acceptance Range** [91]	
Signal window	6.23	**Recommended**	> 2
		Acceptable	> 1
		Unacceptable	< 1
Z' factor	0.3	Excellent	> 0.5
		Acceptable	$0 < Z' < 0.5$
		Unacceptable	< 0
Signal	0.63		
2C7 std	0.09		
Unrelated mAb std	0.06		

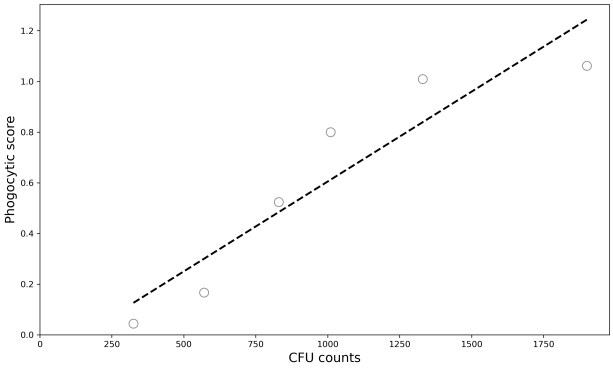

Figure 4.6: **Linear correlation for the 2C7 values measured in vOPA and cOPA**. The values in the two assay result to be linearly correlated with a significant p-value of p=0.007.

correlated with the cOPA CFU count by computing the Pearson correlation coefficient of the 2C7 values. We obtained an R squared value of 0.866 with a significant p-value of 0.007.

4.3.3 The phagocytic score of the vOPA assay show the same dose-response trend of the cOPA CFU counts

To further conferm the Phagocytic score interpretability as an opsonization assay, we compare 2C7 and Unrelated mAb functionality in vOPA and cOPA. Both the 2C7 and the unrelated antibody were tested at a starting concentration of $50\,\mu g/ml$ and diluted 1:10 for 5 steps. The results are reported in Figure 4.7. We

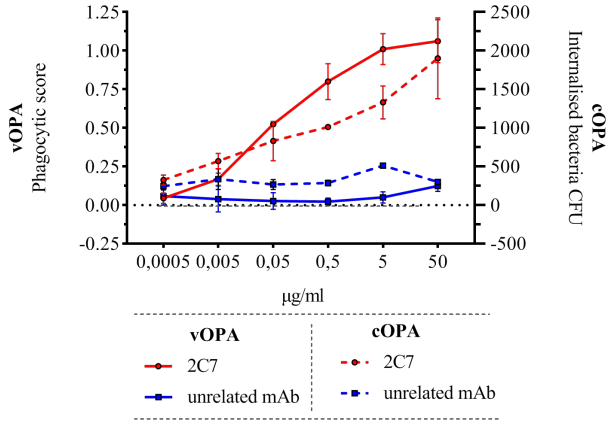

Figure 4.7: **2C7 and unrelated mAb dose response in vOPA and cOPA**. The two antibodies were tested in a 10-fold dilutions experiment. The vOPA read-out is the phagocytic score obtained with the CNN approach described in the method section(left y-axis), the cOPA read-out is the number of internalised bacteria (right y-axis). Concentrations are shown on the x-axis.

can appreciate that the 2C7 and the unrelated mAb displayed similar trends in cOPA, red lines, and vOPA, blue lines. In particular, a dose-dependent promotion of opsono-phagocytosis by 2C7 and no effect mediated by the unrelated antibody were observed.

4.3.4 The vOPA assay with positive control mAb shows ten times higher repeatability than cOPA

To get an estimate of the repeatability of our assay, we evaluated the variability of EC 50 computing the coefficient of variation. By fitting a three-parameter sigmoidal dose-response function to the data shown in Figure 4.7, we obtained the potency (EC50) of the 2C7 measured by the two assays. We report the complete results in Table 4.3. We observe a difference in the EC50 estimate in two assay, in particular vOPA estimate a lower EC50. However, considering the confidence intervals obtained for the estimates, the EC50 values for the two assays are largely compatible. It is worth noting that the confidence intervals associated with determining EC50 values through cOPA procedures appear substantially broader than their counterparts obtained employing vOPA approaches.

Table 4.3: **EC50 estimation and precision measures for vOPA and cOPA**. The EC50 for the two assay has been derived by fitting a three-parameter sigmoidal dose-response function to the 2C7 data. Based on EC50 confidence intervals we derived the coefficient of variation for the 2C7 EC50 in the two assays as a measure of precision.

Measure	**2C7 (vOPA)**	**2C7 (cOPA)**
EC50 Coefficient of variation	2.31%	20.0%
EC50	0.06 $\mu g/ml$	2.031 $\mu g/ml$
EC50 95% CI min	0.025 $\mu g/ml$	0.032 $\mu g/ml$
EC50 95% CI max	0.165 $\mu g/ml$	40.66 $\mu g/ml$

4.3.5 High-throughput screening of 96 TAP-expressed mAbs shows phagocytic score values spanning uniformly the assay range

We then leveraged the high-throughput image acquisition capability of the Opera Phenix system to screen an array of human mAbs for their ability to promote phagocytic activity against *N. gonorrhoeae*. In this experiment, the mAbs, expressed as recombinant protein, were not further purified after supernatant collection, , according to the high-throughput procedure described in 4.2.1. Data generation required five hours, one for sample preparation and infection, three for the image acquisition and one for the analysis. An array of 96 human recombinant anti-*N. gonorrhoeae* mAbs of unknown concentration was tested, at single dilution, to select positive hits. In Figure 4.8, we report the vOPA Phagocytic score for the 96 mAbs and 6 replicates of the positive and negative controls. Based on the assay signal (Table 4.2), we separated the Phagocytic score range in three intervals of equal amplitude representing low-, moderate- and high-phagocytosis-promoting monoclonals. Out of 96 tested mAbs, 51% were low, 38% moderate and only the 11% of the tested mAbs showed high-phagocytosis-promoting activity.

4.3.6 Dose-response curves from purified mAbs support the use of vOPA as a high-throughput screening assay

Based on the Phagocytic score we selected two candidates mAbs, one from the high-phagocytosis promoting group, and one from the moderate phagocytosis-promoting group. Both mAbs were expressed as purified proteins and tested in a 10-fold dilution experiment starting from a concentration of 5 $\mu g/ml$ to confirm their activity. In Figure 4.9 we report the Phagocytic score value for the high and moderate phagocytosis promoting mabs as red and green dots respectivly, the 2C7 and the unrelated mAb Phagocytic score values are reported as red and blue squares respectively. We observed a concentration-dependent effect for the two selected mAbs and the 2C7. For the two selected mAbs

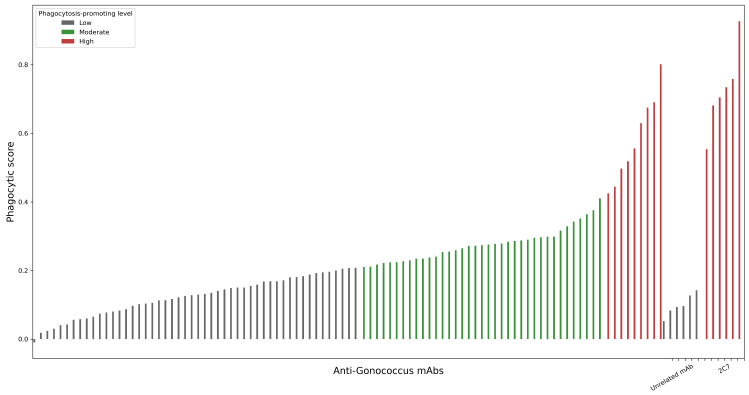

Figure 4.8: **Screening of 96 mAbs using vOPA**. Each mAb was tested at one single dilution, on the y-axis the corresponding phagocytic score is reported. Colours represent groups of phagocytosis-promoting activity, in white no/low activity, in green moderate activity and in red high phagocytosis promoting activity respectively. The groups where defined based on phagocytic score range reported in table 4.2. The negative (Unrelated mAb) and positive (2C7) controls phagocytic scores, reported on the right part of the graph, are coloured according to their phagocytosis-promoting activity.

Table 4.4: **EC50 and Span values for 2C7, high- and moderate-phagocytosis promoting mAbs**.

	2C7	**High**	**Moderate**
EC50	0.049 $\mu g/ml$	0.041 $\mu g/ml$	0.029 $\mu g/ml$
Span	0.905	0.757	0.430

and the 2C7, a 3-parameter sigmoidal dose-response function was fitted to the Phacocytic score values, reported on the y-axis to quantify the difference in the phagocytic-promoting activity in terms of EC50 and Span. The values of EC50 and spans, reported in Table 4.4 supports the groups defined in Section 4.3.5 and Figure 4.8.

4.3.7 vOPA can be used in combination with TAP expression to screen mAbs concentrated as low as 0.5 mug/ml

Figure 4.9 shows that the Phagocytic score is capable of distinguishing active from unrelated mabs, as well as moderate from high-phagocytosis-promoting mAbs, for concentrations as low as 0.05 $\mu g/ml$. For lower concentrations, our score may not be able to discriminate. Figure 4.10 shows that 88 (92%) of

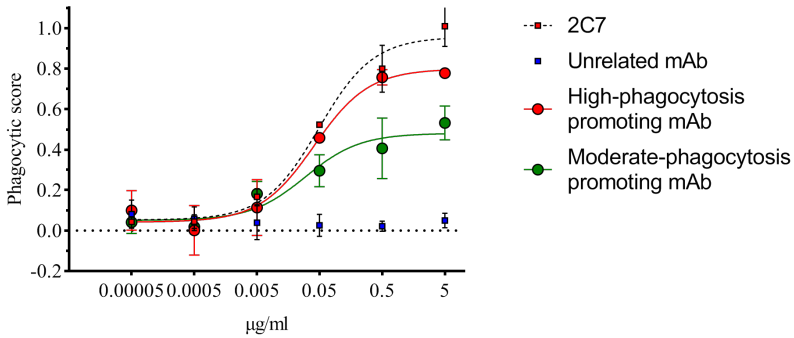

Figure 4.9: **Candidate mAbs dose response evaluation**. 2C7, high- and moderate-phagocytosis promoting mAbs were tested in 1:10 serial dilution experiment, from $5\mu g/ml$ for 5 steps. A 3-parameter sigmoidal dose-response function was fitted to the Phacocytic score values, reported on the y-axis. The resulting EC50 and Span is reported in table 4.4.

the TAPs tested had a concentration higher than 0.05 $\mu g/ml$, supporting the use of the Phagocytic score as a high-throughput screening indicator of human monoclonal antibody mediated phagocytosis activity.

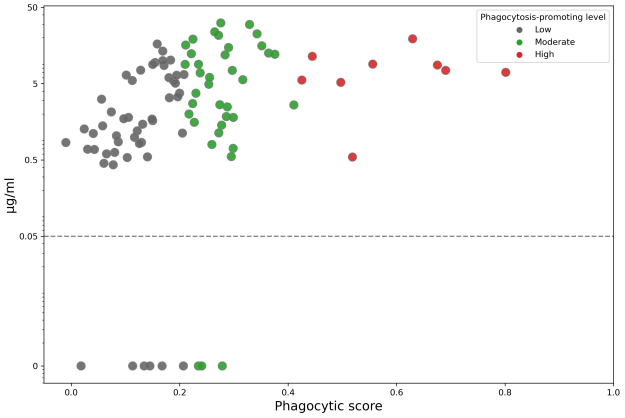

Figure 4.10: **Concentration quantification for the 96 anti-gonococcus mAbs**. The mAbs Phagocytic score is reported on x-axis, and the concentration, measured via quantitative ELISA is reported on the y-axis.

4.4 Discussion

4.4.1 Using high-througput microscopy is possible to reduce the conventional procedure from two days to 5 hours

In this chapter we presented the development and validation of the vOPA, an assay to quantify the opsonophagocytosis-promoting activity of monoclonal antibodies against *N. gonorrhoeae* from fluorescent images. The staining panel, the infection protocol, the image acquisition strategy and the image analysis pipeline were successfully optimised and used for single point dilution screening of human monoclonal antibodies in 96-well format. Using vOPA in combination with an automated confocal microscopy platform, i.e. Opera Phenix, is possible to prepare a 96-well plate of mAbs and acquire the images in less than four hours. We do not have a direct comparison for the CFU based assay (cOPA), in fact we used the cOPA for two monoclonals, the 2C7 and the Unrelated mAb. However, CFU counting is not a one-day procedure, in fact, in case of *N. gonorrhoeae*, the colonies were counted after 24–48 h of incubation time. In this sense, the vOPA is a faster procedure with respect to the conventional CFU counting approach, resulting suitable for high-throughput single point evaluation of human monoclonal antibodies.

4.4.2 The phagocytic score is derived with deep learning and information from the experimental protocol, resulting in an analysis package easy to adapt to new pathogens

For the vOPA we have developed two analysis approaches, one based on segmentation building blocks, available in the Harmony software package, and a second based on fine-tuning a convolutional neural network. Both demonstrated acceptable control separability, as shown in Figure 4.3 and 4.5 respectively. However, the set-up of the segmentation pipeline requires careful parameter setting and adjustment, validated by visual evaluation from the human-user. This strategy is heavily dependent on the image quality and even small sample impurities can alter the segmentation step of bacteria. On the other hand, to fine tune the CNN we only employ the information reported in the experimental protocol, i.e. the label of the sample tested in the well and corresponding concentration, if available. Those two information were used to define the label used to fine-tune the Densenet and train the SVM model used to derive the Phagocytic score. The two analysis approaches can serve also different needs, the deep-learning is more flexible and require less human intervention to obtain a reliable score, however has to be performed with a sufficient number of images at disposal. The segmentation based, if not accurately tuned, would result in a raw estimate of the phagocytosis but can be suitable to visually evaluate images during the image generation process, and early assessing issues during image generation.

Overall, the vOPA assay could be the starting point for image-based high-throughput approaches to investigate the activity of mAbs against bacteria

pathogens. Furthemore, given the feasibility of the staining protocol proposed, and the flexibility of the deep-learning analysis approach, we argue that the assay can be easily extended to other bacterial species and thus identify new mAbs against other bacterial pathogens.

4.4.3 Limitations and further developments

The proposed novel vOPA assay possesses significant biological interpretable features. We first confirmed the previously documented enhancement of pathogen uptake by the 2C7 monoclonal antibody. Furthermore, a positive association between the derived Phagocytic score and the CFU counts supports the reliability of this assay. While this discovery represents substantial progress, expanding upon it by investigating nonlinear relationships or variations unaccounted for in the present analysis may improve future applications. In addition, we obtained the Phagocytic score using all the data from the experiments detailed in Section 4.2.7. Investigation of the contributions made by individual fluorescent channels to the final Phagocytic score could aid our comprehension of the functions played by the stainings utilized within the assay. Additionally, results from such research could point towards ways to optimise or reduce the complexity of the labelling process.

Chapter 5

Conclusions

This thesis discusses two main topics. Firstly, the genetic features of a collection of monoclonal antibodies isolated from convalescent patience and able to neutralise the SARS-CoV-2 virus in vitro are described. Secondly, two computational approaches for extracting important features from biological images are presented.

In the second chapter we described a systematic genetic characterisation of human monoclonal antibodies against SARS-CoV-2. Most potently neutralizing antibodies are extremely rare and recognize the RBD, followed in potency by antibodies recognizing the S1 domain, the trimeric structure and the S2 subunit. We and others found that the antibody repertoire of convalescent patients is mostly germline-like. A potential issue associated with the use of human mAbs against viral pathogens is the potential selection of escape mutants. This is usually addressed by using a combination of antibodies directed against non-overlapping epitopes. While this is an ultimate clear solution, it increases the complexity of development, costs of production, drug availability, and affordability. In our case, we believe that selection of escape mutants upon treatment with a single mAb may be quite difficult as the SARS-CoV-2 RNA-dependent polymerase possesses a proofreading machinery, and the epitope recognized by the antibodies herein described overlaps with the region necessary to bind the hACE2 receptor. In this regard, it took more than 70 days of continuous co-culture of the virus in presence of the antibodies before we were able to detect the first emergence of escape mutants of the WT SARS-CoV-2 (data not shown). The nucleotide and amino acidic sequences of all SARS-CoV-2-neutralising antibodies were deposited in the Italian patent applications n. 102020000015754 filed on June 30th 2020 and 102020000018955 filed on August 3rd 2020.

In the third chapter we presented a self-supervised representation learning method for large scale drug screening initiative. If, in general, SSRL seems a promising solution for those scenarios suffering a paucity of labelled data, the recent work by Wallace et al.[32] has shown how traditional SSRL featurization methodologies fail in several biological downstream tasks. This is mainly imputed on the difficulty in defining a pretext task which can be exploited by traditional contrastive SSRL. On top of these considerations, we have developed GAN-DL, a fully SSRL method leveraging the representation learning acquired by the discriminator of a StyleGAN2 model. Our GAN-DL does not require any task-specific label to obtain the image embedding, as the StyleGAN2 backbone is trained on a generative task based on the competition of a generator and of a discriminator, that is completely independent on the downstream task. By doing so, we address the problem of lack of annotated data. We demonstrated the goodness of our featurization methodology in two downstream supervised tasks: the classification of different cellular models (HRCE versus VERO cells)

and the categorisation of positive versus negative control groups in the RxRx19a benchmark. For this purpose, we trained a simple linear SVM on top of the self-supervised GAN-DL embedding, which does not require a large number of annotated data. The satisfactory results obtained in all the presented scenarios on the one hand demonstrate the goodness and generalisation capability of our approach, on the other hand legitimise the future exploitation of generative SSRL even in other biological applications, where the collection of annotated images is typically a cumbersome task.

In the fourth chapter we presented the development and validation of the vOPA, an assay to quantify the opsonophagocytosis-promoting activity of monoclonal antibodies against *N. gonorrhoeae* from fluorescent images. The staining panel, the infection protocol, the image acquisition strategy and the image analysis pipeline were successfully optimised and used for single point dilution screening of human monoclonal antibodies in 96-well format. The image analysis consist in fine-tuning a convolutional neural network on the image of positive and negative controls, the 2C7 and the Unrelated mAbs respectively. Therefore, to fine tune the CNN we only employ the information reported in the experimental protocol, i.e. the label of the sample tested in the well and corresponding concentration, if available. Those two information were used to define the label used to fine-tune the Densenet and train the SVM model used to derive the Phagocytic score. As a result, the deep learning based score is flexible and require very limited human intervention to obtain a reliable Phagocytic score. Overall, the vOPA assay could be the starting point for image-based high-throughput approaches to investigate the activity of mAbs against bacteria pathogens. Furthemore, given the feasibility of the staining protocol proposed, and the flexibility of the deep-learning analysis approach, we argue that the assay can be easily extended to other bacterial species and thus identify new mAbs against other bacterial pathogens.

Bibliography

[1] Reid, M. J., Arinaminpathy, N., Bloom, A., *et al.*, "Building a tuberculosis-free world: The lancet commission on tuberculosis," *The Lancet*, vol. 393, no. 10178, pp. 1331–1384, 2019.

[2] Kim, H., Webster, R. G., and Webby, R. J., "Influenza virus: Dealing with a drifting and shifting pathogen," *Viral immunology*, vol. 31, no. 2, pp. 174–183, 2018.

[3] Deeks, S. G., Overbaugh, J., Phillips, A., and Buchbinder, S., "Hiv infection," *Nature reviews Disease primers*, vol. 1, no. 1, pp. 1–22, 2015.

[4] Kreuzberger, N., Hirsch, C., Chai, K. L., *et al.*, "Sars-cov-2-neutralising monoclonal antibodies for treatment of covid-19," *Cochrane Database of Systematic Reviews*, no. 9, 2021.

[5] Wang, M.-Y., Zhao, R., Gao, L.-J., Gao, X.-F., Wang, D.-P., and Cao, J.-M., "Sars-cov-2: Structure, biology, and structure-based therapeutics development," *Frontiers in cellular and infection microbiology*, vol. 10, p. 587 269, 2020.

[6] Drożdżal, S., Rosik, J., Lechowicz, K., *et al.*, "An update on drugs with therapeutic potential for sars-cov-2 (covid-19) treatment," *Drug Resistance Updates*, vol. 59, p. 100 794, 2021.

[7] Focosi, D., McConnell, S., Casadevall, A., Cappello, E., Valdiserra, G., and Tuccori, M., "Monoclonal antibody therapies against sars-cov-2," *The Lancet Infectious Diseases*, 2022.

[8] Merad, M., Blish, C. A., Sallusto, F., and Iwasaki, A., "The immunology and immunopathology of covid-19," *Science*, vol. 375, no. 6585, pp. 1122–1127, 2022.

[9] Earle, K. A., Ambrosino, D. M., Fiore-Gartland, A., *et al.*, "Evidence for antibody as a protective correlate for covid-19 vaccines," *Vaccine*, vol. 39, no. 32, pp. 4423–4428, 2021.

[10] Khoury, D. S., Cromer, D., Reynaldi, A., *et al.*, "Neutralizing antibody levels are highly predictive of immune protection from symptomatic sars-cov-2 infection," *Nature medicine*, vol. 27, no. 7, pp. 1205–1211, 2021.

[11] Corti, D., Purcell, L. A., Snell, G., and Veesler, D., "Tackling covid-19 with neutralizing monoclonal antibodies," *Cell*, vol. 184, no. 12, pp. 3086–3108, 2021.

[12] Zohar, T. and Alter, G., "Dissecting antibody-mediated protection against sars-cov-2," *Nature Reviews Immunology*, vol. 20, no. 7, pp. 392–394, 2020.

[13] Rappuoli, R., Bottomley, M. J., D'Oro, U., Finco, O., and De Gregorio, E., "Reverse vaccinology 2.0: Human immunology instructs vaccine antigen design," *Journal of Experimental Medicine*, vol. 213, no. 4, pp. 469–481, 2016.

[14] Zurawski, D. V. and McLendon, M. K., "Monoclonal antibodies as an antibacterial approach against bacterial pathogens," *Antibiotics*, vol. 9, no. 4, p. 155, 2020.

[15] Liu, J. K., "The history of monoclonal antibody development–progress, remaining challenges and future innovations," *Annals of medicine and surgery*, vol. 3, no. 4, pp. 113–116, 2014.

[16] Carpenter, A. E., "Image-based chemical screening," *Nature Chemical Biology*, vol. 3, no. 8, pp. 461–465, 2007.

[17] Caicedo, J. C., Singh, S., and Carpenter, A. E., "Applications in image-based profiling of perturbations," *Current opinion in biotechnology*, vol. 39, pp. 134–142, 2016.

[18] Bagheri, N., Carpenter, A. E., Lundberg, E., Plant, A. L., and Horwitz, R., "The new era of quantitative cell imaging—challenges and opportunities," *Molecular Cell*, vol. 82, no. 2, pp. 241–247, 2022.

[19] Boyd, J., Fennell, M., and Carpenter, A., "Harnessing the power of microscopy images to accelerate drug discovery: What are the possibilities?" *Expert opinion on drug discovery*, vol. 15, no. 6, pp. 639–642, 2020.

[20] Caicedo, J. C., McQuin, C., Goodman, A., Singh, S., and Carpenter, A. E., "Weakly supervised learning of single-cell feature embeddings," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9309–9318.

[21] Liu, X., Zhang, F., Hou, Z., *et al.*, "Self-supervised learning: Generative or contrastive," *IEEE Transactions on Knowledge and Data Engineering*, 2021.

[22] Ponzio, F., Urgese, G., Ficarra, E., and Di Cataldo, S., "Dealing with lack of training data for convolutional neural networks: The case of digital pathology," *Electronics*, vol. 8, no. 3, p. 256, 2019.

[23] Melinščak, M. and Lončarić, S., "Retinal oct image segmentation: How well do algorithms generalize or how transferable are the data?" In *2020 43rd International Convention on Information, Communication and Electronic Technology (MIPRO)*, IEEE, 2020, pp. 192–196.

[24] Ponzio, F., Deodato, G., Macii, E., Di Cataldo, S., and Ficarra, E., "Exploiting "uncertain" deep networks for data cleaning in digital pathology," in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, IEEE, 2020, pp. 1139–1143.

[25] Zhuang, F., Qi, Z., Duan, K., *et al.*, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2020.

[26] Farahani, A., Voghoei, S., Rasheed, K., and Arabnia, H. R., "A brief review of domain adaptation," *Advances in data science and information engineering*, pp. 877–894, 2021.

[27] Sajun, A. R. and Zualkernan, I., "Survey on implementations of generative adversarial networks for semi-supervised learning," *Applied Sciences*, vol. 12, no. 3, p. 1718, 2022.

[28] Van Engelen, J. E. and Hoos, H. H., "A survey on semi-supervised learning," *Machine Learning*, vol. 109, no. 2, pp. 373–440, 2020.

[29] Jing, L. and Tian, Y., "Self-supervised visual feature learning with deep neural networks: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 11, pp. 4037–4058, 2020.

[30] Radford, A., Kim, J. W., Hallacy, C., *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*, PMLR, 2021, pp. 8748–8763.

[31] *Aaai 2020 conference*, Accessed: 2021-04-29.

[32] Wallace, B. and Hariharan, B., "Extending and analyzing self-supervised learning across domains," in *European Conference on Computer Vision*, Springer, 2020, pp. 717–734.

[33] Goldsborough, P., Pawlowski, N., Caicedo, J. C., Singh, S., and Carpenter, A. E., "Cytogan: Generative modeling of cell images," *BioRxiv*, p. 227 645, 2017.

[34] Gildenblat, J. and Klaiman, E., "Self-supervised similarity learning for digital pathology," *arXiv preprint arXiv:1905.08139*, 2019.

[35] Sahasrabudhe, M., Christodoulidis, S., Salgado, R., *et al.*, "Self-supervised nuclei segmentation in histopathological images using attention," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2020, pp. 393–402.

[36] Dmitrenko, A., Masiero, M. M., and Zamboni, N., "Self-supervised learning for analysis of temporal and morphological drug effects in cancer cell imaging data," *arXiv preprint arXiv:2203.04289*, 2022.

[37] Yang, P., Hong, Z., Yin, X., Zhu, C., and Jiang, R., "Self-supervised visual representation learning for histopathological images," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2021, pp. 47–57.

[38] Arjovsky, M., Chintala, S., and Bottou, L., "Wasserstein generative adversarial networks," in *International conference on machine learning*, PMLR, 2017, pp. 214–223.

[39] Andreano, E., Nicastri, E., Paciello, I., *et al.*, "Extremely potent human monoclonal antibodies from covid-19 convalescent patients," *Cell*, vol. 184, no. 7, pp. 1821–1835, 2021.

[40] Lilly, E. *et al.*, *Lilly's neutralizing antibody bamlanivimab (ly-cov555) receives fda emergency use authorization for the treatment of recently diagnosed covid-19*, 2020.

[41] REGENERON, *Regeneron's casirivimab and imdevimab antibody cocktail for covid-19 is first combination therapy to receive fda emergency use authorization*, 2020.

[42] Clausen, T. M., Sandoval, D. R., Spliid, C. B., *et al.*, "Sars-cov-2 infection depends on cellular heparan sulfate and ace2," *Cell*, vol. 183, no. 4, pp. 1043–1057, 2020.

[43] Wang, Q., Zhang, Y., Wu, L., *et al.*, "Structural and functional basis of sars-cov-2 entry by using human ace2," *Cell*, vol. 181, no. 4, pp. 894–904, 2020.

[44] Walls, A. C., Park, Y.-J., Tortorici, M. A., Wall, A., McGuire, A. T., and Veesler, D., "Structure, function, and antigenicity of the sars-cov-2 spike glycoprotein," *Cell*, vol. 181, no. 2, pp. 281–292, 2020.

[45] Zou, X., Chen, K., Zou, J., Han, P., Hao, J., and Han, Z., "Single-cell rna-seq data analysis on the receptor ace2 expression reveals the potential risk of different human organs vulnerable to 2019-ncov infection," *Frontiers of medicine*, vol. 14, no. 2, pp. 185–192, 2020.

[46] Liang, X., Teng, A., Braun, D. M., *et al.*, "Transcriptionally active polymerase chain reaction (tap): High throughput gene expression using genome sequence data," *Journal of Biological Chemistry*, vol. 277, no. 5, pp. 3593–3598, 2002.

[47] Raybould, M. I. J., Kovaltsuk, A., Marks, C., and Deane, C. M., "CoV-AbDab: the Coronavirus Antibody Database," *Bioinformatics*, vol. 37, no. 5, pp. 734–735, 2021.

[48] Yuan, M., Liu, H., Wu, N. C., *et al.*, "Structural basis of a shared antibody response to sars-cov-2," *Science*, vol. 369, no. 6507, pp. 1119–1123, 2020.

[49] Pinto, D., Park, Y.-J., Beltramello, M., *et al.*, "Cross-neutralization of sars-cov-2 by a human monoclonal sars-cov antibody," *Nature*, vol. 583, no. 7815, pp. 290–295, 2020.

[50] Zost, S. J., Gilchuk, P., Chen, R. E., *et al.*, "Rapid isolation and profiling of a diverse panel of human monoclonal antibodies targeting the sars-cov-2 spike protein," *Nature medicine*, vol. 26, no. 9, pp. 1422–1427, 2020.

[51] Rogers, T. F., Zhao, F., Huang, D., *et al.*, "Isolation of potent sars-cov-2 neutralizing antibodies and protection from disease in a small animal model," *Science*, vol. 369, no. 6506, pp. 956–963, 2020.

[52] Cock, P. J., Antao, T., Chang, J. T., *et al.*, "Biopython: Freely available python tools for computational molecular biology and bioinformatics," *Bioinformatics*, vol. 25, no. 11, pp. 1422–1423, 2009.

[53] Ewels, P., Magnusson, M., Lundin, S., and Käller, M., "Multiqc: Summarize analysis results for multiple tools and samples in a single report," *Bioinformatics*, vol. 32, no. 19, pp. 3047–3048, 2016.

[54] Bolger, A. M., Lohse, M., and Usadel, B., "Trimmomatic: A flexible trimmer for illumina sequence data," *Bioinformatics*, vol. 30, no. 15, pp. 2114–2120, 2014.

[55] Bolotin, D. A., Poslavsky, S., Mitrophanov, I., *et al.*, "Mixcr: Software for comprehensive adaptive immunity profiling," *Nature methods*, vol. 12, no. 5, pp. 380–381, 2015.

[56] Desautels, T. A., Arrildt, K. T., Zemla, A. T., *et al.*, "Computationally restoring the potency of a clinical antibody against sars-cov-2 omicron subvariants," *bioRxiv*, pp. 2022–10, 2022.

[57] Kaneko, N., Kuo, H.-H., Boucau, J., *et al.*, "Loss of bcl-6-expressing t follicular helper cells and germinal centers in covid-19," *Cell*, vol. 183, no. 1, pp. 143–157, 2020.

[58] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T., "Analyzing and improving the image quality of stylegan," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8110–8119.

[59] *Rxrx19 dataset*, https://www.rxrx.ai/rxrx19, Accessed: 2021-04-28.

[60] *Rxrx1 dataset*, https://www.rxrx.ai/rxrx1, Accessed: 2021-04-28.

[61] Radford, A., Metz, L., and Chintala, S., "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.

[62] Lin, D., Fu, K., Wang, Y., Xu, G., and Sun, X., "Marta gans: Unsupervised representation learning for remote sensing image classification," *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 11, pp. 2092–2096, 2017.

[63] Zhang, M., Gong, M., Mao, Y., Li, J., and Wu, Y., "Unsupervised feature extraction in hyperspectral images based on wasserstein generative adversarial network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 5, pp. 2669–2688, 2018.

[64] Mao, X., Su, Z., Tan, P. S., Chow, J. K., and Wang, Y.-H., "Is discriminator a good feature extractor?" *arXiv preprint arXiv:1912.00789*, 2019.

[65] Taylor, J., Earnshaw, B., Mabey, B., Victors, M., and Yosinski, J., "Rxrx1: An image set for cellular morphological variation across many experimental batches," in *International Conference on Learning Representations (ICLR)*, 2019.

[66] Cuccarese, M. F., Earnshaw, B. A., Heiser, K., *et al.*, "Functional immune mapping with deep-learning enabled phenomics applied to immunomodulatory and covid-19 drug discovery," *bioRxiv*, 2020.

[67] Zhuang, D. and Ibrahim, A. K., "Deep learning for drug discovery: A study of identifying high efficacy drug compounds using a cascade transfer learning approach," *Applied Sciences*, vol. 11, no. 17, p. 7772, 2021.

[68] Saberian, M. S., Moriarty, K. P., Olmstead, A. D., *et al.*, "Deemd: Drug efficacy estimation against sars-cov-2 based on cell morphology with deep multiple instance learning," *IEEE Transactions on Medical Imaging*, 2022.

[69] Heiser, K., McLean, P. F., Davis, C. T., *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.

[70] Mescheder, L., Geiger, A., and Nowozin, S., "Which training methods for gans do actually converge?" In *International conference on machine learning*, PMLR, 2018, pp. 3481–3490.

[71] Deng, J., Guo, J., Xue, N., and Zafeiriou, S., "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4690–4699.

[72] Pihlgren, G. G., Sandin, F., and Liwicki, M., "Improving image autoencoder embeddings with perceptual loss," in *2020 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2020, pp. 1–7.

[73] Karras, T., Aila, T., Laine, S., and Lehtinen, J., "Progressive growing of gans for improved quality, stability, and variation," *arXiv*, 2018. eprint: 1710.10196 (cs.NE).

[74] Petzka, H., Fischer, A., and Lukovnicov, D., "On the regularization of wasserstein gans," *arXiv*, 2018. eprint: 1709.08894 (stat.ML).

[75] Ko, M., Jeon, S., Ryu, W.-S., and Kim, S., "Comparative analysis of antiviral efficacy of fda-approved drugs against sars-cov-2 in human lung cells," *Journal of medical virology*, vol. 93, no. 3, pp. 1403–1408, 2021.

[76] Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., and Abbeel, P., "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," *Advances in neural information processing systems*, vol. 29, 2016.

[77] Vacca, F., Sala, C., and Rappuoli, R., "Monoclonal antibodies for bacterial pathogens: Mechanisms of action and engineering approaches for enhanced effector functions," *Biomedicines*, vol. 10, no. 9, p. 2126, 2022.

[78] Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q., "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.

[79] Baker, S. J., Payne, D. J., Rappuoli, R., and De Gregorio, E., "Technologies to address antimicrobial resistance," *Proceedings of the National Academy of Sciences*, vol. 115, no. 51, pp. 12 887–12 895, 2018.

[80] Bank, W., *Drug-resistant infections: a threat to our economic future*. World Bank, 2017.

[81] Bloom, D. E., Black, S., Salisbury, D., and Rappuoli, R., "Antimicrobial resistance and the role of vaccines," *Proceedings of the National Academy of Sciences*, vol. 115, no. 51, pp. 12 868–12 871, 2018.

[82] Gordon, S., "Phagocytosis: An immunobiologic process," *Immunity*, vol. 44, no. 3, pp. 463–475, 2016.

[83] Sokolovska, A., Becker, C. E., and Stuart, L. M., "Measurement of phagocytosis, phagosome acidification, and intracellular killing of staphylococcus aureus," *Current protocols in immunology*, vol. 99, no. 1, pp. 14–30, 2012.

[84] Organization, W. H. *et al.*, *Global incidence and prevalence of selected curable sexually transmitted infections-2008*. World Health Organization, 2012.

[85] Micoli, F., Bagnoli, F., Rappuoli, R., and Serruto, D., "The role of vaccines in combatting antimicrobial resistance," *Nature Reviews Microbiology*, vol. 19, no. 5, pp. 287–302, 2021.

[86] Aitolo, G. L., Adeyemi, O. S., Afolabi, B. L., and Owolabi, A. O., "Neisseria gonorrhoeae antimicrobial resistance: Past to present to future," *Current Microbiology*, vol. 78, pp. 867–878, 2021.

[87] Unemo, M., Rio, C. D., and Shafer, W. M., "Antimicrobial resistance expressed by neisseria gonorrhoeae: A major global public health problem in the 21st century," *Emerging Infections 10*, pp. 213–237, 2016.

[88] Diago-Navarro, E., Calatayud-Baselga, I., Sun, D., *et al.*, "Antibody-based immunotherapy to treat and prevent infection with hypervirulent klebsiella pneumoniae," *Clinical and Vaccine Immunology*, vol. 24, no. 1, e00456–16, 2017.

[89] Château, A. and Seifert, H. S., "Neisseria gonorrhoeae survives within and modulates apoptosis and inflammatory cytokine production of human macrophages," *Cellular microbiology*, vol. 18, no. 4, pp. 546–560, 2016.

[90] Paszke, A., Gross, S., Massa, F., Lerer, A., and Bradbury, J. P., "An imperative style, high-performance deep learning library," *Advances in Neural Information Processing Systems*, vol. 32,

[91] Iversen, P. W., Eastwood, B. J., Sittampalam, G. S., and Cox, K. L., "A comparison of assay performance measures in screening assays: Signal window, z'factor, and assay variability ratio," *Journal of biomolecular screening*, vol. 11, no. 3, pp. 247–252, 2006.