UNIVERSITY OF CAMBRIDGE

Accurate estimation of drug combination synergies with uncertainty quantification

Haoting Zhang

Downing College

September 2024

This dissertation is submitted for the degree of Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Haoting Zhang
September 2024

Abstract

Accurate estimation of drug combination synergies with uncertainty quantification

Haoting Zhang

In pharmacology, there exists a range of different quantification frameworks to estimate the synergistic effect of drug combinations. However, the varying assumptions and the absence of a universal gold standard across these frameworks make it difficult to quantitatively prioritise which drug combinations from large screening experiments should advance for further investigation. Furthermore, existing frameworks lack accurate uncertainty quantification and often do not preserve and decouple the potency and efficacy parameters of the combinations, making it difficult to understand and trust the combination candidates selected from such models.

In this work, we first propose SynBa, a flexible Bayesian approach to estimate the uncertainty of the synergistic efficacy and potency of drug combinations, so that actionable decisions can be derived from the model outputs. The actionability is enabled by incorporating the Hill equation into SynBa, so that the parameters representing the potency and the efficacy can be preserved. Existing knowledge may be conveniently inserted due to the flexibility of the prior, as shown by the empirical Beta prior defined for the normalised maximal inhibition. Through experiments on large combination screenings and comparison against benchmark methods, we show that SynBa provides improved accuracy of dose-response predictions and better-calibrated uncertainty estimation for the parameters and the predictions.

We then extend SynBa to SynBa-Batch to explain away the batch effects existing in the noisy measurements so that the noise can be better understood. We provide a route for SynBa and SynBa-Batch to utilise existing monotherapy information and design more informed priors. This approach enhances the reliability and robustness of inferring the synergistic efficacy and potency of drug combinations.

Furthermore, we incorporate SynBa likelihood in the prediction for drug combination responses and develop DeepSynBa, a deep learning model that predicts the complete dose-response matrix of drug pairs instead of relying on an aggregated synergy score. DeepSynBa

separates out efficacy and potency through the incorporation of SynBa likelihood. This enables more informed decision-making, as it offers more detailed view of drug interactions beyond a single synergy measure.

Through SynBa, SynBa-Batch and DeepSynBa, we advocate the use of quantification frameworks in pharmacology that includes a principled uncertainty estimation for model parameters, and preserves and decouples the synergistic efficacy and potency of the combinations to encourage actionability of the model in further decision-making.

Acknowledgements

I would like to start by thanking my three PhD supervisors, Carl Henrik, Marta and Magnus. I am grateful to have such a strong team. Your insights, expertise and guidance have been essential in not only shaping this thesis project, but more importantly, helping me develop as a researcher. You set up examples by being passionate, curious, truth-seeking and rigorous in research. Your encouragement and emotional support have also been special, and gave me the strength to keep going. You are the role model for me.

No words can express how grateful I am to my family for their unconditional love and support. To my partner, Yi, thank you for your love and support throughout not just these four years, but the past ten years during which I pursue my path in research. You are always my biggest supporter and my first audience. Regardless of the situations (COVID, challenges I met in research, to name a few), you are always by my side. Your warmth and positive energy in life have been truly inspiring, and I have learned so much from you.

To my parents and grandparents, thank you for raising me and teaching me the important principles that guide my life. Your life experience has always been a light when I felt lost. It is thanks to all of you that I have had the opportunity to explore the world and pursue what I enjoy. I know I can never fully repay you, but I hope you know how deeply grateful I am for everything you have done for me.

To my friends, thank you so much for your company. A special thanks to Philip - thank you for your firm support since the undergraduate days. You are like my elder brother, who always provides me with a fresh perspective on things in both work and life. We always have great fun whenever we gather, regardless of the location or venue.

Last but not least, to my ginger cat, Biubiu, thank you for keeping me company during the pandemic and the thesis write-up. Whenever I feel stressed or tired, you calm me down by simply being there and doing nothing.

I would also like to acknowledge the receipt of studentship award from the Health Data Research UK-The Alan Turing Institute Wellcome PhD Programme in Health Data Science (Grant Ref: 218529/Z/19/Z) and the Wellcome Cambridge Trust Scholarship. Their financial support has made it possible for this project to happen. A special thanks to Chris - it has been a unique PhD experience thanks to your dedication to the programme.

Contents

1 Introduction and motivation **13**

1.1 Machine learning for drug discovery 14

1.2 Uncertainty and decision-making for drug discovery and combinations . . . 17

1.3 Structure of this thesis 20

2 Methodology background in machine learning **23**

2.1 The frequentist approach 24

2.2 The Bayesian approach 25

2.3 Why Bayesian for this study? 26

2.4 Approximate inference 27

2.4.1 Variational inference 27

2.4.2 Markov chain Monte Carlo 30

2.5 An example: Clustering and latent variable modelling 33

2.6 Conclusion 37

3 Inference in pharmacology **39**

3.1 The dose-response relationship 39

3.1.1 The Hill equation 42

3.2 Monotherapy quantification frameworks 43

3.3 Combination quantification frameworks 45

3.3.1 The classic frameworks: Bliss, Loewe and HSA 46

3.3.2 The deterministic and parametric models 48

3.3.3 The probabilistic models 49

3.4 Synergy and its myth 50

4 SynBa: A new inference method of drug combination synergies **53**

4.1 Methods 55

4.1.1 Overview of SynBa: Monotherapies 56

4.1.2 Overview of SynBa: Combinations 60

4.1.2.1 Inference of the synergy 62

4.1.3 Case studies 63

4.1.4 Training details . . . 65
4.2 Results . . . 66
4.2.1 Prediction of drug combination responses . . . 66
4.2.2 Uncertainty calibration . . . 67
4.3 Discussions and Conclusions . . . 69

5 SynBa-Batch: Modelling the batch effect and utilising information 73
5.1 Batch effect . . . 73
5.2 Methods . . . 75
5.2.1 SynBa-Batch for monotherapies . . . 76
5.2.1.1 Case study . . . 79
5.2.2 SynBa and SynBa-Batch: Utilising existing models or information . 83
5.2.2.1 Case study . . . 85
5.3 Prediction of drug combination responses . . . 87
5.4 Conclusions . . . 90

6 DeepSynBa: Drug combination prediction with complete dose-response profiles 93
6.1 Previous methods . . . 93
6.2 DeepSynBa: Our new approach . . . 97
6.3 Case studies . . . 98
6.4 Results . . . 101
6.4.1 Performance on Cell Lines and Drugs . . . 103
6.4.1.1 Performance Across Tissues . . . 104
6.4.1.2 Performance Across Drug Categories . . . 105
6.4.1.3 Synergy Score Prediction Performance . . . 105
6.5 Discussions and Next Steps . . . 106

7 Conclusions and future directions 109

References 113

Chapter 1

Introduction and motivation

The field of machine learning (ML) has gained significant progress in the past decades. Perhaps most notably, the advancement of ML model architectures, the increasing availability of massive datasets and the increasingly powerful computational resources have enabled the success of deep learning and generative models. Models such as the series of GPT-based models [63] [11], BERT [19], GANs [26] and diffusion models [32] have demonstrated the ability to generate highly realistic content, sometimes indistinguishable from human-created outputs. These models are powerful in tasks such as natural language processing and image and video generation. Their ability to generalise from a vast amount of data has resulted in plenty of successful applications such as multi-language translation, image editing, and recommendation systems.

While it is important to acknowledge that applications such as automated text and image generation can result in issues such as misinformation, scams and identity theft, some of which could lead to catastrophic outcomes, these domains often operate in environments with lower direct risks to human life. In contrast, applications such as healthcare and autonomous vehicles must operate with higher standards of accuracy and reliability, and adhere to strict regulations, where even small errors can lead to significant harm. In these cases, ML models need to offer more than just plausible predictions. They must also provide robust uncertainty estimation and interpretable insights that can be integrated into the established regulatory processes. An accurate estimation of the uncertainty is crucial for a trusted and safe usage of the model outputs, which is still a shortcoming of the large-scale deep learning methods [49].

In drug discovery in particular, diverse sources of uncertainty exist, ranging from establishing the causal relationship between target and disease, to characterising cellular and drug properties, measuring drug responses, and translating findings from lab environment to human beings, to name a few. This becomes problematic when making decisions that could affect human lives, as unreliable predictions may lead to misguided choices and detrimental impact on human well-being. For example, the lack of uncertainty estimation

for the toxicity of a drug may lead to it being more toxic than estimated and kill healthy cells along with the targeted cancer cells.

In the meantime, drug discovery is not only a problem of identifying novel compounds (or a novel combination of compounds) but also a challenge of allocating limited resources. Given the vast space of potential drug candidates and the constrained number of compounds that can be experimentally tested, the reliability of the predictions from machine learning models becomes crucial. To achieve this reliability, it is essential to incorporate uncertainty estimation, enabling informed decision-making by quantifying the confidence in model predictions.

For these reasons, ML models that deal with small data (in terms of size compared to those in e.g. large language models) and incorporate uncertainty estimation are necessary to approach problems in this domain. They also need to be tailored to the specific problem so that the estimated parameters and their uncertainties can be interpreted or explained.

1.1 Machine learning for drug discovery

The drug discovery and development process is a complex and lengthy journey that takes over 10–15 years with an average cost of over 1–2 billion for each new drug to be approved for clinical use [71]. As illustrated in Fig. 1.1 [71], the drug development process is often illustrated as a funnel where a large number of potential compounds are gradually narrowed down to a single or a few FDA-approved drugs.

At the top of the funnel, the first step is to perform an extensive genetic and genomic target identification and validation, usually taking one to two years in practice. This is an important step to confirm that a specific biological target is involved in the disease and can be effectively modulated by a drug. Traditionally, understanding the association between a biological target and a disease requires knowledge of its biological function, pathways or structural properties, derived from biological experiments. With the increasing availability of rich genomic, transcriptomic and proteomic datasets, ML has emerged as a complementary tool for predicting the association between targets and diseases. By leveraging patterns learned from known drug targets, ML models can predict new potential targets based on similarities in biological properties, network interactions, or multi-omics signatures. For example, Jeon et al. developed a support vector machine classifier to distinguish drug targets from non-drug targets in breast, pancreatic, and ovarian cancers, integrating genomic and systematic datasets [35]. Although the causal role between the target and the disease cannot be directly derived, such ML methods have the advantage of being efficient and scalable to large datasets, accelerating the experimental design for biologists.

After that, compound screening is performed, where researchers start with a vast

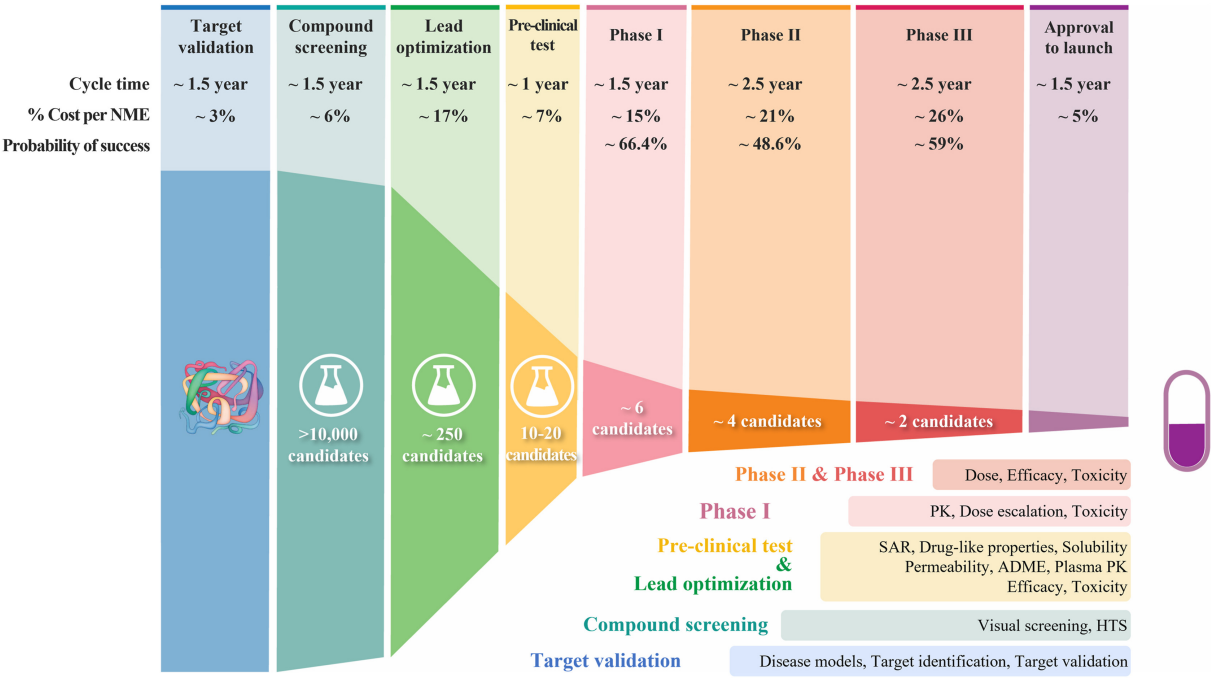

Figure 1.1: The process of drug discovery and development. Figure from [71].

number of compounds, typically around 10,000, and test whether they interact with a target of interest identified in the previous step. These compounds are identified through various methods, including high-throughput screening, computational modelling, and insights from biological pathways. The goal at this stage is to predict which compounds are more likely to achieve desirable therapeutic effects with the target with a relatively low level of toxicity. The process of compound screening takes another one to two years in practice. At this stage, the number of drug candidates is typically reduced to a few hundred. The remaining candidates will proceed to further preclinical testing.

Due to the large number of possible drug candidates and the high-dimensional nature of the biological and chemical data modalities, machine learning methods are increasingly employed to analyse the data and make predictions about which compounds are worth further investigation. The rapid development of deep learning has resulted in increasingly predictive models capable of extracting informative low-dimensional embeddings from the high-dimensional biological and chemical data. Graph neural networks (GNNs) and deep generative models are two classes of ML models that are promising in this scenario, as GNNs can encode the structural and chemical features of molecules, whereas deep generative models can generate novel compounds with the desired properties. For example, DeepGraphMolGen [39] combines GNNs with deep generative models using a reinforcement learning pathway. In this model, GNNs are used to learn crucial properties of compounds, whereas the generative module is used to generate novel compounds that satisfy these properties learnt from the GNNs.

The next stage of the funnel focuses on lead optimisation and pre-clinical testing of

the drugs. In this stage, the pharmacokinetics (how an organism affects the drug) and pharmacodynamics (how the drug affects an organism) of each candidate are evaluated and quantified. In particular, the toxicity and the efficacy of the drugs are carefully examined. At this stage, the complexity of the data increases as more detailed biological and chemical interactions are examined. The large scale, the high dimensions and the diverse modalities of the biomedical data available make ML an appealing tool to learn the drug properties. For example, one important property is how the drug interacts with other drugs in the human system. Predicting drug-drug interactions (DDIs) is a task where ML has become useful. DeepDDI [69] is a deep neural network that predicts the DDIs (in the form of human-readable sentences) from the chemical structures of the compounds, achieving an accuracy of above 90% on the DrugBank gold standard DDI dataset [81].

Unlike the high-throughput phase, where the primary challenge is managing the volume of data, the challenge here lies in dealing with the inherent uncertainties of biological systems. Factors such as variability in the experimental environments, differences in individual responses, and the limitations of translating in-vitro data to animals and humans all contribute to this uncertainty. Understanding this uncertainty is crucial because a successful drug must not exhibit excessive toxicity due to random variation and should consistently reach the expected efficacy when applied to humans. By applying advanced statistical models and machine learning techniques, the aim is to better understand and manage these uncertainties, improving the decision-making process that leads to clinical trials. This will also be the focus of this thesis.

Up to this point, all stages belong to the preclinical phase. The next stage in the funnel enters the clinical phase, where the number of remaining drug candidates is typically very small. From the thousands of drug candidates, a typical drug discovery study may end up with 5 candidates going into the clinical trials. Clinical trials are conducted in multiple phases, starting with a small group of healthy volunteers (Phase 1) to assess safety, followed by larger groups of patients (Phases 2 and 3) to evaluate efficacy and monitor side effects. By the time a drug reaches clinical trials, most of the work follows a predefined and highly regulated framework, with strict protocols and guidelines that govern every aspect of the process. While there is still some room for optimisation, particularly in patient stratification and personalised medicine, the scope for applying novel computational models is more limited compared to earlier stages. In particular, most machine learning models lack the interpretability and explainability required for this stage of the study, and are thus limited to being applied in the preclinical stage which is more exploratory.

In summary, out of the original 5,000 to 10,000 compounds, only one will typically be approved by the FDA and make it to the market [15]. This singular drug represents years of research, development, and testing, with an estimated time frame of 10-15 years from initial discovery to approval. Post-marketing monitoring and research would follow

to ensure the safety and efficacy of the drug in the general population.

The pipeline above describes the development of a new single drug compound, also called *monotherapies*, which we will refer to repeatedly for the remainder of this thesis. Another method of developing new drugs, which has become popular in modern pharmacology, is to combine two drug compounds, due to a number of advantages. First, there is a possibility of an increase in efficacy, especially when the two drugs have different mechanisms of action [18]. Secondly, drug combination allows for decreasing the dosages of each component and thus their toxicity [18]. Thirdly, with the increased use of small molecule drugs in monotherapy, resistance to treatments has become a clear challenge. Drug combinations offer a possible solution to bypassing resistance with alternative targeting [84] [89] [16]. In addition, there is evidence that drug combinations may provide a lower risk of side effects, which directly results in a better quality of life [12] [60] [27]. Combinations are particularly explored in complex diseases such as cancer, where single agents may be insufficient to address the complex nature of the disease. In addition, since the monotherapies are most often FDA approved, the timescale of getting a combination onto the market will be shorter than a brand new drug compound.

The space of potential drug combinations grows exponentially with the number of drugs, creating a vast and complex landscape that is challenging to explore through traditional experimental methods alone. This is where machine learning can offer assistance as a tool. Machine learning models have the potential to analyse large datasets, predict synergistic interactions, and prioritise the most promising combinations for further investigation. Thanks to the development of high-throughput approaches to drug combination screening, the amount of data is now on an unprecedented scale. Examples include the AstraZeneca-Sanger DREAM challenge (DREAM) [54], NCI-ALMANAC [33], DrugComb [85], the screening data from the Wellcome Sanger Institute [34], and the latest AstraZeneca-Sanger data [5]. The availability of these datasets makes it possible to predict the effect of drug combinations from modelling drug-sensitivity data, leveraging information from biological features of the cell lines and chemical characteristics of the drugs [12]. Chapter 6 will include a detailed review of these methods.

1.2 Uncertainty and decision-making for drug discovery and combinations

The drug discovery process described in Section 1.1 is a lengthy and costly process that often takes more than a decade to produce a viable and safe drug from thousands of candidates. Even for the drugs that enter clinical trials, 90% of them will fail to pass the trials. The reasons for this are complex, which include lack of clinical efficacy, unmanageable toxicity and poor drug-like properties [71]. Many of these reasons stem

from the inherent uncertainty that lies along the drug discovery process, from the early stage of target identification to the clinical stage.

In the target identification stage, there are uncertainties in determining the causal relationship between a target and a disease pathway. While a target may appear to be associated with a disease based on genetic, transcriptomic, or proteomic data, this correlation does not necessarily imply causation. Targets may have off-target effects in multiple biological processes, resulting in the indirect association with the disease.

In the stage of lead optimisation and preclinical test, uncertainties arise in the prediction of the efficacy, potency, PK and PD properties of the drugs. A small biological variability or a difference in experimental environments could result in significantly different results for the estimation of these properties. Within the drug properties of interest in this stage, one critical aspect is the toxicity level of the drug, i.e. whether other healthy cells are excessively killed when the drug is applied. If there is a significant uncertainty around toxicity, even if the efficacy appears promising, the potential risks of off-target side effects might outweigh the benefits. It is crucial to prioritise drugs with well-understood and low-toxicity profiles because the consequences of advancing a highly toxic compound are severe, potentially harming patients by damaging healthy cells.

There is also uncertainty in the translation from *in vitro* (i.e. in a controlled environment outside a living organism, typically in a laboratory setting) assays or animal models to the clinical trials on human beings, since the former may not fully capture the complexity of human physiological or disease states, leading to discrepancies in translational outcomes. Even when in vitro studies are rigorously controlled and yield strong results, discrepancies in translational outcomes can still arise.

Due to the presence of multiple sources of uncertainty, it is important to select drug candidates for clinical trials that are not only promising but also supported by robust evidence and low uncertainty. While it is impossible to eliminate uncertainty entirely, carefully modeling and prioritising candidates with lower uncertainties can improve the efficiency and success rate of the drug development pipeline. The goal is to maximise the likelihood of successful drug development and ensure that limited resources are used effectively, ultimately benefiting human well-being.

In this thesis, we focus on the uncertainty that arises in the stage of lead optimisation and preclinical test. In particular, we focus on the uncertainties in predicting how the efficacy and the potency change when two drugs are combined. Over the years, there have been quantification frameworks that measure and interpret the interactions between the drugs, which will be reviewed in detail in Chapter 3. These frameworks generate quantified metrics that provide insights into whether the combined effect of the drugs is greater than, equal to, or less than the sum of their individual effects. Such metrics can be used for subsequent decision-making, guiding researchers and clinicians in determining

whether a particular drug combination is worth pursuing in further preclinical or clinical studies. However, there has been no consensus on which framework is the gold standard. The lack of agreement largely stems from the fact that many existing frameworks rely on assumptions that may not hold true across different biological contexts or therapeutic areas.

Moreover, the existing quantification frameworks for pharmacology do not fully explore the challenge of uncertainty. There are various sources of uncertainty associated with drug-sensitivity modelling. First, the biology of dose-response relationships is still unknown despite existing efforts. This leads to uncertainty associated with insufficient scientific knowledge. Secondly, there are systematic and random errors arising from the experimental procedures. Thirdly, biological variation exists among cell lines of the same disease, resulting in a further source of noise. Finally, uncertainty also stems from the limited information that can be extracted due to the small size of available data (*epistemic uncertainty*).

Given these multiple sources of uncertainty, it is often impossible to reach an accurate estimate of the quantities of interest, e.g potency of a monotherapy, or the synergistic effect (in terms of either potency or efficacy) of a combination. Since it is difficult to reach the ground truth, the next resort would be to quantify the uncertainty of the estimates along with their most probable values. By analysing the estimated uncertainty of the key parameters, researchers can make informed decisions on various aspects of drug development. For example, when evaluating the drug combination candidates, the uncertainty level of the outputs can reveal which combinations are most likely to be synergistic, and whether a high confidence level is accompanied with this prediction. This information can also indicate which combinations require additional experimental validation to reduce uncertainty before advancing. Such data-driven decisions help streamline the development pipeline, ensuring that the finite amount of time and experimental resources are allocated efficiently.

Moreover, examining the uncertainty in the prediction of drug viability across different dosage ranges can aid experimental design. Researchers can strategically select the most informative dosages to explore in subsequent experiments. For example, regions with high uncertainty but potential for efficacy can be targeted to maximise the information gained from additional experiments. By focusing on these areas, researchers can refine their understanding of the dose-response relationship, identifying optimal dosage levels that maximise therapeutic effects while minimising adverse reactions. Compared to traditional trial-and-error methods, conducting experiments sequentially based on the continuously updated uncertainty estimates can improve efficiency, potentially leading to higher information gain with the same number of experiments.

In Chapter 3, we will formally define the key quantities of interest in drug combination

studies, which are the core focus of this thesis. We will review the existing estimation frameworks for these quantities and discuss their strengths and limitations. We will also discuss the role of uncertainty in the estimation of these quantities.

1.3 Structure of this thesis

In this study, we aim to achieve a more accurate estimation of drug combination synergies by developing a principled framework for uncertainty quantification. This framework addresses the uncertainty associated with both the model parameters and the synergy values themselves. By implementing such an approach, we aim to demonstrate the benefits of incorporating structured uncertainty estimation into decision-making processes within drug development.

In Chapter 2, we introduce the necessary machine learning methodologies for this study, with a particular focus on Bayesian inference, the paradigm we use for uncertainty estimation. We will explain why we opt for the Bayesian approach for uncertainty estimation in our problem of interest. The foundation established in this chapter will facilitate a clearer understanding of the methodology presented in Chapters 3 to 6.

Chapter 3 outlines the details of the pharmacological problem we aim to address, i.e. quantifying synergy in drug combinations and the associated uncertainty. We define the parameters of interest and set out the notations. We also discuss whether the existing frameworks for synergy quantification, or the current concept of "synergy" itself, are adequate.

In Chapter 4, we present SynBa, our new Bayesian method to address the problem of quantifying synergy in drug combinations and the associated uncertainty defined in Chapter 3 and improve upon the drawbacks of the existing frameworks. This work is presented as a publication [86]. Its code is available at `https://github.com/HaotingZhang1/SynBa`.

In Chapter 5, we identify the significant impact of the batch effects on dose-response measurements and the lack of adequate approaches to address this critical issue. We also extend SynBa to SynBa-Batch, a new method aimed at addressing the gap in modelling batch effects.

In Chapter 6, we explore how SynBa can be incorporated in predicting the synergy of novel unexplored drug combinations to support the decision-making processes in drug discovery pipelines. As a first step in this direction, we propose DeepSynBa, which incorporates the likelihood function of SynBa and predicts the entire dose-response grid rather than a single synergy score, which has its drawbacks as discussed in Chapters 3 and 4. The preprint of this work is available at bioRxiv [43]. Its code is available at `https://github.com/hikuru/DeepSynBa`. At the end of the chapter, we discuss the next steps for incorporating the full SynBa to the prediction of dose-response grids, including

its Bayesian component.

For the studies in Chapters 4 and 5, I am primarily responsible for the conceptualisation, methodology development, implementation, result analysis and manuscript writing. My co-authors, who are my supervisors, provided guidance, suggestions and feedback through my research and writing.

The study in Chapter 6 is a collaborative effort with Halil Ibrahim Kuru (Bilkent University), A. Ercument Cicek (Bilkent University, Carnegie Mellon University), and Oznur Tastan (Sabanci University). In this collaborative effort, Kuru and myself contribute equally to the work and are labelled as the joint first authors of this work. Myself proposed the initial idea, developed the methodology, assisted with the implementation of the SynBa likelihood incorporation, co-analysed the experimental results, and co-wrote the manuscript. Kuru completed most of the computational implementation, co-analysed the experimental results, and co-wrote the manuscript.

Chapter 2

Methodology background in machine learning

In machine learning, we are interested in finding a model that summarises a dataset $\mathcal{D}$. Two of the most dominant themes in machine learning are supervised learning and unsupervised learning.

In supervised learning, the most common scenario is for each data item $(\mathbf{x}_i, y_i)$ in the dataset $\mathcal{D}$ to contain a potentially multi-dimensional input $\mathbf{x}_i \in \mathbb{R}^d$ (or also called explanatory variables, predictors, or covariates) and a one-dimensional output $y_i \in \mathbb{R}$ (or also called the response, target, or outcome variable). In some cases, the output can also be multi-dimensional. The task in supervised learning is to parameterise the data using a function $f : \mathbb{R}^d \to \mathbb{R}$ as a model. The model can then predict the output $\hat{y}_i = f(\mathbf{x}_i; \boldsymbol{\theta})$ given parameters $\boldsymbol{\theta}$ and some input $\mathbf{x}_i$. The parameter $\boldsymbol{\theta}$ can be one-dimensional or multi-dimensional. In the case of deep learning, for example, $\boldsymbol{\theta}$ can have millions of dimensions.

In unsupervised learning, each data item is not split into input and output. Instead, we directly summarise the data items $\mathbf{x}_i \in \mathcal{D}$ as a whole. For example, one common task is to partition the data into K clusters. In this case, the task is to find a function $f : \mathbb{R}^d \to \{1, 2, ..., K\}$ parameterised by $\boldsymbol{\theta}$ that maps each data item $\mathbf{x}_i$ to one of the clusters $k \in \{1, 2, ..., K\}$. The parameter $\boldsymbol{\theta}$ can be defined in various ways. For example, in clustering, the parameters may include the weight of each cluster, as well as the parameters that define the distribution of each cluster.

Whether it is supervised or unsupervised learning, there are two key aspects to be addressed, namely model specification and estimating the model parameters $\boldsymbol{\theta}$.

Model specification is the process of defining the function space for f, i.e. defining all possible functions that the model can learn. The choice of this function space is crucial because it determines whether the model can capture the underlying relationships in the data. A model with an inappropriate function space for f will yield inaccurate

predictions, regardless of how advanced the computational power and the optimisation techniques are. On the other hand, if the function space is too broad, allowing for excessive flexibility, the model may fall into the issue of identifiability, meaning two or more parameterisations fit the model to the data equally well. For example, if we fit a quadratic function $f(x) = ax^2 + bx + c$ to a dataset with only two points, then there are infinitely many quadratic curves that can pass through these two points and result in zero training error, making the model non-identifiable. Such models may overfit the training data, resulting in poor generalisability to new data from the same underlying population. An effective model specification is essential for machine learning, ensuring that the insights drawn from data have the potential to be accurate and reliable.

After the model is specified, the remaining key question is estimating the parameters $\boldsymbol{\theta}$ given the dataset $\mathcal{D}$. There are two main approaches for estimating $\boldsymbol{\theta}$. One treats it as an unknown fixed value, seeking the best point estimate that explains the data, while the other treats it as a random variable, incorporating uncertainty into the estimation. These approaches form the foundation of the frequentist and Bayesian paradigms in Statistics.

2.1 The frequentist approach

The frequentist approach to parameter estimation focuses on a likelihood function $L(X; \boldsymbol{\theta})$, which represents the probability of observing the data X given the parameters $\boldsymbol{\theta}$. In this framework, the parameter $\boldsymbol{\theta}$ is treated as a fixed, but unknown quantity, and the data is considered random. The likelihood function is defined as:

$$L(X; \boldsymbol{\theta}) = p(X \mid \boldsymbol{\theta}). \tag{2.1}$$

To estimate $\boldsymbol{\theta}$, the frequentist approach maximises the likelihood function. The parameter estimate $\hat{\boldsymbol{\theta}}$ is found by solving the following optimisation problem:

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} L(X; \boldsymbol{\theta}) = \arg\max_{\boldsymbol{\theta}} p(X \mid \boldsymbol{\theta}). \tag{2.2}$$

This method is known as Maximum Likelihood Estimation (MLE). The frequentist interpretation is based on the idea that, with a large amount of data, the MLE will converge to the true parameter value.

In this regime, the standard error (SE) is commonly used as a measure of the uncertainty associated with a sample estimate. The SE quantifies the variability of the sample mean by calculating the standard deviation of the sampling distribution of the mean. Mathematically, the SE is defined as

$$SE = \frac{s}{\sqrt{n}}, \tag{2.3}$$

where s is the sample standard deviation and n is the sample size. The standard error provides an indication of how much the sample mean is likely to deviate from the true population mean.

However, the use of standard error as a measure of uncertainty has several limitations, particularly when dealing with small sample sizes. One significant drawback is that SE relies on the assumption that the sample standard deviation s is an accurate estimate of the population standard deviation. When the sample size n is small, this assumption may not hold. Consequently, the standard error may provide an unreliable measure of uncertainty, giving a false sense of confidence in the sample estimate.

Bootstrap is another common uncertainty estimation method for the frequentist approach. It estimates the sampling distribution of a statistic by resampling with replacement from the original data (nonparametric bootstrap) or by sampling from a model estimated from the data (parametric bootstrap). The estimation is repeated for a specified number of times, and the variability of these resampled statistics provides an empirical distribution that approximates the sampling distribution, enabling more robust estimates of uncertainty.

One advantage of the bootstrap over traditional SE calculations is its flexibility and minimal reliance on assumptions. Since bootstrap does not require the assumption of any specific parametric form for the underlying data distribution, it can be applied to a wide range of problems, including those with non-normal or unknown distributions. However, its effectiveness heavily depends on the size and typicality of the original sample. When the sample size is small, the bootstrap may fail to capture the true variability of the population, leading to biased or excessively optimistic estimates of uncertainty. This is because small samples may not adequately represent the population's underlying distribution, causing the resampling process to replicate the same data points multiple times and fail to account for unobserved variability.

Therefore, although several standard frameworks exist for estimating parameter uncertainty in the frequentist approach, they rely on assumptions that can make the estimated uncertainty unreliable in certain cases. Additional diagnostics are often needed to assess the reliability of the uncertainty estimation, complicating the overall process.

2.2 The Bayesian approach

A different approach to the parameter estimation problem is the Bayesian approach, which treats the parameter $\boldsymbol{\theta}$ as a random variable with its own probability distribution, rather than a fixed quantity as in the frequentist paradigm. Prior beliefs about the parameters $\boldsymbol{\theta}$ are incorporated through a *prior distribution*, and then updated based on the observed

data X to obtain the *posterior distribution*. The update is performed using Bayes' theorem

$$p(\boldsymbol{\theta} \mid X) = \frac{L(X;\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(X)} = \frac{L(X;\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_{\boldsymbol{\theta}'} L(X;\boldsymbol{\theta}')p(\boldsymbol{\theta}')d\boldsymbol{\theta}'}. \tag{2.4}$$

where $p(\boldsymbol{\theta} \mid X)$ is the posterior distribution, $L(X;\boldsymbol{\theta})$ is the likelihood of the observed data given the parameters, and $p(\boldsymbol{\theta})$ is the prior distribution reflecting our prior beliefs about the parameters. The posterior distribution provides a complete summary of our knowledge about the parameters $\boldsymbol{\theta}$ after observing the data X.

The Bayesian approach naturally incorporates uncertainty in parameter estimates through the posterior distribution. Rather than providing a single point estimate, the Bayesian method yields a distribution over possible values of $\boldsymbol{\theta}$, allowing us to quantify the uncertainty associated with the estimates.

The key difference between the frequentist and Bayesian approaches lies in how they treat the parameter $\boldsymbol{\theta}$. Frequentist methods consider $\boldsymbol{\theta}$ as fixed and unknown, focusing solely on the likelihood derived from the data. In contrast, Bayesian methods treat $\boldsymbol{\theta}$ as a random variable, integrating prior information with observed data to form a comprehensive posterior distribution.

The strength of the Bayesian approach is that it offers a direct route to incorporating prior knowledge and quantifying uncertainty. This can be particularly advantageous in fields such as pharmacology, where prior knowledge about biological processes is often available, and understanding the uncertainty associated with model predictions is crucial for decision-making.

2.3 Why Bayesian for this study?

Both SE and bootstrap have been used for uncertainty estimation in the quantification of pharmacological metrics in monotherapies and drug combinations. However, in such applications, the drug viability of interest often has only a small number of measurements, sometimes smaller than 10. For example, in the DREAM Challenge [54], each set of monotherapy experiments has a size of six. This sample size is too small for the SE to adequately capture the variability of the samples. Similarly for bootstrap, the small sample size makes it unlikely to represent the true underlying distribution of the quantify of interest.

For this reason, we opt for the Bayesian approach instead in this study. By using a Bayesian approach, the posterior uncertainty will be continuously updated when new data arrives, and will remain conservative (if a conservative prior distribution is set) when the data size is limited or too noisy, hence naturally avoiding overfitting or being overconfident in the uncertainty.

The second motivation for approaching this problem with the Bayesian framework is the utilisation of priors. With a Bayesian approach, existing knowledge from biology and historical data can be naturally inserted into the model through the prior distributions for the parameters. We aim to design a Bayesian model where the biological knowledge is preserved in the corresponding parameters so that the model can retain interpretability and improve efficiency.

While Bayesian inference offers significant advantages in terms of incorporating prior knowledge and providing a coherent framework for uncertainty quantification, it is not without its challenges [22]. One of the primary difficulties is that Bayesian methods often involve integrals over complex, high-dimensional spaces that do not have closed-form solutions. The integral $\int_{\boldsymbol{\theta}'} L(X;\boldsymbol{\theta}')p(\boldsymbol{\theta}')$ required to compute the marginal likelihood for the data is often intractable, especially for the complex models used for scientific discovery. As a result, the posterior distribution becomes intractable according to Eq. (2.4). Therefore, Bayesian inference typically relies on approximation techniques that will be introduced in Section 2.4.

2.4 Approximate inference

In this section, we introduce two of the most widely used approximate inference schemes for Bayesian inference, Variational Inference (VI) and Markov chain Monte Carlo (MCMC), which are employed to estimate the posterior distribution when an analytical solution is not feasible.

2.4.1 Variational inference

In Bayesian inference, model selection is performed by maximising the model evidence $\log p(X)$. When it is intractable, approximate inference is required. This leads to the motivation of variational inference (VI), which is to estimate the intractable posterior distribution by approximating it with a simpler, tractable class of distribution. We call it the *variational distribution* and denote it as $q(\boldsymbol{\theta})$. The problem now turns into finding a $q(\boldsymbol{\theta})$ that is "good enough" as an approximation for the true posterior $p(\boldsymbol{\theta} \mid X)$. In VI, We do this by finding a tractable lower bound for $\log p(X)$ and finding the $q(\boldsymbol{\theta})$ that maximises this lower bound.

We write $\log p(X)$ as:

$$\begin{aligned}
\log p(X) &= \int q(\boldsymbol{\theta}) \log p(X) d\boldsymbol{\theta} \\
&= \int q(\boldsymbol{\theta}) \log \frac{p(X, \boldsymbol{\theta})}{p(\boldsymbol{\theta} \mid X)} d\boldsymbol{\theta} \\
&= \int q(\boldsymbol{\theta}) \log \frac{p(X, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} \frac{q(\boldsymbol{\theta})}{p(\boldsymbol{\theta} \mid X)} d\boldsymbol{\theta} \\
&= \int q(\boldsymbol{\theta}) \log \frac{p(X, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta} + \int q(\boldsymbol{\theta}) \log \frac{q(\boldsymbol{\theta})}{p(\boldsymbol{\theta} \mid X)} d\boldsymbol{\theta} \\
&\geq \int q(\boldsymbol{\theta}) \log \frac{p(X, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta}
\end{aligned} \tag{2.5}$$

where the second term in the fourth line is the KL divergence between the two distributions $q(\boldsymbol{\theta})$ and $p(\boldsymbol{\theta} \mid X)$. The KL divergence is non-negative for any two distributions, resulting in the inequality sign in the final line, and thus a lower bound for $\log p(X)$.

This lower bound is called the *evidence lower bound* (ELBO):

$$\mathcal{L}(q) := \int q(\boldsymbol{\theta}) \log \frac{p(X, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta} = \mathbb{E}_{q(\boldsymbol{\theta})} \left[\log p(X, \boldsymbol{\theta})\right] - \mathbb{E}_{q(\boldsymbol{\theta})} \left[\log q(\boldsymbol{\theta})\right]. \tag{2.6}$$

This is the quantity that variational inference maximises to find the best approximation $q(\boldsymbol{\theta})$ to the true posterior.

Now the next question is how do we define the function class for $q(\boldsymbol{\theta})$? Since the objective is to find a tractable approximation for the posterior, it is natural to narrow down to distributions that are more computationally convenient. This leads to the *mean-field approximation*, a common variational approach where $q(\boldsymbol{\theta})$ is assumed to factorise over the components of $\boldsymbol{\theta}$:

$$q(\boldsymbol{\theta}) = \prod_{i=1}^{d} q(\theta_i), \tag{2.7}$$

where each θ_i is independent. With this factorisation, the ELBO can now be written as

$$\begin{aligned}
\mathcal{L}(q) &= \mathbb{E}_{q(\boldsymbol{\theta})} \left[\log p(X, \boldsymbol{\theta})\right] - \sum_{i=1}^{n} \mathbb{E}_{q(\theta_i)} \left[\log q(\theta_i)\right] \\
&= \mathbb{E}_{q(\theta_j)} \left[\mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})} \left[\log p(X, \boldsymbol{\theta})\right]\right] - \sum_{i=1}^{n} \mathbb{E}_{q(\theta_i)} \left[\log q(\theta_i)\right]
\end{aligned} \tag{2.8}$$

for each $j \in \{1, ..., D\}$.

One way to approach this optimisation is to iteratively optimise $q_j := q(\theta_j)$ for each j.

Reorganising $\mathcal{L}(q)$, we have

$$\begin{aligned}\mathcal{L}(q_j) &= \mathbb{E}_{q(\theta_j)}\left[\mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right]\right] - \mathbb{E}_{q(\theta_j)}\left[\log q(\theta_j)\right] + \text{constant} \\ &= \mathbb{E}_{q(\theta_j)}\left[\mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right] - \log q(\theta_j)\right] + \text{constant}\end{aligned} \tag{2.9}$$

where some terms are treated as constant because the other distributions q_i are independent to q_j for $i \neq j$. Differentiating this w.r.t. q_j would yield

$$\frac{\partial}{\partial q_j}\mathcal{L}(q_j) = \mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right] - \log q(\theta_j) + \text{constant} \tag{2.10}$$

which equals zero when

$$\log q(\theta_j) = \mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right] + \text{constant}. \tag{2.11}$$

Finally, we have

$$q_j = \text{constant} \cdot \exp(\mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right]). \tag{2.12}$$

If the expectation

$$\mathbb{E}_{q(\boldsymbol{\theta}_{\neg j})}\left[\log p(X, \boldsymbol{\theta})\right] \tag{2.13}$$

is tractable for each j, then we reach the solution for q_j according to Equation (2.12). This can be done by choosing the function class for q_j in clever ways, according to the need for the specific problem of interest.

While VI is "always wrong" in the sense that we cannot write an intractable distribution with a tractable form, it is usually "good enough" for many practical applications. This trade-off is often acceptable, especially when exact inference is computationally intractable and when other approximate inference methods such as MCMC are too computationally costly, as VI can provide a solution with a significantly smaller computational cost.

However, the mean-field approximation often underestimates the true uncertainty by ignoring correlations between parameters in the mean-field formulation, leading to excessively confident estimates [9]. In addition, in cases where the intractable true posterior is multi-modal, it will be challenging for the posterior to be approximated by the variational distribution, which is usually constrained to be unimodal so that the expression (2.13) is tractable. This will be the case to be discussed in Section 4, and also more generally in scientific discovery, where the parameters of interest may not only have one mode.

2.4.2 Markov chain Monte Carlo

As discussed, the main challenge in inferring the posterior distribution is to evaluate the integral

$$p(X) = \int_{\boldsymbol{\theta}'} L(X; \boldsymbol{\theta}') p(\boldsymbol{\theta}') d\boldsymbol{\theta}'.$$

More generally speaking, it has been a common challenge to evaluate the expected value of a function $f(\mathbf{x})$ w.r.t. some probability distribution $p(\mathbf{x})$:

$$\mathbb{E}_{p(\mathbf{x})}[f(\mathbf{x})] = \int f(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}, \tag{2.14}$$

especially when $f(\mathbf{x})$, $p(\mathbf{x})$ or both have a complex form, or when the dimensionality of $\mathbf{x}$ is high. There is a need for a more general scheme to estimate such quantities. This leads to the development of *sampling* methods. The key idea for sampling is that instead of solving the integral analytically, we generate samples from the distribution $p(\mathbf{x})$ and use these samples to approximate the expectation. Monte Carlo (MC) sampling is the most basic and widely used method for approximating such integrals through random sampling. In MC sampling, we generating N independent and identically distributed (i.i.d.) samples $\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N$ from $p(\mathbf{x})$, and then approximating the expectation by the sample mean:

$$\mathbb{E}_{p(\mathbf{x})}[f(\mathbf{x})] \approx \frac{1}{N} \sum_{i=1}^{N} f(\mathbf{x}_i) \tag{2.15}$$

As N increases, the approximation becomes more accurate. By the law of large numbers (LLN), the Monte Carlo estimate converges to the true value of the integral as N goes to infinity.

To compute a sampling estimate, we would need to be able to sample from $p(\mathbf{x})$. In many cases, direct sampling from $p(\mathbf{x})$ is either infeasible or inefficient, especially when $p(\mathbf{x})$ is complex. To address this, advanced sampling techniques have been developed, most of which rely on the idea of drawing samples from a simpler, more tractable distribution $q(x)$, known as the *proposal distribution*. This approach is conceptually similar to VI in that it simplifies the problem by involving a distribution with a simpler form. However, unlike VI, which approximates the posterior and may not converge to the true distribution, MC methods aim for asymptotic exactness, ensuring that the estimated posterior converges to the true posterior as the number of samples increases to infinity.

Among these techniques, one of the simplest is *rejection sampling* [25]. In rejection sampling, we sample from a simpler proposal distribution $q(\mathbf{x})$, which is easy to sample from, and we accept or reject each sample based on a criterion related to the ratio of the target density $p(\mathbf{x})$ to the proposal density $q(\mathbf{x})$. Specifically, we draw a candidate sample

$\mathbf{x}^*$ from $q(\mathbf{x})$, and accept it with probability

$$\alpha(\mathbf{x}^*) = \frac{p(\mathbf{x}^*)}{Mq(\mathbf{x}^*)} \tag{2.16}$$

where $M \geq \max\left(\frac{p(\mathbf{x})}{q(\mathbf{x})}\right)$ is a constant ensuring that $\alpha(\mathbf{x}^*) \leq 1$. This algorithm becomes inefficient when $q(\mathbf{x})$ does not closely resemble $p(\mathbf{x})$, especially in high-dimensional spaces.

Another approach is *importance sampling* [38] [40]. In this method, we sample from the proposal distribution $q(\mathbf{x})$ and reweight the samples using importance weights to correct for the difference between $q(\mathbf{x})$ and $p(\mathbf{x})$. The importance weight for a sample $\mathbf{x}_i$ is given by

$$w(\mathbf{x}_i) = \frac{p(\mathbf{x}_i)}{q(\mathbf{x}_i)}. \tag{2.17}$$

The expectation of a function $f(\mathbf{x})$ under $p(\mathbf{x})$ can then be approximated as

$$\mathbb{E}_{p(\mathbf{x})}[f(\mathbf{x})] \approx \frac{1}{N}\sum_{i=1}^{N} w(\mathbf{x}_i) f(\mathbf{x}_i). \tag{2.18}$$

One advantage of this method is that there is no need to normalise p. When p and q are both unnormalised, denoted as $\tilde{p}$ and $\tilde{q}$, we will have

$$\mathbb{E}_{p(\mathbf{x})}[f(\mathbf{x})] \approx \frac{\sum_{i=1}^{N} \tilde{w}(\mathbf{x}_i) f(\mathbf{x}_i)}{\sum_{i=1}^{N} \tilde{w}(\mathbf{x}_i)} \tag{2.19}$$

where

$$\tilde{w}(\mathbf{x}_i) = \frac{\tilde{p}(\mathbf{x}_i)}{\tilde{q}(\mathbf{x}_i)}. \tag{2.20}$$

However, importance sampling can suffer from high variance if $q(\mathbf{x})$ and $p(\mathbf{x})$ do not overlap well, leading to large weights for some samples and very small weights for others.

To address these inefficiencies, *Markov chain Monte Carlo* (MCMC) methods were developed. MCMC methods construct a Markov chain whose stationary distribution is the desired posterior distribution. In every iteration, the next sample is generated from a distribution that depends only on the current sample, making it a Markov chain. Over time, the samples generated by the chain converge to the posterior distribution, allowing us to estimate expectations and other properties of the posterior.

Conceptually, MCMC is more effective than methods such as rejection sampling and importance sampling due to its ability to adaptively build the sample set, using previously accepted samples to guide future proposals. Over time, the Markov chain spends more time in regions where the target distribution $p(\mathbf{x})$ has higher density. Even if we start from a poor initial proposal, the chain will "adjust" and gradually concentrate on areas with higher probability. This ability to adaptively focus on high-probability regions gives

MCMC an advantage over rejection or importance sampling, which depend on the proposal distribution $q(\mathbf{x})$ being good from the start.

Since the use of MCMC in this thesis is to estimate the posterior $p(\boldsymbol{\theta}' \mid \mathbf{X}, \mathbf{y})$, from this point onward, we use the notation $p(\boldsymbol{\theta}' \mid \mathbf{X}, \mathbf{y})$ in place of the more general notation $p(\mathbf{x})$ used in rejection and importance sampling, but note that the following equations hold for any general distribution, not just for $p(\boldsymbol{\theta}' \mid \mathbf{X}, \mathbf{y})$.

The Metropolis-Hastings algorithm [56] is one of the simplest and most widely used MCMC methods. It generates a sequence of samples $\{\boldsymbol{\theta}^{(t)}\}$ by proposing a new sample $\boldsymbol{\theta}'$ from a proposal distribution $q(\boldsymbol{\theta}' \mid \boldsymbol{\theta}^{(t)})$ and accepting it with a probability:

$$\alpha = \min\left(1, \frac{p(\boldsymbol{\theta}' \mid \mathbf{X}, \mathbf{y}) q(\boldsymbol{\theta}^{(t)} \mid \boldsymbol{\theta}')}{p(\boldsymbol{\theta}^{(t)} \mid \mathbf{X}, \mathbf{y}) q(\boldsymbol{\theta}' \mid \boldsymbol{\theta}^{(t)})}\right). \tag{2.21}$$

If the proposed sample is accepted, $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}'$; otherwise, $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)}$. The Metropolis-Hastings algorithm is general and flexible, but it can suffer from slow convergence and high autocorrelation, especially in high-dimensional parameter spaces.

To address these issues, Hamiltonian Monte Carlo (HMC) [21] introduces auxiliary momentum variables $\boldsymbol{p}$ and uses Hamiltonian dynamics to propose new states. The Hamiltonian $H(\boldsymbol{\theta}, \boldsymbol{p})$ is defined as:

$$H(\boldsymbol{\theta}, \boldsymbol{p}) = -\log p(\boldsymbol{\theta} \mid \mathbf{X}, \mathbf{y}) + \frac{1}{2}\boldsymbol{p}^\top \boldsymbol{M}^{-1}\boldsymbol{p}, \tag{2.22}$$

where $\boldsymbol{M}$ is a mass matrix. By using the momentum as a "push", HMC can make larger, more directed moves. This results in a more efficient exploration of the distribution, leading to faster convergence and lower autocorrelation compared to vanilla MCMC methods. HMC is particularly effective for high-dimensional problems with complex posterior distributions.

The class of MCMC methods has the advantage of being asymptotically convergent to the exact posterior distribution. However, in practice, it is difficult to verify whether the exact correct solution has been reached. While running diagnostics such as trace plots and computing the effective sample size can provide insights into the convergence of the chains, they do not guarantee that the sampler has fully explored the posterior distribution or identified all its modes.

MCMC can also be computationally slow, particularly in high-dimensional spaces, and may suffer from issues such as getting stuck in local modes. The efficiency of MCMC also depends on the choice of the initial values and the proposal distribution. Poor initialisation or poorly chosen proposal strategies can lead to slow mixing and convergence, affecting the reliability of the results.

2.5 An example: Clustering and latent variable modelling

As an example, we now go through the Gaussian Mixture Model (GMM), a classic example of latent variable modelling, that requires approximate inference. Understanding GMMs will provide a solid foundation for the models and inference methods discussed in the subsequent sections and chapters.

Clustering is a fundamental technique in unsupervised learning that aims to partition a set of observations into distinct groups or clusters, such that observations within the same cluster are more similar to each other than to those in different clusters. This process is particularly useful for exploring the structure of data, identifying natural groupings, and summarising large datasets. Clustering techniques are widely applied across various domains, including pharmacology. Due to the complexity of the high-dimensional and multi-modality data structure, as well as the existence of structured noise, it is helpful to identify meaningful patterns and subgroupings within the data [8].

A common approach to clustering involves introducing a latent categorical variable z, which indicates the cluster assignment for each observation. Let $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$ denote the observed data, where each $\mathbf{x}_i \in \mathbb{R}^d$ is a d-dimensional input vector. The latent variable $z_i = k$ represents the assignment of the i-th observation to the k-th cluster, where $k \in \{1, 2, ..., K\}$. In this context, clustering can be viewed as a latent variable modelling problem, where the goal is to infer both the number of clusters K and the cluster assignments $\{z_1, ..., z_n\}$ from the data.

Latent variable models, such as mixture models, provide a probabilistic framework for clustering. In a GMM, for example, the observed data $\mathbf{X}$ is assumed to be generated from a mixture of Gaussian distributions. The generative process for a GMM can be described as follows:

1. Draw the cluster assignment $z_i = k$ from a categorical distribution with probabilities $(\pi_1, \pi_2, ..., \pi_K)$.
2. Given $z_i = k$, draw $\mathbf{x}_i$ from a Gaussian distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$.

The joint distribution of the observed data and the latent variables is given by

$$p(\mathbf{x}_i, z_i = k \mid \boldsymbol{\theta}) = \pi_k \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \tag{2.23}$$

where $\boldsymbol{\theta} = \{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$ are the parameters of the mixture model. The mixture model approach allows for the modelling of complex data distributions by assuming that each cluster is represented by a different component of the mixture, making it a flexible tool for capturing the underlying structure of the data.

The introduction of the latent variable z not only facilitates clustering but also enables the incorporation of uncertainty into the model. By treating the cluster assignments as latent variables, the model can account for the possibility that an observation may belong to multiple clusters with different probabilities. This probabilistic interpretation provides a richer understanding of the data, allowing for soft clustering, where each observation can be associated with more than one cluster to varying degrees. This is particularly helpful in applications with a high level of uncertainty in the data acquisition and measurement process.

To estimate the parameters $\boldsymbol{\theta}$, some specialised computational algorithms are required because the problem does not yield a closed-form solution due to the presence of the latent cluster assignments. This requires the use of iterative optimisation techniques, such as the Expectation-Maximization (EM) algorithm [17], which alternates between estimating the latent variables and optimising the model parameters.

The EM algorithm is particularly well-suited for problems where direct optimisation of the likelihood function is difficult due to the involvement of unobserved latent variables, such as the GMM. By iteratively performing an Expectation step (E-step) to compute the expected log-likelihood and a Maximisation step (M-step) to update the parameter estimates, the EM algorithm provides a practical and efficient means of achieving exact inference in many settings. Through this iterative process, the EM algorithm navigates the complex parameter space, refining the estimates until convergence, thereby finding the parameter values that maximise the likelihood of the observed data.

The EM algorithm consists of two main steps that are iteratively applied:

- **E-step (Expectation step):** In this step, the expectation of the complete-data log-likelihood, $Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(t)})$, is computed with respect to the current estimate of the posterior distribution of the latent variables given the observed data and the current parameter estimates $\boldsymbol{\theta}^{(t)}$. The E-step can be formally written as

$$Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(t)}) = \mathbb{E}_{\mathbf{Z}|\mathbf{X},\mathbf{y},\boldsymbol{\theta}^{(t)}} \left[\log p(\mathbf{y}, \mathbf{Z} \mid \mathbf{X}, \boldsymbol{\theta})\right] \tag{2.24}$$

  where $\mathbf{Z}$ denotes the latent variables, which is the latent cluster assignment in the case of GMM.

- **M-step (Maximisation step):** In this step, the parameter estimates $\boldsymbol{\theta}$ are updated by maximising the expected complete-data log-likelihood computed in the E-step. The updated parameters $\boldsymbol{\theta}^{(t+1)}$ are obtained by solving

$$\boldsymbol{\theta}^{(t+1)} = \arg\max_{\boldsymbol{\theta}} Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(t)}). \tag{2.25}$$

To approach GMM with EM, we first define the responsibility γ_{ik}, which represents

the probability that data point i belongs to cluster k:

$$\gamma_{ik} = \frac{\pi_k \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^{K} \pi_j \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} \tag{2.26}$$

where π_k is the prior probability (mixture proportion) of cluster k, $\mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ is the Gaussian probability density function with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}_k$ for cluster k, and K is the total number of clusters.

We then perform the E-step by computing the expected value of the complete data log-likelihood with respect to the posterior distribution of the latent cluster assignments given the observed data and current parameters:

$$Q(\theta \mid \theta^{(t)}) = \sum_{i=1}^{N} \sum_{k=1}^{K} \gamma_{ik} \left[\log \pi_k + \log \mathcal{N}(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right] \tag{2.27}$$

In the M-step, we update the parameters of the GMM using the responsibilities calculated in the E-step. In particular, we perform the following:

1. Update the mixture proportions π_k:

$$\pi_k = \frac{N_k}{N} \tag{2.28}$$

2. Update the means for each mixture μ_k:

$$\boldsymbol{\mu}_k = \frac{1}{N_k} \sum_{i=1}^{N} \gamma_{ik} \mathbf{x}_i \tag{2.29}$$

3. Update the covariances $\boldsymbol{\Sigma}_k$:

$$\boldsymbol{\Sigma}_k = \frac{1}{N_k} \sum_{i=1}^{N} \gamma_{ik} (\mathbf{x}_i - \boldsymbol{\mu}_k)(\mathbf{x}_i - \boldsymbol{\mu}_k)^\top \tag{2.30}$$

   where $N_k = \sum_{i=1}^{N} \gamma_{ik}$ is the effective number of data points assigned to cluster k, and N is the total number of data points.

This iterative process continues until convergence, resulting in the final parameter estimates for the GMM model.

While the EM algorithm is effective for many problems, it can sometimes be challenging to perform the M-step if the maximisation of $Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(t)})$ with respect to $\boldsymbol{\theta}$ is difficult due to the complexity of the model or the high-dimensional nature of the parameter space. To address this, the Expectation/Conditional Maximisation (ECM) algorithm [55] is used as a variant of the EM algorithm.

The ECM algorithm modifies the M-step by replacing it with a sequence of Conditional Maximisation (CM) steps, in which each parameter θ_i is maximised individually, while conditioning on the other parameters remaining fixed. This approach breaks down the complex optimisation problem into a series of simpler optimisation problems, making it more tractable. Formally, the ECM algorithm proceeds as follows:

- **E-step:** As in the EM algorithm, compute the expected complete-data log-likelihood $Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(t)})$.
- **CM-steps:** Perform a sequence of conditional maximisation steps, each updating a subset of the parameters. For example, if $\boldsymbol{\theta} = (\theta_1, \theta_2, ..., \theta_k)$, the ECM algorithm updates each θ_i as follows:

$$\theta_1^{(t+1)} = \arg\max_{\theta_1} Q(\theta_1, \theta_2^{(t)}, ..., \theta_k^{(t)} \mid \boldsymbol{\theta}^{(t)}),$$

$$\theta_2^{(t+1)} = \arg\max_{\theta_2} Q(\theta_1^{(t+1)}, \theta_2, ..., \theta_k^{(t)} \mid \boldsymbol{\theta}^{(t)}),$$

$$\vdots$$

$$\theta_k^{(t+1)} = \arg\max_{\theta_k} Q(\theta_1^{(t+1)}, ..., \theta_{k-1}^{(t+1)}, \theta_k \mid \boldsymbol{\theta}^{(t)}).$$

By sequentially maximising each parameter while conditioning on the current estimates of the other parameters, the ECM algorithm provides a flexible and efficient approach to parameter estimation, especially in high-dimensional and complex models. This iterative process continues until convergence, providing estimates of $\boldsymbol{\theta}$ that maximise the likelihood function. Although the ECM algorithm is not necessary for GMMs, it becomes valuable when dealing with more complex likelihood models, as will be the case in Chapter 5, where we will revisit the ECM algorithm.

Although GMMs are not traditionally considered Bayesian models because they are typically fit using maximum likelihood estimation rather than full Bayesian inference, they do incorporate elements of Bayesian thinking, such as in how they treat cluster assignments. In GMMs, each data point has an associated probability of belonging to each cluster, represented as $p(z = k)$ and $p(z = k \mid x)$, where z is the latent variable representing cluster membership and x is the observed data. This reflects a form of prior and posterior belief over the cluster assignments. This is a Bayesian concept in the sense that we update our beliefs about which cluster a data point belongs to, based on the observed data.

In reality, most statistical models, even those considered Bayesian, exist on a spectrum in terms of how fully they embrace Bayesian principles. A fully Bayesian model requires priors not just on the parameters but also on the hyperparameters that define those priors. However, in practice, hyperparameters are often empirically set, rather than being

given their own prior distributions. This means that most Bayesian models rely on some non-Bayesian decisions, such as choosing hyperparameters based on heuristics or prior knowledge rather than strict probabilistic reasoning. This also leads to a philosophical debate: Will we ever be fully Bayesian? If we set a prior on hyperparameters, then one could also argue that we should set a prior on those priors, and so on, leading to what is often referred to as an "endless rabbit hole".

Therefore, it is more useful to view models as lying along a spectrum of Bayesian approaches, rather than categorising them as strictly Bayesian or non-Bayesian. The purpose of this chapter is to explore the Bayesian mindset, which involves updating beliefs in the face of uncertainty, rather than focusing on whether a model strictly adheres to Bayesian principles. Whether we apply full Bayesian inference or not, adopting a Bayesian framework helps guide how we incorporate uncertainty, make decisions, and interpret results.

2.6 Conclusion

In conclusion, our preference for the Bayesian approach in this study is driven by its ability to incorporate prior knowledge and quantify uncertainty in a principled manner. It provides a robust framework for making decisions based on probabilistic reasoning, which is valuable in dose-response modelling, where understanding the uncertainty associated with parameter estimates and predictions can directly influence experimental design and drug development strategies.

However, the practical application of Bayesian inference is not without its challenges. The computational difficulty of deriving exact posterior distributions in models with complex parameterisation makes it necessary to apply approximate inference techniques. Methods such as VI and MCMC have been developed to address these challenges, providing feasible solutions for approximating posterior distributions in scenarios where analytical solutions are unattainable. They balance computational efficiency with the need for accurate uncertainty quantification, making Bayesian methods more accessible and applicable to real-world problems. Between them, VI is faster and more computationally efficient but provides an approximation that is limited by the form of the variational distribution. MCMC, on the other hand, is asymptotically exact but can be computationally slow, especially for complex or high-dimensional models.

The next chapter will introduce the specific pharmacological problem that we aim to address using the Bayesian pipeline. After that, Chapter 4 will introduce SynBa, the Bayesian method we developed to address the problem laid out in Chapter 3, utilising the approximation methods described in this chapter. We will explain why we have opted for MCMC instead of VI in our finalised models. By applying the Bayesian inference principles

and approximate inference techniques discussed in this chapter, we aim to demonstrate how these methods can address the challenges of modelling drug combination synergies, thus improving both the theoretical framework and practical approaches for pre-clinical drug development.

Chapter 3

Inference in pharmacology

In this chapter, we introduce and define the parameters and quantification metrics in pharmacology that are relevant to our study. We describe the basics of the modelling of the dose-response relationship in both monotherapies and combinations. We also review the existing quantification methods in both monotherapies and combinations, with emphasis on the ones that have been developed recently.

3.1 The dose-response relationship

We start from introducing the inference of *monotherapy*, which involves the use of a single drug to observe its effects on a biological system in an *in vitro* setting (i.e. in a controlled environment outside a living organism). The dosage, denoted as x, represents the amount of drug administered. This is typically measured in units of concentration, such as micromolar (μM). The drug response, denoted as y, is the observed effect of the drug on the biological system. This response is often quantified as a percentage of inhibition or growth relative to a control, providing a measure of how the drug impacts the system compared to the control condition where no drug is applied.

In this study, we focus on the percentage of inhibition of a cancer cell line relative to a zero-drug control. We start by quantifying the endogenous growth and death of the cells in the culture. Consider two plates, Plate 1 with no drug (i.e. the control plate) and Plate 2 administered with a certain concentration of the drug. Pre-experimental procedures are taken to ensure the number of cells alive are the same for both plates prior to administration of the drug. After the drug is applied for a predetermined number of days, the number of cells alive is counted for each of the two plates. Normalisation for the plate with drug with respect to the control plate is then performed with the following procedure: If the number of cells alive in Plate 1 (the control plate) is N_1 and the number

of cells alive in Plate 2 is N_2, then the *response* y for the Plate 2 will be computed as

$$y_2 := \frac{N_2}{N_1} \times 100\%,$$

and we will have $y_1 = 100\%$ by definition. By this definition, and by the assumption that $N_2 \leq N_1$ for any dosage x (i.e. no drug encourages growth), the value of the *response* y will range from 0 to 100 per cent. This assumption is made based on the biological motivation that cultivating non-cancerous cells in vitro is extremely difficult. This is because most cells induce cell suicide outside of their host environment and there is little room left biologically for a drug to encourage the cancer cell lines to grow even more [74].

Once the responses are collected and normalised, we now model the relationship between the dosage x and the response y. This is modelled by the dose-response curve, represented by some function $y = f(x)$. This curve illustrates how the biological response changes with varying levels of drug exposure and is fundamental in assessing the efficacy and potency of the drug.

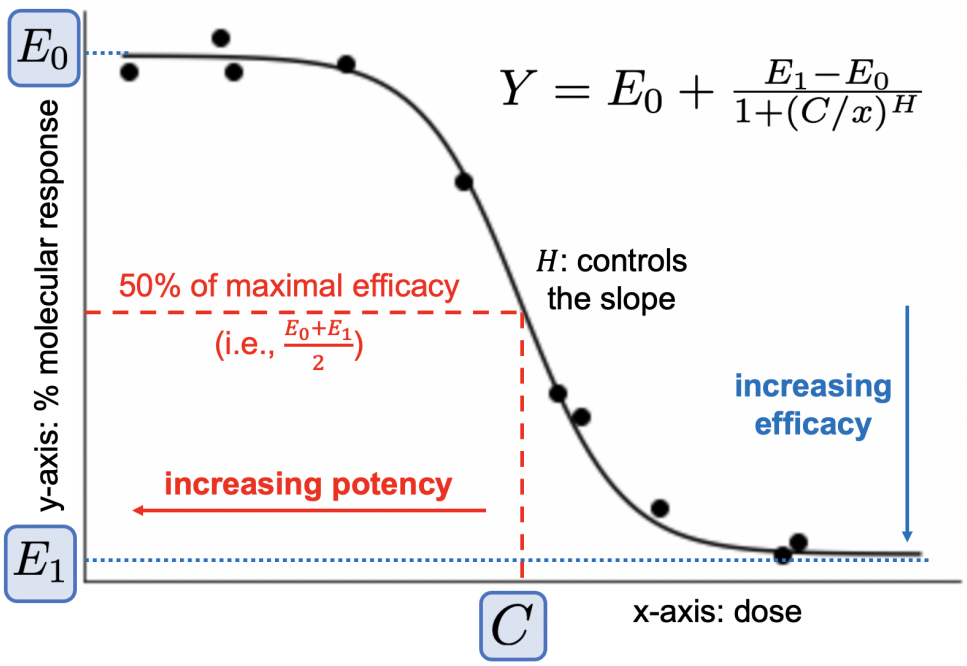

Figure 3.1: A set of monotherapy data for a single drug, with a curve fitted to the data. *Potency* is illustrated in red, while *efficacy* is illustrated in blue. E_0, E_1, C and H are parameters defined in Section 3.1.

Several key parameters are used to characterize the dose-response relationship. E_{inf} represents the maximum effect that can be achieved by the drug, irrespective of the dosage. It is the asymptotic value that the response y approaches as x increases to an infinitely large dosage. In this study, it is defined as the maximum percentage of inhibition relative to the control (zero dosage) observed when a sufficiently high dosage x is applied.

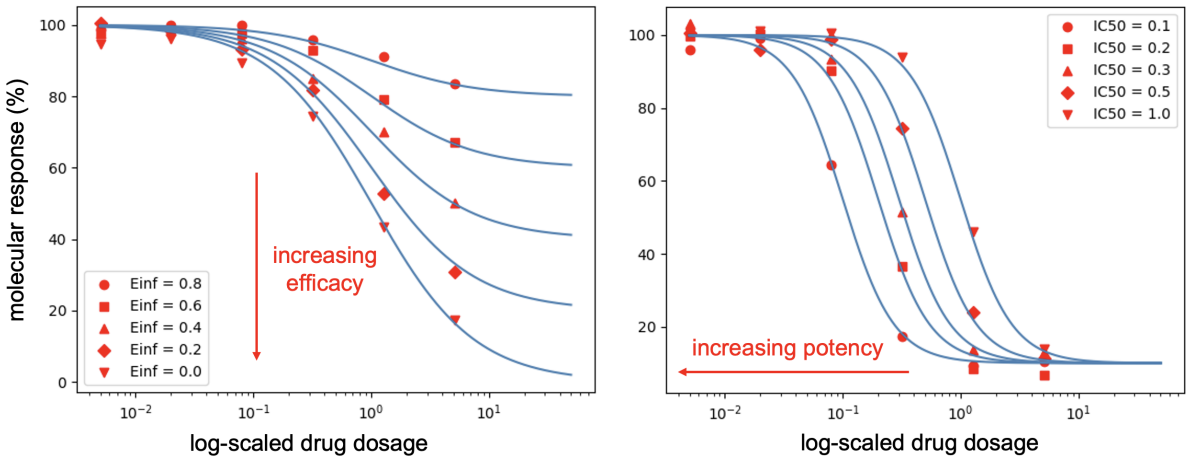

Figure 3.2: (Left) A set of five dose-response curves fitted to five sets of monotherapy data. The five curves have the same underlying potency, but the efficacy increases from the top to the bottom, quantified by a decreasing E_{inf}. (Right) Another set of five dose-response curves fitted to five sets of monotherapy data. The five curves have the same underlying efficacy, but the potency increases from right to left, quantified by a decreasing IC_{50}.

Another crucial parameter is the IC_{50}, or inhibitory concentration, which is defined as the concentration of the drug at which 50% inhibition of the biological activity is observed. This metric serves as an indicator of the *potency* of the drug, providing insight into how much drug is required to achieve half of the maximal inhibitory effect. In the context of cancer studies, the GI_{50}, or growth inhibitory concentration, is used to denote the concentration needed to reduce the growth rate by 50% compared to an untreated control. This parameter is particularly relevant for evaluating the efficacy of anti-proliferative drugs. In this study, we focus on IC_{50} for the consistency of definitions for potency.

Fig. 3.2 illustrates the impact of varying efficacy and potency on the dose-response curves of a drug. The left panel demonstrates how changes in efficacy of the drug are visualised by the dose-response curves. As efficacy increases, which is quantified by a decrease in E_{inf} from 0.8 to 0 (because a more effective inhibition results in a lower c_1), the curves shift downwards, indicating a smaller maximum response level that the drug can achieve.

The right panel focuses on changes in potency, represented by variations in the IC_{50} value, which is the concentration of the drug required to achieve 50% of the maximal effect. As IC_{50} decreases from 1.0 to 0.1, the dose-response curves shift to the left, indicating that lower doses of the drug are needed to achieve the same level of response. This leftward shift signifies an increase in potency, demonstrating that the drug becomes more effective at lower concentrations.

Together, these panels emphasize the importance of both efficacy and potency in defining the pharmacological profile of the drug and highlight the need to consider both

factors when evaluating drug effects. The goodness of a drug cannot be easily summarised by one single number. For example, if a drug is highly efficacious but has a low potency (i.e. a large dosage is required to reach the high efficacy), then it may not be safe due to potential toxicity for healthy cells at high doses. Conversely, if a drug shows high potency, requiring only a small dosage to achieve some therapeutic effect, but does not exhibit strong efficacy on its own, it may still hold potential as part of a combination therapy. Such a drug might engage a pathway that enhances drug action at low dosages, making it a valuable candidate to be combined with other drugs to achieve a synergistic effect. Without delving into both the dimensions of efficacy and potency, critical information about the drug will be partially lost.

A commonly used parameterisation for these dose-response curves is the sigmoidal function. This shape is preferred because it captures the key characteristics of how most drugs interact with their targets at different concentrations, consisting of three phases. In the initial phase, low drug concentrations produce little to no effect. In the middle phase, increasing dosages lead to a rapid increase in effect. In the final phase, further increases in concentration result in minimal additional effect. This pattern is consistent with the behaviour of many biological systems, where receptors or enzymes become saturated at high drug concentrations, limiting the maximal response. This leads to the Hill Equation [31], a widely used mathematical model that describes the sigmoidal nature of the dose-response relationship. It has been the classic choice in modelling the dose-response relationship of monotherapies in pharmacology.

3.1.1 The Hill equation

Before defining the Hill Equation, for the consistency of notations through this thesis, we define the following parameters and keep it the same across the thesis:

- E_0: The base response value, or condition response value, i.e. the response $f(0)$ when no drug is applied.
- E_1: The $\mathrm{E_{inf}}$, which is equivalent to $f(\infty)$. This is a measure of the *efficacy* of a monotherapy.
- C: The $\mathrm{IC_{50}}$, i.e. the dosage x such that the equation $f(x) = \frac{1}{2}(E_0 + E_1)$ is satisfied, or in another word, half of the maximal inhibitory effect is achieved. This is a measure of the *potency* of a monotherapy.
- H: The Hill slope. It parameterises the steepness of the sigmoidal curve. A larger H indicates a steeper curve.

Now we define the Hill equation [31]

$$f_{\text{hill}}(x) = E_0 + \frac{E_1 - E_0}{1 + (\frac{C}{x})^H}. \tag{3.1}$$

In summary, the dose-response relationship in monotherapy provides critical insights into how a drug interacts with a biological system, informing decisions about dosage and potential efficacy. By analyzing parameters such as $E_{\max}$, IC_{50}, and the Hill slope, researchers can characterize the potency and efficacy of a drug, paving the way for its application in clinical settings. Understanding these fundamental concepts and definitions is essential for further exploration of drug interactions and the development of more effective therapeutic strategies.

3.2 Monotherapy quantification frameworks

While the Hill equation is widely used, other mathematical models and packages have been used for modelling dose-response relationships. Ritz et al. developed `drc` [65], a comprehensive R package specifically developed for modelling the dose-response relationship. As listed in Fig. 3.3, there are plenty of parameterisations available as a choice, within which the four-parameter generalised log-logistic model is equivalent to the Hill Equation, and is the most popular option within all listed model types. Larras et al. proposed `DRomics` [45], an R package that chooses from a family of five simple models built to describe a wide variety of monotonic and biphasic dose-response curves for each set of input data.

Despite the abundant choices available, alternative models are often less popular due to their complexity and reduced interpretability compared to the Hill equation. Moreover, their lack of generalisability across different datasets makes them less suitable as a universal framework for dose-response analysis. As a result, the Hill equation remains the preferred choice for modelling the dose-response curves of monotherapies. This preference illustrates the principle of Occam's Razor, where the simplest model that adequately explains the observed phenomena is favoured over more complex alternatives.

However, one important characteristic of the Hill equation, as well as other discussed models and parameterisations, is their inherently deterministic nature. These models typically estimate uncertainty through methods such as standard error or parametric bootstrap. As discussed in Section 2.3, standard error estimates are reliable only when a large sample size is available, a condition often unmet in pharmacological data, whereas parametric bootstrap methods rely on strong assumptions regarding the distribution of the synthetically generated data, which may significantly deviate from the true underlying distribution. Consequently, both approaches may lead to inaccurate uncertainty estimation.

Model type	Model function (*f*)	Function in *drc*
Generalized log-logistic	$c + \frac{d-c}{(1+\exp(b(\log(x)-\log(e))))^f}$	`llogistic()`
Brain-Cousens	$c + \frac{d-c+fx}{1+\exp(b(\log(x)-\log(e)))}$	
Cedergreen-Ritz-Streibig	$c + \frac{d-c+f\exp(-1/(x^\alpha))}{1+\exp(b(\log(x)-\log(e)))}$	`cedergreen()`
(0 < α < 1 is usually fixed in advance)		
Log-logistic fractional polynomial	$c + \frac{d-c}{1+\exp(b(\log(x+1))^{p_1}+e(\log(x+1))^{p_2})}$	`fplogistic()`
Log-normal	$c+(d-c)\Phi(b(\log(x)-\log(e)))$	`lnormal()`
(Φ: distribution function for a normal distribution)		
Weibull I	$c+(d-c)\exp(-\exp(b(\log(x)-\log(e))))$	`weibull1()`
Weibull II	$c+(d-c)(1-\exp(-\exp(b(\log(x)-\log(e)))))$	`weibull2()`
Gamma	$c + (d-c)\tilde{\Gamma}(bx, e, 1)$	`gammadr()`
($\tilde{\Gamma}$: distribution function for a Γ distribution)		
Multistage	$c+(d-c)\exp(-b_1-b_2\,x-b_3\,x^2)$	`multi2()`
NEC	$c+(d-c)\exp(b(x-e))$ for $x > e$ and d otherwise	`NEC.4()`

doi:10.1371/journal.pone.0146021.t001

Figure 3.3: The available parameterisations for dose-response relationship in `drc` [65].

To address these limitations, Bayesian methods have been developed, offering a more robust framework for incorporating and estimating uncertainty.

In particular, Wang et al. developed a Gaussian Process (GP)-based framework to assess pharmacological responses and identify biomarkers, using the inherent uncertainty estimates provided by Gaussian Processes to enhance the robustness of the analysis [78]. The GP framework inherently provides robust uncertainty quantification, which is crucial for assessing the reliability of any derived statistics such as the IC_{50} values. By formulating the model as a mixture, with one component for the drug response and another for outliers, GPs handle outlying measurements effectively, leading to more accurate estimates of drug efficacy and potential biomarkers. By utilising GPs, the model can capture and quantify uncertainty, making it a powerful tool in biomarker detection, surpassing traditional ANOVA-based methods by incorporating uncertainty directly into the statistical analysis.

Despite these advantages, the flexibility of GPs comes with the challenge of unconstrained behaviour outside observed dose ranges, potentially leading to counter-intuitive extrapolations. This study acknowledges this by highlighting increased posterior variance in unobserved regions, indicating high uncertainty rather than relying on potentially misleading mean estimates. Although methods exist to constrain the behaviour of the GPs, the study opts for a less restrictive approach to avoid underestimating uncertainty, proposing that the increased variance effectively signals when predictions are less reliable.

More recently, Tansey et al. introduced an end-to-end Bayesian dose-response generative model that combines prior biological knowledge with observed data to provide more accurate and interpretable estimates of drug efficacy and potency [74]. By adopting an empirical

Bayesian framework, various sources of uncertainty inherent in these experiments are effectively captured through the posterior variance of the parameters. These sources of uncertainty include measurement error, natural variation in cell growth, and heterogeneity in drug response. The generative model is comprehensive, detailing the entire experimental process from the initial raw fluorescent counts. It also integrates cell line and drug properties as auxiliary data, which are incorporated through a neural network. The results demonstrate that this generative model not only outperforms existing state-of-the-art methods for estimating dose-response curves but also successfully captures biologically meaningful interactions between molecular features and drug responses. However, while being sophisticated and well-tailored for modelling the dose-response relationships of monotherapies, it has not yet been extended to combinations, and significant challenges are anticipated in such an extension. As the number of parameters doubles when modelling combinations, the increased complexity can lead to difficulties in achieving convergence during optimisation and risks overparameterisation.

3.3 Combination quantification frameworks

Compared to the one-dimensional case where the Hill equation has been the gold standard, the best approach to modelling two-dimensional dose-response surfaces $Y = f(x_1, x_2)$, where x_1 and x_2 are the dosages of the two drugs and y is the corresponding response, is still an open question. The goal in this context is to determine whether a drug combination is *synergistic*, *antagonistic*, or neither. A *synergistic* combination occurs when the combined effect of the two drugs is greater than applying the drugs individually. Conversely, a *antagonistic* combination refers to the case where the combined effect is less than expected, indicating that one or both drugs inhibit each other's activity. If neither of these conditions is met, the combination is said to be *additive*, where the combined effect of the two drugs is the same as their independent effects.

To model the presence of synergy, various frameworks have been developed based on different assumptions, without a consensus in the field [16]. These models can generally be divided into two types. The first type computes the expected response $Y_e = f(x_1, x_2)$ assuming that the combination is neither synergistic nor antagonistic, and then computes the deviation $D = Y_e - Y$ between the data Y and the expected response Y_e to decide whether the combination is synergistic. The most popular frameworks within this type are the Bliss model and the Loewe model. The second type, on the other hand, directly models the 2-dimensional dose-response surfaces of the drug combinations. These include MuSyC [83], BRAID [76] and the effective dose model [88].

3.3.1 The classic frameworks: Bliss, Loewe and HSA

The Bliss model follows the multiplicative survival principle (MSP), which assumes the probability of a cell surviving treatment by drug 1 is independent of the probability of the same cell surviving treatment by drug 2 [10]. In another word, the MSP assumes that the drugs act independently of each other. If we denote the responses from drug 1 and drug 2 as Y_1 and Y_2 respectively, and denote the response from the combination of the two drugs as Y_{12}, then the null model assumes

$$Y_{12} = \frac{Y_1 Y_2}{100}$$

where the normalisation (division by 100) is due to the scale of the responses being between 0 and 100. The deviation from this expected response is then used to quantify the degree of synergy or antagonism. If

$$Y_{12} < Y_1 Y_2,$$

then synergy is suggested. If

$$Y_{12} > Y_1 Y_2,$$

then antagonism is suggested.

The Loewe model, following the Dose Equivalence Principle (DEP), states that two drugs are synergistic if a lower combined dose achieves the same effect as a higher dose of either drug alone [51]. Mathematically, for a combination of drugs 1 and 2 at dosages x_1 and x_2, the relationship is defined as:

$$\frac{x_1}{C_1} + \frac{x_2}{C_2} = 1$$

where C_1 and C_2 are respectively the doses of drugs 1 and 2 required to achieve the same effect individually. If

$$\frac{x_1}{C_1} + \frac{x_2}{C_2} < 1,$$

then synergy is suggested. If

$$\frac{x_1}{C_1} + \frac{x_2}{C_2} > 1,$$

then antagonism is suggested.

Some other popular choices include the Highest Single Agent (HSA) model [28], which states that if the effect of a combination exceeds the effect of either of the single drugs, i.e.

$$Y_{12} < \min\{Y_1, Y_2\},$$

then the combination is synergistic. Otherwise, if we have

$$Y_{12} > \min\{Y_1, Y_2\},$$

then the combination is antagonistic.

Synergy scores are then obtained by comparing to the null model of Bliss, Loewe or the HSA model. This type of synergy computation is widely used due to its simplicity and low computational cost, as no additional model fitting is required. For example, the DREAM challenge defines its synergy score based on the Loewe model using the Combenefit software [20]. A positive synergy score represents a synergistic effect for the combination compared to applying the two single drugs individually, while a negative synergy score represents an antagonistic effect. The NCI-ALMANAC dataset defines its metric based on a modification of the Bliss model [33].

However, each of these models has its shortcomings. The Bliss model is biased against combinations of drugs with a moderate level of efficacy, compared to those with strong or weak efficacy [83]. This is because the Bliss model is based on the probabilities, or percentages, of the cells affected, which is bounded between 0 and 1. However, in reality, percentage effect data usually saturates before it goes close to 0. For example, if the cell survival percentages for drug 1 and drug 2 are 80% and 90% respectively, then the combination would be synergistic if the survival percentage is lower than $80\% \times 90\% = 72\%$ when the drugs are combined at the corresponding dosages, which is easier than the case where the two single drugs have survival percentages of 40% and 30% respectively, in which case a synergism requires a combined effect of lower than 12%, which would be considerably more difficult than the former case.

For the Loewe model, a subset of the 2D null surface is undefined. The null response is only defined when the response of the combination is not stronger than the weaker drug's maximal effect. However, the area where the response exceeds the weaker drug's maximal effect is crucial for deciding the synergy of a combination.

As for the HSA model, it does not tell the full story either. Unlike the Bliss model and the Loewe model which provide a null hypothesis for multiple dosages on the 2D dose-response surface, the HSA model only models the maximal efficacy, ignoring the dosage required to reach the maximal efficacy. As a result, the HSA model does not reflect whether a combination improves or worsens the toxicity.

These models also do not universally agree with each other, with common contradictions occurring on the same set of data [83]. This is inevitable because they are based on different assumptions, resulting in different interpretations for the definition of synergy and antagonism. One simple example is that if we have $Y_{12} = 40, Y_1 = 50, Y_2 = 50$, then it is considered to be synergistic according to the HSA model, but antagonistic according to the Bliss model because the multiplicative survival principle in the Bliss model requires

at least $Y_{12} \leq \frac{Y_1 Y_2}{100}$ for the combination to be deemed non-antagonistic, but we have $\frac{Y_1 Y_2}{100} = 25 < Y_{12}$.

The recommendation to date is still to apply multiple frameworks and choose the best-fitting model according to some metric or test, such as the χ^2 test [12]. Scientists often make decisions based on experience or practical needs while selecting the frameworks between these options [87] [46].

3.3.2 The deterministic and parametric models

More recently, parametric methods that directly model the 2-dimensional dose-response surfaces of the drug combinations have emerged, including BRAID [76], the effective dose model [88] and MuSyC [83]. These methods seek to better represent the underlying biological processes in combinations by using complex parametric functions $Y = f(x_1, x_2)$ to model the dose-response surface in 3D where the first two axes are the dosages of the two drugs x_1 and x_2 respectively, and the third axis is the drug response Y.

MuSyC is an attempt to unify the synergy metrics for dose-response data in drug combinations. It is the first method that accounts for both potency and efficacy and decouples them in the model. The Hill equation is extended to a generalised 2D Hill equation in MuSyC. In the model outputs, the parameters quantifying synergistic potency and the parameters quantifying synergistic efficacy are separated. As a result, a more complete picture can be said about each combination. It can be inferred whether an improvement comes from a lower toxicity or a stronger efficacy.

The MuSyC framework is a promising step forward in decoupling potency and efficacy and unifying the pharmacology framework. Fig. 3.4 (from [83]) illustrates how MuSyC extends the Hill equation and the transition model from single-drug to two-drug systems so that potency, efficacy and cooperativity (i.e. the steepness or the sensitivity of the dose-response curve) are decoupled by being quantified through distinct parameters.

However, it has its shortcomings in practice. It contains six tiers of model complexity, with the number of parameters increasing from 2 to 12. Each tier of the model is associated to a different level of complexity and assumptions. For each combination, after training models on all six tiers, the best model for the data is chosen based on minimising the deviance information criterion. While this procedure is statistically sound, the associated complexity makes it difficult to be seamlessly applied in real-world scenarios, where a single consistent model instead of six is usually desired.

BRAID is another parametric method for modelling the effect of two drugs. Similarly to MuSyC, the parameterisation in BRAID is also inspired by the Hill equation in one dimension. However, only the efficacy is inferred in BRAID, while the potency is not decoupled. The effective dose model [88], on the other hand, is a method that focuses on generalising the parameterisation of drug effect to more than two drugs. Its drawback is

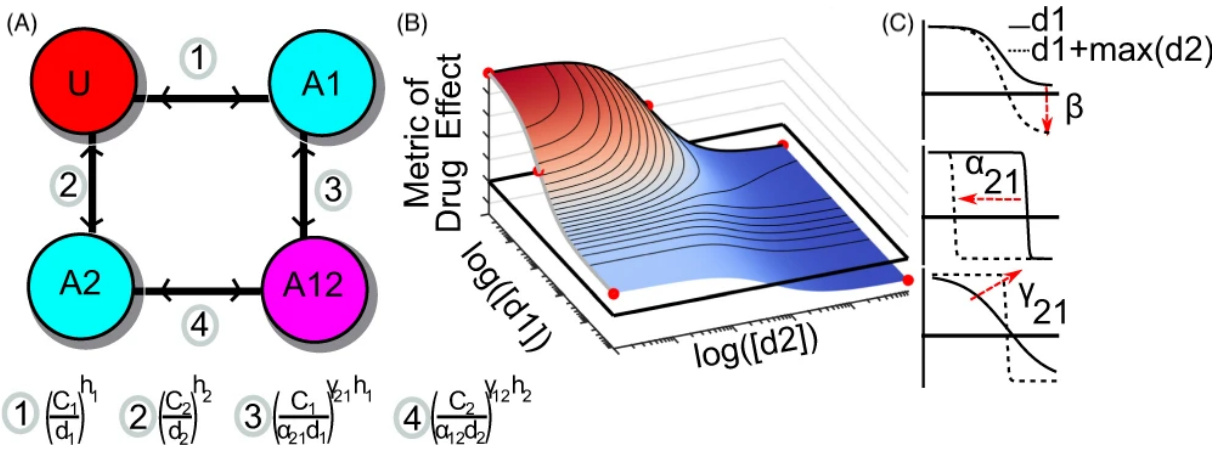

Figure 3.4: The MuSyC model [83]. Subfigure (A) depicts the four-state transition model for a two-drug system. Unaffected cells are indicated in red with the letter "U". Cells affected by Drug 1 or Drug 2 (but not both) are indicated in cyan with "A1" or "A2" respectively. Cells affected by both drugs are indicated in magenta with "A12". d_1 and d_2 denote the dosages of the two drugs. C_1 and C_2 denote the IC50 of the two drugs. h_1 and h_2 denote their Hill slopes. α_{21} and α_{12} quantify how the potency of one drug is affected by the interaction with the other drug, whereas γ_{21} and γ_{12} quantify how the cooperativity of one drug is affected by the interaction with the other drug. Subfigure (B) illustrates the dose-response surface for the combination of two drugs, parameterised by the generalised 2D Hill equation in MuSyC. Subfigure (C) illustrates how the additional parameters in MuSyC correspond to the transformations of the dose-response surface. These transformations separately quantify synergistic efficacy, synergistic potency and synergistic cooperativity, effectively decoupling the different aspects of drug synergy. This figure is adapted from [83].

that potency is not directly parameterised.

Despite the rapid progress in the quantification of drug combinations, all of the aforementioned methods (including those in Section 3.3.1) either do not incorporate uncertainty, or compute the uncertainty in a way that is prone to failure cases. All these methods are deterministic in nature. They estimate the uncertainty through either standard error or parametric bootstrap. However, the standard error is only accurate when a large number of samples is available, which is not the case in pharmacology data. Parametric bootstrap relies on a strong assumption on the distribution of the synthetic data generated during bootstrap, which is far from the ground truth distribution. In either case, the uncertainty estimation is not accurate.

3.3.3 The probabilistic models

Despite most existing methods being deterministic, there has been some recent work on developing probabilistic methods to model the dose-response relationship of drug combinations and incorporate uncertainty in their outputs. Hand-GP ([70]) is a non-parametric model based on the combination of the Hand model with Gaussian processes,

providing more believable uncertainty estimation than MuSyC in some cases. However, Hand-GP does not incorporate the 1D Hill equation that imposed biological constraints useful for providing interpretable model outputs. For example, monotonicity of the monotherapy fitted curve is not enforced in Hand-GP, meaning the model may produce a dose-response surface that is unlikely to occur in the *in vitro* setting ([74]). More importantly, due to the non-parametric structure of Hand-GP, the parameters describing potency and efficacy are lost.

The bayesynergy model ([66]) is a Bayesian framework that models synergistic interaction effects using Gaussian Processes, which provide uncertainty quantification. Although flexible, its formulation is based on the Bliss independence assumption that is biased against drug combinations with a moderate level of efficacy ([83]). The bayesynergy model also does not separate out synergistic potency and efficacy.

3.4 Synergy and its myth

The concept of *synergy* in drug combinations has been a focal point of pharmacological research, aiming to identify combinations of drugs that produce a greater effect together than the sum of their individual effects. However, several challenges exist in defining and quantifying synergy, which can lead to misleading conclusions and limit the effectiveness of the drug development process.

First, a universally agreed gold standard is absent for the quantification of the synergy [72]. Various definitions of synergy are derived from the models in Section 3.3.1, each with its specific assumptions and limitations as discussed in the section. Each model may classify a drug combination as synergistic or antagonistic based on differing criteria, leading to conflicting interpretations of the same data. This results in inconsistencies in synergy quantification across different studies, making it difficult to compare, aggregate or reproduce findings.

Secondly, traditional synergy quantification frameworks do not decouple potency and efficacy. In many synergy models, these two distinct aspects are entangled, leading to ambiguous interpretations of synergy. For example, a combination may appear synergistic due to enhanced potency, meaning that lower doses achieve the desired effect, which is beneficial for reducing toxicity. Conversely, a combination might be synergistic due to enhanced efficacy, increasing the maximum possible therapeutic effect. In some other cases, both could be simultaneously true for a combination. However, the difference between these two aspects of synergy cannot be differentiated if only a single synergy score is computed. Disentangling these factors is crucial, as it allows researchers to better understand the nature of the interaction and the potential benefits of the combination. An ideal combination would minimise the dose while retaining an effective outcome, and the

synergy score shall be capable of covering this aspect. Recently, MuSyC [57] has become the first synergy framework to address this issue. It would be crucial to account for this while developing novel synergy frameworks.

Finally, current synergy quantification methods do not account for noise and uncertainty. Traditional synergy scores such as Bliss and Loewe, the two most widely used frameworks, only yield a single score, without addressing the inherent variability and uncertainty present in the experimental data. This approach can be problematic because a synergy score, for instance, of 10 accompanied by an uncertainty measure of 30, indicates that the ground truth score is likely to be anywhere between -20 and 40 and that the true level of synergy cannot be reliably determined (assuming a positive score represents synergy and a negative score represents antagonism). A high level of uncertainty in the synergy score implies that any conclusions about the combination being synergistic may be unreliable. Ignoring these uncertainties may result in overconfidence in the results and the potential pursuit of drug combinations that lack genuine efficacy.

In this study, we aim to address these three key challenges in synergy quantification by developing a framework that overcomes the limitations of existing models. First, we seek to minimise assumptions about the underlying interactions between drugs to create a more generalisable approach to synergy analysis. Secondly, we aim to decouple potency and efficacy to provide a clearer understanding of how drug combinations interact, enabling a more precise identification of their therapeutic benefits. Finally and most importantly, we will incorporate a Bayesian framework to estimate uncertainty, allowing us to account for variability in experimental data and provide more reliable synergy assessments. By addressing these aspects, our approach will enhance the reliability and actionability of synergy quantification, supporting more effective drug development strategies for the subsequent stages of the drug discovery pipeline, which is an unmet need.

Chapter 4

SynBa: A new inference method of drug combination synergies

As described and discussed in Section 3.4, current frameworks for synergy lack a gold standard, do not decouple potency and efficacy, and lack a principled framework for uncertainty estimation. These lead to misleading conclusions and limit the effectiveness of the drug development process. To define a common framework for drug combination, we need a less ambiguous definition of *synergy* ([72]). When a combination is said to be *synergistic*, it is unclear whether it implies that the combination is desirable in terms of its *potency* or its *efficacy*. Potency is the amount of dosage required for a drug to produce a specified effect, whereas efficacy is the degree of the beneficial effect produced by the drug ([57]). A strong synergistic potency implies that toxicity may be reduced when the drugs are combined, which is crucial for avoiding overdose, whereas a strong synergistic efficacy implies that the combination increases the maximal possible effect. Both aspects are relevant for progressing with pre-clinical and subsequent clinical investigations. For example, there are situations in clinical research where a drug is paired specifically to enhance potency, without necessarily aiming to increase efficacy.

Until recently, in quantification frameworks including Bliss, Loewe, BRAID and the Effective Dose Model, potency and efficacy are entangled within the concept of *synergy* as defined in the traditional quantification frameworks. To tackle this problem, [57] developed MuSyC, a framework that decouples potency and efficacy following the principles of the generalised Hill equation. Although the MuSyC approach is effective for modelling both potency and efficacy, the model does not fully explore the challenge of model-based estimation of uncertainty. There are various sources of uncertainty associated with drug-sensitivity modelling. Firstly, the biology of dose-response relationships is still unknown despite existing efforts. This leads to uncertainty associated with insufficient scientific knowledge. Secondly, there are systematic and random errors arising from the experimental procedures. Thirdly, biological variation exists among cell lines of the same

disease, resulting in a further source of noise. Finally, uncertainty also stems from limited information that can be extracted due to the small size of available data (*epistemic uncertainty*).

Given these multiple sources of uncertainty, it is often impossible to reach an accurate estimate of the quantities of interest, e.g potency of a monotherapy, or the synergistic effect (in terms of either potency or efficacy) of a combination. Since the *best* single deterministic estimates might not be reached, here we focus on accurately quantifying uncertainty in the model as deviation from the true estimate given the noise (e.g. the credible range of model parameter estimates/predictions).

Most existing frameworks for drug combinations either do not compute the uncertainty or estimate it by standard error ([88]) or parametric bootstrap ([83]), each of which contains unrealistic assumptions as discussed in Section 2.1. For this reason, here we approach uncertainty estimation with a Bayesian framework, which incorporates the uncertainty by treating all parameters of interest as probabilistic quantities. This enables us to continuously model the uncertainty in our estimates as the number of measurements grows, without becoming over-confident. Moreover, the estimated quantities of interest are obtained simultaneously with their uncertainties, making this approach computationally efficient.

Recent works including Hand-GP [70] and bayesynergy [66] have also opted for the Bayesian framework, but both of them have modelled the synergistic interactions using GPs which are non-parametric, making it difficult to interpret the uncertainties associated with the model outputs.

In this study, we design a flexible **Ba**yesian framework to infer **syn**ergistic effects of drug combination (**SynBa**) where (1) the classic Hill equation is preserved to produce estimates of efficacy and potency (2) the existing biological knowledge or insight from historical data may be conveniently added through the prior distribution over parameters in the model. We opt for the parametric Bayesian setting so that the model parameters can still be interpreted and explained through the classic Hill equation. In SynBa we will use MuSyC as a baseline framework to decouple synergistic potency and efficacy and add probabilistic inference to provide outputs and their associated uncertainty. This is to design a framework estimating the most favourable scores and efficiently provide optimal candidates to the drug discovery pipeline with actionable decision-making criteria derived from the model outputs. SynBa is also the first synergy framework that simultaneously provides principled uncertainty estimation, preserves the Hill equation and decouples the synergistic efficacy and potency. Our presentation of SynBa has been published as **SynBa: improved estimation of drug combination synergies with uncertainty quantification** in *Bioinformatics 39* [86].

With our approach, when a combination is predicted to be synergistic, a level of

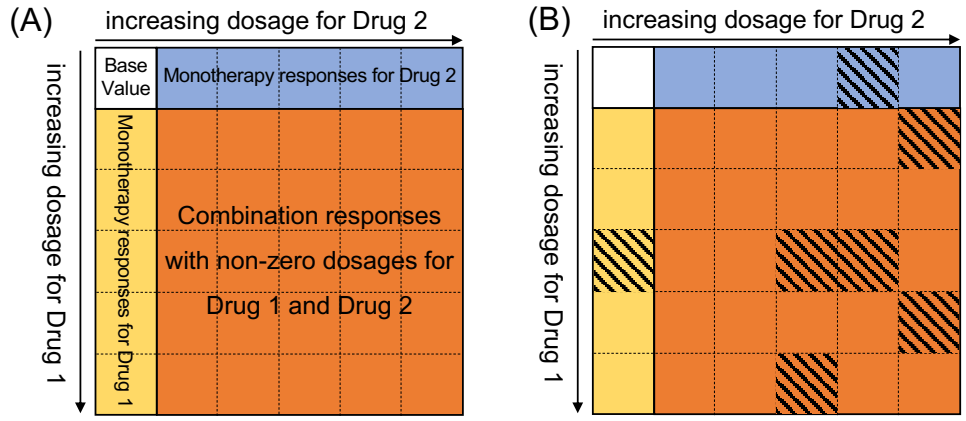

Figure 4.1: Description of the datasets for the proposed framework. (A): A typical dose-response matrix, where the top-left entry is the base value, the first row contains the monotherapy responses for the first agent, the first column contains the monotherapy responses for the second agent. The remaining entries are the responses when a combination of the two drugs is applied. (B): A dose-response matrix where not all responses are available for training. The shaded cells represent the test data.

confidence will be quantified for this prediction, together with a level of improvement in efficacy or in potency. The associated uncertainty can guide decision on further laboratory experiments to proceed to the next phase. This will provide *actionable metrics* for the subsequent stages of the drug discovery pipeline, which is an unmet need.

4.1 Methods

Our proposed method can be used for both analysing existing dose-response data and predicting unseen dose-response data for a given monotherapy $\mathcal{D} = (X, Y) = \{(x_i, y_i)\}$ or a given combination $\mathcal{D} = (X, Y) = \{(\mathbf{x}_i, y_i)\}$. The covariate can be a scalar x_i (in monotherapy) or a vector $\mathbf{x}_i$ (in combination) corresponding to the drug dosages. The response y_i can be defined as cell growth or inhibition of growth, depending on how the data are collected. In this study, we focus on inhibitory datasets, where a large dosage typically results in growth inhibition. In this case, y_i is defined as the percentage of growth-inhibited cells. Nevertheless, our method can be easily modified to accommodate the opposite setting where the drug response is enhancing growth with respect to the dosages.

Fig. 4.1 (A) illustrates a typical dose-response matrix for a combination from a screening, where the first row and column contain monotherapy data and the remaining entries contain combination data. The core aim of our method is to infer the synergistic potency and efficacy given such a matrix (or a vector in the case of monotherapy). To accomplish this, we designed **SynBa**, a **Ba**yesian framework for the inference of **Syn**ergistic effects of drug combinations. SynBa is defined by a prior distribution $p(\boldsymbol{\theta})$

for the parameters $\boldsymbol{\theta}$ and a likelihood function $p(Y \mid \boldsymbol{\theta}, X)$ for the drug responses Y. The likelihood function describes the probability of the responses given the dosages X and the parameters $\boldsymbol{\theta}$. The prior distribution encodes the existing belief or knowledge about the parameters. In monotherapy, we have $\boldsymbol{\theta} = \{E_0, E_1, C, H, \sigma\}$, whereas in combination, $\boldsymbol{\theta} = \{E_0, E_1, E_2, E_3, C_1, C_2, H_1, H_2, \alpha, \sigma\}$. The likelihood function and parameters encode the shape of the dose-response curve (defined in Boxes 1 and 2 and described in more detail below).

In addition, our method provides a way of predicting unseen dose-response data for both monotherapies and combinations. For example in the case of Fig. 4.1 (B), the quantity of interest would be the posterior predictive distribution of the response, i.e.

$$p(\tilde{y} \mid X, Y, \tilde{\mathbf{x}}) = \int_{\boldsymbol{\theta}} p(\tilde{y} \mid \boldsymbol{\theta}, \tilde{\mathbf{x}}) p(\boldsymbol{\theta} \mid X, Y) d\boldsymbol{\theta} \tag{4.1}$$

where $\tilde{\mathbf{x}}$ is the untested dosage of interest, $\tilde{y}$ is the predicted response and $p(\boldsymbol{\theta} \mid X, Y)$ is the posterior distribution over the parameters $\boldsymbol{\theta}$ given training data (X, Y).

4.1.1 Overview of SynBa: Monotherapies

We begin by defining SynBa for monotherapy screens, where the dosage x is a scalar. The likelihood function for the response y is based on the Hill equation ([31]), which has been the classic choice to model pharmacology data:

$$E(x) = E(0) + \frac{E(\infty) - E(0)}{1 + (\frac{C}{x})^H} \tag{4.2}$$

where x is the dosage of the drug, $E(x)$ is the corresponding measured response, and H controls the slope of the curve. The interpretation of C depends on the problem and the dataset of interest. In this study, as the focus is on inhibitory datasets, C represents the dosage required to inhibit the given biological process or biological component by 50%, known as IC_{50}. C quantifies the potency of the monotherapy in the study.

The base level of the monotherapy is denoted by $E_0 := E(0)$, which is the response when no drug is applied. The efficacy is quantified by $E_1 := E(\infty)$, or denoted as $\mathrm{E}_{\mathrm{inf}}$, which is the maximal inhibition when a sufficiently large dosage x is applied.

Box 1 defines the prior distribution for the parameters and the likelihood model for the response given the dosages, which are also illustrated in Fig. 4.2.

To account for observational noise in the data, we define a noise model for y centred around $E(x)$. For each fixed experiment setting (i.e. a fixed cell line treated by a fixed monotherapy or combination in the same laboratory environment), y is modelled as having Gaussian noise that is conditionally independent given the dosage: $y \sim \mathcal{N}(E(x), \sigma^2)$. The assumption of conditional independence is valid because in screenings such as DREAM,

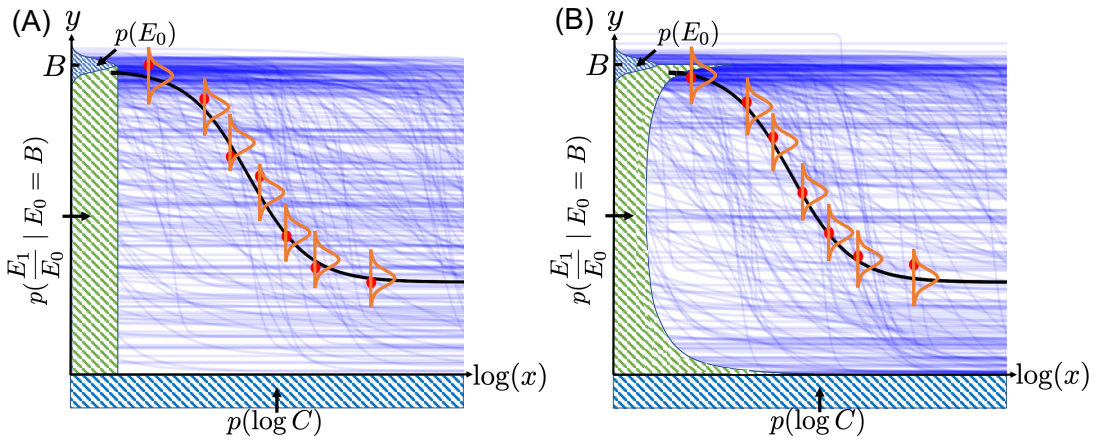

Figure 4.2: Two options (A) and (B) for the monotherapy prior model. The prior for $\log C$ is plotted underneath the x-axis, which is uniform in both options. The priors for E_0 (which is Gaussian in both options) and the normalised E_1 are plotted along the y-axis. Given E_0, the normalised E_1 is uniform in Option (A) and follows a Beta(0.46, 0.58) distribution in Option (B). The 300 blue curves are random samples from the expected prior responses $\mathbb{E}[Y|\boldsymbol{\theta}]$ where $\boldsymbol{\theta}$ are sampled from the prior distributions defined in Box 1. The seven red points illustrate an example set of monotherapy dose-response data $\mathcal{D}$. The black curve is a sample from the expected posterior responses $\mathbb{E}[Y|\boldsymbol{\theta}, \mathcal{D}]$ after the model is fitted to the data $\mathcal{D}$, whilst the orange bell-shaped curves illustrate the i.i.d. Gaussian noise for the responses.

the measurements for each monotherapy or combination are performed independently in different plates, instead of being performed sequentially. Therefore, given a set of measurements for a monotherapy or a combination, the noise level of their corresponding responses is independent. The i.i.d. noise is illustrated as orange bell-shaped curves in Fig. 4.2.

As a result, the joint prior distribution for the response and the parameters of the curve given the dosages is

$$p(Y, E_0, E_1, C, H, \sigma \mid x) = p(Y \mid E_0, E_1, C, H, \sigma, x)p(E_1 \mid E_0)p(E_0)p(C)p(H)p(\sigma). \quad (4.3)$$

where the likelihood model for the responses Y given the dosage x is defined as

$$p(y \mid E_0, E_1, C, H, \sigma, x) = \mathcal{N}(E_0 + \frac{E_1 - E_0}{1 + (\frac{C}{x})^H}, \sigma^2) \quad (4.4)$$

independently for all $y \in Y$, where $E_0 := E(0)$, $E_1 := E(\infty)$ and σ is the standard deviation of the noise level of y,

The datasets are normalised by measuring cell inhibition when no drug is added (at dose zero). However, due to the noise in the biological process, the inhibition at dose zero would not be the same if the experiment is repeated multiple times. Thus, the normalization procedure itself contains uncertainty. Therefore, we define E_0 to be probabilistic instead

of a fixed initial value. A Gaussian prior with mean B and variance $0.03B$ is given for E_0, where B is the normalized inhibition of a dose zero, e.g. $B = 100$ in DREAM. The variance of this prior is defined to be $0.03B$ so that it is flexible enough to allow for errors, but not too conservative.

Next, we define the prior for the normalised maximal response

$$\tilde{E_1} := \frac{E_1}{E_0}.$$

We normalise the maximal response so that its range will always be between 0 and 1. This makes it easier to apply our priors to a wide range of datasets regardless of the original range of their responses. The prior for $\tilde{E_1}$ may be defined in various ways depending on whether to insert existing knowledge or historical information. One option is to remain uninformative and impose a uniform prior, as shown in Fig. 4.2 (A). By doing this, we do not prefer any particular level of drug viability prior to seeing any dose-response measurements.

Alternatively, we may make use of the existing information from the monotherapy data available. According to the single agent datasets in DREAM, the empirical distribution of $\mathrm{E_{inf}}$ has a high density on both extremes of the range, with 17.9% of them smaller than 0.05 and 15.1% equal to 1 (after normalising to the interval $[0, 1]$). Using maximum likelihood estimation to fit a $\mathrm{Beta}(a, b)$ distribution to these $\mathrm{E_{inf}}$ values, we obtain $a = 0.46$ and $b = 0.58$. To account for this information, we define $\mathrm{Beta}(0.46, 0.58)$ as the second option of the prior for the normalised maximal response, as illustrated in Fig. 4.2 (B). This prior is consistent with the biological behaviour that a drug administered with a sufficiently large dose will either kill the majority of the targeted cells if effective, or very few of them if not effective. The choice of this empirical prior shows how existing knowledge or information on the dynamic or kinetic of drugs can be conveniently added to SynBa through its priors.

A uniform prior is imposed for $\log C$ (the logarithm of $\mathrm{IC_{50}}$), with C bounded by δ and M. The values of δ and M depend on the dataset and the unit of the dosages. In this work, we define δ to be smaller than any non-zero dosage in the dataset (i.e. $0 < \delta < \min\{x \mid x > 0, (x, y) \in \mathcal{D}\}$) and M to be larger than any dosage in the dataset (i.e. $M > \max\{x \mid (x, y) \in \mathcal{D}\}$). For example, $\delta = 10^{-10}$ and $M = 10^6$ is a viable choice for DREAM, whereas in NCI-ALMANAC, we may have $\delta = 10^{-15}$ and $M = 10$. The idea is to be uninformative about $\log C$ a priori, since the ideal dosage range for the experiments is unknown and often unsuitable, either too small or too large. It is common that $\mathrm{IC_{50}}$ exceeds the maximum dosage. In other cases, $\mathrm{IC_{50}}$ may lie between zero dosage and the smallest non-zero dosage, due to the tested dosages being too large. Our method includes these possibilities a priori.

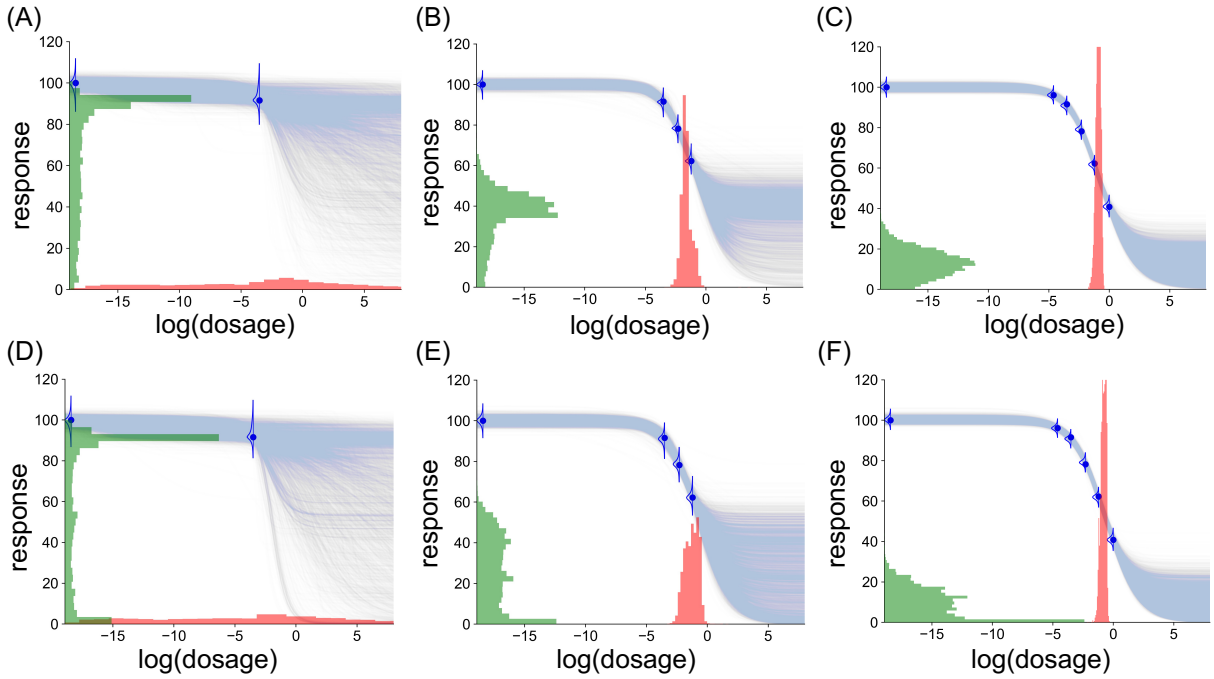

Figure 4.3: An example illustration of SynBa on a set of monotherapy data with six measurements. Each blue curve is a sample from $\mathbb{E}[Y \mid \boldsymbol{\theta}]$ where $\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid \mathcal{D})$. The distribution for IC_{50} is shown in red, whereas the distribution for $\mathrm{E}_{\mathrm{inf}}$ is shown in green. (A)-(C): The prior distribution for $\tilde{E}_1$ is $\mathcal{U}(0,1)$. (D)-(F): The prior distribution for $\tilde{E}_1$ is $\mathrm{Beta}(0.46, 0.58)$.

H and σ are both given a $\mathrm{lognormal}(0,1)$ prior because they are both non-negative and assumed to be moderately small. In addition, previous literature indicated that H is approximately lognormal ([83]). A $\mathrm{lognormal}(0,1)$ prior ensures that $P(H < 5) \approx 0.95$ and $P(\sigma < 5) \approx 0.95$ a priori.

In summary, the priors for SynBa are defined to be

$$\begin{aligned} p(E_0) &= \mathcal{N}(B, 0.03B), \\ p(\tilde{E}_1) &= \mathcal{U}(0,1) \text{ or } \mathrm{Beta}(0.46, 0.58), \\ p(\log C) &= \mathcal{U}(\log(\delta), \log(M)), \\ p(H) &= \mathrm{lognormal}(0,1), \\ p(\sigma) &= \mathrm{lognormal}(0,1), \end{aligned}$$

where $0 = x_1 < x_2 \leq x_3 \leq ... \leq x_n$ are the dosages, and $\delta \in (0, x_2)$ is a small non-zero value to avoid $\log C$ being undefined.

The blue curves in Fig. 4.2 are 300 random samples from the expected prior responses $\mathbb{E}[Y|\boldsymbol{\theta}]$ where $\boldsymbol{\theta}$ are sampled from the prior distributions. It can be observed that the curves cover a wide range of possibilities a priori. Yet, they are not excessively general, as all of them follow the Hill equation.

Note that the priors we have chosen are motivated by general knowledge and previous literature. When more knowledge exists for a specific combination, the prior can be conveniently adjusted to accommodate this, since the inference framework is agnostic to the choice of prior.

Fig. 4.3 is an illustrative example of SynBa trained on a monotherapy (the compound MTOR_1 treated on the cell line MDA-MB-231) with six measurements, taken from the DREAM dataset. The first row shows the resulting model with a uniform prior for the normalised E_{inf} (i.e. $\tilde{E}_1$), whereas the second row shows the model with the Beta$(0.46, 0.58)$ prior for $\tilde{E}_1$. We start with two measurements and add two additional responses each time. It can be observed that the posterior distribution for IC_{50} narrows down quickly for both models because the observed responses span across the range between 40 and 100, which provides sufficient information to estimate IC_{50} with low uncertainty. The posterior for E_{inf}, on the other hand, is more uncertain, due to the dosage range being too small to observe the convergence of the responses. In this case, the two models provide a visibly different posterior distribution for E_{inf}, due to the different priors. As shown in Fig. 4.3 (F), the maximum a posteriori probability estimate for E_{inf} is 0 when the Beta prior is imposed, which results from the inserted prior knowledge that the response is likely to converge to 0 if an effective (but non-zero) response is already observed. This example shows how the prior design affects the inference of the parameter uncertainty. Nevertheless, if we look at the posterior predictive distribution for the responses (with samples illustrated by light blue curves), the two models reach a similar conclusion. We will show in Result that the predictive performance of the model is insensitive to the choices of prior.

4.1.2 Overview of SynBa: Combinations

Extending SynBa to combinations of two drugs, the goal is now to model the dose-response surface $E(x_1, x_2) = f(x_1, x_2)$ where x_1 and x_2 are the dosages for the two drugs, whereas $E(x_1, x_2)$ is the response, and f is some class of function to be defined.

To define our likelihood model for the responses Y, we take inspiration from MuSyC but maintain some flexibility on their model assumptions. The effect of a drug in a system is usually described by the Hill equation that describes the state of equilibrium of a reversible process between an unaffected population and an affected one (the principle of *detailed balance*). To obey to this equation and its effects, in our model we incorporate the assumptions of the principle of detailed balance, of the proliferation rate of unaffected population and of the saturation of the maximum effect of the drug in the affected population. In MuSyC, these assumptions are defined in a nested structure in which levels are called *Tiers* (see Table S5 in [57]). In defining SynBa for combinations, we adopt the same model assumptions as *Tier 4* of the levels specified by MuSyC. This category

encodes the most complex class of models that still maintains the assumption of *detailed balance*. It is worth noting that, conversely to MuSyC, our model posterior covers all four tiers simultaneously. This is because *Tier 4* subsumes all lower tiers and thus for how our model is defined, they will not be eliminated from the posterior distribution, unless the evidence from the data is strongly against them. The concepts of Tiers as described in MuSyC would require a post-learning model selection. The use of a Bayesian approach avoids such a selection procedure whilst still maintaining biologically viable assumptions.

In contrast to MuSyC, in SynBa we choose to maintain the detailed balance assumption to also avoid over-parameterisation, which is likely to occur due to the limited data size for each drug combination set. For example, the majority of the combinations in the NCI-ALMANAC dataset has a data size of 15 (excluding the base level at dose zero), only 3 more than the number of parameters in MuSyC. In a real-world scenario, the data size may often be even smaller. Furthermore, with the detailed balance assumption, matrix multiplication and inversion are avoided, which lowers the computational cost.

As a result, we define the following joint distribution for the response and the parameters of the surface given the dosages $\mathbf{x} = (x_1, x_2)$,

$$\begin{aligned} & p(Y, E_0, E_1, E_2, E_3, C_1, C_2, H_1, H_2, \sigma, \alpha \mid \mathbf{x}) \\ = \; & p(Y \mid E_0, E_1, E_2, E_3, C_1, C_2, H_1, H_2, \sigma, \alpha, \mathbf{x}) \\ \times \; & p(E_1 \mid E_0) p(E_2 \mid E_0) p(E_3 \mid E_0) p(E_0) p(C_1) p(C_2) p(H_1) p(H_2) p(\sigma) p(\alpha), \end{aligned} \tag{4.5}$$

where the likelihood model for responses Y given the dosages is defined as

$$\begin{aligned} & p(y \mid E_0, E_1, E_2, E_3, C_1, C_2, H_1, H_2, \sigma, \alpha, \mathbf{x}) \\ = \; & \mathcal{N}\left(\frac{C_1^{H_1} C_2^{H_2} E_0 + x_1^{H_1} C_2^{H_2} E_1 + C_1^{H_1} x_2^{H_2} E_2 + \alpha x_1^{H_1} x_2^{H_2} E_3}{C_1^{H_1} C_2^{H_2} + x_1^{H_1} C_2^{H_2} + C_1^{H_1} x_2^{H_2} + \alpha x_1^{H_1} x_2^{H_2}}, \sigma^2\right) \end{aligned} \tag{4.6}$$

for all $y \in Y$, where $E_0 := E(0,0)$, $E_1 := E(\infty, 0)$, $E_2 := E(0, \infty)$, $E_3 := E(\infty, \infty)$, C_1 and H_1 are the monotherapeutic parameters associated for Drug 1, C_2 and H_2 are the monotherapeutic parameters associated for Drug 2, and α is an association parameter that controls how the two drugs are affected by the presence of each other.

The priors are

$$\begin{aligned} p(E_0) &= \mathcal{N}(B, 0.03B), \\ p(\tilde{E}_i) &= \mathcal{U}(0,1) \text{ or } \text{Beta}(0.46, 0.58) \quad \text{for } i = 1,2,3, \\ p(\log(C_i)) &= \mathcal{U}(\log(\delta), \log(M)) \quad \text{for } i = 1,2, \\ p(H_i) &= \text{lognormal}(0,1) \quad \text{for } i = 1,2, \\ p(\alpha) &= \text{lognormal}(0,1), \\ p(\sigma) &= \text{lognormal}(0,1), \end{aligned}$$

where $0 = x_{i,1} \leq x_{i,2} \leq ... \leq x_{i,n}$ are the dosages for drug i, and δ is a small non-zero value to avoid $\log C$ being undefined.

The definitions of the priors are a natural extension of the monotherapy model, with the same arguments being followed. The only new parameter is α, which follows a lognormal prior with median 1 because α is non-negative and equals 1 when the combination is neither synergistic nor antagonistic in terms of potency.

4.1.2.1 Inference of the synergy

After inferring the posterior for the parameters and their associated uncertainty, we focus on distinguishing the effect of efficacy and potency in drug combinations. MuSyC has defined metrics for both synergistic efficacy and synergistic potency, which is a promising step in decoupling potency and efficacy. However, the uncertainty for these two quantities has not been quantified systematically. Our model output includes not only quantification for the synergistic efficacy and the synergistic potency, but also a separate uncertainty estimation for each.

For the synergistic efficacy, one simple yet informative quantity is

$$\Delta\text{HSA} = \min(E_1, E_2) - E_3 \tag{4.7}$$

which is the change in the maximal effect between the combination and the more effective single drug of the two ([28]). A positive score indicates synergistic efficacy. As E_1, E_2 and E_3 are probabilistic, the resulting ΔHSA score is also probabilistic. A metric such as

$$P(\Delta\text{HSA} > 0 \mid \mathcal{D}) \tag{4.8}$$

can then be defined to estimate how confident we are about the *synergistic efficacy* of the combination, based on the dataset $\mathcal{D}$. It is possible to have a synergistic combination that is highly uncertain, which would indicate that more data are required to reach confidence in the estimation.

We have opted for ΔHSA instead of the Bliss score or the Loewe score due to its minimal reliance on assumptions. As detailed in Chapter 3, the Bliss score assumes the MSP, while the Loewe score assumes the DEP, each containing its failure cases. In comparison, ΔHSA is a simple measure of whether the maximal efficacy of a combination has improved. Although it does not account for potency, this limitation becomes advantageous in this case, as it allows for the separate evaluation of efficacy and potency.

For the synergistic potency, α contains the required information. $\alpha > 1$ would indicate synergistic potency ([83]), which means the potency of the two drugs has reduced due to

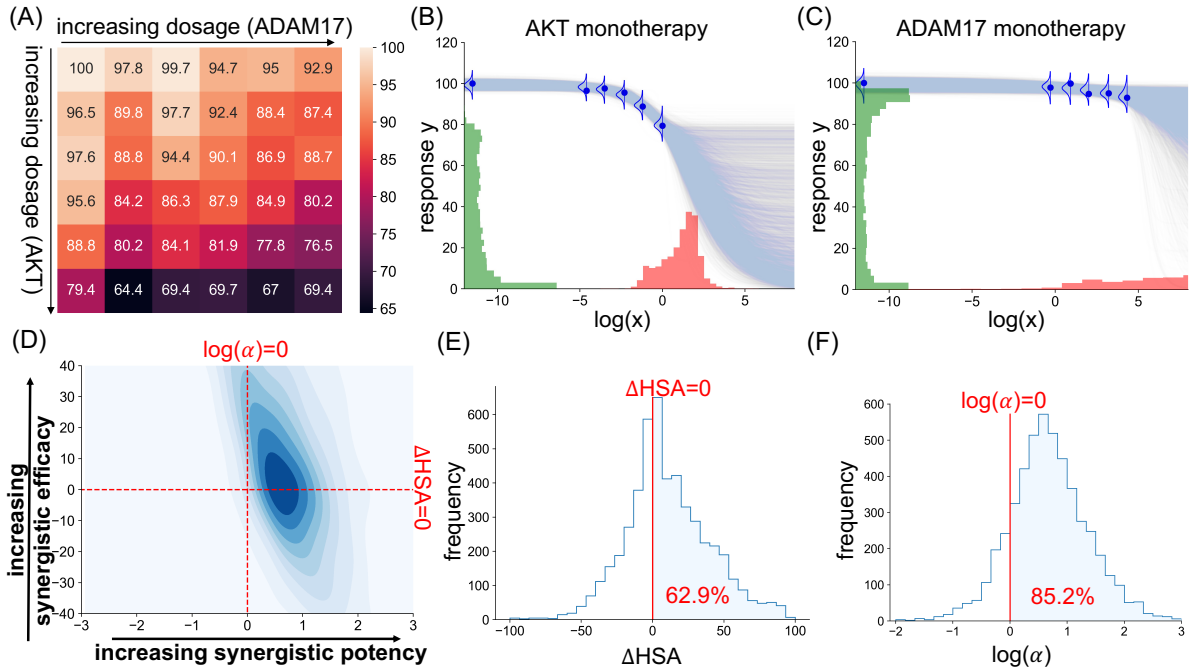

Figure 4.4: The data and the inference outputs of the combination of AKT and ADAM17 applied on the cell line BT-20. (A): The original dose-response matrix. (B)-(C): The monotherapy model outputs. Each blue curve is a sample from $\mathbb{E}[Y \mid \boldsymbol{\theta}]$ where $\boldsymbol{\theta}$ is a sample from the posterior of the respective monotherapy model. The posterior distribution for IC_{50} and E_{inf} are shown in red and green respectively. (D): The contour plot for the joint posterior distribution of the synergistic efficacy (ΔHSA) and the synergistic potency ($\log(\alpha)$). The distribution is smoothed from the empirical posterior with a kernel density estimation for visualisation purpose. (E)-(F): The histogram of the empirical posterior distribution for the synergistic efficacy (ΔHSA) and the synergistic potency ($\log(\alpha)$). The areas on the right-hand side of the red vertical lines are the probability that the combination is synergistic in terms of efficacy (in (E)) and potency (in (F)).

being combined. Consequently, we define

$$P(\alpha > 1 \mid \mathcal{D}) \tag{4.9}$$

to estimate how likely a combination satisfies *synergistic potency*.

4.1.3 Case studies

To illustrate how the uncertainty estimation from our method can be explained and further used for decision-making, we take two combinations from the DREAM dataset as examples.

Fig. 4.4 (A) shows the dose-response matrix for the combination of ADAM17 and AKT acted on the cell line BT-20. The first column is the monotherapy dose-response measurements for AKT (as the dosage for ADAM17 is zero), whilst the first row is the monotherapy measurements for ADAM17 (as the dosage for AKT is zero). It can be

observed that the responses for AKT start to decrease at a higher rate when the dosage increases, but the dosage range is too small to understand its potency (IC_{50}) and efficacy (E_{inf}). The efficacy cannot be determined because the response has not shown any sign of convergence, whereas the potency cannot be determined because it relies on understanding the maximal response, which is itself uncertain. However, if a deterministic Hill equation is fitted to the monotherapy data, it will only provide a point estimate for the IC_{50} and E_{inf}, without acknowledging the above caveats. On the contrary, our method provides an uncertainty estimation for both quantities. As shown in Fig. 4.4 (B)-(C), the posterior distribution of IC_{50} and E_{inf} have large variances, which correspond to large uncertainty. In particular, as shown in Fig. 4.4 (C), the posterior of E_{inf} for ADAM17 has a multimodal shape, which is sensible because it is unclear whether the dosage range is too small (which corresponds to the peak at 0), or the drug is simply ineffective regardless of the dosage (which corresponds to the peak at around 90).

Moving to the inference of the full combination matrix, most existing synergy methods have no means to showcase the uncertainty. Our method, on the contrary, provides the uncertainty around the synergistic potency and the synergistic efficacy, as shown in Fig. 4.4 (D), (E) and (F). According to the model output, the combination is moderately likely to be synergistically potent (with a probability of 85.2%), but it is difficult to conclude its synergistic efficacy (with a probability of 62.9% to be synergistic efficacious). This is reasonable because the excessively small dosage range makes it difficult to conclude anything about efficacy with high certainty, but with the 25 available measurements on the plates where the two drugs have interacted, information can be extracted on whether combining the two drugs may lower the level of toxicity required to reach the same beneficial effect.

This combination is an example where the model implies some potential in the synergy of the combination, but the level of uncertainty in the synergy is still high, which may require more measurements at larger dosages to be narrowed down.

We now consider the combination of AKT and EGFR acted on the cell line MDA-MB-468. Fig. 4.5 shows its dose-response matrix and the inference result from our model for the monotherapies and the combination respectively. All parameters and metrics of interest have low variances, representing low uncertainties. As shown in Fig. 4.5 (B) and (C), the dosages have suitable ranges and approximately follow the sigmoidal shapes of the expected dose-response fit, in particular for AKT. They contain sufficient information for the possibilities for IC_{50} and E_{inf} to be narrowed down. Similarly, the combination data are well-behaved. Fig. 4.5 (C), (D) and (E) show that the probabilities of this combination being synergistic in terms of potency and efficacy are both close to 100%. These are signs that this combination is worth being taken to subsequent steps in the drug development pipeline.

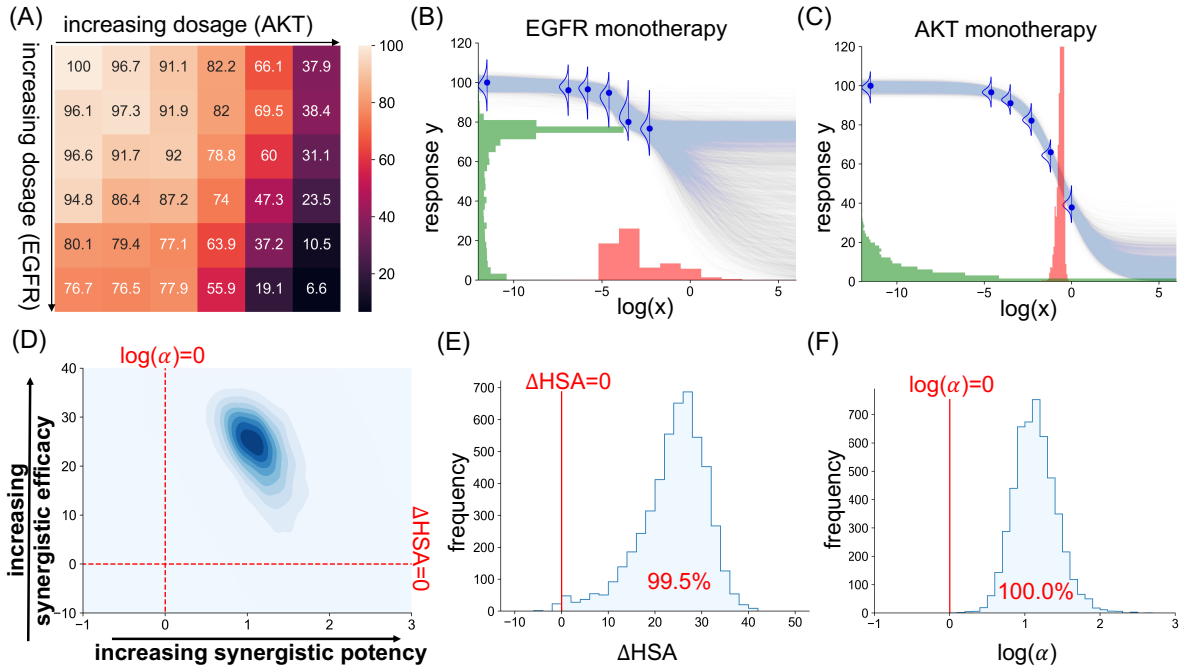

Figure 4.5: The data and the inference outputs of the combination of EGFR and AKT applied on the cell line MDA-MB-468. (A)-(F): The same as the caption of Fig. 4.4.

The two examples above show that concrete decisions can be made based on the posterior distributions (e.g. for IC_{50}, E_{inf}, ΔHSA and α) from our model, and more importantly, the uncertainties associated with these distributions.

4.1.4 Training details

The models are trained by Stan, a state-of-the-art platform for statistical modelling and high-performance statistical computation, particularly for Bayesian computation ([75]). The user specifies the prior model of the parameters and the likelihood model of the data, while Stan performs either full Bayesian statistical inference with MCMC sampling, or approximate Bayesian inference with variational inference. In this study we opt for MCMC due to the importance of the reliability of the output, which is ensured by the asymptotic exactness of the MCMC inference. Despite choosing the slower option of the two, SynBa is still computationally efficient. Running on 4 CPUs of the Intel Xeon Platinum 8276 CPU Processor, the median time taken to fit SynBa (via MCMC with 1000 iterations and 4 chains, including 500 iterations in the warm-up phase) to a 6-by-6 dose-response matrix in DREAM is 10.2 seconds, which is comparable to MuSyC with bootstrap.

With this training pipeline, we can avoid the overhead that occurs during the usage of non-linear optimisation packages in deterministic parametric methods such as MuSyC, BRAID and the Effective Dose model. A different choice of the numerical algorithm (and its hyperparameters) results in a different result for those methods. On the contrary, in

SynBa, the same exact result can be found asymptotically via MCMC with Stan.

For the implementation of MuSyC, BRAID and the Effective Dose model, the Python package `synergy` ([82]) is used. For the implementation of bayesynergy and Hand-GP, we make use of their official repository.

4.2 Results

4.2.1 Prediction of drug combination responses

In this subsection, we show that in addition to providing uncertainty estimations, SynBa is competitive in predicting unseen responses within a dose-response matrix, and is less prone to overfitting compared to the existing methods.

The datasets of interest are DREAM ([54]) and NCI-ALMANAC ([33]), two of the most widely-used publicly-available combination screenings. In DREAM, we focus on all examples in the training set of Challenge 1 that have passed the Quality Assurance and that only contain non-negative responses and one set of replicates, which are 1631 sets of combinations in total. In NCI-ALMANAC, we focus on the subset defined in [37], which is a subset of the data consisting of 50 unique FDA-approved drugs and 36,120 combinations. We remove examples that contain negative measurements, which results in 28,854 remaining combinations.

We leave out 20% of the non-zero dosage combinations for prediction. The models are trained using the remaining 80% dosage combinations, and then evaluated on the left-out points. For the dose-response matrices in DREAM, we leave out 7 of the 35 points with non-zero dosages from the 6-by-6 matrix (see Fig. 4.1(A)) using a specific leave-out strategy. For each of the two monotherapy slices, one point is left out for testing. For the 5-by-5 combination grid (i.e. the orange cells in Fig. 4.1(A)), five points are randomly left out for testing. Fig. 4.1(B) shows an example of such a train-test split. The measurements are left out in this manner so that each monotherapy contains one measurement for testing.

For the dose-response matrices in NCI-ALMANAC, it is not possible to leave out points separately for monotherapies and for combinations because the data size is too small. For most combinations, there are only 3 points for each monotherapy and 9 points for the interactions. Thus, we directly leave out 3 of the 15 points randomly for prediction.

To evaluate the predictive performance of the models, the root-mean-square error (RMSE) is computed using the left-out points. For each combination, its RMSE for the test responses $\{y_1, ..., y_N\}$ is

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{N}} \tag{4.10}$$

where $\hat{y}_i$ is the point estimate for the response that corresponds to dosage $\mathbf{x}_i$. For SynBa, we define $\hat{y}_i$ to be the posterior predictive mean $\mathbb{E}[y_i \mid \mathcal{D}, \mathbf{x}_i]$, which is estimated by Monte Carlo sampling from the posterior.

We compare our prediction results against MuSyC, BRAID, and the Effective Dose model, which are three of the most widely-used synergy models. We also compare against bayesynergy and Hand-GP, the two probabilistic synergy models. For SynBa, we implement both the uniform prior and the empirical Beta prior for the normalised $\mathrm{E}_{\mathrm{inf}}$, which we denote as SynBa-U and SynBa-B respectively. Table 4.1 show the mean and the median of the test RMSE for all six methods. Our method outperforms all other methods in terms of RMSE for both datasets. This is especially encouraging because RMSE is a metric that only considers the quality of point estimates and ignores uncertainty. The upper diagonal panels in Fig. 4.6 show the scatter plots directly comparing the test RMSE values between methods (visualised with the blue colour). Our method is the most competitive, as evidenced by having more points above the diagonal $y = x$ line. These show that our method is not trading off predictive accuracy for uncertainty estimation. By following a principled Bayesian workflow, our model is strong in both prediction and uncertainty estimation.

It is worth noting that at least one of MuSyC, BRAID or the Effective Dose model fail to find a solution for 4.8% of the examples in DREAM and 38.4% of them in NCI-ALMANAC, despite an effort in tuning the bounds, initial values and hyperparameters involved in the optimisation. Most likely this is because these methods rely on external optimisation packages with no guaranteed convergence, which can become a problem when overparameterisation becomes severe due to small data sizes. SynBa does not incur this problem since its priors ensure conservative outputs when data size is too small.

To investigate whether SynBa is prone to overfitting and how it compares to the other three methods, we perform the same prediction experiment on DREAM, but with a train-test split ratio of 40% : 60% instead. As shown in the lower diagonal panels in Fig. 4.6, the test RMSE values (visualised with the red colour) increase significantly for the three deterministic models. For SynBa, however, the test RMSE values have increased on average, but not by much. It can be seen that the mean value and the spread increase more significantly for the other three methods compared to SynBa. It can also be observed that the predictive performance of SynBa is not sensitive to the choice of prior, with the two priors producing very close RMSE values to each other.

4.2.2 Uncertainty calibration

For a model $\mathcal{M}$ with learnt cumulative distribution $F_{\mathcal{M}}$ with well-calibrated uncertainty, it would approximately follow the identity that $F_Y(\mathbf{x}_i) \approx F_{\mathcal{M}}(\mathbf{x}_i)$ for every data point $\{(\mathbf{x}_i, y_i(\mathbf{x}_i)) \mid i = 1, ..., N\}$ in the dataset, where $y_i(\mathbf{x}_i)$ is a sample from the unknown

Table 4.1: The mean and the median of the test root-mean-squared error (RMSE) for MuSyC, BRAID, the Effective Dose model, bayesynergy, Hand-GP and SynBa, computed on a subset of DREAM and NCI-ALMANAC, along with their standard errors. The standard error of the mean is computed by the standard deviation of the metrics across examples divided by the square root of the number of examples. The standard error of the median is estimated by nonparametric bootstrap. SynBa with a uniform prior for the normalised E_{inf} is denoted by SynBa-U. SynBa with the Beta$(0.46, 0.58)$ prior for the normalised E_{inf} is denoted by SynBa-B.

	DREAM, RMSE		NCI-ALMANAC, RMSE	
	mean (± se)	median (± se)	mean (± se)	median (± se)
MuSyC	6.11 ± 0.10	5.01 ± 0.08	14.87 ± 0.10	10.17 ± 0.03
BRAID	5.71 ± 0.09	4.88 ± 0.08	9.57 ± 0.06	7.16 ± 0.03
Effective Dose	6.46 ± 0.09	5.66 ± 0.09	8.47 ± 0.05	6.77 ± 0.09
bayesynergy	8.10 ± 0.03	8.07 ± 0.04	24.18 ± 0.12	17.92 ± 0.20
Hand-GP	10.02 ± 0.16	8.35 ± 0.12	22.91 ± 0.31	14.41 ± 0.37
SynBa-U	5.20 ± 0.07	4.56 ± 0.08	6.72 ± 0.04	5.45 ± 0.05
SynBa-B	$\mathbf{5.15} \pm 0.07$	$\mathbf{4.55} \pm 0.08$	$\mathbf{6.66} \pm 0.04$	$\mathbf{5.43} \pm 0.04$

true cumulative distribution $F_Y(\mathbf{x}_i)$. Equivalently, assuming the measurements $y_i(\mathbf{x}_i)$ for a combination are conditionally independent given the dosages $\mathbf{x}_i$, their cumulative probabilities $F(y_i) := \mathbb{P}(y_i(\mathbf{x}_i) < F_{\mathcal{M}}(\mathbf{x}_i))$ would be approximately uniformly distributed between 0 and 1, if $\mathcal{M}$ is well-calibrated.

In this study, for each combination in DREAM, we split the 35 measurements (excluding the base value) with a 80%:20% ratio in the same way as the prediction evaluation in the previous subsection. We then evaluate the quality of the uncertainty calibration with the Kolmogorov–Smirnov (K-S) uniformity test ([52]) for the empirical cumulative probabilities (or CDF values) across all test data points. If the model is well-calibrated, then the CDF values for the test data points will be approximately uniform for each combination. Otherwise, they will show a non-uniform pattern and the resulting p-value for the K-S test will be statistically significant. This procedure is performed across every combination in DREAM, resulting in a p-value for each combination. The histogram of the p-values for SynBa is illustrated in Fig. 4.7(A), showing that 6.07% of the combinations have not passed the uniformity test, and thus are not well-calibrated. As a comparison, the same procedure is performed using MuSyC. Fig. 4.7(B) shows that 25.1% of the combinations are not well-calibrated when modelled by MuSyC, which is roughly four times as high as the number for SynBa. This shows the estimated uncertainty from SynBa is more reliable and closer to the unknown ground truth on average.

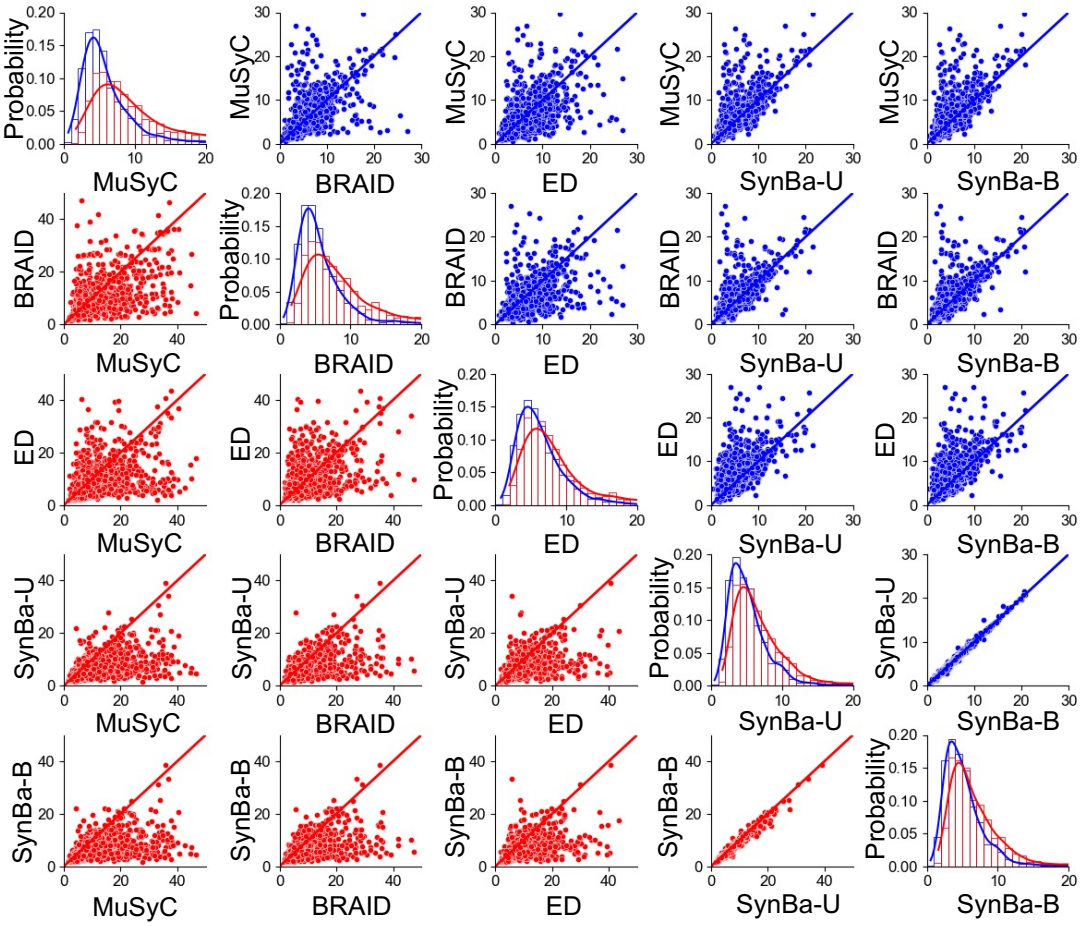

Figure 4.6: Upper diagonal panels (with blue points): Scatter plot of the test RMSE values obtained from different methods on DREAM with a train-test split ratio of 80% : 20%. Lower diagonal panels (with red points): Scatter plot of the test RMSE values obtained from different methods on DREAM with a train-test split ratio of 40% : 60%. Diagonal panels: Histograms of the test RMSE values and their corresponding kernel density estimates, where the train-test split is 4:1 for the blue ones and 2:3 for the red ones.

4.3 Discussions and Conclusions

Machine learning methods have been developed for preclinical modelling and the prediction of drug combinations, thanks to the availability of large screenings ([37, 80]), which are beneficial for discovering and explaining drug combinations. However, a few factors prevent most from being applied to real-world drug discovery projects. One issue is that performance measures rely on synergy scores, which do not have a gold standard and contain a non-trivial amount of uncertainty, as discussed in Introduction. The Spearman correlation of the replicate experiments in DREAM ([54]) and [59] are 0.56 and 0.63 respectively, which show that quantifying a combination with a single synergy score would result in high variance. However, uncertainty measurement is not included in synergy score estimation. This could be one of the reasons that 20% of drug combinations are poorly predicted by all methods in the DREAM challenge. Measuring the uncertainties

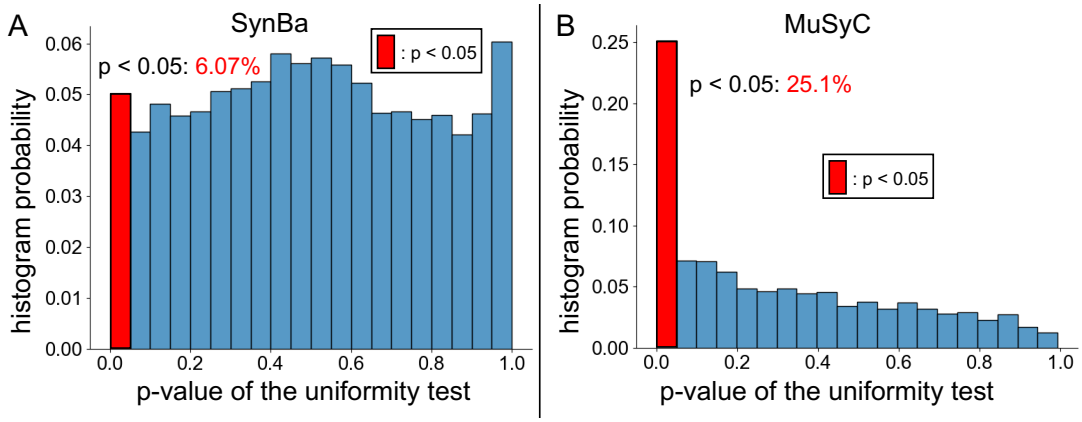

Figure 4.7: Histogram of p-values that represent the calibration quality of combination models. The p-values are derived from the Kolmogorov–Smirnov test between the uniform distribu- tion and the cumulative probability of the data points in their predicted densities. For a well-calibrated model, the p-value will be higher than 0.05.

associated with the estimated scores is important for the subsequent decision-making process based on the model outputs. In real-world scenarios, scientists are often facing the decision to choose amongst a large set of drug combinations that score similarly in terms of synergy. Without quantification of how certain (or uncertain) the estimated scores are, they will have to rely on background knowledge compromising innovation in their choices. SynBa provides a way to implement a ranking strategy in the decision process of a drug discovery pipeline, which is a real-world unmet need. This is illustrated in Fig. 4.8, where the ideal drug candidate will have either a strong synergistic efficacy or a strong synergistic potency (or even better, both) with low associated uncertainties. Drugs in the first quadrant with a small "radius" would represent promising candidates, ranking higher in the decision-making process due to their high potential and low uncertainty. This is more robust than directly ranking deterministic synergy scores such as Bliss or Loewe.

Although SynBa only models the combination of two drugs, while there exist methods such as the Effective Dose Model that consider higher-order combinations, there is nothing in methodology stopping us from extending SynBa to three or even more drugs. However, this would significantly increase the computational cost of evaluating posteriors, which grows exponentially with the number of drugs. It would also be more challenging to ensure that the MCMC algorithms accurately estimate the true posteriors, since a much larger number of posterior samples would be required. While this limitation exists, our primary goal is to provide a method that remains reliable and avoids overparameterisation, aligning with practical needs. Extending to three or more drugs is beyond the scope of immediate real-world application.

In conclusion, we have developed a new framework for quantifying dose-response relationships for monotherapies and combinations that provides a full uncertainty estimation for all parameters that are associated with the monotherapies and the combinations,

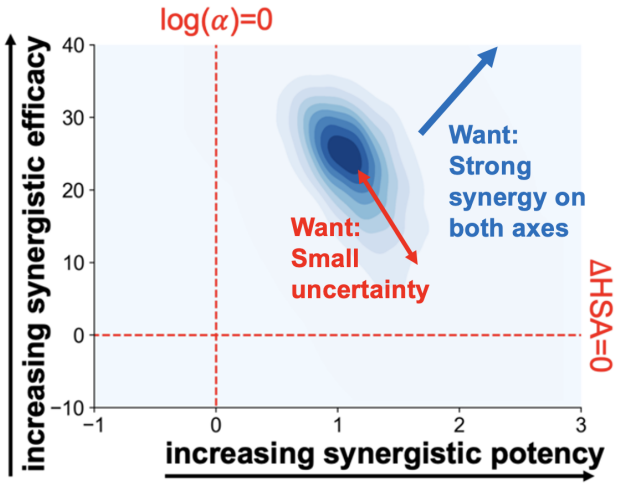

Figure 4.8: Example of a promising combination candidate, with posterior distributions for synergistic efficacy and potency lying in the first quadrant with a relatively small uncertainty.

including information about efficacy, potency and synergy. These uncertainty information would be helpful to the biologists to make further decisions about progressing to the next stages of the drug discovery pipeline, or whether more experiments are required to lower the level of the uncertainty and better understand the drug mechanism of action.

We have also shown that SynBa is competitive in predicting unseen responses within a given dose-response matrix, and outperforms MuSyC, BRAID and the Effective Dose model on DREAM and NCI-ALMANAC. In addition, the prediction performance is not sensitive to the choice of the priors.

In summary, our framework is capable of providing a reliable uncertainty estimation for the potency (e.g. IC_{50}) and the efficacy (e.g. E_{inf}) of a monotherapy, or the synergistic potency and efficacy of a combination, in a decoupled manner, and reliably predicting unseen responses within a dose-response matrix. The parameter uncertainties can be interpreted and used as guidance for further experiments and subsequent decision-making.

Chapter 5

SynBa-Batch: Modelling the batch effect and utilising information

While SynBa provides a robust method for quantifying the decoupled drug synergy as well as estimating the uncertainties, one problem it does not directly address is the variability on the effects of the drugs that stems from the difference in experimental environment and the biological variations, such as the difference in the types of cell culture. For example, the differences in cell numbers seeded in the dish can impact the activity of cells and thus the effect of the drug. Even for the same monotherapy or the same combination (i.e. the same drug or drugs administered to the same cell line), the efficacy and potency derived from different sets of measurements can be significantly different. In such cases, is it still reasonable to assume that a single parametric mean adequately represents all the data like in SynBa? To answer this question, we first need to understand the nature of the batch effects that are exhibited in drug screenings.

5.1 Batch effect

In drug screenings, it is often observed that even for the same monotherapy or combination, drug responses and cell viability often vary significantly across experiments designed to measure the same quantity. These variations can be attributed to biological variances, differences in laboratory environments, and experimental noise ([47, 14, 74]). Even for individual wells on the same plate, spatial bias in microwell assays has been observed among individual wells ([44, 53]).

Fig. 5.1 is an illustrative example that shows the significant variability in dose-response measurements across different plates, indicating a strong noise component that reflects potential batch effects. In this figure, each subfigure contains a set of monotherapies (of the same drug) applied to the cell line 786-0. In each subfigure, it can be observed that the response measurements differ substantially, despite the same drug being applied to

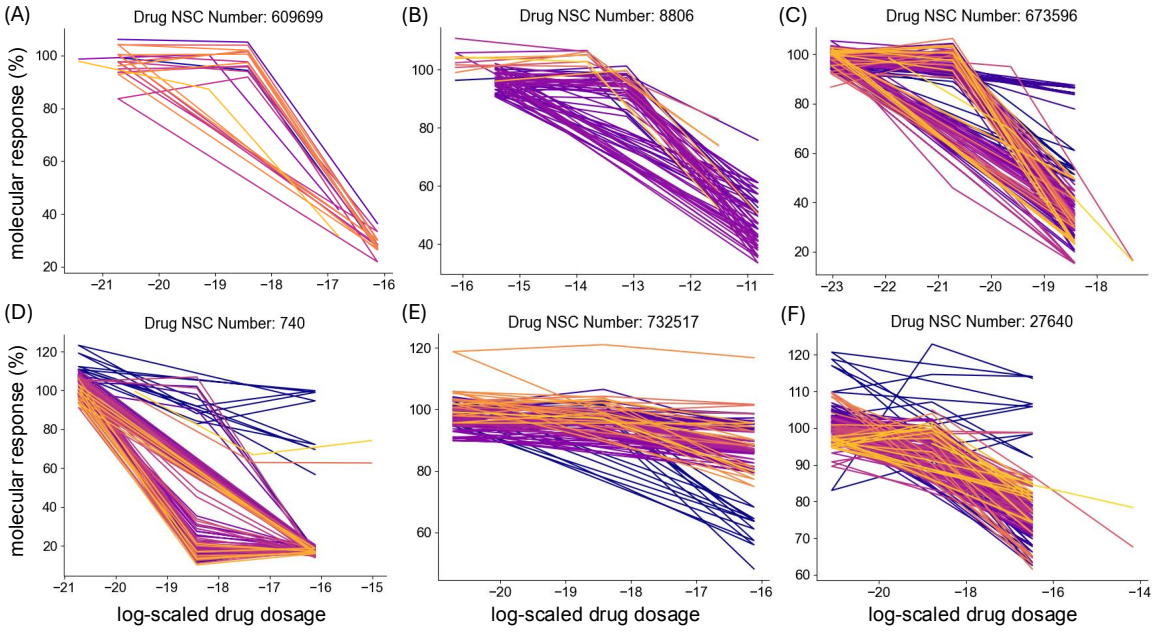

Figure 5.1: Dose-response curves for different drugs applied to the cell line 786-0, coloured by test dates. Each curve represents a set of measurements from the same plate. The x-axis represents log-transformed drug dosages and the y-axis shows the corresponding drug response.

the same cell line. While the measurements in subfigures (A) to (C) show somewhat similar downward pattern, it is not the case for subfigures (D) to (F). In each of the subfigures (D) to (F), it can be observed that some curves are mostly flat around 100, showing little efficacy, whereas some other curves show a steep decline, representing a strong efficacy. The steep declines also occur at different dosages, showing a discrepancy in the potency across plates. From these sets of measurements, it is evident that the dose-response measurements cannot be explained by a single parameterisation of the form $y = f(x) + \epsilon$ where ϵ represents i.i.d. noise.

Despite this, the majority of the current computational models typically aggregate all data into a single framework, neglecting these batch effects. On the experimental side, the outlier measurements are often removed as an empirical approach. These approaches are insufficient as they overlook the inherent variability in the data. Recently, there have been works on dealing with batch effects in the context of monotherapies ([74]). Instead of treating batch effects as a nuisance, the temporal dependencies are utilised to improve the estimation of negative controls, shrinking the differences between experiments conducted on similar days through an empirical Bayes procedure.

In this study, we approach the batch effects by identifying batches that have different dynamics, since each of them reflects different experimental conditions or biological contexts. We then leverage the learned dynamics to enhance the modelling of combination therapies.

No work to date has thoroughly explored this approach in the context of drug combination studies. This study addresses this gap by providing an informed guess for the dynamics of the dose-response matrices, with the aim of enhancing the accuracy of inferring efficacy, potency, and synergy in drug combinations. The informed guess is based on the dynamics learnt from the monotherapy data available for each of the two drugs being combined.

In addition to providing an improved model that distinguishes between batches and learns the different dynamics, we also offer a method for data fusion to utilise existing data and improve predictions for combination responses. Existing methods have aimed to remove or correct variations [48]. In contrast, SynBa-Batch systematically addresses variability and batch effects without discarding data or assuming the presence of a single ground truth distribution for the pharmacological parameters of a monotherapy or a combination. By classifying each combination dose-response matrix into one of the learned batches, the batch effect will be explained away from the modelling if a suitable batch is selected. By doing this, information from diverse experimental conditions and datasets is integrated, hence improving the prediction performance of drug combination viability.

In summary, we improve the existing pharmacology quantification frameworks, in particular SynBa, by (1) modelling batch effects to explain away the uncertainty associated with the batch effects and (2) informing combination models by utilising prior distributions from monotherapy models. We validate our approach through its high correlation with biological metadata, and demonstrate improved predictive performance, offering a robust tool for the inference and prediction of dose-response pharmacology data.

5.2 Methods

In SynBa-Batch, we learn the batch effect by partitioning the data into multiple components and parameterising each dynamic separately using a mixture model framework. We construct our mixture model based on SynBa due to its Bayesian nature, which naturally accommodates probabilistic mixture modelling through latent variable frameworks. This approach performs the quantification of uncertainty within each batch and between batches, providing not merely a clustered model but also the identification of the more significant or informative batches, without ignoring the less significant ones. It allows the reduction of the epistemic uncertainty (i.e. the uncertainty that arises from limited knowledge or data) in the modelling process by explaining away the batch effect through the mixture components. The aleatoric uncertainty (i.e. the inherent uncertainty due to noise) within each batch is better estimated. Consequently, the accuracy of uncertainty quantification is improved for the parameters and the drug response predictions.

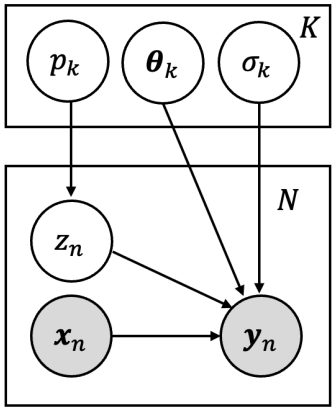

Figure 5.2: The graphical model depicting SynBa-Batch. The shaded shapes indicate known values. The unfilled round shapes indicate unknown random variables. The rectangular shape indicates fixed parameters to be inferred.

5.2.1 SynBa-Batch for monotherapies

To incorporate the batch effect and model measurements of the same monotherapy that potentially have different underlying distributions, we incorporate a probabilistic mixture approach based on SynBa. We begin with a dataset $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n) \in (\mathbb{R}^3_{>0}, \mathbb{R}^3_{\geq 0}) \mid n = 1, \ldots, N\}$, where each $(\mathbf{x}_n, \mathbf{y}_n)$ is a set of experiments from the same experiment setting (i.e. performed simultaneously on the n-th plate, usually labelled with the same experiment ID or plate ID in the metadata). The vector $\mathbf{x}_n$ denotes the set of dosages of the drug applied, whereas $\mathbf{y}_n$ denotes the corresponding set of observed responses.

Under our proposed model the drug responses for a monotherapy are assumed to be generated from a mixture of K components, each characterized by parameters $\boldsymbol{\theta} = (E_0, E_1, C, H)$. We introduce a latent variable $\boldsymbol{z}$ that indicates the component from which the measurements are drawn. It follows a categorical distribution $p(z = k) \sim \text{Cat}(p_1, \ldots, p_K)$, where p_k is the mixture proportion for component k. For each set of measurements $(\mathbf{x}_n, \mathbf{y}_n)$ from the n-th plate, it is assumed to be drawn from one of the K distributions depending on the value of the latent assignment z_n. The likelihood model for the response y given its batch assignment k and the dosage x is defined as

$$p(y \mid z = k, \boldsymbol{\theta}, \sigma, x) = \mathcal{N}(f(\boldsymbol{\theta}, x), \sigma^2) \tag{5.1}$$

independently for all $(x, y) \in (\mathbf{x}_n, \mathbf{y}_n)$, where $E_0 := E(0)$, $E_1 := E(\infty)$ and σ is the standard deviation of the noise level of y.

The joint prior distribution for a set of drug responses in a single study or plate, the

batch they belong to, and the parameters of the curve given the dosages is

$$\begin{aligned} p(\mathbf{y}, z = k, \boldsymbol{\theta}, \sigma \mid \mathbf{x}) &= p(\mathbf{y} \mid z = k, \boldsymbol{\theta}, \sigma, \mathbf{x}) p(z = k) p(\boldsymbol{\theta}) p(\sigma) \\ &= p(z = k) p(\boldsymbol{\theta}) p(\sigma) \prod_{x \in \mathbf{x}} \mathcal{N}(f(\boldsymbol{\theta}, x), \sigma^2) \ . \end{aligned} \tag{5.2}$$

The parameters of the model are learned via the Expectation/ Conditional Maximization (ECM) algorithm ([55]). The ECM algorithm is used instead of the classic EM because the complex parameterisation of the likelihood makes the M-step intractable, whereas ECM simplifies this into easier conditional maximisations.

The E-step involves evaluating the posterior probability $p(z = k \mid \mathbf{y}, \mathbf{x})$ for each set of monotherapy measurements. This is done using the relationship $p(z = k \mid \mathbf{y}) \propto p(\mathbf{y} \mid z = k)p(z = k)$, where the normalization is performed after computing these probabilities for all k. In the CM-step, we maximize the mixture proportions $p_k = p(z = k)$ with respect to the complete log-likelihood while keeping β_k and σ_k fixed. Subsequently, we maximize β_k and σ_k with the mixture proportions p_k fixed. This iterative process allows for the refinement of the model parameters to best fit the observed data.

Finally, for each mixture, the weights $w_i := p(z_i = k \mid \mathbf{y}_i, \mathbf{x}_i)$ are used to train a weighted SynBa model for each mixture component. The overall objective function $L = \sum_{i=1}^{N} w_i p(\mathbf{y}_i \mid \Theta, \mathbf{x}_i)$ is maximized, where Θ represents the SynBa parameters.

This modelling framework provides a robust means of capturing multiple modes from the underlying distribution of the monotherapy measurements. The ECM algorithm ensures that the parameter estimation process is efficient and convergent, balancing between statistical guarantee and computational efficiency.

We now write down SynBa-Batch for monotherapy inference in full as Algorithm 5.1.

Algorithm 5.1: **SynBa-Batch** for monotherapy inference.

Step 1: Initialisation

- Define $p_k := p(z = k) = \frac{1}{K}$ for all $k = 1, ..., K$.
- For each plate n, compute the optimal deterministic parameters $\Theta_n = (C_n, E_{0n}, E_{1n}, h_n)$ by fitting to the 1D Hill equation.
- Initialise $\boldsymbol{\theta}_k$ by training a standard Gaussian Mixture Model on $\Theta = \{\Theta_n \mid n = 1, ..., N\}$ and taking the means $\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_K$.
- Initialise $\sigma_k = 1$ for all k.

Step 2: Iterate using the Expectation/Conditional Maximization algorithm ([55])

- E step: Evaluate $p(z_n = k \mid Y_n)$ for plate $n \in \{1, ..., N\}$. We have

$$\begin{aligned} p(z_n = k \mid Y_n) &\propto p(Y_n \mid z_n = k)p(z_n = k) \\ &= \prod_{i=1}^{L_n} p(y_{n,i} \mid z_n = k)p(z_n = k) \\ &= \prod_{i=1}^{L_n} \mathcal{N}(y_{n,i} \mid f(\boldsymbol{\theta}_k, x), \sigma_k^2)p(z_n = k) \end{aligned} \tag{5.3}$$

- CM step: Iteratively maximise p_k, $\boldsymbol{\theta}_k$ and σ_k w.r.t. the complete data log-likelihood

$$Q(\Theta|\Theta_t) = \mathbb{E}[\log p(Y \mid \Theta) \mid Y, \Theta_t]$$

for all $k = 1, ..., K$. Empirically, this is estimated by $\sum_{n=1}^{N} \sum_{k=1}^{K} p(z_n = k \mid Y_n) \log p_k p(Y_n \mid z_n = k, \boldsymbol{\theta}_k, \sigma_k)$ ([77]).

Compute $p_k \leftarrow \frac{1}{N} \sum_{n=1}^{N} p(z_n = k \mid Y_n)$.

Optimise $\boldsymbol{\theta}_k, \sigma_k \leftarrow \arg\max_{\boldsymbol{\theta}_k, \sigma_k} \sum_{n=1}^{N} \sum_{k=1}^{K} p(z_n = k \mid Y_n) \log p_k p(Y_n \mid z_n = k, \boldsymbol{\theta}_k, \sigma_k)$.

Step 3: Train the weighted SynBa separately for each mixture

- For each mixture $k \in \{0, ..., K\}$ and each plate n, define the weight for each set of measurements as

$$w_{k,n} := p(z_n = k \mid \mathbf{y}_n, \mathbf{x}_n), \tag{5.4}$$

which is the posterior probability of this set belonging to the k-th batch.

- For each mixture $k \in \{0, ..., K\}$, train the weighted SynBa, where the overall objective function to be maximised is modified to be

$$L_k = \sum_{n=1}^{N} w_{k,n} \log p(\mathbf{y}_n \mid \Theta_k, \mathbf{x}_n), \tag{5.5}$$

where Θ_k represents the SynBa parameters for mixture k.

Once all steps in Algorithm 5.1 are followed, we end up with k separately trained SynBa models, one for each mixture. These models now account for distinct batch effects, explaining away the variability that arises from different biological or experimental conditions. Each SynBa model represents a specific monotherapy profile, reflecting the unique dynamics of its batch. This provides biologists with a clearer and more interpretable understanding of the drug response across various contexts, facilitating more informed

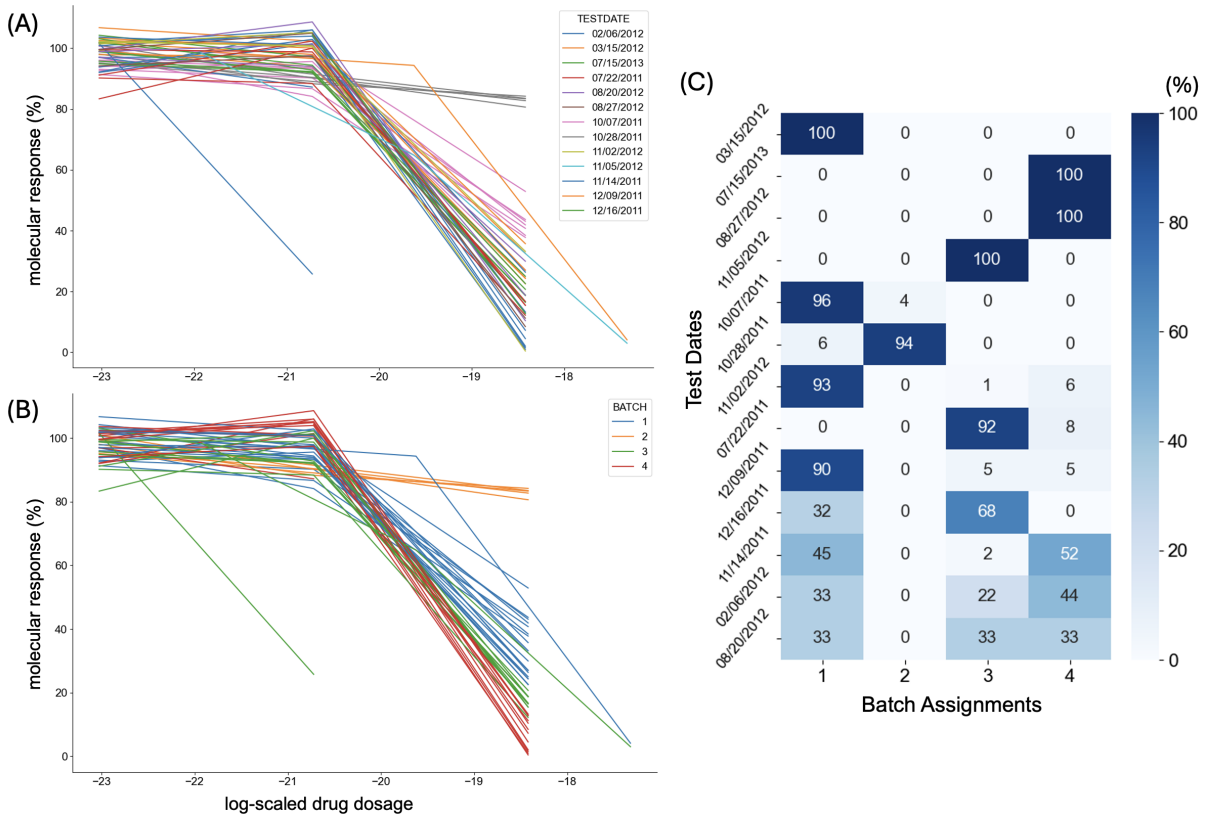

Figure 5.3: Visualisation of monotherapy measurements for the drug 7-Ethyl-10-hydroxycamptothecin applied to the cell line 786-0. (A): Coloured by test dates. (B): Coloured by the batches learnt by SynBa-Batch. (C): The confusion matrix showing the percentage of measurements assigned to each batch for each test date.

experimental decisions based on specific experimental settings. We now go through a case study to illustrate the usefulness of SynBa-Batch in monotherapy modelling.

5.2.1.1 Case study

We illustrate the benefits of SynBa-Batch using an example from the NCI-ALMANAC dataset [33]. Fig. 5.3(A) illustrates all monotherapy measurements for the drug 7-Ethyl-10-hydroxycamptothecin applied to the cell line 786-0, coloured by the different test dates. The percent growth inhibition is plotted against the logarithm of the applied dosages, where each line represents measurements from the same plate. The screener of the experiments is illustrated by the line style. It can be observed that measurements performed on the same date from the same screener tend to have a similar pattern as dosage increases. On the contrary, measurements from different dates vary significantly even when the same dosage of the same drug is applied on a fixed cell line. In this particular example, measurements from the date 10/28/2011 (coloured in grey) are particularly distinct from the others, showing a much weaker efficacy. The mutual existence of strong similarity of measurements from the same environment and strong variability of measurements

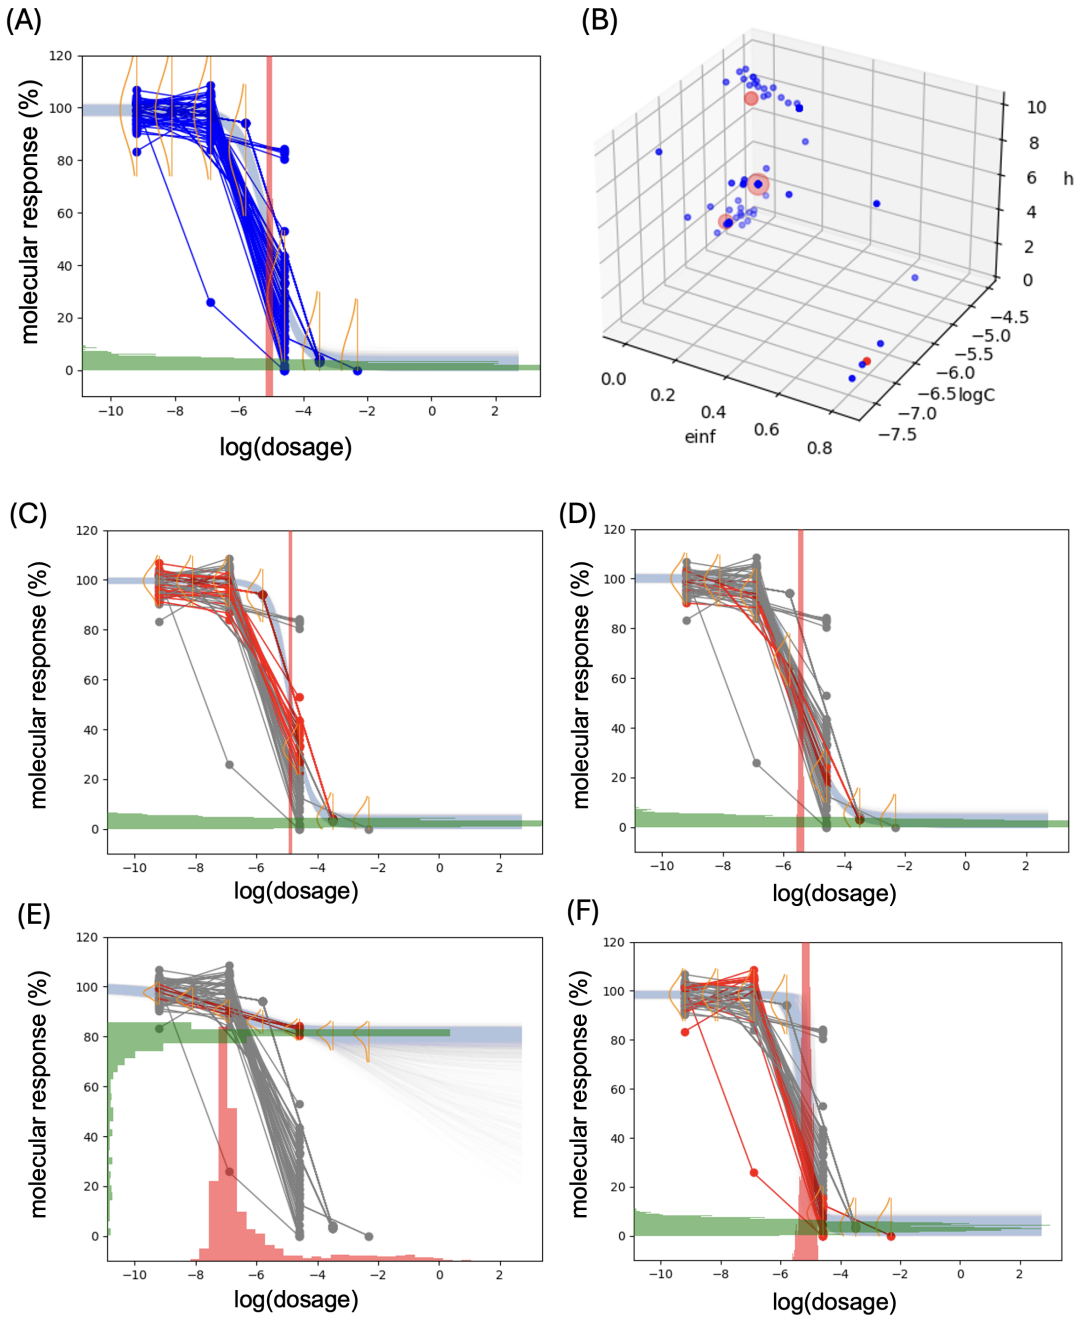

Figure 5.4: Comparison of SynBa and SynBa-Batch outputs for 7-Ethyl-10-hydroxycamptothecin administered on the 786-0 cell line from the NCI-ALMANAC dataset. (A): Output from the original SynBa. (B): The E_{inf}, $\log C$ and H across multiple sets of monotherapy measurements. Each blue dot is a set of (E_{inf}, $\log C$ and H) learnt from a single set of monotherapy measurements. Each red dot illustrates the mean of a mixture learnt from the GMM trained using the blue dots as the data. (C)-(F): Outputs from the four learnt batches of SynBa-Batch. Measurements are coloured in red if their probability of being assigned to this batch is larger than the others. The distribution for IC_{50} is shown in red, whereas the distribution for E_{inf} is shown in green.

from different environments suggests the presence of batch effects. This variability shows the importance of incorporating batch-specific factors into the modelling framework to accurately capture the underlying biological response and improve predictive performance.

Fig. 5.3(B) shows the same measurements as (A), but now coloured by the batch assignments learnt by SynBa-Batch as described in Algorithm 5.1. As the model is probabilistic, the assignments are inferred according to

$$\tilde{z}_n = \arg\max_k p(z_n = k|\mathcal{D}) \tag{5.6}$$

for each set of experiments $(\mathbf{x}_n, \mathbf{y}_n)$ in the same plate.

From Fig. 5.3(A) and (B), a high correlation is observed between the cluster assignments from SynBa-Batch and the test dates, indicating that despite the inherent batch effects being unknown, our model successfully captures significant patterns. In particular, the measurements with the test date 10/28/2011 are successfully identified as a separate batch (labelled as Batch 2). Furthermore, Fig. 5.3(C) computes the confusion matrix showing the percentage of measurements assigned to each batch for each test date. Of the 13 test dates, 9 of them have at least 90% of the measurements clustered to the same batch, showing the correlation between the test dates and the batch assignments, which have been inferred without knowing the test dates. Although the ground truth for batch effects is unknown and thus cannot be directly tested, the close alignment between our batch assignments and the test dates is still significant evidence that suggests our approach provides valuable insights into the underlying variability, making the results informative and helpful for further biological investigations.

Fig. 5.4(A) illustrates the posterior distributions for the predicted responses and the pharmacology parameters learnt from the original SynBa model (described in Chapter 4). A single trend is learnt and averages out all scenarios. This averaging leads to a posterior distribution that overlooks significant differences between batches, reducing the model's capacity to accurately summarise the data. For example, the posteriors for the predicted responses and the maximal efficacy E_1 give negligible weight to the measurements from test date 10/28/2011 (identified as Batch 2 in Fig. 5.3(B)). This is because the drug responses from that specific date differ significantly from those of other dates. As shown in Fig. 5.4(B), the pharmacological parameters of the sets of measurements obtained from the 1D Hill Equation vary vastly from each other. However, SynBa's single-mode likelihood model restricts the ability to simultaneously learn multiple plausible scenarios, preventing this specific mode from being captured. Such outlier modes should be summarised by the model instead of being discarded as an abnormality, since an outlier is not necessarily a measurement error, but could stem from legitimate factors including biological variation and laboratory environment.

In contrast, as shown in Fig. 5.4(C)-(F), various scenarios are successfully separated

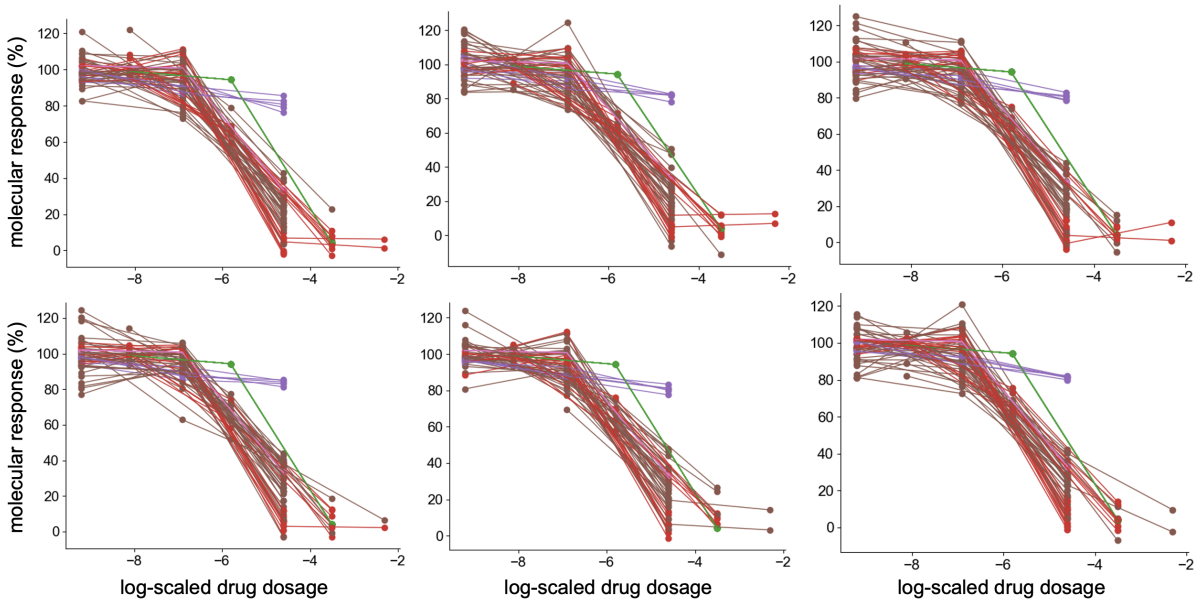

Figure 5.5: Synthetic drug responses generated from the learnt SynBa-Batch model. Points with the same colour are generated from the same batch. The dosages in each subfigure are set to be the same as the measurements in the training set.

by SynBa-Batch. The aforementioned outlier batch from the test date 10/28/2011 is illustrated in Fig. 5.4(E), where the posterior distribution for the maximal efficacy is significantly different from the others. The major mode is also captured by SynBa-Batch, as shown in Fig. 5.4(C)(D)(F). Although these three mixtures have similar posteriors for maximal efficacy, their posteriors for IC_{50} are less similar, either showing different means or a different level of variance, meaning these three mixtures are not repetitive.

Fig. 5.5 shows the synthetic drug response levels generated from the learned SynBa-Batch model. This is obtained by first drawing $z = k$ from the posterior categorical distribution $p(z = k)$ for the latent batch assignment, and then drawing

$$y \mid z = k, x \sim \mathcal{N}(f(\boldsymbol{\theta}_k, x), \sigma_k^2) \tag{5.7}$$

for a specified dosage x. The dosages are set to be the same as those in the original dataset so that the synthetic responses can be directly compared with the measurements. This procedure is repeated six times, resulting in six subfigures. Across these subfigures, the generated responses closely resemble the real data, demonstrating similar patterns and variability, which validates the effectiveness of the SynBa-Batch model in capturing the underlying distribution of the original measurements. This similarity indicates that SynBa-Batch can accurately replicate and model the complex dynamics observed in real-world pharmacology screenings.

5.2.2 SynBa and SynBa-Batch: Utilising existing models or information

Now that we have introduced SynBa-Batch for modelling batch effects of monotherapies, we think about how learning the batch effects could aid the inference of combination measurements. More generally speaking, can we leverage the rich information available from monotherapy models to better infer combination outcomes? In Bayesian methods, this can be achieved by incorporating the posterior distributions from the monotherapy models into the priors for the combination model. By doing so, we not only inform the combination model with meaningful prior knowledge but also potentially improve its prediction accuracy on untested dosages. Algorithms 5.2 and 5.3 outline the detailed method for incorporating both SynBa and SynBa-Batch models in combination prediction. The former is denoted as **SynBa-Informed** to differentiate with the original SynBa.

Algorithm 5.2: **SynBa-Informed** for combinations with existing monotherapy information.

Step 1: Train SynBa on each of the two monotherapies including historical measurements. For Drug 1, estimate the posterior distribution for E_1 and C_1. For Drug 2, estimate the posterior distribution for E_2 and C_2. These contain the efficacy and potency information for the two drugs.

Step 2: Define the prior for E_1, E_2, C_1 and C_2 in the combination model using the posteriors for the parameters from the k-th batch. For h_1 and h_2, the priors are set to be lognormal$(0, 1)$ as in the original SynBa. The prior for E_3 is set to be $p(E_1E_2)$, assuming Bliss independence.

Step 3: Train the combination model as in Box 2 of [86], also outlined in Section 4.1.2.

Algorithm 5.3: **SynBa-Batch** for combinations with existing monotherapy information.

Step 1: Train SynBa-Batch on each of the two monotherapies including historical measurements. These contain the efficacy and potency information for the two drugs.

Step 2: For each of the two drugs, for each plate n:

Extract the *batch assignment probabilities* $p(z_n = k \mid \mathcal{D})$ from the outputs of SynBa-Batch.

Step 3: Pick the *batch assignments* k for each of the two monotherapy slices $(\mathbf{x}_1, \mathbf{x}_2)$

associated with the combination matrix (i.e. the first row $\mathbf{x}_1$ and the first column $\mathbf{x}_2$ of the matrix). For each of the two drugs, pick the batch with the largest *batch assignment probabilities* from the previous step:

$$k_1 := \arg\max_k p(z_1 = k \mid \mathcal{D}), \qquad k_2 := \arg\max_k p(z_2 = k \mid \mathcal{D}).$$

Step 4: Define the prior for E_1, E_2, C_1 and C_2 in the combination model using the posteriors for the monotherapy parameters from the k_1-th batch of Drug 1 and the k_2-th batch of Drug 2. For h_1 and h_2, the priors are set to be lognormal$(0, 1)$ as in the original SynBa. The prior for E_3 is set to be $p(E_1E_2)$, assuming Bliss independence.
Step 5: Train the combination model as in Box 2 of [86], also outlined in Section 4.1.2.

In summary, if the training is performed using Batch-SynBa, then there is an additional step that extracts the posterior probabilities for each set of dosages and responses. This extraction allows for identifying the most likely *batch assignments* for the monotherapy measurements within the combination matrix, which are the first row and the first column of the matrix in our setting. This ensures that the combination model is informed by the batch-specific characteristics captured during the monotherapy training phase, which is expected to improve the model's ability to account for batch effects and variability.

In the prior specification step (i.e. Step 3 in Algorithm 5.2 and Step 4 in Algorithm 5.3), the prior distributions for the parameters E_1, E_2, C_1, C_2 in the combination model are defined using the posteriors from the pre-trained monotherapy models. For the parameters h_1 and h_2, the priors remain lognormal$(0, 1)$ as in the original SynBa framework to ensure numerical stability. The prior for E_3 is set to $p(E_1E_2)$, assuming a multiplicative (but not synergistic) interaction under the Bliss independence model [10]. This conservative prior allows the posterior to start from a null assumption and shift towards synergistic or antagonistic interactions only if supported by sufficient evidence. The final step involves training the combination model following the methodology outlined in Box 2 of [86].

This method allows utilising historical information from the existing large screenings or additional available data to enhance the inference or prediction of combination drug responses. By leveraging trained SynBa-Batch monotherapy models, it accounts for batch effects and informs the identification of significant batches, providing more accurate and robust predictions.

Note that when uncertainty is high, the posterior batch assignment probability, $p(z_i = k \mid \mathcal{D})$, will be lower, contributing less to the priors used in the combination model. By doing this, the model ensures that strong priors are only imposed when the historical

data provides information with low uncertainty. This process allows us to more carefully handle uncertainty in noisy data, avoiding overconfident predictions when the data does not provide a clear signal.

5.2.2.1 Case study

To illustrate how modelling monotherapies with SynBa-Batch helps inform combination models, we consider an example in DREAM.

Take the combination of FGFR and MTOR_1 applied to the cell line HCC70 as an example, as illustrated in Fig. 5.6. In this figure, subfigures (A) and (B) show the available monotherapy measurements for each of the single drugs, i.e. FGFR and MTOR_1 respectively. Each set of measurements, indicated by connected lines, comes from different experimental plates. Each plate contains five measurements at non-zero dosages. Historical measurements are coloured in blue, with data points from the same plate connected by lines. The set of monotherapy measurements associated with the combination (i.e. measured as a part of the combination, in the same plate and thus under the same controlled experimental condition) are coloured in red. They correspond to the first row and the first column of the dose-response matrix respectively.

For the drug FGFR illustrated in the left part of subfigure (A), if we only focus on the red data points, it is difficult to infer whether this monotherapy will eventually show efficacy at higher dosages, since the response remains relatively flat at the tested dosages. However, when considering the blue measurements, we observe two distinct patterns. One batch shows an efficacious response (steep decline in response), while the other exhibits flat lines, indicating little to no efficacy across a wide dosage range. This suggests that the batch effects from biological or experimental variation are influencing the potency of the drug. Given that the red line closely resembles the batch with flat lines, it can be inferred that the red measurements likely belong to the same non-efficacious batch, sharing the same parametric mean and minimal drug efficacy.

For the drug MTOR_1 illustrated in the right part of subfigure (A), if we focus solely on the red points, it remains ambiguous whether the observed decline in response is due to a true pharmacological effect or simply a result of measurement noise. However, when considering the historical data in blue, we can observe that there exists a batch where the drug demonstrates significant efficacy. This allows us to reasonably infer that the red measurements are likely part of the efficacious batch, though it is still fairly uncertain without further measurements.

Fig. 5.6(B) and (C) illustrate the outputs of Algorithm 5.1 trained on the existing monotherapy measurements of the two drugs. These are the historical information that we can utilise. Now the task is to assign the two sets of monotherapy measurements associated with the combination (i.e. the red measurements in Subfigure (A)) to one of these batches.

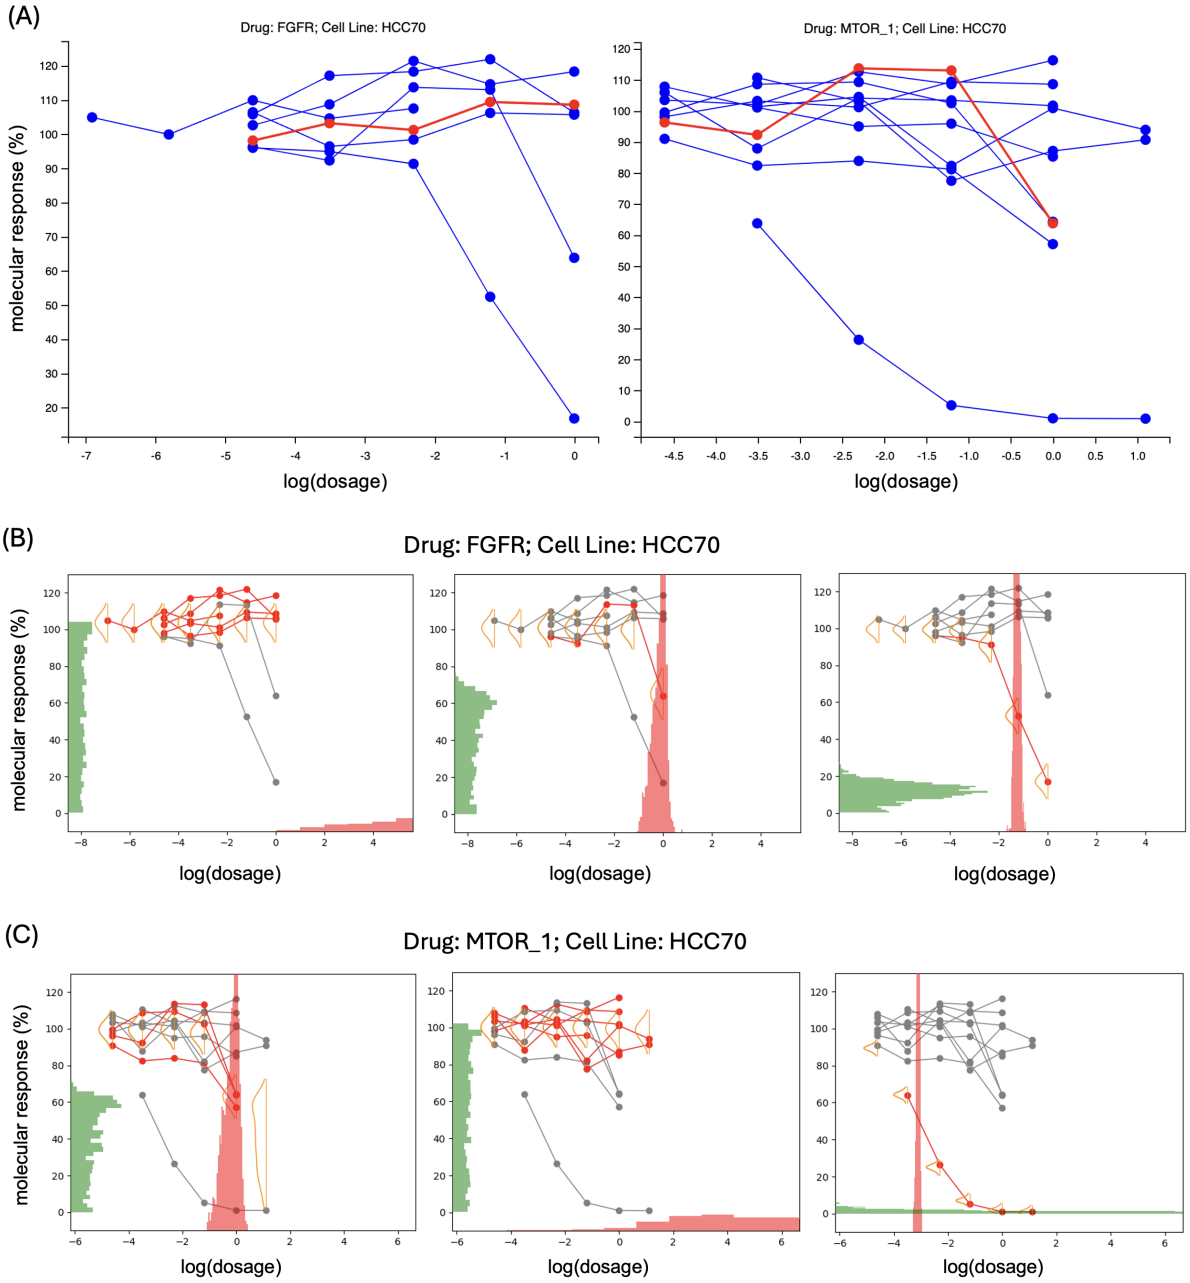

Figure 5.6: (A): Illustration of the monotherapy slices (from the combination dose-response matrix) among all historical monotherapy measurements. (B) and (C): Outputs from Algorithm 5.1 for the three learnt batches of Drug FGFR and MTOR_1 respectively. Measurements are coloured in red if their probability of being assigned to this batch is larger than the others. The distribution for IC_{50} is shown in red, whereas the distribution for E_{inf} is shown in green.

For FGFR, it is assigned to the first of the three batches in Subfigure (B), which is the "constant batch". It provides little prior information on the efficacy of the drug, as more measurements at larger dosages would be required. For MTOR_1, the monotherapy is assigned to the first of the three batches in Subfigure (C). This is a batch that includes sets of measurements with a significant level of efficacy. There is also strong information on the potency.

By having these informative priors, we will be able to better infer the combination parameters related to synergistic efficacy and potency. This is especially the case when the size of the dose-response matrix is small which could result in a large epistemic uncertainty if the priors are conservative (and could lead to over-parameterisation if deterministic parametric approaches are used instead of the Bayesian approaches).

5.3 Prediction of drug combination responses

In this section, we apply SynBa and SynBa-Batch informed with monotherapy information to the task of predicting unseen responses within a dose-response matrix using the DREAM challenge ([54]). The aim is two-fold. First, we assess whether the prediction accuracy is improved by integrating the existing monotherapy measurements. We utilise historical information by defining the priors of the combination model using the posteriors learnt from the monotherapy models. This data utilisation step is performed by Steps 1 and 2 in Algorithm 5.2 if the original SynBa is used for monotherapy modelling, and by Steps 1 to 4 in Algorithm 5.3 if SynBa-Batch is used for modelling the monotherapies. The objective is to provide a more informed starting point for the combination model, which would be particularly helpful when the number of available measurements is small.

Secondly, we investigate whether the prediction is improved by learning the batch effects, which is done by incorporating a mixture model through SynBa-Batch as opposed to the single-mean likelihood model in SynBa. By capturing and explaining away the batch effects present in the data, the correct batch can be identified for each of the monotherapy sets, rather than being averaged out with the other batches.

In DREAM, due to the limited number of measurements per monotherapy, we focus on a subset of Challenge 1 where the number of monotherapies is sufficiently high for batch effect modelling to be meaningful. This reflects the real-world scenario in which historical data on monotherapy measurements for each of the two drugs to be combined are available, since combinations are in most cases not randomly selected but are informed by empirical evidence and established knowledge. Fig. 5.7 illustrates the number of monotherapies in DREAM across drugs and cell lines. The rows and columns are sorted by the hierarchical clustering algorithm based on the counts for each (drug, cell line) pair. The top-left part

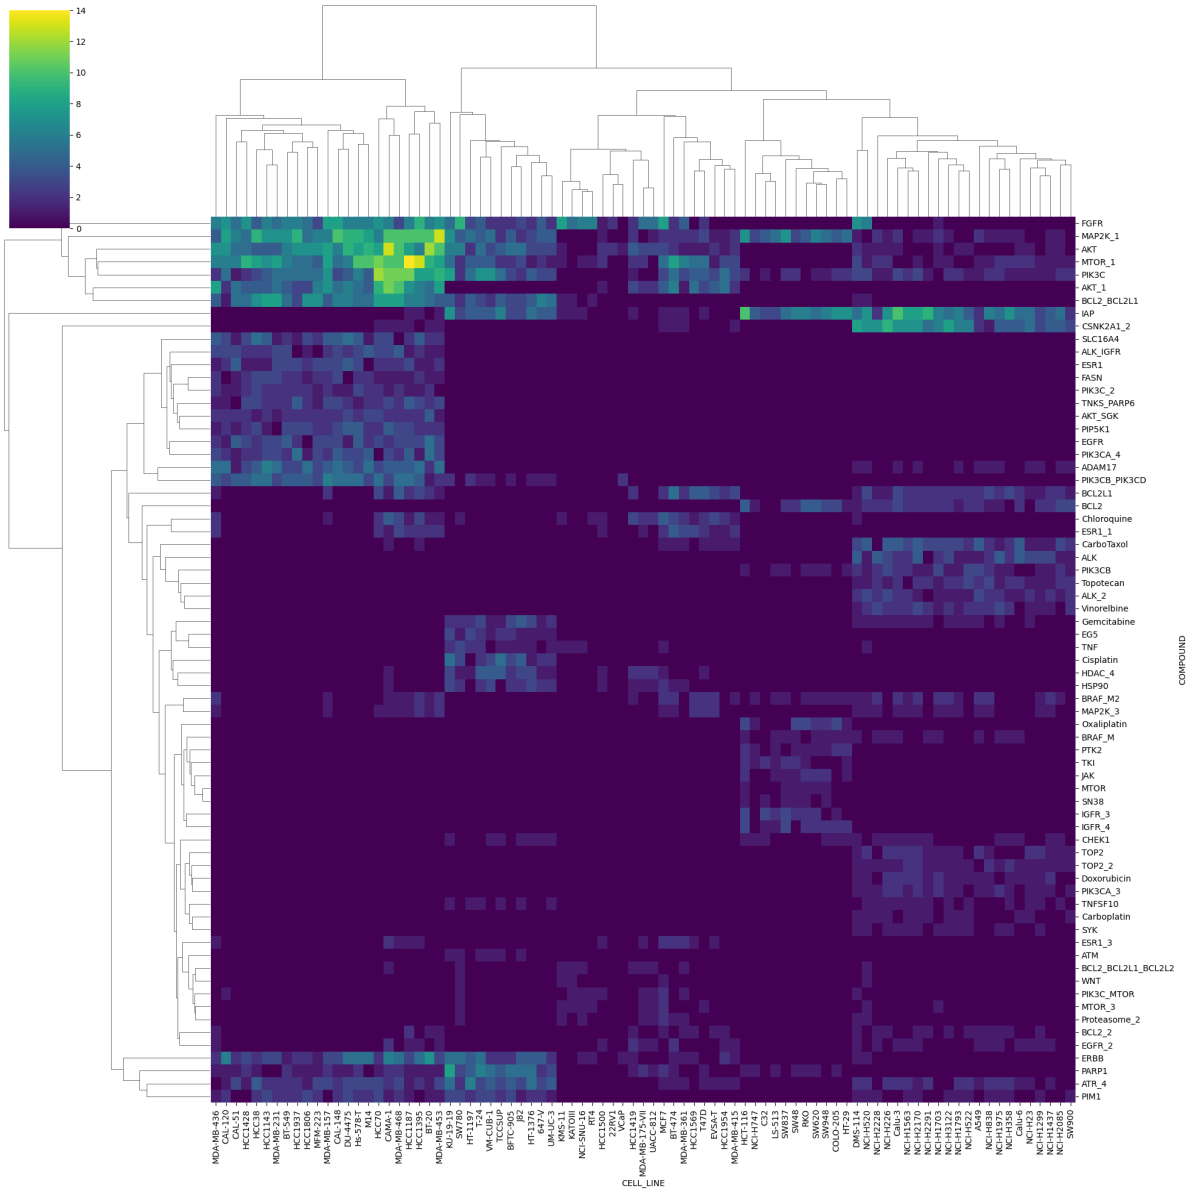

Figure 5.7: The counts for the sets of available monotherapies in DREAM Challenge 1.

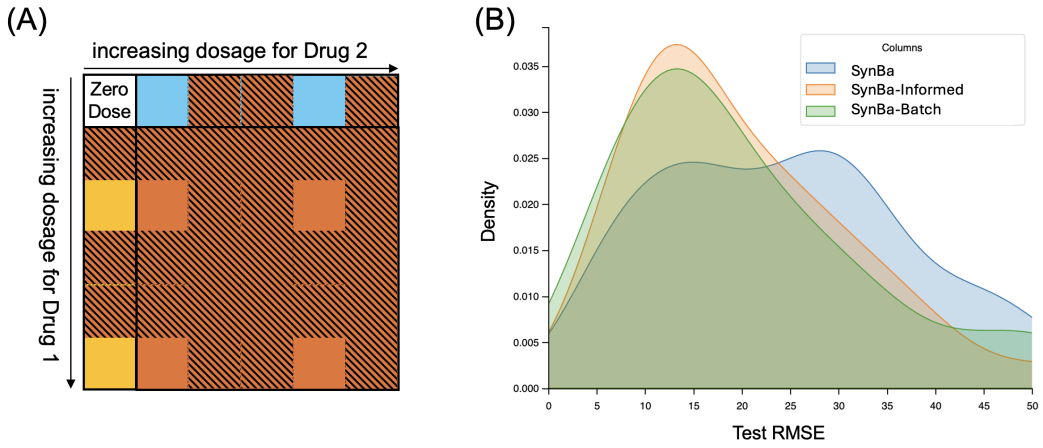

Figure 5.8: (A) The dataset and the prediction task in Section 5.3. The shaded rows and columns represent the test data, which are randomly selected from non-zero dosages for each drug. (B) Visualisation of the test performance for filling up the grid. The smoothed densities for SynBa-Batch and SynBa are illustrated.

of the heatmap consists of monotherapies with significantly more measurements available, which are mostly breast cancer cell lines. For the analysis, we have chosen the following compounds:

- FGFR, MAP2K_1, AKT, MTOR_1, PIK3C, AKT_1, BCL2_BCL2L1

and the following cell lines:

- MDA-MB-436, CAL-120, CAL-51, HCC1428, HCC38, HCC1143, MDA-MB-231, BT-549, HCC1937, HCC1806, MFM-223, MDA-MB-157, CAL-148, DU-4475, Hs-578-T, M14, HCC70, CAMA-1, MDA-MB-468, HCC1187, HCC1395, BT-20, MDA-MB-453.

Except for M14 which belongs to soft tissue, all are breast cancer cell lines.

For each drug, we randomly leave out three of the five non-zero dosages from the 6-by-6 dose-response matrix, but not the first row or the first column as they are the monotherapy measurements. Fig. 5.8(A) illustrates the leave-out strategy. The train-test split is done in this manner so that the training data simulates a scenario where we have a small 3-by-3 dose-response matrix, and need to infer the synergistic potency and efficacy of the combination from the limited number of measurements. The test data evaluates whether the model can accurately infer these quantities. If the predictions are accurate, then this model can support decision-making by (1) suggesting the optimal dosage, taking into account both efficacy and toxicity, and (2) if uncertainty remains high (which is very likely the case, due to the small size of the dose-response measurements for combinations), recommending the range of dosages for the next batch of experimentation to reduce the uncertainty.

Table 5.1: The mean, 1st quartile (Q1), median and the 3rd quartile (Q3) of the test root-mean-squared error (RMSE) for SynBa, SynBa-Informed and SynBa-Batch, computed on a subset of DREAM, along with their standard errors. The standard error of the mean is computed by the standard deviation of the metrics across examples divided by the square root of the number of examples. The standard error of the median, Q1 and Q3 are estimated by nonparametric bootstrap.

	mean (± se)	Q1 (± se)	median (± se)	Q3 (± se)
SynBa	24.62 ± 1.20	13.71 ± 2.28	25.43 ± 2.21	32.80 ± 3.14
SynBa-Informed	$\mathbf{20.24} \pm 1.08$	11.82 ± 1.39	17.39 ± 1.56	$\mathbf{27.31} \pm 3.78$
SynBa-Batch	20.39 ± 1.21	$\mathbf{10.86} \pm 2.15$	$\mathbf{16.49} \pm 1.65$	27.40 ± 4.85

Similar to SynBa, the weighted SynBa in Step 3 of Algorithm 5.1 is trained by Stan, a state-of-the-art platform for Bayesian computation ([75]), in which we specify the prior model of the parameters and the likelihood model of the data, whereas Stan performs full Bayesian statistical inference with Markov chain Monte Carlo (MCMC) sampling.

Table 5.1 shows the mean, 1st quartile (Q1), median and the 3rd quartile (Q3) of test RMSEs for SynBa without utilising monotherapy information (i.e. the original version in Chapter 4), SynBa-Informed (i.e. Algorithm 5.2, SynBa with monotherapy information utilised) and SynBa-Batch with monotherapy information utilised (i.e. Algorithm 5.3). These are also illustrated by Fig. 5.8(B), where the smoothed densities are plotted for the test RMSEs for each of the methods. Results show an improvement in RMSEs by incorporating historical data into the priors for combination models, as shown by the improvement in RMSEs for SynBa-Informed and SynBa-Batch compared to the original SynBa. While the RMSEs do not show a clear advantage of modelling the batch effect in SynBa-Batch compared to SynBa-Informed in terms of prediction accuracy, the visualisations generated by SynBa-Batch still offer additional insights, as described in the case study earlier in this chapter. These visualisations allow for a deeper understanding of drug behaviour under varying laboratory conditions or in the presence of biological variations, such as those caused by different cell culture types. This added interpretability is particularly useful when investigating the impact of external factors on drug efficacy and potency.

5.4 Conclusions

In summary, by leveraging the prior information from historical monotherapies and addressing the batch variability present in monotherapy, more informed inferences can be made about the combinations, leading to better predictive performance on untested dosages compared to models that do not account for historical measurements or batch-specific effects. This could be particularly helpful when we work with multiple datasets containing

different sources of variations. In these cases, SynBa-Batch can be applied to harmonise the measurements from different sources. By integrating diverse datasets, we can leverage a broader range of experimental data, leading to more accurate predictions and deeper insights into drug behaviour across varying conditions.

More broadly speaking, this chapter not only presents a new method, but also highlights the significance of batch effects in pharmacological modelling, an issue that has been largely overlooked. While SynBa-Informed and SynBa-Batch have provided some improvements in the specific task of predicting left-out measurements, the real value lies in how it enables a more detailed inference and more informed decision-making. This shows the benefits of explicitly modelling batch effects. With this chapter, we advocate for future methods to take batch effects into the modelling for more reliable and actionable predictions in pharmacology.

Chapter 6

DeepSynBa: Drug combination prediction with complete dose-response profiles

In the previous chapters, we have focused on the inference of drug combinations where dose-response measurements are available. More specifically, we introduced SynBa and SynBa-Batch to quantify the synergistic potency and synergistic efficacy of the combinations, along with their uncertainty estimates. SynBa and SynBa-Batch have addressed the scenario (S1) in Fig. 6.1 with a promising performance. However, when dealing with a large number of possible chemical compounds and a diverse set of cancer cell lines, it is impractical to experimentally test all possible triplets of drug 1, drug 2 and cell lines. This brings us to the next challenge: How do we predict the synergistic effects of untested drug combinations? More specifically, how do we approach the scenarios (S2) and (S3) in Fig. 6.1, in which significant extrapolation would be required? In (S2), we focus on untested combinations with existing monotherapy responses, whereas in (S3), we have no monotherapy measurements either, and the only data we have are the chemical structures of the two drugs and the biological profile of the cell line.

6.1 Previous methods

The DREAM challenge, which is the dataset of focus in this project, was initially designed to address such questions. More specifically, two challenges were posed. The first is to predict synergies in cancer cell lines for untested drug combinations. The second is to derive mechanistic models or rules, such as biomarkers, that discriminate between synergistic and non-synergistic behaviours. Apart from the 6-by-6 dose-response matrices, the dataset also includes molecular data for the untreated (baseline) cell lines and chemical information for the respective drugs. 160 teams participated in this challenge. Synergy is predicted

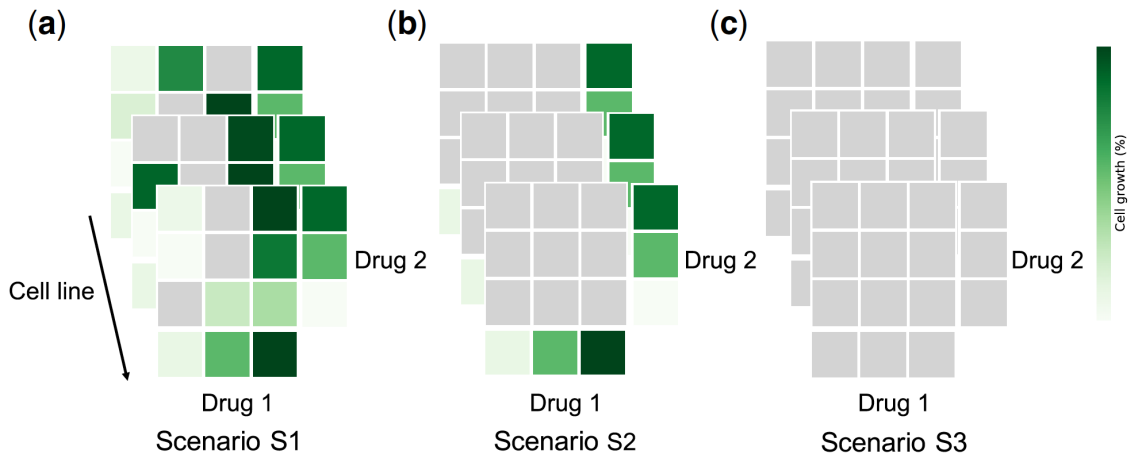

Figure 6.1: Illustration of different drug combination response prediction scenarios. (a) Filling in the gaps in partially measured dose–response matrices (S1); predicting dose–response matrices of new drug combinations (b) with monotherapy responses (S2) and (c) without monotherapy responses (S3). Figure from [80].

with an accuracy matching biological replicates for more than 60% of combinations [54]. Amongst these methods, some identified genomic rationale for synergy predictions, such as ADAM17 inhibitor antagonism when combined with PIK3CB/D inhibition contrasting to synergy when combined with other PI3K-pathway inhibitors in PIK3CA mutant cells [54].

Apart from the DREAM challenge, which stimulated the methodology development for predicting drug combinations, a wide variety of methods have been developed in this growing research area, making use of the wide range of available screenings in the past decade [2]. More specifically, they focus on predicting the synergy scores of drug combinations using drug and cell line related inputs such as the transcriptomic profiles of the cell lines and the chemical structures of the drugs [62, 41, 23, 79, 42, 36, 7, 1]. Preuer et al. developed DeepSynergy [62], in which the chemical structures and cancer cell line gene expression data are taken as input to a fully connected neural network to predict the Loewe synergy scores of drug combinations. Ling et al. developed IDACombo [50], which predicts the combination synergy by assuming the principle of Independent Drug Action and predicting that the efficacy of a drug combination in a given cell line or patient will be equal to the effect of the single best drug in that combination. This simple assumption provides a somewhat surprisingly good prediction performance on both in vitro and clinical trial datasets. Jin et al. designed a neural network architecture that jointly learns drug-target interaction and drug-drug synergy through a drug-target interaction module and a target-disease association module, which effectively identified synergistic drug combinations for treating COVID-19 [36]. Bertin et al. developed RECOVER [7], a sequential model optimisation search platform, which was applied to a deep learning model to quickly discover synergistic drug combinations. Kuru et al. designed MatchMaker [41],

a multimodal architecture that employs three fully connected neural networks to learn cell line-specific representations of drugs and predicts the Bliss synergy scores of drug combinations. More recently, Kuru et al. developed PDSP [42], in which MatchMaker was extended to include the patient-specific single drug response data for customising patient drug synergy predictions.

The development of these methods is beneficial for discovering and explaining drug combinations. However, a few factors prevent many of them from being widely applied in real-world drug discovery projects [6]. First, many of these methods are trained to perform well on data of the same format, some of which are due to the nature of the DREAM challenge. However, the resulting models are often not generalisable to the real-world scenario, where the data have a diverse format and have very different choices of endpoints.

Secondly, the prediction labels are synergy scores in most cases (e.g. Loewe synergy score in DeepSynergy [62] and Bliss synergy score in MatchMaker [41] and RECOVER [7]), which is a proxy measure that contains its shortcomings as discussed in the previous chapters. It cannot, for example, distinguish between combinations that are highly potent at low doses and those that are efficacious but require higher doses, nor can it account for the varying levels of uncertainty associated with these outcomes. By using single synergy scores such as Bliss or Loewe as the response variable, we are attempting to predict something that is vague on its own and does not paint a complete picture for the combination profiles.

Moreover, due to the limited data size for each combination and the inherent noise in the data, any label defined using a single synergy score would be associated with a non-trivial amount of uncertainty, but uncertainty measurements are not included as part of the labels. The Spearman correlation of the replicate experiments performed in DREAM [54] and the O'Neil et al. dataset [59] are 0.56 and 0.63 respectively, which show that quantifying a combination with a single synergy score would result in a high variance. As a result, the models are encouraged to overfit to a proxy measure that is inaccurate and uncertain. A high level of uncertainty is ignored during the modelling process. This could be one of the reasons that 20% of drug combinations are poorly predicted by all methods in the DREAM challenge [54]. Predicting uncertainties in the outputs would be important for the subsequent decision-making based on the model outputs.

Due to the shortcomings of the traditional synergy score measures such as Bliss and Loewe, a natural alternative would be to directly predict the entire dose-response matrices instead of a single synergy score measure. By predicting the entire grids, we obtain a more detailed profile for the combinations at different dosages. Once the grids are predicted, we can then compute the synergy scores as a follow-up step after the model is trained. There have been some recent methods focusing on this alternative route. Julkunen et al. developed comboFM [37], which models the drug interactions through higher-order

tensors, and efficiently learns latent factors of the tensor with factorization machines. Wang et al. developed comboLTR [80], which is also based on polynomial regression through higher-order tensors like comboFM, but different to comboFM, in comboLTR the latent tensor reconstruction (LTR) is employed as an alternative of the factorization machines used in comboFM, hence extending the range of functions that can be learnt and allowing much larger datasets to be processed compared to comboFM. ComboFM and comboLTR are shown to be competitive in all three prediction scenarios illustrated in Fig. 6.1. Rønneberg et al. developed PIICM [67], an extension from bayesynergy [66] discussed in Chapter 4, which directly predicts the dose-response surfaces of untested drug combinations by utilising a permutation invariant version of the intrinsic co-regionalisation model for multi-output Gaussian process (MOGP) regression. More recently, Gutierrez et al. [30] developed a probabilistic multi-output model to simultaneously predict all dose–responses and uncover their biomarkers. MOGP models are employed to describe the relationship between genomic features and chemical properties to every response at every dose, different from PIICM which incorporates MOGP for extrapolating curves only.

Despite recent methods such as comboFM, comboLTR and PIICM providing a more complete picture for the combination profiles, the lack of interpretability for the model outputs is still a major barrier for existing methods to be applied in real-world drug discovery settings. In particular, for combinations, most of these models typically do not offer an explicit separation of key synergy parameters such as the synergistic potency and the synergistic efficacy. In PDSP, the potency of drug 1 and drug 2 (quantified by a binarised IC50) are separated out, which is a forward step towards providing actionable insights for further decision-making on the combination candidates. However, most other methods have not done this. There is also no principled uncertainty estimation for these predicted outputs.

Incorporating SynBa into these predictive models could offer a more interpretable and quantitative approach by explicitly providing posterior distributions over the key synergy parameters such as synergistic efficacy and synergistic potency. This allows for a clearer understanding of whether a drug combination is effective due to improved potency (lower toxicity) or increased efficacy (greater maximal effect). In addition, SynBa quantifies the uncertainty around these estimates, giving researchers a sense of confidence in these predictions. If each prediction for a combination can be accompanied with a decoupled estimation of its potency and efficacy, along with an uncertainty estimation for each of them, then the model could potentially be more helpful for biologists to make informed decisions based on the predictions.

Building on these motivations, we develop DeepSynBa, a prediction model that incorporates the likelihood formulation of SynBa to directly predict the dose-response grids. We formulate the dose-response matrix prediction task as a regression problem, where the

model is trained with dose-response matrices to predict matrix entries across a dosage grid. DeepSynBa employs the Matchmaker's cell line conditioned drug representation modules. To parameterise dose-response surfaces and disentangle the synergistic potency and efficacy, DeepSynBa follows the approach outlined in SynBa [86] and adopts a simplified version of the two-dimensional generalised Hill Equation derived in MuSyC [83]. DeepSynBa's novel intermediate layer learns the parameters in generalised Hill equation during training. Using these predicted parameters, the complete dose-response surface can be generated for a given drug combination as shown in Fig. 6.3 (a) and (e) panels. This model design also enables the computation of traditional synergy scores like Loewe and Bliss post-hoc, ensuring compatibility with existing drug synergy prediction methods while remaining flexible without being constrained to any specific synergy score.

Our contributions can be summarised as follows: DeepSynBa (i) provides a better dose-response prediction capability than existing dose-response predictors with its novel framework, (ii) allows separating out efficacy and potency which is not possible with existing models, and (iii) can assist experimental design by generating the dose-response surface across all dosages and thus identifying the suitable dosage range for the next batch of experiments.

6.2 DeepSynBa: Our new approach

From our discussions, we hope to construct a model that (1) directly predicts the dose-response grid instead of a single synergy score, (2) provides a decoupled estimation for synergistic efficacy and synergistic potency, and (3) includes some level of uncertainty estimation for the quantities of interest including efficacy, potency and the drug responses.

We now introduce a model that achieves (1) and (2) by adding a SynBa likelihood layer before the dose-response matrix itself, as illustrated in Fig. 6.2. To also achieve (3), an (at least partially) probabilistic model would be required so that the potency and efficacy parameters are treated as random variables and the uncertainties can be estimated through their posteriors. This will be the next step, but before that, it would be helpful to first investigate whether our model can make some improvements on the prediction tasks by incorporating the SynBa likelihood, which is the focus of this work.

The proposed model architecture, illustrated in Fig. 6.2, is designed to predict the effects of drug combinations on cell lines by incorporating both drug and cell line features. Each drug combination consists of two drugs applied to a single cell line, and the model includes two dedicated drug-cell encoder sub-networks. Each encoder focuses on fusing the properties of one drug with the gene expression profiles of the cell line, enabling a detailed capture of interactions between the drugs and the cellular gene expression profiles.

The model uses feature representations for drugs (d_i and d_j) and cell lines (c_k). Drug

features are generated as vector embeddings from the MoLFormer model [68], and the cell line features are represented by vectors of untreated gene expression data. These features are concatenated into a combined vector ($[d_i; c_k]$ or $[d_j; c_k]$) and then processed through a series of fully connected layers. These layers transform the combined drug-cell line features into a comprehensive representation that captures the interactions between the drugs and the cell line. The outputs of the drug-cell encoders are drug-induced gene expression feature maps, which encapsulate the interaction between each drug and the cell line at a molecular level. The feature maps from each of the encoders are then concatenated to produce a unified representation, effectively summarising the combined impact of both drugs on the cell line.

Subsequently, the unified representation is inputted into eight distinct prediction heads, each corresponding to one parameter in the SynBa likelihood function [86]. Each prediction head comprises fully connected layers that process the unified representation to predict a single scalar value for each parameter. By using these prediction heads, the model can estimate the SynBa parameters.

Specifically, the prediction heads are $E_1, E_2, E_3, C_1, C_2, H_1, H_2$ and α from the likelihood function defined in SynBa, which are the parameters that characterise the expected value of the drug responses y. Given the dosages x_1 for Drug 1 and x_2 for Drug 2, the expected value for a response y is defined as

$$\begin{aligned}&\mathbb{E}[Y \mid E_0, E_1, E_2, E_3, C_1, C_2, H_1, H_2, \alpha, x_1, x_2] \\ =&\frac{C_1^{H_1} C_2^{H_2} E_0 + x_1^{H_1} C_2^{H_2} E_1 + C_1^{H_1} x_2^{H_2} E_2 + \alpha x_1^{H_1} x_2^{H_2} E_3}{C_1^{H_1} C_2^{H_2} + x_1^{H_1} C_2^{H_2} + C_1^{H_1} x_2^{H_2} + \alpha x_1^{H_1} x_2^{H_2}},\end{aligned} \tag{6.1}$$

where the formula and the definitions of the parameters follow Chapter 4. For the stability of training, E_0 is set to be a fixed value in DeepSynBa and determined by the definition of the dataset. For example, the base value is 100 in NCI-ALMANAC.

This architecture enables the model to focus on directly predicting dose-response matrices, whereas the synergy scores, regardless of their definitions, can be computed as a follow-up step.

6.3 Case studies

To demonstrate how DeepSynBa provides deeper insights by predicting the full dose-response matrix or surface, we examine two specific examples in the NCI-ALMANAC dataset. The first example focuses on the combination of Bleomycin Sulfate (denoted as Drug 1) and Cabazitaxel (denoted as Drug 2) administered on the renal cancer cell line 786-0. Drug 1 is an antibiotic with a chemotherapy effect on fast-growing cells, while Drug 2 is a strong chemotherapy agent used in the treatment of prostate cancer.

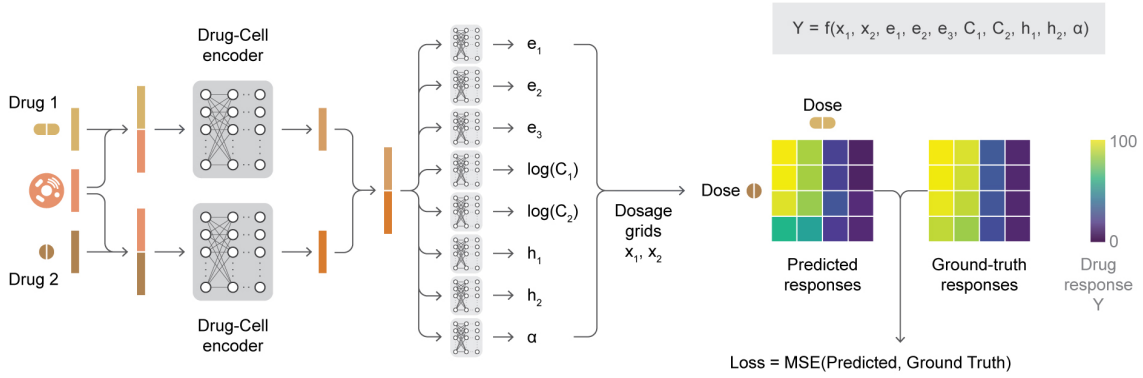

Figure 6.2: Our improved model integrates the SynBa likelihood function. It uses the unified drug-cell line representation to predict SynBa parameters, directly generating dose-response matrices, supervised by dose-response matrix labels.

Fig. 6.3(a)-(d) illustrates the predicted and measured effects of the drug combination. Panel (a) presents the predicted dose-response surface, with the actual measured responses coloured in white. The surface predicted by DeepSynBa matches the downward pattern of the measured responses as the dosages of both drugs increase. The killing effect of the drugs is quantified as a drop along the y-axis.

Panel (b) presents the corresponding contour plot for the dose-response surface, highlighting the combinations of dosage levels that reach the predicted drug effects. The plot reveals a rapid decline along the y-axis from 100% to 60%, indicating fast cell killing as the dosages for both drugs increase. More interestingly, this drop is accompanied by evidence of synergistic potency, as indicated by the curvature of the contour lines for response levels 90%, 80%, 70%, and 60%. For a combination that is neither synergistic nor antagonistic, the contour lines for dose-response levels typically resemble a circle or an oval. In panel (b), the curvature bends inward towards the origin, showing that a smaller dosage is required when the drugs are combined to achieve the same level of efficacy as single-drug treatments. However, beyond the 60% response level, increasing the dosage of Drug 1 (right side of the x-axis) does not significantly improve the treatment response. In contrast, Drug 2 continues to exhibit increasing effectiveness as the dosage increases. This aligns with the fact that Drug 2 is a strong chemotherapy agent, and it is expected to become dominant at higher dosage levels for both drugs.

Panel (b) also demonstrates that even for the same drug combination, the synergy profile can vary across different dosage ranges. Synergistic potency may only be present within specific dosage subsets. This further highlights the importance of predicting the full dose-response surface to identify actionable combinations and their respective dose ranges.

The ability to estimate the full drug combination profile surface allows us to take slices of the surface for further inspection of the drug's interaction and effects. For panels (c)

and (d), slices were taken from the dose-response surface at the dosage levels where actual measurements were available, enabling an evaluation of predictions against the observed data. The measured responses (represented with the dots) and predicted dose-response curves are shown for each of the two drugs.

In both panels (c) and (d), it can be observed that at lower doses of one drug, the curves shift downward as the dose of the other drug increases. However, when both drugs are administered at higher doses, the difference between the curves narrows. This is particularly evident in (d), where the predicted maximal efficacy does not improve as the dose of Drug 1 increases. Nevertheless, its drug response reaches the IC50 at a lower dosage when combined with the other drug. This shows the potential benefit of combining these two drugs at a medium-to-low dose range. Such a combination could reduce the risk of off-target toxicity associated with high doses of the strong chemotherapy agent, achieving the desired therapeutic effect with lower doses of each drug compared to their monotherapy counterparts.

In the second example, we look at the combination of Azacitidine (Drug 1) and Triethylenemelamine (Drug 2), two strong chemotherapy agents, administered on MDA-MB-468, a triple-negative breast cancer (TNBC) cell line. Azacitidine is used in leukaemia, while Triethylenemelamine is particularly effective against rapidly growing cancer cells. In Fig. 6.3, panels (e)-(h) illustrate their predicted behaviour and the comparisons with the actual measurements. In panel (f), it is evident that there is no clear synergistic interaction between these two drugs, as the contour lines from 80% to 40% mostly bend outwards, away from the origin (i.e. the bottom-left corner of the plot). Despite both drugs being effective on MDA-MB-468, combining them does not offer a clear advantage for this particular tumour indication. This is an example showing how our model helps to identify the lack of beneficial interactions in a specific setting, providing insight into why certain combinations might not be suitable for further testing in that tumour type.

These two examples demonstrate the ability of DeepSynBa to model and predict how two drugs interact while separately quantifying two key combination metrics: synergistic potency and maximal efficacy. This approach provides not only clearer biological insights into their combined effect but also valuable guidance for decision-making when evaluating drug combination dose regimes.

Fig. 6.3 illustrates that the performed analysis would not be possible if the entire combination profile were summarised into a single synergy score. The richer information provided by the full-dose response matrix is crucial for optimising combination therapies, as it identifies regions in the dose space where drugs can achieve significant benefit through enhanced potency. This could help guide experimental design in collecting further evidence for the combination regimes and ultimately supporting the development of safer and more effective treatment strategies.

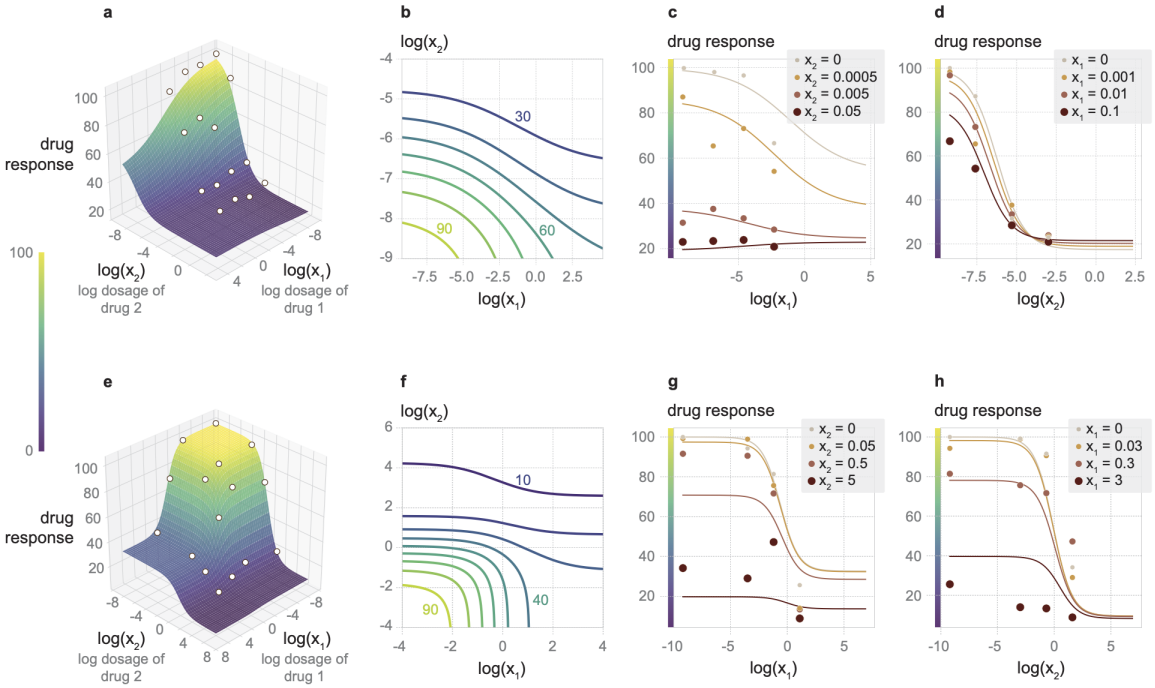

Figure 6.3: The predicted dose-response outputs from DeepSynBa for two different drug-cell line combinations and their comparisons with the ground-truth experimental measurements. (a)-(d) focus on the combination of Drug 1 (Bleomycin Sulfate) and Drug 2 (Cabazitaxel) administered on cell line 786-0. (e)-(h) correspond to the combination of Drug 1 (Azacitidine) and Drug 2 (Triethylenemelamine) administered on the cell line MDA-MB-468. (a)(e): Predicted dose-response surfaces and the measured responses (highlighted in white). (b)(f): Contour plots for the predicted drug response at various dosage levels of the two drugs. (c)(g): Dose-response curves for fixed concentrations of Drug 2 (in unit μM), comparing predictions with measurements (in μM). (d)(h): Dose-response curves for fixed concentrations of Drug 1 (in unit μM), comparing predictions with measurements (in μM).

6.4 Results

We evaluate the performance of our models across three stratification scenarios: (i) new drug combinations, (ii) new cell lines, and (iii) new drugs. In the case of new drug combinations, a specific drug pair $< i, j >$ is restricted to appearing in only one of the three datasets: training, validation, or testing. For the new cell scenario, drug combinations are grouped by cell line types, ensuring that any given cell line k is included exclusively in one of the splits, i.e. training, validation, or test. In the new drug scenario, drug combinations in the test set must include at least one drug that is absent from any combinations used in the training or validation sets. Similarly, the validation set must contain drug combinations with at least one drug not present in the training set. In all three scenarios, the data is split into train, validation and test with 60%, 20% and 20%, respectively.

We evaluated our model's performance against comboFM [37] and comboLTR [80], two models designed for predicting dose-response outcomes in drug combination studies by modelling complex multi-dimensional interactions among drugs, dosages, and cell lines. comboFM leverages factorisation machines to create a high-order tensor of drug features, dosage levels, and cell line responses, facilitating the capture of unique interactions across various cell contexts. In contrast, comboLTR uses latent tensor reconstruction (LTR), offering an alternative approach to factorisation for modelling intricate dose-response relationships within these datasets. Among existing methods, comboFM and comboLTR are the most similar to DeepSynBa in terms of the data modalities used and the outputs predicted.

While PIICM also predicts dose-response surfaces, its primary focus is on modelling correlations between dose-response functions across different experiments. Cell or drug properties are not incorporated, restricting its scope to cell lines and drugs that have already been tested. Therefore, it cannot be directly applied to evaluation scenarios which involve new drugs and new cell lines. This is why we did not include it in our experiments. For scenario (i), PIICM has shown a similar predictive performance to comboLTR when tested on the O'Neil dataset [59, 67], which is an indication of how its predictive performance compares with ours. Moreover, one major advantage of PIICM is that it accounts for the experimental uncertainty, which is also a crucial aspect of dose-response prediction, but not our focus in this work.

To comprehensively evaluate models' performances, we assess each model across various scenarios using three core metrics: root mean squared error (RMSE), Pearson correlation, and Spearman correlation. These metrics quantify the alignment between predicted dose-response matrices and the ground-truth values. The metrics are calculated over all entries of the dose-response matrices. Table 6.1 summarises each model's performance in terms of RMSE, Pearson, and Spearman correlations across these scenarios.

DeepSynBa consistently outperforms all competing models across each experimental setting. Specifically, DeepSynBa achieves RMSE values of 7.72 and 7.71 in the new drug combination and new cell line settings, respectively (Table 6.1a). This performance represents a substantial improvement, approximately 65%, relative to the next-best model, comboFM (Table 6.1b). Such results underscore DeepSynBa's ability to capture complex drug-drug-cell interactions more effectively.

While all models experience performance drops in the more challenging new drug evaluation scenario, DeepSynBa consistently leads. DeepSynBa achieves an RMSE of 18.69 compared to 40.22 from comboLTR and 42.39 from comboFM. The same holds for the correlation metrics (Table 6.1c). For both Pearson and Spearman correlations, DeepSynBa demonstrates higher scores than other models. This consistent advantage across RMSE, Pearson, and Spearman metrics illustrates the effectiveness of DeepSynBa's architecture

	RMSE	Pearson	Spearman
comboFM	12.76	0.87	0.68
comboLTR	16.11	0.78	0.64
DeepSynBa	**7.72**	**0.95**	**0.82**

(a) New drug combination scenario

	RMSE	Pearson	Spearman
comboFM	12.86	0.87	0.69
comboLTR	15.23	0.81	0.67
DeepSynBa	**7.71**	**0.95**	**0.82**

(b) New cell line scenario

	RMSE	Pearson	Spearman
comboFM	42.39	0.35	0.24
comboLTR	40.22	0.40	0.40
DeepSynBa	**18.69**	**0.71**	**0.64**

(c) New drug scenario

Table 6.1: Performance comparison of all methods in three different scenarios: (a) new drug combination, (b) new cell line and (c) new drug.

in accurately predicting drug combination responses.

6.4.1 Performance on Cell Lines and Drugs

We also calculated the performance metrics separately for each cell line and drug, enabling a detailed view of model performance across various tumour indications and drug pharmacokinetics (PK). These results are presented as distributions of performance metrics obtained for cell lines and drugs in Fig. 6.4. A pairwise t-test is applied to evaluate the statistical significance of performance differences between methods. All pairwise comparisons are statistically significant, with p-values $< 10^{-10}$. As shown in Figure 6.4, the RMSE distributions for cell line-based and drug-based results are concentrated at lower values for DeepSynBa, while the Pearson and Spearman correlation distributions are concentrated at higher values. These concentrated distributions reflect DeepSynBa's superior predictive performance. Results from the t-test confirm that these performance improvements are statistically significant, highlighting DeepSynBa's advantage in predicting dose-response outcomes. Moreover, the shorter height of DeepSynBa's violin plots suggests its performance is consistent across different cell lines and drugs, indicating greater robustness.

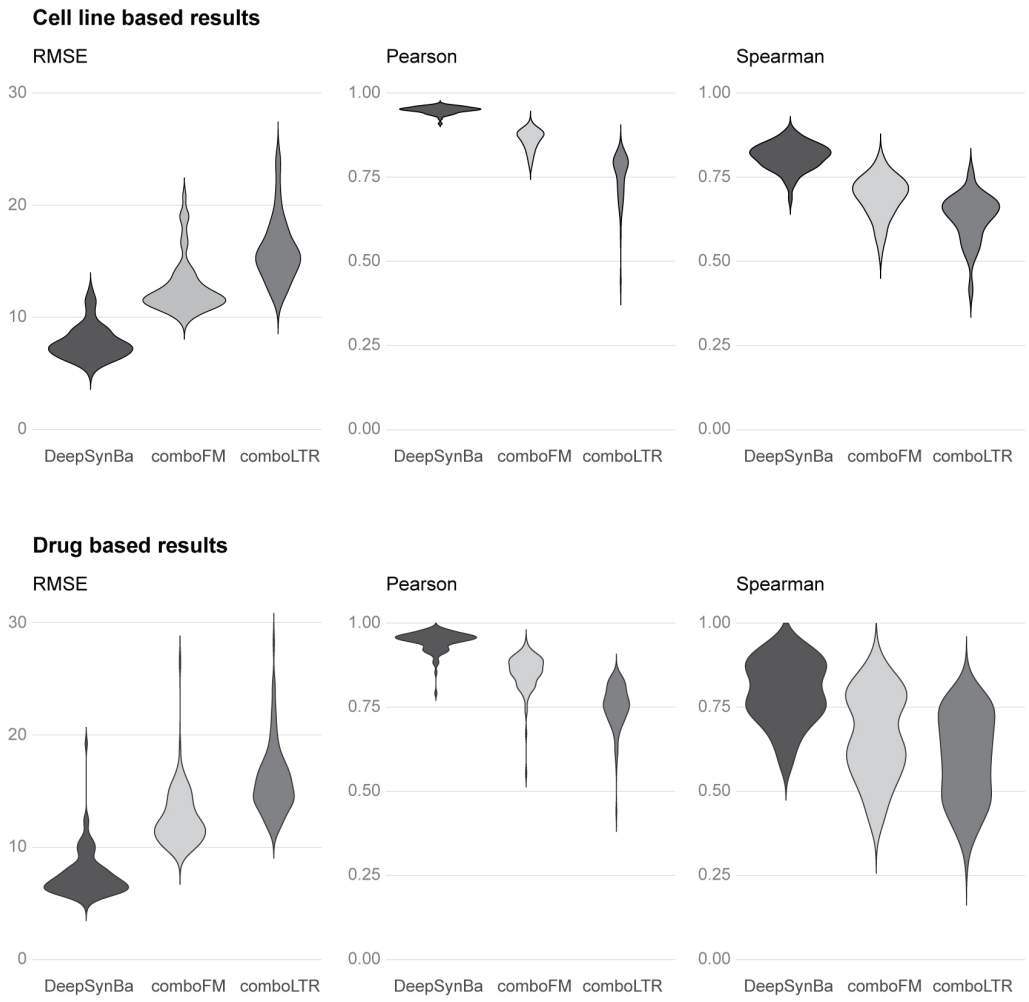

Figure 6.4: The distributions of RMSE, Pearson and Spearman correlations of dose-response surface predictions on different cell lines (upper panel) and drugs (lower panel).

6.4.1.1 Performance Across Tissues

We also evaluate model performance across both tissue types. The cell lines are grouped by their tissue of origin, and performance metrics are calculated within each group to assess variations in predictive accuracy across different tissues. Although performance improvements vary across tissues (see Table 6.2), DeepSynBa consistently outperforms other models across all examined tissue types. The model achieves the highest predictive accuracy for dose-response predictions in renal, breast, and ovarian cancers while exhibiting lower performance in leukaemia. This discrepancy likely arises from the distinct biological characteristics of these cancers and the nature of the screening assays, which tend to be more effective on solid tumours. Breast and ovarian cancers are characterised by well-defined genetic markers, such as BRCA1/2 mutations, which may facilitate model learning by providing clearer molecular patterns [61, 64, 4]. Similarly, renal cancers are

	RMSE			Pearson			Spearman		
	DeepSynBa	**comboFM**	**comboLTR**	**DeepSynBa**	**comboFM**	**comboLTR**	**DeepSynBa**	**comboFM**	**comboLTR**
Breast	**6.67**	11.39	14.05	**0.95**	0.86	0.77	**0.82**	0.63	0.68
CNS	**7.14**	12.35	14.20	**0.96**	0.87	0.81	**0.80**	0.65	0.68
Colon	**8.55**	13.48	17.25	**0.95**	0.87	0.77	**0.82**	0.64	0.71
Leukemia	**11.09**	18.62	22.34	**0.96**	0.87	0.80	**0.86**	0.72	0.77
Melanoma	**7.30**	11.76	15.49	**0.95**	0.87	0.75	**0.81**	0.61	0.68
Non-Small Cell Lung	**7.15**	11.76	15.11	**0.95**	0.87	0.77	**0.80**	0.61	0.68
Ovarian	**6.97**	11.76	15.46	**0.95**	0.85	0.72	**0.80**	0.59	0.66
Prostate	**7.09**	11.73	15.09	**0.96**	0.88	0.78	**0.82**	0.65	0.68
Renal	**6.93**	11.16	14.87	**0.95**	0.87	0.75	**0.82**	0.64	0.70

Table 6.2: Performances based on tissue types with new drug combination scenario.

	RMSE			Pearson			Spearman		
	DeepSynBa	**comboFM**	**comboLTR**	**DeepSynBa**	**comboFM**	**comboLTR**	**DeepSynBa**	**comboFM**	**comboLTR**
Chemotherapy-Chemotherapy	**8.07**	13.03	16.57	**0.96**	0.88	0.80	**0.84**	0.75	0.70
Chemotherapy-Targeted	**8.33**	12.39	16.16	**0.94**	0.86	0.73	**0.80**	0.67	0.59
Chemotherapy-Other	**7.29**	13.48	16.23	**0.96**	0.86	0.79	**0.81**	0.67	0.63
Targeted-Targeted	**6.32**	11.77	15.55	**0.94**	0.79	0.54	**0.78**	0.58	0.44
Targeted-Other	**5.69**	10.45	12.95	**0.94**	0.81	0.68	**0.75**	0.53	0.43
Other-Other	**5.77**	11.29	14.69	**0.97**	0.88	0.79	**0.71**	0.50	0.47

Table 6.3: Performances based on drug combination categories in the new drug combination scenario.

often associated with specific signalling pathways, such as PI3K/AKT/mTOR, which are targets for therapeutic interventions and support more reliable prediction of therapeutic responses [58, 29]. In contrast, leukaemia lacks a solid tumour structure and displays significant genetic heterogeneity, posing challenges for predictive models. [73, 24, 13].

6.4.1.2 Performance Across Drug Categories

We further analyse the results by categorising the drugs in the combinations according to NCI-ALMANAC's classification: chemotherapy, targeted, and other. For two-drug combinations, six distinct pairings emerge from these categories. Performance metrics are computed separately for these pairings. Across all combinations, our model, DeepSynBa, consistently outperforms the competing methods, as demonstrated in Table 6.3. The most accurate dose-response predictions are observed when one drug belongs to the "other" category, paired with either a "targeted" drug or another "other" drug. This phenomenon can likely be attributed to the distribution of responses within these categories. Specifically, the "Targeted-Other" and "Other-Other" combinations show minimal impact on the overall response, as their distributions are concentrated within a narrow range (see Fig. 6.5), making the prediction task more tractable.

6.4.1.3 Synergy Score Prediction Performance

For the final evaluation, we employed the Loewe and Bliss synergy models to assess the performance of each approach. Synergy scores were computed for the ground truth and predicted dose-response matrices across all models. The results, presented in Table 6.4, demonstrate that DeepSynBa consistently outperforms the other drug combination

	Loewe			Bliss		
	RMSE	**Pearson**	**Spearman**	**RMSE**	**Pearson**	**Spearman**
DeepSynBa	6.37	0.93	0.93	7.06	0.52	0.35
comboFM	9.70	0.83	0.83	9.32	0.24	0.15
comboLTR	11.46	0.75	0.76	10.97	-0.01	-0.01
MatchMaker	**5.69**	**0.94**	**0.95**	**5.22**	**0.78**	**0.73**
DeepSynergy	11.10	0.85	0.85	6.77	0.60	0.53

Table 6.4: Performances with Loewe and Bliss synergy scores. Ground-truth Loewe/Bliss scores are calculated with dose-response matrices. Predicted Loewe/Bliss scores for DeepSynBa, comboFM and comboLTR models are calculated with predicted dose-response matrices. MatchMaker and DeepSynergy models are directly trained with Loewe/Bliss synergy scores.

matrix prediction methods in terms of each metric. These findings highlight the superior performance of DeepSynBa relative to alternative approaches in capturing drug synergy, providing robust evidence for its efficacy in predicting synergistic interactions. We also assess the performance of two synergy score predictor models, MatchMaker [41] and DeepSynergy [62], which directly learn synergy scores. Although DeepSynBa is not optimised to predict Loewe or Bliss synergy scores directly, it achieves comparable synergy score prediction performance, especially in the case of Loewe scores.

6.5 Discussions and Next Steps

In this chapter, we introduced DeepSynBa which predicts full dose-response matrices of drug combinations while disentangling efficacy and potency. The model also enables the prediction of full dose-response surfaces. We assessed the performance of our model under three evaluation scenarios: (i) previously unseen drug combinations drug combinations, (ii) new cell lines, and (iii) new drugs. In all settings, DeepSynBa outperforms existing models in predicting dose-response matrices.

By capturing the complex interactions between drug and cell line features with a deep architecture, our model shows the capability of generalising across previously unseen treated cell lines. This is evident from the close performance in RMSE and correlation metrics when comparing new drug combination and new cell line scenario, as shown in Table 1a and 1b. DeepSynBa does not drop in performance when approaching the difficult task of predicting combination effects in cell lines that have not been tested before.

One limitation of our model is in the prediction of previously unseen new drug combinations. Table 1c shows RMSE and correlation metrics degrade compared to the previous two scenarios, Table 1a and 1b. This suggests that while DeepSynBa effectively

utilises cell line specificity, further improvements are needed to enhance its ability to learn drug characteristics to generalise better across previously unseen drugs.

As a secondary analysis, we transformed DeepSynBa predicted dose-response matrices into synergy scores using the Loewe and Bliss synergy models. Despite not being directly trained on synergy scores, we show that DeepSynBa predictions are highly competitive with the models that are designed and trained for synergy score predictions. Although our model has a small drop in performance in predicting traditional synergy score metrics compared to others, the additional insights gained into the dynamics of drug interactions and the tumour indication dependency make this trade-off valuable for real-world applications in drug discovery and experimental design.

By developing DeepSynBa, we seek to encourage a broader perspective beyond using traditional synergy scores as the primary objective for drug combination prediction. Converting a dose-response grid into a single synergy score inherently results in a loss of information given the complexity of drug combinations. Such conversion oversimplifies the interaction profile of combined drugs and hides the potentially key pharmacodynamic patterns across different dosage pairs that can be discovered across the entire dose-response surface, providing a richer representation of drug dynamics. By predicting the full dose-response surface, DeepSynBa can reveal how the effect of one drug changes in response to increasing doses of the other. More importantly, the intermediate layer that learns the parameters of the surface allows a clearer understanding of whether the observed interaction is driven by an enhanced potency or an increased maximal efficacy.

Overall, DeepSynBa supports decision-making processes in drug discovery pipelines by predicting actionable drug combination response profiles. The case studies presented in this work provide an example of how the model provides enchanced interpretability and guides the use of the combination regimes in practical applications.

The next steps would include making the SynBa parameters probabilistic, with the aim of allowing uncertainty estimation for the synergy-related parameters. One potential approach is to introduce a variational distribution for each of the parameters in the SynBa parameter layer so that they are treated as latent variables instead of fixed values. We may try techniques such as the Variational Information Bottleneck (VIB) [3]. This class of method would allow the sampling of a stochastic representation from the variational distribution, which is then used to predict the drug responses and compute the loss.

This chapter shows that the SynBa framework could be extended to address various challenges where decoupling synergy estimation, incorporating uncertainty quantification and providing interpretable pharmacological parameters are beneficial. There is a broad range of tasks where SynBa can be applied to enrich the inference of synergy profiles of combinations.

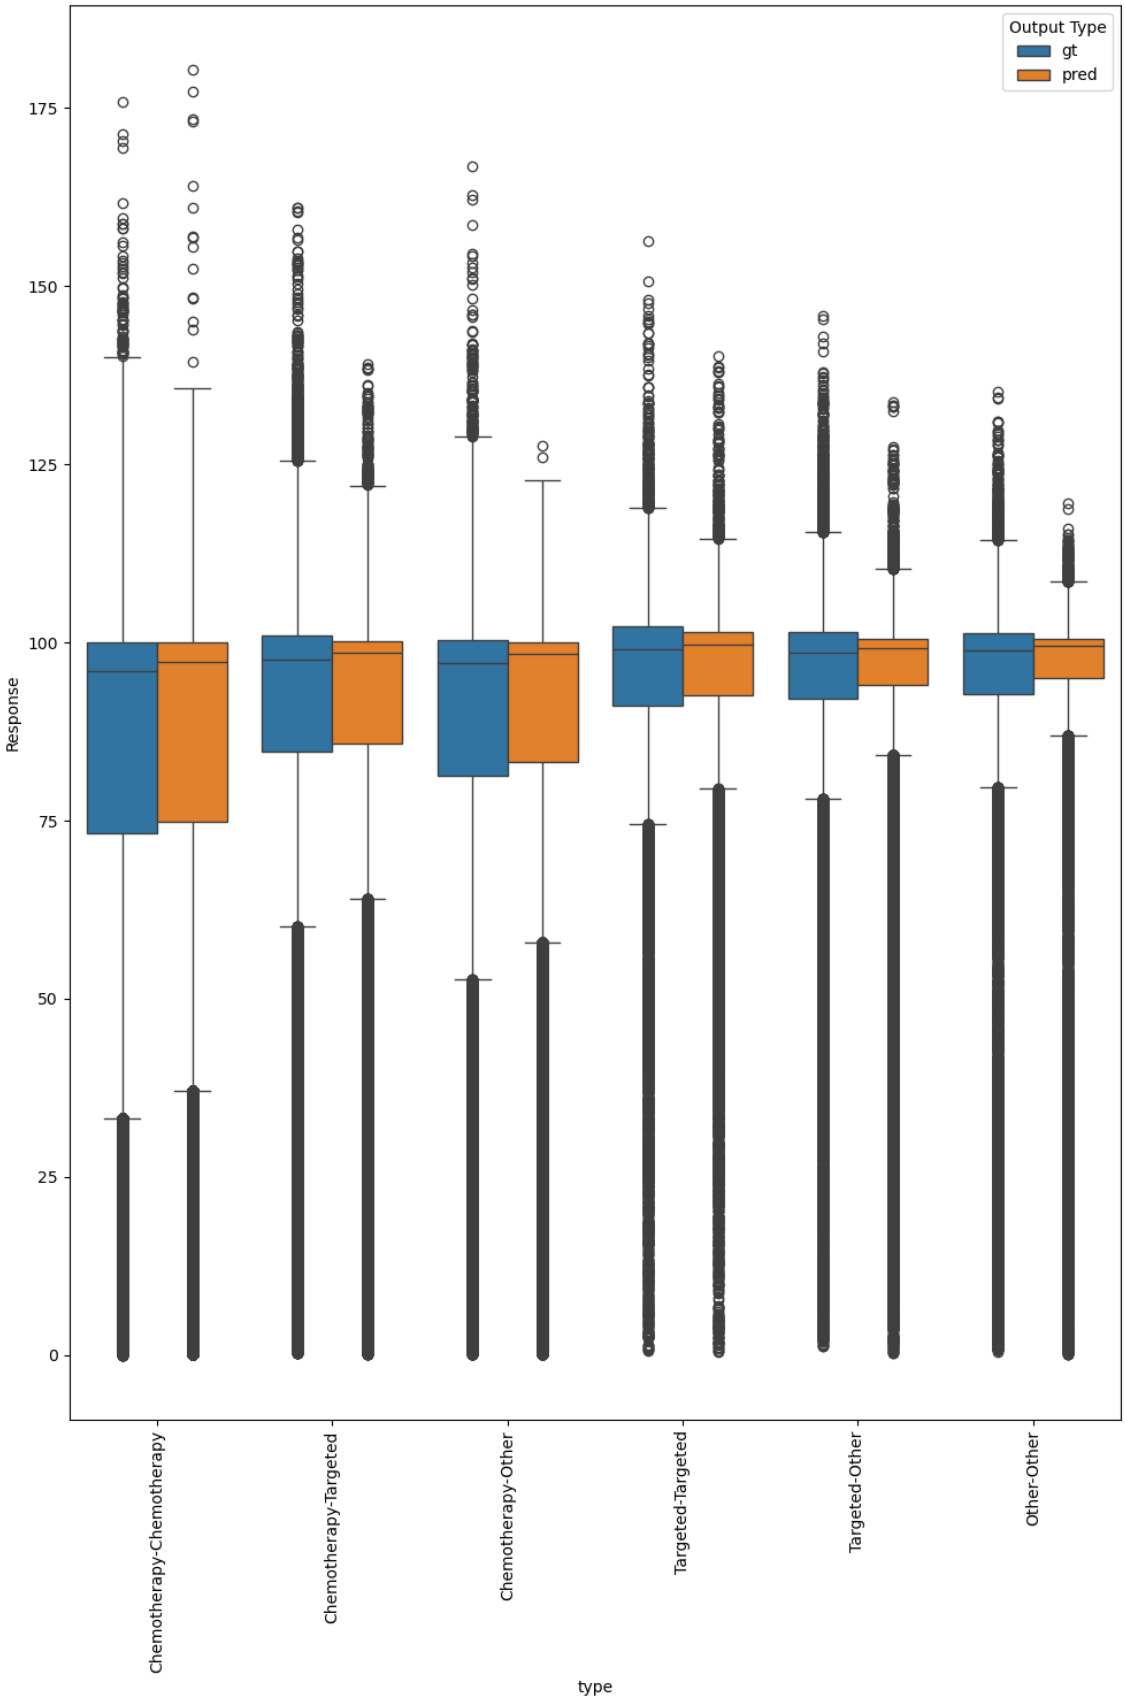

Figure 6.5: The ground-truth and the predicted response distributions for drug types in the combination.

Chapter 7

Conclusions and future directions

In this thesis, we introduced SynBa [86], a new Bayesian framework, and SynBa-Batch, its extension, designed to quantify drug combination synergies while also addressing the uncertainty associated with the pharmacological data. This framework provides potential routes to rank drug combinations based on both efficacy and potency, ultimately supporting the goal of safer and more effective drug treatments. By leveraging the posterior distributions to quantify uncertainty and decoupling the synergistic potency and efficacy, more informed decision-making can be achieved. In addition to introducing SynBa and SynBa-Batch, we developed DeepSynBa [43], a deep learning method that predicts the complete dose-response matrix or surface, achieved by using the SynBa parameterisation to describe the dose-response surface and predicting the SynBa parameters as an intermediate layer in the model.

One core strength of SynBa is its ability to provide posterior distributions for drug synergy parameters. These distributions represent the uncertainty surrounding the predicted efficacy and toxicity of drug combinations. Unlike traditional methods, which output a single, deterministic synergy score, SynBa provides a probabilistic view, allowing biologists or decision-makers to have an idea of how likely a particular drug combination is to be effective or toxic. This uncertainty quantification is particularly valuable when experimental data is limited or noisy, as is often the case in early-stage drug discovery. By providing a range of possible outcomes rather than a single estimate, SynBa can help researchers prioritise drug combinations that simultaneously have high efficacy and low toxicity, enabling a more efficient and safer drug development process.

The insights gained from SynBa are not just useful for ranking drug combinations but can also inform the design of future experiments. For example, the posterior distributions can highlight dosage ranges where uncertainty is highest, directing experimental efforts to areas where additional data collection could significantly reduce this uncertainty. By focusing on areas of the dose-response curve where the uncertainty is greatest, researchers can more efficiently find the most effective and safe doses, reducing the number of

experimental trials required. This has the potential to speed up and improve the safety of the drug discovery and development process.

Moreover, SynBa represents a significant departure from the "synergy score"-based methods that are often used as the response variable in the prediction methods for drug combination studies. These traditional methods reduce complex drug interactions into a single synergy score, a reduction that is too simplistic and can be misleading. As discussed in this thesis, a single synergy score does not adequately capture the multifaceted nature of drug interactions. This is an example showing that one of the main challenges of AI in drug discovery lies in asking the right question and modelling the right endpoint in the first place [6]. This is also the motivation of the development of DeepSynBa. Moving forward, we argue that it is necessary to move away from single score-based methods and develop frameworks like SynBa that depict a more complete picture for the combination profile and simultaneously account for potency, efficacy, and uncertainty.

Despite the advantages of SynBa, one challenge that we have identified in this thesis is the potential overparameterisation in the 1D and 2D Hill equations, which serve as the foundation for many dose-response models, including SynBa. While these models provide interpretability for every parameter involved, making it easy to understand the pharmacological consequence of the parameters, this level of interpretability comes at a cost. As discussed in Chapter 4, the overparameterisation makes it challenging to extend the model to the combination of more than two drugs, as the computational cost would rise exponentially with the number of drugs combined. More importantly, overparameterisation can lead to issues with model identifiability, making it difficult to determine whether the model parameters genuinely reflect underlying biological processes or are simply fitting noise in the data. In cases where data is scarce or noisy, which often occurs in pharmacology, overparameterised models can lead to misleading conclusions, as the model may overfit the data and fail to generalise to new drug combinations or dose ranges. The Hill equation also assumes a functional form that may not always be appropriate for complex biological systems, particularly when considering interactions between multiple drugs.

Given these limitations, future work may need to explore entirely new likelihood models for the dose-response surfaces that are more robust to overparameterisation, without losing the interpretability that comes with the parameters. While nonparametric methods that use Gaussian Processes provide more flexibility, they sacrifice the interpretability of key parameters, making it challenging to communicate the results to biologists. This will limit the actionability of these methods as a decision-support system in practical settings. Balancing flexibility and actionability of such models still an open question. Nevertheless, SynBa has provided a framework for an effective uncertainty estimation for drug combinations that interprets and decouples synergistic efficacy and potency for the

first time. If a likelihood model with better flexibility is developed, it can directly replace the likelihood model in SynBa without altering its core philosophy, i.e. providing robust uncertainty estimation while decoupling synergistic efficacy and potency.

In conclusion, SynBa offers a robust framework for drug synergy assessment by providing a probabilistic approach to uncertainty quantification, improving decision-making and experimental design. While it addresses key limitations of traditional synergy score-based methods, there are also challenges such as overparameterisation, high computational cost and the need for more flexible models. More importantly, while SynBa provides a new framework for drug synergy assessment, its true value lies not only in the specifics of the method itself but also in the broader mindset of estimating uncertainties in a principled manner to aid decision-making and improve safety and efficiency in application. As we refine the technical details and explore new models, it is this philosophy that will continue to drive progress in pharmacology, drug discovery and beyond.

References

[1] Fatemeh Abbasi and Juho Rousu. New methods for drug synergy prediction: A mini-review. *Current Opinion in Structural Biology*, 86:102827, 2024.

[2] George Adam, Ladislav Rampášek, Zhaleh Safikhani, Petr Smirnov, Benjamin Haibe-Kains, and Anna Goldenberg. Machine learning approaches to drug response prediction: challenges and recent progress. *NPJ precision oncology*, 4(1):19, 2020.

[3] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

[4] Valeria Barili, Enrico Ambrosini, Beatrice Bortesi, Roberta Minari, Erika De Sensi, Ilenia Rita Cannizzaro, Antonietta Taiani, Maria Michiara, Angelica Sikokis, Daniela Boggiani, et al. Genetic basis of breast and ovarian cancer: Approaches and lessons learnt from three decades of inherited predisposition testing. *Genes*, 15(2):219, 2024.

[5] Azadeh C Bashi, Elizabeth A Coker, Krishna C Bulusu, Patricia Jaaks, Claire Crafter, Howard Lightfoot, Marta Milo, Katrina McCarten, David F Jenkins, Dieudonne van der Meer, et al. Large-scale pan-cancer cell line screening identifies actionable and effective drug combinations. *Cancer discovery*, 14(5):846–865, 2024.

[6] Andreas Bender and Isidro Cortés-Ciriano. Artificial intelligence in drug discovery: what is realistic, what are illusions? part 1: ways to make an impact, and why we are not there yet. *Drug discovery today*, 26(2):511–524, 2021.

[7] Paul Bertin, Jarrid Rector-Brooks, Deepak Sharma, Thomas Gaudelet, Andrew Anighoro, Torsten Gross, Francisco Martinez-Pena, Eileen L Tang, Cristian Regep, Jeremy Hayter, et al. RECOVER: sequential model optimization platform for combination drug repurposing identifies novel synergistic compounds in vitro. *arXiv preprint arXiv:2202.04202*, 2022.

[8] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer: New York, 2006.

[9] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.

[10] Chester I Bliss. The toxicity of poisons applied jointly 1. *Annals of applied biology*, 26(3):585–615, 1939.

[11] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[12] Krishna C Bulusu, Rajarshi Guha, Daniel J Mason, Richard PI Lewis, Eugene Muratov, Yasaman Kalantar Motamedi, Murat Cokol, and Andreas Bender. Modelling of compound combination effects and applications to efficacy and toxicity: state-of-the-art, challenges and perspectives. *Drug discovery today*, 21(2):225–238, 2016.

[13] Alexander Calderon, Cuijuan Han, Sadik Karma, and Eric Wang. Non-genetic mechanisms of drug resistance in acute leukemias. *Trends in Cancer*, 10(1):38–51, 2024.

[14] Iurie Caraus, Abdulaziz A Alsuwailem, Robert Nadon, and Vladimir Makarenkov. Detecting and overcoming systematic bias in high-throughput screening technologies: a comprehensive review of practical issues and methodological solutions. *Briefings in bioinformatics*, 16(6):974–986, 2015.

[15] Christine Cassidy, Danielle Dever, Laura Stanbery, Gerald Edelman, Lance Dworkin, and John Nemunaitis. Fda efficiency for approval process of covid-19 therapeutics. *Infectious Agents and Cancer*, 15:1–13, 2020.

[16] Ting-Chao Chou. Theoretical basis, experimental design, and computerized simulation of synergism and antagonism in drug combination studies. *Pharmacological reviews*, 58(3):621–681, 2006.

[17] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977.

[18] Vincent T Devita Jr, Robert C Young, and George P Canellos. Combination versus single agent chemotherapy: a review of the basis for selection of drug treatment of cancer. *Cancer*, 35(1):98–110, 1975.

[19] Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[20] Giovanni Y Di Veroli, Chiara Fornari, Dennis Wang, Séverine Mollard, Jo L Bramhall, Frances M Richards, and Duncan I Jodrell. Combenefit: an interactive platform for the analysis and visualization of drug combinations. *Bioinformatics*, 32(18):2866–2868, 2016.

[21] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.

[22] Bradley Efron. Why isn't everyone a bayesian? *The American Statistician*, 40(1):1–5, 1986.

[23] Mohamed Reda El Khili, Safyan Aman Memon, and Amin Emad. Marsy: a multitask deep-learning framework for prediction of drug combination synergy scores. *Bioinformatics*, 39(4), April 2023. ISSN 1367-4811. doi: 10.1093/bioinformatics/btad177. URL `http://dx.doi.org/10.1093/bioinformatics/btad177`.

[24] Elihu Estey, Ross L Levine, and Bob Löwenberg. Current challenges in clinical development of "targeted therapies": the case of acute myeloid leukemia. *Blood, The Journal of the American Society of Hematology*, 125(16):2461–2466, 2015.

[25] George E Forsythe. Von neumann's comparison method for random sampling from the normal and other distributions. *Mathematics of Computation*, 26(120):817–826, 1972.

[26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[27] Alan H Gradman, Jan N Basile, Barry L Carter, George L Bakris, American Society of Hypertension Writing Group, et al. Combination therapy in hypertension. *Journal of the American Society of Hypertension*, 4(2):90–98, 2010.

[28] William R Greco. The search for synergy: a critical review from a response surface perspective. *Pharmacol Rev*, 47:331–385, 1995.

[29] Huifang Guo, Peter German, Shanshan Bai, Sean Barnes, Wei Guo, Xiangjie Qi, Hongxiang Lou, Jiyong Liang, Eric Jonasch, Gordon B Mills, et al. The pi3k/akt pathway and renal cell carcinoma. *Journal of genetics and genomics*, 42(7):343–353, 2015.

[30] Juan-José Giraldo Gutierrez, Evelyn Lau, Subhashini Dharmapalan, Melody Parker, Yurui Chen, Mauricio A Álvarez, and Dennis Wang. Multi-output prediction of dose–response curves enables drug repositioning and biomarker discovery. *npj Precision Oncology*, 8(1):209, 2024.

[31] Archibald Vivian Hill. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *J. Physiol.*, 40:4–7, 1910.

[32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[33] Susan L. Holbeck, Richard Camalier, James A. Crowell, Jeevan Prasaad Govindharajulu, Melinda Hollingshead, Lawrence W. Anderson, Eric Polley, Larry Rubinstein, Apurva Srivastava, Deborah Wilsker, Jerry M. Collins, and James H. Doroshow. The national cancer institute ALMANAC: A comprehensive screening resource for the detection of anticancer drug pairs with enhanced therapeutic activity. *Cancer Research*, 77(13):3564–3576, 2017.

[34] Patricia Jaaks, Elizabeth A Coker, Daniel J Vis, Olivia Edwards, Emma F Carpenter, Simonetta M Leto, Lisa Dwane, Francesco Sassi, Howard Lightfoot, Syd Barthorpe, et al. Effective drug combinations in breast, colon and pancreatic cancer cells. *Nature*, 603(7899):166–173, 2022.

[35] Jouhyun Jeon, Satra Nim, Joan Teyra, Alessandro Datti, Jeffrey L Wrana, Sachdev S Sidhu, Jason Moffat, and Philip M Kim. A systematic approach to identify novel cancer drug targets using machine learning, inhibitor design and high-throughput screening. *Genome medicine*, 6:1–18, 2014.

[36] Wengong Jin, Jonathan M Stokes, Richard T Eastman, Zina Itkin, Alexey V Zakharov, James J Collins, Tommi S Jaakkola, and Regina Barzilay. Deep learning identifies synergistic drug combinations for treating covid-19. *Proceedings of the National Academy of Sciences*, 118(39), 2021.

[37] Heli Julkunen, Anna Cichonska, Prson Gautam, Sandor Szedmak, Jane Douat, Tapio Pahikkala, Tero Aittokallio, and Juho Rousu. Leveraging multi-way interactions for systematic prediction of pre-clinical drug combination effects. *Nature communications*, 11(1):1–11, 2020.

[38] Herman Kahn and Andy W Marshall. Methods of reducing sample size in monte carlo computations. *Journal of the Operations Research Society of America*, 1(5): 263–278, 1953.

[39] Yash Khemchandani, Stephen O'Hagan, Soumitra Samanta, Neil Swainston, Timothy J Roberts, Danushka Bollegala, and Douglas B Kell. Deepgraphmolgen, a multi-objective, computational strategy for generating molecules with desirable properties: a graph convolution and reinforcement learning approach. *Journal of cheminformatics*, 12:1–17, 2020.

[40] Teun Kloek and Herman K Van Dijk. Bayesian estimates of equation system parameters: an application of integration by monte carlo. *Econometrica: Journal of the Econometric Society*, pages 1–19, 1978.

[41] Halil Ibrahim Kuru, Oznur Tastan, and A Ercument Cicek. Matchmaker: a deep learning framework for drug synergy prediction. *IEEE/ACM transactions on computational biology and bioinformatics*, 19(4):2334–2344, 2021.

[42] Halil Ibrahim Kuru, A Ercument Cicek, and Oznur Tastan. From cell lines to cancer patients: personalized drug synergy prediction. *Bioinformatics*, 40(5):btae134, 2024.

[43] Halil Ibrahim Kuru, Haoting Zhang, Magnus Rattray, Carl Henrik Ek, A Ercument Cicek, Oznur Tastan, and Marta Milo. Deepsynba: Actionable drug combination prediction with complete dose-response profiles. *bioRxiv*, pages 2025–01, 2025.

[44] Alexander Lachmann, Federico M Giorgi, Mariano J Alvarez, and Andrea Califano. Detection and removal of spatial bias in multiwell assays. *Bioinformatics*, 32(13): 1959–1965, 2016.

[45] Floriane Larras, Elise Billoir, Vincent Baillard, Aurélie Siberchicot, Stefan Scholz, Tesfaye Wubet, Mika Tarkka, Mechthild Schmitt-Jansen, and Marie-Laure Delignette-Muller. Dromics: a turnkey tool to support the use of the dose–response framework for omics data in ecological risk assessment. *Environmental science & technology*, 52 (24):14461–14468, 2018.

[46] J Jack Lee, Maiying Kong, Gregory D Ayers, and Reuben Lotan. Interaction index and different methods for determining drug interaction in combination therapy. *Journal of biopharmaceutical statistics*, 17(3):461–480, 2007.

[47] Jeffrey T Leek, Robert B Scharpf, Héctor Corrada Bravo, David Simcha, Benjamin Langmead, W Evan Johnson, Donald Geman, Keith Baggerly, and Rafael A Irizarry. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics*, 11(10):733–739, 2010.

[48] Jeffrey T Leek, W Evan Johnson, Hilary S Parker, Andrew E Jaffe, and John D Storey. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics*, 28(6):882–883, 2012.

[49] Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. Explainable ai: A review of machine learning interpretability methods. *Entropy*, 23(1):18, 2020.

[50] Alexander Ling and R Stephanie Huang. Computationally predicting clinical drug combination efficacy with cancer cell line screens and independent drug action. *Nature communications*, 11(1):1–13, 2020.

[51] S Loewe. The problem of synergism and antagonism of combined drugs. *Arzneimittelforschung*, 3:285–290, 1953.

[52] F. J. Massey Jr. The Kolmogorov-Smirnov test for goodness of fit. *J. Am. Stat. Assoc.*, 46(253):68–78, 1951.

[53] Bogdan Mazoure, Robert Nadon, and Vladimir Makarenkov. Identification and correction of spatial bias are essential for obtaining quality data in high-throughput screening technologies. *Scientific reports*, 7(1):11921, 2017.

[54] Michael P Menden, Dennis Wang, Mike J Mason, Bence Szalai, Krishna C Bulusu, Yuanfang Guan, Thomas Yu, Jaewoo Kang, Minji Jeon, Russ Wolfinger, et al. Community assessment to advance computational prediction of cancer drug combinations in a pharmacogenomic screen. *Nature communications*, 10(1):1–17, 2019.

[55] Xiao-Li Meng and Donald B Rubin. Maximum likelihood estimation via the ecm algorithm: A general framework. *Biometrika*, 80(2):267–278, 1993.

[56] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.

[57] Christian T Meyer, David J Wooten, B Bishal Paudel, Joshua Bauer, Keisha N Hardeman, David Westover, Christine M Lovly, Leonard A Harris, Darren R Tyson, and Vito Quaranta. Quantifying drug combination synergy along potency and efficacy axes. *Cell systems*, 8(2):97–108, 2019.

[58] Daniela Miricescu, Daniela Gabriela Balan, Adrian Tulin, Ovidiu Stiru, Ileana Adela Vacaroiu, Doina Andrada Mihai, Cristian Constantin Popa, Raluca Ioana Papacocea, Mihaly Enyedi, Nedelea Andrei Sorin, et al. Pi3k/akt/mtor signalling pathway involvement in renal cell carcinoma pathogenesis. *Experimental and Therapeutic Medicine*, 21(5):1–7, 2021.

[59] Jennifer O'Neil, Yair Benita, Igor Feldman, Melissa Chenard, Brian Roberts, Yaping Liu, Jing Li, Astrid Kral, Serguei Lejnine, Andrey Loboda, et al. An unbiased oncology compound screen to identify novel combination strategies. *Molecular cancer therapeutics*, 15(6):1155–1162, 2016.

[60] Shilpa M Patel and Louis D Saravolatz. Monotherapy versus combination therapy. *Medical Clinics*, 90(6):1183–1195, 2006.

[61] Nancie Petrucelli, Mary B Daly, and Tuya Pal. Brca1-and brca2-associated hereditary breast and ovarian cancer. *GeneReviews*, 2022.

[62] Kristina Preuer, Richard PI Lewis, Sepp Hochreiter, Andreas Bender, Krishna C Bulusu, and Günter Klambauer. Deepsynergy: predicting anti-cancer drug synergy with deep learning. *Bioinformatics*, 34(9):1538–1546, 2018.

[63] Alec Radford. Improving language understanding by generative pre-training. *Preprint*, 2018.

[64] Mihaela Raluca Radu, Alina Prădatu, Florentina Duică, Romeo Micu, Sanda Maria Creţoiu, Nicolae Suciu, Dragoş Creţoiu, Valentin Nicolae Varlas, and Viorica Elena Rădoi. Ovarian cancer: biomarkers and targeted therapy. *Biomedicines*, 9(6):693, 2021.

[65] Christian Ritz, Florent Baty, Jens C Streibig, and Daniel Gerhard. Dose-response analysis using r. *PloS one*, 10(12):e0146021, 2015.

[66] Leiv Rønneberg, Andrea Cremaschi, Robert Hanes, Jorrit M Enserink, and Manuela Zucknick. Bayesynergy: flexible bayesian modelling of synergistic interaction effects in in vitro drug combination experiments. *Briefings in bioinformatics*, 22(6):bbab251, 2021.

[67] Leiv Rønneberg, Paul DW Kirk, and Manuela Zucknick. Dose–response prediction for in-vitro drug combination datasets: a probabilistic approach. *BMC bioinformatics*, 24(1):161, 2023.

[68] Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, 2022.

[69] Jae Yong Ryu, Hyun Uk Kim, and Sang Yup Lee. Deep learning improves prediction of drug–drug and drug–food interactions. *Proceedings of the national academy of sciences*, 115(18):E4304–E4311, 2018.

[70] Yuliya Shapovalova, Tom Heskes, and Tjeerd Dijkstra. Non-parametric synergy modeling with gaussian processes. *bioRxiv*, 2021.

[71] Duxin Sun, Wei Gao, Hongxiang Hu, and Simon Zhou. Why 90% of clinical drug development fails and how to improve it? *Acta Pharmaceutica Sinica B*, 12(7): 3049–3062, 2022.

[72] Jing Tang, Krister Wennerberg, and Tero Aittokallio. What is synergy? the saariselkä agreement revisited. *Frontiers in pharmacology*, 6:181, 2015.

[73] Lu Tang, Zhongpei Huang, Heng Mei, and Yu Hu. Immunotherapy in hematologic malignancies: achievements, challenges and future prospects. *Signal Transduction and Targeted Therapy*, 8(1):306, 2023.

[74] W. Tansey et al. Dose–response modeling in high-throughput cancer drug screenings: an end-to-end approach. *Biostatistics*, 23(2):643–665, 2022.

[75] Stan Development Team. Stan modeling language users guide and reference manual, version 2.29, 2022. URL `https://mc-stan.org`.

[76] Nathaniel R Twarog, Elizabeth Stewart, Courtney Vowell Hammill, and Anang A Shelat. Braid: a unifying paradigm for the analysis of combined drug action. *Scientific reports*, 6(1):1–15, 2016.

[77] Naonori Ueda, Ryohei Nakano, Zoubin Ghahramani, and Geoffrey E Hinton. Smem algorithm for mixture models. *Neural computation*, 12(9):2109–2128, 2000.

[78] Dennis Wang, James Hensman, Ginte Kutkaite, Tzen S Toh, Ana Galhoz, GDSC Screening Team Lightfoot Howard Yang Wanjuan Soleimani Maryam Barthorpe Syd Mironenko Tatiana Beck Alexandra Richardson Laura Lleshi Ermira Hall James Tolley Charlotte Barendt William, Jonathan R Dry, Julio Saez-Rodriguez, Mathew J Garnett, Michael P Menden, et al. A statistical framework for assessing pharmacological responses and biomarkers using uncertainty estimates. *Elife*, 9:e60352, 2020.

[79] Jinxian Wang, Xuejun Liu, Siyuan Shen, Lei Deng, and Hui Liu. Deepdds: deep graph neural network with attention mechanism to predict synergistic drug combinations. *Briefings in Bioinformatics*, 23(1), September 2021. ISSN 1477-4054. doi: 10.1093/bib/bbab390. URL `http://dx.doi.org/10.1093/bib/bbab390`.

[80] Tianduanyi Wang, Sandor Szedmak, Haishan Wang, Tero Aittokallio, Tapio Pahikkala, Anna Cichonska, and Juho Rousu. Modeling drug combination effects via latent tensor reconstruction. *Bioinformatics*, 37(Supplement_1):i93–i101, 2021.

[81] David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids research*, 46(D1): D1074–D1082, 2018.

[82] David J Wooten and Réka Albert. synergy: a python library for calculating, analyzing and visualizing drug combination synergy. *Bioinformatics*, 37(10):1473–1474, 2021.

[83] David J Wooten, Christian T Meyer, Alexander LR Lubbock, Vito Quaranta, and Carlos F Lopez. Musyc is a consensus framework that unifies multi-drug synergy metrics for combinatorial drug discovery. *Nature communications*, 12(1):1–16, 2021.

[84] Roberta J Worthington and Christian Melander. Combination approaches to combat multidrug-resistant bacteria. *Trends in biotechnology*, 31(3):177–184, 2013.

[85] Bulat Zagidullin, Jehad Aldahdooh, Shuyu Zheng, Wenyu Wang, Yinyin Wang, Joseph Saad, Alina Malyutina, Mohieddin Jafari, Ziaurrehman Tanoli, Alberto Pessia, et al. Drugcomb: an integrative cancer drug combination data portal. *Nucleic acids research*, 47(W1):W43–W51, 2019.

[86] Haoting Zhang, Carl Henrik Ek, Magnus Rattray, and Marta Milo. Synba: improved estimation of drug combination synergies with uncertainty quantification. *Bioinformatics*, 39(Supplement_1):i121–i130, 2023.

[87] Wei Zhao, Kris Sachsenmeier, Lanju Zhang, Erin Sult, Robert E Hollingsworth, and Harry Yang. A new bliss independence model to analyze drug combination data. *Journal of biomolecular screening*, 19(5):817–821, 2014.

[88] Anat Zimmer, Itay Katzir, Erez Dekel, Avraham E Mayo, and Uri Alon. Prediction of multidimensional drug dose responses based on measurements of drug pairs. *Proceedings of the National Academy of Sciences*, 113(37):10442–10447, 2016.

[89] Grant R Zimmermann, Joseph Lehar, and Curtis T Keith. Multi-target therapeutics: when the whole is greater than the sum of the parts. *Drug discovery today*, 12(1-2): 34–42, 2007.