



HAL
open science

Machine learning applications in drug development

Clémence Réda, Emilie Kaufmann, Andrée Delahaye-Duriez

► **To cite this version:**

Clémence Réda, Emilie Kaufmann, Andrée Delahaye-Duriez. Machine learning applications in drug development. Computational and Structural Biotechnology Journal, 2020, 18, pp.241-252. 10.1016/j.csbj.2019.12.006 . hal-02533303

HAL Id: hal-02533303

<https://hal.science/hal-02533303v1>

Submitted on 7 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Machine learning applications in drug development

Clémence Réda^{a,b,*}, Emilie Kaufmann^c, Andrée Delahaye-Duriez^{a,d,e,*}



^a NeuroDiderot, UMR 1141, Inserm, Université de Paris, Sorbonne Paris Cité, Hôpital Robert Debré, 48, boulevard Sérurier, Paris 75019, France

^b Université Paris Diderot, Université de Paris, Sorbonne Paris Cité, 5, rue Thomas Mann, Paris 75013, France

^c Univ. Lille, CNRS, Centrale Lille, Inria, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille, F-59000 Lille, France

^d Université Paris 13, Sorbonne Paris Cité, UFR de santé, médecine et biologie humaine, Bobigny 93000, France

^e Service histologie-embryologie-cytogénétique-biologie de la reproduction-CECOS, Hôpital Jean Verdier, AP-HP, Bondy 93140, France

ARTICLE INFO

Article history:

Received 29 September 2019

Received in revised form 10 December 2019

Accepted 10 December 2019

Available online 26 December 2019

Keywords:

Drug discovery
Drug repurposing
Multi-armed bandit
Collaborative filtering
Bayesian optimization
Adaptive clinical trial

ABSTRACT

Due to the huge amount of biological and medical data available today, along with well-established machine learning algorithms, the design of largely automated drug development pipelines can now be envisioned. These pipelines may guide, or speed up, drug discovery; provide a better understanding of diseases and associated biological phenomena; help planning preclinical wet-lab experiments, and even future clinical trials. This automation of the drug development process might be key to the current issue of low productivity rate that pharmaceutical companies currently face. In this survey, we will particularly focus on two classes of methods: sequential learning and recommender systems, which are active biomedical fields of research.

© 2019 The Authors. Published by Elsevier B.V. on behalf of Research Network of Computational and Structural Biotechnology. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A great variety of experimental data, at a chemical, transcriptomic, or genomic-level is available to readily use for drug development. Summarizing the huge amount of biological data at hand into meaningful models, to grasp the full mechanism of diseases, seems harder and harder. However, systems biology and machine learning approaches are continuously enhanced in order to accelerate the path to efficient drug development. We will focus on three significant related and intermingled questions, that can be subject to automation: drug discovery, drug testing, and drug repurposing. Firstly, this review briefly dwells on the current context in drug development. Later, we will review generic machine learning algorithms, and more specifically, we will focus on sequential learning algorithms and recommender systems. These algorithms have also proven themselves useful in other research fields, and are active biomedical fields of research.

1.1. Drug development

1.1.1. Current context in drug development

Development of new drugs is a time-consuming and costly process. Indeed, in order to ensure both the patients' safety and drug effectiveness, prospective drugs must undergo a competitive and long procedure. Drug development is roughly split into four major stages, called phases. Phase 0 comprises basic research/drug discovery and preclinical tests, which aim at assessing the efficiency and body processing of the drug candidate. The last three stages are clinical trials: study of dose-toxicity, short-lived side effects, and kinetic relationships (Phase I); determination of drug performance (Phase II); and comparison of the molecule to the standard-of-care (Phase III). An optional Phase IV can be post-drug marketing to monitor long-lasting side effects and drug combination with other therapies. See Fig. 1 for the whole drug development timeline. This pipeline takes at least 5 years to be completed [1], and can last up to 15 years [2]. The minimal amount of time covers the setup of preclinical and clinical tests (Phases 0 up to III), that is, the time to ponder upon and write down the study design, to recruit and select patients, to analyze the results, and so on, let alone to perform the actual wet-lab experiments.

Clinical development time (that is, from Phase I) has steadily increased. For drugs approved in 2005–2006, the average clinical

* Corresponding author at: NeuroDiderot, UMR 1141, Inserm, Université de Paris, Sorbonne Paris Cité, Hôpital Robert Debré, 48, boulevard Sérurier, Paris 75019, France (A. Delahaye-Duriez).

E-mail addresses: clemence.reda@inserm.fr (C. Réda), emilie.kaufmann@univ-lille.fr (E. Kaufmann), andree.delahaye@inserm.fr (A. Delahaye-Duriez).

IND: Investigational New Drug **BLA:** Biologics License Application
NDA: New Drug Application

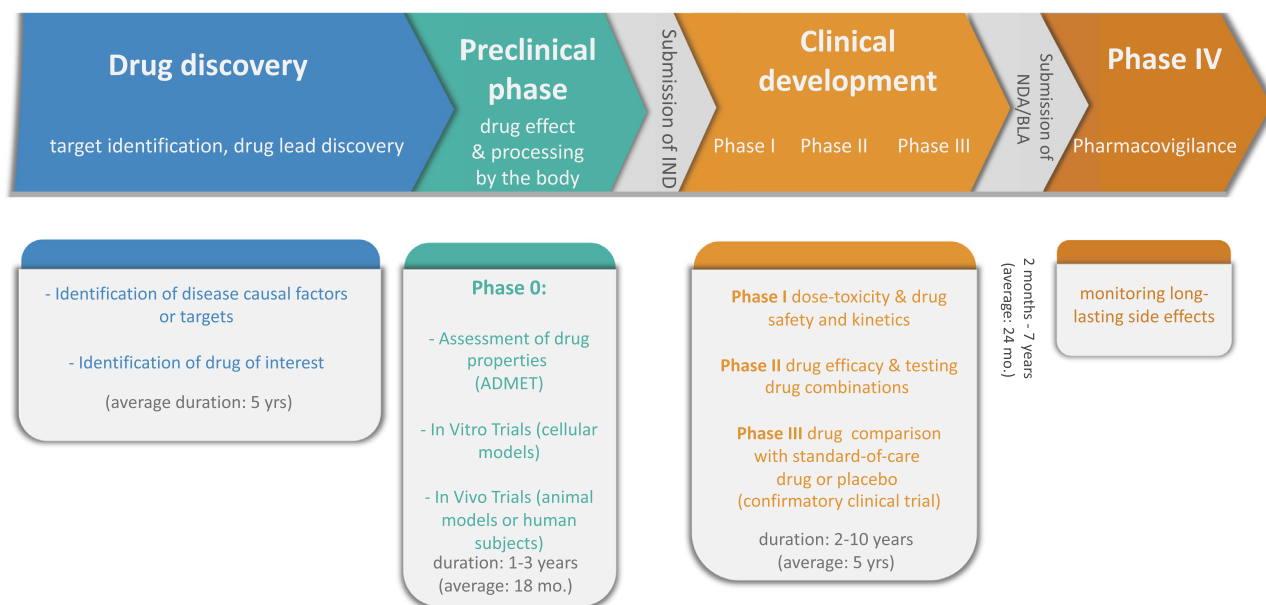


Fig. 1. Representation of the four stages of drug development, along with Phase IV, which occurs after the start of drug marketing.

development time was 6.4 years, whereas it increased up to 9.1 years for 2008–2012 drug candidates [3]. This might denote an issue in assessing drug effects and benefits. Conversely, the high failure rate of drug development pipelines, often at late stages of clinical testing, has always been a critical issue [4]. In clinical trials occurred between 1998 and 2008 (in Phases II and III), [5] have reported a failure rate of 54%. Main reasons for failure were the lack of efficacy (57% of the failing drug candidates), and safety concerns (17%). Among safety concerns were increased risk of death or of serious side effects, which were still the main reasons of failure in Phases II and III in 2012 [3], and in 2019 [7]. For drug pipelines starting in 2007–2009, the gap in estimated success rates in 2012 was particularly steep between Phase II (first patient-dose) and Phase III (first clinical trial-dose). This means that Phase II, which is related to drug performance assessment, is particularly discriminatory: only 14% of the drug candidates that reached Phase II, compared to 64% of the drug pipelines reaching Phase III, were eventually marketed [7]. This can still be observed for drug pipelines starting in 2015–2017 [6]. 25% of the drug candidates that reached Phase II, compared to 62% of the drug pipelines reaching Phase III, were approved (estimation made in 2019).

Meanwhile, total capitalized expected cost of drug development was estimated at \$868 million for approved drugs in 2006, with an average clinical development cost of \$487 million, according to the public Pharma projects database [9]. There are large variations due to drug type (\$479 million for a HIV drug, compared to \$936 million for a rheumatoid arthritis drug) [8]. The recent study of a cohort of 12 large pharmaceutical labs led by [10] shows that total development cost per approved drug has skyrocketed from 2010 to 2018 (from \$1,477 to \$2,168 million), and almost increased two-fold in eight years (from 2010 to 2018) (see Fig. 2). It is worth noticing that, in the meantime, the number of discovered molecules that have reached the late clinical test stage for this cohort dropped by 22% [10]. Moreover, clinical trials pose a barrier to rapid drug development. The cost and time involved in patient recruitment has been increasing. There is a high failure rate and consequent financial loss in product development. Drug efficiency




Fig. 2. Evolution of average development cost in a cohort of 12 major pharmaceutical labs, in millions of dollars, between 2010 and 2018 [10].

assessment might not be carried to term because of prematurely-stopped clinical testing, due to the lack of funding [11]. These factors contribute to the decrease in approved drugs.

In a nutshell, all these figures show an expensive, time-consuming, and frustrating R&D context. Although efforts have been made in order to tackle these issues (as shown by the 2019 figures of success rates between Phases II and III), there is still room for improvement in terms of study planning, as suggested in [11], or designing more insightful preclinical testing [1].

1.2. The future of drug development

This context has indeed transformed the pharmaceutical industry in the span of ten years. Even the biggest pharmaceutical companies encounter productivity issues, in terms of number of

approved molecules with regard to the number of drug candidates [4]. Although a few political efforts have been made to promote orphan disease research [12], this situation has led the pharmaceutical industry to focus on the most profitable diseases. Between 2017 and 2018, the number of active drug pipelines for cancer therapy has increased by 7.6%, whereas the number of anti-infective drugs has dropped by 9.3% [9]. The most studied diseases in 2018, in terms of number of active drug pipelines, are cancer subtypes (breast, lung), diabetes, and Alzheimer's disease [9]. This observation raises the issue of finding therapies for rarer, complex diseases, where the limited number of patients might hinder meaningful studies to be carried on; or for tropical diseases, where the drug development cost might be too prohibitive with respect to the estimated selling profits [13].

As highlighted by many articles [14–16], one rather inexpensive way to improve these numbers might be to automate some important but repetitive data processing and analysis tasks, more especially, through robotics [17] and Machine Learning (ML) methods [18]. Indeed, a lot of blossoming collaborations between Artificial Intelligence (AI) and Machine Learning (ML) companies and pharmaceutical labs, as well as universities and research centres [18–20], slowly bridge the gap in bioinformatics between applied mathematics, computer sciences and biology. This would allow to accelerate drug development pipelines as they might be computationally, thus automatically, performed and less prone to human-related technical mistakes. Authors of [21] estimate their use would shrink the drug candidate identification phase from a few months to one year. Still, one should remain cautious, and not expect computational methods to solve completely the failure rate problem [22]. Nonetheless, integrating ML methods into drug development pipelines might also decrease drug development cost and time [23], and make therapies more patient-oriented, as the easier integration of multiview data might allow implementation or enhancement of precision medicine techniques [24,25]. Conversely, systematic methods allow study replication and reusability, and enable standardized, transparent data quality control and sharing [26], and *in silico* identification of promising targets. These methods could also provide quantitative values to assess and compare the efficiency of candidate molecules, before any wet-lab experiment or preclinical test.

1.3. Towards an automated search for therapies

In order to recontextualize research in drug development, we suggest reading the introductory part of the following referenced papers [13,23,27,28]. Even though a fully-automated drug development pipeline seems out of reach for now [13], the combined efforts from biology, medicine, bioinformatics, computer science and mathematics communities have been spent on improving each part of the drug development pipeline – for example, drug discovery through high-throughput screening (HTS) of drugs via genomic and transcriptomic data [29], genome-wide association studies (GWAS) to uncover new relevant drug targets [30], and increasing application of generic algorithms from machine learning. The use of these methods makes sense in a context where a large quantity and diversity of curated information is available about drugs and their therapeutic indications, disease/aggravating factor targets, disease pathways, and gene/protein regulatory interactions.

2. Sequential learning and recommender systems in machine learning

Machine learning (ML) is a subfield of artificial intelligence (AI) in computer science. Here, a ML algorithm designates any computational method where results from past actions or decisions, or

past observations, are used to improve predictions or future decision-making. ML techniques are now extremely popular in drug development (see [13,27,31] for recent surveys) as they allow automation of highly-dimensional, noisy biological data analysis.

Many different machine learning tasks have been studied, which fall broadly into three categories. The first one is supervised learning, in which the goal is to predict the label of new observations given a large database of labelled examples. Several supervised learning algorithms have been applied in a biological context, such as Support Vector Machines [32] or (Deep) Neural Networks [33]. The second task is unsupervised learning, and it aims at detecting underlying relationships or patterns in unlabeled data. Dimension reduction methods, like Principal Component Analysis (PCA), fall in this class. But other unsupervised problems are also studied in the context of drug development, such as density estimation, clustering (grouping data) or even collaborative filtering. We shall elaborate on some examples below. The third type of task is sequential learning, where algorithms rely on trial-and-error, and iteratively use external observations in order to find the best decision with respect to the environment they interact with.

A large literature has dwelled on the use of sequential learning algorithms, where an agent, that is, a goal-oriented entity interacting with its environment, must make one choice at a time according to previous observations of the environment (from which the input data originate) they are interacting with. While offline –or batch– methods use batches of data in order to learn, online –or sequential learning– algorithms process one data point at a time (receiving a stream of data), and update their prediction or decision accordingly.

Multi-Armed Bandit (MAB) algorithms [34] constitute a popular and versatile family of sequential decision-making algorithms, and were actually motivated by clinical trials, as we shall see. In MABs, a fixed set of actions, called arms, is available. An agent sequentially interacts with the environment by selecting arms, as illustrated in Fig. 3 below. Each arm selection produces some noisy observation, often interpreted as a *reward*. However, the average reward associated with each arm is initially unknown to the agent, and has to be learnt in the process while achieving a certain objective. Typically, this objective could be to discover the most efficient arm(s), that is, the arm(s) with highest average reward, or to maximize the total reward accumulated across iterated arm selections [35]. See the following referenced paper [36] for a comparison between bandit problems.

Recommender systems are closely linked to MABs, as MABs can be used to design sequential recommender systems, see for instance the following referenced examples [37,38]. Recommender systems actually belong to different families of ML methods, since a recommender system broadly designates an algorithm which aims at predicting rating of a given user which tests a given object. Refer to [39] for a review of the topic. A large part of the literature about recommender systems is motivated by commercial purposes, see for instance [40–42]. However, we will show that this flexible class of algorithms can actually be applied to solve drug development-related problems.

In the next section, we will review interesting applications of ML in a pseudo-chronological order of appearance in a drug development pipeline –namely, drug discovery, drug repurposing and drug testing– in which established ML algorithms of the three described classes are involved. We will focus on a subset of ML methods, which comprises sequential learning algorithms and recommender systems. Although they are rarely reviewed in a biomedical setting, they have been investigated for tackling drug development-related problems.

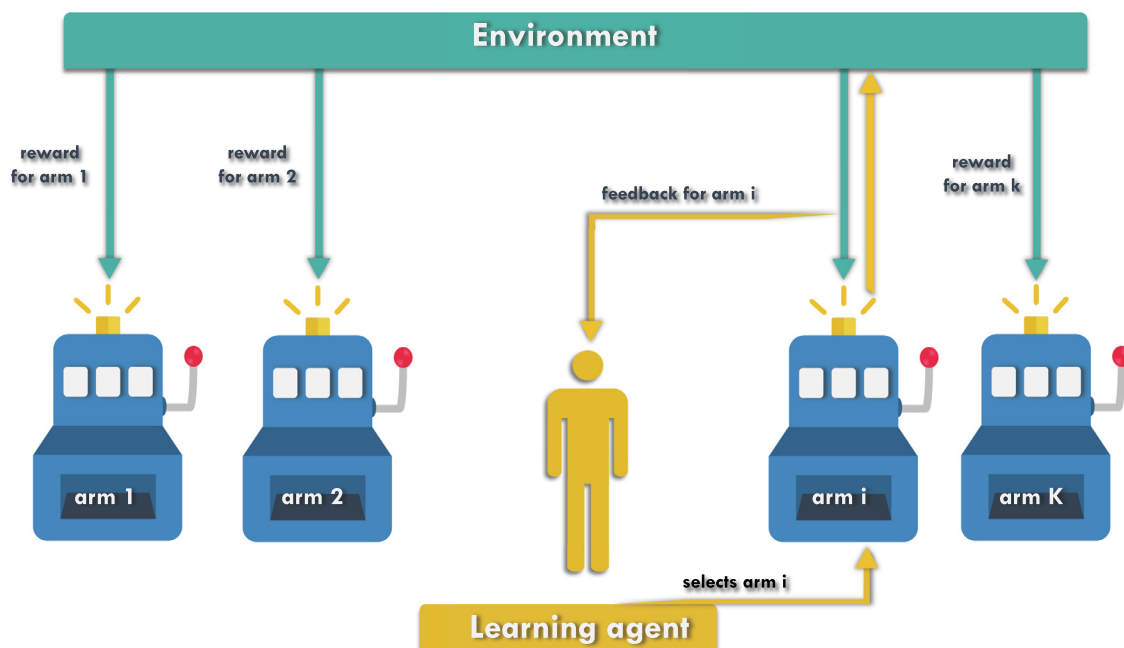


Fig. 3. A K-armed bandit, where the learning agent interacts with its environment by sequentially selecting arms, and updating its strategy using the observations it obtains.

3. Examples of machine learning applied to drug development

3.1. Drug discovery

Drug discovery is usually considered the first stage of a drug development pipeline [43], and is an exploratory step which aims at uncovering putative drug candidates or gene targets, or causal factors, of a given disease or a given chemical compound. A variety of supervised learning methods (for instance, Support Vector Machines and Deep Learning [44–46], regression methods [47,48]) and unsupervised learning methods [13] applied to biomedical problems have been thoroughly reviewed in the last decade, with a growing interest in Deep Learning (DL). For further review for DL methods applied to drug discovery, please refer to [33,49]. Applications may tackle interesting problems in drug discovery: for example, drug candidate identification via molecule docking, in order to predict and preselect interesting drug-target interactions for further research [43]; and protein engineering, that is, *de novo* molecular design of proteins with specific expected binding or motif functions [44].

A fairly recent breakthrough in protein design uses generative DL, more precisely, Generative Adversarial Networks (GANs [50]). A GAN is made of two simultaneously trained neural networks (NN) with distinct roles: a Generator, which is trained to sample instances (“generated instances”), and a Discriminator. The goal of the latter is to recognize training instances from generated ones, by assigning them a probability value of the considered instance being sampled from the training set. The objective is to train the Generator to create fake instances that are able to “fool” the Discriminator, that is, which are convincing enough with respect to the underlying “goodness” function represented by the Discriminator. See Fig. 4 which illustrates a GAN. For instance, [51] have applied GANs to generate DNA sequences matching specific DNA motifs, that is, short DNA sequence variants which are associated with a specific function. As a proof-of-concept, highly-rated samples obtained via this procedure, exhibiting one or several copies of the desired motif, were shown. Another paper [52] uses, as “predictor” network, a NN which predicts the probability of the (generated) DNA sequence of coding for an antimicrobial peptide (AMP),

and have succeeded in training a GAN which returns 77.08% of the time AMP-coding sequences.

Drug discovery problems have also motivated research in black-box optimization, mostly using Bayesian Optimization (BO), see for instance the following referenced examples [53,54]. Bayesian Optimization is a field of research for finding global optimum (i.e., either maximizers or minimizers) of a black-box function by sequentially selecting where to evaluate this function. Indeed, the so-called black-box objective function is accessible only through its values at selected points, and might be costly to evaluate. The interest in a sequential strategy is to adaptively choose where to collect information next (which defines a so-called acquisition function). The motivation for viewing some drug discovery tasks as a black-box optimization problem comes from the fact that they usually rely on expensive simulations (at chemical-level), for instance, protein-folding [55], making automatic drug screening and assessment via these operations time-consuming.

The specificity of BO is to choose a prior probability distribution (embedding *a priori* knowledge, simply called prior) on the objective function: e.g., one assumes that the objective function is a sample of a given probability distribution over functions. This prior is then updated after each new function evaluation into a posterior distribution, which in turns guides the selecting of the new point to evaluate. For instance, in [56], the authors apply BO to design gene sequences which maximize transcription and translation rates, from initial sequences. The optimization is performed on feature vectors of fixed length associated with the gene sequences, and the objective function f is the function which, given a gene sequence feature vector, returns the associated transcription and translation rate functions. The prior upon the objective function is a Gaussian Process, classically used in BO [57]. The authors then use as acquisition function the average objective, which will maximize the average of the transcription and translation rates. The feature vector maximizing this acquisition function will then define an optimal gene design rule (for instance, frequencies of amino acid codons) for the maximization of both transcription and translation rates. Since several codons can code for the same amino acid, given a set of sequences coding for a protein of interest, these sequences can be ranked according to the similarity of their




Fig. 4. Generative Adversarial Networks for Drug Discovery. A Generative Adversarial Network is a set of two neural networks, the Generator and the Discriminator. These two networks are trained at the same time.

corresponding feature vector with the optimal gene design rules that have been derived. The idea is to produce then the protein of interest with lower costs, since the protein production rate is maximized.

Another example of applications of BO to drug discovery is to stimulate the discovery of new chemical compounds [58], that is, finding small molecules that might optimize for a property of interest, while being chemically different from known compounds. This approach might cope with the caveats of previous methods [59].

3.2. Drug testing

Once one or several drug candidates are selected, preclinical (Phase 0) and clinical development (Phases I up to III) start. Drug properties, related to body processing of the candidate molecule, should be assessed in early phases: e.g., the ADME properties: absorption, distribution, metabolism and excretion, along with the toxicity levels. Evaluation of their efficacy is performed later during Phases II and III. Automation and development of *in silico* prediction models might save time and money on later testing stages, and subsequent *in vitro* and *in vivo* experiments. Moreover, these methods might come to sometimes replace experiments on animal models, since some agencies start banning animal testing [60]. Such instances of potential *in silico* guiding of wet-lab experiments can be found in [61–63]; in particular, [62,63] describe the application of Bayesian multi-armed bandit methods to Phase I studies.

For instance, in [61], a graph-based framework is developed in order to build a prediction model for perturbation experiments, provided time-series expression data and putative gene regulatory interactions between genes of interest. The resulting model can then predict expression levels for the selected set of genes after Knock Out and Over Expression perturbations. This method might help assessing the drug effects on pathways after perturbation of the drug targets. The authors have validated their method on a mouse pluripotency model, with *in vitro* experiments, and report that 60.7% of the predicted phenotypes could be reproduced *in vitro*.

Moreover, given the intrinsic sequential nature of a clinical trial, in which patients are given treatments, one (group of) patient(s) after the other, MAB algorithms would be natural candidates to be used in further phases of drug testing. However, nowadays, the motivation for developing new bandit algorithms has entirely

shifted to applications to online content optimization, such as sequential recommender systems [42]. Bandit designs appear to have been seldomly used in human clinical trials. Traditional randomized clinical trial (RCT) (where, at the beginning of the trial, each patient from the pool of subjects is randomly assigned to and treated with a random treatment until the end of the clinical trial) have been the gold standard since the 1960's. Notwithstanding, the past years have witnessed an increased interest in all kind of Adaptive Clinical Trials (ACT), in which the next allocated treatment could be dependent of the outcome of previously allocated treatments. We will now elaborate on the latter. See Fig. 5 for a comparison between RCTs and ACTs.

Seen as a MAB, a Phase III clinical trial proceeds as follows: given K treatments (arms), where each of them has an unknown probability of success p_1, p_2, \dots, p_K , one of the treatments I_t is chosen for the t^{th} patient, and its efficiency (the associated reward value) is subsequently measured. Under the simplest MAB model for a clinical trial, a reward $X_t = 1$ is obtained if the treatment is successful, and $X_t = 0$ if the treatment fails. The allocation is adaptive in that the selection of I_t may depend on $I_1, X_1, I_2, X_2, \dots, I_{t-1}, X_{t-1}$, that is, on the previously given treatments and their observed outcomes. This adaptivity could lead to a smaller sample size to attain a given power, or early stopping for toxicity or futility of a treatment [64], by automatically performing interim analyses [65]. The problem of maximizing rewards in such a bandit model was extensively studied from the 1950's, either from a frequentist view [66,67] (where the success probabilities are treated as unknown parameters to be inferred); or from a Bayesian view [68,69] (where success probabilities are assumed to come from some prior distribution, similarly to Bayesian optimization). Maximizing rewards amounts to maximizing the number of cured patients, which is (arguably) not the purpose of a clinical trial. Yet, the statistical community has also looked at the different problem of finding, as quickly and accurately as possible, the treatment with the largest probability of success. This was studied for example under the name "ranking and selection" [70,71] and later best arm identification [72,73]. An interesting take-out from the bandit literature is that the two objectives of treatment identification and curing patients cannot be achieved (optimally) by the same allocation strategy [74].

However, as attractive as the idea of determining the best treatment while treating properly as many patients as possible [75] is, the use of ACTs remains quite rare. This might be due to intrinsic differences between traditional clinical trial methodology and

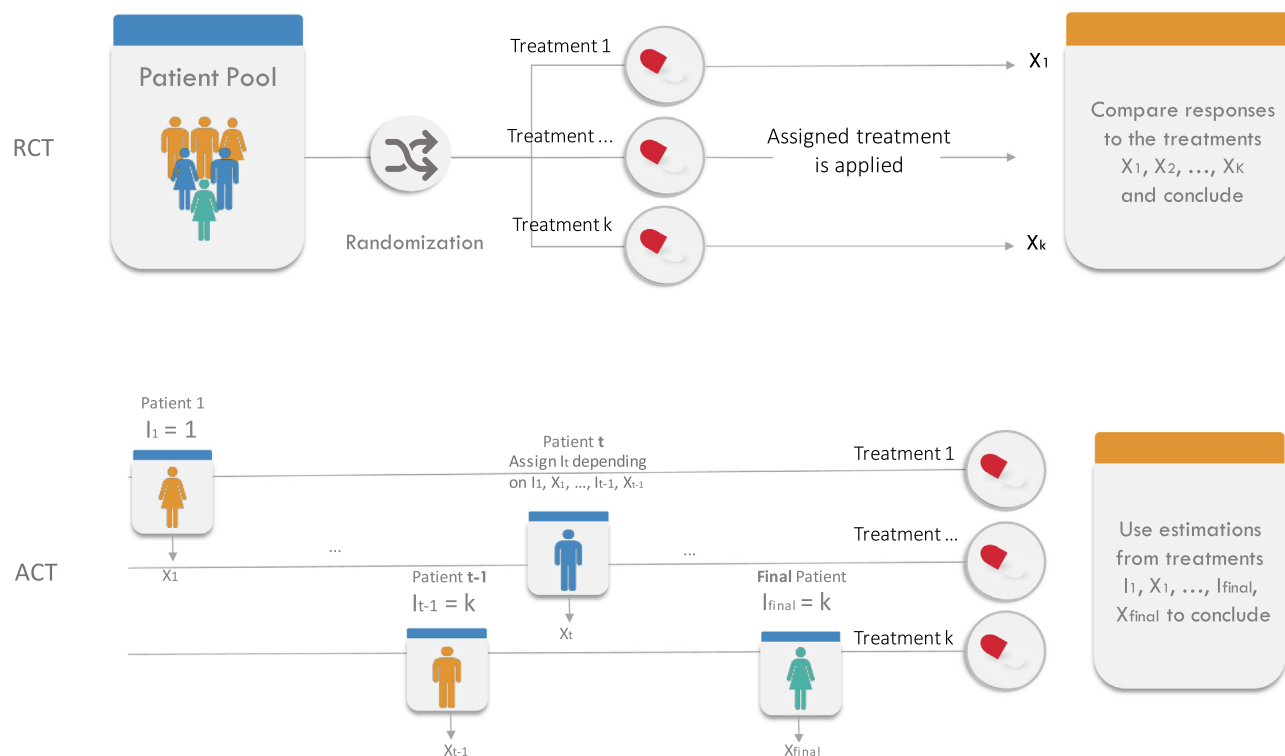


Fig. 5. Randomized Clinical Trial (RCT) versus Adaptive Clinical Trial (ACT) for Phase III. A Randomized Clinical Trial (RCT) “randomly” assign patients to treatment arms (ensuring balance of covariates of interest) before testing, whereas an Adapted Clinical Trial sequentially assigns patients to treatment arms according to previous testing results.

data-dependent allocation, as suggested by [124]. For example, balance between prognostic covariates in each treated group of patients is required for statistical relevance [76]. The use of ACTs might also be hindered by practical reasons in the clinical trial setting, for instance, when one must deal with considerably delayed feedback, as reported by [77]. Nonetheless, these drawbacks might be mitigated by the benefit-risk ratio of the treatment. For life-threatening diseases, an adaptive clinical trial might be a hope for the patients to improve their condition [78], for best empirical treatments, using previous observations, can be assigned to patients.

Despite this initial hostility to ACTs, there has been a recent surge of interest in promoting their actual use. As a notable sign of this evolution, the FDA has updated a draft of guidelines concerning adaptive clinical trials [79], listing in particular several concrete examples of successful adaptive trials. Similarly, [124] presents some examples and promotes some good practice for using ACTs. In the meantime, the authors in [77] have performed multiple simulations illustrating the characteristics of usual bandit algorithms (mostly aimed at maximizing rewards) in terms of allocation and final selection (statistical power), in order to popularize their use.

Among existing bandit algorithms, Bayesian algorithms have achieved a certain popularity. Interestingly, the very first bandit algorithm can be traced back to the work of Thompson in 1933 [35], who suggests randomizing the treatments according to their posterior probability of being optimal. This principle (now sometimes called Thompson Sampling or posterior sampling) was rediscovered around ten years ago in the ML community for its excellent empirical performance in complex models [80]. Interestingly, the use of (variants of) this principle appears to have been proposed as well in different context of drug testing. For example, for phase II trials [81] present a compromise between RCT and Thompson Sampling, while [64] present a proof-of-concept of a phase II trial for the Alzheimer disease that rely on such posterior

sampling ideas. More broadly, several examples of successful Bayesian adaptive designs have emerged over the last 20 years, and we refer the reader to [63,82] for a survey.

Furthermore, research is ongoing in order to tackle for instance the issue of delayed feedback, see for instance [83]. ACTs might allow clinical trials to take explicitly into account inter-patient variability [84], or additional information about the patient [85].

3.3. Drug repurposing

The challenges of designing new molecular entities, and testing them through all clinical phases, has generated research interest in a more profitable and efficient technique, called drug repurposing, or drug repositioning. This approach aims at studying already available drugs and chemical compounds to find them new therapeutic indications. This strategy is useful when repurposed drugs have well-documented safety-profiles (that is, side effects and treatments are known), and known mechanism of action.

Different approaches have been used to tackle the drug repurposing problem. For example, some rely on automatic processing of Electronic Health Records (EHR), clinical trial data, and text mining methods to identify correlations between drug molecules and gene or protein targets in literature [86–88]. However, this approach might be sensitive, but not really specific, since text interpretation is still a hard problem, and the relationship between disease factors and drugs might not be clear. The current state-of-the-art methods seem to have turned to different paradigms of repurposing, see for instance the following reviews for a classification of these different methods [27,89,90].

However, most of these methods rely on a rather strong hypothesis, which is that similarity between elements – for instance, chemical composition of drug molecules – implies correlation at therapeutic effect level, or at drug target level. Nonetheless, counter-examples to this hypothesis have been shown to

lead to disastrous events: for instance, thalidomide exists as two chiral forms (same chemical composition but having mirrored structures). One of these forms can treat morning sickness; the other form can have teratogen effects [91].

An attempt to quantify more accurately drug effects is signature reversion, also called connectivity mapping, which focuses on expression measurements: given a pathological phenotype (“query signature”) associated with the disease at study, the objective is to identify which treatments are most able to revert this signature. This operation is performed via comparisons of the query signature with so-called drug signatures, that is, vectorized summaries of genewise expression changes due to the considered drugs. This approach has recorded some successes in drug repurposing; see for instance [29] for a comprehensive review of this type of method. However, note that relying only on transcriptomic measurements to understand the mechanism behind a disease might lead to wrong directions when these cannot account for its causal factors. Moreover, drug repurposing procedures that directly use drug signatures extracted from LINCS database [92], either to compute a similarity measure, or to predict treatment effects at transcriptome level, are biased: indeed, the drug signatures that have been computed in LINCS are measured in cancerous or immortalized or pluripotent cells – which (post) transcriptomic regulation might differ from healthy cells.

Furthermore, especially in matrix factorization and some deep learning methods that perform a dimension reduction or, more generally, feature learning on the drug signatures, the learned features often hardly make sense biologically speaking, which prevents easy interpretation of the results and sanity checks.

Another way of solving the drug repurposing problem, as emphasized by [27], is to see it as a recommender system problem, where an agent should “recommend” best available options. Here, the agent should select the most promising drug candidates with respect to the disease or the target at study. Some recent papers have adopted this approach: for instance, in both [89,93], the authors have designed a graph-based method to predict drug-target or drug-disease interactions. Given a drug, the model predicts a list of fixed length, that contains disease-related targets most likely affected by the chemical compound.

In [89], the algorithm tackles the problem of predicting drug-target interactions (DTI). It relies on a given input bipartite graph of drugs on the one side, and disease-associated gene targets on the other side, which adjacency matrix is denoted A , of size $n \times m$ (n is the number of targets, and m the number of considered drugs). An edge connects a drug and a target if and only if the drug targets this gene (meaning, for a pair of nodes (i, j) , $A(i, j) = 1$ if and only if j is a drug targeting gene i , else $A(i, j) = 0$). In a recommender system point of view, the goal of the algorithm is to determine how probable (high) the considered drug (user) will target (rate) a given gene (object). The authors of [89] compute a weight matrix W , of size $n \times n$, which depends on A and on drug-drug and target-target similarities (importance of each type of similarity might be parametrized): W_{ij} is the coefficient corresponding to the probability that a drug will also target j knowing that it targets gene i . In order to make the inference about the missing edges between drugs and targets, one then computes matrix $R = WA$.

4. Use of data relevant for drug development

In the next section, we give a summary of the fairly new data types that are openly available as of 2019, and that might be useful in regard to drug development challenges, as many datasets can be integrated to training, validation and feature data for the related algorithms. Indeed, what allows ML techniques to really be efficient is publicly available, curated, annotated data. Multiple types of datasets might be relevant with respect to drug development and drug repositioning questions: information about drug candidates, that are, for instance, the chemical structure of the active molecule, disease gene/protein targets, mechanism of action of drugs, but also their documented side effects. One might also be interested in deducing interesting drug candidates by comparing pairs of diseases, of drugs, of protein targets, and applying the principle of “guilt-by-association”. Integrating multi-view data in a drug development method has been shown to increase its accuracy [94–96]. In this section and in Table 1, we will review currently a few publicly available datasets according to their type which are of interest in a drug development pipeline.

Table 1

List of datasets that are relevant for drug development, ordered according to their type. *When provided by the contributors to the database.

| Data Type | Description | Databases (name, reference, date of last data update*, URL, size*) | API |
|--------------|---|--|-----------|
| Genomic data | (1) Compilation of disease-gene associations; different species are represented in CTD, while the other two databases refer to the human. In CTD, some interactions are manually curated instead of being computationally inferred. OpenTargets and DisGeNET gather data from several curated sources. All of these databases provide a coefficient for each disease-gene association quantifying its corresponding level of evidence. | OpenTargets [97], (2019–11) https://www.opentargets.org/ 27,069 targets \times 13,579 diseases | Yes |
| | | Comparative Toxicogenomics Database (CTD) [98], (2019–11) https://ctdbase.org/ Curated: 8,637 \times 5816 Inferred: 48,634 \times 3168 | Yes |
| | | DisGeNET [99], (2019–07) http://www.disgenet.org/ 17,549 targets \times 24,166 diseases/traits | Yes |
| | (2) SNP reporting; COSMIC reports expert manually-curated data. | COSMIC [100], (2019–09) https://cancer.sanger.ac.uk/cosmic 1,207,190 copy number variants 9,197,630 gene expression variants 7,929,832 differentially methylated CpGs 13,099,101 non coding variants | Yes |
| | (3) Regulatory system (e.g., cis-regulatory modules) data: in CisView, the focus is on the mouse (<i>Mus musculus</i>), and data is collected using a TF binding motif analysis on ChIP-seq experiments. It reports several measures of interest, such as conservation scores and quality assessment of the inferred bindings. UK BioBank collects various types of information (genomics, imaging) in a huge anonymous human cohort (around 500,000 people). | CisView [101], (2016–12) https://lgsun.irp.nia.nih.gov/geneindex/cisview.html UK BioBank [102], (2019–09) https://www.ukbiobank.ac.uk/ | No Yes |

(continued on next page)

Table 1 (continued)

| Data Type | Description | Databases (name, reference, date of last data update*, URL, size*) | API | |
|---------------------------|---|--|--|-----|
| Interaction data | (1) Protein-protein or pathway information; STRING reports PPIs (protein-protein interactions) for thousands of organisms, classified according to their level of evidence: computationally inferred (via functional enrichment analysis), experimentally-proven or extracted from curated databases. A score combining all this information is associated to each PPI. KEGG gathers manually assembled biological (signaling and metabolic) pathways. (2) Biological models of gene and pathway interactions; CausalBioNet collects manually curated rat, mouse and human models which are machine readable (encoded into BEL language, convertible into SBML). BioModels lists literature-based (some of them being manually curated) models, and computationally inferred ones, mostly in SBML format. (3) Drug signatures (genewise expression changes due to treatment) in human immortalized cell lines, from standardized experiments. CMap is a preliminary version of LINCS L1000, and is not supported anymore. | STRING database [103], (2019–01) https://string-db.org/ 24,584,628 proteins and 3,123,056,667 interactions | Yes | |
| | | KEGG Pathway database [104] (2019–11) https://www.genome.jp/kegg/pathway.html | Yes | |
| | | Causal BioNet [105] http://causalbionet.com/ BioModels [106], (2017–06) https://www.ebi.ac.uk/biomodels/ Manually curated: 831 models Literature-based: 1640 models | No | |
| | | Connectivity Map (CMap) [107] https://portals.broadinstitute.org/cmap/ 1309 compounds × 4 cell lines × 154 concentrations LINCS [92] https://clue.io/lincs 51,423 perturbation types 2570 cell lines 4 doses | Yes | |
| Drug-Disease associations | These databases provide information about disease potential therapeutic targets, along with interacting chemical compounds. PROMISCUOUS reports text-mining (from literature) based associations, however some of the texts are manually curated. | Therapeutic Target Database (TTD) [108], (2019–07) http://bidd.nus.edu.sg/group/cjttd/ 3419 targets × 37316 drugs | No | |
| | | PROMISCUOUS [109] http://bioinformatics.charite.de/promiscuous/ 10,208,308 proteins × 25,170 compounds | No | |
| Clinical trials | Repositories of clinical trial settings, status, and results. ClinicalTrials.gov is a large database which mostly collects information about US-located trials (formatted in XML), whereas RepoDB provides visualization and data querying. Clinical trial data is a good source of information for Machine Learning methods, because it lists negative results as well (that is, drugs that failed to prove to be of use in treatment), and potentially the reasons for failure. | RepoDB [110], (2017–07) http://apps.chiragjgroup.org/repoDB/ 1571 approved drugs × 2051 diseases | No | |
| | | ClinicalTrials.gov https://clinicaltrials.gov 323,890 studies | Yes | |
| Chemical & Drug data | (1) Protein-related; automatic annotations. | UniProt [111], (2019–11) https://www.uniprot.org/561,356 proteins (Swiss-Prot dataset)181,787,788 proteins (TrEMBL) | Yes | |
| | | (2) Drug-related; comprises approved, withdrawn drugs, as well as tool chemical compounds, and reports their potential indications. | Drug Bank [112], (2019-07) https://www.drugbank.ca/ 13,450 drugs | Yes |
| | | (3) ADMET drug properties (among other types of relevant drug information). | ChEMBL [113], (2018-12) https://www.ebi.ac.uk/chembl/ 1,879,206 compounds × 12,482 targets | Yes |

4.1. Data types

We mainly focused on datasets which are publicly available online, where the associated data is either easy to download and automatically read, or easy to get access to via an API (Application Programming Interface). This criterion is crucial as it allows data to be easily integrated into a computational method. We required as well that they contain high-quality data – meaning, expert-curated or processed in a relevant way that discards purely correlative assumptions and provides a confidence score – in large quantities, so to avoid having to gather data from different sources which are preprocessed in different, non comparable, ways. Moreover, whenever possible, we also focused on databases which benefited from recent updates (less than one year), in order to ensure that they are still maintained and relevant.

4.2. Feature selection

In ML and statistics, feature selection aims at trimming or transforming (raw) input data in order to only feed valuable information to the prediction model. This step should not be considered optional when designing ML algorithms. If not applied to the input

data, the algorithm might actually learn artifacts, be biased, or even only learn rubbish. As such, feature selection is of paramount importance in order to ensure study replication and to guarantee that the developed prediction model will be useful [114]. Feature selection allows data denoising and thus reducing the batch effect.

This step can either be manually performed, by carefully selecting biologically-relevant features, and then proceed to test them sequentially and assessing their usefulness; or either automatically, by designing an algorithm which will learn these useful features by itself.

The advantage of the manual way is that good model interpretability usually follows, provided one can quantify from the trained model the strength of the influence of each feature (for instance, it is ensured when using linear regression models). Multiple ML methods exist in order to select a subset of preselected features that satisfies interesting properties, such as being predictive of the expected outcome, and being non redundant. For instance, in a drug development setting, statistical methods, random forests, and gradient boosting algorithms helped to successfully predict therapeutic targets from selected features linking genes and diseases of the Open Targets platform [97], or based on gene expression profiles of the LINCS database [92]. The algo-

gorithms used for feature selection are usually classified into three major categories: filter, wrapper and embedded methods, where the latter is a hybrid of the former two. See [115] for a comprehensive introduction to feature selection.

Automatic feature selection algorithms can take advantage of the great flexibility and learning power of Deep Neural Networks, that, provided the raw input data and the expected outcome, create the most discriminant features [116,117]. For instance, the deepDR approach developed in [116] uses an Auto-Encoder (AE) to generate informative features from heterogeneous drug-related data, in order to predict new drug-disease pairs. An AE comprises of two neural networks, one Encoder, and one Decoder. The Encoder will project the raw data onto a latent space of features, such that the Decoder is able to reproduce the expected outcomes given the feature vectors associated with the raw input. However, as in any application of Deep Learning, careful training and regularization of these networks should be performed in order to ensure the relevance of the learned features.

5. Perspectives

5.1. Seeing the organism as a whole system: integration of system biology-related methods to drug development

The power of systems biology and network-based approaches comes from the analysis of multiple genes in functionally enriched pathway, as opposed to traditional single gene and single target approaches. Integration of system biology-related methods to drug development has been implemented for epilepsy in [118], and has allowed the identification of drug target candidates in a systematic way [29], and of a gene module which global expression is highly anti-correlated to epileptic phenotypes [118]. A whole set of genes (a gene module) can be targeted for treatment instead of screening drugs against a single relevant target. Indeed, papers have emphasized on the importance of small-effect gene in a system biology model, as their belonging to highly interconnected gene regulatory networks (GRNs) implies that any slight perturbation on these genes might impact significantly “core” disease genes [24]. A large and active literature [119] has emerged about the formalization, the building and the validation of such GRNs, along with the identification of gene modules highly correlated with pathological phenotype. As GRNs are assumed to mirror gene activity with regard to other genes’ expression, building them usually require (time-series) expression data. These data can be extracted from databases recording measurements of expression after genewise perturbations – for instance, in [120], which relies on Knock Down gene expression measurements collected in LINCS. Even if not yet described today, it can be expected that future methods will take advantage of the use of a restricted part of a GRN to predict *in silico* the effect of a chemical compound on a set of genes of interest.

5.2. Precision, or personalized, medicine

Papers have underlined the variability between patients in terms of disease outcome, drug side effect or even drug action [121] due to genomic variation. Precision medicine aims at tailoring a therapy for a specific patient, by taking into account their transcriptomic profiles, genotype, somatic mutations, etc. [25]. If the integration of multi-view data often allows the prediction model to be more accurate [94,95], it raises the issue of processing, denoising high-dimensional and heterogeneous data (and how to perform feature selection in this case). Practical issues must also be faced, as the data needed to run the model might not be routinely obtained at any hospital, as noticed in [114].

6. Summary and outlook

A few takeaway messages can be highlighted from the vast literature about drug development-related methods.

Firstly, since the 2010’s, there is a widely acknowledged decrease in productivity of the drug development industry (the extent of which may vary according to the disease at study), that is likely due to the increasing complexity of diseases to tackle, and to the high average drug development cost. This observation has led several pharmaceutical groups and labs to be interested in ML techniques, along with robotics, in order to decrease drug development time, and also to share observational data and clinical trial results [26,122]. Even if these efforts only result in a small decrease in drug failure rate during clinical development, this would be a both financially and scientifically profitable improvement for drug development.

Secondly, data availability and quality are key ingredients for the success of ML methods. There are hundreds of online publicly, curated databases which, provided some scripting efforts, can be integrated modularly to drug development pipelines. The feature selection step, which selects and transforms raw data to generate useful model inputs, has been shown to be of paramount importance in order to understand, and to obtain relevant results from ML applications. A variety of methods is now available in order to tackle this problem, and select valid, discriminatory features for prediction or decision model inputs.

Finally, the use of statistical learning algorithms is not short of challenges and should be handled with care. Nonetheless, the field of research seems ripe enough to be applied in at least semi-automated drug development pipelines; there is a growing number of papers published on both relevant data processing, and on algorithms applied for biomedical purposes. Subtle and powerful ML generic algorithms, such as refined DL architectures and sequential algorithms (that have been proven to be useful in a number of non-biology related fields of research), are becoming more and more prominent in biomedical research. They have been applied, with some success, to inherently complex problems, for instance, for adapted clinical trials, *in vitro* experiment prediction or guiding, or *de novo* protein design. Standardization, systematic validation and comparison of drug development methods on independent datasets are anticipated in the future [123].

In the light of new challenging problems, such as designing algorithms generating targeted recommendations for precision medicine, and modelling drug responses as outputs of a much larger system than a handful of genes, modern ML methods might be a useful tool to further enhance drug development.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors acknowledge the help of Pierre Gressens (who has proofread the manuscript) and Amazigh Mokhtari (who has redesigned figures 1, 3 and 5, the graphical abstract, and has provided useful feedback).

Funding

This work was supported by the Institut National pour la Santé et la Recherche Médicale (INSERM, France), the Centre National de la Recherche Scientifique (CNRS, France), the Université Paris 13,

the Université de Paris, NeuroDiderot, France, the French Ministry of Higher Education and Research [ENS.X19RDTME-SACLAY19-22] (C. R.), the French National Research Agency [ANR-18-CE17-0009-01] (A.D.-D.) [ANR-18-CE37-0002-03] (A.D.-D and A.M.) and by the «Digital health challenge» Inserm CNRS joint program (C.R., E. K. and A. D.-D.). The funders had no role in review design, decision to publish, or preparation of the manuscript.

Roles of the authors

C.R. and A.D.-D. conceived the short review. C. R., E. K. and A. D.-D. wrote the manuscript. C.R. conceived the figures and the graphical abstract. A.D.-D coordinated the work. All authors have approved the final article.

References

- [1] Eliopoulos H, Giranda V, Carr R, Tiehen R, Leahy T, Gordon G. Phase 0 trials: an industry perspective. *Clin Cancer Res* 2008;14(12):3683–8.
- [2] Xue H, Li J, Xie H, Wang Y. Review of drug repositioning approaches and resources. *Int J Biol Sci* 2018;14(10):1232.
- [3] Schuhmacher A, Gassmann O, Hinder M. Changing R&D models in research-based pharmaceutical companies. *J Transl Med* 2016;14(1):105.
- [4] Khanna I. Drug discovery in pharmaceutical industry: productivity challenges and trends. *Drug Discovery Today* 2012;17(19–20):1088–102.
- [5] Hwang TJ, Carpenter D, Lauffenburger JC, Wang B, Franklin JM, Kesselheim AS. Failure of investigational drugs in late-stage clinical development and publication of trial results. *JAMA Internal Med* 2016;176(12):1826–33.
- [6] Thomson Reuters (2014) CMR International Pharmaceutical R&D Executive Summary. Available at: citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.675.2511. Accessed [September 17, 2019].
- [7] Lowe D. The Latest on Drug Failure and Approval Rates. Available at: blogs.sciencemag.org/pipeline/archives/2019/05/09/the-latest-on-drug-failure-and-approval-rates. Accessed [September 16, 2019].
- [8] Wong CH, Siah KW, Lo AW. Estimation of clinical trial success rates and related parameters. *Biostatistics* 2019;20(2):273–86.
- [9] Pharamaprojects. Pharma r&d annual review 2018. Available at: pharmaintelligence.informa.com/~media/informa-shop-window/pharma/files/infographics/pharamaprojects-2018-pharma-rd-infographic.pdf. Accessed [September 16, 2019].
- [10] Deloitte Centre for Health Solutions. Embracing the future of work to unlock R&D productivity. Available at: deloitte.com/content/dam/Deloitte/uk/Documents/life-sciences-health-care/deloitte-uk-measuring-roi-pharma.pdf. Accessed [December 25, 2018].
- [11] Fogel DB. Factors associated with clinical trials that fail and opportunities for improving the likelihood of success: a review. *Contemp Clinical Trials Commun* 2018;11:156–64.
- [12] Meekings KN, Williams CS, Arrowsmith JE. Orphan drug development: an economically viable strategy for biopharma R&D. *Drug Discovery Today* 2012;17(13–14):660–4.
- [13] Ekins S, Puhl AC, Zorn KM, Lane TR, Russo DP, Klein JJ, et al. Exploiting machine learning for end-to-end drug discovery and development. *Nat Mater* 2019;18(5):435.
- [14] Zucchelli P. Lab Automation Increases Repeatability, Reduces Errors in Drug Development; 2018. Available at: technologynetworks.com/drug-discovery/articles/lab-automation-increases-repeatability-reduces-errors-in-drug-development-310034. Accessed [September 16, 2019].
- [15] Synced (Medium user). How AI Can Speed Up Drug Discovery; 2018. Available at: medium.com/syncedreview/how-ai-can-speed-up-drug-discovery-3c7f01654625. Accessed [September 16, 2019].
- [16] Sciforce (Medium user). AI in Pharmacy: Speeding up Drug Discovery. Available at: medium.com/sciforce/ai-in-pharmacy-speeding-up-drug-discovery-c7ca252c51bc. Accessed [September 16, 2019].
- [17] Meath P. How the AI Revolution Is Speeding Up Drug Discovery. Available at: jpmorgan.com/commercial-banking/insights/ai-revolution-drug-discovery. Accessed [September 16, 2019].
- [18] Walker, J. Machine Learning Drug Discovery Applications – Pfizer, Roche, GSK, and More. Available at: emerj.com/ai-sector-overviews/machine-learning-drug-discovery-applications-pfizer-roche-gsk. Accessed [September 16, 2019].
- [19] Waterfield, P. How Is Machine Learning Accelerating Drug Development? Available at: journal.binarydistrict.com/how-is-machine-learning-accelerating-drug-development. Accessed [September 16, 2019].
- [20] Budek K, Kornakiewicz A. Machine learning in drug discovery. Available at: deepense.ai/machine-learning-in-drug-discovery. Accessed [September 16, 2019].
- [21] Chan HS, Shan H, Dahoun T, Vogel H, Yuan S. Advancing drug discovery via artificial intelligence. *Trends Pharmacol Sci* 2019.
- [22] Armstrong M. Big pharma piles into machine learning, but what will it get out of it?; 2018. Available at: evaluate.com/vantage/articles/analysis/vantage-points/big-pharma-piles-machine-learning-what-will-it-get-out-of-it. Accessed [September 16, 2019].
- [23] Dutton G. Automation cuts drug development to 5 years. Available at: lifescienceleader.com/doc/automation-cuts-drug-development-to-years-0001. Accessed [September 16, 2019].
- [24] Dugger SA, Platt A, Goldstein DB. Drug development in the era of precision medicine. *Nat Rev Drug Discovery* 2018;17(3):183.
- [25] West M, Ginsburg GS, Huang AT, Nevins JR. Embracing the complexity of genomic data for personalized medicine. *Genome Res* 2006;16(5):559–66.
- [26] McKinsey Company. The 'big data' revolution in healthcare; 2013. Available at: mckinsey.com/~media/mckinsey/industries/healthcare%20systems%20and%20services/our%20insights/the%20big%20data%20revolution%20in%20us%20health%20care/the_big_data_revolution_in_healthcare.aspx. Accessed [September 16, 2019].
- [27] Hodos RA, Kidd BA, Khader S, Readhead BP, Dudley JT. Computational approaches to drug repurposing and pharmacology. *Wiley Interdiscip Rev Syst Biol Med* 2016;8(3):186.
- [28] Pope A. The Evolution of Automation for Pharmaceutical Lead Discovery; 2010. Available at: case2010.org/Automation%20for%20Pharmaceutical%20Lead%20Discovery.pdf. Accessed [September 16, 2019].
- [29] Musa A, Ghorai LS, Zhang SD, Glazko G, Yli-Harja O, Dehmer M, et al. A review of connectivity map and computational approaches in pharmacogenomics. *Brief Bioinform* 2017;19(3):506–23.
- [30] Srivastava PK, van Eyll J, Godard P, Mazzuferi M, Delahaye-Duriez A, Van Steenwinckel J, et al. A systems-level framework for drug discovery identifies Csf1R as an anti-epileptic drug target. *Nat Commun* 2018;9(1):3561.
- [31] Aliper A, Plis S, Artemov A, Ulloa A, Mamoshina P, Zhavoronkov A. Deep learning applications for predicting pharmacological properties of drugs and drug repurposing using transcriptomic data. *Mol Pharm* 2016;13(7):2524–30.
- [32] Schölkopf B, Tsuda K, Vert JP. Support vector machine applications in computational biology. MIT Press; 2004.
- [33] Chen H, Engkvist O, Wang Y, Olivecrona M, Blaschke T. The rise of deep learning in drug discovery. *Drug Discovery Today* 2018;23(6):1241–50.
- [34] Lattimore T, Szepesvári C. Bandit algorithms; 2018. preprint.
- [35] Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 1933;25(3/4):285–94.
- [36] Kaufmann E, Garivier A. Learning the distribution with largest mean: two bandit frameworks. *ESAIM: Proceed Surveys* 2017;60:114–31.
- [37] Mary J, Gaudel R, Preux P. Bandits and recommender systems. In: *International Workshop on Machine Learning, Optimization and Big Data*; 2015. Springer, Cham., p. 325–36.
- [38] Guillou F, Gaudel R, Preux P. Large-scale bandit recommender system. In: *International workshop on machine learning, optimization, and big data*. Cham: Springer; 2016. p. 204–15.
- [39] Adomavicius G, Tuzhilin A. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Trans Knowl Data Eng* 2005;6:734–49.
- [40] Brynjolfsson E, Hu YJ, Smith MD. The longer tail: The changing shape of Amazon's sales distribution curve; 2010. Available at SSRN 1679991.
- [41] Smith B, Linden G. Two decades of recommender systems at Amazon. *com. IEEE Internet Comput* 2017;21(3):12–8.
- [42] Li L, Chu W, Langford J, Schapire RE. A contextual-bandit approach to personalized news article recommendation. In: *Proceedings of the 19th international conference on world wide web*; 2010. ACM. p. 661–70.
- [43] Rifaioğlu AS, Atas H, Martin MJ, Cetin-Atalay R, Atalay V, Dogan T. Recent applications of deep learning and machine intelligence on in silico drug discovery: methods, tools and databases. *Brief Bioinform*; 2018, 10.
- [44] Zhang L, Tan J, Han D, Zhu H. From machine learning to deep learning: progress in machine intelligence for rational drug discovery. *Drug Discovery Today* 2017;22(11):1680–5.
- [45] Ferrero E, Dunham I, Sanseau P. In silico prediction of novel therapeutic targets using gene-disease association data. *J Transl Med* 2017;15(1):182.
- [46] Celesti F, Celesti A, Carnevale L, Galletta A, Campo S, Romano A et al. Big data analytics in genomics: The point on Deep Learning solutions. In: *2017 IEEE Symposium on Computers and Communications (ISCC)*; 2017. IEEE. p. 306–9.
- [47] Riniker S, Wang Y, Jenkins JL, Landrum GA. Using information from historical high-throughput screens to predict active compounds. *J Chem Inf Model* 2014;54(7):1880–91.
- [48] Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. *Nat Rev Drug Discovery* 2019;1.
- [49] Zou J, Huss M, Abid A, Mohammadi P, Torkamani A, Telenti A. A primer on deep learning in genomics. *Nat Genet* 2018;1.
- [50] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. In: *Advances in neural information processing systems*; 2014. p. 2672–80.
- [51] Killoran N, Lee LJ, Delong A, Duvenaud D, Frey BJ. Generating and designing DNA with deep generative models; 2017. arXiv preprint arXiv:1712.06148.

- [52] Gupta A, Zou J. Feedback GAN (FBGAN) for DNA: A novel feedback-loop architecture for optimizing protein functions; 2018. arXiv preprint arXiv:1804.01694.
- [53] Kandasamy K, Krishnamurthy A, Schneider J, Póczos B. Parallelised bayesian optimisation via Thompson sampling. In: International Conference on Artificial Intelligence and Statistics; 2018. p. 133–42.
- [54] Griffiths RR, Hernández-Lobato JM. Constrained bayesian optimization for automatic chemical design; 2017. arXiv preprint arXiv:1709.05501.
- [55] Anand N, Huang P. Generative modeling for protein structures. In: Advances in neural information processing systems; 2018. p. 7494–7505.
- [56] Gonzalez J, Longworth J, James DC, Lawrence ND. Bayesian optimization for synthetic gene design; 2015. arXiv preprint arXiv:1505.01627.
- [57] Williams CK, Rasmussen, CE. (2006). Gaussian processes for machine learning (vol. 2, No. 3, p. 4). Cambridge, MA: MIT Press.
- [58] Pyzner-Knapp EO. Bayesian optimization for accelerated drug discovery. *IBM J Res Dev* 2018;62(6):2–11.
- [59] Benhenda M. ChemGAN challenge for drug discovery: can AI reproduce natural chemical diversity?; 2017. arXiv preprint arXiv:1708.08227.
- [60] Grimm D. U.S. EPA to eliminate all mammal testing by 2035. Available at: sciencemag.org/news/2019/09/us-epa-eliminate-all-mammal-testing-2035. Accessed [September 13, 2019].
- [61] Dunn SJ, Martello G, Yordanov B, Emmott S, Smith AG. Defining an essential transcription factor program for naive pluripotency. *Science* 2014;344(6188):1156–60.
- [62] Aziz M, Kaufmann E, Riviere MK. On Multi-Armed Bandit Designs for Phase I Clinical Trials; 2019. arXiv preprint arXiv:1903.07082.
- [63] Berry SM, Carlin BP, Lee JJ, Muller P. Bayesian adaptive methods for clinical trials. CRC Press; 2010.
- [64] Satlin A, Wang J, Logovinsky V, Berry S, Swanson C, Dhadda S, et al. Design of a Bayesian adaptive phase 2 proof-of-concept trial for BAN2401, a putative disease-modifying monoclonal antibody for the treatment of Alzheimer's disease. *Alzheimer's Dementia: Transl Res Clin Intervent* 2016;2(1):1–12.
- [65] Berry DA. Interim analysis in clinical trials: the role of the likelihood principle. *Am Statist* 1987;41(2):117–22.
- [66] Robbins H. Some aspects of the sequential design of experiments. *Bull Am Math Soc* 1952;58(5):527–35.
- [67] Lai TL, Robbins H. Asymptotically efficient adaptive allocation rules. *Adv Appl Math* 1985;6(1):4–22.
- [68] Berry DA, Fristedt B. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). London: Chapman and Hall; 1985. p. 71–87.
- [69] Gittins JC. Bandit processes and dynamic allocation indices. *J Roy Stat Soc: Ser B (Methodol)* 1979;41(2):148–64.
- [70] Jennison C, Johnstone IM, Turnbull BW. Asymptotically optimal procedures for sequential adaptive selection of the best of several normal means. In: Statistical decision theory and related topics III. Academic Press; 1982. p. 55–86.
- [71] Bechhofer RE. A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Ann Math Stat* 1954;16–39.
- [72] Even-Dar E, Mannor S, Mansour Y. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *J Machine Learn Res* 2006;7(June):1079–105.
- [73] Audibert JY, Bubeck S. Best arm identification in multi-armed bandits; 2010.
- [74] Bubeck S, Munos R, Stoltz G. Pure exploration in finitely-armed and continuous-armed bandits. *Theoret Comput Sci* 2011;412(19):1832–52.
- [75] Hardwick J, Stout QF. Bandit strategies for ethical sequential allocation. *Comp Sci Stat* 1991;23(6.1):421–4.
- [76] Armitage P. The search for optimality in clinical trials. *Int Statist Rev/Rev Int Statist* 1985;15–24.
- [77] Villar SS, Bowden J, Wason J. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Stat Sci* 2015;30(2):199.
- [78] Villar SS. Bandit strategies evaluated in the context of clinical trials in rare life-threatening diseases. *Probab Eng Inf Sci* 2018;32(2):229–45.
- [79] Food and Drug Administration (FDA). Adaptive Design Clinical Trials for Drugs and Biologics; 2018. Available at: fda.gov/regulatory-information/search-fda-guidance-documents/adaptive-design-clinical-trials-drugs-and-biologics. Accessed [September 12, 2019].
- [80] Chapelle O, Li L. An empirical evaluation of Thompson sampling. In: Advances in neural information processing systems; 2011. p. 2249–57.
- [81] Thall PF, Wathen JK. Practical Bayesian adaptive randomisation in clinical trials. *Eur J Cancer* 2007;43(5):859–66.
- [82] Berry DA. Bayesian clinical trials. *Nat Rev Drug Discovery* 2006;5(1):27.
- [83] Ginsbourger D, Janusevskis J, Le Riche R. Dealing with asynchronicity in parallel Gaussian process based global optimization; 2011.
- [84] Varatharajah Y, Berry B, Koyejo S, Iyer R. A Contextual-bandit-based Approach for Informed Decision-making in Clinical Trials; 2018. arXiv preprint arXiv:1809.00258.
- [85] Durand A, Achilleos C, Iacovides D, Strati K, Mitsis GD, Pineau J. Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis. In: Machine learning for healthcare conference; 2018. p. 67–82.
- [86] Andronis C, Sharma A, Virvilis V, Deftereos S, Persidis A. Literature mining, ontologies and information visualization for drug repurposing. *Briefings Bioinf* 2011;12(4):357–68.
- [87] Bisgin H, Liu Z, Kelly R, Fang H, Xu X, Tong W. Investigating drug repositioning opportunities in FDA drug labels through topic modeling. In *BMC bioinformatics*. BioMed Central 2012;13(15):S6.
- [88] Tari LB, Patel JH. Systematic drug repurposing through text mining. In: *Biomedical Literature Mining*. New York, NY: Humana Press; 2014. p. 253–67.
- [89] Alaimo S, Giugno R, Pulvirenti A. Recommendation techniques for drug-Target interaction prediction and drug repositioning. In: *Data mining techniques for the life sciences*. New York, NY: Humana Press; 2016. p. 441–62.
- [90] Sardana D, Zhu C, Zhang M, Gudivada RC, Yang L, Jegga AG. Drug repositioning for orphan diseases. *Briefings Bioinf* 2011;12(4):346–56.
- [91] Vargesson N. Thalidomide-induced teratogenesis: history and mechanisms. *Birth Defects Res Part C: Embryo Today: Rev* 2015;105(2):140–56.
- [92] Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. *Cell* 2017;171(6):1437–52.
- [93] Zhang W, Yue X, Huang F, Liu R, Chen Y, Ruan C. Predicting drug-disease associations and their therapeutic function based on the drug-disease association bipartite network. *Methods* 2018;145:51–9.
- [94] Napolitano F, Zhao Y, Moreira VM, Tagliaferri R, Kere J, D'Amato M, et al. Drug repositioning: a machine-learning approach through data integration. *J Cheminf* 2013;5(1):30.
- [95] Gilvary C, Madhukar N, Elkhader J, Elemento O. The missing pieces of artificial intelligence in medicine. *Trends Pharmacol Sci* 2019.
- [96] Zitnik M, Nguyen F, Wang B, Leskovec J, Goldenberg A, Hoffman MM. Machine learning for integrating data in biology and medicine: Principles, practice, and opportunities. *Inform Fusion* 2019;50:71–91.
- [97] Carvalho-Silva D, Pierleoni A, Pignatelli M, Ong C, Fumis L, Karamanis N, et al. Open targets platform: new developments and updates two years on. *Nucleic Acids Res* 2018;47(D1):D1056–65.
- [98] Davis AP, Grondin CJ, Johnson RJ, Sciaky D, McMorran R, Wieggers J, et al. The comparative toxicogenomics database: update 2019. *Nucleic Acids Res* 2018;47(D1):D948–54.
- [99] Piñero J, Bravo A, Gutiérrez-Sacristán A, Centeno E, Deu-Pons J, García J. DisGeNET, a centralized repository of the genetic basis of human diseases. *F1000Research* 2017;6.
- [100] Forbes SA, Bindal N, Bamford S, Cole C, Kok CY, Beare D, et al. COSMIC: mining complete cancer genomes in the Catalogue of Somatic Mutations in Cancer. *Nucleic Acids Res* 2010;39(suppl_1):D945–50.
- [101] Sharov AA, Dudekula DB, Ko MS. CisView: a browser and database of cis-regulatory modules predicted in the mouse genome. *DNA Res* 2006;13(3):123–34.
- [102] Bycroft C, Freeman C, Petkova D, Band G, Elliott LT, Sharp K, et al. The UK Biobank resource with deep phenotyping and genomic data. *Nature* 2018;562(7726):203.
- [103] Szklarczyk D, Gable AL, Lyon D, Junge A, Wyder S, Huerta-Cepas J, et al. STRING v11: protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. *Nucleic Acids Res* 2018;47(D1):D607–13.
- [104] Kanehisa M, Furumichi M, Tanabe M, Sato Y, Morishima K. KEGG: new perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Res* 2016;45(D1):D353–61.
- [105] Boué S, Talikka M, Westra JW, Hayes W, Di Fabio A, Park J, et al. Causal biological network database: a comprehensive platform of causal biological network models focused on the pulmonary and vascular systems. *Database* 2015;2015.
- [106] Chelliah V, Juty N, Ajmera I, Ali R, Dumousseau M, Glont M, et al. BioModels: ten-year anniversary. *Nucleic Acids Res* 2014;43(D1):D542–8.
- [107] Lamb J, Crawford ED, Peck D, Modell JW, Blat IC, Wrobel MJ, et al. The Connectivity Map: using gene-expression signatures to connect small molecules, genes, and disease. *Science* 2006;313(5795):1929–35.
- [108] Li YH, Yu CY, Li XX, Zhang P, Tang J, Yang Q, et al. Therapeutic target database update 2018: enriched resource for facilitating bench-to-clinic research of targeted therapeutics. *Nucleic Acids Res* 2017;46(D1):D1121–7.
- [109] Von Eichborn J, Murgueitio MS, Dunkel M, Koerner S, Bourne PE, Preissner R. PROMISCUOUS: a database for network-based drug-repositioning. *Nucleic Acids Res* 2010;39(suppl_1):D1060–6.
- [110] Brown AS, Patel CJ. A standard database for drug repositioning. *Sci Data* 2017;4:170029.
- [111] UniProt Consortium. UniProt: the universal protein knowledgebase. *Nucleic Acids Res* 2018;46(5):2699.
- [112] Wishart DS, Feunang YD, Guo AC, Lo EJ, Marcu A, Grant JR, et al. DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic Acids Res* 2017;46(D1):D1074–82.
- [113] Gaulton A, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, et al. The ChEMBL database in 2017. *Nucleic Acids Res* 2016;45(D1):D945–54.
- [114] Koscielny S. Why most gene expression signatures of tumors have not been useful in the clinic. *Sci Transl Med* 2010;2(14).
- [115] Guyon I, Elisseeff A. An introduction to variable and feature selection. *J Machine Learn Res* 2003;3(Mar):1157–82.
- [116] Zeng X, Zhu S, Liu X, Zhou Y, Nussinov R, Cheng F. deepDR: a network-based deep learning approach to in silico drug repositioning. *Bioinformatics* 2019.
- [117] Miotto R, Wang F, Wang S, Jiang X, Dudley JT. Deep learning for healthcare: review, opportunities and challenges. *Briefings Bioinf* 2017;19(6):1236–46.
- [118] Delahaye-Duriez A, Srivastava P, Shkura K, Langley SR, Laaniste L, Moreno-Moral A, et al. Rare and common epilepsies converge on a shared gene regulatory network providing opportunities for novel antiepileptic drug discovery. *Genome Biol* 2016;17(1):245.

- [119] Ahmed SS, Roy S, Kalita JK. Assessing the effectiveness of causality inference methods for gene regulatory networks. *IEEE/ACM transactions on computational biology and bioinformatics*; 2018.
- [120] Young WC, Yeung KY, Raftery AE. A posterior probability approach for gene regulatory network inference in genetic perturbation data; 2016. arXiv preprint arXiv:1603.04835.
- [121] Tardif JC, Rhainds D, Brodeur M, Feroz Zada Y, Fouodjio R, Provost S, et al. Genotype-dependent effects of dalcetrapib on cholesterol efflux and inflammation: concordance with clinical outcomes. *Circulation: Cardiovascular Genetics* 2016;9(4):340–8.
- [122] Burki T. Pharma blockchains AI for drug development. *Lancet* 2019;393(10189):2382.
- [123] Brown AS, Kong SW, Kohane IS, Patel CJ. ksRepo: a generalized platform for computational drug repositioning. *BMC Bioinf* 2016;17(1):78.
- [124] Pallmann P, Bedding AW, Choodari-Oskooei B, Dimairo M, Flight L, Hampson LV, et al. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC medicine* 2018;16:29. <https://doi.org/10.1186/s12916-018-1017-7>.