PORTLAND PRESS

Review Article

Molecular-based precision oncology clinical decision making augmented by artificial intelligence

Jia Zeng and Md Abu Shufean

Sheikh Khalifa Bin Zayed Al Nahyan Institute for Personalized Cancer Therapy, The University of Texas MD Anderson Cancer Center, Houston, TX, U.S.A.

Correspondence: Jia Zeng (jzeng@mdanderson.org)

OPEN ACCESS

The rapid growth and decreasing cost of Next-generation sequencing (NGS) technologies have made it possible to conduct routine large panel genomic sequencing in many disease settings, especially in the oncology domain. Furthermore, it is now known that optimal disease management of patients depends on individualized cancer treatment guided by comprehensive molecular testing. However, translating results from molecular sequencing reports into actionable clinical insights remains a challenge to most clinicians. In this review, we discuss about some representative systems that leverage artificial intelligence (AI) to facilitate some processes of clinicians' decision making based upon molecular data, focusing on their application in precision oncology. Some limitations and pitfalls of the current application of AI in clinical decision making are also discussed.

Introduction

Clinical decision making is a contextual, continuous, and evolving process, where data are gathered, interpreted, and evaluated in order to select an evidence-based choice of action [1]. It consists of three integrated stages: diagnosis, prognosis assessment and disease management.

Molecular tests that check for certain changes at the genomic level that may cause or affect the chance of developing a specific disease or disorder have been routinely used to assist diagnosis and treatment of infectious diseases, inherited diseases and cancer. As the world's second largest cause of deaths, cancer refers to a group of genetic diseases that are caused by disruptive changes to the genes that control the way our cells function. The high throughput capacity of next-generation sequencing (NGS) technologies have transformed genomic testing and our understanding of cancer [2,3], making practicing precision oncology at a large scale a reality.

The field of precision oncology is concerned with developing treatments targeting the molecular characteristic of an individual's tumor [4]. There is a growing recognition that optimal disease management of patients depends on individualized cancer treatment guided by comprehensive molecular testing [5]. A 2017 survey revealed that 75.6% of the 1281 United States oncologists surveyed utilized NGS tests to guide their treatment decisions in the 12-month period preceding the survey [6]. As large panel genomic sequencing becomes routine in many disease settings, the challenge of translating the molecular sequencing reports to actionable clinical insights is not to be underestimated.

Artificial Intelligence (AI) refers to a broad field in computer science that studies the theory, algorithms and architectures that enable machines to perform tasks that would otherwise require human intelligence. Machine learning (ML) is a subfield of AI that focuses on the development of computer systems that are able to learn and adapt without explicit instructions, but instead leverage examples to draw inferences and identify patterns computationally. Thanks to the advancement and decreasing costs of computer hardware, the past decade has seen exponential growth in the use and development of ML applications. Deep sequencing of exomes or even whole genomes has also been rapidly
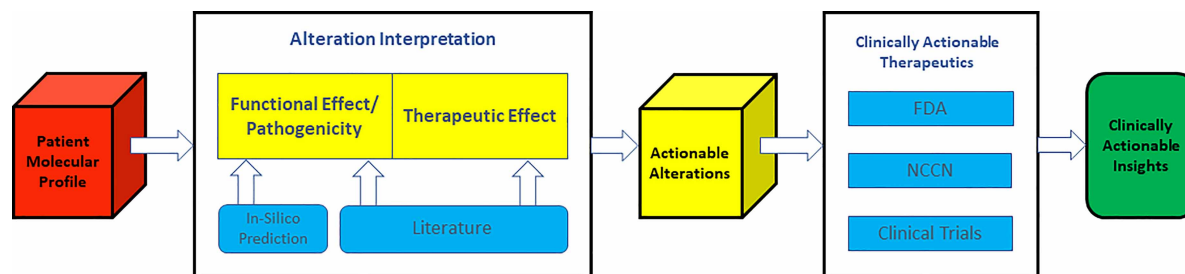
Figure 1. Operational processes of translating patient molecular profile into clinically actionable insights.

integrated into clinical practice, giving rise to well-curated databases such as The Cancer Genome Atlas (TCGA) which provides annotations of clinically relevant metadata which are suitable to serve as training examples in AI systems [7,8].

In this review, we provide a general overview of the operational processes involved in using molecular-based results to guide treatment planning in precision oncology (Figure 1) and discuss about how ML technologies can be applied to facilitate some of these processes in order to augment clinicians in their decision making. We focus on the state-of-the-art systems that have been explored in clinical settings, or those having high translational potential and providing publicly-accessible resources (we have verified all of the URLs shared in this review to ensure they are functional as of the time of writing). We also discuss some current limitations in applying AI in clinical processes and present some suggestions for future work to address these concerns.

Alteration interpretation

Not all mutations lead to cancer progression [9]. From clinical perspective, it is now known that the efficacy of targeted therapies depends on the genetic alterations of individual patients [10]. Understanding the functional and therapeutic significance of individual alterations is therefore critical to support clinicians' decision making.

Computational prediction of variants' functional effect and pathogenicity

Since it is cost prohibitive to functionally evaluate all known mutations using validated assay, it is routine to utilize benchmark *in silico* tools such as PolyPhen-2, SIFT, FASMIC [11–13], to provide computational prediction on the functional effects of variants of interest. PolyPhen-2 leverages a supervised learning method, called Naïve Bayes classifier, to train their model using sequence-based feature profile derived from two pairs of datasets that contain comprehensive repository of disease-causing mutations known to be either damaging or neutral [14] (URL: http://genetics.bwh.harvard.edu/pph2/). While tools such as PolyPhen-2, SIFT, VEST and MutationAssessor predict individual mutation's pathogenicity on a population level [15,16], which may provide limited value to clinicians that are specialized to ascertain somatic mutations in cancer patients, several tools such as CHASM, CanDrA, fathmm and transFIC have been developed to identify cancer-specific driver mutations [17–20]. Specifically, Mao and colleagues developed a meta predictor called CanDrA by leveraging predictions on 95 structural and evolutionary features from 10 existing functional prediction algorithms including CHASM, SIFT and MutationAssessor. Through the utilization of feature optimization and support vector machine, CanDrA was shown to combat the curse of dimensionality, an issue that plagued the oncology domain, better than its counterparts and yielded better overall performance on real-world data (https://bioinformatics.mdanderson.org/public-software/candra/). Lawrence and colleagues observed that with cancer genome studies, as sample size increases, extensive false-positive findings ensue. They hypothesized that this was due to mutational heterogeneity and incorporated this aspect into their analysis. Their method, named MutSigCV, was applied to exome sequences of 3083 tumor-normal pairs and was found to eliminate most of the artefactual findings (https://software.broadinstitute.org/cancer/cga/mutsig) [21]. Dietlein et al. provided another novel methodology that was built on the observation that mutations in certain contexts provide a signal in favor of driver genes thus by combining this nucleotide context feature with signals traditionally used for driver-gene identification, their method identified 460 driver genes related to 21 cancer-related pathways when being applied on whole-exome sequencing data from 11 873 tumor-normal pairs [22]. To leverage the synergy created by utilizing multiple prediction tools, an ensemble strategy has been explored in some systems.

Bailey et al. reported a PanCancer and PanSoftware analysis of data provided by all the projects in TCGA and 26 computational tools. The ensemble tool identified more than 3400 potential missense driver mutations supported by multiple lines of evidence [23]. Using a training set of variants across multiple genes enables ML systems such as PolyPhen-2 to be applicable as a general model, however, for specific genes, assuming well-characterized data is available (such as *BRCA1, BRCA2*), gene-specific models likely outperform general models [24]. Using results from validated functional assays, Hart et al. [24] iteratively trained and evaluated hundreds of ML algorithms along with the associated hyperparameters, resulting in a new optimal BRCA-ML model which yielded a system that significantly outperformed existing approaches. More recently, Martinez-Jimenez et al. systematically combined several prediction tools that aim to identify signals of positive selection of cancer driver genes in tumorigenesis and created an Integrative OncoGenomics (IntOGen) pipeline (https://www.intogen.org/search). They applied their ensemble system to somatic mutations of more than 28 000 tumors across 66 cancer types and identified 568 cancer genes that are likely contributors of tumorigenesis [25].

Literature mining of alteration actionability

While computational prediction of functional effects of alterations offers some insights about variants of unknown significance (VUSs), different prediction tools may reach different conclusions and sometimes they may not align with experimental observations [2]. Whenever available, publications reporting functionally characterized mutations are still seen as a higher level of evidence than *in silico* predictions. To achieve this goal at a large scale, a systematic literature mining process is expected. Furthermore, the actionability of an alteration is not only related to its functional but also therapeutic effect. Identifying evidence of therapeutic significance of mutations is largely dependent on reviewing the published biomedical literature. A combination of factors including the growing gene panel size, the enormous and rapidly expanding body of literature, and the labor-intensive nature of manual review makes utilization of informatics systems highly desirable.

To computationally assist with an assessment of an alteration's actionability using literature mining, it is critical to accomplish two tasks: (1) named entity recognition (NER), which recognizes the key entities such as biomarkers (including genes, mutations), the effects of interest, as well as therapies and conditions (for therapeutic contexts); and (2) relationship extraction. Traditionally, NER tasks may benefit from dictionary and rule-based NLP systems, especially on specific domains where expert input is available for defining customized rules. But they may be difficult to extend to other types of entities; iterative updates of dictionaries and rules may also increase the maintenance cost.

Machine learning technologies may enable the systems to adapt and evolve with little human intervention. A sequence labeling method called conditional random fields (CRFs) is considered the state-of-the-art ML algorithm for NER tasks and is leveraged by popular systems such as ABNER, BANNER where gold standard annotations of pre-defined features were used in training the model [26,27]. More recently, deep learning methods have been explored to expand the ML models to consider non-linear features such as word embedding. Habibi et al. [28] reported impressive performance of a generic model by combining long short-term memory (LSTM) network with traditional CRFs and leveraged word embeddings to help the deep neural network to capture semantic similarities among related terms (source code and trained models are available on https://github.com/glample/tagger).

In the context of alteration actionability assessment, simply recognizing the key entities is not sufficient and the ability to automatically extract biomarker-drug relationship is critical. Lee et al. [29] developed an exceptional dictionary-based NER called Biomedical Entity Search Tool (BEST) that offers real-time retrieval of PubMed articles that mentioned the concepts in ten types of biomedical entities including mutations, genes, drugs, diseases, etc. Thanks to their creative design of the inverted indices which include entity-document pairs, BEST is able to deliver dynamic results to users with little response latency and also can accommodate conjunctive queries. In the next iteration of the system, Lee et al. leveraged the score produced by BEST as one of the features and applied deep learning technology coupled with word embedding features obtained from word2vec and achieved promising results in extracting mutation-gene-drug relationships from the PubMed literature [30]. The baseline BEST tool is publicly available at http://best.korea.ac.kr/ and is updated on a regular basis (according to the creators). To facilitate a team of professional curators' work in reviewing clinical evidence for precision oncology, Lever et al. have developed a text mining tool which prioritizes literature predicted to be highly relevant to their mission. Logistic regression classifier was trained by pre-labeled sentences with desired relationships [31]. The data is made publicly available at http://bionlp.bcgsc.ca/civicmine/. It is worth noting that the development team prioritized precision over recall in their parameter tuning process so

while extracting some relationship (such as associated variant) yielded a recall of 0.794, the recall of extracting predictive evidence, a very important factor of alteration interpretation in precision oncology, had much room for improvement (0.141). In the commercial domain, the cognitive computing powerhouse IBM Watson developed a comprehensive literature mining pipeline called Watson for Genomics (WfG) that leveraged training data curated by subject matter experts to identify genomic-based actionability evidence [32]. Compared with a traditional molecular tumor board, WfG was shown to identify new evidence not previously discovered and do so with much higher efficiency than the manual counterpart.

Prioritization

In the case when multiple actionable alterations are identified in a patient's molecular profile, how to prioritize is a practical concern for clinicians. The bioinformatics community has explored using supervised learning methods to identify driver mutations but cohort-level solutions often cannot sufficiently address the needs at single patient level. iCAGES leveraged curated drug association database to support a patient-level analysis on somatic mutations and copy-number data without the need for users to provide sophisticated configuration data as in tools such as OncoIMPACT and DriverNet [33–35]. More recently, PANOPLY incorporated clinical features in addition to omics data and applied random forest analysis to identify prioritized treatment given a patient's clinical and molecular profile (http://kalarikrlab.org/Software/Panoply.html) and some anecdotal success was reported [36]. Nulsen et al. [37] developed a one-class support vector machine called sysSVM that was trained on TCGA data to create a pan-cancer detection tool for identifying driver genes at the granularity of single patients (https://github.com/ciccalab/sysSVM2). Computational validation has shown promising results in terms of low false positive rate which is very essential for the clinical utilization but none of the applications has been formally evaluated in clinical contexts. Overall, extensive clinical validation is needed before such type of informatics systems can be integrated into clinical care workflow.

Clinical trial matching and pre-screening

After identifying mutations in the patient molecular profile that are actionable (if any), the next operational task for treatment planning is to find matching therapeutics that are clinically available, i.e, approved by Food and Drug Administration (FDA), recommended by professional guidelines or under clinical development via clinical trials.

In terms of data access, FDA provides an open-access database that catalogs all pharmaceutical agents which gained the approval from this agency [38]. National Comprehensive Cancer Network (NCCN), a national leader in professional guidelines for oncologists, offers subscribers access to manually curated structured data about drugs that are referenced in their clinical guidelines, along with the corresponding biomarkers (NCCN Drugs and Biologics Compendium and Biomarker Compendium) [39]. Therefore, there is little ambiguity in data extraction from these high-level data sources and most commercial vendors include this information in their molecular sequencing reports.

However, while seventy or so molecular-matched therapies have gained FDA approval, many more are still in clinical development phase and thus are only clinically available in experimental settings as part of an ongoing clinical trial. As of the time of writing, there are over 21 000 open oncology clinical trials that are listed on ClinicalTrials.gov, the world's largest clinical trial databank [40]. As comprehensive as ClinicalTrials. gov's coverage is, much pertinent information such as diseases, drugs, eligibility criteria is still largely captured in unstructured format [41]. Especially for precision oncology trials which aim to assess genomically matched therapies, there are many unique characteristics and nuances in the eligibility criteria that need to be accommodated. Dedicated frameworks and workflows have been proposed in the literature to help facilitate the curation of clinical trial knowledgebase which in turn is used for patient matching but most of the curation tasks remain manual thus scalability is a practical concern [42,43]. Being able to comprehensively characterize a patient's clinical and molecular profile for accurate trial matching is also often a manual process that is laborious.

Artificial intelligence techniques have been actively explored in clinical trial matching and patient pre-screening. Generally speaking, these systems apply a combination of NLP and ML techniques to extract pertinent demographic, clinical and molecular information from patient medical records, clinical trial documents, or a combination of both. Tools like Antidote, Synergy are patient-centric while Deep6.ai, Mendel.ai and Watson for Clinical Trial Matching (WCTM) are more clinician-oriented [44–48].

Formal evaluation of these commercial systems is rare but there are some exceptions. Mendel.ai published a retrospective study using their AI-empowered tool to facilitate clinical trial matching and reported 24%–50% increase over standard practices in terms of number of patients that were correctly identified as eligible in two out of three trials but failed to identify eligible patients for a trial that was closed due to lack of accrual [49]. Their system combined ML and NLP techniques in tasks such as text recognition and clinical language understanding. A central knowledge base stores all knowledge extracted from real world data regarding patients' profiles. WCTM published several studies where cognitive computing was used in real clinical settings, where WCTM was found to outperform manual counterparts in accuracy and yield significant increase in efficiency [47,50,51]. WCTM used NLP to process unstructured data from patient's electronic medical records (whenever available) and further leveraged its ML component to populate patient model with pertinent disease related attributes. A similar process was utilized in intaking the clinical trial information using lessons learnt from several rounds of trial ingestion in the pilot phase.

In academic settings, utilizing NLP in conjunction with ML to automate certain processes that facilitate clinical trial matching is also extensively studied. On one hand, some efforts are related to automatically convert unstructured text describing clinical trial eligibility criteria into structured conditions that are query-able [52,53]. Notably, Liu et al. [54] have developed a tool called DQueST which creates dynamic questionnaire for clinical trial searches. They developed an integrated NLP pipeline including negation detection such as NegEx [55] and entity recognition module that applied CRFs as the ML model in their NLP module to identify clinical entity, attribute recognitions and domains leveraging annotated corpus [54]. Miotto and Weng [56] developed a novel solution which used a case-based reasoning method to establish a 'target patient profile' using electronic records data from patients' who have already enrolled in a given clinical trial and to compare new patients' profiles with the said 'target profile' in order to computationally predict their eligibility. While the real-life application of this approach remains to be seen, the creative use of vector-based ML method bypasses the need to literally process clinical trial documents and may inspire other related applications.

Discussion

With the rapid development in computer hardware and software, utilizing artificial intelligence to provide near real-time assistance to clinicians in their decision making has become a reachable goal from a purely technical sense. This review discussed some representative systems that have either been explored in some clinical settings or exhibited high translational potential. While the progress is encouraging, many challenges remain that limit a wider utilization of AI systems in clinical decision making. Here we outline a few examples of such challenges and some potential mitigation tactics.

Data availability

The performance of a data-driven model is directly correlated with the quality and volume of the data it analyzes. While technologies help generate a myriad of data, data do not equate knowledge. Supervised learning methods, which constitute the majority of the AI methods explored today, specifically require pre-labeled data for training their models. While publicly-accessible datasets exist, most of the granular clinical decision making processes warrant more domain-specific gold standards. Manually curating these gold standards for specific operational needs is not a trivial process. Efforts for making more of these curated datasets publicly available will help the community to grow.

Validation

While all studies mentioned in this review were evaluated, usually via some retrospective studies, more extensive validation is often needed before they can be embraced by the broad clinical community. Randomized trials provide more systematic evaluation and ultimate proof of clinical utility [57] Real-world validation that applies the AI systems in actual clinical care process is also important to assess the real-world benefits and limitations of such systems.

Stakeholder acceptance

To broaden the impact of AI in medicine, it is critical that the development of these methods is done with the human experts, who are the stakeholders, in the loop and with the intention of augmenting instead of replacing human intelligence. Some basic literacy of ML methods is critical for the clinicians and decision makers to arrive at a critical yet reasonable assessment of the technologies involved. It is important to acknowledge that

there is no universal solution and the choice and design of technologies are usually context-dependent. Generalizability, explainability, resource overhead and cost-reduction potential are some examples of important factors for stakeholders to consider in addition to the raw performance reported on some validation experiments. Interested readers can review Faes et al. [58] work which outlined a clinician's guide to critically assess ML studies for more details.

Regulatory development

Due to the ethical and safety impact of all clinical decision making processes, it is important for any technology maker to consider the regulatory component early in the development phase. While the related regulatory processes remain largely elusive, some recent clinical trial guidelines for protocols involving AI systems such a SPIRIT-AI are a positive step towards the right direction [57].

Summary

- Next-generation sequencing technologies transformed our understanding of cancer and are actively being utilized in many disease settings to guide treatment planning for the purpose of optimizing outcomes.

- The past decade has seen exponential growth in the development of artificial intelligence (AI), machine learning (ML) applications which have the potential to augment human intelligence in a scalable fashion.

- This review discussed some state of the art systems that employed ML technologies to facilitate some processes involved in clinicians' treatment planning using molecular data. Some challenges of AI applications in this domain are also discussed.

Competing Interests

The authors declare that there are no competing interests associated with the manuscript.

Author Contributions

Conception and design: J.Z.; Collection and assembly of data: J.Z., M.A.S.; Data analysis and interpretation: J.Z.; Manuscript writing: J.Z., M.A.S.; Final approval of manuscript: J.Z., M.A.S.; Accountable for all aspects of the work: J.Z., M.A.S.

Abbreviations

AI, artificial intelligence; BEST, Biomedical Entity Search Tool; CRFs, conditional random fields; FDA, Food and Drug Administration; IntOGen, Integrative OncoGenomics; ML, machine learning; NCCN, National Comprehensive Cancer Network; NER, named entity recognition; NGS, next-generation sequencing; WCTM, Watson for Clinical Trial Matching.

References

1 Tiffen, J., Corbridge, S.J. and Slimmer, L. (2014) Enhancing clinical decision making: development of a contiguous definition and conceptual framework. *J. Prof. Nurs.* **30**, 399–405 https://doi.org/10.1016/j.profnurs.2014.01.006

2 Zeng, J., Johnson, A., Shufean, M.A., Kahle, M., Yang, D., Woodman, S.E. et al. (2019) Operationalization of next-Generation sequencing and decision support for precision oncology. *JCO Clin. Cancer Inform.* **3**, 1–12 https://ascopubs.org/doi/10.1200/CCI.19.00089

3   Zeng, J., Johnson, A., Abu Shufean, M., Kahle, M., Yang, D., Woodman, S.E. et al. (2019) Operationalization of next-Generation sequencing and decision support for precision oncology. *JCO Clin. Cancer Info.* **3**, 1–12 https://doi.org/10.1200/CCI.19.00089

4   Brody, H., Hodson, R. and Rooke, J. (2020) Precision oncology. *Nature* **585**, S1

5   Tamborero, D., Dienstmann, R., Rachid, M.H., Boekel, J., Baird, R., Brana, I. et al. (2020) Support systems to guide clinical decision-making in precision oncology: the cancer core Europe molecular tumor board portal. *Nat. Med.* **26**, 992–994 https://doi.org/10.1038/s41591-020-0969-2

6   Freedman, A.N., Klabunde, C.N., Wiant, K., Enewold, L., Gray, S.W., Filipski, K.K. et al. (2018) Use of next-generation sequencing tests to guide cancer treatment: Results from a nationally representative survey of oncologists in the United States. *JCO Precis. Oncol.* **2**, 1–13 https://doi.org/10.1200/PO.18.00169

7   Adlung, L., Cohen, Y., Mor, U. and Elinav, E. (2021) Machine learning in clinical decision making. *Med* **2**, 642–665 https://doi.org/10.1016/j.medj.2021.04.006

8   Cancer Genome Atlas Research Network, Weinstein, J.N., Collisson, E.A., Mills, G.B., Shaw, K.R., Ozenberger, B.A. et al. (2013) The Cancer Genome Atlas Pan-Cancer analysis project. *Nat. Genet.* **45**, 1113–1120 https://doi.org/10.1038/ng.2764

9   Jordan, E.J., Patil, K., Suresh, K., Park, J.H., Mosse, Y.P., Lemmon, M.A. et al. (2019) Computational algorithms for in silico profiling of activating mutations in cancer. *Cell. Mol. Life Sci.* **76**, 2663–2679 https://doi.org/10.1007/s00018-019-03097-2

10   Normanno, N., Rachiglio, A.M., Roma, C., Fenizia, F., Esposito, C., Pasquale, R. et al. (2013) Molecular diagnostics and personalized medicine in oncology: challenges and opportunities. *J. Cell. Biochem.* **114**, 514–524 https://doi.org/10.1002/jcb.24401

11   Adzhubei, I.A., Schmidt, S., Peshkin, L., Ramensky, V.E., Gerasimova, A., Bork, P. et al. (2010) A method and server for predicting damaging missense mutations. *Nat. Methods* **7**, 248–249 https://doi.org/10.1038/nmeth0410-248

12   Kumar, P., Henikoff, S. and Ng, P.C. (2009) Predicting the effects of coding non-synonymous variants on protein function using the SIFT algorithm. *Nat. Protoc.* **4**, 1073–1082 https://doi.org/10.1038/nprot.2009.86

13   Ng, P.K.S., Li, J., Jeong, K.J., Shao, S., Chen, H., Tsang, Y.H. et al. (2018) Systematic functional annotation of somatic mutations in cancer. *Cancer Cell* **33**, 450 https://doi.org/10.1016/j.ccell.2018.01.021

14   Adzhubei, I., Jordan, D.M. and Sunyaev, S.R. (2013) Predicting functional effect of human missense mutations using polyPhen-2. *Curr. Protoc. Hum. Genet.* January 2013, Supplement 76, 7.20.1-7.20.41 https://doi.org/10.1002/0471142905.hg0720s76

15   Carter, H., Douville, C., Stenson, P.D., Cooper, D.N. and Karchin, R. (2013) Identifying Mendelian disease genes with the variant effect scoring tool. *BMC Genomics* **14**, S3 https://doi.org/10.1186/1471-2164-14-S3-S3

16   Reva, B., Antipin, Y. and Sander, C. (2011) Predicting the functional impact of protein mutations: application to cancer genomics. *Nucleic Acids Res.* **39**, E118–EU85 https://doi.org/10.1093/nar/gkr407

17   Carter, H., Chen, S.N., Isik, L., Tyekucheva, S., Velculescu, V.E., Kinzler, K.W. et al. (2009) Cancer-specific high-throughput annotation of somatic mutations: computational prediction of driver missense mutations. *Cancer Res.* **69**, 6660–6667 https://doi.org/10.1158/0008-5472.CAN-09-1133

18   Gonzalez-Perez, A., Deu-Pons, J. and Lopez-Bigas, N. (2012) Improving the prediction of the functional impact of cancer mutations by baseline tolerance transformation. *Genome Med.* **4**, 89 https://doi.org/10.1186/gm390

19   Mao, Y., Chen, H., Liang, H., Meric-Bernstam, F., Mills, G.B. and Chen, K. (2013) Candra: cancer-specific driver missense mutation annotation with optimized features. *PLoS ONE* **8**, e77945 https://doi.org/10.1371/journal.pone.0077945

20   Shihab, H.A., Gough, J., Cooper, D.N., Stenson, P.D., Barker, G.L.A., Edwards, K.J. et al. (2013) Predicting the functional, molecular, and phenotypic consequences of amino acid substitutions using hidden markov models. *Hum. Mutat.* **34**, 57–65 https://doi.org/10.1002/humu.22225

21   Lawrence, M.S., Stojanov, P., Polak, P., Kryukov, G.V., Cibulskis, K., Sivachenko, A. et al. (2013) Mutational heterogeneity in cancer and the search for new cancer-associated genes. *Nature* **499**, 214–218 https://doi.org/10.1038/nature12213

22   Dietlein, F., Weghorn, D., Taylor-Weiner, A., Richters, A., Reardon, B., Liu, D. et al. (2020) Identification of cancer driver genes based on nucleotide context. *Nat. Genet.* **52**, 208–218 https://doi.org/10.1038/s41588-019-0572-y

23   Bailey, M.H., Tokheim, C., Porta-Pardo, E., Sengupta, S., Bertrand, D., Weerasinghe, A. et al. (2018) Comprehensive characterization of cancer driver genes and mutations (vol 173, 371.e1, 2018). *Cell* **174**, 1034–1035 https://doi.org/10.1016/j.cell.2018.07.034

24   Hart, S.N., Polley, E.C., Shimelis, H., Yadav, S. and Couch, F.J. (2020) Prediction of the functional impact of missense variants in BRCA1 and BRCA2 with BRCA-ML. *NPJ Breast Cancer.* **6**, 13 https://doi.org/10.1038/s41523-020-0159-x

25   Martinez-Jimenez, F., Muinos, F., Sentis, I., Deu-Pons, J., Reyes-Salazar, I., Arnedo-Pac, C. et al. (2020) A compendium of mutational cancer driver genes. *Nat. Rev. Cancer* **20**, 555–572 https://doi.org/10.1038/s41568-020-0290-x

26   Leaman, R. and Gonzalez, G. (2008) BANNER: an executable survey of advances in biomedical named entity recognition. *Pac. Symp. Biocomput.* **13**, 652–663

27   Settles, B. (2005) ABNER: an open source tool for automatically tagging genes, proteins and other entity names in text. *Bioinformatics* **21**, 3191–3192 https://doi.org/10.1093/bioinformatics/bti475

28   Habibi, M., Weber, L., Neves, M., Wiegandt, D.L. and Leser, U. (2017) Deep learning with word embeddings improves biomedical named entity recognition. *Bioinformatics* **33**, i37–i48 https://doi.org/10.1093/bioinformatics/btx228

29   Lee, S., Kim, D., Lee, K., Choi, J., Kim, S., Jeon, M. et al. (2016) BEST: next-generation biomedical entity search tool for knowledge discovery from biomedical literature. *PLoS ONE* **11**, e0164680 https://doi.org/10.1371/journal.pone.0164680

30   Lee, K., Kim, B., Choi, Y., Kim, S., Shin, W., Lee, S. et al. (2018) Deep learning of mutation-gene-drug relations from the literature. *BMC Bioinformatics* **19**, 21 https://doi.org/10.1186/s12859-018-2029-1

31   Lever, J., Jones, M.R., Danos, A.M., Krysiak, K., Bonakdar, M., Grewal, J.K. et al. (2019) Text-mining clinically relevant cancer biomarkers for curation into the CIViC database. *Genome Med.* **11**, 78 https://doi.org/10.1186/s13073-019-0686-y

32   Patel, N.M., Michelini, V.V., Snell, J.M., Balu, S., Hoyle, A.P., Parker, J.S. et al. (2018) Enhancing next-generation sequencing-guided cancer care through cognitive computing. *Oncologist* **23**, 179–185 https://doi.org/10.1634/theoncologist.2017-0170

33   Bashashati, A., Haffari, G., Ding, J., Ha, G., Lui, K., Rosner, J. et al. (2012) Drivernet: uncovering the impact of somatic driver mutations on transcriptional networks in cancer. *Genome Biol.* **13**, R124 https://doi.org/10.1186/gb-2012-13-12-r124

34   Bertrand, D., Chng, K.R., Sherbaf, F.G., Kiesel, A., Chia, B.K., Sia, Y.Y. et al. (2015) Patient-specific driver gene prediction and risk assessment through integrated network analysis of cancer omics profiles. *Nucleic Acids Res.* **43**, e44 https://doi.org/10.1093/nar/gku1393

35   Dong, C., Guo, Y., Yang, H., He, Z., Liu, X. and Wang, K. (2016) iCAGES: integrated CAncer GEnome score for comprehensively prioritizing driver genes in personal cancer genomes. *Genome Med.* **8**, 135 https://doi.org/10.1186/s13073-016-0390-0

36   Kalari, K.R., Sinnwell, J.P., Thompson, K.J., Tang, X., Carlson, E.E., Yu, J. et al. (2018) PANOPLY: Omics-guided drug prioritization method tailored to an individual patient. *JCO Clin. Cancer Inform.* **2**, 1–11 https://doi.org/10.1200/CCI.18.00012

37   Nulsen, J., Misetic, H., Yau, C. and Ciccarelli, F.D. (2021) Pan-cancer detection of driver genes at the single-patient resolution. *Genome Med.* **13**, 12 https://doi.org/10.1186/s13073-021-00830-0

38   Administration FaD. FDA

39   Network NCC. NCCN Compendia

40   ClinicalTrials.gov. ClinicalTrials.gov [Available from: https://clinicaltrials.gov]

41   Miron, L., Goncalves, R.S. and Musen, M.A. (2020) Obstacles to the reuse of study metadata in clinicalTrials.gov. *Sci. Data* **7**, 443 https://doi.org/10.1038/s41597-020-00780-z

42   Jain, N., Mittendorf, K.F., Holt, M., Lenoue-Newton, M., Maurer, I., Miller, C. et al. (2020) The My cancer genome clinical trial data model and trial curation workflow. *J. Am. Med. Inform. Assoc.* **27**, 1057–1066 https://doi.org/10.1093/jamia/ocaa066

43   Zeng, J., Shufean, M.A., Khotskaya, Y., Yang, D., Kahle, M., Johnson, A. et al. (2019) OCTANE: oncology clinical trial annotation engine. *JCO Clin. Cancer Inform.* **3**, 1–11 https://ascopubs.org/doi/10.1200/CCI.18.00145

44   Antidote. Antidote [Available from: https://www.antidote.me]

45   Bio M. Synergy

46   Deep6.ai. Deep6

47   Beck, J.T., Rammage, M., Jackson, G.P., Preininger, A.M., Dankwa-Mullan, I., Roebuck, M.C. et al. (2020) Artificial intelligence tool for optimizing eligibility screening for clinical trials in a large community cancer center. *JCO Clin. Cancer Inform.* **4**, 50–59 https://doi.org/10.1200/CCI.19.00079

48   Mendel.ai. Mendel.ai

49   Calaprice-Whitty, D., Galil, K., Salloum, W., Zariv, A. and Jimenez, B. (2020) Improving clinical trial participant prescreening with artificial intelligence (AI): a comparison of the results of AI-assisted vs standard methods in 3 oncology trials. *Ther. Innov. Regul. Sci.* **54**, 69–74 https://doi.org/10.1007/s43441-019-00030-4

50   Alexander, M., Solomon, B., Ball, D.L., Sheerin, M., Dankwa-Mullan, I., Preininger, A.M. et al. (2020) Evaluation of an artificial intelligence clinical trial matching system in Australian lung cancer patients. *JAMIA Open* **3**, 209–215 https://doi.org/10.1093/jamiaopen/ooaa002

51   Haddad, T., Helgeson, J.M., Pomerleau, K.E., Preininger, A.M., Roebuck, M.C., Dankwa-Mullan, I. et al. (2021) Accuracy of an artificial intelligence system for cancer clinical trial eligibility screening: retrospective pilot study. *JMIR Med. Inform.* **9**, e27767 https://doi.org/10.2196/27767

52   Yuan, C., Ryan, P.B., Ta, C., Guo, Y., Li, Z., Hardin, J. et al. (2019) Criteria2query: a natural language interface to clinical databases for cohort definition. *J. Am. Med. Inform. Assoc.* **26**, 294–305 https://doi.org/10.1093/jamia/ocy178

53   Kang, T., Zhang, S., Tang, Y., Hruby, G.W., Rusanov, A., Elhadad, N. et al. (2017) EliIE: An open-source information extraction system for clinical trial eligibility criteria. *J. Am. Med. Inform. Assoc.* **24**, 1062–1071 https://doi.org/10.1093/jamia/ocx019

54   Liu, C., Yuan, C., Butler, A.M., Carvajal, R.D., Li, Z.R., Ta, C.N. et al. (2019) DQueST: dynamic questionnaire for search of clinical trials. *J. Am. Med. Inform. Assoc.* **26**, 1333–1343 https://doi.org/10.1093/jamia/ocz121

55   Chapman, W.W., Bridewell, W., Hanbury, P., Cooper, G.F. and Buchanan, B.G. (2001) A simple algorithm for identifying negated findings and diseases in discharge summaries. *J. Biomed. Inform.* **34**, 301–310 https://doi.org/10.1006/jbin.2001.1029

56   Miotto, R. and Weng, C. (2015) Case-based reasoning using electronic health records efficiently identifies eligible patients for clinical trials. *J. Am. Med. Inform. Assoc.* **22**, e141–e150 https://doi.org/10.1093/jamia/ocu050

57   Cruz Rivera, S., Liu, X., Chan, A.W., Denniston, A.K., Calvert, M.J., Spirit, A.I. et al. (2020) Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. *Nat. Med.* **26**, 1351–1363 https://doi.org/10.1038/s41591-020-1037-7

58   Faes, L., Liu, X., Wagner, S.K., Fu, D.J., Balaskas, K., Sim, D.A. et al. (2020) A clinician's guide to artificial intelligence: how to critically appraise machine learning studies. *Transl. Vis. Sci. Technol.* **9**, 7 https://doi.org/10.1167/tvst.9.2.7