



Perspective

Precision Medicine 3.0: Integrating Genomic Stratification with Adaptive Therapeutic Trajectories

Hui-Qi Qu ¹  and Hakon Hakonarson ^{1,2,3,4,5,*}

¹ The Center for Applied Genomics, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA; quh@chop.edu

² Department of Pediatrics, The Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA

³ Division of Human Genetics, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

⁴ Division of Pulmonary Medicine, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

⁵ Faculty of Medicine, University of Iceland, 101 Reykjavik, Iceland

* Correspondence: hakonarson@chop.edu; Tel.: +1-267-426-0088

Abstract

Complex diseases require therapeutic decisions that adapt as the disease evolves over time and in response to treatment. Current precision medicine remains largely centered on selecting an individual drug or predefined regimen from molecular characteristics measured at a particular time. The next evolution of precision medicine is toward adaptive therapeutic trajectories in which treatment decisions change over the course of treatment. Germline genomic information remains constant and contributes baseline susceptibility and pharmacogenomic information. By contrast, somatic alterations, multi-omic profiles, clinical state, and treatment response may change during disease progression and therapy and can inform subsequent treatment decisions. The volume and complexity of longitudinal molecular and clinical data exceed what can be integrated reliably by unaided human cognition. Artificial intelligence (AI) may help integrate these high-dimensional longitudinal data. The scientific challenge is whether AI-derived treatment recommendations are reliable and improve patient outcomes. This model raises a regulatory problem beyond conventional drug approval: how to evaluate a bounded therapeutic system in which patient-specific treatment recommendations change as the patient's condition evolves.

Keywords: precision medicine; genomics; multi-omics; artificial intelligence; adaptive therapeutics; dynamic treatment regimes; causal inference



Academic Editor: Maria Gazouli

Received: 18 August 2026

Revised: 12 September 2026

Accepted: 15 September 2026

Published: 16 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Precision Medicine Beyond Static Molecular Profiling

1.1. Limits of Static Molecular Profiling

Precision medicine uses biological differences among patients to guide treatment [1]. Germline genetics can identify genetic susceptibility and variation that affects drug response [2]. Somatic sequencing can identify acquired disease drivers [3]. Transcriptomic, proteomic, metabolomic, epigenomic, and immune profiling can reveal dynamic molecular changes beyond the genome that are not apparent from the clinical phenotype [4]. In practice, however, treatment usually remains static: a molecular feature is identified and used to select a drug or predefined regimen. This approach works best when one mechanism largely determines treatment response. It is less suited to complex diseases in which several mechanisms contribute or change during the course of disease and treatment.

For complex diseases, a susceptibility locus may contain several plausible genes and regulatory elements. Fine-mapping can reduce the number of candidate variants. Colocalization with eQTL signals, chromatin accessibility, and enhancer-promoter maps can connect noncoding variants to candidate target genes [5]. In the retinal pigment epithelium, integrated analysis of chromatin accessibility, transcription-factor binding, and gene-regulatory networks has identified regulatory programs associated with cell fate and disease-related pathways that may provide targets for pharmacogenomic or genome-editing strategies [6]. CRISPR perturbation can test whether disrupting the implicated sequence changes gene expression or a disease-associated cellular phenotype [7]. For coding variants, AI-based models such as AlphaMissense can prioritize missense substitutions likely to alter protein function by integrating sequence conservation with structural context [8]. AlphaFold 3 predicts biomolecular complexes involving proteins, nucleic acids, and small molecules and can be used to examine candidate target structures and molecular interactions before experimental testing [9]. Functional validation can inform the choice among small-molecule, RNA-based, and genome-editing approaches [10]. When no established pharmacologically targetable driver is present, transcriptomic and single-cell analyses can identify differentially regulated genes for functional prioritization. In epithelial ovarian cancer, Xu et al. integrated single-cell and bulk transcriptomic data to identify *ITGA1* as associated with neoadjuvant chemotherapy response. Knockdown of *ITGA1* restored cisplatin sensitivity in resistant cells, providing functional evidence for prioritizing a differentially regulated gene before therapeutic targeting [11]. Genes whose perturbation reverses a disease-associated or treatment-resistant phenotype may be considered for modality-specific development, including RNA-based targeting. Recent generative-AI studies have produced model-designed compounds that were subsequently synthesized and tested experimentally [12].

The relative contribution of disease mechanisms differs between patients with the same diagnosis and changes within the same patient over time. Disease mechanisms reflect the interaction between genetic predisposition and the patient's changing physiological state. Treatment alters this balance by suppressing some disease processes while potentially allowing others to become more important over time. An actionable molecular finding does not necessarily identify the treatment that will provide the greatest clinical benefit. The same molecular finding may have different treatment implications at different stages of disease or after prior therapy. The clinical value of molecularly matched treatment has been tested in precision oncology. The prospective I-PREDICT precision-oncology study demonstrated the feasibility of constructing individualized multidrug regimens according to the molecular alterations identified in each patient's tumor [13]. The WINTHER precision-oncology trial broadened molecular matching by incorporating tumor transcriptomic information alongside genomic alterations [14]. In contrast, the randomized phase 2 SHIVA trial found that molecularly matched off-label therapy did not improve progression-free survival over physician-selected treatment in heavily pretreated patients [15]. Molecular matching can identify plausible interventions without establishing greater clinical benefit.

1.2. From Molecular Profiling to Sequential Decision-Making

Dynamic treatment regimes (DTRs) provide a statistical framework for studying sequential treatment decisions in which later treatment depends on the patient's evolving clinical information [16]. Q-learning, A-learning, and related methods have been developed to estimate such treatment rules from clinical data [17]. These approaches provide an established basis for adaptive treatment, although their use in routine clinical care remains limited. Extending DTRs to genome-enabled precision medicine introduces an additional challenge because the longitudinal molecular states used to guide treatment must be

informative and measured reliably enough. Treatment rules based on these states and prior response must show reproducible performance and clinical validity across patients and clinical settings.

Precision medicine should extend treatment selection from a single time point to repeated decisions informed by longitudinal molecular and clinical information. Genome biotechnology now allows repeated characterization of treatment-relevant changes that were previously difficult to observe. Sequential decision methods and artificial intelligence (AI) provide tools for incorporating these data into subsequent treatment choices. The question addressed in this Perspective is how genome-enabled precision medicine can move beyond molecularly guided treatment selection toward treatment strategies that adapt over the course of disease. The molecular basis for such adaptation is currently most developed in oncology, where somatic evolution and repeated tumor profiling can inform subsequent treatment. In other clinical settings, adaptive treatment may rely more heavily on changing physiological information combined with fixed germline genomic information.

1.3. Precision Medicine 1.0, 2.0, and 3.0

We propose three stages in the evolution of precision medicine based on the structure of the clinical decision (Figure 1). Precision Medicine 1.0 represents molecularly informed selection of a single intervention for an individual patient. The defining decision is which intervention to select from available alternatives at a given decision point. This model underlies much of pharmacogenomics and molecularly targeted therapy, in which germline or somatic characteristics are used to guide drug selection [1].




Figure 1. Evolution from static molecular matching to adaptive precision medicine. Precision Medicine 1.0 selects an individual intervention from molecular characteristics measured at a decision point. Precision Medicine 2.0 extends this approach to individualized combination therapy. Precision Medicine 3.0 links longitudinal molecular and clinical states with prior treatment history and response to guide sequential clinical decisions over time. AI-assisted integration, causal inference, physician oversight, prospective validation, and defined regulatory boundaries support implementation of the adaptive decision system. The arrows indicate the sequential decision process. Created in BioRender. Qu, H. (2026). <https://BioRender.com/gd0kyut>, accessed on 8 August 2026.

Pharmacogenomics is an example of how fixed germline information can remain actionable across a treatment course. Germline genotype does not change, but its implications for drug selection or dosing may recur as different therapies are considered. Variants

affecting drug metabolism, transport, or toxicity can affect treatment decisions at multiple points in response to longitudinal clinical signals. Examples are genotype-informed drug selection or dosing for fluoropyrimidines according to *DPYD*, thiopurines according to *TPMT* and *NUDT15*, and clopidogrel according to *CYP2C19* [18–20].

Precision Medicine 2.0 extends individualization from a single intervention to a patient-specific combination of interventions. The clinical problem is which interventions should be used together when multiple actionable mechanisms are present in the same patient. I-PREDICT exemplifies Precision Medicine 2.0 and provides a clinical antecedent for this concept: patients with advanced cancers received individualized combination regimens constructed according to the molecular alterations identified in their tumors [13].

Precision Medicine 3.0 changes the decision structure from individualization at one decision point to individualization across a sequence of decisions. The appropriate action at a later point depends on the patient's current molecular and clinical state, previous treatments, and response to those treatments. Dynamic treatment regimes enable this type of sequential decision-making by mapping updated patient information and treatment history to subsequent treatment choices [16]. In particular, this adaptive precision medicine framework incorporates longitudinal somatic genomic and multi-omic measurements into the information used to guide subsequent clinical actions, with germline information serving as a baseline genomic input. The proposed taxonomy is a clinical framework instead of a new statistical formulation of DTRs or reinforcement learning. It is similar to the distinction between evidence-based medicine as a clinical framework and meta-analysis as one major methodological tool. Its central feature is the individualization of the treatment trajectory, which consequently creates requirements for prospective validation and regulatory evaluation of the bounded decision system. More than new terminology, it defines an area and direction of precision medicine.

2. Longitudinal Molecular State in Precision Medicine

In this Perspective, molecular state refers to molecular features that may change during disease progression or treatment [21,22]. Germline genomic information remains fixed over the treatment course. Gene expression, protein abundance, metabolism, and immune activity vary with disease progression, treatment, and other physiological changes. Repeated molecular characterization is technically feasible, but the clinical value of these changes depends on whether they inform subsequent treatment. The prospective MATCH-R study shows that repeated tumor biopsies obtained after acquired resistance identified molecular changes that emerged during therapy and could inform subsequent treatment choices [23]. In precision oncology, germline variants can contribute inherited susceptibility and pharmacogenomic information, while the baseline tumor genome defines somatic alterations that may guide initial treatment. Structural variation may be an important component of this baseline genomic state. In retinoblastoma, whole-genome sequencing has identified structural rearrangements disrupting *RB1*, including complex rearrangements and chromothripsis-like events that may be missed by conventional mutation testing [24,25]. Long-read sequencing can improve detection of *RB1* structural variants [26]. Subsequent clonal evolution and acquired resistance can change the somatic genomic profile used for later treatment decisions. The evidence for longitudinal molecular profiling is currently strongest in oncology. Serial tumor biopsy, circulating tumor DNA, and related measurements can capture somatic evolution during treatment and acquired resistance. These repeated measurements are not routinely feasible for many non-oncologic diseases. Adaptive treatment may instead combine fixed germline information with longitudinal clinical measurements, treatment response, and toxicity. For instance, a pharmacogenomic variant may become actionable only when a corresponding drug is considered.

The same molecular finding may carry different implications depending on when it is measured. A variant or expression pattern detected before treatment is interpreted in the context of untreated disease, whereas the same finding after therapy must be considered together with response, toxicity, and changes induced by prior treatment. The main difficulty is determining which longitudinal changes matter for subsequent treatment. A marker may change with disease activity without presenting a targetable mechanism, and treatment-induced changes may reflect response without indicating what should be given next. While a molecular finding suggests a treatment option, its clinical benefit must still be established, as shown by the contrasting results of I-PREDICT, WINTHER, and SHIVA [13–15].

Repeated measurements introduce an additional requirement: the same state must be recognized reliably across time and settings. High-throughput assays remain sensitive to batch effects and other technical variation [27]. Epigenomic measurements introduce additional challenges in assay standardization and interpretation, as continuous, tissue- and cell-state-specific measurements can vary across platforms, laboratories, and analytical pipelines [28]. Tissue source and sampling interval further influence what is measured. Comprehensive serial multi-omic profiling remains costly, particularly when repeated sampling and multiple assay platforms are required. In many non-oncologic diseases, repeated sampling of the affected tissue is not practical or needed, and peripheral samples may not represent disease processes in less accessible tissues, such as the brain in neurodegenerative diseases. Longitudinal monitoring in these settings will depend more often on clinically accessible molecular assays and physiological or clinical measurements.

3. Building Adaptive Treatment Trajectories

Building an adaptive treatment trajectory requires linking the patient's current molecular and clinical state with prior treatment history to determine the next clinical action. Causal inference is needed to estimate which available action is expected to improve outcome. Clinical use requires repeated application of these decisions within defined treatment options and safety limits.

3.1. Causal Inference in Sequential Treatment Decisions

A model may predict how a patient will respond to the treatment they receive, but this does not show whether another treatment would have produced a better outcome. This problem becomes more difficult when treatment decisions are made repeatedly over time. A biomarker may identify patients at higher risk, but changing therapy on the basis of that biomarker requires evidence that the alternative treatment improves outcome. Sequential causal inference must account for time-varying confounding. A patient's current state may influence the next treatment choice while also reflecting the effects of prior treatment. G-methods can estimate causal effects under these conditions [29]. Naimi et al. illustrated this problem using a hypothetical HIV treatment example in which therapy was assigned at baseline and again during follow-up. HIV viral load measured before the second treatment was affected by prior therapy and also influenced subsequent treatment and final CD4 count. Standard regression did not recover the true combined treatment effect, whereas inverse probability weighting, g-estimation, and the g-formula each recovered the specified effect of 50 CD4 cells/mm³ [30]. For the same treatment sequence, a target-trial analysis would define the treatment strategies being compared and follow patients from the corresponding decision points [31].

Q-learning and A-learning are reinforcement-learning approaches that can estimate DTRs from updated patient information [16,17]. The Artificial Intelligence Clinician study used retrospective intensive-care data to derive fluid and vasopressor strategies for patients

with sepsis [32]. The model identified treatment policies associated with better observed outcomes, but the analysis could not establish that applying those policies prospectively would improve survival. Reviews of reinforcement learning in healthcare have highlighted the same limitation: the observed data contain only the treatments that were actually given, while outcomes under alternative choices remain unobserved [33].

This limitation is compounded in observational data because treatment choice reflects the patient's condition. A clinician may intensify therapy because the patient is deteriorating or reduce treatment because the patient is improving. The resulting association between treatment and outcome may therefore reflect the reason the treatment was chosen instead of the effect of the treatment itself. Target-trial methodology addresses this problem by defining the treatment strategies, eligibility criteria, follow-up, outcomes, and causal comparison that an observational analysis is intended to emulate [31]. Even when the causal comparison is well defined, the objective of the treatment policy must be specified. A policy that improves a molecular endpoint may worsen toxicity or quality of life.

3.2. Operationalizing Adaptive Treatment Trajectories

An adaptive trajectory requires each treatment decision to be linked to what happened after the previous one. A patient may continue the same therapy when disease control is maintained, change therapy when response is lost, or avoid an otherwise active drug when toxicity becomes limiting. The next decision is based on current findings together with prior treatment history and response.

This creates a practical role for AI. Longitudinal molecular and clinical data may exceed what can be integrated consistently by unaided clinical reasoning, particularly when several previous treatments influence the available choices. AI is needed to integrate these data and compare candidate treatment options. High-dimensional learning methods can incorporate large numbers of genetic, clinical, and longitudinal variables when estimating dynamic treatment regimes [34]. At each decision point, the model would combine the patient's current state with prior treatment history and response to compare the available therapeutic options. Candidate interventions should be prioritized according to the strength of clinical and pharmacological evidence and their consistency with the patient's current molecular and clinical state. Molecular plausibility alone is insufficient to justify a treatment change when comparative clinical benefit cannot be estimated reliably. Recent I-PREDICT data suggest that molecular coverage and dose intensity should be treated as separate dimensions of combination therapy. The matching score quantified the proportion of pathogenic alterations therapeutically targeted, and greater molecular matching was associated with higher disease control and longer progression-free and overall survival. Outcomes were not associated with the number of drugs or their doses. Dose should be optimized according to pharmacological activity and tolerability [35].

Baseline tumor profiling defines the dominant driver architecture, but treatment imposes selective pressure that can reshape the clonal composition of the tumor. At progression, re-biopsy or circulating tumor DNA may identify newly selected resistance alterations, loss of the original dependency, or activation of a bypass pathway that was not dominant at diagnosis [23,36]. These molecular findings should be interpreted together with prior drug exposure and the duration of response to determine whether the next treatment should continue to target the original pathway, add a second target, or move to a different therapeutic strategy.

The available treatment choices must be defined before clinical use. The system should operate within specified drugs, permitted modifications, and safety limits. Essential physician review includes interpretation of model recommendations, assessment of missing or atypical clinical information, and evaluation of expected benefit and risk.

4. Prospective Validation of the Decision Policy

Validation must test the decision policy as a whole. Patients guided by the same policy may receive different treatments because their molecular and clinical states differ. A prospective trial could compare policy-guided care with standard clinical decision-making while allowing the treatment sequence within the policy-guided group to vary. The main outcome is whether use of the policy improves clinical outcomes.

Sequential decisions create a specific validation problem because an early error may affect later decisions. An incorrect assessment of patient state may lead to the wrong treatment, and that treatment may change the information used at the next decision point. Validation must examine performance across the full treatment sequence. This becomes particularly important when genomic information contributes to treatment selection, because genomic datasets remain uneven across populations [37].

In addition, a previously validated policy may require revalidation as the clinical setting changes. New drugs may alter the available choices, new assays may change how patient state is measured, and changes in clinical practice may shift the population in which the policy is used. Revalidation is needed when these changes affect the policy's recommendations or clinical performance.

5. FDA Regulation of a Bounded Adaptive System

As longitudinal molecular and clinical information becomes more complex, AI-assisted decision systems will be needed to integrate these data and support sequential treatment decisions. The individual drugs would remain subject to their existing regulatory requirements. Additional evidence would be needed to show that the decision system selects and modifies treatment safely within its defined scope. Regulatory evaluation of such an adaptive system requires the decision space to be defined prospectively. The molecular findings that can alter treatment, the permitted therapeutic changes, and the evidence supporting those changes should be specified before clinical use. Changes to the policy, available drugs, molecular inputs, assays, or safety boundaries that alter the validated scope of the system would require reassessment of clinical performance and safety. Physician review remains necessary when the observed resistance mechanism falls outside the validated decision space.

Current U.S. Food and Drug Administration (FDA) regulation evaluates drugs, biological products, devices, and combination products under established pathways. FDA's combination-product framework addresses products that combine drugs, devices, or biological products [38]. AI-enabled medical-device software is regulated through existing device pathways, and FDA's January 2025 draft guidance addresses lifecycle risk management for such software [39]. An adaptive treatment system creates a different regulatory problem because patient-specific pharmacological recommendations may change with the patient's state. FDA's guidance on Predetermined Change Control Plans for Artificial Intelligence-Enabled Device Software Functions provides a relevant precedent [40]. A predetermined change control plan allows specified software modifications to be described and evaluated before those modifications are implemented. The regulatory object remains the AI-enabled device. In an adaptive treatment system, the treatment recommendation may change while the software remains unchanged. This distinction becomes important when approved drugs are combined, sequenced, or modified in ways that were not evaluated as fixed regimens.

Independent regulatory review of every possible patient-specific sequence would be impractical. This Perspective proposes that regulation should apply to the adaptive decision system as an integrated clinical unit. Evidence-based evaluation should establish the comparative effectiveness of the policy across the patient states for which it is intended.

Prospective comparative trials should determine whether policy-guided care improves clinically meaningful outcomes relative to standard clinical decision-making while allowing recommendations to change with the patient's evolving state. With prospective validation and an appropriate regulatory framework, Precision Medicine 3.0 could move from a conceptual model to clinical implementation.

Author Contributions: H.-Q.Q.: Conceptualization, Investigation, Writing—original draft, Writing—review and editing. H.H.: Conceptualization, Funding acquisition, Writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: The study was supported by the Institutional Development Funds from the Children's Hospital of Philadelphia to the Center for Applied Genomics, and the Children's Hospital of Philadelphia Endowed Chair in Genomic Research to H.H.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
DTRs	Dynamic treatment regimes
FDA	U.S. Food and Drug Administration

References

1. Hamburg, M.A.; Collins, F.S. The path to personalized medicine. *N. Engl. J. Med.* **2010**, *363*, 301–304. [[CrossRef](#)]
2. Evans, W.E.; Relling, M.V. Moving towards individualized medicine with pharmacogenomics. *Nature* **2004**, *429*, 464–468. [[CrossRef](#)]
3. Bailey, M.H.; Tokheim, C.; Porta-Pardo, E.; Sengupta, S.; Bertrand, D.; Weerasinghe, A.; Colaprico, A.; Wendl, M.C.; Kim, J.; Reardon, B.; et al. Comprehensive Characterization of Cancer Driver Genes and Mutations. *Cell* **2018**, *173*, 371–385.e18. [[CrossRef](#)]
4. Babu, M.; Snyder, M. Multi-Omics Profiling for Health. *Mol. Cell. Proteom.* **2023**, *22*, 100561. [[CrossRef](#)]
5. Strober, B.J.; Zhang, M.J.; Amariuta, T.; Rossen, J.; Price, A.L. Fine-mapping causal tissues and genes at disease-associated loci. *Nat. Genet.* **2025**, *57*, 42–52. [[CrossRef](#)]
6. Zibetti, C. Molecular Determinants of the Human Retinal Pigment Epithelium Cell Fate and Potential Pharmacogenomic Targets for Precision Medicine. *Int. J. Mol. Sci.* **2025**, *26*, 5817. [[CrossRef](#)]
7. Kim, H.S.; Kweon, J.; Kim, Y. Recent advances in CRISPR-based functional genomics for the study of disease-associated genetic variants. *Exp. Mol. Med.* **2024**, *56*, 861–869. [[CrossRef](#)]
8. Cheng, J.; Novati, G.; Pan, J.; Bycroft, C.; Žemgulytė, A.; Applebaum, T.; Pritzel, A.; Wong, L.H.; Zielinski, M.; Sargeant, T.; et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. *Science* **2023**, *381*, eadg7492. [[CrossRef](#)]
9. Abramson, J.; Adler, J.; Dunger, J.; Evans, R.; Green, T.; Pritzel, A.; Ronneberger, O.; Willmore, L.; Ballard, A.J.; Bambrick, J.; et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **2024**, *630*, 493–500. [[CrossRef](#)]
10. Zheng, Y.; Li, Y.; Zhou, K.; Li, T.; VanDusen, N.J.; Hua, Y. Precise genome-editing in human diseases: Mechanisms, strategies and applications. *Signal Transduct. Target. Ther.* **2024**, *9*, 47. [[CrossRef](#)]
11. Xu, X.; Yang, R.; Fang, L.; Sun, M.; Ma, Y.; Zhang, Y.; Wu, C. A Multi-Omics and Single-Cell Framework Identifies ITGA1 as a Candidate Predictive Biomarker of Neoadjuvant Chemotherapy Response in Epithelial Ovarian Cancer. *BioFactors* **2026**, *52*, e70135. [[CrossRef](#)]
12. Swanson, K.; Liu, G.; Catacutan, D.B.; Arnold, A.; Zou, J.; Stokes, J.M. Generative AI for designing and validating easily synthesizable and structurally novel antibiotics. *Nat. Mach. Intell.* **2024**, *6*, 338–353. [[CrossRef](#)]
13. Sicklick, J.K.; Kato, S.; Okamura, R.; Schwaederle, M.; Hahn, M.E.; Williams, C.B.; De, P.; Krie, A.; Piccioni, D.E.; Miller, V.A. Molecular profiling of cancer patients enables personalized combination therapy: The I-PREDICT study. *Nat. Med.* **2019**, *25*, 744–750. [[CrossRef](#)]
14. Rodon, J.; Soria, J.-C.; Berger, R.; Miller, W.H.; Rubin, E.; Kugel, A.; Tsimberidou, A.; Saintigny, P.; Ackerstein, A.; Braña, I. Genomic and transcriptomic profiling expands precision cancer medicine: The WINTHER trial. *Nat. Med.* **2019**, *25*, 751–758. [[CrossRef](#)]

15. Le Tourneau, C.; Delord, J.-P.; Gonçalves, A.; Gavoille, C.; Dubot, C.; Isambert, N.; Campone, M.; Trédan, O.; Massiani, M.-A.; Mauborgne, C. Molecularly targeted therapy based on tumour molecular profiling versus conventional therapy for advanced cancer (SHIVA): A multicentre, open-label, proof-of-concept, randomised, controlled phase 2 trial. *Lancet Oncol.* **2015**, *16*, 1324–1334. [\[CrossRef\]](#)
16. Murphy, S.A. Optimal dynamic treatment regimes. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2003**, *65*, 331–355. [\[CrossRef\]](#)
17. Schulte, P.J.; Tsiatis, A.A.; Laber, E.B.; Davidian, M. Q- and A-learning Methods for Estimating Optimal Dynamic Treatment Regimes. *Stat. Sci.* **2014**, *29*, 640–661. [\[CrossRef\]](#)
18. Amstutz, U.; Henricks, L.M.; Offer, S.M.; Barbarino, J.; Schellens, J.H.M.; Swen, J.J.; Klein, T.E.; McLeod, H.L.; Caudle, K.E.; Diasio, R.B.; et al. Clinical Pharmacogenetics Implementation Consortium (CPIC) Guideline for Dihydropyrimidine Dehydrogenase Genotype and Fluoropyrimidine Dosing: 2017 Update. *Clin. Pharmacol. Ther.* **2018**, *103*, 210–216. [\[CrossRef\]](#)
19. Maillard, M.; Schwab, M.; Whirl-Carrillo, M.; Moyer, A.M.; Suarez-Kurtz, G.; Pui, C.H.; Stein, C.M.; Klein, T.E.; Spahn, C.; Kwon, S.; et al. Clinical Pharmacogenetics Implementation Consortium (CPIC) Guideline for Thiopurine Dosing Based on TPMT and NUDT15 Genotypes: 2025 Update. *Clin. Pharmacol. Ther.* **2026**, *119*, 916–927. [\[CrossRef\]](#)
20. Lee, C.R.; Luzum, J.A.; Sangkuhl, K.; Gammal, R.S.; Sabatine, M.S.; Stein, C.M.; Kisor, D.F.; Limdi, N.A.; Lee, Y.M.; Scott, S.A.; et al. Clinical Pharmacogenetics Implementation Consortium Guideline for CYP2C19 Genotype and Clopidogrel Therapy: 2022 Update. *Clin. Pharmacol. Ther.* **2022**, *112*, 959–967. [\[CrossRef\]](#)
21. Chen, R.; Mias, G.I.; Li-Pook-Than, J.; Jiang, L.; Lam, H.Y.; Chen, R.; Miriami, E.; Karczewski, K.J.; Hariharan, M.; Dewey, F.E.; et al. Personal omics profiling reveals dynamic molecular and medical phenotypes. *Cell* **2012**, *148*, 1293–1307. [\[CrossRef\]](#)
22. Schüssler-Fiorenza Rose, S.M.; Contrepois, K.; Moneghetti, K.J.; Zhou, W.; Mishra, T.; Mataraso, S.; Dagan-Rosenfeld, O.; Ganz, A.B.; Dunn, J.; Hornburg, D.; et al. A longitudinal big data approach for precision health. *Nat. Med.* **2019**, *25*, 792–804. [\[CrossRef\]](#)
23. Recondo, G.; Mahjoubi, L.; Maillard, A.; Lorient, Y.; Bigot, L.; Facchinetti, F.; Bahleda, R.; Gazzah, A.; Hollebecque, A.; Mezquita, L.; et al. Feasibility and first reports of the MATCH-R repeated biopsy trial at Gustave Roussy. *npj Precis. Oncol.* **2020**, *4*, 27. [\[CrossRef\]](#)
24. McEvoy, J.; Nagahawatte, P.; Finkelstein, D.; Richards-Yutz, J.; Valentine, M.; Ma, J.; Mullighan, C.; Song, G.; Chen, X.; Wilson, M.; et al. RB1 gene inactivation by chromothripsis in human retinoblastoma. *Oncotarget* **2014**, *5*, 438–450. [\[CrossRef\]](#)
25. Davies, H.R.; Broad, K.D.; Onadim, Z.; Price, E.A.; Zou, X.; Sheriff, I.; Karaa, E.K.; Scheimberg, I.; Reddy, M.A.; Sagoo, M.S.; et al. Whole-Genome Sequencing of Retinoblastoma Reveals the Diversity of Rearrangements Disrupting RB1 and Uncovers a Treatment-Related Mutational Signature. *Cancers* **2021**, *13*, 754. [\[CrossRef\]](#)
26. Zheng, J.; Li, T.; Ye, H.; Jiang, Z.; Jiang, W.; Yang, H.; Wu, Z.; Xie, Z. Comprehensive identification of pathogenic variants in retinoblastoma by long- and short-read sequencing. *Cancer Lett.* **2024**, *598*, 217121. [\[CrossRef\]](#)
27. Leek, J.T.; Scharpf, R.B.; Bravo, H.C.; Simcha, D.; Langmead, B.; Johnson, W.E.; Geman, D.; Baggerly, K.; Irizarry, R.A. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev. Genet.* **2010**, *11*, 733–739. [\[CrossRef\]](#)
28. Carter, A.C.; Chang, H.Y.; Church, G.; Dombkowski, A.; Ecker, J.R.; Gil, E.; Giresi, P.G.; Greely, H.; Greenleaf, W.J.; Hacohen, N.; et al. Challenges and recommendations for epigenomics in precision health. *Nat. Biotechnol.* **2017**, *35*, 1128–1132. [\[CrossRef\]](#)
29. Robins, J.M.; Hernán, M.A.; Brumback, B. Marginal structural models and causal inference in epidemiology. *Epidemiology* **2000**, *11*, 550–560. [\[CrossRef\]](#)
30. Naimi, A.I.; Cole, S.R.; Kennedy, E.H. An introduction to g methods. *Int. J. Epidemiol.* **2017**, *46*, 756–762. [\[CrossRef\]](#)
31. Hernán, M.A.; Robins, J.M. Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available. *Am. J. Epidemiol.* **2016**, *183*, 758–764. [\[CrossRef\]](#)
32. Komorowski, M.; Celi, L.A.; Badawi, O.; Gordon, A.C.; Faisal, A.A. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat. Med.* **2018**, *24*, 1716–1720. [\[CrossRef\]](#)
33. Gottesman, O.; Johansson, F.; Komorowski, M.; Faisal, A.; Sontag, D.; Doshi-Velez, F.; Celi, L.A. Guidelines for reinforcement learning in healthcare. *Nat. Med.* **2019**, *25*, 16–18. [\[CrossRef\]](#)
34. Shi, C.; Fan, A.; Song, R.; Lu, W. High-Dimensional A-Learning for Optimal Dynamic Treatment Regimes. *Ann. Stat.* **2018**, *46*, 925–957. [\[CrossRef\]](#)
35. Sicklick, J.K.; Nishizaki, D.; Miyashita, H.; Okamura, R.; Hahn, M.E.; Nikanjam, M.; Fanta, P.T.; Piccioni, D.E.; Patel, H.; Eskander, R.N.; et al. Investigation of Profile-Related Evidence Determining Individualized Cancer Therapy (I-PREDICT) N-of-1 Precision Oncology Study: Molecular Profiling to Match Individually Dosed, Personalized Drug Combinations. *J. Clin. Oncol.* **2026**, *44*, 540–552. [\[CrossRef\]](#)
36. Chmielecki, J.; Gray, J.E.; Cheng, Y.; Ohe, Y.; Imamura, F.; Cho, B.C.; Lin, M.C.; Majem, M.; Shah, R.; Rukazenzov, Y.; et al. Candidate mechanisms of acquired resistance to first-line osimertinib in EGFR-mutated advanced non-small cell lung cancer. *Nat. Commun.* **2023**, *14*, 1070. [\[CrossRef\]](#)
37. Sirugo, G.; Williams, S.M.; Tishkoff, S.A. The Missing Diversity in Human Genetic Studies. *Cell* **2019**, *177*, 26–31. [\[CrossRef\]](#)
38. U.S. Food and Drug Administration. *Combination Product Definition Combination Product Types*; Office of Combination Products: Silver Spring, MD, USA, 2026.

39. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations. In *Draft Guidance for Industry and Food and Drug Administration Staff*; U.S. Food and Drug Administration: Silver Spring, MD, USA, 2025.
40. U.S. Food and Drug Administration. Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions. In *Guidance for Industry and Food and Drug Administration Staff*; U.S. Food and Drug Administration: Silver Spring, MD, USA, 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.