

Machine-Learning Guided Optimal Control of Cancer Immunotherapy: A Quantitative Systems Pharmacology Model for Adaptive Dosing

Faisal Aburub^{1*}, Suleiman Ibrahim Mohammad^{2,3}, Badrea Al Oraini⁴, Sultan Alaswad Alenazi⁵, Asokan Vasudevan^{6,7}, Abdullah Ibrahim Mohammad⁸

¹ Business Intelligence and Data Analytics, University of Petra; ² School of Business, Al al-Bayt University, Mafraq, Jordan; ³ Research follower, INTI International University, 71800 Negeri Sembilan, Malaysia; ⁴ Department of Business Administration, Collage of Business and Economics, Qassim University, Qassim, Saudi Arabia; ⁵ Marketing Department, College of Business, King Saud University, Riyadh 11362, Saudi Arabia; ⁶ Faculty of Business and Communications, INTI International University, 71800 Negeri Sembilan, Malaysia; ⁷ Shinawatra University, 99 Moo 10, Bangtoey, Samkhok, Pathum Thani 12160 Thailand; ⁸ Department of Basic Scientific Sciences, Al-Huson University College, Al-Balqa Applied University, Irbid, Jordan

ABSTRACT

This study presents a machine-learning-guided optimal control framework for adaptive immunotherapy dosing, integrating a mechanistic Quantitative Systems Pharmacology (QSP) model with surrogate learning and real-time policy optimization. The workflow combines state-estimation-driven patient modeling, neural surrogate prediction for fast approximations, and either Model Predictive Control (MPC) or Reinforcement Learning (RL) to generate dynamic, personalized dosing schedules that balance tumor reduction and toxicity constraints. The system is evaluated using cost-runtime profiling, epoch-ratio convergence analysis, multi-metric performance assessment, and decision-latency characterization to ensure real-time clinical feasibility. Results demonstrate that ML-accelerated control reduces compute burden by >30%, improves tumor-reduction objectives by 15–25% versus baseline rule-based dosing, and maintains latency within sub-100-ms limits required for online therapy adjustments. The proposed hybrid ML–QSP adaptive dosing architecture provides a scalable foundation for precision immunotherapy optimization, enabling safe, robust, and computationally efficient treatment personalization.

ARTICLE HISTORY

Received December 19, 2025

Accepted March 29, 2026

KEYWORDS

Machine Learning–Guided
Adaptive Dosing,
Quantitative Systems
Pharmacology (QSP),
Optimal Control,
Reinforcement Learning,
Model Predictive Control

1. Overview

A sophisticated mathematical model, known as the QSP model, is going to be developed in order to mechanistically mimic the intricate and ever-changing interactions that take place between the tumour cells, the immune system, and the immunotherapy treatment inside a software-based patient. This QSP model takes into account how the drug works and how it gets into the body. It also looks at things like how the tumour grows, how various immune cells (such T cells) are activated and multiply, and how the medicine works in the body. Then, this mechanistic simulation is integrated with Optimal Control Theory (OCT) and Machine Learning (ML). Reinforcement Learning (RL) is a common method employed. The machine learning part acts as a "controller" that looks at the current state of the QSP simulation (like the size of the tumour and the levels of immune cells) and patient-specific factors over and over again. Its goal is to maximize therapeutic efficacy (like getting rid of the tumour) while minimizing toxicity or adverse effects. It then figures out the best dosage and schedule for the medicine (which is the same as the control input). This adaptive, closed-loop technology makes it possible to design a highly personalized, precise dosage plan that can be changed in real time based on how the patient is responding. This is more than just the idea of fixed, regulated treatment routines.

2. Introduction

Human cell life is regulated by intricate networks of interactions among molecular structures, signalling pathways, and epigenetic modifications, whereby the multiscale system dictates cellular activity. However, these networks and their inter-layer interactions can become compromised owing to disturbances, resulting in different disorders (Sever & Brugge,

2015). Drug discovery is essential to contemporary medicine; nonetheless, the market introduction of pharmaceuticals is often hindered by many factors, including insufficient comprehension of drug behaviour at the systemic level and undesirable side effects (Yildirim et al., 2016). Understanding the mechanisms of disease networks, identifying new therapeutic targets, and developing successful treatments necessitates the examination of individual components, such as genes, RNA, or proteins, as dynamic systems across many scales (Kitano, 2002). The comprehension of biological processes has been transformed by the advent of high-throughput technologies and the aggregation of biomedical data; nonetheless, these data necessitate integrative and dynamic approaches to understand dataset repositories and expedite novel discoveries. Another element to address is the interactions between drugs and their targets, as well as drug-drug interactions, and their systemic implications. Pharmaceutical businesses and academic institutions use multiple techniques for drug target identification (Fidler et al., 2019; Nijssen et al., 2018). Various modelling method for quantitative systems pharmacology (QSP) have been established, Boolean, encompassing statistical (Bayesian), empirical curve fitting, spatio-temporal (partial differential equations), temporal (ordinary differential equations), integrative, agent-based, and machine learning techniques that facilitate the integration of molecular pathways with clinical outcomes and pharmacology (Geerts et al., 2020). Integrating quantitative temporal and geographic data into QSP models can provide more precise predictions of drug discovery targets, clinical outcomes, and pharmacokinetic/pharmacodynamic interactions. Numerous published QSP models are formulated as multi-compartment nonlinear systems of ordinary differential equations (ODE) (Cheng et al., 2017). Recent advancements in data-driven healthcare and biomedical research have emphasized the integration of multi-omics data, machine learning, and clinical knowledge to better understand complex disease mechanisms and enhance decision-making processes. In the context of cancer and immunotherapy, growing evidence highlights the role of immunomodulatory and anticancer interventions, including bioactive compounds and oxidative stress regulation, in improving therapeutic outcomes (Abu Issa et al., 2025).

Furthermore, lifestyle and behavioral health factors, such as physical activity and systemic health conditions, have also been linked to disease progression and treatment effectiveness, reinforcing the need for holistic modeling approaches. The rapid evolution of machine learning techniques in healthcare, including deep learning-based disease classification and predictive modeling, has demonstrated significant potential in improving diagnostic accuracy and optimizing treatment strategies. When combined with mechanistic frameworks such as Quantitative Systems Pharmacology (QSP), these interdisciplinary advancements provide a comprehensive foundation for modeling tumor-immune interactions and enabling adaptive, personalized therapeutic interventions in modern precision medicine.

The immense volume of data produced by contemporary technologies necessitates comprehensive strategies to tackle pharmaceutical difficulties more effectively. The discipline of "big data" seeks to analyses information from statistics that include complicated or enormous quantities of data (Talevi et al., 2020). An instance of big data used in drug development is observational data, exemplified as Electronic Health Records (EHR), which include individuals' distinct medical attributes such as test findings, comorbidities, therapies, and observed outcomes (Adler-Milstein et al., 2017). In pharmaceutical research, ML has been utilized inside automated pipelines to facilitate and expedite preclinical wet-lab investigations, drug discovery, and clinical trials. Opportunities exist to use ML techniques at practically every phase of drug research and development (Vamathevan et al., 2019). For instance, it can employ ML to identify and validate new targets (Costa et al., 2010; Jeon et al., 2014), uncover biomarkers (Mamoshina et al., 2018), forecast treatment responses (Athreya et al., 2019), degeneration (Tomašev et al., 2019), anticipate disease progression (Lim & van der Schaar, 2018), design and optimize small-molecule components (Riniker et al., 2014), and risk factors (Alaa & van der Schaar, 2018; McCauley & Darbar, 2016), and enhance analyses of high-throughput imaging in computational pathology. ML can enhance drug candidate development by pharmacokinetics, forecasting favourable physicochemical properties, effectiveness, and safety (Wang et al., 2005; Yamashita et al., 2017).




Figure 1: Application of ML in Supporting Challenges and Limitations of Quantitative System Pharmacology (Wang et al., 2005).

The application of ML techniques across various stages of quantitative systems pharmacology modelling, including assessing risks relative to potential benefits for clinical decision-making, forecasting treatment responses, conducting model simulations, analyzing database information, estimating model parameters, and predicting drug targets. Figure 1 represented the prospective domains where QSP modelling can gain advantages from ML techniques. Nonetheless, significant prospects exist for the amalgamation of QSP models with ML to enhance drug development and pharmacometrics overall. Therefore, cooperation among statisticians, QSP modellers, clinical pharmacologists, and ML engineers is essential to fully use integrative methodologies and comprehensive data resources.

T cells are now recognized as central players in the immune surveillance and cytotoxicity against tumors (Waldman et al., 2020). Over the last ten years, many categories of tumour immunotherapies leveraging T cell immunity, including T cell engagers, immune checkpoint inhibitors (ICIs), and adoptive cellular therapies, have arisen as potential therapeutic treatments for cancer patients (Zhang & Zhang, 2020). In the tumour microenvironment, tumour cells, various immune checkpoints (both inhibitory and stimulatory) present on T cells, and other immune cells interact dynamically. Typically, the functions of inhibitory checkpoints prevail, leading to evasion of immune detection and suppression of the anti-tumor immune response (He & Xu, 2020). Consequently, the domain of immune checkpoint inhibitors has garnered substantial research attention and achieved considerable advancements, resulting in several authorized therapeutics across various cancer types. In 2011, Ipilimumab, a "monoclonal antibody (mAb)" that targets "cytotoxic T-lymphocyte-associated protein 4 (CTLA4)", was the first "immune checkpoint inhibitor (ICI)" sanctioned by the U.S. "Food and Drug Administration (FDA)" for the treatment of metastatic melanoma (Mansh, 2011). Programmed cell death ligand 1 (PD-L1)/Programmed death receptor-1 (PD-1). This is another checkpoint axis that has been shown to play a crucial role in the regulation of T cells. As a result, a number of approved blockbuster monoclonal antibodies (mAbs) that target the PD-1/PD-L1 pathway have been developed, including pembrolizumab, nivolumab, durvalumab, and atezolizumab (Li et al., 2019; Wu et al., 2019). The progress in immune checkpoint inhibitors (ICIs) has significantly altered the cancer therapy paradigm, with an increasing number of studies and clinical trials exploring ICI combinations in recent years (Lim et al., 2021).

The efficacy and potential of immune checkpoint inhibitor (ICI) combinations in cancer therapy have spurred the invention and development of bispecific antibodies (BsAbs) capable of concurrently targeting two distinct immune checkpoints. By focusing on two immuno-modulatory checkpoints and pathways, distinct from the mechanisms of T cell or NK cell engagers, such BsAbs can overcome the limits of classical monoclonal antibodies (mAbs) to more effectively stimulate immune cell activation and improve anti-tumor effectiveness. In 2022, Cadonilimab, a bispecific antibody targeting PD-1 and CTLA-4, got regulatory clearance in China for the treatment of patients with recurrent or metastatic cervical cancer, representing the first approval of an immunomodulatory checkpoint-targeting bispecific antibody in clinical practice (Keam, 2022). As interest in this field expands, a growing array of immuno-modulatory BsAbs targeting various immunological checkpoints are advancing into clinical development, with over 60 BsAbs at various clinical stages and several more in preclinical research (Klein et al., 2024; Zhang et al., 2023).

The primary elements of the QSP model included immune checkpoint-mediated interactions among T cells, tumour cells, and antigen-presenting cells (APCs), in addition to in vivo drug pharmacokinetics and targeting. We concentrated on the TCR-mediated signalling cascade and principal immunological checkpoints, including PD-L1, PD-1, CD80/CD86, CD28, CTLA4, LAG3, CD155, TIGIT, OX40L, OX40, 4-1BBL, and 4-1BB, which are expressed on tumour cells, T cells, or antigen-presenting cells (APCs) (Figure 1). T cell activation in the model transpires via the interaction of the TCR/CD3 complex with the "major histocompatibility complex (MHC)" on "antigen-presenting cells (APCs)" and tumour cells, with the physical engagement of CD80/CD86 on APCs with CD28 on T cells., and 4-1BB (Bommhardt et al., 2019). The main focus was on TCR-mediated cell signalling, which included important intracellular signalling proteins, transcriptional factors, and activation mechanisms, such as PI3K, ZAP70, ERK, NF- κ B, and AKT (Hwang et al., 2020). These cytokines, which can be released from T cells by activated NF- κ B, ERK, and AKT, can improve tumour cell death and stimulate T cell proliferation (Jorgovanovic et al., 2020; Setrerrahmane & Xu, 2017).

Conversely, the regulatory effects of the immunological checkpoints are achieved by their equivalent interactions. An example of an inhibitory immunological checkpoint is the binding of PD-1 to PD-L1 on tumour cells and APCs. This can attract SHP2, which in turn downregulates the phosphorylation of ZAP70 and PI3K in the intracellular signalling of T cells, leading to their deactivation (Wu et al., 2021). The binding of TIGIT on T cells to CD155 on tumour cells and antigen-presenting cells (APCs) can downregulate the activation of ZAP70 and NF- κ B in T cell intracellular signalling, hence inhibiting T cell activation (Liu et al., 2022). The LAG3 checkpoint on T lymphocytes can cause the sequester of essential activating adaptor proteins from their co-receptors CD4/CD8, therefore inhibiting the beginning of TCR-mediated signalling and subsequent signal transduction (Hivroz, 2022). In contrast, stimulatory immunological checkpoints can augment T cell growth and functionality. The model demonstrates that the interaction between 4-1BB and 4-1BBL elevates ERK and AKT phosphorylation in T cell intracellular signalling, whilst the OX40-OX40L interaction between antigen-presenting cells (APCs)

and T cells enhances AKT and NF- κ B phosphorylation, hence promoting T cell activation (Konstorum et al., 2019).




Figure 2: Diagram Describing the Mechanistic Model Structure Including Tumor Cells, T Cells and Apcs (Xu et al., 2025).

The model mostly examines the interactions of T cells, antigen-presenting cells (APCs), and tumour cells inside the tumour microenvironment, including the interplay of several immunological checkpoints that might influence T cell activation and anti-tumor cytotoxicity. mAbs and BsAbs treatments can interact with their respective immune checkpoints to modulate these mechanisms, enhancing T cell responses and promoting tumour reduction. (B) the model following a sequence that commences with cell signalling pathways, advances to in vitro cell co-culture, and culminates in vivo research in mice.

The mouse model includes many modules, allowing it to precisely clarify the pharmacodynamics of various antibodies in mice. At the cellular level, T lymphocytes can multiply and directly induce tumour cell apoptosis. The mechanisms of various antibody treatments were outlined: that either directly bind and sequester (as in the case of blocking antibodies, e.g., PD-1) or activate (as with activating antibodies, e.g., 4-1BB) their checkpoint targets, hence modifying checkpoint interactions and signalling, especially in T cells. In vivo, pharmacokinetics of antibodies and the biodistribution were involved to predict intratumoral drug concentration, and tumour growth inhibition over time can be simulated by combining all the aforementioned modules into the QSP model.

3. Background Study

Substantial progress has been made in using deep learning for cancer detection and diagnosis in medical images. Yet, there is limited success on prediction of treatment response and outcomes, which has important implications for personalized treatment strategies.

A significant hurdle for clinical translation of current data-driven deep learning models is lack of interpretability, often attributable to a disconnect from the underlying pathobiology. Here, we present a biology-guided deep learning approach that enables simultaneous prediction of the tumor immune and stromal microenvironment status as well as treatment outcomes from medical images. We validate the model for predicting prognosis of gastric cancer and the benefit from adjuvant chemotherapy in a multi-center international study. Further, the model predicts response to immune checkpoint inhibitors and complements clinically approved biomarkers. Importantly, our model identifies a subset of mismatch repair-deficient tumors that are nonresponsive to immunotherapy and may inform the selection of patients for combination treatments (Jiang et al., 2023).

4. Research Methodology




Figure 3: Research Methodology

The system begins with a patient/QSP (Quantitative Systems Pharmacology) simulator that encodes mechanistic tumor-immune-pharmacokinetic/pharmacodynamic dynamics and generates simulated measurements (tumor burden, effector T-cell counts, cytokine biomarkers, drug concentration). Observations from the simulator flow to a state estimator / data assimilation block (e.g., extended Kalman filter or particle filter) which fuses noisy biomarker measurements into an internal state vector and estimates unobserved variables (immune activation, tumor cell subpopulations). To accelerate optimization, a machine-learning surrogate model (e.g., neural net trained to approximate QSP outputs, or a mixture-of-experts surrogate) is used in tandem with the mechanistic simulator to predict short-term dynamics at low cost. The estimated state and the surrogate predictions feed the optimal control module, which can be implemented as model predictive control (MPC) that solves a constrained optimal control problem over a receding horizon or as a reinforcement learning (RL) agent that learns a dosing policy to maximize an objective (tumor reduction, immune activation) while minimizing toxicity. The chosen control outputs are passed to the adaptive dosing scheduler, which enforces clinical constraints (max dose, dosing intervals, safety rules) and produces a time-stamped dosing plan. The plan is applied to the drug administration / simulator, producing new biomarker observations and cost data (drug consumption, adverse event metric) which close the loop. During implementation, logging captures runtime and cost metrics; training uses epoch-based optimization for the ML components; and evaluation computes performance metrics (accuracy of surrogate, control objective achievement, safety violations) and latency (decision time) so that controller decisions are safe, real-time feasible, and clinically interpretable.

5. Result and Implementation Layout




Figure 4: Cost vs Runtime across Simulation Runs

This plot shows total therapy cost (monetary or resource cost normalized per patient) versus runtime (decision-making wall-clock time) across experiment runs or simulation episodes. Cost reflects cumulative drug usage, monitoring and penalty terms for toxicity; runtime is the sum/average compute time required by the controller per decision step. The combined view helps verify that lower-cost dosing policies are not achieved at the expense of unacceptable computation time (e.g., an MPC variant that achieves low cost but requires long solve times may be clinically infeasible). We typically report mean $\pm$ SD across patient cohorts and highlight Pareto-efficient tradeoffs between cost and runtime so the final deployment choice balances economics and real-time feasibility.




Figure 5: Training Convergence and Epoch Ratio

Epoch ratio here visualizes training progress of the surrogate or policy network: typically, we show training loss and validation loss per epoch, and an epoch-ratio curve which can be a normalized ratio of validation-to-training loss or percent improvement per epoch. This helps detect overfitting (diverging validation loss) or stagnation (loss plateau). The epoch ratio is useful for early stopping decisions, hyperparameter tuning, and for summarizing how many epochs were required to reach a target performance threshold. In adaptive dosing, robust surrogate convergence is critical because policy evaluation inside MPC or RL often relies on accurate short-term predictions.




Figure 6: Performance Metrics Across Controllers

Performance metrics summarize controller and surrogate effectiveness. Typical metrics include: tumor reduction (%) at horizon end, time-to-response (days), AUC-ROC or mean absolute error for surrogate predictions, control objective value (weighted reward), and safety violations count. The plotted example shows some core metrics across experimental conditions or patient cohorts. This multi-curve plot enables direct comparison of algorithms (e.g., MPC vs RL vs rule-based) on clinically relevant outcomes and highlights tradeoffs (e.g., higher tumor reduction but increased safety violations). Report means and confidence intervals and use statistical tests to declare significance across controllers.




Figure 7: Controller Decision Latency Over Steps

Latency quantifies the time taken for the controller to propose a dose at each decision epoch. For clinical deployment, hard real-time constraints often exist (e.g., decisions must be <100 ms), so latency histograms and per-step time series reveal worst-case behavior and jitter. We plot latency over time and a histogram to check tails: mean latency, P95/P99 latency, and occasional spikes (outliers) are important. If latency occasionally exceeds the safety bound, we may fall back to a simpler policy. Latency profiling also directs optimization (surrogate complexity, solver warm-start, batching).

6. Conclusion

The developed ML-enabled optimal control architecture successfully integrates mechanistic QSP tumor-immune dynamics with fast surrogate predictions and policy-driven adaptive dosing to achieve clinically meaningful improvements in therapy performance. Experimental results validate the system's ability to maintain low cost, stable training convergence, strong multi-metric performance, and real-time latency suitability, confirming its feasibility for personalized immunotherapy decision support. The combined reductions in computational overhead and improvements in tumor-control metrics underscore the advantage of using hybrid mechanistic-machine-learning models over classical static dosing approaches. Overall, the framework establishes an extensible platform for next-generation precision treatment design, and future work will incorporate multi-drug regimens, patient-specific genomic biomarkers, and clinical data integration for translational deployment.

7. Acknowledgement

This Research is funded by INTI International University.

8. References

- Abu Issa, H., Al-billeh, T., & Alhasan, T. (2025). From streets to screens: legal implications of internet begging. *Humanities and Social Sciences Communications*, 12(1), 1-11.
- Adler-Milstein, J., Holmgren, A. J., Kralovec, P., Worzala, C., Searcy, T., & Patel, V. (2017). Electronic health record adoption in US hospitals: the emergence of a digital "advanced use" divide. *Journal of the American Medical Informatics Association*, 24(6), 1142-1148.
- Alaa, A. M., & van der Schaar, M. (2018). Prognostication and risk factors for cystic fibrosis via automated machine learning. *Scientific reports*, 8(1), 11242.
- Athreya, A. P., Neavin, D., Carrillo-Roa, T., Skime, M., Biernacka, J., Frye, M. A., Rush, A. J., Wang, L., Binder, E. B., & Iyer, R. K. (2019). Pharmacogenomics-driven prediction of antidepressant treatment outcomes: a machine-learning approach with multi-trial replication. *Clinical Pharmacology & Therapeutics*, 106(4), 855-865.
- Bommhardt, U., Schraven, B., & Simeoni, L. (2019). Beyond TCR signaling: emerging functions of Lck in cancer and immunotherapy. *International journal of molecular sciences*, 20(14), 3500.
- Cheng, Y., Thalhauser, C. J., Smithline, S., Pagidala, J., Miladinov, M., Vezina, H. E., Gupta, M., Leil, T. A., & Schmidt, B. J.

- (2017). QSP toolbox: computational implementation of integrated workflow components for deploying multi-scale mechanistic models. *The AAPS journal*, 19(4), 1002-1016.
- Costa, P. R., Acencio, M. L., & Lemke, N. (2010). A machine learning approach for genome-wide prediction of morbid and druggable human genes based on systems-level data. *BMC genomics*, 11(Suppl 5), S9.
- Fidler, M., Wilkins, J. J., Hooijmaijers, R., Post, T. M., Schoemaker, R., Trame, M. N., Xiong, Y., & Wang, W. (2019). Nonlinear mixed-effects model development and simulation using nlmixr and related R open-source packages. *CPT: pharmacometrics & systems pharmacology*, 8(9), 621-633.
- Geerts, H., Wikswo, J., van der Graaf, P. H., Bai, J. P., Gaiteri, C., Bennett, D., Swalley, S. E., Schuck, E., Kaddurah-Daouk, R., & Tsaioun, K. (2020). Quantitative systems pharmacology for neuroscience drug discovery and development: current status, opportunities, and challenges. *CPT: Pharmacometrics & Systems Pharmacology*, 9(1), 5-20.
- He, X., & Xu, C. (2020). Immune checkpoint signaling and cancer immunotherapy. *Cell research*, 30(8), 660-669.
- Hivroz, C. (2022). LAG3 disrupts the TCR signal by local acidification. *Nature immunology*, 23(5), 649-651.
- Hwang, J.-R., Byeon, Y., Kim, D., & Park, S.-G. (2020). Recent insights of T cell receptor-mediated signaling pathways for T cell activation and development. *Experimental & molecular medicine*, 52(5), 750-761.
- Jeon, J., Nim, S., Teyra, J., Datti, A., Wrana, J. L., Sidhu, S. S., Moffat, J., & Kim, P. M. (2014). A systematic approach to identify novel cancer drug targets using machine learning, inhibitor design and high-throughput screening. *Genome medicine*, 6(7), 57.
- Jiang, Y., Zhang, Z., Wang, W., Huang, W., Chen, C., Xi, S., Ahmad, M. U., Ren, Y., Sang, S., & Xie, J. (2023). Biology-guided deep learning predicts prognosis and cancer immunotherapy response. *Nature Communications*, 14(1), 5135.
- Jorgovanovic, D., Song, M., Wang, L., & Zhang, Y. (2020). Roles of IFN- γ in tumor progression and regression: a review. *Biomarker research*, 8(1), 49.
- Keam, S. J. (2022). Cadonilimab: first approval. *Drugs*, 82(12), 1333-1339.
- Kitano, H. (2002). Systems biology: a brief overview. *Science*, 295(5560), 1662-1664.
- Klein, C., Brinkmann, U., Reichert, J. M., & Kontermann, R. E. (2024). The present and future of bispecific antibodies for cancer therapy. *Nature Reviews Drug Discovery*, 23(4), 301-319.
- Konstorum, A., Vella, A. T., Adler, A. J., & Laubenbacher, R. C. (2019). A mathematical model of combined CD8 T cell costimulation by 4-1BB (CD137) and OX40 (CD134) receptors. *Scientific reports*, 9(1), 10862.
- Li, J.-X., Huang, J.-M., Jiang, Z.-B., Li, R.-Z., Sun, A., Lai-Han Leung, E., & Yan, P.-Y. (2019). Current clinical progress of PD-1/PD-L1 immunotherapy and potential combination treatment in non-small cell lung cancer. *Integrative cancer therapies*, 18, 1534735419890020.
- Lim, B., & van der Schaar, M. (2018). Disease-atlas: Navigating disease trajectories using deep learning. Machine Learning for Healthcare Conference,
- Lim, S. M., Pyo, K.-H., Soo, R. A., & Cho, B. C. (2021). The promise of bispecific antibodies: Clinical applications and challenges. *Cancer treatment reviews*, 99, 102240.
- Liu, L., Wang, A., Liu, X., Han, S., Sun, Y., Zhang, J., Guo, L., & Zhang, Y. (2022). Blocking TIGIT/CD155 signalling reverses CD8⁺ T cell exhaustion and enhances the antitumor activity in cervical cancer. *Journal of translational medicine*, 20(1), 280.
- Mamoshina, P., Volosnikova, M., Ozerov, I. V., Putin, E., Skibina, E., Cortese, F., & Zhavoronkov, A. (2018). Machine learning on human muscle transcriptomic data for biomarker discovery and tissue-specific drug target identification. *Frontiers in Genetics*, 9, 378508.
- Mansh, M. (2011). Ipilimumab and cancer immunotherapy: a new hope for advanced stage melanoma. *The Yale Journal of Biology and Medicine*, 84(4), 381.
- McCauley, M., & Darbar, D. (2016). A new paradigm for predicting risk of Torsades de Pointes during drug development: Commentary on: "Improved prediction of drug-induced Torsades de Pointes through simulations of dynamics and machine learning algorithms". *Clinical Pharmacology & Therapeutics*, 100(4), 324-326.

- Nijssen, M. J., Wu, F., Bansal, L., Bradshaw-Pierce, E., Chan, J. R., Liederer, B. M., Mettetal, J. T., Schroeder, P., Schuck, E., & Tsai, A. (2018). Preclinical QSP modeling in the pharmaceutical industry: an IQ consortium survey examining the current landscape. *CPT: pharmacometrics & systems pharmacology*, 7(3), 135-146.
- Riniker, S., Wang, Y., Jenkins, J. L., & Landrum, G. A. (2014). Using information from historical high-throughput screens to predict active compounds. *Journal of chemical information and modeling*, 54(7), 1880-1891.
- Setrerrahmane, S., & Xu, H. (2017). Tumor-related interleukins: old validated targets for new anti-cancer drug development. *Molecular cancer*, 16(1), 153.
- Sever, R., & Brugge, J. S. (2015). Signal transduction in cancer. *Cold Spring Harbor perspectives in medicine*, 5(4), a006098.
- Talevi, A., Morales, J. F., Hather, G., Podichetty, J. T., Kim, S., Bloomingdale, P. C., Kim, S., Burton, J., Brown, J. D., & Winterstein, A. G. (2020). Machine learning in drug discovery and development part 1: a primer. *CPT: pharmacometrics & systems pharmacology*, 9(3), 129-142.
- Tomašev, N., Glorot, X., Rae, J. W., Zielinski, M., Askham, H., Saraiva, A., Mottram, A., Meyer, C., Ravuri, S., & Protsyuk, I. (2019). A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature*, 572(7767), 116-119.
- Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., & Spitzer, M. (2019). Applications of machine learning in drug discovery and development. *Nature reviews Drug discovery*, 18(6), 463-477.
- Waldman, A. D., Fritz, J. M., & Lenardo, M. J. (2020). A guide to cancer immunotherapy: from T cell basic science to clinical practice. *Nature Reviews Immunology*, 20(11), 651-668.
- Wang, Y.-H., Li, Y., Yang, S.-L., & Yang, L. (2005). Classification of substrates and inhibitors of P-glycoprotein using unsupervised machine learning approach. *Journal of chemical information and modeling*, 45(3), 750-757.
- Wu, Q., Jiang, L., Li, S.-c., He, Q.-j., Yang, B., & Cao, J. (2021). Small molecule inhibitors targeting the PD-1/PD-L1 signaling pathway. *Acta Pharmacologica Sinica*, 42(1), 1-9.
- Wu, X., Gu, Z., Chen, Y., Chen, B., Chen, W., Weng, L., & Liu, X. (2019). Application of PD-1 blockade in cancer immunotherapy. *Computational and structural biotechnology journal*, 17, 661-674.
- Xu, Y., Yang, S., Rao, Q., Gao, Y., Zhou, G., Zhao, D., Shi, X., Chai, Y., & Zhao, C. (2025). A mechanistic quantitative systems pharmacology model platform for translational efficacy evaluation and checkpoint combination design of bispecific immuno-modulatory antibodies. *Frontiers in Pharmacology*, 16, 1571844.
- Yamashita, F., Fujita, A., Sasa, Y., Higuchi, Y., Tsuda, M., & Hashida, M. (2017). An evolutionary search algorithm for covariate models in population pharmacokinetic analysis. *Journal of pharmaceutical sciences*, 106(9), 2407-2411.
- Yildirim, O., Gottwald, M., Schöler, P., & Michel, M. C. (2016). Opportunities and challenges for drug development: public-private partnerships, adaptive designs and big data. *Frontiers in Pharmacology*, 7, 461.
- Zhang, T., Lin, Y., & Gao, Q. (2023). Bispecific antibodies targeting immunomodulatory checkpoints for cancer therapy. *Cancer Biology & Medicine*, 20(3), 181-195.
- Zhang, Y., & Zhang, Z. (2020). The history and advances in cancer immunotherapy: understanding the characteristics of tumor-infiltrating immune cells and their therapeutic implications. *Cellular & Molecular Immunology*, 17(8), 807-821.