

ORIGINAL RESEARCH

Open access

A Causal Transformer Framework for Counterfactual Estimation of Antihypertensive Dose Responses from Observational Electronic Health Records

Jose Martinez^{1*}, Carmen Lopez¹

Abstract

Hypertension affects over 1.4 billion adults worldwide, and antihypertensive dose titration is a common but complex clinical decision. Although electronic health records contain longitudinal data on medication adjustments and blood pressure outcomes, determining optimal individualized dosing remains challenging due to confounding in observational data, where patients receiving higher doses often have worse baseline health. We propose a transformer-based model with causal attention masking to estimate counterfactual blood pressure outcomes under alternative dose regimens. The architecture ensures temporal validity by preventing information leakage from future events and encodes medication dose changes in a continuous representation. It includes a dose encoder, outcome predictor, and counterfactual contrastive loss to distinguish between competing treatment paths. This framework learns patient-specific dose–response relationships and enables personalized predictions for antihypertensive adjustments. While it supports individualized treatment planning from observational EHR data, prospective validation is still required before clinical deployment.

Keywords Electronic health records, Causal transformer, Counterfactual estimation, Antihypertensive dosing, Attention masking, Treatment effect heterogeneity

*Correspondence:

Jose Martinez
jose.martinez@gmail.com

¹ Department of AI in Healthcare Analytics, University of Salamanca, Salamanca, Spain

Introduction

Hypertension affects approximately 1.4 billion adults globally and is a primary modifiable risk factor for cardiovascular disease, stroke, and renal failure [1, 2]. Clinical guidelines recommend target blood pressure below 130/80 mmHg, but achieving this target frequently requires dose titration of first-line agents (ACE inhibitors, ARBs, calcium channel blockers, thiazides) or the addition of multiple drug classes [3, 4]. For an individual patient on lisinopril 10mg, the clinician faces a counterfactual question: Would increasing to 20mg lower blood pressure more effectively than adding amlodipine 5mg? Observational EHR data contain historical records of such dose changes and subsequent blood pressure

measurements, yet extracting causal answers remains challenging because treatment decisions depend on patient characteristics [5, 6].

Standard machine learning approaches—including recurrent neural networks and conventional transformers—are confounded by indication when applied to treatment effect estimation [7, 8]. Patients who receive dose escalations are systematically sicker than those maintained on stable doses: they have higher baseline blood pressure, more comorbidities such as diabetes or chronic kidney disease, and poorer responses to initial therapy [9, 10]. A naive transformer predicting blood pressure from observed dose histories would learn that higher doses are associated with worse outcomes, precisely because higher doses are

given to sicker patients. This bias cannot be resolved by simply including all measured covariates, as the same clinical logic generates the confounding structure [11, 12].

Causal inference methods for time-series data, including marginal structural models, G-computation, and instrumental variable approaches, have been developed to address confounding by indication [13, 14]. However, these methods rely on strong assumptions such as sequential ignorability (no unmeasured confounders) and positivity (non-zero probability of any treatment at any time), and they often require manual specification of propensity score models that become impractical when treatment options include multiple drugs and continuous doses [15, 16]. Furthermore, traditional causal methods do not naturally leverage the representational learning capabilities of modern sequence models, limiting their ability to capture complex temporal patterns in blood pressure trajectories and medication responses [17].

This paper proposes a causal transformer architecture with attention masking that directly encodes the temporal ordering required for counterfactual estimation of antihypertensive dose responses [5, 18]. The framework extends the standard transformer by introducing a causal mask that prevents information leakage from future time steps when predicting potential outcomes under alternative dose regimens [19]. Unlike conventional causal inference pipelines that separate confounder adjustment from outcome modeling, our approach learns a causally-valid representation of patient history end-to-end.

Background

Hypertension pharmacology

First-line antihypertensive drug classes include ACE inhibitors (e.g., lisinopril, ramipril), angiotensin receptor blockers (e.g., losartan, valsartan), calcium channel blockers (e.g., amlodipine, nifedipine), thiazide diuretics (e.g., hydrochlorothiazide, chlorthalidone), and beta-blockers (e.g., metoprolol, atenolol) [1, 2]. Dose-response curves vary across both drugs and patient subgroups: for ACE inhibitors, blood pressure reduction follows a log-linear relationship with dose from 5mg to 40mg, but the incremental benefit diminishes above 20mg in many patients. Current hypertension guidelines from the European Society of Cardiology and Hypertension Canada recommend target blood pressure below 130/80 mmHg for most adults, with treatment intensification when readings

exceed this threshold on three consecutive measurements [20, 21].

Treatment escalation patterns

Clinical practice proceeds through well-documented escalation patterns: monotherapy at low dose → titration to maximum tolerated dose → addition of a second agent from a complementary class → dual therapy titration → triple therapy with fixed-dose combination pills [6, 7]. The decision to increase an existing dose versus adding a new agent depends on multiple factors including the patient's current blood pressure, side effect profile, adherence, age, renal function, and prior medication failures [19, 22]. Fixed-dose combination therapies combining perindopril, indapamide, and amlodipine in a single pill have demonstrated improved adherence and blood pressure control compared to free combinations, suggesting that regimen complexity directly influences outcomes [23-26].

Confounding by indication

Confounding by indication arises when the treatment assignment (dose increase) is causally influenced by the same factors that predict the outcome (future blood pressure), creating non-causal associations [12, 13]. In hypertension management, clinicians escalate doses precisely when patients exhibit uncontrolled blood pressure despite current therapy, meaning that the indication for dose increase (high measured BP) is also a strong predictor of future high BP. Standard regression adjustment fails when confounders are measured with error, when interactions are misspecified, or when unmeasured factors such as dietary sodium intake or medication adherence drive both treatment decisions and outcomes [14, 27]. The problem is exacerbated in time-series settings where time-varying confounders (e.g., interim blood pressure readings) are themselves affected by prior treatments, creating feedback loops that bias naive estimators [15, 16].

Causal inference from time-series data

Marginal structural models with inverse probability of treatment weighting address time-varying confounding by estimating stabilized weights that adjust for both baseline and time-dependent confounders [14, 17]. G-computation provides an alternative by directly modeling the outcome distribution under hypothetical treatment regimes, but both methods require correct specification of the propensity score and outcome models [18, 28]. Instrumental variable methods can handle unmeasured confounding when a valid

instrument exists, but such instruments (e.g., clinician prescribing preferences) are difficult to identify in hypertension dose titration [29]. Recurrent marginal structural networks and adversarially balanced representations have been proposed to learn balanced representations of treatment history, yet these methods still rely on the sequential ignorability assumption that all confounders are measured [6-8].

Framework Overview

High-level architecture

The causal transformer framework takes as input a sequence of patient observations across time steps (typically weekly intervals): systolic blood pressure, diastolic blood pressure, current medication type, current dose in milligrams, co-medications, age, weight, and serum creatinine [17, 18]. The transformer encoder processes this sequence with a causal attention mask that prevents information from future time steps from influencing predictions at the current time. The output layer produces counterfactual predictions for blood pressure measurements at specified future horizons under alternative dose regimens that may differ from the dose actually administered [5, 19].

Figure 1 presents the proposed causal transformer architecture as a directional framework linking longitudinal EHR history, temporal attention restriction, counterfactual dose-regimen prediction, and individualized antihypertensive decision support.

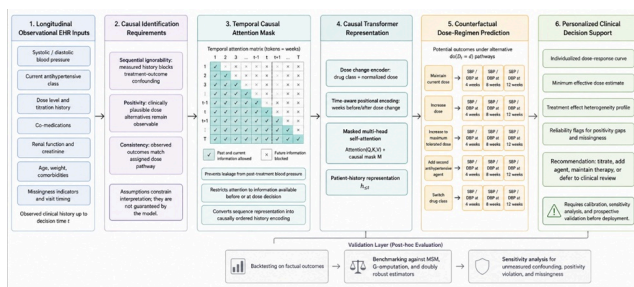


Figure 1. Causal Transformer Framework for Counterfactual Antihypertensive Dose-Response Estimation from Observational EHR Data

Core assumptions

The framework operates under three standard causal assumptions adapted to the time-series setting. Sequential

ignorability requires that, conditional on the observed history up to time t (including all past treatments, outcomes, and covariates), the treatment assigned at time t is independent of the potential outcomes under any treatment sequence [6, 14]. Positivity requires that for any patient history observed in the data, there is a non-zero probability of receiving each possible dose level at that time. Consistency requires that the observed outcome for a patient who received a particular dose equals the potential outcome under that dose [7, 15, 16].

Design principles

The framework adheres to four design principles motivated by clinical hypertension management. Causal validity requires that all predictions respect temporal ordering and do not use future information when estimating counterfactual outcomes under past dose decisions [8, 9]. Patient-specificity requires that the model produce individualized dose-response curves rather than population-average effects, enabling personalized treatment recommendations [10, 11]. Time-awareness requires that the architecture explicitly model the timing of dose changes, including the duration since the last dose adjustment. Uncertainty-awareness, though not fully addressed here, would require probabilistic outputs that reflect the increased uncertainty for counterfactual regimens far from the observed treatment path [1, 2].

Table 1 clarifies how each architectural element contributes to causal validity, counterfactual estimation, and clinical interpretability beyond conventional sequence prediction.

Table 1. Causal Design Logic of the Proposed Transformer Framework

Framework element	Causal problem addressed	Operational role in the model	Analytical contribution beyond standard prediction
Longitudinal EHR history	Treatment assignment depends on prior BP, comorbidity, renal function, and medication response	Encodes observed patient history before each dose decision	Shifts prediction from cross-sectional associations to temporal condition estimation

Causal attention mask	Future blood pressure or covariates may leak into treatment-effect prediction	Blocks attention from post-decision time steps	Enforce "no future information condition" required for counterfactual estimation
Dose change encoder	Medication class and dose intensity are heterogeneous and not purely categorical	Represents drug type, normalized dose, and dose-transition structure	Allows comparison across time switches and add-on treatment decisions
Time-aware positional encoding	Dose response depends on timing since treatment change	Encodes pre-dose and post-dose intervals relative to clinical decision points	Captures pharmacological response latency and clinical plateau
Counterfactual prediction head	Observed outcomes exist only for the treatment actually received	Generates potential BP outcomes under alternative dose paths	Enables patient-specific counterfactual estimation using unobserved dose-response trajectories
Contrastive counterfactual loss	Factual prediction alone may ignore dose variation	Separates representations for clinically distinct dose alternatives	Encourages treatment-sensitive representations rather than prognostic representations
Positivity diagnostics	Some dose alternatives may be rare or absent for specific patient profiles	Flags unreliable counterfactual predictions in unsupported regions	Prevents overconfident extrapolation beyond observed clinical practice
Sensitivity analysis	Sequential ignorability cannot be proven from observational EHR data	Tests robustness to unmeasured adherence, diet, socioeconomic status, and missingness	Makes uncertainty explicit rather than hidden inside performance metrics

Causal Attention Masking

Standard vs causal attention

The standard transformer attention mechanism computes attention weights between all pairs of positions in a sequence, allowing information to flow from future tokens to current predictions [5]. This design is appropriate for language modeling where the full context is known at inference but violates causal validity for counterfactual estimation: predicting the effect of a dose change at time t should not incorporate blood pressure readings from time $t+1$ that occur after the dose was administered. The causal attention mask modifies the attention matrix by setting weights to negative infinity for all pairs where the key position (source) is greater than the query position (target), ensuring that predictions for time t depend only on information from times $\leq t$ [22, 23].

Masking structure

The masking structure is defined as follows: for a sequence of length T , the attention weight from position i (query) to position j (key) is masked when $j > i$. When predicting the potential outcome at time $t+1$ under a dose change that occurred at time t , the transformer can attend to all pre-treatment information including baseline covariates, prior blood pressure readings, medication history, and the dose change decision itself—but cannot attend to any post-treatment outcomes or future covariate values [24, 25]. This masking structure operationalizes the "no future information" condition required for counterfactual identification in time-series settings, directly encoding the temporal ordering that standard causal inference methods impose through separate modeling steps [26, 27].

Positional encoding for treatment timing

Standard positional encodings in transformers encode absolute or relative position in the sequence, but counterfactual dose estimation requires distinguishing periods relative to treatment changes. The framework uses relative positional encodings that explicitly encode the time since the most recent dose change and the time until the next scheduled blood pressure measurement [5, 28]. For a patient whose dose was increased from lisinopril 10mg to 20mg at week 4, the encoding for week 5 should indicate "1 week post-dose increase," while the encoding for week 3 should indicate "1 week pre-dose increase." This temporal structuring enables the attention mechanism to learn patterns such as "blood pressure changes typically plateau

4 weeks after dose adjustment" across patients with different absolute treatment timings [20, 21].

Transformer Architecture

Input sequence

Each time step in the input sequence (typically weekly intervals, though irregularly sampled data can be handled with time-aware positional encodings) comprises a vector of clinical variables: systolic blood pressure (mmHg), diastolic blood pressure (mmHg), medication type encoded as a categorical variable with levels for ACEi, ARB, CCB, thiazide, beta-blocker, and combinations thereof, current dose in milligrams normalized by the maximum approved dose, co-medication indicators, age in years, weight in kilograms, and estimated glomerular filtration rate from serum creatinine [25, 26]. Missing values are common in EHR data; the framework uses forward imputation for blood pressure (carrying last observation forward) and indicator flags for missingness, though more sophisticated approaches such as missingness attention masks could be incorporated [27, 29].

Causal transformer encoder

The encoder consists of L stacked layers, each containing multi-head self-attention with the causal mask described in Section 4, followed by layer normalization and position-wise feedforward networks [1]. Each attention head computes scaled dot-product attention:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}} + M\right)V, \text{ where } M \text{ is the causal}$$

mask matrix with zeros for allowed positions and $-\infty$ for masked positions. The feedforward network applies two linear transformations with a ReLU activation:

$$\text{FFN}(x) = \max(0, xW_1 + b_1)W_2 + b_2$$

applied before each sub-layer (pre-normalization architecture), with residual connections around each sub-layer to facilitate gradient flow. The output at each time step is a hidden representation that summarizes the causally-permissible history up to that time, which is then passed to the counterfactual prediction head described in Section 6 [26, 28].

Counterfactual Estimation

Counterfactual prediction objective

The counterfactual prediction task is defined as follows: given a patient's observed history up to time t , predict the blood pressure at time $t+k$ (typically 4 weeks for antihypertensive dose response) under an alternative dose path " $do(D_t = d)$ " that may differ from the dose actually administered [27, 28]. During training, the model learns from observed transitions where the actual dose and actual outcome are known; the objective is to minimize prediction error for factual outcomes while enforcing constraints that counterfactual predictions differ from factual predictions when the alternative dose differs from the observed dose. For a patient who received a dose increase from 10mg to 20mg at time t and had a subsequent blood pressure of 135/85 at $t+4$, the model must learn to predict that outcome for the factual scenario while also predicting what the blood pressure would have been under counterfactual scenarios (remain at 10mg, increase to 40mg, or switch to amlodipine) [8-10].

Contrastive loss for counterfactuals

Standard training with only factual prediction error does not ensure that the model learns a valid causal mapping, as the model could ignore the dose variable entirely and predict based solely on baseline characteristics. The framework incorporates a contrastive loss that explicitly regularizes the representation to distinguish between different dose regimens [1, 11]. For a given patient and time point, the model generates predictions under multiple dose alternatives; the loss encourages the representations for different doses to diverge when the predicted outcomes differ, while remaining similar when the predicted outcomes are clinically equivalent. This approach builds on adversarial balanced representations and generative counterfactual estimation methods, adapted to the transformer architecture with causal masking [2, 3, 6]. The complete loss function combines factual mean squared error (between predicted and observed blood pressure for the actual dose), contrastive loss (encouraging dose-specific representation separation), and a regularization term that penalizes violations of the causal hierarchy constraints derived from the causal graph [4, 7].

Treatment Effect Heterogeneity

Individualized dose response

For each patient at each clinical decision point, the causal transformer estimates individualized dose-response curves by feeding multiple counterfactual dose regimens through the encoder with the same historical context [12, 13]. The output includes predicted systolic and diastolic blood pressure at 4, 8, and 12 weeks post-dose-change for each candidate dose: for lisinopril, predictions are generated for 5mg, 10mg, 20mg, and 40mg (or up to the maximum approved dose). The predicted dose-response curve enables identification of the minimum effective dose—the lowest dose that achieves target blood pressure (below 130/80)—as well as estimation of the incremental benefit of each dose escalation step. For a patient with baseline systolic pressure of 148 mmHg on lisinopril 10mg, the model might predict that increasing to 20mg achieves 130 mmHg (10 mmHg reduction) but increasing to 40mg achieves only 128 mmHg (2 mmHg additional reduction), suggesting diminishing returns beyond 20mg [14, 20].

Table 2 translates the model's counterfactual outputs into clinically interpretable decision options while identifying the principal causal validity threat associated with each option.

Table 2. Counterfactual Decision Matrix for Individualized Antihypertensive Dose Selection

Clinical decision option	Counterfactual question estimated by the model	Required model output	Main validity threat
Maintain current dose	What would BP be if the patient remained on the present dose?	Predicted SBP/DBP at 4, 8, and 12 weeks under no titration	Confounding from patients' stable disease being more likely remain untreated
Increase current dose	What would BP be if the current medication were titrated upward?	Dose-response curve across approved dose levels	Confounding by indication because escalation given to patients with uncontrolled BP

Increase to maximum tolerated dose	Would maximal dose provide additional benefit beyond moderate titration?	Marginal BP reduction from intermediate to high dose	Positive violation: patient's rare prescription high dose
Add second agent	Would combination therapy outperform dose escalation alone?	Predicted BP under add-on therapy versus higher monotherapy dose	Treatment selection bias from comorbidity side effects and clinical preference
Switch drug class	Would an alternative drug class produce better BP control?	Predicted BP under replacement regimen	Sparingly used switch: data on unmeasured intolerance, historical
Defer automated recommendation	Are predictions insufficiently supported by observed data?	Uncertainty estimate, overlap warning, or missingness warning	Weakened positive: missing covariates, suspected unmeasured confounding

Subgroup analysis

Beyond individual predictions, the framework supports subgroup analysis to identify patient characteristics associated with differential treatment benefit [15, 16]. By aggregating individualized dose-response estimates across patient cohorts, the model can query which subgroups show large responses to dose escalation versus which subgroups benefit more from add-on therapy. Potential effect modifiers include age (older adults often show greater sensitivity to ACE inhibitors but higher risk of adverse effects), race (Black patients typically have smaller renin-angiotensin system-mediated responses and greater benefit from thiazides or CCBs), baseline blood pressure severity, presence of chronic kidney disease (which alters drug pharmacokinetics and contraindicates certain agents at high doses), and genetic polymorphisms in drug-

metabolizing enzymes [21, 29]. The transformer's attention weights themselves can be analyzed to identify which clinical features the model relies upon when making dose-specific counterfactual predictions, providing a form of explainable AI that complements the causal framework [17, 18].

Evaluation Strategy

Backtesting on observational data

In the absence of a gold-standard counterfactual dataset (where both the observed outcome and the outcome under an alternative dose are known for the same patient), evaluation relies on backtesting procedures that assess factual prediction accuracy and calibration [19, 22]. The model's factual predictions (blood pressure following the observed dose) are compared to actual outcomes using metrics including mean absolute error, root mean squared error, and calibration of prediction intervals across clinically relevant subgroups (e.g., defined by baseline BP strata, age groups, and comorbidity categories). A well-specified model should show no systematic bias across dose levels: overprediction of blood pressure in patients who received high doses would suggest residual confounding [23, 24]. Temporal cross-validation (training on earlier time periods, testing on later periods) assesses whether the causal structure remains stable over calendar time, which is particularly important given changes in hypertension treatment guidelines and the introduction of fixed-dose combination pills during the study period [25, 26].

Benchmark against causal methods

Performance is benchmarked against established causal inference methods for time-series treatment effect estimation, including inverse probability of treatment weighting marginal structural models, G-computation with recurrent neural networks, and doubly-robust estimators that combine outcome modeling with propensity score weighting [14, 27]. Since ground truth counterfactuals are never observed in real EHR data, benchmarking requires semi-synthetic datasets where the data-generating process is known and the true counterfactual outcomes can be computed. The simulation framework generates realistic hypertension patient trajectories based on clinical parameters extracted from published trial data (e.g., the SPRINT trial for intensive blood pressure lowering), with known dose-response functions and specified confounding structures [28, 29]. The causal transformer's mean squared

error for counterfactual predictions is compared to benchmark methods across varying degrees of confounding and different patterns of treatment discontinuation, with particular attention to performance under positivity violations (e.g., no patient with stage 2 hypertension remains on the lowest dose for extended periods) [5, 20].

Sensitivity analysis

Three sensitivity analyses assess the robustness of conclusions to violations of the causal assumptions underlying the framework. Unmeasured confounding analysis introduces realistic unmeasured confounders (e.g., medication adherence estimated from pharmacy refill data but often missing in EHRs, dietary sodium intake, physical activity levels, socioeconomic status) into the simulation and quantifies how severely the causal transformer's estimates are biased as a function of the unmeasured confounder's strength [6, 21]. Positivity violation analysis examines performance in regions of the covariate space where certain dose levels are rarely prescribed, such as very high ACE inhibitor doses in patients with advanced chronic kidney disease, using propensity score overlap diagnostics to identify unreliable predictions. Missingness analysis evaluates the impact of non-random missing blood pressure measurements—for example, patients with poorly controlled hypertension may have more frequent clinic visits (more measurements) while those with excellent control may have fewer—by artificially inducing missingness patterns and comparing complete-case analysis to the proposed forward imputation with missingness indicators [7, 8].

Limitations

Technical limitations

The causal transformer inherits several technical limitations from its parent architecture and causal framework. The sequential ignorability assumption remains fundamentally untestable from observational data; although the causal attention mask enforces temporal ordering, it cannot detect or correct for unmeasured confounders that affect both treatment decisions and outcomes [9-11]. Attention patterns may capture spurious correlations that happen to satisfy the causal mask but do not correspond to genuine causal mechanisms, particularly when time-varying confounders are measured with error or at irregular intervals. The computational complexity of multi-head self-attention scales

quadratically with sequence length $O(T^2)$, which becomes prohibitive for patients with very long hypertension histories (e.g., 10+ years of weekly observations), though this can be mitigated with sparse attention mechanisms or sequence truncation at the cost of losing early history information [1, 2, 5]. Finally, the framework models dose as a categorical or discrete-continuous variable, but continuous dose optimization (e.g., finding the exact milligram dose that achieves target BP) would require either extensive discretization or a fundamentally different approach such as a dose-conditioned neural network [3, 4].

Clinical limitations

From a clinical perspective, several important limitations constrain the framework's immediate applicability. Unmeasured confounders common in hypertension research—including medication adherence (patients who take 80% of prescribed doses vs those who take 50%), dietary sodium intake (which modifies BP response to ACE inhibitors and diuretics), physical activity, alcohol consumption, and socioeconomic status—remain major concerns that no purely observational method can fully resolve [12-14]. The framework assumes that treatment decisions occur at discrete time points (typically clinic visits), but in reality, patients may adjust doses in response to home blood pressure monitoring between visits, and such adjustments are rarely recorded in structured EHR fields. Validation in prospective randomized trials would be required before clinical deployment, as the framework cannot replace the evidence standard of a well-conducted randomized controlled trial such as those evaluating fixed-dose triple combinations or single-pill combinations [15, 16, 26, 27]. Additionally, the framework currently models blood pressure as the only outcome, but clinical decision-making also balances efficacy against adverse effects (e.g., cough with ACE inhibitors, edema with CCBs, electrolyte disturbances with thiazides), and extending the framework to multi-outcome counterfactual estimation with trade-offs remains an open challenge [17, 18].

Conclusion

This manuscript has presented a causal transformer architecture with attention masking for estimating counterfactual antihypertensive dose responses from observational electronic health records. The framework extends standard sequence models by explicitly encoding the temporal ordering required for causal identification,

ensuring that predictions of potential outcomes under alternative dose regimens use only information available before the treatment decision. The causal attention mask, contrastive loss for counterfactual separation, and individualized dose-response estimation together provide a principled approach to personalized treatment recommendation from observational data.

The key advantages of this framework over conventional methods include its end-to-end learning of causally-valid representations without manual specification of propensity score models, its ability to handle both dose titration and add-on therapy decisions within a unified architecture, and its patient-specific predictions that capture treatment effect heterogeneity across clinically relevant subgroups. Unlike marginal structural models that require separate modeling steps, the causal transformer directly encodes the data-generating process of sequential decision-making in hypertension management.

Nevertheless, important limitations remain. The sequential ignorability assumption that all confounders are measured cannot be validated from observational data alone, and unmeasured confounders such as adherence, diet, and socioeconomic status may bias estimates in ways that no purely statistical method can detect. Missing data patterns in EHRs are often non-random, and the computational demands of full self-attention limit applicability to very long patient histories. Prospective validation in randomized controlled trials or targeted trial emulations would be required before clinical deployment.

We call for implementation of the proposed causal transformer framework on large-scale hypertension EHR cohorts, including the UK Biobank, the All of Us Research Program, and national hypertension registries from healthcare systems with comprehensive structured data on drug dosing and blood pressure measurements. Such implementations would enable empirical evaluation of the framework's performance relative to existing causal inference methods, characterization of the types of patients for whom counterfactual dose predictions are most reliable, and ultimately, the development of clinical decision support tools that provide evidence-based, personalized antihypertensive dosing recommendations. Translating this framework from proof-of-concept to clinical practice will require interdisciplinary collaboration among machine learning researchers, causal inference methodologists, clinical pharmacologists, and practicing hypertension specialists.

Acknowledgements

None

None

Conflict of interest

None

Ethics statement

None

Financial support

Received: 27 Nov 2024 Revised: 23 Jan 2025 Accepted: 25 Feb 2025

Published online: 20 July 2025

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Schwab P, Linhardt L, Bauer S, Buhmann JM, Karlen W. Learning counterfactual representations for estimating individual dose-response curves. *Proc AAAI Conf Artif Intell*. 2020;34(4):5612–9.
- Bica I, Alaa A, van der Schaar M. Time series deconfounder: estimating treatment effects over time in the presence of hidden confounders. *Proc Int Conf Mach Learn (ICML)*. 2020;119:884–95.
- Yoon J, Jordon J, van der Schaar M. GANITE: estimation of individualized treatment effects using generative adversarial nets. *Int Conf Learn Represent (ICLR)*. 2018.
- Hu Y, Huerta J, Cordella N, Mishuris RG, Paschalidis IC. Personalized hypertension treatment recommendations by a data-driven model. *BMC Med Inform Decis Mak*. 2023;23(1):44.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst*. 2017;30:5998–6008.
- Melnychuk V, Frauen D, Feuerriegel S. Causal transformer for estimating counterfactual outcomes. *Proc Int Conf Mach Learn (ICML)*. 2022;162:15293–329.
- Bica I, Alaa AM, Jordon J, van der Schaar M. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *arXiv preprint arXiv:2002.04083*. 2020.
- Lim B. Forecasting treatment responses over time using recurrent marginal structural networks. *Adv Neural Inf Process Syst*. 2018;31:7493–503.
- Curth A, Svensson D, Weatherall J, van der Schaar M. Really doing great at estimating CATE? A critical look at ML benchmarking practices in treatment effect estimation. *NeurIPS Datasets and Benchmarks Track*. 2021:1–12.
- Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc Natl Acad Sci USA*. 2019;116(10):4156–65.
- Shalit U, Johansson FD, Sontag D. Estimating individual treatment effect: generalization bounds and algorithms. *Proc Int Conf Mach Learn (ICML)*. 2017;70:3076–85.
- Oikonomou EK, Spatz ES, Suchard MA, Khera R. Individualising intensive systolic blood pressure reduction in hypertension using computational trial phenomaps and machine learning: a post-hoc analysis of randomised clinical trials. *Lancet Digit Health*. 2022;4(11):e796–e805.
- Bress AP, Greene T, Derington CG, Shen J, Xu Y, Zhang Y, et al. Patient selection for intensive blood pressure management based on benefit and adverse events. *J Am Coll Cardiol*. 2021;77(16):1977–90.

Collier DJ, Taylor M, Godec T, Shiel J, James R, Chowdury Y, et al. Personalized antihypertensive treatment optimization with smartphone-enabled remote precision dosing of amlodipine during the COVID-19 pandemic (PERSONAL-CovidBP Trial). *J Am Heart Assoc.* 2024;13(4):e030749.

Ye X, Zeng QT, Facelli JC, Brixner DI, Conway M, Bray BE. Predicting optimal hypertension treatment pathways using recurrent neural networks. *Int J Med Inform.* 2020;139:104122.

Yi J, Wang L, Song J, Liu Y, Liu J, Zhang H, et al. Development of a machine learning-based model for predicting individual responses to antihypertensive treatments. *Nutr Metab Cardiovasc Dis.* 2024;34(7):1660–9.

Layton AT. AI, machine learning, and ChatGPT in hypertension. *Hypertension.* 2024;81(4):709–16.

Cavero-Redondo I, Martinez-Rodrigo A, Saz-Lara A, Moreno-Herraiz N, Casado-Vicente V, Gomez-Sanchez L, et al. Antihypertensive drug recommendations for reducing arterial stiffness in patients with hypertension: machine learning–based multicohort (RIGIPREV) study. *J Med Internet Res.* 2024;26:e54357.

Verma AA, Khuu W, Tadrous M, Gomes T, Mamdani MM. Fixed-dose combination antihypertensive medications, adherence, and clinical outcomes: a population-based retrospective cohort study. *PLoS Med.* 2018;15(6):e1002584.

Williams B, Mancia G, Spiering W, Agabiti Rosei E, Azizi M, Burnier M, et al. 2018 ESC/ESH Guidelines for the management of arterial hypertension. *Eur Heart J.* 2018;39(33):3021–104.

Mills KT, Stefanescu A, He J. The global epidemiology of hypertension. *Nat Rev Nephrol.* 2020;16(4):223–37.

Webster R, Salam A, De Silva HA, Selak V, Stepien S, Rajapakse S, et al. Fixed low-dose triple combination antihypertensive medication vs usual care for blood pressure

control in patients with mild to moderate hypertension in Sri Lanka: a randomized clinical trial. *JAMA.* 2018;320(6):566–79.

Salam A, Atkins ER, Hsu B, Webster R, Patel A, Rodgers A. Efficacy and safety of triple versus dual combination blood pressure-lowering drug therapy: a systematic review and meta-analysis of randomized controlled trials. *J Hypertens.* 2019;37(8):1567–73.

Mourad JJ, Amodeo C, de Champvallins M, Brzozowska-Villatte R, Asmar R. Blood pressure-lowering efficacy and safety of perindopril/indapamide/amlodipine single-pill combination in patients with uncontrolled essential hypertension. *J Hypertens.* 2017;35(7):1481–95.

Fleig SV, Weger B, Haller H, Limbourg FP. Effectiveness of a fixed-dose, single-pill combination of perindopril and amlodipine in patients with hypertension: a non-interventional study. *Adv Ther.* 2018;35(3):353–66.

DiPette DJ, Skeete J, Ridley E, Campbell NR, Lopez-Jaramillo P, Kishore SP, et al. Fixed-dose combination pharmacologic therapy to improve hypertension control worldwide: clinical perspective and policy implications. *J Clin Hypertens.* 2019;21(1):4–15.

Burnier M, Egan BM. Adherence in hypertension: a review of prevalence, risk factors, impact, and management. *Circ Res.* 2019;124(7):1124–40.

Eadon MT, Kanuri SH, Chapman AB. Pharmacogenomic studies of hypertension: paving the way for personalized antihypertensive treatment. *Expert Rev Precis Med Drug Dev.* 2018;3(1):33–47.

Rabi DM, McBrien KA, Sapir-Pichhadze R, Nakhla M, Ahmed SB, Dumanski SM, et al. Hypertension Canada's 2020 comprehensive guidelines for the prevention, diagnosis, risk assessment, and treatment of hypertension in adults and children. *Can J Cardiol.* 2020;36(5):596–624.