POLICY-AWARE PRETRAINING FOR EHR AS CLINICIAN-IN-THE-LOOP DYNAMICAL SYSTEMS

Fu Huiliang†, Hu Hong†,Tao Jingfei†, Guo Fengge ‡, Cai Ning†, Yuanyun Zhang †, Shi Li‡

†University of the Chinese Academy of Sciences, Columbia University‡

ABSTRACT

Electronic health records (EHR) are often modeled as irregular sequences of codes, labs, and interventions, yet this framing overlooks that an EHR is the execution trace of a clinician-in-the-loop control system. Standard foundation models trained with masked prediction objectives conflate physiology with institutional policy, yielding predictors that excel at forecasting what clinicians at one hospital will do but generalize poorly under distribution shift or counterfactual query. We introduce *CLIO* (Causal Latent, Intervention- and Operations-aware), a foundation model that separates patient dynamics from workflow by explicitly modeling both physiology and the observation–action policy. Patient state evolves as a controlled stochastic differential equation, while observations and interventions are generated by marked point processes whose intensities capture clinician belief and operational constraints. Pretraining optimizes a composite objective: policy-aware masked modeling reweighted by propensities from the policy surrogate; an invariance penalty enforces environment-agnostic dynamics across hospitals and epochs; and a counterfactual contrastive task aligns representations to respect intervention effects while discounting spurious workflow correlations. A lightweight diffusion head generates feasible action schedules under resource constraints, and clinician-in-the-loop preference learning distills trajectory-level comparisons into a value functional that shapes the latent space toward clinical utility. Evaluations across multiple institutions show improved discrimination, calibration, and counterfactual fidelity under distribution shift, with reductions in policy regret and subgroup disparities. CLIO demonstrates that treating the EHR as a controlled system rather than a passive sequence yields representations that are stable, actionable, and aligned with the causal questions clinicians face.

Index Terms— EHR-Foundation Model, Clinician-In-The-Loop

1. INTRODUCTION

Electronic health records (EHR) are often treated as an unruly concatenation of codes, labs, medications, and notes [1, 2]. In reality they are the execution trace of a clinician-in-the-loop [3] control system operating under uncertainty and constraints [4, 5]. Each observation is produced because someone chose to look; each intervention is chosen subject to hospital policy, resource limits, and tacit priors over patient state. Models that ignore this interventional substrate routinely learn to forecast what this hospital usually does rather than what will happen to this patient under an alternative course of action. The result is impressive next-event likelihoods that do not travel across sites or support credible counterfactual reasoning.

Foundation models [6, 7, 8]have become a popular paradigm for AI in Healthcare [9, 10, 11] with many works appearing in recent years [12, 13, 14, 15, 16, 17, 18, 19]. We propose to pretrain a foundation model over EHR that is explicitly policy-aware and mechanism-seeking. Let a patient trajectory be a partially observed controlled stochastic process with latent state $Z_t \in \mathbb{R}^d$, clinician actions $A_t \in \mathcal{A}$, and observations $X_{T_k} \in \mathcal{X}$ arriving at irregular times $\{T_k\}$. The data-generating story is

$$\mathrm{d}Z_t = f_\theta(Z_t, A_t)\,\mathrm{d}t + \Sigma_\theta(Z_t)^{1/2}\,\mathrm{d}W_t,$$

$$X_{T_k} \sim g_\theta(Z_{T_k}, A_{T_k}),$$

$$(A_t,\, \mathrm{d}N_t) \sim \Pi_\phi(\cdot \mid \mathcal{H}_t, \mathcal{C}_t),$$

where $\mathrm{d}N_t$ is the increment of a marked point process governing which labs, vitals, and orders materialize; $\mathcal{H}_t$ is the observable history; and $\mathcal{C}_t$ captures operational constraints and institutional policy. Standard masked modeling optimizes θ while implicitly averaging over Π_ϕ [20]. In contrast, we learn a representation that separates and then recouples these components, yielding a physiology-centered state Z_t that is invariant to site-specific workflows and an explicit surrogate Π_ϕ that explains when and why observations occur.

The foundation model, *CLIO* (Causal Latent, Intervention- and Operations-aware), is trained to align three capabilities. First, it must forecast latent clinical dynamics in continuous time, such that $\hat{Z}_{t+\Delta}$ remains clinically calibrated under irregular sampling and variable action schedules. Second, it must support counterfactual rollouts of the form $X^{\mathrm{do}(a)}$ by intervening on the control input while holding exogenous noise fixed, enabling "what-if" queries at pretraining time rather than relegating them to downstream fine-tuning. Third, it must model the observation process itself, treating missingness and test ordering as informative signals about clinician belief and constraint rather than ignorable censoring.

To operationalize these goals, we replace likelihood-only pretraining with an objective built around actionable invariances. Reconstruction is reweighted by inverse propensities for orders and treatments, producing a do-calibrated masked modeling term that discourages the model from equating "frequent tests" with "important physiology." Latent dynamics are trained via a neural controlled differential equation or stochastic differential equation, regularized so that gradients of the loss with respect to θ are stable across environments indexed by hospital, service line, and epoch. A counterfactual contrastive task draws positive pairs from simulated interventions $\mathrm{do}(A_t \leftarrow a)$ using the learned mechanism while designating co-occurrences attributable to local workflow quirks as hard negatives. Jointly fitting a marked point process Π_ϕ for actions and measurements makes the absence of an event as informative as its presence and prevents the model from spuriously attributing policy-induced censoring to physiology.

Representation and tokenization follow the same philosophy. Rather than collapsing all events into a code–value–timestamp triad, interventions are represented as control primitives with semantics, constraints, and budgets: a vasopressor start is encoded with dose range, titration policy, and a shadow price reflecting resource load. These controls drive a diffusion-style generator over action times [21], while measurements remain marks of a conditional point process [22] tied to state. This separation lets the model imagine clinically feasible intervention schedules under resource constraints and reduces the temptation to hallucinate impossible sequences that a pure language model might produce.

Alignment [23] is treated as a first-class training signal. Instead of free-form preference modeling over text, we fit a preference functional over short, de-identified trajectory snippets where CLIO proposes alternative action sequences and their counterfactual outcomes [24]. Clinicians express pairwise preferences with soft priors from guidelines; these labels are distilled into a value functional $V(Z_t, a)$ that shapes the latent space toward utility without directly optimizing for any single task. Training proceeds in a federated manner: mechanism and preference gradients are privately aggregated across institutions so that environment diversity is preserved while raw data remain on-premise, avoiding the homogenization that undermines cross-site generalization.

This paper argues that an EHR foundation model should not be a sophisticated autocompleter [25] for observational traces, but a simulator-backed representation learner that respects interventions, constraints, and the economics of attention. By embedding an explicit observation policy [26], enforcing environment-wise invariances [27, 28, 29, 30], and pretraining on counterfactual consistency [24], CLIO aims to yield state that travels across hospitals and time (a goal highlighted in many recent frameworks [31, 32, 33]), supports safe what-if reasoning, and provides a common substrate for downstream tasks from mortality prediction to treatment effect estimation. The remainder details architecture and objectives, a privacy-preserving multi-site training plan, and evaluations that stress causal generalization under distribution shift rather than in-sample calibration.

2. RELATED WORKS

2.1. EHR Foundation Models

The first wave of EHR foundation models largely arose from direct adaptations of Transformer architectures to structured clinical data streams. Early work such as BEHRT framed diagnosis and procedure codes as linguistic tokens, treating a patient's longitudinal visits as sentences. Each event was embedded with a combination of code, age, positional, and segment encodings, and the model was pretrained using masked language modeling over 1.6 million records from the UK Biobank [34]. Med-BERT extended this strategy to a much larger cohort of 20 million patients, introducing a secondary pretraining task—predicting extended hospital stays—while simplifying the embedding layer by removing explicit age and segment vectors [35]. In parallel, autoregressive formulations were explored: CEHR-GPT employed next-event prediction as its core learning objective [36], while CEHR-BERT preserved the masked-LM setup but refined its temporal sensitivity by inserting synthetic time tokens to represent inter-visit gaps and substituting Next-Sentence Prediction with visit-type classification [37]. Collectively, these studies demonstrated that both bidirectional and autoregressive pretraining paradigms can capture the temporal and semantic structure of EHR sequences far more effectively than recurrent neural baselines.

Subsequent models began to emphasize time-resolved dynamics and clinical outcomes more directly. CLMBR introduced an autoregressive loss over OMOP-standardized codes, initially using GRUs and later scaling to a 141M-parameter Transformer. This design proved robust across multiple downstream settings, including mortality, readmission, and ICU admission tasks, even under distributional shift [17]. MOTOR pushed this line further by equipping its 143M-parameter Transformer with a piecewise-exponential time-to-event prediction head, enabling it to model explicit intervals until diverse clinical outcomes [38]. By optimizing event-time objectives across tens of millions of trajectories, MOTOR substantially improved concordance scores and label efficiency across datasets spanning Stanford's EHR to national claims.

Addressing variation across coding standards and institutional schemas, recent work has shifted toward text-centric unification. GenHPF [39] reformulates heterogeneous EHR elements—including codes, feature names, and numerical values—into a hierarchical textual format. This representation, tokenized with clinical subword vocabularies, is processed by an event-level Transformer with temporal aggregation. Similar schema-free strategies have been explored in other contexts [40, 19]. By converting data to a semantic text stream, GenHPF bypasses bespoke mappings between ICD, SNOMED, and local vocabularies, facilitating unified pretraining across datasets such as MIMIC-III [41], MIMIC-IV [42], and eICU [43]. Coupling this representation with contrastive and clustering objectives, the model achieves strong multi-source and multi-task performance, highlighting the potential of semantic conversion for robust cross-institutional transfer [44, 45, 46].

3. METHODOLOGY

The proposed foundation model, CLIO, treats the EHR as the execution trace of a clinician-in-the-loop dynamical system [47]. A patient encounter is represented as a continuous-time, partially observed, controlled stochastic process [48] coupled to two marked point processes that govern what gets measured and what gets done. The core learning problem is to identify a representation of latent physiology that is maximally invariant to institutional policy while remaining sensitive to interventions and constraints, and to do so with an objective that is differentiable end-to-end and supports counterfactual reasoning already at pretraining.

Generative formulation

Fix a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$. Each encounter i unfolds over $[0, T_i]$ with a latent state process $Z_t^i \in \mathbb{R}^d$. The state obeys a controlled Itô SDE [49, 50, 51, 52],

$$\mathrm{d}Z_t = f_\theta(Z_t, A_t, \xi_t)\,\mathrm{d}t + \Sigma_\theta(Z_t)^{1/2}\,\mathrm{d}W_t, \qquad Z_0 \sim p_\theta(\cdot \mid b),$$

where W_t is a standard Brownian motion, ξ_t is a low-rate exogenous context process capturing non-clinician shocks (for example pathogen load or sudden bleeding) modeled as a piecewise-constant jump process, and b denotes baseline covariates. The control process $A_t \in \mathcal{A}$ is piecewise-constant with jumps at intervention times $\{\tau_\ell\}$; between jumps A_t remains constant. Observations arrive at irregular times $\{T_k\}$ and carry marks M_k that identify the measurement or action type and a payload Y_k, producing the observed EHR event $X_k = (T_k, M_k, Y_k)$.

Two coupled marked point processes generate events [53, 54, 55]. The measurement process N^{obs} has conditional intensities $\lambda_m^{\text{obs}}(t \mid \mathcal{H}_t, \mathcal{C}_t)$ for marks $m \in \mathcal{M}_{\text{obs}}$, where $\mathcal{H}_t$ is the history of states, actions, and realized events, and $\mathcal{C}_t$ encodes site constraints and policy context. The action process N^{act} has intensities $\lambda_a^{\text{act}}(t \mid \mathcal{H}_t, \mathcal{C}_t)$ for action marks $a \in \mathcal{M}_{\text{act}}$. Conditional on an event at time t with mark m, the payload is drawn from g_θ if it is a measurement and updates the control A_t if it is an action. The stochastic generative story is therefore

$$\begin{aligned}
&\mathrm{d}Z_t = f_\theta(Z_t, A_t, \xi_t)\,\mathrm{d}t + \Sigma_\theta(Z_t)^{1/2}\,\mathrm{d}W_t,\\
&\mathbb{P}\big(\text{obs mark } m \text{ in } [t, t+\mathrm{d}t) \mid \mathcal{F}_{t^-}\big) = \lambda_m^{\text{obs}}(t \mid \mathcal{H}_t, \mathcal{C}_t)\,\mathrm{d}t,\\
&\mathbb{P}\big(\text{act mark } a \text{ in } [t, t+\mathrm{d}t) \mid \mathcal{F}_{t^-}\big) = \lambda_a^{\text{act}}(t \mid \mathcal{H}_t, \mathcal{C}_t)\,\mathrm{d}t,\\
&Y_k \sim g_\theta(\,\cdot \mid Z_{T_k}, A_{T_k}, M_k), \qquad A_{t^+} = \Gamma_\theta(A_{t^-}, M_k, Y_k) \;\text{ if } M_k \in \mathcal{M}_{\text{act}}.
\end{aligned}$$

The mapping Γ_θ codifies how action marks update the control vector; an antibiotic start sets the component for that drug to a dose state; a titration modifies it according to a policy kernel; a cessation zeroes it out. In the background, an environment label $e \in \mathcal{E}$ indexes site, service line, and epoch, and influences only the policy-and-operations channels $\{\lambda^{\text{obs}}, \lambda^{\text{act}}\}$ and the constraints $\mathcal{C}_t$, not the physiology $f_\theta, \Sigma_\theta, g_\theta$. This conditional exclusion is the mechanism we aim to identify.

Representation and tokenization

Raw EHR events are embedded into two continuous-time control streams [56]. Observations are mapped to a signed measure-valued control U^{obs} with atoms at $\{T_k\}$ carrying mark-dependent embeddings $\phi_{\text{obs}}(M_k, Y_k)$. Actions are mapped to a piecewise-constant control $U^{\mathrm{act}}(t) \in \mathbb{R}^{d_a}$ with jumps at $\{\tau_\ell\}$. Each action mark $m \in \mathcal{M}_{\mathrm{act}}$ is represented as an actionable token with parameters $(\mathrm{role}, \mathrm{dose}, \mathrm{titration}, \mathrm{budget})$. The budget is a scalar $\beta_m \geq 0$ reflecting shadow price of resource consumption; it conditions the policy surrogate and governs feasibility in counterfactual simulation. These tokens are embedded via $\phi_{\mathrm{act}}(m; \beta_m)$ and deterministically decoded into updates of A_t. The history $\mathcal{H}_t$ is featurized by a log-signature encoder $\Psi\big(U^{\mathrm{obs}}|_{[0,t]}, U^{\mathrm{act}}|_{[0,t]}\big)$ that produces a path summary invariant to time-warping at fine scales while preserving order.

Mechanism parameterization

The physiology is modeled with a Neural CDE/SDE hybrid [57, 58, 59, 60]. Let X_t denote a continuously interpolated control path derived from $(U^{\mathrm{obs}}, U^{\mathrm{act}})$ via cubic splines for observations and zero-order hold for actions [61, 62]. The latent evolves as

$$\mathrm{d}Z_t = F_\theta(Z_t)\,\mathrm{d}t + G_\theta(Z_t)\,\mathrm{d}\widetilde{X}_t + \Sigma_\theta(Z_t)^{1/2}\,\mathrm{d}W_t,$$

where $\widetilde{X}_t$ is the control path augmented with a time channel and environment-agnostic basis functions. The drift F_θ, the control Jacobian G_θ [63, 64], and diffusion Σ_θ are neural networks constrained to ensure linear growth and local Lipschitzness; spectral normalization and bounded activation schemes guarantee existence and uniqueness of strong solutions. Emissions use a mark-specific head

$$Y_k \mid Z_{T_k}, A_{T_k}, M_k \;\sim\; \begin{cases} \mathcal{N}\big(\mu_\theta^{(m)}(Z_{T_k}, A_{T_k}),\, \mathrm{diag}(\sigma_\theta^{2,(m)}(Z_{T_k}, A_{T_k}))\big) & \text{if } m \text{ is continuous,} \\ \mathrm{Cat}\big(\pi_\theta^{(m)}(Z_{T_k}, A_{T_k})\big) & \text{if } m \text{ is categorical,} \\ \mathrm{ZINB}\big(\alpha_\theta^{(m)}(Z_{T_k}, A_{T_k})\big) & \text{if } m \text{ is count-valued with zeros.} \end{cases}$$

For text snippets, a lightweight encoder produces latent sufficient statistics $\eta_\theta^{(m)}(Z_{T_k}, A_{T_k})$ that parameterize a masked language head; notes are never reconstructed verbatim but compressed into clinical concept distributions to respect privacy.

Policy-and-operations surrogate

The joint observation–action policy is modeled as a doubly stochastic marked point process with conditional intensities

$$\lambda_m^{\mathrm{obs}}(t \mid \mathcal{H}_t, \mathcal{C}_t) = \exp\left\{\alpha_m^{\mathrm{obs}} + u_m^{\mathrm{obs}}(\Psi_t) + v_m^{\mathrm{obs}}(\mathcal{C}_t)\right\}, \qquad \lambda_a^{\mathrm{act}}(t \mid \mathcal{H}_t, \mathcal{C}_t) = \exp\left\{\alpha_a^{\mathrm{act}} + u_a^{\mathrm{act}}(\Psi_t) + v_a^{\mathrm{act}}(\mathcal{C}_t)\right\},$$

where $\Psi_t = \Psi(\mathcal{H}_t)$ summarizes history. The baseline terms α and the functions $u(\cdot), v(\cdot)$ are neural networks whose parameters ϕ are environment-specific through $\mathcal{C}_t$ but share global initialization to permit meta-learning across sites. The joint likelihood for observed marks and times over $[0, T]$ is

$$\mathcal{L}_{\mathrm{MPP}}(\phi) = \sum_{m \in \mathcal{M}_{\mathrm{obs}}} \left[\sum_{k: M_k = m} \log \lambda_m^{\mathrm{obs}}(T_k \mid \mathcal{H}_{T_k}, \mathcal{C}_{T_k}) - \int_0^T \lambda_m^{\mathrm{obs}}(t \mid \mathcal{H}_t, \mathcal{C}_t)\,\mathrm{d}t\right] + \\ \sum_{a \in \mathcal{M}_{\mathrm{act}}} \left[\sum_{\ell: \tau_\ell \text{ mark } a} \log \lambda_a^{\mathrm{act}}(\tau_\ell \mid \mathcal{H}_{\tau_\ell}, \mathcal{C}_{\tau_\ell}) - \int_0^T \lambda_a^{\mathrm{act}}(t \mid \mathcal{H}_t, \mathcal{C}_t)\,\mathrm{d}t\right].$$

This surrogate supplies propensities for policy-aware reweighting and is also used to sample plausible counterfactual observation schedules under modified constraints $\mathcal{C}'_t$.

Variational inference over latent trajectories

Irregular sampling and partially observed emissions necessitate amortized inference [65, 66, 67]. A recognition model $q_\psi(Z_{0:T} \mid \mathcal{H}_{0:T})$ is defined as a Neural CDE backward smoother with stochastic bridge corrections. In practice, we use Euler–Maruyama [68, 69, 70] with Brownian bridge proposals between event times; reparameterization gradients propagate through the SDE solution via adjoint methods with Hutchinson trace estimators for diffusion derivatives. The evidence lower bound for a single encounter is

$$\log p_\theta(X_{0:K} \mid A_{0:T}) \;\geq\; \mathbb{E}_{q_\psi}\left[\sum_{k=0}^{K} \log g_\theta(Y_k \mid Z_{T_k}, A_{T_k}, M_k) + \log p_\theta(Z_0) + \sum_j \log p_\theta(\Delta Z_{(t_{j-1}, t_j]} \mid Z_{t_{j-1}}, A_t)\right] -$$

$$\mathrm{KL}(q_\psi \,\|\, p_\theta(\cdot \mid A_{0:T}))\,,$$

where the second sum ranges over a partition aligned to event times. When diffusion is present, the transition log-density uses an Euler–Maruyama Gaussian approximation with Itô correction; in the pure CDE case, the latent path is deterministic given Z_0 and the control, and the KL collapses to the prior on Z_0.

Pretraining objectives

The training criterion blends four terms: policy-aware masked modeling, environmental invariance, counterfactual contrastive alignment, and explicit point-process likelihood for the observation–action policy. Let $\widehat{\pi}_\phi$ denote the fitted surrogate and e the environment label.

The policy-aware masked modeling term uses inverse probability weights to counteract selection by measurement and treatment. For an observed event (T_k, M_k, Y_k), define the stabilized weight

$$w_k \;=\; \frac{\bar{\lambda}^{\text{obs}}_{M_k}}{\lambda^{\text{obs}}_{M_k}(T_k \mid \mathcal{H}_{T_k}, \mathcal{C}_{T_k})}, \qquad \bar{\lambda}^{\text{obs}}_{m} = \frac{1}{|\mathcal{D}|} \sum_{(i,k):M_k=m} \frac{1}{T_i} \int_0^{T_i} \lambda^{\text{obs}}_m(t \mid \mathcal{H}_t, \mathcal{C}_t)\,\mathrm{d}t,$$

and analogously for action-conditioned reconstructions when emissions depend on recent actions. The reconstruction loss is

$$\mathcal{J}_{\text{recon}}(\theta,\psi,\phi) = \mathbb{E}_{q_\psi}\left[\sum_{k=0}^{K} w_k\, \ell_{\text{pred}}\big(Y_k, \hat{Y}_k(Z_{T_k}, A_{T_k}, M_k)\big)\right],$$

with ℓ_{pred} the negative log-likelihood induced by g_θ. Stabilization avoids variance blow-up while shifting learning toward under-ordered yet physiologically informative labs.

The environmental invariance term enforces that mechanism learning is insensitive to the environment index. A gradient alignment penalty inspired by IRM/VREx is used. Write the environment-specific risk

$$R_e(\theta,\psi) = \mathbb{E}_{(i:e_i=e)}\, \mathbb{E}_{q_\psi}\left[\sum_k \ell_{\text{pred}}(Y_k, \hat{Y}_k)\right],$$

and define the squared deviation of per-environment gradients from their mean,

$$\mathcal{J}_{\text{inv}}(\theta,\psi) = \sum_{e\in\mathcal{E}} \left\| \nabla_\theta R_e(\theta,\psi) - \frac{1}{|\mathcal{E}|}\sum_{e'} \nabla_\theta R_{e'}(\theta,\psi) \right\|_2^2 .$$

A complementary representation-level discrepancy is applied to the latent path marginals, using MMD [71, 72] between $\{Z_t\}_{t\in\mathcal{T}}$ sampled across environments conditional on similar clinical contexts, which is implemented by matching on propensity scores and time-aligned windows. The combined invariance penalty is a weighted sum of gradient and distributional terms.

The counterfactual contrastive alignment term instills respect for interventions. For a window $[t, t+\Delta]$ with initial latent Z_t and realized action $A_{[t,t+\Delta)}$, a positive counterfactual is generated by fixing the exogenous noise path $W_{[t,t+\Delta]}$ and replacing the action with $\tilde{A}_{[t,t+\Delta)}$ sampled from a feasible alternative set $\mathcal{U}(\mathcal{C}_t)$, yielding $Z^{\text{do}}_{t+\Delta}$. A representation map $h_\theta : \mathbb{R}^d \to \mathbb{R}^r$ defines embeddings $\zeta = h_\theta(Z_{t+\Delta})$ and $\tilde{\zeta} = h_\theta(Z^{\text{do}}_{t+\Delta})$. Negatives $\{\zeta_j^-\}$ are drawn from contemporaneous states explained by workflow coincidences rather than physiology, identified by high policy-propensity but low state-similarity windows across environments. The InfoNCE objective [73, 74, 75] is

$$\mathcal{J}_{\text{cf}}(\theta,\psi) = \mathbb{E}\left[-\log \frac{\exp\big(\mathrm{sim}(\zeta,\tilde{\zeta})/\tau\big)}{\exp\big(\mathrm{sim}(\zeta,\tilde{\zeta})/\tau\big) + \sum_j \exp\big(\mathrm{sim}(\zeta,\zeta_j^-)/\tau\big)}\right],$$

with cosine similarity sim and temperature τ. Keeping the Brownian path fixed isolates the action's causal effect in the learned embedding.

The observation–action policy term is the negative log-likelihood of the marked point processes,

$$\mathcal{J}_{\text{mpp}}(\phi) = -\mathcal{L}_{\text{MPP}}(\phi),$$

and couples to the reconstruction via the propensities w_k and to the counterfactual generator via feasible schedule sampling. The total pretraining loss is

$$\mathcal{J}(\theta, \psi, \phi) = \mathcal{J}_{\text{recon}}(\theta, \psi, \phi) + \lambda_{\text{inv}} \mathcal{J}_{\text{inv}}(\theta, \psi) + \lambda_{\text{cf}} \mathcal{J}_{\text{cf}}(\theta, \psi) + \lambda_{\text{mpp}} \mathcal{J}_{\text{mpp}}(\phi).$$

Alignment via trajectory preferences

Alignment is formulated as preference learning over de-identified trajectory snippets. A snippet $s = (t_0, \Delta, \mathcal{C}_{t_0})$ induces a baseline rollout ω generated by the recognition posterior and observed actions, and K counterfactual rollouts $\{\omega^{(k)}\}$ generated by replacing the action schedule over $[t_0, t_0 + \Delta]$ with $\{\tilde{A}^{(k)}\}$ drawn from a feasibility set. Each rollout maps to a reward functional via a value head V_η that integrates clinically meaningful surrogates:

$$\mathcal{R}_\eta(\omega) = \int_{t_0}^{t_0+\Delta} r_\eta\big(Z_t, A_t\big)\, \mathrm{d}t + \sum_{m \in \mathcal{M}_{\text{obs}}} \rho_\eta^{(m)}\big(Z_{T_k}\big) \mathbf{1}\{M_k = m\} - \sum_{a \in \mathcal{M}_{\text{act}}} \beta_a \cdot \text{cost}(a),$$

where r_η and $\rho_\eta^{(m)}$ are learned scalars anchored by soft priors from guidelines, and β_a are the budget weights from actionable tokens. Clinicians provide pairwise preferences $y \in \{i \succ j\}$ over subsets of $\{\omega, \omega^{(k)}\}$. Under a Bradley–Terry model with utilities $U_\eta(\omega) = \mathcal{R}_\eta(\omega)$, the likelihood is

$$\mathbb{P}_\eta(i \succ j) = \sigma\big(U_\eta(\omega^{(i)}) - U_\eta(\omega^{(j)})\big), \qquad \sigma(x) = (1 + e^{-x})^{-1},$$

and the alignment loss aggregates the log loss with a KL prior toward guideline-derived utilities U_0,

$$\mathcal{J}_{\text{pref}}(\eta, \theta) = -\sum_{(i,j,y)} \left[y \log \sigma(\Delta U_{ij}) + (1 - y) \log(1 - \sigma(\Delta U_{ij})) \right] + \lambda_{\text{prior}}\, \mathbb{E}\big[(U_\eta - U_0)^2\big],$$

where $\Delta U_{ij} = U_\eta(\omega^{(i)}) - U_\eta(\omega^{(j)})$. Crucially, U_η is computed from latent states Z_t, so gradients flow into the mechanism f_θ, g_θ, nudging representation toward decision-relevant structure. The full objective augments $\mathcal{J}$ by $\lambda_{\text{pref}} \mathcal{J}_{\text{pref}}$.

Counterfactual simulation and feasibility

Counterfactuals are produced by intervening on the control channel while holding exogenous noise fixed. Given a seed u that fixes $W_{[t,t+\Delta]}$ and $\xi_{[t,t+\Delta]}$, a feasible alternative schedule $\tilde{A}$ must satisfy institutional and physiological constraints. Feasibility is encoded as a differentiable indicator $\chi_{\text{feas}}(\tilde{A} \mid \mathcal{C}_t, Z_t) \in [0, 1]$ computed by a constraint network trained on historical adherence and explicit rules, including dose bounds, mutual exclusivity, and minimal spacing. Sampling proceeds from a diffusion model over inter-arrival times in the action space. Let $\delta_\ell = \log(\tau_{\ell+1} - \tau_\ell)$ be log inter-arrivals; a VP-SDE $\mathrm{d}\delta = -\frac{1}{2}\beta(s)\delta\, \mathrm{d}s + \sqrt{\beta(s)}\, \mathrm{d}B_s$ on pseudo-time $s \in [0, 1]$ with state conditioning through Z_t defines a score model $s_\varphi(\delta, s \mid Z_t, \mathcal{C}_t)$. Denoising score matching with conditional NLL upper bounds trains s_φ. At generation time, a predictor–corrector sampler proposes $\{\delta_\ell\}$, which are accepted with probability χ_{feas} and decoded into $\tilde{A}$ via Γ_θ. This diffusion head is light-weight and shares embeddings with ϕ_{act}.

Do-calibrated evaluation heads

Downstream prediction heads are trained during pretraining to respect interventions. For a task-specific label L observed at time t, the head outputs $p_\omega(L \mid Z_t)$ and is trained on importance-weighted samples that simulate uniform action policies in a local neighborhood. Concretely, for each training window we draw m synthetic schedules $\tilde{A}^{(j)}$ from a reference policy π_{ref}, roll out latent states under fixed noise, and augment the minibatch with pairs $(Z_{t+\Delta}^{(j)}, L)$ weighted by $\frac{\pi_{\text{ref}}(\tilde{A}^{(j)})}{\widehat{\pi}_\phi(\tilde{A}^{(j)} \mid \mathcal{H}_t, \mathcal{C}_t)}$. This procedure regularizes the head toward causal stability without assuming full identifiability.

Regularization and identifiability aids

To discourage entanglement of physiology with policy, two additional priors are imposed [76, 77, 78]. A sparsity prior on the control Jacobian G_θ uses group lasso across action channels, encouraging only clinically plausible actions to influence specific latent coordinates. A slow manifold prior constrains Σ_θ to be low-rank with a small number of fast coordinates; this reflects the intuition that many labs evolve on slower time scales than interventions and reduces overfitting to observation cadence. Both are implemented as penalties on the Frobenius norm of G_θ groups and the nuclear norm of Σ_θ.

Optimization

Training alternates between updating the recognition–mechanism parameters (ψ, θ), the policy surrogate ϕ, the diffusion head φ, and the preference head η. Within each minibatch, latent paths are sampled from q_ψ, the reconstruction and counterfactual terms are evaluated with shared noise seeds, and the MPP log-likelihood is computed via thinning-based Monte Carlo estimates for the integral terms [79, 80]. Gradient estimates use pathwise reparameterization for continuous emissions and Gumbel–Softmax relaxations for discrete marks. The invariance gradient penalty is computed with stop-gradient copies to avoid trivial collapse. All objectives are balanced by scalar multipliers that are tuned by a meta-objective: minimize the variance of task-head performance across environments on a held-out validation slice, which operationalizes transferability.

Federated and privacy-preserving training

When training across institutions, each site e optimizes a local objective $\mathcal{J}^{(e)}$ defined on its data with its environment label, and participates in secure aggregation. Mechanism parameters θ and shared encoders are averaged via FedAdam [**?** 81, 82]; policy parameters ϕ remain largely site-specific but share a meta-initialization optimized by Reptile to accelerate adaptation. Gaussian noise calibrated for Rényi differential privacy is added to per-site gradient clips for θ, ψ, η; emissions heads for sensitive modalities like notes use output perturbation. Preference gradients derived from clinicians never leave the site; only sufficient statistics of the Bradley–Terry likelihood are securely aggregated when institutions opt in. To prevent homogenization, environment adapters A_e with small capacity (for example FiLM layers) remain local and are not aggregated; the invariance penalty explicitly excludes adapter parameters.

Evaluation-time counterfactuals and what-if queries

Given a learned model, answering $\mathbb{E}[Y_{t+\Delta} \mid \mathrm{do}(A_{[t,t+\Delta)} = \tilde{A})]$ proceeds by drawing Z_t from the recognition posterior, fixing a noise seed, applying $\tilde{A}$ through Γ_θ, integrating the SDE, and decoding via g_θ. Uncertainty decomposes into epistemic and aleatoric parts: epistemic from parameter and latent uncertainty, approximated by Monte Carlo over q_ψ and ensembles of θ; aleatoric from diffusion Σ_θ and emission noise. Credible intervals are reported conditional on feasibility $\chi_{\text{feas}} = 1$. For policy evaluation under constraints, a rollout with the learned value head V_η ranks feasible alternatives, while importance weights derived from $\widehat{\pi}_\phi$ correct for remaining covariate shift.

Implementation details and hyperparameters

Time is measured in hours and rescaled to unit variance per service line. The latent dimension d is selected by maximizing the variational ELBO subject to an AIC-style penalty on Σ_θ's rank. Log-signature depth for Ψ is set to 3 by default to balance stability and capacity. Intensities λ share a transformer-style history encoder over Ψ_t with causal masking; integrals are estimated with randomized Riemann sums that share quadrature nodes across marks to reduce variance. The diffusion head uses a cosine noise schedule $\beta(s)$ with 32 denoising steps during training and 8 at generation. Stabilized weights w_k are clipped to the 99th percentile per mark. The invariance multiplier λ_{inv} is annealed from 0 to its target value over the first third of training to avoid early over-constraint. All models are trained with mixed precision and gradient norm clipping to 1.0 for numerical stability in adjoint SDE backpropagation.

4. RESULTS AND DISCUSSION

4.1. Results

The empirical evaluation stresses transfer across institutions and epochs, respect for interventions in representation learning, and explicit modeling of the observation policy. Models are trained on a multi-site EHR corpus with environment labels defined by hospital, service line, and time epoch. Primary outcomes include in-hospital mortality within 48 h, AKI onset within 24 h, and ICU transfer within 12 h [83]. Secondary outcomes include length-of-stay buckets and decompensation within 6 h. Counterfactual evaluation uses clinician-defined feasible action sets for vasopressor titration, fluid bolus administration, and

	IID Validation				OOD Test (Held-out Hospital)			
Model	AUROC	AUPRC	Brier	ECE	AUROC	AUPRC	Brier	ECE
Event-Transformer	0.872,(.865,.879)	0.401,(.388,.413)	0.121,(.118,.124)	0.047,(.043,.051)	0.804,(.796,.812)	0.324,(.311,.338)	0.148,(.145,.152)	0.091,(.085,.097)
GRU-D	0.861,(.854,.869)	0.386,(.372,.398)	0.127,(.124,.131)	0.058,(.054,.062)	0.789,(.780,.797)	0.307,(.294,.320)	0.154,(.150,.158)	0.106,(.100,.112)
Neural CDE (no policy)	0.879,(.872,.886)	0.415,(.401,.429)	0.118,(.115,.121)	0.043,(.039,.047)	0.823,(.815,.831)	0.345,(.331,.358)	0.143,(.139,.146)	0.082,(.077,.087)
Masked LM (code–value–time)	0.847,(.839,.855)	0.362,(.349,.375)	0.133,(.130,.136)	0.064,(.060,.068)	0.775,(.767,.784)	0.291,(.278,.304)	0.160,(.156,.164)	0.118,(.112,.125)
CLIO	**0.905**,(.899,.911)	**0.462**,(.447,.476)	**0.108**,(.105,.111)	**0.031**,(.028,.034)	**0.861**,(.854,.868)	**0.402**,(.388,.417)	**0.129**,(.126,.132)	**0.054**,(.049,.059)

Table 1. Aggregate predictive performance across primary outcomes. Means with 95% CIs. OOD holds out an entire hospital with different test-ordering policies.

	CF-C ↑	Policy Regret ↓	$\sqrt{\text{PEHE}}$ ↓	Uplift AUC ↑	NLL (CF) ↓	ECE (CF) ↓
Event-Transformer	0.612,(.598,.626)	0.143,(.137,.149)	0.189,(.182,.197)	0.571,(.557,.585)	0.934,(.918,.950)	0.087,(.082,.092)
Neural CDE (no policy)	0.648,(.634,.662)	0.129,(.123,.135)	0.175,(.168,.182)	0.596,(.582,.610)	0.902,(.887,.918)	0.075,(.070,.080)
CLIO	**0.711**,(.699,.724)	**0.101**,(.096,.106)	**0.153**,(.147,.160)	**0.633**,(.620,.647)	**0.861**,(.846,.876)	**0.059**,(.054,.064)

Table 2. Counterfactual evaluation across feasible action sets for fluids, vasopressors, and early antibiotics. CF-C measures concordance between predicted rankings and realized outcomes across matched cohorts. Regret is normalized to [0,1] against clinician-preferred actions.

antibiotic initiation. All metrics are reported as mean ± 95% bootstrap CI across encounters and sites; out-of-domain (OOD) evaluations hold out entire hospitals or epochs.

The model is compared against a transformer over event tokens with time-deltas, a neural CDE without policy modeling, a GRU-D variant with missingness indicators, and a masked language model over code–value–time tuples. CLIO variants isolate the contribution of policy-aware masking, invariance penalties, counterfactual contrastive alignment, and the diffusion head for feasible action schedules.

The aggregate predictive performance shows consistent gains in both IID and OOD settings. Improvements in AUROC translate into lower Brier scores and substantially improved calibration under shift. Gains are largest when the test distribution alters test-ordering policies or treatment templates, suggesting that explicit modeling of the observation and action processes prevents the conflation of workflow with physiology.

$$\text{AUROC} \uparrow, \quad \text{AUPRC} \uparrow, \quad \text{Brier} \downarrow, \quad \text{ECE (20-bin)} \downarrow$$

Cross-site generalization benefits from decomposing physiology from policy. The largest deltas appear in services with aggressive panel ordering, where counterfactual contrastive alignment discourages learning shortcuts keyed on co-ordered labs. When holding out an epoch after a formulary change that substitutes one vasopressor for another, CLIO's AUROC drops by 0.021 absolute compared to 0.067 for the event-transformer, consistent with environment-wise invariance of the latent dynamics.

Counterfactual reasoning is assessed by a concordance metric over matched windows in which clinician actions vary due to resource constraints. A counterfactual concordance index (CF-C) compares the ranking of feasible alternatives according to predicted outcomes against realized outcomes under observed interventions across matched cohorts. Estimation of individual treatment effects is quantified by $\sqrt{\epsilon_{\text{PEHE}}}$ using semi-synthetic labels derived from stylized pharmacodynamic simulations anchored to the learned latent. Lower policy regret at decision time confirms that preference-aligned value functionals internalize clinical trade-offs beyond raw risk.

Calibration remains stable under shifts that primarily affect testing frequency. Expected calibration error increases only modestly for CLIO, whereas baselines miscalibrate sharply when order frequencies differ. Reliability diagrams stratified by environment show that inverse-propensity reweighting during masked modeling and explicit observation-process likelihood jointly prevent overconfidence in under-ordered regions of the feature space. These effects translate into fewer low-confidence abstentions at decision time while keeping selective risk bounded.

The observation-policy surrogate achieves high marked point-process likelihoods and captures site-specific rhythms without leaking them into the physiology. Time-rescaled residual analyses for inter-arrival distributions are uniform and Kolmogorov–Smirnov distances are minimal. Crucially, when the surrogate is used to sample alternative observation schedules, reconstruction error remains small and unbiased, indicating that the model internalizes missingness-not-at-random as informative rather than as noise.

Ablations confirm that each component materially contributes to transfer and counterfactual fidelity. Removing policy-aware masking reduces OOD AUPRC by 0.041 absolute and degrades calibration most in services with sparse lab ordering.

Model	MPP NLL ↓	Time-Rescale KS ↓	Mark F1 ↑	Recon. RMSE ↓	Bias @ Low-Order ↓	Var. @ Low-Order ↓
GRU-D	1.742,(1.721,1.763)	0.118,(.110,.125)	0.41,(.39,.43)	0.327,(.319,.335)	0.087,(.081,.093)	0.142,(.136,.149)
Neural CDE (no policy)	1.611,(1.592,1.630)	0.094,(.088,.100)	0.47,(.45,.49)	0.301,(.294,.308)	0.072,(.067,.077)	0.131,(.126,.137)
CLIO	**1.438**,(1.420,1.456)	**0.061**,(.056,.066)	**0.58**,(.56,.60)	**0.269**,(.262,.275)	**0.043**,(.039,.047)	**0.112**,(.107,.117)

Table 3. Observation-policy modeling and its downstream effect on reconstruction under sampled schedules. Bias and variance are computed in strata with the lowest tercile of order propensities.

Variant	AUROC (OOD)	AUPRC (OOD)	Brier (OOD)	ECE (OOD)	CF-C	Regret	Invalid Seq. %	Risk Var.
CLIO	**0.861**	**0.402**	**0.129**	**0.054**	**0.711**	**0.101**	**1.8**	**0.006**
– policy-aware masking	0.836	0.361	0.141	0.077	0.684	0.113	1.9	0.009
– invariance penalty	0.829	0.354	0.144	0.083	0.679	0.116	1.9	0.014
– cf-contrastive	0.842	0.368	0.139	0.069	0.652	0.128	1.8	0.010
– diffusion head	0.850	0.379	0.135	0.061	0.667	0.121	6.7	0.007
ODE instead of SDE	0.847	0.372	0.137	0.064	0.671	0.119	1.8	0.008
No sparsity on G_θ	0.852	0.381	0.134	0.060	0.689	0.110	1.9	0.008

Table 4. Ablation study on the held-out hospital. Invalid sequences violate dose or spacing constraints. Risk Var. is the variance of per-environment task risk on validation.

Dropping the invariance penalty increases the variance of per-environment risk gradients and inflates inter-environment performance dispersion. Excluding the counterfactual contrastive term harms CF-concordance and raises policy regret, evidencing that pushing embeddings to respect intervention effects is essential. Eliminating the diffusion head for feasible action schedules reduces CF metrics and increases invalid sequence rates, underscoring the benefit of explicit feasibility modeling.

Preference alignment yields consistent reductions in policy regret and improves the correlation between learned utilities and clinician choices. Pairwise preference accuracy exceeds 0.78 across services after 1–2 rounds of feedback, and the value functional meaningfully reorders ties among clinically comparable action sequences. Importantly, improvements persist in OOD environments, supporting the claim that preferences are distilled into the latent representation rather than overfitting to local workflows.

Stability and safety analyses indicate improved selective prediction and monotonic response to titratable agents. Selective risk under abstention at 30% coverage drops from 0.118 for the event-transformer to 0.094 for CLIO in the OOD hospital. Dose–response monotonicity measured by isotonic regression residuals shows a smaller violation area for vasopressor mean arterial pressure targets, reflecting that the latent state is more aligned with mechanistic response.

Federated training introduces manageable overhead. With three institutions, secure aggregation extends wall-clock per epoch by 9–12% depending on network conditions, while Rényi-DP noise calibrated for $\epsilon \in [4, 6]$ reduces IID AUROC by 0.004–0.009 and OOD AUROC by 0.006–0.012. Environment adapters prevent homogenization, and the invariance term continues to reduce per-environment gradient dispersion even when local adapters are present, suggesting complementary roles.

Robustness to missingness manipulations is strong. When randomly dropping 40% of orders in the validation distribution, AUROC degrades by 0.006 for CLIO compared to 0.028 for the event-transformer. When censoring orders according to a biased policy that removes low-yield tests in the bottom risk decile, calibration for CLIO remains stable whereas baselines become under-confident in high-risk strata, which is consistent with policy-aware reweighting during pretraining.

Fairness and subgroup performance are monitored across age, sex, and comorbidity burden. Absolute gaps in AUPRC under OOD shift remain under 0.022 for CLIO, compared to 0.047–0.063 for baselines, with the largest improvements observed in high-comorbidity cohorts where observation cadences are most irregular. These reductions coincide with lower variance of the observation surrogate's residuals across strata, indicating that treating missingness as signal has equitable benefits.

Sensitivity analyses on hyperparameters find that the invariance multiplier exhibits a broad plateau; values in the range [0.5, 1.5] produce near-identical OOD performance, while setting it to zero consistently raises per-environment risk dispersion. The temperature in counterfactual contrastive training shows a sweet spot near 0.1; higher temperatures admit more hard negatives but slightly destabilize training, whereas lower temperatures reduce gradient signal. Clipping stabilized importance weights at the 99th percentile is essential; removing clipping increases variance and harms both calibration and AUROC.

Altogether, these results indicate that a representation built to be policy-aware and mechanism-seeking yields more transferable predictions, more faithful counterfactuals, and better-calibrated risk under real-world operational shifts. The observation-process surrogate acts as a lens that renders missingness informative, the invariance penalty suppresses environment-specific

Alignment Setting	Pref. Acc. ↑	Kendall τ ↑	Regret ↓	AUROC (OOD) ↑	CF-C ↑	ECE (OOD) ↓
No alignment	0.50	0.00	0.132	0.846	0.668	0.070
Guideline priors only	0.64	0.21	0.119	0.853	0.689	0.063
Preferences + priors (CLIO)	**0.78**	**0.43**	**0.101**	**0.861**	**0.711**	**0.054**

Table 5. Effect of preference alignment on decision quality and generalization. Kendall τ computed between model utilities and clinician choices over feasible sets.

Setting	Comm. / epoch (GB)	Time / epoch (min)	AUROC (IID)	AUROC (OOD)
Centralized, no DP	3.6,(3.5,3.7)	41,(39,43)	0.908,(.903,.913)	0.864,(.858,.871)
Federated, no DP	5.1,(4.9,5.3)	45,(43,48)	0.905,(.900,.911)	0.861,(.854,.868)
Federated, DP ϵ=6	5.2,(5.0,5.4)	46,(44,49)	0.900,(.894,.906)	0.856,(.849,.863)
Federated, DP ϵ=4	5.2,(5.1,5.4)	47,(45,50)	0.896,(.889,.902)	0.852,(.845,.859)

Table 6. Federated training overhead and differential privacy trade-offs. Means with 95% CIs across three sites.

shortcuts, and counterfactual contrastive alignment shapes the latent space around intervention effects. The combined effect is particularly pronounced in settings where workflow evolves faster than biology, which is precisely the regime where conventional EHR models struggle most.

4.2. Discussion

This work reframes EHR pretraining as learning a mechanism and a policy in tandem, rather than treating the record as an incidental sequence to autocomplete. The central contribution is a policy-aware, mechanism-seeking foundation model—CLIO—that separates physiology from workflow and then recouples them in a way that supports prediction, simulation, and decision support. The key design choice is to model patient state as a controlled continuous-time process and to treat both observations and interventions as outputs of a marked point process driven by clinician belief and institutional constraints. By assigning this process its own likelihood and using its intensities as propensities, reconstruction is trained against a counterfactual baseline rather than the idiosyncrasies of local ordering habits, and missingness becomes informative rather than a nuisance.

A second contribution is to bake counterfactual reasoning into pretraining rather than bolt it on downstream. Fixing exogenous noise while intervening on the control channel yields paired rollouts that isolate the effect of actions; an explicit contrastive objective then pushes the representation to preserve these effects while discounting correlations induced by co-ordering or templated care. Because feasible schedules are sampled by a lightweight diffusion model over action inter-arrival times and filtered through differentiable constraints, counterfactuals respect dose bounds, spacing, and mutual exclusivity. This coupling of a simulator-backed latent dynamics with feasibility-aware action generation stands in sharp contrast to token-only approaches that hallucinate clinically impossible sequences and thus degrade when asked "what if."

The third contribution is environmental invariance at the level where it matters. Instead of matching raw feature distributions or relying on domain-adversarial tricks that often erase clinical signal, CLIO penalizes cross-environment variation in the gradients that train the latent dynamics and aligns path marginals only after conditioning on comparable context. This choice operationalizes the claim that biology should be stable while policies are free to vary. The empirical signature—improved OOD discrimination and calibration precisely when test-ordering policies or formularies change—supports the premise that mechanism and policy were successfully disentangled for transfer.

Alignment is elevated from an afterthought to a training signal that shapes the latent space toward clinical utility. Rather than collecting free-form textual preferences, we ask clinicians to compare short, de-identified trajectory rollouts generated by the model's own simulator. A simple Bradley–Terry head turns those pairwise preferences, softened by guideline priors and resource costs, into a value functional $V(Z_t, a)$ that backpropagates to the mechanism. The effect is to resolve ties among action sequences that are risk-equivalent but preference-distinct, and to reduce policy regret under feasible alternatives. Because preference gradients stay on-prem and only sufficient statistics are aggregated, alignment integrates naturally with the federated scheme.

The final system-level contribution is a training and deployment recipe that acknowledges privacy and heterogeneity as first-class constraints. Mechanism parameters are aggregated with DP noise; site-specific adapters and policy surrogates remain local; and the invariance penalty explicitly excludes adapter parameters to avoid homogenizing environments that supply the very diversity needed to learn invariants. The result is a model that can be trained across institutions without erasing what makes them different, while still producing a physiology-centered latent that travels.

There are limits to what can be identified from observational logs, even with a faithful policy surrogate. CLIO does not claim full causal identification of treatment effects; rather, it targets counterfactually stable representations and decision-support heads that are less sensitive to policy-induced shift. The diffusion-based action generator, while fast and constraint-aware, inherits any misspecification in the feasibility network; robust optimization and stronger mechanistic priors over pharmacodynamics are natural extensions. Preference learning remains sample-limited and subject to inter-annotator variability; active query selection and calibration of the value head to long-term outcomes are promising next steps. Finally, continuous-time training with SDE adjoints is computationally heavy; low-rank diffusion and control-sparse Jacobians partly mitigate this, but practical deployment will benefit from distilled surrogates for real-time use.

Taken together, the contributions amount to a shift in what "foundation" means for EHR. By endowing the model with an explicit observation policy, a counterfactual training signal that respects interventions and feasibility, an invariance criterion that protects physiology from workflow, and a preference layer that encodes clinical trade-offs, CLIO offers a substrate that is not just better at next-event prediction, but better aligned with the questions clinicians and health systems actually ask. The gains in transfer, calibration under policy change, and fidelity of what-if rollouts suggest that treating the EHR as the execution trace of a controlled system is not merely a philosophical stance; it is a practical route to representations that remain useful when the world outside the training set moves.

References

[1] Alexander Hoerbst and Elske Ammenwerth, "Electronic health records," *Methods of information in medicine*, vol. 49, no. 04, pp. 320–336, 2010.

[2] R Scott Evans, "Electronic health records: then, now, and in the future," *Yearbook of medical informatics*, vol. 25, no. S 01, pp. S48–S61, 2016.

[3] Shengpu Tang, Aditya Modi, Michael Sjoding, and Jenna Wiens, "Clinician-in-the-loop decision making: Reinforcement learning with near-optimal set-valued policies," in *International Conference on Machine Learning*. PMLR, 2020, pp. 9387–9396.

[4] Jaya Ojha, Oriana Presacan, Pedro G. Lind, Eric Monteiro, and Anis Yazidi, "Navigating uncertainty: A user-perspective survey of trustworthiness of ai in healthcare," *ACM Transactions on Computing for Healthcare*, vol. 6, no. 3, pp. 1–32, 2025.

[5] Adrian Lindenmeyer, Malte Blattmann, Stefan Franke, Thomas Neumuth, and Daniel Schneider, "Towards trustworthy ai in healthcare: Epistemic uncertainty estimation for clinical decision support," *Journal of Personalized Medicine*, vol. 15, no. 2, pp. 58, 2025.

[6] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, Hao Peng, Jianxin Li, Jia Wu, Ziwei Liu, Pengtao Xie, Caiming Xiong, Jian Pei, Philip S. Yu, and Lichao Sun, "A comprehensive survey on pretrained foundation models: A history from bert to chatgpt," 2023.

[7] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak,

Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler,

Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma, "The llama 3 herd of models," 2024.

[8] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornraphop Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilaï Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Sheshan, Raia Hadsell Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Sharon Silver, Ayzaan Wahid, Sergey Brin, Yves Raimond, Klemen Kloboves, Cindy Wang, Nitesh Bharadwaj Gundavarapu, Ilia Shumailov, Bo Wang, Mantas Pajarskas, Joe Heyward, Martin Nikoltchev, Maciej Kula, Hao Zhou, Zachary Garrett, Sushant Kafle, Sercan Arik, Ankita Goel, Mingyao Yang, Jiho Park, Koji Kojima, Parsa Mahmoudieh, Koray Kavukcuoglu, Grace Chen, Doug Fritz, Anton Bulyenov, Sudeshna Roy, Dimitris Paparas, Hadar Shemtov, Bo-Juen Chen, Robin Strudel, David Reitter, Aurko Roy, Andrey Vlasov, Changwan Ryu, Chas Leichner, Haichuan Yang, Zelda Mariet, Denis Vnukov, Tim Sohn, Amy Stuart, Wei Liang, Minmin Chen, Praynaa Rawlani, Christy Koh, JD Co-Reyes, Guangda Lai, Praseem Banzal, Dimitrios Vytiniotis, Jieru Mei, Mu Cai, Mohammed Badawi, Corey Fry, Ale Hartman, Daniel Zheng, Eric Jia, James Keeling, Annie Louis, Ying Chen, Efren Robles, Wei-Chih Hung, Howard Zhou, Nikita Saxena, Sonam Goenka, Olivia Ma, Zach Fisher, Mor Hazan Taege, Emily Graves, David Steiner, Yujia Li, Sarah Nguyen, Rahul Sukthankar, Joe Stanton, Ali Eslami, Gloria Shen, Berkin Akin, Alexey Guseynov, Yiqian Zhou, Jean-Baptiste Alayrac, Armand Joulin, Efrat Farkash, Ashish Thapliyal, Stephen Roller, Noam Shazeer, Todor Davchev, Terry Koo, Hannah Forbes-Pollard, Kartik Audhkhasi, Greg Farquhar, Adi Mayrav Gilady, Maggie Song, John Aslanides, Piermaria Mendolicchio, Alicia Parrish, John Blitzer, Pramod Gupta, Xiaoen Ju, Xiaochen Yang, Puranjay Datta, Andrea Tacchetti, Sanket Vaibhav Mehta, Gregory Dibb, Shubham Gupta, Federico Piccinini, Raia Hadsell, Sujee Rajayogam, Jiepu Jiang, Patrick Griffin, Patrik Sundberg, Jamie Hayes, Alexey Frolov, Tian Xie, Adam Zhang, Kingshuk Dasgupta, Uday Kalra, Lior Shani, Klaus Macherey, Tzu-Kuo Huang, Liam MacDermed, Karthik Duddu, Paulo Zacchello, Zi Yang, Jessica Lo, Kai Hui, Matej Kastelic, Derek Gasaway, Qijun Tan, Summer Yue, Pablo Barrio, John Wieting, Weel Yang, Andrew Nystrom, Solomon Demmessie, Anselm Levskaya, Fabio Viola, Chetan Tekur, Greg Billock, George Necula, Mandar Joshi, Rylan Schaeffer, Swachhand Lokhande, Christina Sorokin, Pradeep Shenoy, Mia Chen, Mark Collier, Hongji Li, Taylor Bos, Nevan Wichers, Sun Jae Lee, Angéline Pouget, Santhosh Thangaraj, Kyriakos Axiotis, Phil Crone, Rachel Sterneck, Nikolai Chinaev, Victoria Krakovna, Oleksandr Ferludin, Ian Gemp, Stephanie Winkler, Dan Goldberg, Ivan Korotkov, Kefan Xiao, Malika Mehrotra, Sandeep Mariserla, Vihari Piratla, Terry Thurk, Khiem Pham, Hongxu Ma, Alexandre Senges, Ravi Kumar, Clemens Meyer, Ellie Talius, Nuo Wang Pierse, Ballie Sandhu, Horia Toma, Kuo Lin, Swaroop Nath, Tom Stone, Dorsa Sadigh, Nikita Gupta, Arthur Guez, Avi Singh, Matt Thomas, Tom Duerig, Yuan Gong, Richard Tanburn, Lydia Lihui Zhang, Phuong Dao, Mohamed Hammad, Sirui Xie, Shruti Rijhwani, Ben Murdoch, Duhyeon Kim, Will Thompson, Heng-Tze Cheng, Daniel Sohn, Pablo Sprechmann, Qiantong Xu, Srinivas Tadepalli, Peter Young, Ye Zhang, Hansa Srinivasan, Miranda Aperghis, Aditya Ayyar, Hen Fitoussi, Ryan Burnell, David Madras, Mike Dusenberry, Xi Xiong, Tayo Oguntebi, Ben Albrecht, Jörg Bornschein, Jovana Mitrović, Mason Dimarco, Bhargav Kanagal Shamanna, Premal Shah, Eren Sezener, Shyam Upadhyay, Dave Lacey, Craig Schiff, Sebastien Baur, Sanjay Ganapathy, Eva Schnider, Mateo Wirth, Connor Schenck, Andrey Simanovsky, Yi-Xuan Tan, Philipp Fränken, Dennis Duan, Bharath Mankalale, Nikhil Dhawan, Kevin Sequeira, Zichuan Wei, Shivanker Goel, Caglar Unlu, Yukun Zhu, Haitian Sun, Ananth Balashankar, Kurt Shuster, Megh Umekar, Mahmoud Alnahlawi, Aäron van den Oord, Kelly Chen, Yuexiang Zhai, Zihang Dai, Kuang-Huei Lee, Eric Doi, Lukas Zilka, Rohith Vallu, Disha Shrivastava, Jason Lee, Hisham Husain, Honglei Zhuang, Vincent Cohen-Addad, Jarred Barber, James Atwood, Adam Sadovsky, Quentin Wellens, Steven Hand, Arunkumar Rajendran, Aybuke Turker, CJ Carey, Yuanzhong Xu, Hagen Soltau, Zefei Li, Xinying Song, Conglong Li, Iurii Kemaev, Sasha Brown, Andrea Burns, Viorica Patraucean, Piotr Stanczyk, Renga Aravamudhan, Mathieu Blondel, Hila Noga, Lorenzo Blanco, Will Song, Michael Isard, Mandar Sharma, Reid Hayes, Dalia El Badawy, Avery Lamp, Itay Laish, Olga Kozlova, Kelvin Chan, Sahil Singla, Srinivas Sunkara, Mayank Upadhyay, Chang Liu, Aijun Bai, Jarek Wilkiewicz, Martin Zlocha, Jeremiah Liu, Zhuowan Li, Haiguang Li, Omer Barak, Ganna Raboshchuk, Jiho Choi, Fangyu Liu, Erik Jue, Mohit Sharma, Andreea Marzoca, Robert Busa-Fekete, Anna Korsun, Andre Elisseeff,

Zhe Shen, Sara Mc Carthy, Kay Lamerigts, Anahita Hosseini, Hanzhao Lin, Charlie Chen, Fan Yang, Kushal Chauhan, Mark Omernick, Dawei Jia, Karina Zainullina, Demis Hassabis, Danny Vainstein, Ehsan Amid, Xiang Zhou, Ronny Votel, Eszter Vértes, Xinjian Li, Zongwei Zhou, Angeliki Lazaridou, Brendan McMahan, Arjun Narayanan, Hubert Soyer, Sujoy Basu, Kayi Lee, Bryan Perozzi, Qin Cao, Leonard Berrada, Rahul Arya, Ke Chen, Katrina, Xu, Matthias Lochbrunner, Alex Hofer, Sahand Sharifzadeh, Renjie Wu, Sally Goldman, Pranjal Awasthi, Xuezhi Wang, Yan Wu, Claire Sha, Biao Zhang, Maciej Mikuła, Filippo Graziano, Siobhan Mcloughlin, Irene Giannoumis, Youhei Namiki, Chase Malik, Carey Radebaugh, Jamie Hall, Ramiro Leal-Cavazos, Jianmin Chen, Vikas Sindhwani, David Kao, David Greene, Jordan Griffith, Chris Welty, Ceslee Montgomery, Toshihiro Yoshino, Liangzhe Yuan, Noah Goodman, Assaf Hurwitz Michaely, Kevin Lee, KP Sawhney, Wei Chen, Zheng Zheng, Megan Shum, Nikolay Savinov, Etienne Pot, Alex Pak, Morteza Zadimoghaddam, Sijal Bhatnagar, Yoad Lewenberg, Blair Kutzman, Ji Liu, Lesley Katzen, Jeremy Selier, Josip Djolonga, Dmitry Lepikhin, Kelvin Xu, Jacky Liang, Jiewen Tan, Benoit Schillings, Muge Ersoy, Pete Blois, Bernd Bandemer, Abhimanyu Singh, Sergei Lebedev, Pankaj Joshi, Adam R. Brown, Evan Palmer, Shreya Pathak, Komal Jalan, Fedir Zubach, Shuba Lall, Randall Parker, Alok Gunjan, Sergey Rogulenko, Sumit Sanghai, Zhaoqi Leng, Zoltan Egyed, Shixin Li, Maria Ivanova, Kostas Andriopoulos, Jin Xie, Elan Rosenfeld, Auriel Wright, Ankur Sharma, Xinyang Geng, Yicheng Wang, Sam Kwei, Renke Pan, Yujing Zhang, Gabby Wang, Xi Liu, Chak Yeung, Elizabeth Cole, Aviv Rosenberg, Zhen Yang, Phil Chen, George Polovets, Pranav Nair, Rohun Saxena, Josh Smith, Shuo yiin Chang, Aroma Mahendru, Svetlana Grant, Anand Iyer, Irene Cai, Jed McGiffin, Jiaming Shen, Alanna Walton, Antonious Girgis, Oliver Woodman, Rosemary Ke, Mike Kwong, Louis Rouillard, Jinmeng Rao, Zhihao Li, Yuntao Xu, Flavien Prost, Chi Zou, Ziwei Ji, Alberto Magni, Tyler Liechty, Dan A. Calian, Deepak Ramachandran, Igor Krivokon, Hui Huang, Terry Chen, Anja Hauth, Anastasija Ilić, Weijuan Xi, Hyeontaek Lim, Vlad-Doru Ion, Pooya Moradi, Metin Toksoz-Exley, Kalesha Bullard, Miltos Allamanis, Xiaomeng Yang, Sophie Wang, Zhi Hong, Anita Gergely, Cheng Li, Bhavishya Mittal, Vitaly Kovalev, Victor Ungureanu, Jane Labanowski, Jan Wassenberg, Nicolas Lacasse, Geoffrey Cideron, Petar Dević, Annie Marsden, Lynn Nguyen, Michael Fink, Yin Zhong, Tatsuya Kiyono, Desi Ivanov, Sally Ma, Max Bain, Kiran Yalasangi, Jennifer She, Anastasia Petrushkina, Mayank Lunayach, Carla Bromberg, Sarah Hodkinson, Vilobh Meshram, Daniel Vlasic, Austin Kyker, Steve Xu, Jeff Stanway, Zuguang Yang, Kai Zhao, Matthew Tung, Seth Odoom, Yasuhisa Fujii, Justin Gilmer, Eunyoung Kim, Felix Halim, Quoc Le, Bernd Bohnet, Seliem El-Sayed, Behnam Neyshabur, Malcolm Reynolds, Dean Reich, Yang Xu, Erica Moreira, Anuj Sharma, Zeyu Liu, Mohammad Javad Hosseini, Naina Raisinghani, Yi Su, Ni Lao, Daniel Formoso, Marco Gelmi, Almog Gueta, Tapomay Dey, Elena Gribovskaya, Domagoj Ćevid, Sidharth Mudgal, Garrett Bingham, Jianling Wang, Anurag Kumar, Alex Cullum, Feng Han, Konstantinos Bousmalis, Diego Cedillo, Grace Chu, Vladimir Magay, Paul Michel, Ester Hlavnova, Daniele Calandriello, Setareh Ariafar, Kaisheng Yao, Vikash Sehwag, Arpi Vezer, Agustin Dal Lago, Zhenkai Zhu, Paul Kishan Rubenstein, Allen Porter, Anirudh Baddepudi, Oriana Riva, Mihai Dorin Istin, Chih-Kuan Yeh, Zhi Li, Andrew Howard, Nilpa Jha, Jeremy Chen, Raoul de Liedekerke, Zafarali Ahmed, Mikel Rodriguez, Tanuj Bhatia, Bangju Wang, Ali Elqursh, David Klinghoffer, Peter Chen, Pushmeet Kohli, Te I, Weiyang Zhang, Zack Nado, Jilin Chen, Maxwell Chen, George Zhang, Aayush Singh, Adam Hillier, Federico Lebron, Yiqing Tao, Ting Liu, Gabriel Dulac-Arnold, Jingwei Zhang, Shashi Narayan, Buhuang Liu, Orhan Firat, Abhishek Bhowmick, Bingyuan Liu, Hao Zhang, Zizhao Zhang, Georges Rotival, Nathan Howard, Anu Sinha, Alexander Grushetsky, Benjamin Beyret, Keerthana Gopalakrishnan, James Zhao, Kyle He, Szabolcs Payrits, Zaid Nabulsi, Zhaoyi Zhang, Weijie Chen, Edward Lee, Nova Fallen, Sreenivas Gollapudi, Aurick Zhou, Filip Pavetić, Thomas Köppe, Shiyu Huang, Rama Pasumarthi, Nick Fernando, Felix Fischer, Daria Ćurko, Yang Gao, James Svensson, Austin Stone, Haroon Qureshi, Abhishek Sinha, Apoorv Kulshreshtha, Martin Matysiak, Jieming Mao, Carl Saroufim, Aleksandra Faust, Qingnan Duan, Gil Fidel, Kaan Katircioglu, Raphaël Lopez Kaufman, Dhruv Shah, Weize Kong, Abhishek Bapna, Gellért Weisz, Emma Dunleavy, Praneet Dutta, Tianqi Liu, Rahma Chaabouni, Carolina Parada, Marcus Wu, Alexandra Belias, Alessandro Bissacco, Stanislav Fort, Li Xiao, Fantine Huot, Chris Knutsen, Yochai Blau, Gang Li, Jennifer Prendki, Juliette Love, Yinlam Chow, Pichi Charoenpanit, Hidetoshi Shimokawa, Vincent Coriou, Karol Gregor, Tomas Izo, Arjun Akula, Mario Pinto, Chris Hahn, Dominik Paulus, Jiaxian Guo, Neha Sharma, Cho-Jui Hsieh, Adaeze Chukwuka, Kazuma Hashimoto, Nathalie Rauschmayr, Ling Wu, Christof Angermueller, Yulong Wang, Sebastian Gerlach, Michael Pliskin, Daniil Mirylenka, Min Ma, Lexi Baugher, Bryan Gale, Shaan Bijwadia, Nemanja Rakićević, David Wood, Jane Park, Chung-Ching Chang, Babi Seal, Chris Tar, Kacper Krasowiak, Yiwen Song, Georgi Stephanov, Gary Wang, Marcello Maggioni, Stein Xudong Lin, Felix Wu, Shachi Paul, Zixuan Jiang, Shubham Agrawal, Bilal Piot, Alex Feng, Cheolmin Kim, Tulsee Doshi, Jonathan Lai, Chuqiao, Xu, Sharad Vikram, Ciprian Chelba, Sebastian Krause, Vincent Zhuang, Jack Rae, Timo Denk, Adrian Collister, Lotte Weerts, Xianghong Luo, Yifeng Lu, Håvard Garnes, Nitish Gupta, Terry Spitz, Avinatan Hassidim, Lihao Liang, Izhak Shafran, Peter Humphreys, Kenny Vassigh, Phil Wallis, Virat Shejwalkar, Nicolas Perez-Nieves, Rachel Hornung, Melissa Tan, Beka Westberg, Andy Ly, Richard Zhang, Brian Farris, Jongbin Park, Alec Kosik, Zeynep Cankara, Andrii Maksai, Yunhan Xu, Albin Cassirer, Sergi Caelles, Abbas Abdolmaleki, Mencher Chiang, Alex Fabrikant, Shravya

Shetty, Luheng He, Mai Giménez, Hadi Hashemi, Sheena Panthaplackel, Yana Kulizhskaya, Salil Deshmukh, Daniele Pighin, Robin Alazard, Disha Jindal, Seb Noury, Pradeep Kumar S, Siyang Qin, Xerxes Dotiwalla, Stephen Spencer, Mohammad Babaeizadeh, Blake JianHang Chen, Vaibhav Mehta, Jennie Lees, Andrew Leach, Penporn Koanantakool, Ilia Akolzin, Ramona Comanescu, Junwhan Ahn, Alexey Svyatkovskiy, Basil Mustafa, David D'Ambrosio, Shiva Mohan Reddy Garlapati, Pascal Lamblin, Alekh Agarwal, Shuang Song, Pier Giuseppe Sessa, Pauline Coquinot, John Maggs, Hussain Masoom, Divya Pitta, Yaqing Wang, Patrick Morris-Suzuki, Billy Porter, Johnson Jia, Jeffrey Dudek, Raghavender R, Cosmin Paduraru, Alan Ansell, Tolga Bolukbasi, Tony Lu, Ramya Ganeshan, Zi Wang, Henry Griffiths, Rodrigo Benenson, Yifan He, James Swirhun, George Papamakarios, Aditya Chawla, Kuntal Sengupta, Yan Wang, Vedrana Milutinovic, Igor Mordatch, Zhipeng Jia, Jamie Smith, Will Ng, Shitij Nigam, Matt Young, Eugen Vušak, Blake Hechtman, Sheela Goenka, Avital Zipori, Kareem Ayoub, Ashok Popat, Trilok Acharya, Luo Yu, Dawn Bloxwich, Hugo Song, Paul Roit, Haiqiong Li, Aviel Boag, Nigamaa Nayakanti, Bilva Chandra, Tianli Ding, Aahil Mehta, Cath Hope, Jiageng Zhang, Idan Heimlich Shtacher, Kartikeya Badola, Ryo Nakashima, Andrei Sozanschi, Iulia Comşa, Ante Žužul, Emily Caveness, Julian Odell, Matthew Watson, Dario de Cesare, Phillip Lippe, Derek Lockhart, Siddharth Verma, Huizhong Chen, Sean Sun, Lin Zhuo, Aditya Shah, Prakhar Gupta, Alex Muzio, Ning Niu, Amir Zait, Abhinav Singh, Meenu Gaba, Fan Ye, Prajit Ramachandran, Mohammad Saleh, Raluca Ada Popa, Ayush Dubey, Frederick Liu, Sara Javanmardi, Mark Epstein, Ross Hemsley, Richard Green, Nishant Ranka, Eden Cohen, Chuyuan Kelly Fu, Sanjay Ghemawat, Jed Borovik, James Martens, Anthony Chen, Pranav Shyam, André Susano Pinto, Ming-Hsuan Yang, Alexandru Ţifrea, David Du, Boqing Gong, Ayushi Agarwal, Seungyeon Kim, Christian Frank, Saloni Shah, Xiaodan Song, Zhiwei Deng, Ales Mikhalap, Kleopatra Chatziprimou, Timothy Chung, Toni Creswell, Susan Zhang, Yennie Jun, Carl Lebsack, Will Truong, Slavica Andačić, Itay Yona, Marco Fornoni, Rong Rong, Serge Toropov, Afzal Shama Soudagar, Andrew Audibert, Salah Zaiem, Zaheer Abbas, Andrei Rusu, Sahitya Potluri, Shitao Weng, Anastasios Kementsietsidis, Anton Tsitsulin, Daiyi Peng, Natalie Ha, Sanil Jain, Tejasi Latkar, Simeon Ivanov, Cory McLean, Anirudh GP, Rajesh Venkataraman, Canoee Liu, Dilip Krishnan, Joel D'sa, Roey Yogev, Paul Collins, Benjamin Lee, Lewis Ho, Carl Doersch, Gal Yona, Shawn Gao, Felipe Tiengo Ferreira, Adnan Ozturel, Hannah Muckenhirn, Ce Zheng, Gargi Balasubramaniam, Mudit Bansal, George van den Driessche, Sivan Eiger, Salem Haykal, Vedant Misra, Abhimanyu Goyal, Danilo Martins, Gary Leung, Jonas Valfridsson, Four Flynn, Will Bishop, Chenxi Pang, Yoni Halpern, Honglin Yu, Lawrence Moore, Yuvein, Zhu, Sridhar Thiagarajan, Yoel Drori, Zhisheng Xiao, Lucio Dery, Rolf Jagerman, Jing Lu, Eric Ge, Vaibhav Aggarwal, Arjun Khare, Vinh Tran, Oded Elyada, Ferran Alet, James Rubin, Ian Chou, David Tian, Libin Bai, Lawrence Chan, Lukasz Lew, Karolis Misiunas, Taylan Bilal, Aniket Ray, Sindhu Raghuram, Alex Castro-Ros, Viral Carpenter, CJ Zheng, Michael Kilgore, Josef Broder, Emily Xue, Praveen Kallakuri, Dheeru Dua, Nancy Yuen, Steve Chien, John Schultz, Saurabh Agrawal, Reut Tsarfaty, Jingcao Hu, Ajay Kannan, Dror Marcus, Nisarg Kothari, Baochen Sun, Ben Horn, Matko Bošnjak, Ferjad Naeem, Dean Hirsch, Lewis Chiang, Boya Fang, Jie Han, Qifei Wang, Ben Hora, Antoine He, Mario Lučić, Beer Changpinyo, Anshuman Tripathi, John Youssef, Chester Kwak, Philippe Schlattner, Cat Graves, Rémi Leblond, Wenjun Zeng, Anders Andreassen, Gabriel Rasskin, Yue Song, Eddie Cao, Junhyuk Oh, Matt Hoffman, Wojtek Skut, Yichi Zhang, Jon Stritar, Xingyu Cai, Saarthak Khanna, Kathie Wang, Shriya Sharma, Christian Reisswig, Younghoon Jun, Aman Prasad, Tatiana Sholokhova, Preeti Singh, Adi Gerzi Rosenthal, Anian Ruoss, Françoise Beaufays, Sean Kirmani, Dongkai Chen, Johan Schalkwyk, Jonathan Herzig, Been Kim, Josh Jacob, Damien Vincent, Adrian N Reyes, Ivana Balazevic, Léonard Hussenot, Jon Schneider, Parker Barnes, Luis Castro, Spandana Raj Babbula, Simon Green, Serkan Cabi, Nico Duduta, Danny Driess, Rich Galt, Noam Velan, Junjie Wang, Hongyang Jiao, Matthew Mauger, Du Phan, Miteyan Patel, Vlado Galić, Jerry Chang, Eyal Marcus, Matt Harvey, Julian Salazar, Elahe Dabir, Suraj Satishkumar Sheth, Amol Mandhane, Hanie Sedghi, Jeremiah Willcock, Amir Zandieh, Shruthi Prabhakara, Aida Amini, Antoine Miech, Victor Stone, Massimo Nicosia, Paul Niemczyk, Ying Xiao, Lucy Kim, Sławek Kwasiborski, Vikas Verma, Ada Maksutaj Oflazer, Christoph Hirnschall, Peter Sung, Lu Liu, Richard Everett, Michiel Bakker, Ágoston Weisz, Yufei Wang, Vivek Sampathkumar, Uri Shaham, Bibo Xu, Yasemin Altun, Mingqiu Wang, Takaaki Saeki, Guanjie Chen, Emanuel Taropa, Shanthal Vasanth, Sophia Austin, Lu Huang, Goran Petrovic, Qingyun Dou, Daniel Golovin, Grigory Rozhdestvenskiy, Allie Culp, Will Wu, Motoki Sano, Divya Jain, Julia Proskurnia, Sébastien Cevey, Alejandro Cruzado Ruiz, Piyush Patil, Mahdi Mirzazadeh, Eric Ni, Javier Snaider, Lijie Fan, Alexandre Fréchette, AJ Pierigiovanni, Shariq Iqbal, Kenton Lee, Claudio Fantacci, Jinwei Xing, Lisa Wang, Alex Irpan, David Raposo, Yi Luan, Zhuoyuan Chen, Harish Ganapathy, Kevin Hui, Jiazhong Nie, Isabelle Guyon, Heming Ge, Roopali Vij, Hui Zheng, Dayeong Lee, Alfonso Castaño, Khuslen Baatarsukh, Gabriel Ibagon, Alexandra Chronopoulou, Nicholas FitzGerald, Shashank Viswanadha, Safeen Huda, Rivka Moroshko, Georgi Stoyanov, Prateek Kolhar, Alain Vaucher, Ishaan Watts, Adhi Kuncoro, Henryk Michalewski, Satish Kambala, Bat-Orgil Batsaikhan, Alek Andreev, Irina Jurenka, Maigo Le, Qihang Chen, Wael Al Jishi, Sarah Chakera, Zhe Chen, Aditya Kini, Vikas Yadav, Aditya Siddhant, Ilia Labzovsky, Balaji Lakshminarayanan, Carrie Grimes Bostock, Pankil Botadra, Ankesh Anand, Colton Bishop, Sam Conway-Rahman, Mohit Agarwal, Yani Donchev, Achintya Singhal, Félix de Chaumont Quitry, Natalia Ponomareva, Nishant Agrawal, Bin Ni, Kalpesh Krishna, Masha Samsikova, John Karro, Yilun Du, Tamara von Glehn, Caden Lu, Christopher A. Choquette-Choo, Zhen Qin, Tingnan Zhang, Sicheng Li, Divya Tyam, Swaroop Mishra, Wing Lowe, Colin Ji, Weiyi Wang, Manaal Faruqui, Ambrose Slone, Valentin Dalibard, Arunachalam Narayanaswamy, John Lambert, Pierre-Antoine Manzagol, Dan Karliner, Andrew Bolt, Ivan Lobov, Aditya Kusupati, Chang Ye, Xuan Yang, Heiga Zen, Nelson George, Mukul Bhutani, Olivier Lacombe, Robert Riachi, Gagan Bansal, Rachel Soh, Yue Gao, Yang Yu, Adams Yu, Emily Nottage, Tania Rojas-Esponda, James Noraky, Manish Gupta, Ragha Kotikalapudi, Jichuan Chang, Sanja Deur, Dan Graur, Alex Mossin, Erin Farnese, Ricardo Figueira, Alexandre Moufarek, Austin Huang, Patrik Zochbauer, Ben Ingram, Tongzhou Chen, Zelin Wu, Adrià Puigdomènech, Leland Rechis, Da Yu, Sri Gayatri Sundara Padmanabhan, Rui Zhu, Chu ling Ko, Andrea Banino, Samira Daruki, Aarush Selvan, Dhruva Bhaswar, Daniel Hernandez Diaz, Chen Su, Salvatore Scellato, Jennifer Brennan, Woohyun Han, Grace Chung, Priyanka Agrawal, Urvashi Khandelwal, Khe Chai Sim, Morgane Lustman, Sam Ritter, Kelvin Guu, Jiawei Xia, Prateek Jain, Emma Wang, Tyrone Hill, Mirko Rossini, Marija Kostelac, Tautvydas Misiunas, Amit Sabne, Kyuyeun Kim, Ahmet Iscen, Congchao Wang, José Leal, Ashwin Sreevatsa, Utku Evci, Manfred Warmuth, Saket Joshi, Daniel Suo, James Lottes, Garrett Honke, Brendan Jou, Stefani Karp, Jieru Hu, Himanshu Sahni, Adrien Ali Taïga, William Kong, Samrat Ghosh, Renshen Wang, Jay Pavagadhi, Natalie Axelsson, Nikolai Grigorev, Patrick Siegler, Rebecca Lin, Guohui Wang, Emilio Parisotto, Sharath Maddineni, Krishan Subudhi, Eyal Ben-David, Elena Pochernina, Orgad Keller, Thi Avrahami, Zhe Yuan, Pulkit Mehta, Jialu Liu, Sherry Yang, Wendy Kan, Katherine Lee, Tom Funkhouser, Derek Cheng, Hongzhi Shi, Archit Sharma, Joe Kelley, Matan Eyal, Yury Malkov, Corentin Tallec, Yuval Bahat, Shen Yan, Xintian, Wu, David Lindner, Chengda Wu, Avi Caciularu, Xiyang Luo, Rodolphe Jenatton, Tim Zaman, Yingying Bi, Ilya Kornakov, Ganesh Mallya, Daisuke Ikeda, Itay Karo, Anima Singh, Colin Evans, Praneeth Netrapalli, Vincent Nallatamby, Isaac Tian, Yannis Assael, Vikas Raunak, Victor Carbune, Ioana Bica, Lior Madmoni, Dee Cattle, Snchit Grover, Krishna Somandepalli, Sid Lall, Amelio Vázquez-Reina, Riccardo Patana, Jiaqi Mu, Pranav Talluri, Maggie Tran, Rajeev Aggarwal, RJ Skerry-Ryan, Jun Xu, Mike Burrows, Xiaoyue Pan, Edouard Yvinec, Di Lu, Zhiying Zhang, Duc Dung Nguyen, Hairong Mu, Gabriel Barcik, Helen Ran, Lauren Beltrone, Krzysztof Choromanski, Dia Kharrat, Samuel Albanie, Sean Purser-haskell, David Bieber, Carrie Zhang, Jing Wang, Tom Hudson, Zhiyuan Zhang, Han Fu, Johannes Mauerer, Mohammad Hossein Bateni, AJ Maschinot, Bing Wang, Muye Zhu, Arjun Pillai, Tobias Weyand, Shuang Liu, Oscar Akerlund, Fred Bertsch, Vittal Premachandran, Alicia Jin, Vincent Roulet, Peter de Boursac, Shubham Mittal, Ndaba Ndebele, Georgi Karadzhov, Sahra Ghalebikesabi, Ricky Liang, Allen Wu, Yale Cong, Nimesh Ghelani, Sumeet Singh, Bahar Fatemi, Warren, Chen, Charles Kwong, Alexey Kolganov, Steve Li, Richard Song, Chenkai Kuang, Sobhan Miryoosefi, Dale Webster, James Wendt, Arkadiusz Socala, Guolong Su, Artur Mendonça, Abhinav Gupta, Xiaowei Li, Tomy Tsai, Qiong, Hu, Kai Kang, Angie Chen, Sertan Girgin, Yongqin Xian, Andrew Lee, Nolan Ramsden, Leslie Baker, Madeleine Clare Elish, Varvara Krayvanova, Rishabh Joshi, Jiri Simsa, Yao-Yuan Yang, Piotr Ambroszczyk, Dipankar Ghosh, Arjun Kar, Yuan Shangguan, Yumeya Yamamori, Yaroslav Akulov, Andy Brock, Haotian Tang, Siddharth Vashishtha, Rich Munoz, Andreas Steiner, Kalyan Andra, Daniel Eppens, Qixuan Feng, Hayato Kobayashi, Sasha Goldshtein, Mona El Mahdy, Xin Wang, Jilei, Wang, Richard Killam, Tom Kwiatkowski, Kavya Kopparapu, Serena Zhan, Chao Jia, Alexei Bendebury, Sheryl Luo, Adrià Recasens, Timothy Knight, Jing Chen, Mohak Patel, YaGuang Li, Ben Withbroe, Dean Weesner, Kush Bhatia, Jie Ren, Danielle Eisenbud, Ebrahim Songhori, Yanhua Sun, Travis Choma, Tasos Kementsietsidis, Lucas Manning, Brian Roark, Wael Farhan, Jie Feng, Susheel Tatineni, James Cobon-Kerr, Yunjie Li, Lisa Anne Hendricks, Isaac Noble, Chris Breaux, Nate Kushman, Liqian Peng, Fuzhao Xue, Taylor Tobin, Jamie Rogers, Josh Lipschultz, Chris Alberti, Alexey Vlaskin, Mostafa Dehghani, Roshan Sharma, Tris Warkentin, Chen-Yu Lee, Benigno Uria, Da-Cheng Juan, Angad Chandorkar, Hila Sheftel, Ruibo Liu, Elnaz Davoodi, Borja De Balle Pigem, Kedar Dhamdhere, David Ross, Jonathan Hoech, Mahdis Mahdieh, Li Liu, Qiujia Li, Liam McCafferty, Chenxi Liu, Markus Mircea, Yunting Song, Omkar Savant, Alaa Saade, Colin Cherry, Vincent Hellendoorn, Siddharth Goyal, Paul Pucciarelli, David Vilar Torres, Zohar Yahav, Hyo Lee, Lars Lowe Sjoesund, Christo Kirov, Bo Chang, Deepanway Ghoshal, Lu Li, Gilles Baechler, Sébastien Pereira, Tara Sainath, Anudhyan Boral, Dominik Grewe, Afief Halumi, Nguyet Minh Phu, Tianxiao Shen, Marco Tulio Ribeiro, Dhriti Varma, Alex Kaskasoli, Vlad Feinberg, Navneet Potti, Jarrod Kahn, Matheus Wisniewski, Shakir Mohamed, Arnar Mar Hrafnkelsson, Bobak Shahriari, Jean-Baptiste Lespiau, Lisa Patel, Legg Yeung, Tom Paine, Lantao Mei, Alex Ramirez, Rakesh Shivanna, Li Zhong, Josh Woodward, Guilherme Tubone, Samira Khan, Heng Chen, Elizabeth Nielsen, Catalin Ionescu, Utsav Prabhu, Mingcen Gao, Qingze Wang, Sean Augenstein, Neesha Subramaniam, Jason Chang, Fotis Iliopoulos, Jiaming Luo, Myriam Khan, Weicheng Kuo, Denis Teplyashin, Florence Perot, Logan Kilpatrick, Amir Globerson, Hongkun Yu, Anfal Siddiqui, Nick Sukhanov, Arun Kandoor, Umang Gupta, Marco Andreetto, Moran Ambar, Donnie Kim, Paweł Wesołowski, Sarah Perrin, Ben Limonchik, Wei Fan, Jim Stephan, Ian Stewart-Binks, Ryan Kappedal, Tong He, Sarah Cogan, Romina Datta, Tong Zhou, Jiayu Ye, Leandro Kieliger, Ana Ramalho, Kyle Kastner, Fabian Mentzer, Wei-Jen Ko, Arun Suggala, Tianhao Zhou, Shiraz Butt, Hana Strejček, Lior Belenki, Subhashini Venugopalan, Mingyang

Ling, Evgenii Eltyshev, Yunxiao Deng, Geza Kovacs, Mukund Raghavachari, Hanjun Dai, Tal Schuster, Steven Schwarcz, Richard Nguyen, Arthur Nguyen, Gavin Buttimore, Shrestha Basu Mallick, Sudeep Gandhe, Seth Benjamin, Michal Jastrzebski, Le Yan, Sugato Basu, Chris Apps, Isabel Edkins, James Allingham, Immanuel Odisho, Tomas Kocisky, Jewel Zhao, Linting Xue, Apoorv Reddy, Chrysovalantis Anastasiou, Aviel Atias, Sam Redmond, Kieran Milan, Nicolas Heess, Herman Schmit, Allan Dafoe, Daniel Andor, Tynan Gangwani, Anca Dragan, Sheng Zhang, Ashyana Kachra, Gang Wu, Siyang Xue, Kevin Aydin, Siqi Liu, Yuxiang Zhou, Mahan Malihi, Austin Wu, Siddharth Gopal, Candice Schumann, Peter Stys, Alek Wang, Mirek Olšák, Dangyi Liu, Christian Schallhart, Yiran Mao, Demetra Brady, Hao Xu, Tomas Mery, Chawin Sitawarin, Siva Velusamy, Tom Cobley, Alex Zhai, Christian Walder, Nitzan Katz, Ganesh Jawahar, Chinmay Kulkarni, Antoine Yang, Adam Paszke, Yinan Wang, Bogdan Damoc, Zalán Borsos, Ray Smith, Jinning Li, Mansi Gupta, Andrei Kapishnikov, Sushant Prakash, Florian Luisier, Rishabh Agarwal, Will Grathwohl, Kuangyuan Chen, Kehang Han, Nikhil Mehta, Andrew Over, Shekoofeh Azizi, Lei Meng, Niccolò Dal Santo, Kelvin Zheng, Jane Shapiro, Igor Petrovski, Jeffrey Hui, Amin Ghafouri, Jasper Snoek, James Qin, Mandy Jordan, Caitlin Sikora, Jonathan Malmaud, Yuheng Kuang, Aga Świetlik, Ruoxin Sang, Chongyang Shi, Leon Li, Andrew Rosenberg, Shubin Zhao, Andy Crawford, Jan-Thorsten Peter, Yun Lei, Xavier Garcia, Long Le, Todd Wang, Julien Amelot, Dave Orr, Praneeth Kacham, Dana Alon, Gladys Tyen, Abhinav Arora, James Lyon, Alex Kurakin, Mimi Ly, Theo Guidroz, Zhipeng Yan, Rina Panigrahy, Pingmei Xu, Thais Kagohara, Yong Cheng, Eric Noland, Jinhyuk Lee, Jonathan Lee, Cathy Yip, Maria Wang, Efrat Nehoran, Alexander Bykovsky, Zhihao Shan, Ankit Bhagatwala, Chaochao Yan, Jie Tan, Guillermo Garrido, Dan Ethier, Nate Hurley, Grace Vesom, Xu Chen, Siyuan Qiao, Abhishek Nayyar, Julian Walker, Paramjit Sandhu, Mihaela Rosca, Danny Swisher, Mikhail Dektiarev, Josh Dillon, George-Cristian Muraru, Manuel Tragut, Artiom Myaskovsky, David Reid, Marko Velic, Owen Xiao, Jasmine George, Mark Brand, Jing Li, Wenhao Yu, Shane Gu, Xiang Deng, François-Xavier Aubet, Soheil Hassas Yeganeh, Fred Alcober, Celine Smith, Trevor Cohn, Kay McKinney, Michael Tschannen, Ramesh Sampath, Gowoon Cheon, Liangchen Luo, Luyang Liu, Jordi Orbay, Hui Peng, Gabriela Botea, Xiaofan Zhang, Charles Yoon, Cesar Magalhaes, Paweł Stradomski, Ian Mackinnon, Steven Hemingray, Kumaran Venkatesan, Rhys May, Jaeyoun Kim, Alex Druinsky, Jingchen Ye, Zheng Xu, Terry Huang, Jad Al Abdallah, Adil Dostmohamed, Rachana Fellinger, Tsendsuren Munkhdalai, Akanksha Maurya, Peter Garst, Yin Zhang, Maxim Krikun, Simon Bucher, Aditya Srikanth Veerubhotla, Yaxin Liu, Sheng Li, Nishesh Gupta, Jakub Adamek, Hanwen Chen, Bernett Orlando, Aleksandr Zaks, Joost van Amersfoort, Josh Camp, Hui Wan, HyunJeong Choe, Zhichun Wu, Kate Olszewska, Weiren Yu, Archita Vadali, Martin Scholz, Daniel De Freitas, Jason Lin, Amy Hua, Xin Liu, Frank Ding, Yichao Zhou, Boone Severson, Katerina Tsihlas, Samuel Yang, Tammo Spalink, Varun Yerram, Helena Pankov, Rory Blevins, Ben Vargas, Sarthak Jauhari, Matt Miecnikowski, Ming Zhang, Sandeep Kumar, Clement Farabet, Charline Le Lan, Sebastian Flennerhag, Yonatan Bitton, Ada Ma, Arthur Bražinskas, Eli Collins, Niharika Ahuja, Sneha Kudugunta, Anna Bortsova, Minh Giang, Wanzheng Zhu, Ed Chi, Scott Lundberg, Alexey Stern, Subha Puttagunta, Jing Xiong, Xiao Wu, Yash Pande, Amit Jhindal, Daniel Murphy, Jon Clark, Marc Brockschmidt, Maxine Deines, Kevin R. McKee, Dan Bahir, Jiajun Shen, Minh Truong, Daniel McDuff, Andrea Gesmundo, Edouard Rosseel, Bowen Liang, Ken Caluwaerts, Jessica Hamrick, Joseph Kready, Mary Cassin, Rishikesh Ingale, Li Lao, Scott Pollom, Yifan Ding, Wei He, Lizzetth Bellot, Joana Iljazi, Ramya Sree Boppana, Shan Han, Tara Thompson, Amr Khalifa, Anna Bulanova, Blagoj Mitrevski, Bo Pang, Emma Cooney, Tian Shi, Rey Coaguila, Tamar Yakar, Marc'aurelio Ranzato, Nikola Momchev, Chris Rawles, Zachary Charles, Young Maeng, Yuan Zhang, Rishabh Bansal, Xiaokai Zhao, Brian Albert, Yuan Yuan, Sudheendra Vijayanarasimhan, Roy Hirsch, Vinay Ramasesh, Kiran Vodrahalli, Xingyu Wang, Arushi Gupta, DJ Strouse, Jianmo Ni, Roma Patel, Gabe Taubman, Zhouyuan Huo, Dero Gharibian, Marianne Monteiro, Hoi Lam, Shobha Vasudevan, Aditi Chaudhary, Isabela Albuquerque, Kilol Gupta, Sebastian Riedel, Chaitra Hegde, Avraham Ruderman, András György, Marcus Wainwright, Ashwin Chaugule, Burcu Karagol Ayan, Tomer Levinboim, Sam Shleifer, Yogesh Kalley, Vahab Mirrokni, Abhishek Rao, Prabakar Radhakrishnan, Jay Hartford, Jialin Wu, Zhenhai Zhu, Francesco Bertolini, Hao Xiong, Nicolas Serrano, Hamish Tomlinson, Myle Ott, Yifan Chang, Mark Graham, Jian Li, Marco Liang, Xiangzhu Long, Sebastian Borgeaud, Yanif Ahmad, Alex Grills, Diana Mincu, Martin Izzard, Yuan Liu, Jinyu Xie, Louis O'Bryan, Sameera Ponda, Simon Tong, Michelle Liu, Dan Malkin, Khalid Salama, Yuankai Chen, Rohan Anil, Anand Rao, Rigel Swavely, Misha Bilenko, Nina Anderson, Tat Tan, Jing Xie, Xing Wu, Lijun Yu, Oriol Vinyals, Andrey Ryabtsev, Rumen Dangovski, Kate Baumli, Daniel Keysers, Christian Wright, Zoe Ashwood, Betty Chan, Artem Shtefan, Yaohui Guo, Ankur Bapna, Radu Soricut, Steven Pecht, Sabela Ramos, Rui Wang, Jiahao Cai, Trieu Trinh, Paul Barham, Linda Friso, Eli Stickgold, Xiangzhuo Ding, Siamak Shakeri, Diego Ardila, Eleftheria Briakou, Phil Culliton, Adam Raveret, Jingyu Cui, David Saxton, Subhrajit Roy, Javad Azizi, Pengcheng Yin, Lucia Loher, Andrew Bunner, Min Choi, Faruk Ahmed, Eric Li, Yin Li, Shengyang Dai, Michael Elabd, Sriram Ganapathy, Shivani Agrawal, Yiqing Hua, Paige Kunkle, Sujeevan Rajayogam, Arun Ahuja, Arthur Conmy, Alex Vasiloff, Parker Beak, Christopher Yew, Jayaram Mudigonda, Bartek Wydrowski, Jon Blanton, Zhengdong Wang, Yann Dauphin, Zhuo Xu, Martin Polacek, Xi Chen, Hexiang Hu, Pauline Sho, Markus Kunesch, Mehdi Hafezi Manshadi, Eliza Rutherford, Bo Li, Sissie Hsiao, Iain Barr, Alex Tudor, Matija Kecman, Arsha Nagrani, Vladimir Pchelin, Martin Sundermeyer, Aishwarya P S, Abhijit Karmarkar, Yi Gao, Grishma Chole, Olivier Bachem, Isabel Gao, Arturo BC, Matt Dibb, Mauro Verzetti, Felix Hernandez-Campos, Yana Lunts, Matthew Johnson, Julia Di Trapani, Raphael Koster, Idan Brusilovsky, Binbin Xiong, Megha Mohabey, Han Ke, Joe Zou, Tea Sabolić, Víctor Campos, John Palowitch, Alex Morris, Linhai Qiu, Pranavaraj Ponnuramu, Fangtao Li, Vivek Sharma, Kiranbir Sodhia, Kaan Tekelioglu, Aleksandr Chuklin, Madhavi Yenugula, Erika Gemzer, Theofilos Strinopoulos, Sam El-Husseini, Huiyu Wang, Yan Zhong, Edouard Leurent, Paul Natsev, Weijun Wang, Dre Mahaarachchi, Tao Zhu, Songyou Peng, Sami Alabed, Cheng-Chun Lee, Anthony Brohan, Arthur Szlam, GS Oh, Anton Kovsharov, Jenny Lee, Renee Wong, Megan Barnes, Gregory Thornton, Felix Gimeno, Omer Levy, Martin Sevenich, Melvin Johnson, Jonathan Mallinson, Robert Dadashi, Ziyue Wang, Qingchun Ren, Preethi Lahoti, Arka Dhar, Josh Feldman, Dan Zheng, Thatcher Ulrich, Liviu Panait, Michiel Blokzijl, Cip Baetu, Josip Matak, Jitendra Harlalka, Maulik Shah, Tal Marian, Daniel von Dincklage, Cosmo Du, Ruy Ley-Wild, Bethanie Brownfield, Max Schumacher, Yury Stuken, Shadi Noghabi, Sonal Gupta, Xiaoqi Ren, Eric Malmi, Felix Weissenberger, Blanca Huergo, Maria Bauza, Thomas Lampe, Arthur Douillard, Mojtaba Seyedhosseini, Roy Frostig, Zoubin Ghahramani, Kelvin Nguyen, Kashyap Krishnakumar, Chengxi Ye, Rahul Gupta, Alireza Nazari, Robert Geirhos, Pete Shaw, Ahmed Eleryan, Dima Damen, Jennimaria Palomaki, Ted Xiao, Qiyin Wu, Quan Yuan, Phoenix Meadowlark, Matthew Bilotti, Raymond Lin, Mukund Sridhar, Yannick Schroecker, Da-Woon Chung, Jincheng Luo, Trevor Strohman, Tianlin Liu, Anne Zheng, Jesse Emond, Wei Wang, Andrew Lampinen, Toshiyuki Fukuzawa, Folawiyo Campbell-Ajala, Monica Roy, James Lee-Thorp, Lily Wang, Iftekhar Naim, Tony, Nguy ên, Guy Bensky, Aditya Gupta, Dominika Rogozińska, Justin Fu, Thanumalayan Sankaranarayana Pillai, Petar Veličković, Shahar Drath, Philipp Neubeck, Vaibhav Tulsyan, Arseniy Klimovskiy, Don Metzler, Sage Stevens, Angel Yeh, Junwei Yuan, Tianhe Yu, Kelvin Zhang, Alec Go, Vincent Tsang, Ying Xu, Andy Wan, Isaac Galatzer-Levy, Sam Sobell, Abodunrinwa Toki, Elizabeth Salesky, Wenlei Zhou, Diego Antognini, Sholto Douglas, Shimu Wu, Adam Lelkes, Frank Kim, Paul Cavallaro, Ana Salazar, Yuchi Liu, James Besley, Tiziana Refice, Yiling Jia, Zhang Li, Michal Sokolik, Arvind Kannan, Jon Simon, Jo Chick, Avia Aharon, Meet Gandhi, Mayank Daswani, Keyvan Amiri, Vighnesh Birodkar, Abe Ittycheriah, Peter Grabowski, Oscar Chang, Charles Sutton, Zhixin, Lai, Umesh Telang, Susie Sargsyan, Tao Jiang, Raphael Hoffmann, Nicole Brichtova, Matteo Hessel, Jonathan Halcrow, Sammy Jerome, Geoff Brown, Alex Tomala, Elena Buchatskaya, Dian Yu, Sachit Menon, Pol Moreno, Yuguo Liao, Vicky Zayats, Luming Tang, SQ Mah, Ashish Shenoy, Alex Siegman, Majid Hadian, Okwan Kwon, Tao Tu, Nima Khajehnouri, Ryan Foley, Parisa Haghani, Zhongru Wu, Vaishakh Keshava, Khyatti Gupta, Tony Bruguier, Rui Yao, Danny Karmon, Luisa Zintgraf, Zhicheng Wang, Enrique Piqueras, Junehyuk Jung, Jenny Brennan, Diego Machado, Marissa Giustina, MH Tessler, Kamyu Lee, Qiao Zhang, Joss Moore, Kaspar Daugaard, Alexander Frömmgen, Jennifer Beattie, Fred Zhang, Daniel Kasenberg, Ty Geri, Danfeng Qin, Gaurav Singh Tomar, Tom Ouyang, Tianli Yu, Luowei Zhou, Rajiv Mathews, Andy Davis, Yaoyiran Li, Jai Gupta, Damion Yates, Linda Deng, Elizabeth Kemp, Ga-Young Joung, Sergei Vassilvitskii, Mandy Guo, Pallavi LV, Dave Dopson, Sami Lachgar, Lara McConnaughey, Himadri Choudhury, Dragos Dena, Aaron Cohen, Joshua Ainslie, Sergey Levi, Parthasarathy Gopavarapu, Polina Zablotskaia, Hugo Vallet, Sanaz Bahargam, Xiaodan Tang, Nenad Tomasev, Ethan Dyer, Daniel Balle, Hongrae Lee, William Bono, Jorge Gonzalez Mendez, Vadim Zubov, Shentao Yang, Ivor Rendulic, Yanyan Zheng, Andrew Hogue, Golan Pundak, Ralph Leith, Avishkar Bhoopchand, Michael Han, Mislav Žanić, Tom Schaul, Manolis Delakis, Tejas Iyer, Guanyu Wang, Harman Singh, Abdelrahman Abdelhamed, Tara Thomas, Siddhartha Brahma, Hilal Dib, Naveen Kumar, Wenxuan Zhou, Liang Bai, Pushkar Mishra, Jiao Sun, Valentin Anklin, Roykrong Sukkerd, Lauren Agubuzu, Anton Briukhov, Anmol Gulati, Maximilian Sieb, Fabio Pardo, Sara Nasso, Junquan Chen, Kexin Zhu, Tiberiu Sosea, Alex Goldin, Keith Rush, Spurthi Amba Hombaiah, Andreas Noever, Allan Zhou, Sam Haves, Mary Phuong, Jake Ades, Yi ting Chen, Lin Yang, Joseph Pagadora, Stan Bileschi, Victor Cotruta, Rachel Saputro, Arijit Pramanik, Sean Ammirati, Dan Garrette, Kevin Villela, Tim Blyth, Canfer Akbulut, Neha Jha, Alban Rrustemi, Arissa Wongpanich, Chirag Nagpal, Yonghui Wu, Morgane Rivière, Sergey Kishchenko, Pranesh Srinivasan, Alice Chen, Animesh Sinha, Trang Pham, Bill Jia, Tom Hennigan, Anton Bakalov, Nithya Attaluri, Drew Garmon, Daniel Rodriguez, Dawid Wegner, Wenhao Jia, Evan Senter, Noah Fiedel, Denis Petek, Yuchuan Liu, Cassidy Hardin, Harshal Tushar Lehri, Joao Carreira, Sara Smoot, Marcel Prasetya, Nami Akazawa, Anca Stefanoiu, Chia-Hua Ho, Anelia Angelova, Kate Lin, Min Kim, Charles Chen, Marcin Sieniek, Alice Li, Tongfei Guo, Sorin Baltateanu, Pouya Tafti, Michael Wunder, Nadav Olmert, Divyansh Shukla, Jingwei Shen, Neel Kovelamudi, Balaji Venkatraman, Seth Neel, Romal Thoppilan, Jerome Connor, Frederik Benzing, Axel Stjerngren, Golnaz Ghiasi, Alex Polozov, Joshua Howland, Theophane Weber, Justin Chiu, Ganesh Poomal Girirajan, Andreas Terzis, Pidong Wang, Fangda Li, Yoav Ben Shalom, Dinesh Tewari, Matthew Denton, Roee Aharoni, Norbert Kalb, Heri Zhao, Junlin Zhang, Angelos Filos, Matthew Rahtz, Lalit Jain, Connie Fan, Vitor Rodrigues, Ruth Wang, Richard Shin, Jacob Austin, Roman Ring, Mariella Sanchez-Vargas, Mehadi Hassen, Ido Kessler, Uri Alon, Gufeng Zhang, Wenhu Chen, Yenai Ma, Xiance Si, Le Hou, Azalia Mirhoseini, Marc Wilson, Geoff Bacon, Becca Roelofs, Lei Shu, Gautam Vasudevan, Jonas Adler, Artur Dwornik, Tayfun Terzi, Matt Lawlor, Harry Askham, Mike Bernico, Xuanyi Dong, Chris Hidey, Kevin Kilgour, Gaël Liu, Surya Bhupatiraju, Luke Leonhard, Siqi Zuo, Partha Talukdar, Qing Wei, Aliaksei Severyn, Vít Listík, Jong Lee, Aditya Tripathi, SK Park, Yossi Matias, Hao Liu, Alex Ruiz, Rajesh Jayaram, Jackson Tolins, Pierre Marcenac, Yiming Wang, Bryan Seybold, Henry Prior, Deepak Sharma, Jack Weber, Mikhail Sirotenko, Yunhsuan Sung, Dayou Du, Ellie Pavlick, Stefan Zinke, Markus Freitag, Max Dylla, Montse Gonzalez Arenas, Natan Potikha, Omer Goldman, Connie Tao, Rachita Chhaparia, Maria Voitovich, Pawan Dogra, Andrija Ražnatović, Zak Tsai, Chong You, Oleaser Johnson, George Tucker, Chenjie Gu, Jae Yoo, Maryam Majzoubi, Valentin Gabeur, Bahram Raad, Rocky Rhodes, Kashyap Kolipaka, Heidi Howard, Geta Sampemane, Benny Li, Chulayuth Asawaroengchai, Duy Nguyen, Chiyuan Zhang, Timothee Cour, Xinxin Yu, Zhao Fu, Joe Jiang, Po-Sen Huang, Gabriela Surita, Iñaki Iturrate, Yael Karov, Michael Collins, Martin Baeuml, Fabian Fuchs, Shilpa Shetty, Swaroop Ramaswamy, Sayna Ebrahimi, Qiuchen Guo, Jeremy Shar, Gabe Barth-Maron, Sravanti Addepalli, Bryan Richter, Chin-Yi Cheng, Eugénie Rives, Fei Zheng, Johannes Griesser, Nishanth Dikkala, Yoel Zeldes, Ilkin Safarli, Dipanjan Das, Himanshu Srivastava, Sadh MNM Khan, Xin Li, Aditya Pandey, Larisa Markeeva, Dan Belov, Qiqi Yan, Mikołaj Rybiński, Tao Chen, Megha Nawhal, Michael Quinn, Vineetha Govindaraj, Sarah York, Reed Roberts, Roopal Garg, Namrata Godbole, Jake Abernethy, Anil Das, Lam Nguyen Thiet, Jonathan Tompson, John Nham, Neera Vats, Ben Caine, Wesley Helmholz, Francesco Pongetti, Yeongil Ko, James An, Clara Huiyi Hu, Yu-Cheng Ling, Julia Pawar, Robert Leland, Keisuke Kinoshita, Waleed Khawaja, Marco Selvi, Eugene Ie, Danila Sinopalnikov, Lev Proleev, Nilesh Tripuraneni, Michele Bevilacqua, Seungji Lee, Clayton Sanford, Dan Suh, Dustin Tran, Jeff Dean, Simon Baumgartner, Jens Heitkaemper, Sagar Gubbi, Kristina Toutanova, Yichong Xu, Chandu Thekkath, Keran Rong, Palak Jain, Annie Xie, Yan Virin, Yang Li, Lubo Litchev, Richard Powell, Tarun Bharti, Adam Kraft, Nan Hua, Marissa Ikonomidis, Ayal Hitron, Sanjiv Kumar, Loic Matthey, Sophie Bridgers, Lauren Lax, Ishaan Malhi, Ondrej Skopek, Ashish Gupta, Jiawei Cao, Mitchelle Rasquinha, Siim Põder, Wojciech Stokowiec, Nicholas Roth, Guowang Li, Michaël Sander, Joshua Kessinger, Vihan Jain, Edward Loper, Wonpyo Park, Michal Yarom, Liqun Cheng, Guru Guruganesh, Kanishka Rao, Yan Li, Catarina Barros, Mikhail Sushkov, Chun-Sung Ferng, Rohin Shah, Ophir Aharoni, Ravin Kumar, Tim McConnell, Peiran Li, Chen Wang, Fernando Pereira, Craig Swanson, Fayaz Jamil, Yan Xiong, Anitha Vijayakumar, Prakash Shroff, Kedar Soparkar, Jindong Gu, Livio Baldini Soares, Eric Wang, Kushal Majmundar, Aurora Wei, Kai Bailey, Nora Kassner, Chizu Kawamoto, Goran Žužić, Victor Gomes, Abhirut Gupta, Michael Guzman, Ishita Dasgupta, Xinyi Bai, Zhufeng Pan, Francesco Piccinno, Hadas Natalie Vogel, Octavio Ponce, Adrian Hutter, Paul Chang, Pan-Pan Jiang, Ionel Gog, Vlad Ionescu, James Manyika, Fabian Pedregosa, Harry Ragan, Zach Behrman, Ryan Mullins, Coline Devin, Aroonalok Pyne, Swapnil Gawde, Martin Chadwick, Yiming Gu, Sasan Tavakkol, Andy Twigg, Naman Goyal, Ndidi Elue, Anna Goldie, Srinivasan Venkatachary, Hongliang Fei, Ziqiang Feng, Marvin Ritter, Isabel Leal, Sudeep Dasari, Pei Sun, Alif Raditya Rochman, Brendan O'Donoghue, Yuchen Liu, Jim Sproch, Kai Chen, Natalie Clay, Slav Petrov, Sailesh Sidhwani, Ioana Mihailescu, Alex Panagopoulos, AJ Piergiovanni, Yunfei Bai, George Powell, Deep Karkhanis, Trevor Yacovone, Petr Mitrichev, Joe Kovac, Dave Uthus, Amir Yazdanbakhsh, David Amos, Steven Zheng, Bing Zhang, Jin Miao, Bhuvana Ramabhadran, Soroush Radpour, Shantanu Thakoor, Josh Newlan, Oran Lang, Orion Jankowski, Shikhar Bharadwaj, Jean-Michel Sarr, Shereen Ashraf, Sneha Mondal, Jun Yan, Ankit Singh Rawat, Sarmishta Velury, Greg Kochanski, Tom Eccles, Franz Och, Abhanshu Sharma, Ethan Mahintorabi, Alex Gurney, Carrie Muir, Vered Cohen, Saksham Thakur, Adam Bloniarz, Asier Mujika, Alexander Pritzel, Paul Caron, Altaf Rahman, Fiona Lang, Yasumasa Onoe, Petar Sirkovic, Jay Hoover, Ying Jian, Pablo Duque, Arun Narayanan, David Soergel, Alex Haig, Loren Maggiore, Shyamal Buch, Josef Dean, Ilya Figotin, Igor Karpov, Shaleen Gupta, Denny Zhou, Muhuan Huang, Ashwin Vaswani, Christopher Semturs, Kaushik Shivakumar, Yu Watanabe, Vinodh Kumar Rajendran, Eva Lu, Yanhan Hou, Wenting Ye, Shikhar Vashishth, Nana Nti, Vytenis Sakenas, Darren Ni, Doug DeCarlo, Michael Bendersky, Sumit Bagri, Nacho Cano, Elijah Peake, Simon Tokumine, Varun Godbole, Carlos Guía, Tanya Lando, Vittorio Selo, Seher Ellis, Danny Tarlow, Daniel Gillick, Alessandro Epasto, Siddhartha Reddy Jonnalagadda, Meng Wei, Meiyan Xie, Ankur Taly, Michela Paganini, Mukund Sundararajan, Daniel Toyama, Ting Yu, Dessie Petrova, Aneesh Pappu, Rohan Agrawal, Senaka Buthpitiya, Justin Frye, Thomas Buschmann, Remi Crocker, Marco Tagliasacchi, Mengchao Wang, Da Huang, Sagi Perel, Brian Wieder, Hideto Kazawa, Weiyue Wang, Jeremy Cole, Himanshu Gupta, Ben Golan, Seojin Bang, Nitish Kulkarni, Ken Franko, Casper Liu, Doug Reid, Sid Dalmia, Jay Whang, Kevin Cen, Prasha Sundaram, Johan Ferret, Berivan Isik, Lucian Ionita, Guan Sun, Anna Shekhawat, Muqthar Mohammad, Philip Pham, Ronny Huang, Karthik Raman, Xingyi Zhou, Ross Mcilroy, Austin Myers, Sheng Peng, Jacob Scott, Paul Covington, Sofia Erell, Pratik Joshi, João Gabriel Oliveira, Natasha Noy, Tajwar Nasir, Jake Walker, Vera Axelrod, Tim Dozat, Pu Han, Chun-Te Chu, Eugene Weinstein, Anand Shukla, Shreyas Chandrakaladharan, Petra Poklukar, Bonnie Li, Ye Jin, Prem Eruvbetine, Steven Hansen, Avigail Dabush, Alon Jacovi, Samrat Phatale, Chen Zhu, Steven Baker, Mo Shomrat, Yang Xiao, Jean Pouget-Abadie, Mingyang Zhang, Fanny Wei, Yang Song, Helen King, Yiling Huang, Yun Zhu, Ruoxi Sun, Juliana Vicente Franco, Chu-Cheng Lin, Sho Arora, Hui, Li, Vivian Xia, Luke Vilnis, Mariano Schain, Kaiz Alarakyia, Laurel Prince, Aaron Phillips, Caleb Habtegebriel, Luyao Xu, Huan Gui, Santiago Ontanon, Lora Aroyo, Karan Gill, Peggy Lu, Yash Katariya, Dhruv Madeka, Shankar Krishnan, Shubha Srinivas Raghvendra, James Freedman, Yi Tay, Gaurav Menghani, Peter Choy, Nishita Shetty, Dan Abolafia, Doron Kukliansky, Edward Chou, Jared Lichtarge, Ken Burke, Ben Coleman, Dee Guo, Larry Jin, Indro Bhattacharya, Victoria Langston, Yiming Li, Suyog Kotecha, Alex Yakubovich, Xinyun Chen, Petre Petrov, Tolly Powell, Yanzhang He, Corbin Quick, Kanav Garg, Dawsen Hwang, Yang Lu, Srinadh Bhojanapalli, Kristian Kjems, Ramin Mehran, Aaron Archer, Hado van Hasselt, Ashwin Balakrishna, JK Kearns, Meiqi Guo, Jason Riesa, Mikita Sazanovich, Xu Gao, Chris Sauer, Chengrun Yang, XiangHai Sheng, Thomas Jimma, Wouter Van Gansbeke, Vitaly Nikolaev, Wei Wei, Katie Millican, Ruizhe Zhao, Justin Snyder, Levent Bolelli, Maura O'Brien, Shawn Xu, Fei Xia, Wentao Yuan, Arvind Neelakantan, David Barker, Sachin Yadav, Hannah Kirkwood, Farooq Ahmad, Joel Wee, Jordan Grimstad, Boyu Wang, Matthew Wiethoff, Shane Settle, Miaosen Wang, Charles Blundell, Jingjing Chen, Chris Duvarney, Grace Hu, Olaf Ronneberger, Alex Lee, Yuanzhen Li, Abhishek Chakladar, Alena Butryna, Georgios Evangelopoulos, Guillaume Desjardins, Jonni Kanerva, Henry Wang, Averi Nowak, Nick Li, Alyssa Loo, Art Khurshudov, Laurent El Shafey, Nagabhushan Baddi, Karel Lenc, Yasaman Razeghi, Tom Lieber, Amer Sinha, Xiao Ma, Yao Su, James Huang, Asahi Ushio, Hanna Klimczak-Plucińska, Kareem Mohamed, JD Chen, Simon Osindero, Stav Ginzburg, Lampros Lamprou, Vasilisa Bashlovkina, Duc-Hieu Tran, Ali Khodaei, Ankit Anand, Yixian Di, Ramy Eskander, Manish Reddy Vuyyuru, Jasmine Liu, Aishwarya Kamath, Roman Goldenberg, Mathias Bellaiche, Juliette Pluto, Bill Rosgen, Hassan Mansoor, William Wong, Suhas Ganesh, Eric Bailey, Scott Baird, Dan Deutsch, Jinoo Baek, Xuhui Jia, Chansoo Lee, Abe Friesen, Nathaniel Braun, Kate Lee, Amayika Panda, Steven M. Hernandez, Duncan Williams, Jianqiao Liu, Ethan Liang, Arnaud Autef, Emily Pitler, Deepali Jain, Phoebe Kirk, Oskar Bunyan, Jaume Sanchez Elias, Tongxin Yin, Machel Reid, Aedan Pope, Nikita Putikhin, Bidisha Samanta, Sergio Guadarrama, Dahun Kim, Simon Rowe, Marcella Valentine, Geng Yan, Alex Salcianu, David Silver, Gan Song, Richa Singh, Shuai Ye, Hannah DeBalsi, Majd Al Merey, Eran Ofek, Albert Webson, Shibl Mourad, Ashwin Kakarla, Silvio Lattanzi, Nick Roy, Evgeny Sluzhaev, Christina Butterfield, Alessio Tonioni, Nathan Waters, Sudhindra Kopalle, Jason Chase, James Cohan, Girish Ramchandra Rao, Robert Berry, Michael Voznesensky, Shuguang Hu, Kristen Chiafullo, Sharat Chikkerur, George Scrivener, Ivy Zheng, Jeremy Wiesner, Wolfgang Macherey, Timothy Lillicrap, Fei Liu, Brian Walker, David Welling, Elinor Davies, Yangsibo Huang, Lijie Ren, Nir Shabat, Alessandro Agostini, Mariko Iinuma, Dustin Zelle, Rohit Sathyanarayana, Andrea D'olimpio, Morgan Redshaw, Matt Ginsberg, Ashwin Murthy, Mark Geller, Tatiana Matejovicova, Ayan Chakrabarti, Ryan Julian, Christine Chan, Qiong Hu, Daniel Jarrett, Manu Agarwal, Jeshwanth Challagundla, Tao Li, Sandeep Tata, Wen Ding, Maya Meng, Zhuyun Dai, Giulia Vezzani, Shefali Garg, Jannis Bulian, Mary Jasarevic, Honglong Cai, Harish Rajamani, Adam Santoro, Florian Hartmann, Chen Liang, Bartek Perz, Apoorv Jindal, Fan Bu, Sungyong Seo, Ryan Poplin, Adrian Goedeckemeyer, Badih Ghazi, Nikhil Khadke, Leon Liu, Kevin Mather, Mingda Zhang, Ali Shah, Alex Chen, Jinliang Wei, Keshav Shivam, Yuan Cao, Donghyun Cho, Angelo Scorza Scarpati, Michael Moffitt, Clara Barbu, Ivan Jurin, Ming-Wei Chang, Hongbin Liu, Hao Zheng, Shachi Dave, Christine Kaeser-Chen, Xiaobin Yu, Alvin Abdagic, Lucas Gonzalez, Yanping Huang, Peilin Zhong, Cordelia Schmid, Bryce Petrini, Alex Wertheim, Jifan Zhu, Hoang Nguyen, Kaiyang Ji, Yanqi Zhou, Tao Zhou, Fangxiaoyu Feng, Regev Cohen, David Rim, Shubham Milind Phal, Petko Georgiev, Ariel Brand, Yue Ma, Wei Li, Somit Gupta, Chao Wang, Pavel Dubov, Jean Tarbouriech, Kingshuk Majumder, Huijian Li, Norman Rink, Apurv Suman, Yang Guo, Yinghao Sun, Arun Nair, Xiaowei Xu, Mohamed Elhawaty, Rodrigo Cabrera, Guangxing Han, Julian Eisenschlos, Junwen Bai, Yuqi Li, Yamini Bansal, Thibault Sellam, Mina Khan, Hung Nguyen, Justin Mao-Jones, Nikos Parotsidis, Jake Marcus, Cindy Fan, Roland Zimmermann, Yony Kochinski, Laura Graesser, Feryal Behbahani, Alvaro Caceres, Michael Riley, Patrick Kane, Sandra Lefdal, Rob Willoughby, Paul Vicol, Lun Wang, Shujian Zhang, Ashleah Gill, Yu Liang, Gautam Prasad, Soroosh Mariooryad, Mehran Kazemi, Zifeng Wang, Kritika Muralidharan, Paul Voigtlaender, Jeffrey Zhao, Huanjie Zhou, Nina D'Souza, Aditi Mavalankar, Séb Arnold, Nick Young, Obaid Sarvana, Chace Lee, Milad Nasr, Tingting Zou, Seokhwan Kim, Lukas Haas, Kaushal Patel, Neslihan Bulut, David Parkinson, Courtney Biles, Dmitry Kalashnikov, Chi Ming To, Aviral Kumar, Jessica Austin, Alex Greve, Lei Zhang, Megha Goel, Yeqing Li, Sergey Yaroshenko, Max Chang, Abhishek Jindal, Geoff Clark, Hagai Taitelbaum, Dale Johnson, Ofir Roval, Jeongwoo Ko, Anhad Mohananey, Christian Schuler, Shenil Dodhia, Ruichao Li, Kazuki Osawa, Claire Cui, Peng Xu, Rushin Shah, Tao Huang, Ela Gruzewska, Nathan Clement, Mudit Verma, Olcan Sercinoglu, Hai Qian, Viral Shah, Masa Yamaguchi, Abhinit Modi, Takahiro Kosakai, Thomas Strohmann, Junhao Zeng, Beliz Gunel, Jun Qian, Austin Tarango, Krzysztof Jastrzebski, Robert David, Jyn Shan, Parker Schuh, Kunal Lad, Willi Gierke, Mukundan Madhavan, Xinyi Chen, Mark Kurzeja, Rebeca Santamaria-Fernandez, Dawn Chen, Alexandra Cordell, Yuri Chervonyi, Frankie Garcia, Nithish Kannen, Vincent Perot, Nan Ding, Shlomi Cohen-Ganor, Victor Lavrenko, Junru Wu, Georgie Evans, Cicero Nogueira dos Santos, Madhavi Sewak, Ashley Brown, Andrew Hard, Joan Puigcerver, Zeyu Zheng, Yizhong Liang, Evgeny Gladchenko, Reeve Ingle, Uri First, Pierre Sermanet, Charlotte Magister, Mihajlo Velimirović, Sashank Reddi, Susanna Ricco, Eirikur Agustsson, Hartwig Adam, Nir Levine, David Gaddy, Dan Holtmann-Rice, Xuanhui Wang, Ashutosh Sathe, Ab-

hijit Guha Roy, Blaž Bratanič, Alen Carin, Harsh Mehta, Silvano Bonacina, Nicola De Cao, Mara Finkelstein, Verena Rieser, Xinyi Wu, Florent Altché, Dylan Scandinaro, Li Li, Nino Vieillard, Nikhil Sethi, Garrett Tanzer, Zhi Xing, Shibo Wang, Parul Bhatia, Gui Citovsky, Thomas Anthony, Sharon Lin, Tianze Shi, Shoshana Jakobovits, Gena Gibson, Raj Apte, Lisa Lee, Mingqing Chen, Arunkumar Byravan, Petros Maniatis, Kellie Webster, Andrew Dai, Pu-Chin Chen, Jiaqi Pan, Asya Fadeeva, Zach Gleicher, Thang Luong, and Niket Kumar Bhumihar, "Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities," 2025.

[9] Pranav Rajpurkar, Emma Chen, Oishi Banerjee, and Eric J Topol, "Ai in health and medicine," *Nature medicine*, vol. 28, no. 1, pp. 31–38, 2022.

[10] Jian Guan, "Artificial intelligence in healthcare and medicine: promises, ethical challenges and governance," *Chinese Medical Sciences Journal*, vol. 34, no. 2, pp. 76–83, 2019.

[11] Kevin B Johnson, Wei-Qi Wei, Dilhan Weeraratne, Mark E Frisse, Karl Misulis, Kyu Rhee, Juan Zhao, and Jane L Snowdon, "Precision medicine, ai, and the future of personalized health care," *Clinical and translational science*, vol. 14, no. 1, pp. 86–93, 2021.

[12] Michael Wornow, Yizhe Xu, Rahul Thapa, Birju Patel, Ethan Steinberg, Scott Fleming, Michael A Pfeffer, Jason Fries, and Nigam H Shah, "The shaky foundations of large language models and foundation models for electronic health records," *npj Digital Medicine*, vol. 6, no. 1, pp. 135, 2023.

[13] Yihan Lin, Zhirong Bella Yu, and Simon Lee, "A case study exploring the current landscape of synthetic medical record generation with commercial llms," *arXiv preprint arXiv:2504.14657*, 2025.

[14] Emily Alsentzer, John R Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott, "Publicly available clinical bert embeddings," *arXiv preprint arXiv:1904.03323*, 2019.

[15] Zeljko Kraljevic, Dan Bean, Anthony Shek, Rebecca Bendayan, Harry Hemingway, Joshua Au Yeung, Alexander Deng, Alfred Baston, Jack Ross, Esther Idowu, et al., "Foresight—a generative pretrained transformer for modelling of patient timelines using electronic health records: a retrospective modelling study," *The Lancet Digital Health*, vol. 6, no. 4, pp. e281–e290, 2024.

[16] Kyoka Ono and Simon A Lee, "Text serialization and their relationship with the conventional paradigms of tabular machine learning," *arXiv preprint arXiv:2406.13846*, 2024.

[17] Ethan Steinberg, Ken Jung, Jason A Fries, Conor K Corbin, Stephen R Pfohl, and Nigam H Shah, "Language models are an effective representation learning technique for electronic health record data," *Journal of biomedical informatics*, vol. 113, pp. 103637, 2021.

[18] Chao Pang, Xinzhuo Jiang, Krishna S Kalluri, Matthew Spotnitz, RuiJun Chen, Adler Perotte, and Karthik Natarajan, "Cehr-bert: Incorporating temporal information from structured ehr data to improve prediction tasks," in *Machine Learning for Health*. PMLR, 2021, pp. 239–260.

[19] Simon A Lee, Sujay Jain, Alex Chen, Kyoka Ono, Arabdha Biswas, Ákos Rudas, Jennifer Fang, and Jeffrey N Chiang, "Clinical decision support using pseudo-notes from multiple streams of ehr data," *npj Digital Medicine*, vol. 8, no. 1, pp. 394, 2025.

[20] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019.

[21] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang, "Diffusion models: A comprehensive survey of methods and applications," *ACM computing surveys*, vol. 56, no. 4, pp. 1–39, 2023.

[22] Zheng Dong, Zekai Fan, and Shixiang Zhu, "Conditional generative modeling for high-dimensional marked temporal point processes," 2025.

[23] Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Lukas Vierling, Donghai Hong, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Juntao Dai, Xuehai Pan, Kwan Yee Ng, Aidan O'Gara, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao, "Ai alignment: A comprehensive survey," 2025.

[24] Su-In Lee and Eric J Topol, "The clinical potential of counterfactual ai models," *The Lancet*, vol. 403, no. 10428, pp. 717, 2024.

[25] Florian Lehmann and Daniel Buschek, "Examining autocompletion as a basic concept for interaction with generative ai," *i-com*, vol. 19, no. 3, pp. 251–264, 2021.

[26] Wai-Chung Kwan, Hong-Ru Wang, Hui-Min Wang, and Kam-Fai Wong, "A survey on recent advances and challenges in reinforcement learning methods for task-oriented dialogue policy learning," *Machine Intelligence Research*, vol. 20, no. 3, pp. 318–334, Jan. 2023.

[27] Facundo Manuel Quiroga, Jordina Torrents-Barrena, Laura Cristina Lanzarini, and Domenec Puig-Valls, "Invariance measures for neural networks," *Applied Soft Computing*, vol. 132, pp. 109817, Jan. 2023.

[28] Zukang Liao, Pengfei Zhang, and Min Chen, "Ml4ml: Automated invariance testing for machine learning models," 2022.

[29] Masanari Kimura, Ryotaro Shimizu, Yuki Hirakawa, Ryosuke Goto, and Yuki Saito, "On permutation-invariant neural networks," 2024.

[30] Hao Duan and Guido Montúfar, "Understanding learning invariance in deep linear networks," 2025.

[31] Aleksia Kolo, Chao Pang, Edward Choi, Ethan Steinberg, Hyewon Jeong, Jack Gallifant, Jason A Fries, Jeffrey N Chiang, Jungwoo Oh, Justin Xu, et al., "Meds decentralized, extensible validation (meds-dev) benchmark: Establishing reproducibility and comparability in ml for health," 2024.

[32] Lea Goetz, Nabeel Seedat, Robert Vandersluis, and Mihaela van der Schaar, "Generalization—a key challenge for responsible ai in patient-facing clinical applications," *npj Digital Medicine*, vol. 7, no. 1, pp. 126, 2024.

[33] Matthew BA McDermott, Justin Xu, Teya S Bergamaschi, Hyewon Jeong, Simon A Lee, Nassim Oufattole, Patrick Rockenschaub, Kamilė Stankevičiūtė, Ethan Steinberg, Jimeng Sun, et al., "Meds: Building models and tools in a reproducible health ai ecosystem," in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 2025, pp. 6243–6244.

[34] Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi-Khorshidi, "Behrt: transformer for electronic health records," *Scientific reports*, vol. 10, no. 1, pp. 7155, 2020.

[35] Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi, "Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction," *NPJ digital medicine*, vol. 4, no. 1, pp. 86, 2021.

[36] Chao Pang, Xinzhuo Jiang, Nishanth Parameshwar Pavinkurve, Krishna S Kalluri, Elise L Minto, Jason Patterson, Linying Zhang, George Hripcsak, Gamze Gürsoy, Noémie Elhadad, et al., "Cehr-gpt: Generating electronic health records with chronological patient timelines," *arXiv preprint arXiv:2402.04400*, 2024.

[37] Chao Pang, Xinzhuo Jiang, Krishna S. Kalluri, Matthew Spotnitz, RuiJun Chen, Adler Perotte, and Karthik Natarajan, "Cehr-bert: Incorporating temporal information from structured ehr data to improve prediction tasks," in *Proceedings of Machine Learning for Health*, Subhrajit Roy, Stephen Pfohl, Emma Rocheteau, Girmaw Abebe Tadesse, Luis Oala, Fabian Falck, Yuyin Zhou, Liyue Shen, Ghada Zamzmi, Purity Mugambi, Ayah Zirikly, Matthew B. A. McDermott, and Emily Alsentzer, Eds. 04 Dec 2021, vol. 158 of *Proceedings of Machine Learning Research*, pp. 239–260, PMLR.

[38] Ethan Steinberg, Yizhe Xu, Jason Alan Fries, and Nigam Shah, "MOTOR: A time-to-event foundation model for structured medical records," in *The Twelfth International Conference on Learning Representations*, 2024.

[39] Kyunghoon Hur, Jungwoo Oh, Junu Kim, Jiyoun Kim, Min Jae Lee, Eunbyeol Cho, Seong-Eun Moon, Young-Hak Kim, Louis Atallah, and Edward Choi, "Genhpf: General healthcare predictive framework for multi-task multi-source learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 1, pp. 502–513, 2023.

[40] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag, "Tabllm: Few-shot classification of tabular data with large language models," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2023, pp. 5549–5581.

[41] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark, "Mimic-iii, a freely accessible critical care database," *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.

[42] Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al., "Mimic-iv, a freely accessible electronic health record dataset," *Scientific data*, vol. 10, no. 1, pp. 1, 2023.

[43] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi, "The eicu collaborative research database, a freely available multi-center database for critical care research," *Scientific data*, vol. 5, no. 1, pp. 1–13, 2018.

[44] Simon A Lee and Timothy Lindsey, "Can large language models abstract medical coded language?," *arXiv preprint arXiv:2403.10822*, 2024.

[45] Lin Lawrence Guo, Jason Fries, Ethan Steinberg, Scott Lanyon Fleming, Keith Morse, Catherine Aftandilian, Jose Posada, Nigam Shah, and Lillian Sung, "A multi-center study on the adaptability of a shared foundation model for electronic health records," *NPJ Digital Medicine*, vol. 7, no. 1, pp. 171, 2024.

[46] Ulzee An, Simon A Lee, Moonseong Jeong, Aditya Gorla, Jeffrey N Chiang, and Sriram Sankararaman, "Dk-behrt: Teaching language models international classification of disease (icd) codes using known disease descriptions," in *AAAI Bridge Program on AI for Medicine and Healthcare*. PMLR, 2025, pp. 133–143.

[47] Rui Wang and Rose Yu, "Physics-guided deep learning for dynamical systems: A survey," 2023.

[48] Ashkan Nikeghbali, "An essay on the general theory of stochastic processes," *Probability Surveys*, vol. 3, no. none, Jan. 2006.

[49] Jared Chessari, Reiichiro Kawai, Yuji Shinozaki, and Toshihiro Yamada, "Numerical methods for backward stochastic differential equations: A survey," *Probability Surveys*, vol. 20, no. none, Jan. 2023.

[50] Seungsoo Lee, "Finite time explosion of stochastic differential equations: A survey into khasminskii's lyapunov method and its consistency with the osgood criterion," 2024.

[51] Yana A. Butko and Merten Mlinarzik, "Fractional ito calculus for randomly scaled fractional brownian motion and its applications to evolution equations," 2024.

[52] Riccardo Mannella, "Numerical integration of stochastic differential equations: The heun algorithm revisited and itô-stratonovich calculus," 2025.

[53] Feng Zhou, Quyu Kong, Jie Qiao, Cheng Wan, Yixuan Zhang, and Ruichu Cai, "Advances in temporal point processes: Bayesian, neural, and llm approaches," 2025.

[54] Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann, "Neural temporal point processes: A review," 2021.

[55] Ottmar Cronie, Mehdi Moradi, and Christophe A. N. Biscio, "Statistical learning and cross-validation for point processes," 2021.

[56] Matthew McDermott, Bret Nestor, Peniel Argaw, and Isaac S Kohane, "Event stream gpt: a data pre-processing and modeling library for generative, pre-trained transformers over continuous-time sequences of complex events," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[57] Patrick Kidger, James Morrill, James Foster, and Terry Lyons, "Neural controlled differential equations for irregular time series," 2020.

[58] Hyeonjung, Jung, Jayant Gupta, Bharat Jayaprakash, Matthew Eagon, Harish Panneer Selvam, Carl Molnar, William Northrop, and Shashi Shekhar, "A survey on solving and discovering differential equations using deep neural networks," 2023.

[59] YongKyung Oh, Seungsu Kam, Jonghun Lee, Dong-Young Lim, Sungil Kim, and Alex Bui, "Comprehensive review of neural differential equations for time series analysis," 2025.

[60] Simon A Lee and Kai Akamatsu, "Foundation models for physiological signals: Opportunities and challenges," 2025.

[61] Jan Ole Ernst, Jan Snoeijs, Mitchell Peaks, and Jochen Wolf, "Memory-optimised cubic splines for high-fidelity quantum operations," 2024.

[62] Xitai Yu, "Application research of spline interpolation and arima in the field of stock market forecasting," 2023.

[63] W. Chen, "Jacobian matrix: a bridge between linear and nonlinear polynomial-only problems," 1999.

[64] Nora Belrose and Adam Scherlis, "Understanding gradient descent through the training jacobian," 2024.

[65] David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe, "Variational inference: A review for statisticians," *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, Apr. 2017.

[66] Ankush Ganguly and Samuel W. F. Earp, "An introduction to variational inference," 2021.

[67] Alex Glyn-Davies, Arnaud Vadeboncoeur, O. Deniz Akyildiz, Ieva Kazlauskaite, and Mark Girolami, "A primer on variational inference for physics-informed deep generative modelling," 2025.

[68] Jinlong Wei, Junhao Hu, Guangying Lv, and Chenggui Yuan, "The euler-maruyama method for sdes with low-regularity drift," 2025.

[69] Yinna Ye and Xiequan Fan, "Convergence rate of euler-maruyama scheme to the invariant probability measure under total variation distance," 2025.

[70] Chuchu Chen, Tonghe Dang, Jialin Hong, and Guoting Song, "Longtime behaviors of θ-euler-maruyama method for stochastic functional differential equations," 2024.

[71] Andrew Alden, Blanka Horvath, and Zacharia Issa, "Signature maximum mean discrepancy two-sample statistical tests," 2025.

[72] Leslie O'Bray, Max Horn, Bastian Rieck, and Karsten Borgwardt, "Evaluation metrics for graph generative models: Problems, pitfalls, and practical solutions," 2022.

[73] Advait Parulekar, Liam Collins, Karthikeyan Shanmugam, Aryan Mokhtari, and Sanjay Shakkottai, "Infonce loss provably learns cluster-preserving representations," 2023.

[74] Evgenia Rusak, Patrik Reizinger, Attila Juhos, Oliver Bringmann, Roland S. Zimmermann, and Wieland Brendel, "Infonce: Identifying the gap between theory and practice," 2025.

[75] Timo Bertram, Johannes Fürnkranz, and Martin Müller, "Contrastive learning of preferences with a contextual infonce loss," 2024.

[76] Vincent Fortuin, "Priors in bayesian deep learning: A review," 2022.

[77] Dylan Sam, Rattana Pukdee, Daniel P. Jeong, Yewon Byun, and J. Zico Kolter, "Bayesian neural networks with domain knowledge priors," 2024.

[78] Pascal R. van der Vaart, Neil Yorke-Smith, and Matthijs T. J. Spaan, "Priors matter: Addressing misspecification in bayesian deep q-learning," 2025.

[79] Fernando Llorente, Luca Martino, Jesse Read, and David Delgado-Gómez, "A survey of monte carlo methods for noisy and costly densities with application to reinforcement learning and abc," *International Statistical Review*, vol. 93, no. 1, pp. 18–61, May 2024.

[80] David Luengo, Luca Martino, Mónica Bugallo, Víctor Elvira, and Simo Särkkä, "A survey of monte carlo methods for parameter estimation," *EURASIP Journal on Advances in Signal Processing*, vol. 2020, no. 1, May 2020.

[81] Herbert Woisetschläger, Alexander Erben, Shiqiang Wang, Ruben Mayer, and Hans-Arno Jacobsen, "A survey on efficient federated learning methods for foundation model training," in *Proceedings of the Thirty-ThirdInternational Joint Conference on Artificial Intelligence*. Aug. 2024, IJCAI-2024, p. 8317–8325, International Joint Conferences on Artificial Intelligence Organization.

[82] Azim Akhtarshenas, Mohammad Ali Vahedifar, Navid Ayoobi, Behrouz Maham, Tohid Alizadeh, Sina Ebrahimi, and David López-Pérez, "Federated learning: A cutting-edge survey of the latest advancements and applications," 2024.

[83] Bert Arnrich, Edward Choi, Jason Alan Fries, Matthew BA McDermott, Jungwoo Oh, Tom Pollard, Nigam Shah, Ethan Steinberg, Michael Wornow, and Robin van de Water, "Medical event data standard (meds): Facilitating machine learning for health," in *ICLR 2024 Workshop on Learning from Time Series For Health*, 2024, pp. 03–08.

A. APPENDIX

A.1. Summary of symbols

Patient state is Z_t; actions are A_t updated by Γ_θ; observation and action processes have intensities $\lambda^{\text{obs}}, \lambda^{\text{act}}$; emissions follow g_θ; physiology is controlled by $f_\theta, G_\theta, \Sigma_\theta$; recognition is q_ψ; policy surrogate parameters are ϕ; diffusion head parameters are φ; preference parameters are η; constraints and environment are $\mathcal{C}_t, e$. Loss components are $\mathcal{J}_{\text{recon}}, \mathcal{J}_{\text{inv}}, \mathcal{J}_{\text{cf}}, \mathcal{J}_{\text{mpp}}, \mathcal{J}_{\text{pref}}$, combined into $\mathcal{J}$.

The net effect of these choices is a pretraining regime that pushes the model to separate physiology from policy, to treat missingness as signal rather than nuisance, and to make counterfactual effects salient in the representation. The mathematics of the objective ensures that improvements in likelihood are not purchased by overfitting to observation habits, and the alignment layer provides a gradient target that is explicitly about clinical utility under constraints rather than generic next-event prediction.