Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Digital Twin–Enabled Robust MPC for Personalized Ultrafiltration Scheduling in Hemodialysis: A Proof-of-Concept Study

Rammah Abohtyra

abohtyra_r@utpb.edu

The University of Texas Permian Basin https://www.utpb.edu The University of Texas Permian Basin
https://orcid.org/0000-0003-0945-4570

Method Article

Keywords: Digital Twin, Hemodialysis, Personalized Ultrafiltration, Robust MPC and Uncertainty Tightening

Posted Date: September 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11119867/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Digital Twin–Enabled Robust MPC for Personalized Ultrafiltration Scheduling in Hemodialysis: A Proof-of-Concept Study

Rammah M. Abohtyra[1], Omar A. Beg, and Faisal A. Mohamed

Abstract— **Background: Safe fluid removal during hemodialysis (HD) is challenging when the ultrafiltration rate (UFR) exceeds vascular refilling capacity, reducing circulating blood volume and increasing the risk of intradialytic hypotension (IDH). Methods: The study develops a digital twin (DT)–enabled robust model predictive control (MPC) framework for personalized UFR scheduling. A mechanistic fluid-volume model and unscented Kalman filter estimate absolute blood volume (ABV) and vascular refilling from hematocrit measurements and propagate estimation uncertainty into tightened safety constraints. A residual machine-learning model is evaluated separately for longer-horizon hematocrit prediction. Results: In a clinically parameterized virtual-patient simulation, DT–MPC achieved the prescribed fluid-removal target while maintaining ABV above its critical threshold, whereas constant UFR crossed both hematocrit and ABV safety limits. Residual learning reduced mean 20-min-ahead hematocrit RMSE by 42.4% across 18 held-out trajectories and by 27.9% across 30 stress-test trajectories. Conclusion: The results demonstrate the simulation feasibility of uncertainty-based UFR control and hybrid hematocrit prediction across heterogeneous virtual patients. Multi-patient clinical validation is needed to establish safety, efficacy, and generalizability.**

I. INTRODUCTION

Hemodialysis (HD) is a life-sustaining therapy for patients with kidney failure; however, safe fluid removal remains a major clinical challenge [1]–[4]. Excessive or inadequately prescribed ultrafiltration can cause hemodynamic instability and intradialytic hypotension (IDH), which is associated with increased morbidity, treatment interruptions, and long-term cardiovascular risk [5]–[7].

A key mechanism contributing to IDH is reduced absolute blood volume (ABV) when the ultrafiltration rate (UFR) exceeds vascular refilling capacity. Because ABV and vascular refilling are not routinely measured directly, they must be inferred from indirect measurements [8]. Conventional indicator-dilution methods are unsuitable for continuous monitoring, while online dialysate-dilution methods provide only intermittent estimates and require specialized equipment [9], [10].

Despite advances in dialysis technology, ultrafiltration is often prescribed using fixed or empirically determined profiles that may not account for patient-specific variability or evolving hemodynamic responses. This motivates adaptive strategies for personalized UFR scheduling using patient dynamics and real-time measurements [11]–[13].

Digital twin (DT) technology provides a dynamic virtual representation of a physical system, continuously synchronized with measured data to support monitoring, simulation, prediction, and decision-making [14], [15]. Initially developed for engineering applications, DTs are increasingly being applied in medicine [16]–[18].

In the context of HD, a patient-specific DT can combine mechanistic models with dialysis measurements to estimate unmeasured fluid states, capture intradialytic fluid dynamics, and support prediction and decision-making. In principle, such a framework could enable prospective evaluation of alternative ultrafiltration strategies and adaptive personalization of fluid removal during HD.

This study develops a dialysis-specific DT framework for personalized UFR scheduling that integrates mechanistic modeling, UKF-based state estimation, robust MPC, and residual learning. The primary contribution is a covariance-informed MPC strategy that converts fluid-state estimation uncertainty into constraint tightening for adaptive UFR control while satisfying predefined physiological constraints and the prescribed fluid-removal target. Residual learning further improves hematocrit prediction across heterogeneous virtual patients. The DT framework is evaluated across multiple virtual-patient phenotypes and an out-of-distribution stress cohort, with closed-loop control demonstrated in a proof-of-concept simulation.

The central technical challenge is that ABV and vascular refilling are unmeasured and must be estimated from hematocrit measurements, making the resulting fluid-state estimates inherently uncertain. The key challenge is therefore to translate this estimation uncertainty into sufficiently conservative MPC constraints to maintain the predefined physiological constraints while avoiding excessive constraint tightening that could make the UFR optimization infeasible or prevent the prescribed fluid-removal target from being achieved.

This paper is organized as follows. Section II presents the proposed HD-specific DT framework. Section III reports the prediction and closed-loop control results. Section IV discusses the results, and Section V concludes the paper.

II. HEMODIALYSIS DIGITAL TWIN

As illustrated in Fig. 1, the proposed DT comprises:

A. Data

Historical clinical data were used to identify the patient-specific mechanistic-model parameters used to construct the virtual patient. During the simulated HD session, UFR and hematocrit measurements update the fluid-state estimates and UFR optimization; the residual-learning evaluation uses the virtual-patient cohorts described below.

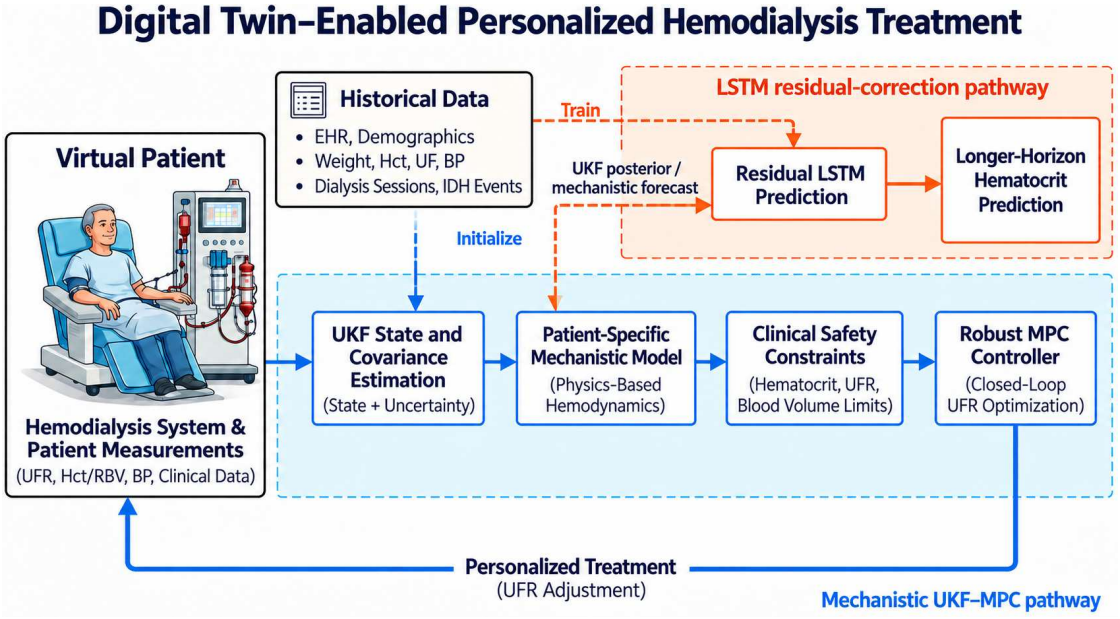


Fig. 1: DT framework for personalized HD. The UKF–MPC pathway estimates fluid states and uncertainty for UFR control within predefined limits, while the LSTM improves longer-horizon hematocrit prediction.

B. Fluid Volume Mechanistic Model

The discrete-time mechanistic model is written as

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, u_k, \boldsymbol{\theta}), \tag{1}$$

$$H_k = h_H(\mathbf{x}_k, \boldsymbol{\theta}) = \frac{V_{\mathrm{rbc}}}{V_{\mathrm{rbc}} + V_{p,k}}, \tag{2}$$

where k is the sampling index, $\mathbf{x}_k = [V_{p,k}, V_{i,k}]^\top$ contains the plasma and interstitial fluid volumes, u_k is the UFR input, and $\boldsymbol{\theta}$ is the patient-specific parameter vector, including the red-blood-cell volume V_{rbc}. The functions $f(\cdot)$ and $h_H(\cdot)$ denote the discrete state-transition and hematocrit-output maps, respectively. Hematocrit (H) is expressed as a percentage. The underlying fluid volume model, its discretization and the parameter-identification procedure were previously developed and evaluated [8], [19]–[23].

C. Parameter Identification

The patient-specific parameters are obtained by solving

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \sum_{k=1}^{N} \left\| H_k^{\mathrm{meas}} - \hat{H}_k^{\mathrm{M}}(\boldsymbol{\theta}) \right\|_{W_H}^2, \tag{3}$$

subject to the state dynamics in (1). Here, H_k^{meas} is the measured H, $\hat{H}_k^{\mathrm{M}}(\boldsymbol{\theta})$ is the corresponding mechanistic model output, N is the number of identification samples, $\boldsymbol{\Theta}$ is the physiologically admissible parameter set, and $\|e\|_{W_H}^2 = e^\top W_H e$ is the weighted squared error [24]. The resulting parameter estimate, $\hat{\boldsymbol{\theta}}$, is fixed during fluid-state estimation and residual learning.

D. UKF State Estimation

For online fluid-state estimation, the unscented Kalman filter (UKF) is used with process and measurement uncertainties modeled as follows [25]:

$$\mathbf{x}_{k+1} = f\left(\mathbf{x}_k, u_k, \hat{\boldsymbol{\theta}}\right) + \mathbf{w}_k, \qquad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \tag{4}$$

$$H_k^{\mathrm{meas}} = h_H\left(\mathbf{x}_k, \hat{\boldsymbol{\theta}}\right) + v_k, \qquad v_k \sim \mathcal{N}(0, R), \tag{5}$$

where $\mathbf{w}_k$ and v_k are zero-mean process and measurement noise, with covariance $\mathbf{Q}$ and variance R, respectively. The UKF propagates the sigma points through $f(\cdot)$ and adds $\mathbf{Q}$ to the predicted covariance. The measurement update is

$$\nu_k = H_k^{\mathrm{meas}} - \hat{H}_{k|k-1}, \tag{6}$$

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}^{xH} S_k^{-1}, \tag{7}$$

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \nu_k, \tag{8}$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k S_k \mathbf{K}_k^\top. \tag{9}$$

Here, $\hat{H}_{k|k-1}$ is the predicted H, ν_k is the innovation, S_k is its variance, $\mathbf{P}_{k|k-1}^{xH}$ is the state–hematocrit cross-covariance, $\mathbf{K}_k$ is the UKF gain, and $\hat{\mathbf{x}}_{k|k}$ and $\mathbf{P}_{k|k}$ are the posterior state mean and covariance.

E. Virtual-Patient UKF–LSTM Prediction

Starting from the UKF posterior state estimate $\hat{\mathbf{x}}_{k|k}$, the mechanistic model is propagated over the planned UFR sequence:

$$\hat{\mathbf{x}}_{k|k}^{\mathrm{M}} = \hat{\mathbf{x}}_{k|k}, \tag{10}$$

$$\hat{\mathbf{x}}_{k+j+1|k}^{\mathrm{M}} = f\left(\hat{\mathbf{x}}_{k+j|k}^{\mathrm{M}}, u_{k+j|k}, \hat{\boldsymbol{\theta}}\right), \tag{11}$$

$$\hat{H}_{k+N_p|k}^{\mathrm{M}} = h_H\left(\hat{\mathbf{x}}_{k+N_p|k}^{\mathrm{M}}, \hat{\boldsymbol{\theta}}\right). \tag{12}$$

The state forecast in (11) is evaluated for $j = 0, \ldots, N_p - 1$. Equations (10)–(12) define the nominal mean forecast; thus, the zero-mean noise terms are omitted, while their effects are retained through the UKF covariance. Here, $k+j|k$ denotes

a prediction for time $k+j$ issued at time k, the superscript M denotes the mechanistic model, and $N_p = 20$ corresponds to a 20-min prediction horizon.

To evaluate prediction under heterogeneous simulated conditions, 120 synthetic virtual-patient trajectories were generated across six physiological, treatment, and measurement phenotypes. Complete trajectories were partitioned into 84 training, 18 validation, and 18 held-out test cases without patient-level overlap. An additional 30-trajectory out-of-distribution (OOD) cohort was used for stress testing without retraining.

These trajectories were used to train and evaluate a residual Long Short-Term Memory (LSTM) learner [26] that corrects persistent discrepancies between the mechanistic forecast and the reference trajectory H. The target residual is defined as

$$r_{k+N_p|k} = H^{\mathrm{ref}}_{k+N_p} - \hat{H}^{\mathrm{M}}_{k+N_p|k}, \tag{13}$$

where $H^{\mathrm{ref}}_{k+N_p}$ is the simulated reference H for the virtual-patient experiments; in a clinical implementation, the corresponding target would be the measured H at the forecast target time. Because this target becomes available only after the prediction horizon has elapsed, it is used only for retrospective training. The sequence-to-one network uses ten input features, a 16-unit LSTM layer, and a dropout layer to predict the future scalar residual. Dropout ($p = 0.30$) was used for regularization [27].

At time k, the LSTM receives the following causal feature vector:

$$\begin{aligned}\boldsymbol{\phi}_k = \big[&\widetilde{H}^{\mathrm{meas}}_k,\ \hat{H}_{k|k},\ \hat{H}^{\mathrm{M}}_{k+N_p|k},\ u_k,\ \bar{u}_{k:k+N_p-1},\ \Delta u_k,\\ &d_k,\ \nu_k,\ r^{\mathrm{mat}}_k,\ \bar{r}^{\mathrm{E}}_k\big]^\top,\end{aligned} \tag{14}$$

where $\widetilde{H}^{\mathrm{meas}}_k$ denotes the available H measurement or its UKF estimate during dropout, $\hat{H}_{k|k}$ is the UKF posterior H, $\bar{u}_{k:k+N_p-1}$ is the mean planned UFR, and $\Delta u_k = u_k - u_{k-1}$. Furthermore, d_k is the binary indicator of measurement-dropout $d_k \in \{0, 1\}$, where $d_k = 1$ denotes a missing measurement and $d_k = 0$ denotes an available measurement, ν_k is the UKF innovation, and r^{mat}_k is the most recently matured forecast residual, $r^{\mathrm{mat}}_k = H^{\mathrm{meas}}_k - \hat{H}^{\mathrm{M}}_{k|k-N_p}$, with $\bar{r}^{\mathrm{E}}_k$ denoting its exponentially weighted average.

Using the causal feature history $\boldsymbol{\phi}_{1:k}$, the LSTM predicts the residual at the future target time:

$$\hat{r}^{\mathrm{LSTM}}_{k+N_p|k} = g_{\boldsymbol{\psi}}(\boldsymbol{\phi}_{1:k}), \tag{15}$$

where $g_\psi(\cdot)$ denotes the trained LSTM mapping with parameters ψ, which are estimated by minimizing the mean-squared residual error over the training trajectories. Validation trajectories are used for early stopping, hyperparameter selection, and prediction-interval calibration, whereas the test trajectories are reserved for final evaluation.

The LSTM correction is causally smoothed and combined with the mechanistic forecast:

$$\tilde{r}_{k+N_p|k} = \beta \hat{r}^{\mathrm{LSTM}}_{k+N_p|k} + (1-\beta)\tilde{r}_{k+N_p-1|k-1}, \tag{16}$$

$$\hat{H}^{\mathrm{DT}}_{k+N_p|k} = \hat{H}^{\mathrm{M}}_{k+N_p|k} + \alpha \tilde{r}_{k+N_p|k}, \tag{17}$$

where $\alpha \in [0, 1]$ and $\beta \in (0, 1]$ are selected using validation data. If residual correction does not improve the validation RMSE, α is set to zero, reverting to the mechanistic predictor.

Using the empirical validation-error quantiles $q_{0.025}$ and $q_{0.975}$, the pointwise validation-calibrated empirical 95% prediction interval is

$$\mathcal{I}^{95\%}_{k+N_p|k} = \Big[\hat{H}^{\mathrm{DT}}_{k+N_p|k} + q_{0.025},\ \hat{H}^{\mathrm{DT}}_{k+N_p|k} + q_{0.975}\Big]. \tag{18}$$

Thus, the UKF synchronizes the mechanistic model with current measurements, while the LSTM corrects systematic long-horizon prediction errors.

The LSTM is evaluated only as a prediction enhancement; it is not used for mechanistic parameter identification or activated in the reported UKF–MPC closed-loop comparison.

F. UKF-Based Robust MPC and Uncertainty Tightening

At each sampling time k, the UKF uses the measured H to obtain the posterior state estimate $\hat{\mathbf{x}}_{k|k}$ and covariance $\mathbf{P}_{k|k}$. Starting from this estimate, the patient-specific mechanistic model is propagated over an N-step prediction horizon:

$$\hat{\mathbf{x}}^{\mathrm{M}}_{k|k} = \hat{\mathbf{x}}_{k|k}, \tag{19}$$

$$\hat{\mathbf{x}}^{\mathrm{M}}_{k+i+1|k} = f\Big(\hat{\mathbf{x}}^{\mathrm{M}}_{k+i|k}, u_{k+i|k}, \hat{\boldsymbol{\theta}}\Big). \tag{20}$$

Equation (20) is evaluated for $i = 0, \ldots, N-1$, where $\hat{\mathbf{x}}^{\mathrm{M}}_{k+i|k} = [\hat{V}_{p,k+i|k}, \hat{V}_{i,k+i|k}]^\top$ is the model state prediction, $u_{k+i|k}$ is the candidate UFR, and $\hat{\boldsymbol{\theta}}$ is the identified patient-specific parameter vector. The notation $k+i|k$ denotes a prediction for time $k+i$ issued at time k.

The corresponding H and ABV predictions are

$$\hat{H}^{\mathrm{M}}_{k+i|k} = \frac{\hat{V}_{\mathrm{rbc}}}{\hat{V}_{\mathrm{rbc}} + \hat{V}_{p,k+i|k}}, \tag{21}$$

$$\hat{V}_{b,k+i|k} = \hat{V}_{\mathrm{rbc}} + \hat{V}_{p,k+i|k} = \frac{\hat{V}_{\mathrm{rbc}}}{\hat{H}^{\mathrm{M}}_{k+i|k}}, \tag{22}$$

where H is expressed in percent and $\hat{V}_{\mathrm{rbc}}$ is the identified red-blood-cell volume. Since $V_b = V_{\mathrm{rbc}}/H$, the upper H limit implies the corresponding lower ABV limit.

The MPC decision vector is

$$\mathbf{U}_k = \begin{bmatrix} u_{k|k} & u_{k+1|k} & \cdots & u_{k+N-1|k}\end{bmatrix}^\top \in \mathbb{R}^N. \tag{23}$$

At each sampling time, the controller solves a constrained nonlinear MPC problem with a quadratic control objective and uncertainty-tightened physiological constraints:

$$\begin{aligned}\underset{\mathbf{U}_k}{\text{minimize}}\quad & \frac{1}{2}\mathbf{U}_k^\top \mathbf{R}_U \mathbf{U}_k - \lambda_v T_s \mathbf{1}_N^\top \mathbf{U}_k\\ \text{subject to}\quad & \hat{\mathbf{x}}^{\mathrm{M}}_{k|k} = \hat{\mathbf{x}}_{k|k},\\ & \hat{\mathbf{x}}^{\mathrm{M}}_{k+i+1|k} = f\Big(\hat{\mathbf{x}}^{\mathrm{M}}_{k+i|k}, u_{k+i|k}, \hat{\boldsymbol{\theta}}\Big),\\ & \hat{H}^{\mathrm{M}}_{k+i|k} + \gamma\, \sigma^{H}_{k+i|k} \le H_{\mathrm{safe}},\\ & \hat{H}^{\mathrm{M}}_{k+i+1|k} - \hat{H}^{\mathrm{M}}_{k+i|k} \le \Delta H_{\max},\\ & u_{\min} \le u_{k+i|k} \le u_{\max},\\ & T_s \mathbf{1}_N^\top \mathbf{U}_k = V^{\mathrm{tar}}_{\mathrm{UF},k}.\end{aligned} \tag{24}$$

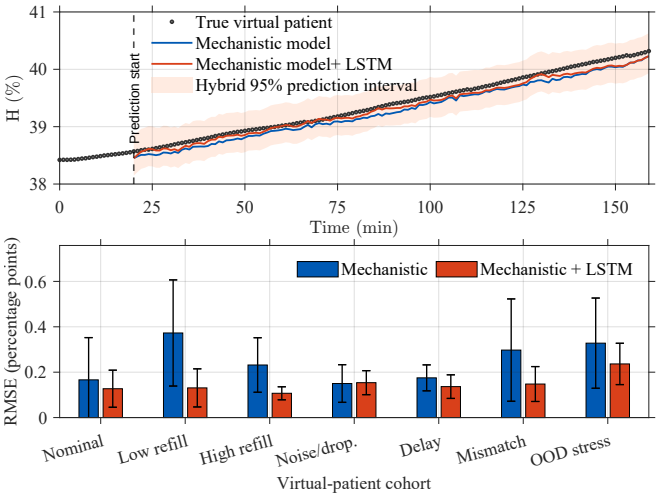


Fig. 2: Twenty-minute-ahead H prediction: (a) held-out trajectory with 95% interval and (b) RMSE across held-out phenotypes and the OOD stress cohort (mean $\pm$ SD).

The state, safety, and input constraints are imposed for $i = 0, \dots, N-1$, while the hematocrit-rate constraint is imposed for $i = 0, \dots, N-2$. Here, T_s is the sampling period, H_{safe} is the maximum admissible H, $\sigma^H_{k+i|k}$ is the predicted uncertainty (standard deviation) of H, and γ is the uncertainty-margin multiplier, $\Delta H_{\max}$ is the maximum permitted H increase per sampling interval, and $[u_{\min}, u_{\max}]$ defines the prescribed UFR bounds used in the simulation. The quantity $V^{\text{tar}}_{\text{UF},k}$ is the prescribed fluid volume to be removed over the optimization horizon. The input-weighting matrix is

$$\mathbf{R}_U = \lambda_{\Delta u}\mathbf{D}^\top\mathbf{D} + \lambda_u \mathbf{I}_N, \tag{25}$$

where $\lambda_{\Delta u} \geq 0$ penalizes UFR variation, $\lambda_u \geq 0$ penalizes UFR magnitude, and $\lambda_v \geq 0$ rewards fluid removal. The first-difference matrix is

$$\mathbf{D} = \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & -1 & 1 \end{bmatrix} \in \mathbb{R}^{(N-1)\times N}, \tag{26}$$

such that $\|\mathbf{D}\mathbf{U}_k\|_2^2$ penalizes changes between successive UFR values. The first move can also be linked to the previously applied UFR by adding the penalty $\lambda_{\Delta u}(u_{k|k} - u_{k-1})^2$.

The UKF covariance is propagated along the predicted trajectory and mapped to H and ABV uncertainty. For $\mathbf{x} = [V_p, V_i]^\top$, the corresponding output-sensitivity vectors are

$$\mathbf{g}_{H,k+i|k} = \left.\frac{\partial h_H(\mathbf{x}, \hat{\boldsymbol{\theta}})}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}^{\text{M}}_{k+i|k}}, \qquad \mathbf{g}_{V_b} = \frac{\partial V_b}{\partial \mathbf{x}} = \begin{bmatrix} 1 & 0 \end{bmatrix}^\top. \tag{27}$$

The predicted standard deviations are approximated by

$$\sigma^H_{k+i|k} = \sqrt{\mathbf{g}^\top_{H,k+i|k}\mathbf{P}_{k+i|k}\mathbf{g}_{H,k+i|k}}, \tag{28}$$

$$\sigma^{V_b}_{k+i|k} = \sqrt{\mathbf{g}^\top_{V_b}\mathbf{P}_{k+i|k}\mathbf{g}_{V_b}}. \tag{29}$$

Accordingly, the uncertainty bounds are

$$\underline{V}_{b,k+i|k} = \hat{V}_{b,k+i|k} - \gamma\sigma^{V_b}_{k+i|k},$$
$$\overline{V}_{b,k+i|k} = \hat{V}_{b,k+i|k} + \gamma\sigma^{V_b}_{k+i|k}, \tag{30}$$

$$\underline{H}_{k+i|k} = \hat{H}^{\text{M}}_{k+i|k} - \gamma\sigma^H_{k+i|k},$$
$$\overline{H}_{k+i|k} = \hat{H}^{\text{M}}_{k+i|k} + \gamma\sigma^H_{k+i|k}. \tag{31}$$

Therefore, the safety constraint in (24) is imposed on the upper H bound, $\overline{H}_{k+i|k} \leq H_{\text{safe}}$. For $\gamma = 2$, these bounds approximate a pointwise 95% interval under Gaussian assumptions. Thus, increasing UKF uncertainty directly tightens the admissible H region and produces more conservative UFR actions.

Only the first optimized input,

$$u_k = u^*_{k|k}, \tag{32}$$

is applied to the patient model. At the next sampling time, the UKF assimilates the new measurement, and the estimation and optimization steps are repeated in a receding-horizon manner.

TABLE I: Identified parameters $\hat{\theta}$.

$V_p(0)$ (L)	$V_i(0)$ (L)	V_{rbc} (L)	K_f (L/min·mmHg)	d_1	d_2
3.25	13.75	2	0.0057	0.0096	1.46

TABLE II: 20-min-ahead H prediction on held-out virtual-patient trajectories.

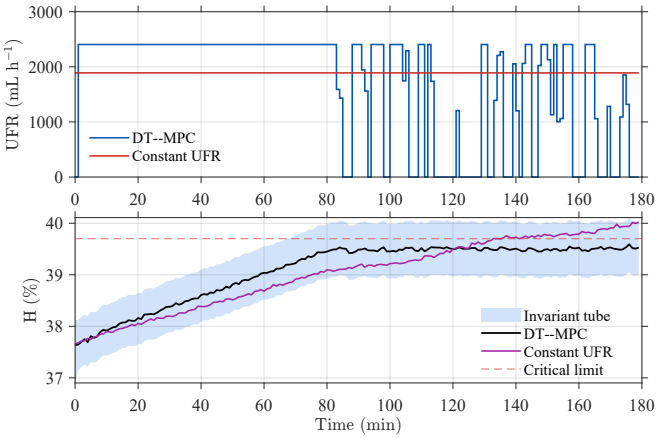
Model	Mean RMSE $\pm$ SD (%)	Median RMSE (%)	Mean MAE (%)
Mechanistic model	0.232 $\pm$ 0.162	0.214	0.219
Mechanistic model + LSTM	0.134 $\pm$ 0.057	0.111	0.110

95% interval coverage: 96.9%; mean width: 0.698 percentage points. Calibrated using validation trajectories only.

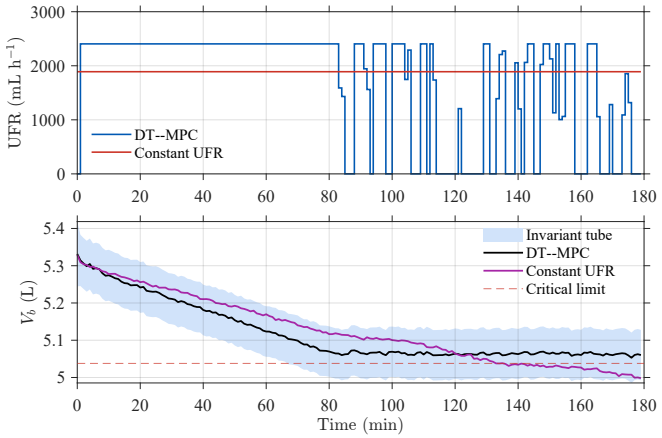
III. Results

The patient-specific mechanistic model was initialized using the identified parameters listed in Table I [23]. The estimated initial plasma and interstitial volumes were 3.25 and 13.75 L, respectively, and the red-blood-cell volume was 2.00 L. The remaining identified parameters were $K_f = 0.0057$ L/(min·mmHg), $d_1 = 0.0096$, and $d_2 = 1.46$. These parameters were fixed during the subsequent UKF state estimation, LSTM residual learning, and closed-loop simulations. All model parameters are reported in [8].

The residual LSTM improved 20-min-ahead H prediction across the held-out synthetic virtual-patient test trajectories. As summarized in Table II, mean RMSE decreased from 0.232 ± 0.162 to 0.134 ± 0.057 percentage points, a 42.4% reduction; median RMSE decreased from 0.214 to 0.111, and mean MAE from 0.219 to 0.110 percentage points. The validation-calibrated empirical 95% prediction interval achieved 96.9% test coverage with a mean width of 0.698 percentage points. Across the six test phenotypes, 13 of 18 trajectories improved; in the additional 30-trajectory OOD stress cohort, mean RMSE decreased from 0.328 to 0.236



(a) H trajectories under closed-loop DT–MPC and standard constant-UFR profiles.



(b) ABV trajectories and the corresponding safety threshold under closed-loop DT–MPC and standard constant UFR profiles.

Fig. 3: Closed-loop DT–MPC and constant-UFR comparison: (a) UFR and H trajectories and (b) UFR and ABV trajectories. Shaded regions represent the UKF-derived uncertainty tubes, and dashed lines indicate the physiological safety limits.

percentage points (27.9%), with 18 of 30 trajectories improved. For the representative held-out trajectory in Fig. 2, the LSTM correction reduced the mechanistic RMSE from 0.217 to 0.063 percentage points. It also reduced the persistent forecast bias while preserving the overall mechanistic trend. Because the forecasts are displayed at their target times, the first 20-min-ahead prediction appears at $t = 20$ min, although the simulation begins at $t = 0$.

The closed-loop simulation was conducted over 180 min using a virtual patient whose mechanistic parameters were identified from clinical data [8]. The simulation prescribed a 5.67 L ultrafiltration target over 3 h for a 140 kg virtual patient (13.5 mL/kg/h). The constant-UFR strategy therefore delivered 1.89 L/h. DT–MPC achieved the target while enforcing the UFR and H constraints. As shown in Fig. 3, DT–MPC initially applied a higher UFR and subsequently introduced reductions and temporary interruptions as the trajectory approached the safety region. The DT–MPC trajectory remained within the prescribed H and ABV limits, whereas constant UFR crossed the corresponding critical limits near the end of treatment.

The optimized UFR profile altered the temporal balance between ultrafiltration and vascular refilling. Figure 4 shows greater accumulated vascular refilling under DT–MPC during much of the treatment, followed by a lower terminal accumulated refilling volume than with constant UFR. Despite the same prescribed fluid-removal target, DT–MPC maintained the H and ABV trajectories within their critical limits, whereas constant UFR did not. Together, these results demonstrate the algorithmic feasibility of the proposed DT framework.

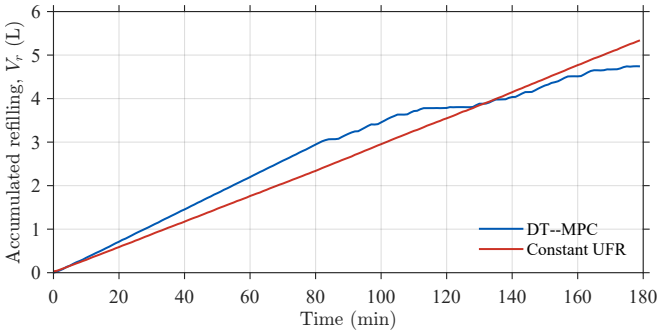


Fig. 4: Accumulated vascular refilling volume (V_r) under DT–MPC and constant-UFR profiles for the same fluid-removal target. DT–MPC yields greater cumulative refilling during most of the session, whereas constant UFR achieves a higher terminal accumulated refilling volume.

IV. Discussion

This study demonstrates the feasibility of a dialysis-specific DT integrating state estimation, prediction, and constrained UFR optimization. Its main contribution is using UKF covariance to tighten H constraints, directly linking blood-volume uncertainty to UFR decisions. Over the 180-min session, DT–MPC achieved the prescribed 5.67-L fluid-removal target while maintaining H and ABV within their critical limits. In contrast, the constant-UFR profile, prescribed to remove the same volume at 1.89 L/h, crossed the critical H and ABV limits near the end of treatment.

The optimized UFR profile initially increased fluid removal and then reduced or temporarily interrupted ultrafiltration near the safety region. Accumulated refilling was greater under DT–MPC during much of the session but lower at treatment completion than under constant UFR. Thus, the relevant control benefit is not greater terminal refilling volume, but redistribution of ultrafiltration over time to achieve the prescribed removal target while maintaining the H and ABV constraints.

The residual-learning results show improved longer-horizon prediction across heterogeneous simulated conditions. Across 18 held-out test trajectories, the hybrid predictor reduced mean 20-min-ahead hematocrit RMSE from

0.232 to 0.134 percentage points (42.4%), with improvement in 13 trajectories. The largest gains occurred under altered refilling and strong model mismatch, whereas performance was essentially unchanged under noisy/dropout sensing. Without retraining, the hybrid predictor retained a 27.9% aggregate RMSE reduction in the 30-trajectory OOD stress cohort, with 18 trajectories improved (Fig. 2b). The LSTM was evaluated solely as a prediction enhancement and was not activated in the closed-loop comparison; therefore, the control results arise from the mechanistic UKF–MPC pathway.

Several limitations should be recognized. Phenotype-stratified held-out testing and a separate OOD stress cohort provide evidence of robustness within the simulated environment but cannot reproduce the full structural mismatch and measurement variability encountered clinically. Performance also depends on the prescribed safety threshold, noise models, and prediction-interval calibration. Furthermore, H and ABV do not represent all IDH mechanisms, including vascular resistance, cardiac compensation, and blood-pressure responses. Independent multi-patient clinical validation and closed-loop integration of the residual predictor remain necessary.

V. Conclusion

This study developed a dialysis-specific DT framework integrating UKF-based fluid-state estimation, covariance-informed MPC, and residual learning for hematocrit prediction. In a closed-loop simulation, DT–MPC achieved the prescribed fluid-removal target while maintaining hematocrit and ABV within their critical limits, whereas constant UFR crossed both safety limits near the end of treatment. Across 18 held-out virtual-patient trajectories, residual learning reduced mean 20-min-ahead hematocrit RMSE by 42.4%, with a 27.9% reduction under OOD stress testing. These results demonstrate algorithmic feasibility under heterogeneous simulated conditions; independent multi-patient clinical validation is needed to establish clinical safety and benefit.

References

[1] W. Wathanavasin, C. Thongprayoon, P. W. Davis, and W. Cheungpasitporn, "The efficacy and safety of technology-guided dry weight adjustment among dialysis patients: A meta-analysis of randomized controlled trials," *Kidney Medicine*, vol. 7, no. 8, p. 101052, 2025.

[2] S. Patel, S. Ashokkumar, and A. Green, "Modern tools for optimizing fluid management in dialysis patients: A comprehensive review," *BMC Nephrology*, vol. 25, no. 1, p. 464, 2024.

[3] S. Hamada, M. Hata, S. Furukawa, and S. Yamamoto, "Effectiveness of blood volume change-guided ultrafiltration control (BV-UFC) in hemodialysis: A crossover comparative study," *Clinical Kidney Journal*, vol. 18, no. 5, p. sfaf141, 2025.

[4] M. Bossola, I. Mariani, C. T. Strizzi, C. P. Piccinni, and E. D. Stasio, "How to limit interdialytic weight gain in patients on maintenance hemodialysis: State of the art and perspectives," *Journal of Clinical Medicine*, vol. 14, no. 6, p. 1846, 2025.

[5] J. Yu, X. Chen, Y. Li, Y. Wang, Z. Liu, B. Shen, J. Teng, J. Zou, and X. Ding, "High ultrafiltration rate induced intradialytic hypotension is a predictor for cardiac remodeling: A 5-year cohort study," *Renal Failure*, vol. 43, no. 1, pp. 40–48, 2021.

[6] J. A. Chou, E. Streja, D. V. Nguyen, C. M. Rhee, Y. Obi, J. K. Inrig, A. Amin, C. P. Kovesdy, J. J. Sim, and K. Kalantar-Zadeh, "Intradialytic hypotension, blood pressure changes and mortality risk in incident hemodialysis patients," *Nephrology Dialysis Transplantation*, vol. 33, no. 1, pp. 149–159, 2018.

[7] A. Douvris, K. Zeid, S. Hiremath, S. M. Bagshaw, R. Wald, W. Beaubien-Souligny, J. Kong, C. Ronco, and E. G. Clark, "Mechanisms for hemodynamic instability related to renal replacement therapy: A narrative review," *Intensive Care Medicine*, vol. 45, no. 10, pp. 1333–1346, 2019.

[8] R. Abohtyra, T. Vincent, and D. Schneditz, "Magnitude and precision of absolute blood volume estimated during hemodialysis," *Renal Failure*, vol. 46, no. 2, p. 2377781, 2024.

[9] S. Thijssen, F. Kappel, and P. Kotanko, "Absolute blood volume in hemodialysis patients: Why is it relevant, and how to measure it?" *Blood Purification*, vol. 35, no. 1–3, pp. 63–71, 2013.

[10] J. Kron, D. Schneditz, T. Leimbach, S. Aign, and S. Kron, "A simple and feasible method to determine absolute blood volume in hemodialysis patients in clinical practice," *Blood Purification*, vol. 38, no. 3–4, pp. 180–187, 2015.

[11] J. E. Flythe, S. E. Kimmel, and S. M. Brunelli, "Rapid fluid removal during dialysis is associated with cardiovascular morbidity and mortality," *Kidney International*, vol. 79, no. 2, pp. 250–257, 2011.

[12] M. M. Assimon and J. E. Flythe, "Rapid ultrafiltration rates and outcomes among hemodialysis patients: Re-examining the evidence base," *Current Opinion in Nephrology and Hypertension*, vol. 24, no. 6, pp. 525–530, 2015.

[13] A. Davenport, "Intradialytic complications during hemodialysis," *Hemodialysis International*, vol. 10, no. 2, pp. 162–167, 2006.

[14] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," *IEEE Access*, vol. 8, pp. 108 952–108 971, 2020.

[15] D. Jones, C. Snider, A. Nassehi, J. Yon, and B. Hicks, "Characterising the digital twin: A systematic literature review," *CIRP Journal of Manufacturing Science and Technology*, vol. 29, pp. 36–52, 2020.

[16] K. P. Venkatesh, M. M. Raza, and J. C. Kvedar, "Health digital twins as tools for precision medicine: Considerations for computation, implementation, and regulation," *npj Digital Medicine*, vol. 5, no. 1, p. 150, 2022.

[17] T. Sun, X. He, and Z. Li, "Digital twin in healthcare: Recent updates and challenges," *Digital Health*, vol. 9, p. 20552076221149651, 2023.

[18] E. Katsoulakis, Q. Wang, H. Wu, L. Shahriyari, R. Fletcher, J. Liu, L. Achenie, H. Liu, P. Jackson, Y. Xiao, T. Syeda-Mahmood, R. Tuli, and J. Deng, "Digital twins for health: A scoping review," *npj Digital Medicine*, vol. 7, no. 1, p. 77, 2024.

[19] R. M. Abohtyra, C. Hollot, J. Horowitz, M. Germain, and Y. Chait, "Designing robust ultrafiltration rate profiles based on identifying fluid volume model parameters during hemodialysis," in *Dynamic Systems and Control Conference*, vol. 58271. American Society of Mechanical Engineers, 2017, p. V001T08A006.

[20] R. Abohtyra, Y. Chait, M. J. Germain, C. V. Hollot, and J. Horowitz, "Individualization of ultrafiltration in hemodialysis," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 8, pp. 2174–2181, 2019.

[21] R. M. Abohtyra and Y. Chait, "New algorithm to design real time optimal and robust ultrafiltration rates in chronic kidney disease to prevent cardiovascular morbidity and mortality," in *Dynamic Systems and Control Conference*, vol. 51890. American Society of Mechanical Engineers, 2018, p. V001T11A005.

[22] R. Abohtyra, C. Hollot, M. G. Germain, Y. Chait, and J. Horowitz, "Personalized ultrafiltration profiles to minimize intradialytic hypotension in end-stage renal disease," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 309–314.

[23] R. M. Abohtyra and O. A. Beg, "Body fluid estimation during standard ultrafiltration in chronic kidney disease," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 4659–4664.

[24] R. M. Abohtyra, C. L. Chan, D. J. Albers, and B. J. Gluckman, "Inferring insulin secretion rate from sparse patient glucose and insulin measures," *Frontiers in Physiology*, vol. 13, p. 893862, 2022.

[25] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.

[26] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[27] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.