MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Bayesian Forecasting with Deep Generative Disturbance Models in Stochastic MPC for Building Energy Systems

Sorouifar, Farshud; Paulson, Joel A.; Wang, Ye; Quirynen, Rien; Laughman, Christopher R.; Chakrabarty, Ankush

TR2024-110 August 22, 2024

Abstract

Predictive control strongly depends on the quality of disturbance predictions. While recent disturbance modeling efforts have adopted a probabilistic perspective to protect against unreliable deterministic predictions, such probabilistic models are often applicable only in data-rich settings or involve making simplifying assumptions on the underlying distributions. Generative models, such as conditional variational autoencoders (CVAEs), provide an expressive and automated approach for learning distributions from data. By sampling the learned latent space, one can generate unseen disturbance realizations. In this paper, we develop methods to leverage these generative models for the design of economic stochastic model predictive control (SMPC) that utilizes disturbance signals generated by a CVAE for online adaptation. Scenarios generated by the CVAE can be transformed to conditional probabilities on learned latent vectors, wherein the conditioning is with respect to factors that affect the disturbance signal shape itself (e.g., effect of workday/weekend on internal heat loads) along with the observed data (i.e., how likely the latent is, based on the observed data). We can consequently generate the most relevant disturbance signals for use in a scenario tree-based SMPC approach to reduce conservativeness of the control policy while satisfying constraints.

IEEE Conference on Control Technology and Applications (CCTA) 2024

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Bayesian Forecasting with Deep Generative Disturbance Models in Stochastic MPC for Building Energy Systems

Farshud Sorouifar[2†], Joel A. Paulson[2],
Ye Wang[1], Rien Quirynen[1], Christopher R. Laughman[1], and Ankush Chakrabarty[1*]

***Abstract*—Predictive control strongly depends on the quality of disturbance predictions. While recent disturbance modeling efforts have adopted a probabilistic perspective to protect against unreliable deterministic predictions, such probabilistic models are often applicable only in data-rich settings or involve making simplifying assumptions on the underlying distributions. Generative models, such as conditional variational autoencoders (CVAEs), provide an expressive and automated approach for learning distributions from data. By sampling the learned latent space, one can generate unseen disturbance realizations. In this paper, we develop methods to leverage these generative models for the design of economic stochastic model predictive control (SMPC) that utilizes disturbance signals generated by a CVAE for online adaptation. Scenarios generated by the CVAE can be transformed to conditional probabilities on learned latent vectors, wherein the conditioning is with respect to factors that affect the disturbance signal shape itself (e.g., effect of workday/weekend on internal heat loads) along with the observed data (i.e., how likely the latent is, based on the observed data). We can consequently generate the most relevant disturbance signals for use in a scenario tree-based SMPC approach to reduce conservativeness of the control policy while satisfying constraints.**

***Index Terms*—Building energy management, deep learning, learning for control, adaptive control, model predictive control.**

I. Introduction

Model-based predictive control (MPC) has demonstrated potential for constrained control of next-generation engineering systems due to their ability to leverage disturbance forecasts and take anticipatory actions that reduce operational costs or optimize user-specified performance objectives [1]–[3]. MPC involves iteratively solving of an open-loop constrained optimal control problem (OCP) at a given time to obtain a trajectory of control actions based on model-based predictions, implementing the first control action on the system [1], and then repeating this optimization step at the next control step. The solution of this OCP allows the controller to make anticipatory actions using model-based predictions of the system behavior and exogenous disturbance inputs acting on the system, while re-solving at each iteration enables flexibility by accounting for feedback. The control actions obtained by solving this OCP depend strongly on the quality of predictions made by both the disturbance forecasting model and dynamical models.

Given that any disturbance model will not be perfectly accurate and that the internal dynamical model will contain unmodeled components, controllers such as stochastic MPC (SMPC) that account for uncertainty have demonstrated success on applications ranging from building HVAC systems [3]–[6] to chemical processes [7]–[9], in terms of both closed-loop performance and constraint enforcement. However, most SMPC formulations assume that the disturbances are generated by simplified stochastic processes that may not reflect the observed data [10], [11]. While these simplifications may result in tractable controller design, they often deteriorate performance for disturbance inputs that cannot be represented by simple disturbance models. Our work is motivated by occupant-centric buildings, where, for example, energy use in building zones depend on human occupancy patterns and appliance use, both of which cannot be represented by simple stochastic processes as they are affected by multiple complex components in a building system, and human behavior does not conform to first-principles modeling. Even the utility of non-homogeneous Markov models is limited because most disturbance signals are continuous not discrete, and equipping these models with domain-informed conditional inputs requires storing large numbers of transition matrices.

Deep probabilistic networks offer a scalable and automated framework for learning complex conditional distributions directly from data. Two common approaches include the use of generative networks that construct disturbance signals over a fixed time horizon, or forecasting networks that take a previous time window of disturbances as inputs and predict the disturbances over a future time window. While both classes of models have been demonstrated to be accurate [12]–[14], we have empirically noted that the forecasting approach performs better for data-rich applications but often exhibits low-variance predictions when data is limited and conditional inputs are plentiful, possibly due to more complex relationships arising from the recursive model structure, although further work is needed to fully understand the data requirements for these two methods. In comparison, generative models appear to perform well even in data-limited settings, *e.g.*, when designing controllers for buildings that have recently been constructed or have limited archival data, and for more non-standard distributions [15]. While generative models have been used for scenario generation [12], [15] and in the closed-

[1]Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA.
[2]Department of Chemical and Biomolecular Engineering, The Ohio State University, Columbus, OH 43210, USA.
[†]This work was completed during F. Sorourifar's internship at MERL.
[*]Corresponding author. Phone: +1 (617) 758-6175. Fax: +1 (617)-621-7550. Email: achakrabarty@ieee.org.

loop verification of control policies [13], there is no extant work on leveraging generative models explicitly for controller *design*.

In this paper, we propose τ-SMPC, a **T**rajectory **A**dapting **U**ncertainty (TAU$\equiv \tau$) framework which endows scenario-tree based SMPC with the property of adapting the uncertainty distributions based on measured disturbances. Our SMPC is informed by a conditional variational autoencoder (CVAE) whose latent space can be sampled to give realistic and out-of-sample exogenous disturbance scenarios. The primary novelty of this work is in developing a strategy to generate disturbance samples based on a partially revealed disturbance sequence and a learned prior, to generate more accurate scenarios that improve control performance in practice. We show via simulation experiments on a simplified building thermal dynamics model that τ-SMPC can reduce operational costs and does a better job at handling constraints by providing more accurate disturbance forecasts by learning approximate conditional distributions from data.

II. Problem Statement

We consider a discrete-time dynamic system of the form

$$x_{k+1} = f(x_k, u_k, w_k), \tag{1}$$

where $x_k \in \mathbb{R}^{n_x}$ denotes the current state of the system, $u_k \in \mathbb{R}^{n_u}$ denotes a set of control inputs, and $w_k \in \mathbb{R}^{n_w}$ denotes the exogenous disturbances which affect the system at the current time k. The control inputs are computed using a two-stage scenario-tree SMPC framework, which is cast as the optimization problem

$$\min_{u_{0|t},\dots,u_{n_p-1|t}} \quad \sum_{s=1}^{n_s} \omega^{(s)} \sum_{k=0}^{n_p} \ell_{k|t}\left(\hat{x}^{(s)}_{k|t}, u^{(s)}_{k|t}, \hat{w}^{(s)}_{k|t}\right), \tag{2a}$$

$$\text{s.t.} \quad \hat{x}^{(s)}_{0|t} = x_t, \tag{2b}$$

$$\hat{x}^{(s)}_{k+1|t} = f\left(\hat{x}^{(s)}_{k|t}, u^{(s)}_{k|t}, \hat{w}^{(s)}_{k|t}\right), \tag{2c}$$

$$\left(\hat{x}^{(s)}_{k|t}, u^{(s)}_{k|t}\right) \in \mathcal{Z}, \tag{2d}$$

$$\forall k = 0, \cdots, n_p - 1, \quad \forall s = 1, \cdots, n_s,$$

where $\omega^{(s)} > 0$ denotes the likelihood weight for scenario s, $\hat{x}^{(s)}_{k|t}$, $u^{(s)}_{k|t}$, and $\hat{w}^{(s)}_{k|t}$ denote the predicted state, input, and disturbance value k steps ahead of current time step t for scenario s, n_p is the prediction horizon, $\mathcal{Z}$ denotes the set of joint state-input constraints, and $\ell_{k|t}(\cdot)$ denotes the stage cost function evaluated k steps ahead of current time t. The SMPC policy can be represented generally by

$$u^*_t = \pi_{\mathsf{SMPC}}(x_t, \hat{W}_t), \text{ where } \hat{W}_t := \{\hat{w}^{(s)}_{0|t}, \dots, \hat{w}^{(s)}_{n_p|t}\}_{s=1}^{n_s}$$

is a disturbance sequence forecast. The implication is that the quality of the control input for any given state depends entirely on the accuracy of the predicted disturbances. Scenarios are generated by considering several disturbance realizations, weighted by their respective probabilities. For example, a common strategy is to uses three branches at each state, where one branch is the expected outcome, and the other two branches are lower probability realizations that bound the expectation. As a result, SMPC frameworks assume a known probability distribution function (PDF).

In this paper, we discuss how probabilistic deep generative networks can learn complex distributions from real data whose mean and variance vary across time-series, thus providing disturbances which can effectively be used for constructing and forecasting SMPC scenarios required to obtain an effective SMPC policy. For example, in building energy systems (which motivates this work), deep generative networks can model internal heat loads generated due to appliance use, hot water and ventilation, all of which are difficult to model accurately with first-principles models because they depend on human occupancy: a signal that is very challenging to forecast based on 'physics'. Instead, we can learn time-series distributions of heat loads from real building usage data and construct synthetic scenarios along with assigned probability weights indicating how likely they are to occur: this information can be critical to building performance management via SMPC.

We also recognize that when disturbance signals are strongly time-correlated, generating scenarios without accounting for past observations will result in overly conservative scenarios in the best cases, and failure to obey constraints in the worst cases. An example in building energy modeling is the ambient weather conditions: past observations can inform good future predictions of ambient temperature throughout the day. Ignoring the past two hours (say) of temperatures can be detrimental when forecasting temperature profile for the rest of the day. In such cases, to encourage our forecast to be cognizant of past disturbance signal values, we design an adaptive forecasting strategy that uses a generative model to approximate a predictive posterior, conditioned on past observations.

III. Methodology

Conditional Variational Autoencoders (CVAEs): We use a CVAE [16] to generate the disturbance sequence $W \equiv W_{0:T} := (w_0, \dots, w_T)$ for a given day (e.g., $T = 24$ for an hourly disturbance profile), conditioned on an environmental variable $c \in [0,1]^{n_c}$, which captures the external factors (e.g., season, workday, etc.). The CVAE consists of an encoder that compresses W, given c, to a latent representation $z \in \mathbb{R}^{n_z}$, and a decoder that reconstructs W, given c, from the learned latent representation; see Fig. 1. The decoder models the conditional distribution $\pi(W|z, c)$, and together with the assumed isotropic Gaussian prior distribution of the latent representation $\pi(z)$, implicitly (and intractably) models the conditional data distribution through $\pi(W|c) = \int \pi(W|z, c)\pi(z)dz$. The encoder models an approximation of the posterior distribution $\pi(z|W, c)$ that is consistent with the generative model distribution implied by the decoder. The CVAE, parameterized by θ, is trained by optimizing the standard evidence-based lower bound (ELBO) loss function; c.f. [12], [16] for details.

Forecasting with a CVAE: For building control, which motivates this work, at a given time $t \in \mathbb{N}$ we have partial sequences of measured disturbances $W_{0:t} \triangleq (w_0, \dots, w_t)$, and

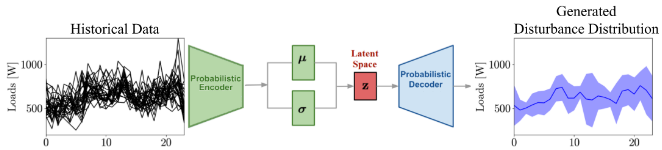

Fig. 1: Structure of the CVAE trained to generate disturbance sequences from a latent space depending on the context (environmental variable).

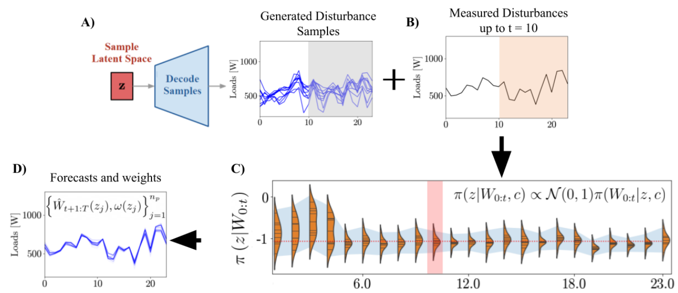

Fig. 2: Overview of the proposed $\tau-$SMPC framework.

will leverage a trained CVAE model to provide a forecasts of the remaining unrealized sequence $W_{t+1:T} \triangleq (w_{t+1}, ..., w_T)$, along with quantified uncertainty, e.g. confidence intervals or standard deviations. More formally, our aim is to extract and sample the conditional distribution $\pi(W_{t+1:T}|W_{0:t}, c)$ from the model learned by the CVAE. However, this specific conditional dependency structure is not directly provided by the CVAE. Instead, our approach aims to first extract the latent representation distribution conditioned on only the partially revealed perturbations, $\pi(z|W_{0:t}, c)$. Our aim is to try to ascertain which latent vector most likely resulted in the revealed perturbations. Then, from sampled representations $z \sim \pi(z|W_{0:t}, c)$, we apply the CVAE decoder model $\pi(W|z, c)$ to sample the corresponding completed sequences, including the unseen portions $W_{t+1:T}$.

The latent posterior $\pi(z|W_{0:t}, c)$, although not known in closed form, can be evaluated using Bayes' rule,

$$\pi(z|W_{0:t}, c) \propto \mathcal{N}(0, I)\, \pi(W_{0:t}|z, c). \tag{3}$$

using a Gaussian prior $\pi(z|c) = \mathcal{N}(0, I)$ and likelihood $\pi(W_{0:t}|z, c)$, noting that the marginal $\pi(W_{0:t}|c)$ is a constant.

Although $\pi(W_{0:t}|z, c)$ is also not available in closed-form, it can be numerically approximated. First, we note that $\pi(W|z, c)$ is defined by the decoder, which uses the learned distribution and reparameterization trick to generate samples $\hat{W}$. By assuming a Gaussian prior, we can numerically evaluate the conditional probability of latent samples using

$$\pi(W_{0:t}|z, c) = \frac{1}{\beta} \exp\left[\frac{-\delta_M\left(\mu_{\theta,0:t}, \Sigma_{\theta,0:t}^{-1}\right)^2}{2}\right], \tag{4}$$

where δ_M is the Mahalanobis distance, and β is the pre-exponential factor for a multivariate Gaussian. We can now generate forecasts $\hat{W}_{t+1:T}$ and compute its respective probability by jointly and normally sampling the decoder and probability function. We present this conditional forecast graphically in Fig. 2. In the following section, we discuss how these forecasts can be used in a scenario-tree SMPC.

Note that a tight (small variance) prior distribution may prevent the forecaster from producing a scenarios that are consistent with the evidence, which is problematic when generating scenarios that are used to enforce constraints. To alleviate this issue, a tuning parameter $\alpha \in [0, 1]$ can be used to loosen the variance of the prior, such that (3) becomes

$$\pi(z|W_{0:t}, c) \propto \pi(z)^{\alpha} \pi(W_{0:t}|z, c).$$

In the case study, we use $\alpha = 0$, due to the strong short term temporal correlations in weather patterns that make a seasonal averages (i.e., the prior) effectively inconsequential when forecasting over a given day. Furthermore, this likelihood based prediction facilitates the downstream task of robust control by encouraging outlier sampling. That is, we can determine an SMPC policy that is cognizant of low-probability events that could prove to have a debilitating effect on the application; e.g., the 2021 cold snap in Texas.

Forecast-based Scenario Tree: Given the conditional density $\pi(z|W_{0:t}, c)$ and disturbance forecasts $\hat{W}_{t+1:T}(z, c)$, we require a scenario selection strategy to select the forecasts to use in a scenario-tree MPC. Construction of the scenario tree for non-i.i.d. disturbances requires the transition probabilities between any two consecutive states in a scenario. However, when building a scenario tree from generated forecasts, state-transition probabilities $\pi(W_{i+1}|W_i, c)\ \forall\ i \in [t : T-1]$ are not known, and would be expensive to compute. Instead, generated samples can be directly used in a single-stage robust horizon decision tree, where each scenario is a generated forecast.

Assuming that the learned distribution can produce realistic scenarios, sampling the distribution can produce scenarios of a tree with arbitrarily long robust horizons, without the need for explicitly defining branches and transition probabilities. Thus taking a small subset of forecasts can be likened to a pruned scenario tree. We do not attempt to formally investigate the myriad strategies which one could use for scenario selection [5], [8], [17]–[20], and instead use a simple but effective strategy for illustrating the importance of forecasting. Here, the scenarios are generated by using the most probable forecasts $\mathcal{S} \subset \{\hat{W}_{t:T}(z, c)\}$ (e.g., take the top λ sequences), and generating scenarios

$$\hat{W}_{t:T}^{(s)} = \mathbb{E}[\mathcal{S}] + \xi_s \sigma(\mathcal{S}), \tag{5}$$

where $s = 1, \ldots, n_s$. The scalars $\xi = [\xi_1, ..., \xi_{n_s}]$ can are used to select the conservativeness of the scenarios, using the normalized probabilities of a standard Gaussian PDF $\phi(\xi_i)$ as the weights i.e., $\omega^{(s)} = \phi(\xi_s) / \sum_i \phi(\xi_i)$, Thus using a single scenario where $\xi_1 = 0$ is equivalent to a standard MPC implementation, and increasing the number of scenarios and/or values of ξ_i results in more conservative control. Additionally, we note that for non-adaptive strategies, the scenarios can be generated exactly as outlined above, expect where $\mathcal{S}$ is the set of most probable scenarios generated by sampling $z \sim \mathcal{N}(0, I)$, i.e., the unconditioned prior.

Algorithm 1 τ-SMPC

Require: $W_{0:t}$, $\hat{G}_\theta$, m, λ, and $\xi = [\xi_1, ..., \xi_{n_s}]$
for i in m **do**
Take random latent samples: $z_i \leftarrow \mathcal{N}(0, 1)$
Compute probabilities: $\Pi_i = \pi(z_i | W_{0:t}, c)$
Generate uncertainty sequences: $\bar{W}_i = \hat{G}_\theta(z_i)$
end for
Collect most probable sequences:

$$\mathcal{S} = \{\bar{W}_i \mid \max_{i \in 1:\lambda} \Pi_i\}$$

Select scenarios for OCP: $\hat{W}^{(s)}_{t:T} = \mathbb{E}[\mathcal{S}] + \xi_s \sigma(\mathcal{S})$,
Normalize probabilities: $\omega^{(s)} = \phi(\xi_s) / \sum_i^{n_s} \phi(\xi_i)$
Solve OCP with scenarios and weights $\{\hat{W}^{(s)}_{t:T}, \omega^{(s)}\}^{n_s}_{s=1}$

The proposed forecasting strategy produces a set of scenarios $\{\hat{W}_{t:T}\}^{n_s}_{s=1}$, which is directly applicable for finite horizon MPC strategies; to accommodate receding horizon problems, we must extend the forecast $\hat{W}_{t:T}$, to $\hat{W}_{t:t+n_p}$. This can be accomplished similarly to the non-adaptive case, where forecasts over $n_p > T - t$ are generated by sampling the unconditioned prior, $W_{0:T}(z \sim \mathcal{N}(0, I), c)$. The forecast we use for the receding horizon problem can be written as

$$\hat{W}^{(s)}_{t:t+n_p} = \left[\hat{W}^{(s)}_{t:T}(z, c_0), \hat{W}^{(s)}_{0:T}(z, c_1), ..., \hat{W}^{(s)}_{0:n_r}(z, c_d)\right] \quad (6)$$

where $\hat{W}^{(s)}_{t:T}(z, c_0)$ is the forecast conditioned on $W_{0:t}$, with each forecast scenario beginning with the measured disturbances at the current time t, i.e., $\hat{W}^{(s)}_t = W_t \; \forall s$, and $\hat{W}_{0:T}(z, c_i)$ and $\hat{W}_{0:n_r}(z, c_i)$ $\forall i \in [1 : d]$ are forecasts generated by sampling the unconditioned prior. Here, c_i represents environmental variables for day i, where $i = 0$ corresponds to the current day, and $d \geq 0$, $n_r \in [0, T]$ are selected at each iteration to ensure the forecast spans the prediction horizon; i.e., $T - t + dT + n_r = n_p$. We summarize the $\tau-$SMPC framework in Algorithm 1.

IV. Case Study

A. Simplified Model of Thermal Dynamics

For this study, we use the previously studied building envelope model (see [11] for details and parameter values)

$$\dot{x} = Ax + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} u + \begin{bmatrix} K_3/C_1 & 1/C_1 & 1/C_1 \\ 0 & 1/C_2 & 0 \\ 0 & 0 & K_4/C_3 \end{bmatrix} w, \quad (7)$$

where

$$A = \begin{bmatrix} \frac{-(K_1+K_2+K_3+K_5)}{C_1} & \frac{K_2}{C_1} & \frac{K_5}{C_1} \\ \frac{K_1+K_2}{C_2} & \frac{-(K_1+K_2}{C_2} & 0 \\ \frac{K_5}{C_3} & 0 & \frac{-(K_4+K_5)}{C_3} \end{bmatrix}.$$

Note that $x, w \in \mathbb{R}^3$, and $u \in \mathbb{R}$, and explicit time dependence, *e.g.*: x_t, w_t has been dropped for convenience.

The parameters C and K denote heat capacities and heat gain coefficients. The states of the system x include room air temperature, interior wall surface temperature, and exterior wall core temperatures, all in ^oC. The control input u is the net heating and cooling power of a heat pump $[kW]$, while the disturbance input vector w is comprised of outside air temperature $[^oC]$, solar radiation $[kW]$, and internal heat loads $[kW]$. We assume that the initial conditions are sampled from the normal distribution $\mathcal{N}([22, 22, 23], \text{diag}[1.5, 1.5, 0.5])$ in summer, and the normal distribution $\mathcal{N}([20, 20, 19], \text{diag}[1.5, 1.5, 0.5])$ in winter.

We assume the net heating and cooling power of the heat pump system is represented by box constraints on the control input, given by $-800 \leq u \leq 800$. Additionally, we assume that the power sourced to operate the HVAC comes from a tiered grid, where the electricity price is defined as $p_t = 0.025[\$/kW]$ from 6am to 10pm, and $0.010[\$/kW]$ otherwise. Finally, given the system is an office building space, we assume that comfort constraints should be tightened during business hours, and that violations during work hours results in a greater economic impact. We formulate the soft time-varying comfort constraints on the room temperature x_1, given by $20 \leq x_1 \leq 24$ from 8AM to 6PM, and $17 \leq x_1 \leq 27$ otherwise; we refer to the upper and lower time-varying bounds by $\bar{x}_t$ and $\underline{x}_t$ at time t, respectively. The economic impact of comfort constraint violations are defined as $\rho_t = 1000$ from 8AM to 6PM, otherwise $\rho_t = 0.1$. Our *objective* is to design a control policy that minimizes this economic cost while enforcing temperature-based comfort, and heat-pump input, constraints.

B. Stochastic MPC Formulation

The scenario tree OCP is given by

$$\min_{u_{0|t}, ..., u_{n_p|t}} \sum_{s=1}^{ns} \omega^{(s)} \sum_{k=0}^{n_p} p_{k+t} |u_{k|t}| + \rho_{k+t} \zeta_{k|t} \quad (8a)$$

$$\text{s.t.} \;\; \hat{x}^{(s)}_{0|t} = x_t, \quad (8b)$$

$$\hat{x}^{(s)}_{k+1|t} = \hat{x}^{(s)}_{k|t} + \Delta t \cdot \dot{x}\left(\hat{x}^{(s)}_{k|t}, u_{k|t}, \hat{w}^{(s)}_{k|t}\right), \quad (8c)$$

$$\zeta_{k|t} = \left[\hat{x}^{(s)}_{1,k|t} - \bar{x}_{k+t}\right]^+ + \left[\underline{x}_{k+t} - \hat{x}^{(s)}_{1,k|t}\right]^+, \quad (8d)$$

$$-800 \leq u_{k|t} \leq 800, \quad \forall k = 0, ..., n_p, \quad (8e)$$

where $[\alpha]^+ = \max(\alpha, 0)$, and $\alpha^{(s)}$ indicates values for the s-th scenario in the tree. For computational efficiency, we compute a single control action u_k over the weighted average across all scenarios; this can be extended to accommodate a full scenario tree. We use $n_s = 3$ for the scenarios used in SMPC, generated using equation (5) for $\xi = [0, 3, -3]$, and weighted using the normalized probability of $\hat{W}^{(s)}_{t+1:T}$. For the standard MPC, we use $n_s = 1$ generated using (5) for $\xi = [0]$. Classically, SMPC uses *offline* disturbance forecasts $\hat{w}^{(s)}_t$ and weights $\omega^{(s)}$, representing the uncertainty which can result in conservative decision-making. Instead, we aim to reduce conservativeness by *online adaptation* of the scenarios and probabilities, under current conditions and given the measurements. Next, we compare the adaptive, non-adaptive, and perfect forecasting, and shows the $\tau-$based methods improve control performance.

C. Results

Here, we compare τ-SPMC performance to a non-adaptive SMPC and to a MPC with perfect disturbance predictions. We use a CVAE with conditioning on whether the day is in summer ($c_1 = 1$) or winter ($c_1 = 0$), and if it's a workday ($c_2 = 1$) or holiday ($c_2 = 0$). In this preliminary study, the disturbance forecasts $\hat{w}_t$ are the only source of mismatch between the building dynamics and the SMPC internal model. As a result, the MPC with perfect disturbance predictions serves as the best performance possible for the given system.

All simulations are done in `Python3`. The generative models are built with the `PyTorch` machine-learning framework, while the SMPC optimization problem is solved using `CasADi` and `IPOPT`. We use the CityLearn dataset [21] to obtain one year's worth of hourly disturbance inputs for a medium office building in Peoria, IL, USA, which has wide seasonal variation in temperature and solar radiation. Internal heat loads are generated using a modified ASHRAE 90.1 occupancy load scheduale. Based on previous studies, we use seasonal and workday/weekend as binary conditioning inputs to the CVAE [3], [13]. Note that even though we have given one year of measurements, some conditioning inputs such as weekend/winter yields a much smaller subset of data points for training; therefore, the learning problem is posed in a limited-data setting. Both the encoder and decoder have 5 layers, with hidden dimension 128-128-64-64-32 for the encoder and in reverse order for the decoder. The latent dimension is fixed to 4. The training of the CVAE is performed using the Adamax solver with a learning rate of 0.0001, and a batch-size of 64.

First, we provide an example how the adaptive method can generate forecasts for nominal and anomalous scenarios. In Figure 3, we show the three ambient temperature scenarios used in the SMPC forecast at various times throughout the day for the adaptive (in blue) and non-adaptive strategies (in magenta) relative to the true disturbance sequence (in black) for a normal summer day (top), and an unseasonably cold summer day (bottom). In the normal day, the true disturbances are captured in the forecasted scenarios for both methods, however the heavily weighted trajectory (i.e., highest probability forecasted scenario) from the adaptive method much more closely approximates the true disturbance. However, the true disturbance on the unseasonably cold day closely resembles a low weighted scenario generated by the non-adaptive forecaster, owing to the fact that the day in question represents a low probability event. On the other hand, the adaptive strategy strategy generates a heavily weighted forecast in close agreement to the true trajectory.

We present the results for 10 closed loop experiments using each for the three controllers over a three day simulation horizon in both winter and summer in Table I. The reported values consist of a mean and 2 standard deviation confidence interval for the energy cost term $\Sigma_0^{72} p_t|u_t|$, and the cumulative total degrees of peak-hour constraint violations

TABLE I: Closed-loop performance comparison for τ-SMPC vs. non-adaptive SMPC. Improvement results are presented using the mean and 95% confidence interval (CI) for ten 3-day closed-loop simulations. Reported costs are the average total energy spent on HVAC operations over three days, and violations are the cumulative total degrees of constraint violation during office hours over the three days.

Method	Winter Cost	Winter Violations	Summer Cost	Summer Violations
Ideal Forecast	509 ± 185	$.00 \pm .00$	445 ± 65	$.00 \pm .00$
τ-MPC	514 ± 167	$.05 \pm .07$	471 ± 64	$.03 \pm .05$
τ-SMPC	**509 ± 161**	**.00±.00**	489±50	**.00±.00**
MPC	519 ± 157	$.45 \pm .52$	**457±51**	$.43 \pm .31$
SMPC	512 ± 147	$.06 \pm .10$	458 ± 49	$.03 \pm .08$

$\Sigma_0^{72} \hat{\rho}_t \zeta_t$, where $\hat{\rho}_k = 1$ from 8AM to 6PM, otherwise $\hat{\rho}_k = 0$. The results show that, although the $\tau-$MPC strategy ($n_s = 1$) does exhibit a minor degree of constraint violation, it beats both the MPC and SMPC in terms of constraint handling across both seasons. Moreover, the winter scenarios for $\tau-$SMPC performs similarly to the perfect forecast in terms of constraint handling and average operational cost. In the summer scenarios, the $\tau-$MPC and $\tau-$SMPC strategies result in higher operational costs, however, incurs over 10-fold reduction in constraint violation between the $\tau-$MPC and SMPC. In all cases, we see that $\tau-$MPC results in less constraint violations than the SMPC, and only the perfect forecasting and $\tau-$SMPC strategies are capable of robustly controlling the system. Additionally, the computational cost of implementing $\tau-$SMPC over the non-adaptive SMPC is only dependent on the number disturbance samples m. The average time to implement the trajectory adaptation is 3.27 seconds, in addition to 0.34 and 2.45 seconds for the standard SMPC using 1 and 3 scenarios respectively. These computational details were generated using an i7-5500U with two 2.4 GHz processors and 7.7 GB memory.

Lastly, we present the state profiles for the indoor temperature and control actions for the closed loop experiments in summer discussed above, along side the realized disturbance profiles in Fig. 4. The state profile satisfies the comfort constraints at all times for the perfect forecast and $\tau-$SMPC formulations, with the latter being notably more conservative, due to forecast uncertainty. In all cases, we see the controllers making use of the relaxed comfort constraints in the off-peak hours to store thermal energy during off-peak pricing. On the other hand, the non-adaptive forecasting strategy is similarly conservative during peak-hours, but attempts to save energy by aggressively cooling the system during the off-peak hours, which results in constraint violations during the early morning hours when comfort constraints are tightened.

V. Conclusions

In this paper, we develop an SMPC framework capable of leveraging real-world disturbance inputs modeled using conditioned variational autoencoders. Based on prior measurements of disturbance inputs, we provide a tractable method to forecast future disturbances based on Bayesian estimation

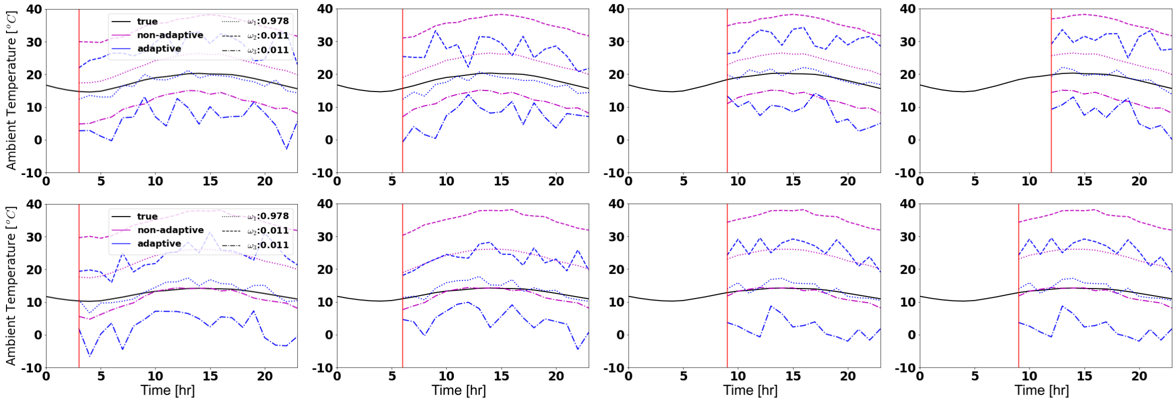

Fig. 3: Adaptive and non-adaptive scenario forecast examples for a normal (top) and cold (bottom) summer day.

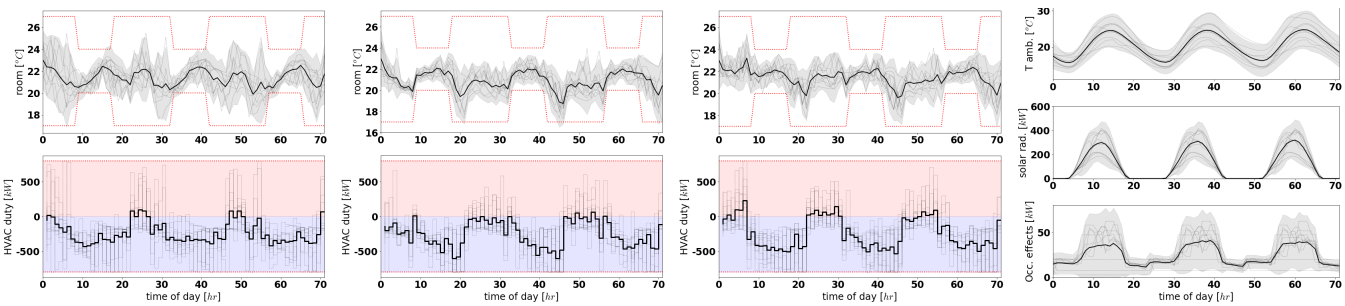

Fig. 4: From left to right, control trajectories using a perfect forecast, non-adaptive, adaptive, and the disturbance profiles.

of conditional probabilities using the generative model. We show that our methodology, while suitable for typical operations, is also well-prepared for out-of-sample scenarios such as unexpected climatic events, therefore demonstrating its potential in combating climate change. In future work, we will investigate alterations in the generative model training procedure that can help the adaptation process, alternative strategies for selecting forecast scenarios, and evaluate the control scheme on high-fidelity simulators.

References

[1] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.

[2] D. Kim *et al.*, "System identification for building thermal systems under the presence of unmeasured disturbances in closed loop operation: Theoretical analysis and application," *Energy and Buildings*, 2018.

[3] P. Kumar *et al.*, "Grey-box model and neural network disturbance predictor identification for economic mpc in building energy systems," *Energy and Buildings*, vol. 286, p. 112936, 2023.

[4] F. Oldewurtel *et al.*, "Stochastic model predictive control for building climate control," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 3, pp. 1198–1205, 2014.

[5] M. Maiworm, T. Bäthge, and R. Findeisen, "Scenario-based model predictive control: Recursive feasibility and stability," 2015. 9th IFAC Symposium on Advanced Control of Chemical Processes ADCHEM.

[6] R. Kumar *et al.*, "Stochastic model predictive control for central hvac plants," *Journal of Process Control*, vol. 90, pp. 1–17, 2020.

[7] J. A. Paulson *et al.*, "Stochastic model predictive control with joint chance constraints," *International Journal of Control*, 2020.

[8] S. Thangavel, R. Paulen, and S. Engell, "Multi-stage nmpc using sigma point principles," 2020. 6th Conference on Advances in Control and Optimization of Dynamical Systems ACODS.

[9] J. Pohlodek *et al.*, "Stochastic model predictive control utilizing bayesian neural networks," 2023.

[10] E. D. Klenske *et al.*, "Approximate dual control maintaining the value of information with an application to building control," in *European Control Conference*, 2016.

[11] R. Gondhalekar, F. Oldewurtel, and C. N. Jones, "Least-restrictive robust mpc of periodic affine systems with application to building climate control," in *49th Conference on Decision and Control*, 2010.

[12] A. Salatiello *et al.*, "Synthesizing building operation data with generative models: VAEs, GANs, or Something In Between?," in *Proc. of the 14th ACM Int. Conf. on Future Energy Systems*.

[13] F. Khayatian, Z. Nagy, and A. Bollinger, "Using generative adversarial networks to evaluate robustness of reinforcement learning agents against uncertainties," *Energy and Buildings*, vol. 251, p. 111334, 2021.

[14] C. A. Hans *et al.*, "Risk-averse model predictive operation control of islanded microgrids," *IEEE Transactions on Control Systems Technology*, 2020.

[15] Y. Chen *et al.*, "Model-free renewable scenario generation using generative adversarial networks," *IEEE Transactions on Power Systems*, vol. 33, pp. 3265–3275, 2018.

[16] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *Int. Conf. on Learning Representations*, 2014.

[17] M. Wytock, N. Moehle, and S. Boyd, "Dynamic energy management with scenario-based robust mpc," in *2017 ACC*, 2017.

[18] G. C. Goodwin and A. M. Medioli, "Scenario-based, closed-loop model predictive control with application to emergency vehicle scheduling," *International Journal of Control*, vol. 86, no. 8, 2013.

[19] A. D. Bonzanini *et al.*, "Fast approximate learning-based multistage nonlinear model predictive control using gaussian processes and deep neural networks," *Computers & Chemical Engineering*.

[20] A. B. Kordabad *et al.*, "Reinforcement learning based on scenario-tree mpc for asvs," in *American Control Conference*, pp. 1985–1990, 2021.

[21] J. R. Vázquez-Canteli *et al.*, "Citylearn v1.0: An openai gym environment for demand response with deep reinforcement learning," in *Proc. 6th ACM Int. Conf. on Systems for Energy-Efficient Buildings, Cities, and Transportation*, BuildSys '19, Association for Computing Machinery, 2019.