

Can a Transformer Represent a Kalman Filter?

Gautam Goel

Peter Bartlett

Simons Institute, UC Berkeley

Abstract

Transformers are a class of autoregressive deep learning architectures which have recently achieved state-of-the-art performance in various vision, language, and robotics tasks. We revisit the problem of Kalman Filtering in linear dynamical systems and show that Transformers can approximate the Kalman Filter in a strong sense. Specifically, for any observable LTI system we construct an explicit causally-masked Transformer which implements the Kalman Filter, up to a small additive error which is bounded uniformly in time; we call our construction the Transformer Filter. Our construction is based on a two-step reduction. We first show that a softmax self-attention block can exactly represent a Nadaraya–Watson kernel smoothing estimator with a Gaussian kernel. We then show that this estimator closely approximates the Kalman Filter. We also investigate how the Transformer Filter can be used for measurement-feedback control and prove that the resulting nonlinear controllers closely approximate the performance of standard optimal control policies such as the LQG controller.

1 Introduction

Transformers are a class of autoregressive deep learning architectures designed for various sequence modelling tasks, first introduced in [8]. Transformers have quickly emerged as the best performing class of deep learning models across a variety of challenging domains, including computer vision, natural language processing, and robotics, and have also been studied in the context of reinforcement learning and decision-making (e.g. [1, 4, 5, 9]). While the empirical successes of Transformers are exciting, we still lack a formal theory that explains what Transformers can do and why they work. In this paper, we study how Transformers can be used for filtering and control in linear dynamical systems. We ask perhaps the most basic question one could ask: can a Transformer be used for Kalman Filtering? The Kalman Filter is foundational in optimal control and a crucial component of the Linear-Quadratic-Gaussian (LQG) controller. If Transformers were unable to perform Kalman Filtering, then the use of Transformers in signal processing and control would be suspect; conversely, establishing that Transformers can indeed perform Kalman Filtering is a crucial first step towards establishing the viability of Transformers in these domains.

In the mathematical theory of deep learning, three questions naturally arise. First, which functions can a given deep learning architecture represent? Second, when trained on data, what function does the deep learning system actually learn? Lastly, how well does this learned function generalize on new data? We focus on the first of these questions and leave the other two for future work. Specifically, we investigate the following questions. First, is the nonlinear structure of a Transformer compatible with a Kalman Filter at all? This is not obvious; it is possible *a priori* that no matter how a Transformer is implemented, the softmax nonlinearity in the self-attention block will cause the state estimates of the Transformer and the Kalman Filter to diverge over time. Second, if it is possible to represent the Kalman Filter with a Transformer, what would that Transformer look like? How should the states and observations be represented within the Transformer? It is known that *positional encoding* improves the performance of Transformers in some tasks - is it necessary for Kalman Filtering? How large must the Transformer must be, e.g., how large must the embedding dimension be, and how many self-attention blocks are required?

1.1 Key contributions

We construct an explicit Transformer which implements the Kalman Filter, up to a small additive error; we call our construction the Transformer Filter. Our construction is based on a two-step reduction. First, we show that a self-attention block can exactly represent a Nadaraya–Watson kernel smoothing estimator with a Gaussian kernel. We select a specific covariance in our Gaussian kernel with a system-theoretic interpretation: it measures how closely a

previous state estimate matches the most recent state estimate, where the measure of “closeness” is the ℓ_2 distance between the one-step Kalman Filter updates using each of the state estimates. The kernel takes as inputs nonlinear embeddings of the previous state estimates and observations; these embeddings have quadratic dependence on the size of the underlying state-space model. In particular, if the state-space model has an n -dimensional state and p -dimensional observations, the kernel we construct takes as input embeddings of dimension $O((n+p)^2)$. The second step in our construction is to show that this kernel smoothing algorithm approximates the Kalman Filter in a strong sense. Specifically, for every $\varepsilon > 0$, we show that by increasing a temperature parameter β in our kernel, we can ensure that the sequence of state estimates generated by the Transformer Filter is ε -close to the sequence of state estimates generated by the Kalman Filter. A noteworthy aspect of our construction is that it does not use any positional embedding; permuting the history of state estimates and observations has no effect on the state estimates generated in subsequent timesteps.

We next investigate how the Transformer Filter can be incorporated into a measurement-feedback control system. A key technical challenge is to understand the closed-loop dynamics that are induced by the Transformer Filter; since the state-estimates produced by the Transformer Filter are a nonlinear function of the observations, the resulting closed-loop map is also nonlinear. This means that standard techniques for establishing stability of the system, such as bounding the eigenvalues of the closed-loop map, cannot be used. We show that the Transformer Filter can closely approximate an LQG controller, in the following sense: for every $\varepsilon > 0$, we construct a controller using the Transformer Filter which generates a state sequence that is ε -close to the state sequence generated by the LQG controller. A consequence of this result is that the controllers we construct are weakly stabilizing in the following sense; while they may not drive the state all the way to zero, they are guaranteed to drive the state into a small ball centered at zero. Our result also implies that the cost incurred by our new controller can be driven arbitrarily close to the optimal cost achieved by the LQG controller. All of our approximation results also hold when the reference algorithm is taken to be an H_∞ filter or H_∞ controller.

2 Preliminaries

2.1 Filtering and Control

The first problem we consider is *Filtering in Linear Dynamical Systems*. In this problem, we consider a partially observed linear system

$$x_{t+1} = Ax_t + w_t, \quad y_t = Cx_t + v_t,$$

where $x_t \in \mathbb{R}^n$ is an unknown state and $y_t \in \mathbb{R}^p$ is a noisy linear observation of the state; the variables w_t v_t are exogenous disturbances which perturb the state and observation. The state is initialized at time $t = 0$ to some fixed state x_0 . The task of filtering is to sequentially estimate the state sequence given the observation sequence. We focus on the strictly causal setting, where the filtering algorithm estimates the state x_t after observing $y_0, \dots, y_{t-1}$. The best-known algorithm for filtering in linear dynamical systems is undoubtedly the *Kalman Filter*, which is the mean-square-optimal linear filter when the disturbances are stochastic. More precisely, if $\{w_t\}_{t \geq 0}$ and $\{v_t\}_{t \geq 0}$ are assumed to be independent, white noise processes, then the estimate $\hat{x}_t^*$ produced by the Kalman Filter satisfies

$$\hat{x}_t^* = \inf_z \mathbb{E} [\|z - x_t\|^2],$$

where the infimum is taken over all linear functions $z(y_0, \dots, y_t)$ of the observations. In the special case where the disturbances are Gaussian, the Kalman Filter estimate is also a maximum likelihood estimate of the state conditioned on the observations. The Kalman Filter has the following recursive form: the prediction of the next state given the observations $y_0, \dots, y_{t-1}$ is

$$\hat{x}_t^* = (A - LC)\hat{x}_{t-1}^* + Ly_{t-1}, \tag{1}$$

where L is a fixed matrix called the *Kalman gain* and we initialize $x_0^* = x_0$. We also note that the H_∞ filter has an identical recursive form to the Kalman Filter, except with a different gain matrix L [2].

The second problem we consider is *Measurement-Feedback Control in Linear Dynamical Systems*. In this problem, we again consider a partially observed linear system, but the system is now augmented to include a control input $u_t \in \mathbb{R}^m$:

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad y_t = Cx_t + v_t.$$

The goal of the controller is to select the control action u_t to regulate the state using only the observations $y_0, \dots, y_t$. In the Linear-Quadratic-Gaussian (LQG) model, the disturbances $\{w_t\}_{t \geq 0}$ and $\{v_t\}_{t \geq 0}$ are once again assumed to be independent, white noise processes, and the control actions are selected to minimize the infinite-horizon cost

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\{w_t, v_t\}_{t \geq 0}} \left[\sum_{t=0}^T x_t^\top Q x_t + u_t^\top R u_t \right].$$

It is known that in this case the optimal policy is to use the Kalman Filter to produce a state estimate $\hat{x}_t^*$ and then to pick the control actions as a linear function of the estimate [2]. The Kalman Filter estimate is adjusted to account for the influence of the control input, so the LQG policy is

$$\hat{x}_t^* = (A + BK - LC)\hat{x}_{t-1}^* + Ly_{t-1}, \quad u_t = K\hat{x}_t^*, \quad (2)$$

where K is called the *state-feedback matrix*. Other measurement-feedback controllers of the general form (2) include the H_∞ measurement-feedback controller, which uses a different choice of L and K [2].

We assume that the pair (A, C) is observable, and the pair (A, B) is controllable; we refer to [3] for background on linear systems. We let $\|A\|$ denote the spectral norm of a matrix A . We make repeated use of the following facts.

Fact 1. *Let $A \in \mathbb{R}^{q \times q}$ be any stable matrix, i.e., any matrix with spectral radius strictly less than 1. There exist matrices M, θ such that $A = M\theta M^{-1}$ and $\|\theta\| < 1$.*

Fact 2. *Let (L, K) represent any stabilizing linear measurement-feedback controller. The matrices $A - LC$ and $A + BK$ are both stable.*

2.2 Transformers and Softmax Self-Attention

A Transformer is a deep learning architecture which alternates between self-attention blocks and Multilayer Perceptron (MLP) blocks. In this paper we focus on Transformers with a single self-attention block, followed by a single MLP block; furthermore, we always assume that the weights of the MLP block are chosen so that the MLP block represents the identity function. The interesting part of our construction hence lies in how we choose the parameters of the self-attention block.

A general softmax self-attention block has the following form. It takes as input a series of tokens $q_0, \dots, q_N$ and a query token q , and outputs

$$F(q_0, \dots, q_N; q) = \frac{\sum_{i=0}^N \exp(q^\top A q_i) M q_i}{\sum_{j=0}^N \exp(q^\top A q_j)},$$

where A and M are parameters of the Transformer; we refer to [7] for an excellent overview of Transformers. In our paper we consider *causally masked* Transformers, which means that we think of the tokens as being indexed by time and at each timestep t we drop all the tokens which have not yet been observed, only keeping those up until time t . In our results, we also drop all but the last H observed tokens, to obtain the self-attention block

$$F(q_{t-H+1}, \dots, q_t; q) = \frac{\sum_{i=t-H+1}^t \exp(q^\top A q_i) M q_i}{\sum_{j=t-H+1}^t \exp(q^\top A q_j)}.$$

In our construction, the tokens q_i are embeddings of the i -th state-estimate and the i -th observation, i.e., $q_i = \phi(\hat{x}_i, y_i)$, where ϕ is a nonlinear embedding map. The Transformer Filter generates state estimates recursively; it takes as input the past H state estimates and observations $(\hat{x}_{t-H}, y_{t-H}), \dots, (\hat{x}_{t-1}, y_{t-1})$, embeds them as tokens using the map ϕ , feeds these tokens $q_{t-H+1}, \dots, q_t$ into the self-attention block, and outputs a new state estimate $\hat{x}_t = F(q_{t-H+1}, \dots, q_t; q)$, where we take $q = q_t$. We note that the Kalman Filter has a similar recursive form; it uses the previous estimate $\hat{x}_{t-1}$ and the previous observation y_{t-1} to generate the new estimate $\hat{x}_t$. In fact, in the special case when $H = 1$, the Transformer Filter exactly coincides with the Kalman Filter.

3 Nadaraya–Watson Kernel Smoothing via Softmax Self-Attention

Our first result is that the class of Transformers we study is capable of representing a Nadaraya–Watson estimator with a Gaussian kernel. Intuitively, given data $\{z_i\}_{i=0}^N$ and a query point z , a Gaussian kernel smoothing estimator outputs

a linear combination of the data, weighted by how close each datapoint z_i is to z , where the measure of “closeness” is determined by a fixed covariance matrix Σ . We refer to [6] for more background on kernel smoothing and the Nadaraya–Watson estimator.

Theorem 1. Fix $\Sigma \in \mathbb{R}^{d \times d}$ and $W \in \mathbb{R}^{k \times d}$. Suppose we are given $z_0, \dots, z_N \in \mathbb{R}^d$ and $z \in \mathbb{R}^d$. Define the Nadaraya–Watson estimator

$$F(z_0, \dots, z_N; z) = \frac{\sum_{i=0}^N \exp(-(z - z_i)^\top \Sigma (z - z_i)) W z_i}{\sum_{j=0}^N \exp(-(z - z_j)^\top \Sigma (z - z_j))}. \quad (3)$$

The function F can be represented by a softmax self-attention block of size $O(d^2 H)$. In particular, there exists a nonlinear embedding map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^\ell$ and matrices $M \in \mathbb{R}^{k \times \ell}$ and $A \in \mathbb{R}^{\ell \times \ell}$ such that

$$F(z_0, \dots, z_N; z) = \frac{\sum_{i=0}^N \exp(q^\top A q_i) M q_i}{\sum_{j=0}^N \exp(q^\top A q_j)},$$

where we define $q_i = \phi(z_i)$ and $q = \phi(z)$ and set $\ell = \binom{n}{2} + n + 1$.

Proof. We first show that the function $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ given by

$$f(u, v) = \exp(-(u - v)^\top \Sigma (u - v))$$

can be represented as

$$f(u, v) = \exp(\phi(u)^\top A \phi(v))$$

for some embedding map ϕ and matrix A of appropriate dimensions. Observe that

$$(u - v)^\top \Sigma (u - v) = \sum_{i,j=1}^n \Sigma_{i,j} (u_i u_j - 2u_i v_j + v_i v_j). \quad (4)$$

Define

$$\phi(u) = [1 \quad u_1 \quad \dots \quad u_n \quad u_1 u_1 \quad u_1 u_2 \quad \dots \quad u_n u_{n-1} \quad u_n u_n]^\top.$$

It is easy to see that the set

$$S(u, v) = \{\phi(u)^\top A \phi(v) \mid A \in \mathbb{R}^{\ell \times \ell}, A = A^\top\}$$

represents all polynomials in $(u_1, \dots, u_n, v_1, \dots, v_n)$ of degree at most four such that within each monomial, the degree of the variables appearing in u (resp. v) is at most 2. In particular, there exists a choice of A such that $\phi(u)^\top A \phi(v)$ is exactly the polynomial (4). We take the matrix M to be the $k \times \ell$ matrix which contains W as a submatrix from columns 2 to $n + 1$ and is zero elsewhere; notice that $M \phi(u) = W u$ for all $u \in \mathbb{R}^d$. $\square$

4 Filtering

We ask:

Can a Transformer implement the Kalman Filter?

Naturally, since a Transformer is a complicated nonlinear function of its inputs, it is too much to expect a Transformer to exactly represent a Kalman Filter. We instead ask the following approximation-theoretic question: for any $\varepsilon > 0$, does there exist a Transformer which generates state estimates which are ε -close to the state estimates generated by the Kalman Filter, uniformly in time?

We consider the one-layer Transformer whose MLP block is the identity function and whose self-attention block takes as input embeddings of the past H state estimates and observations

$$\begin{bmatrix} \hat{x}_{t-H} \\ y_{t-H} \end{bmatrix}, \dots, \begin{bmatrix} \hat{x}_{t-1} \\ y_{t-1} \end{bmatrix}$$

and outputs the estimate

$$\hat{x}_t = \sum_{i=t-H+1}^t \alpha_{i,t} \tilde{x}_i,$$

where we define

$$\alpha_{i,t} = \frac{\exp(-\beta \|\tilde{x}_i - \tilde{x}_t\|^2)}{\sum_{j=t-H+1}^t \exp(-\beta \|\tilde{x}_j - \tilde{x}_t\|^2)}, \quad \tilde{x}_i = \begin{bmatrix} A - LC & L \end{bmatrix} \begin{bmatrix} \hat{x}_{i-1} \\ y_{i-1} \end{bmatrix},$$

for all $t \geq 1$ and set $\hat{x}_0, \tilde{x}_0 = x_0$. We adopt the convention that $\hat{x}_i, y_i = 0$ for all $i < 0$. We call this filter the Transformer Filter; it is easy to check that this filter is a special case of the Gaussian kernel smoothing estimator (3) and hence by Theorem 1 can be represented by a Transformer. The variables $\tilde{x}_i$ have the following interpretation; they are the estimates that would be generated by the Kalman Filter recursion (1) if the previous Kalman Filter estimate $\hat{x}_{i-1}^*$ were replaced by the Transformer estimate $\hat{x}_{i-1}$. In that sense, the variables $\tilde{x}_i$ interpolate between the Kalman Filter and the Transformer Filter. We prove:

Theorem 2. *For each $\varepsilon > 0$, there exists a $\beta > 0$ such that the state estimates $\{\hat{x}_t\}_{t \geq 0}$ generated by the Transformer Filter satisfy*

$$\|\hat{x}_t - \hat{x}_t^*\| \leq \varepsilon$$

at all times $t \geq 0$, where $\{\hat{x}_t^*\}_{t \geq 0}$ are the state-estimates generated by the Kalman Filter (1). In particular, it suffices to take

$$\beta \geq \frac{H^2 \kappa^2}{2e(1 - \|\theta\|)^2 \varepsilon^2},$$

where M, θ are $n \times n$ matrices such that $A - LC = M\theta M^{-1}$ and $\|\theta\| < 1$, and we define $\kappa = \|M\| \|M^{-1}\|$.

Proof. We first show that for all $\varepsilon_1 > 0$, there exists a β such that

$$\|\hat{x}_t - \tilde{x}_t\| \leq \varepsilon_1$$

at all times $t \geq 0$. Fix any $\varepsilon_1 > 0$ and any $t \geq 1$. Notice that for each $i \in \{t-H+1, \dots, t\}$, the following inequality holds:

$$\alpha_{i,t} < \exp(-\beta \|\tilde{x}_i - \tilde{x}_t\|^2).$$

This is because

$$\sum_{j=t-H+1}^t \exp(-\beta \|\tilde{x}_j - \tilde{x}_t\|^2) > 1.$$

It follows that

$$\begin{aligned} \|\hat{x}_t - \tilde{x}_t\| &= \left\| \sum_{i=t-H+1}^t \alpha_{i,t} (\tilde{x}_i - \tilde{x}_t) \right\| \\ &\leq \sum_{i=t-H+1}^t \alpha_{i,t} \|\tilde{x}_i - \tilde{x}_t\| \\ &< \sum_{i=t-H+1}^t \exp(-\beta \|\tilde{x}_i - \tilde{x}_t\|^2) \|\tilde{x}_i - \tilde{x}_t\| \\ &\leq H \max_{\gamma \geq 0} \exp(-\beta \gamma^2) \gamma, \end{aligned}$$

where we used the fact that $\sum_{i=t-H+1}^t \alpha_{i,t} = 1$ in the first step. It is easy to check that the function $f(\gamma) = H e^{-\beta \gamma^2} \gamma$ is strictly increasing in the interval $(0, (2\beta)^{-1/2})$ and strictly decreasing in the interval $((2\beta)^{-1/2}, \infty)$ and hence takes its maximum value of $H e^{-1/2} (2\beta)^{-1/2}$ at $\gamma = (2\beta)^{-1/2}$. It follows that $\|\hat{x}_t - \tilde{x}_t\| \leq \varepsilon_1$ as long as $\beta \geq \frac{H^2}{2e\varepsilon_1^2}$.

We now show that this result implies Theorem 2. Fix $\varepsilon > 0$ and any $t \geq 0$. Set $\varepsilon_1 = \frac{(1-\gamma)\varepsilon}{\kappa}$ and $\beta \geq \frac{H^2}{2\varepsilon\varepsilon_1^2}$. Using the preceding argument, this suffices to ensure that $\|\hat{x}_t - \tilde{x}_t\| \leq \varepsilon_1$ for all $t \geq 0$. We see that

$$\begin{aligned}\|\hat{x}_t - \hat{x}_t^*\| &\leq \|\hat{x}_t - \tilde{x}_t\| + \|\tilde{x}_t - \hat{x}_t^*\| \\ &\leq \|\hat{x}_t - \tilde{x}_t\| + \|(A - LC)(\hat{x}_{t-1} - \hat{x}_{t-1}^*)\|\end{aligned}$$

Proceeding recursively, we obtain the bound

$$\begin{aligned}\|\hat{x}_t - \hat{x}_t^*\| &\leq \sum_{i=0}^t \|(A - LC)^i\| \|\hat{x}_{t-i} - \tilde{x}_{t-i}\| \\ &\leq \varepsilon_1 \sum_{i=0}^t \|(A - LC)^i\|.\end{aligned}$$

Using the fact that $A - LC = M\theta M^{-1}$ with $\|M\|\|M^{-1}\| = \kappa$ and $\|\theta\| < 1$, we see that

$$\begin{aligned}\|\hat{x}_t - \hat{x}_t^*\| &\leq \varepsilon_1 \sum_{i=0}^t \|(M\theta M^{-1})^i\| \\ &\leq \kappa\varepsilon_1 \sum_{i=0}^t \|\theta\|^i \\ &\leq \frac{\kappa\varepsilon_1}{1 - \|\theta\|} \\ &= \varepsilon.\end{aligned}$$

□

We note that the only property of the gain matrix L we used is that $A - LC$ is stable; since this property also holds for the H_∞ -optimal choice of L , our proof also shows that a Transformer can approximate an H_∞ -optimal filter.

5 Control

We ask:

Can the Transformer Filter be used in place of the Kalman Filter in the LQG controller?

Since the Transformer Filter only represents the Kalman Filter approximately, we cannot hope to implement the LQG controller exactly. Instead, we ask if the closed-loop dynamics generated by the Transformer can closely approximate the closed-loop dynamics generated by the LQG controller in the following sense: for any $\varepsilon > 0$, can we guarantee that the states generated by the Transformer are ε -close to the states generated by the Transformer, uniformly in time? We emphasize that this is far from obvious, and in particular does not follow directly from Theorem 2. Even if the state-estimates generated by the Transformer Filter are close to those generated by the Kalman Filter, it does not automatically follow that the resulting control policies will generate similar state trajectories. This is because any difference in state estimates will lead to a difference in the control actions, which in turn affects future states, future observations, and so on; in effect, minute deviations between the two state estimates could be amplified over time, leading to diverging trajectories. In order to show that this scenario does not occur, we need to analyze the stability of the closed-loop map induced by the Transformer Filter. This is challenging, because this map is nonlinear, and hence we cannot use standard techniques from linear systems theory.

We consider the controller given by

$$u_t = K\hat{x}_t$$

where we set

$$\hat{x}_t = \sum_{i=t-H+1}^t \alpha_{i,t} \tilde{x}_i,$$

and we define

$$\alpha_{i,t} = \frac{\exp(-\beta \|\tilde{x}_i - \tilde{x}_t\|^2)}{\sum_{j=t-H+1}^t \exp(-\beta \|\tilde{x}_j - \tilde{x}_t\|^2)}, \quad \tilde{x}_i = (A + BK - LC)\hat{x}_{i-1} + Ly_{i-1}.$$

We initialize the state of the system driven by the Transformer system to match the state of the system driven by the LQG policy (i.e., $x_0 = x_0^*$) and similarly initialize the state estimates to be the same ($\hat{x}_0 = \hat{x}_0^*$). We also initialize $\tilde{x}_0 = \hat{x}_0$. We prove:

Theorem 3. *For each $\varepsilon > 0$, there exists a $\beta > 0$ such that the states $\{x_t\}_{t \geq 0}$ generated by the Transformer Filter satisfy*

$$\|x_t - x_t^*\| \leq \varepsilon$$

at all times $t \geq 0$, where $\{x_t^\}_{t \geq 0}$ are the states generated by the optimal LQG control policy (2). In particular, it suffices to take*

$$\beta \geq \frac{CH^2\kappa^2}{2e(1 - \|\Theta\|)^2\varepsilon^2}$$

where we define

$$C = 2\|BK\|^2 + 2\|A + BK - LC\|^2, \quad \kappa = \|\mathbb{M}\|\|\mathbb{M}^{-1}\|,$$

and $\mathbb{M}, \Theta$ are $4n \times 4n$ matrices such that $\mathbb{A} = \mathbb{M}\Theta\mathbb{M}^{-1}$ and $\|\Theta\| < 1$, and we define

$$\mathbb{A} = \begin{bmatrix} A & BK & 0 & 0 \\ LC & A + BK - LC & 0 & 0 \\ 0 & 0 & A & BK \\ 0 & 0 & LC & A + BK - LC \end{bmatrix}.$$

Before we turn to the proof, we note an interesting consequence of this result: the controller induced by the Transformer Filter is *weakly stabilizing* in the sense that no matter how x_0 is chosen, if the disturbances are zero then the states generated by the controller will eventually be confined to a ball of radius ε centered at the origin. This follows from the fact that the LQG controller is stabilizing (i.e., it drives the state to zero in the absence of noise).

Proof. An identical argument to that appearing in the proof of Theorem 2 establishes that for all $\varepsilon_1 > 0$, choosing

$$\beta \geq \frac{H^2}{2e\varepsilon_1^2}$$

guarantees that

$$\|\hat{x}_t - \tilde{x}_t\| \leq \varepsilon_1$$

at all times $t \geq 0$. The closed-loop dynamics can be written as

$$\begin{bmatrix} x_{t+1} \\ \tilde{x}_{t+1} \\ x_{t+1}^* \\ \hat{x}_{t+1}^* \end{bmatrix} = \mathbb{A} \begin{bmatrix} x_t \\ \tilde{x}_t \\ x_t^* \\ \hat{x}_t^* \end{bmatrix} + \eta_t + \nu_t \quad (5)$$

where we set

$$\eta_t = \begin{bmatrix} BK(\hat{x}_t - \tilde{x}_t) \\ (A + BK - LC)(\hat{x}_t - \tilde{x}_t) \\ 0 \\ 0 \end{bmatrix}, \quad \nu_t = \begin{bmatrix} w_t \\ Lv_t \\ w_t \\ Lv_t \end{bmatrix}$$

for all $t \geq 0$. We emphasize that while the dynamics (5) may superficially appear linear, the variable η_t depends on $\hat{x}_t$, which itself is a nonlinear function of the H observations $y_{t-H+1}, \dots, y_t$, so that the overall behavior of the closed-loop system is nonlinear. In effect, we have pushed all of the nonlinearity and memory in the closed-loop system into the variables $\{\eta_t\}_{t \geq 0}$.

The matrix $\mathbb{A}$ is stable. To see this, notice that

$$\mathbb{A} = \mathbb{Q}^{-1}\mathbb{S}\mathbb{Q},$$

where we define the $4n \times 4n$ block matrices

$$\mathbb{Q} = \begin{bmatrix} I & -I & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & -I \\ 0 & 0 & 0 & I \end{bmatrix}, \quad \mathbb{S} = \begin{bmatrix} A - LC & 0 & 0 & 0 \\ LC & A + BK & 0 & 0 \\ 0 & 0 & A - LC & 0 \\ 0 & 0 & LC & A + BK \end{bmatrix}.$$

It is clear that $\mathbb{S}$ is stable, because $\mathbb{S}$ is block lower-triangular and the matrices $A - LC$ and $A + BK$ appearing on the diagonal of $\mathbb{S}$ are both stable. Since $\mathbb{A}$ is similar to $\mathbb{S}$ and $\mathbb{S}$ is stable, $\mathbb{A}$ must also be stable. It follows that $\mathbb{A} = \mathbb{M}\Theta\mathbb{M}^{-1}$ for some $4n \times 4n$ matrices $\mathbb{M}$ and Θ such that $\|\Theta\| < 1$.

Fix $t \geq 1$ and $\varepsilon > 0$. Set

$$\varepsilon_1 = \frac{\varepsilon(1 - \|\Theta\|)}{\kappa\sqrt{2\|BK\|^2 + 2\|A + BK - LC\|^2}}, \quad \beta \geq \frac{H^2}{2e\varepsilon_1^2}.$$

The closed-loop dynamics (5) imply that

$$\begin{bmatrix} x_t \\ \tilde{x}_t \\ x_t^* \\ \hat{x}_t^* \end{bmatrix} = \mathbb{A}^t \begin{bmatrix} x_0 \\ \tilde{x}_0 \\ x_0^* \\ \hat{x}_0^* \end{bmatrix} + \sum_{i=0}^{t-1} \mathbb{A}^{t-1-i} (\eta_i + \nu_i).$$

It follows that

$$x_t - x_t^* = \begin{bmatrix} I & 0 & -I & 0 \end{bmatrix} \left(\mathbb{A}^t \begin{bmatrix} x_0 \\ \tilde{x}_0 \\ x_0^* \\ \hat{x}_0^* \end{bmatrix} + \sum_{i=0}^{t-1} \mathbb{A}^{t-1-i} \nu_i \right) + \begin{bmatrix} I & 0 & -I & 0 \end{bmatrix} \sum_{i=0}^{t-1} \mathbb{A}^{t-1-i} \eta_i.$$

Notice that the first term is zero; this follows from the assumption that we initialize $x_0 = x_0^*$ and $\tilde{x}_0 = \hat{x}_0 = \hat{x}_0^*$ and the block-diagonal structure of $\mathbb{A}$. We see that

$$\begin{aligned} \|x_t - x_t^*\| &\leq \left\| \begin{bmatrix} I & 0 & -I & 0 \end{bmatrix} \right\| \cdot \sum_{i=0}^{t-1} \|\mathbb{A}^{t-1-i}\| \|\eta_i\| \\ &= \left\| \begin{bmatrix} I & 0 & -I & 0 \end{bmatrix} \right\| \cdot \sum_{i=0}^{t-1} \left\| (\mathbb{M}\Theta\mathbb{M}^{-1})^{t-1-i} \right\| \|\eta_i\| \\ &\leq \varepsilon_1 \cdot \kappa \sqrt{2\|BK\|^2 + 2\|A + BK - LC\|^2} \sum_{i=0}^{t-1} \|\Theta\|^{t-1-i} \\ &\leq \frac{\varepsilon_1 \cdot \kappa \sqrt{2\|BK\|^2 + 2\|A + BK - LC\|^2}}{1 - \|\Theta\|} \\ &\leq \varepsilon, \end{aligned}$$

where in the third step we used the fact that $\|\tilde{x}_i - \hat{x}_i\| \leq \varepsilon_1$ for all $i \geq 0$. $\square$

We note that the only property of the gain matrices L and K we used is that $A - LC$ and $A + BK$ are stable; since this property also holds for the H_∞ -optimal choice of L and K , our proof also shows that a Transformer can approximate an H_∞ -optimal measurement-feedback controller.

References

- [1] Lili Chen et al. “Decision transformer: Reinforcement learning via sequence modeling”. In: *Advances in neural information processing systems* 34 (2021), pp. 15084–15097.
- [2] Babak Hassibi, Ali H Sayed, and Thomas Kailath. *Indefinite-Quadratic estimation and control: a unified approach to H2 and H-Infinity theories*. SIAM, 1999.

- [3] Thomas Kailath, Ali H Sayed, and Babak Hassibi. *Linear estimation*. Prentice Hall, 2000.
- [4] Jonathan N Lee et al. “Supervised Pretraining Can Learn In-Context Reinforcement Learning”. In: *arXiv preprint arXiv:2306.14892* (2023).
- [5] Licong Lin, Yu Bai, and Song Mei. “Transformers as Decision Makers: Provable In-Context Reinforcement Learning via Supervised Pretraining”. In: *arXiv preprint arXiv:2310.08566* (2023).
- [6] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [7] Mary Phuong and Marcus Hutter. “Formal algorithms for transformers”. In: *arXiv preprint arXiv:2207.09238* (2022).
- [8] Ashish Vaswani et al. “Attention is all you need”. In: *Advances in neural information processing systems* 30 (2017).
- [9] Qinqing Zheng, Amy Zhang, and Aditya Grover. “Online decision transformer”. In: *international conference on machine learning*. PMLR. 2022, pp. 27042–27059.