JOINT ESTIMATION OF TRAJECTORY AND DYNAMICS FROM PAIRED COMPARISONS

Gregory H. Canal, Matthew R. O'Shaughnessy*, Christopher J. Rozell, and Mark A. Davenport*

School of Electrical & Computer Engineering
Georgia Institute of Technology
{gregory.canal, moshaughnessy6, crozell, mdav}@gatech.edu

ABSTRACT

Recent literature has developed methods for localizing a low-dimensional vector $\boldsymbol{w}$ from paired comparisons of the form "$\boldsymbol{w}$ is closer to $\boldsymbol{p}$ than $\boldsymbol{q}$," where $\boldsymbol{p}$ and $\boldsymbol{q}$ are selected from a fixed set of landmark points and $\boldsymbol{w}$ does not change over time. In this work, we consider a *time-varying* extension of this problem, in which $\boldsymbol{w}$ evolves according to some unknown dynamics model. We consider the task of actively selecting informative paired comparisons between landmark points to jointly estimate the state trajectory and identify the true dynamics model from a finite set of candidate models. Leveraging information-theoretic insights, we propose selecting pairs that simultaneously maximize information gain about both the trajectory and dynamics model, and propose a Bayesian method for tracking and system identification. We demonstrate the efficacy of our approach with numerical simulations, showing that our method is able to jointly estimate the state trajectory and identify the correct dynamical model.

***Index Terms*—** paired comparisons, time-varying models, dynamical systems, Bayesian inference

1. INTRODUCTION

The problem of estimating a vector using *paired comparisons* between a fixed set of landmark points has been studied in several settings where measurements are only available as comparisons between distances. For example, in a simple two- or three-dimensional setting one might wish to triangulate an object's location using an array of sensors, where the only available information is which of any two given sensors the object is closer to (since exact sensor range measurements might be unavailable or too noisy to be utilized directly). In higher dimensions, paired comparisons have been studied in the context of recommender systems, where users provide preferences between pairs of items [1, 2, 3, 4, 5, 6, 7, 8]. This idea underlies the *ideal point model* [9], in which each user's preferences are encoded in a feature space as a vector that represents the user's "ideal" item, and responses to paired comparisons between items are determined by which of the two items is closer to the user's ideal point. Similar models have been used for learning rankings over items [10] and non-metric multidimensional scaling [11, 12, 13].

Many of these applications can be naturally extended to include time-varying dynamics and system identification. For instance, when performing triangulation using pairs of sensors, the localized object may be one of several vehicles navigating along a path. In this case, one may wish to jointly estimate the vehicle's position, velocity, and acceleration (state estimation), as well as identify the vehicle type from its dynamics (system identification). In a recommender system, a user's preferences may change with time, and these changes may be characteristic of one of several user phenotypes. While the task of selecting pairs and estimating the state vector has been studied in the static case [14, 15, 16, 17], the time-varying setting has not been addressed.

Mathematically, we consider the problem of tracking the evolution of a vector $\boldsymbol{w} \in \mathbb{R}^d$ which varies over time according to an unknown dynamics model f as

$$\boldsymbol{x}_t = \begin{bmatrix} \boldsymbol{w}_t \\ \boldsymbol{v}_t \end{bmatrix}, \quad \boldsymbol{x}_{t+1} = f\left(\boldsymbol{x}_t\right) + \boldsymbol{\nu}_{t+1}, \quad \boldsymbol{x}_0 \sim P_0, \tag{1}$$

where $\boldsymbol{v}_t \in \mathbb{R}^l$ is a vector of latent state variables (e.g., velocity and acceleration) which together with $\boldsymbol{w}_t$ comprise the state vector $\boldsymbol{x}_t \in \mathbb{R}^{d+l}$. The dynamics are perturbed by *innovation noise* $\boldsymbol{\nu}_{t+1} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R})$ with known covariance $\boldsymbol{R}$. We assume that f is drawn from a finite set of candidate dynamics models $\mathcal{F} = \{f_1, f_2, \dots f_K\}$ with prior distribution $p_f(f_i)$, and that the initial state is drawn from a known prior distribution P_0.

At each time step $t = 0, \dots, T$, we can access the state only through binary measurements consisting of paired comparisons that indicate which of two landmark points $\boldsymbol{w}_t$ is closer to. Our task is to actively select a sequence of paired comparisons between landmark pairs to jointly estimate the state trajectory $\boldsymbol{x}_{0:T}$ and identify the true dynamics model f.

In Section 2, we extend ideas for active selection of paired comparisons from the static setting to the time-varying

*Equal contribution. This work was supported in part by NSF grant numbers CCF-1350954, CCF-1350616, and CCF-1409422, ONR grant number N00014-15-1-2619, a gift from the Alfred P. Sloan Foundation, and the National Defense Science and Engineering Graduate (NDSEG) fellowship.

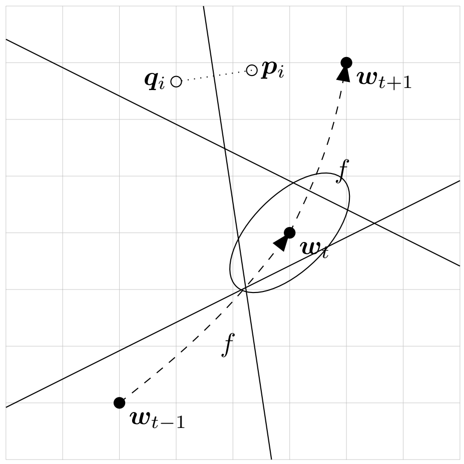

Fig. 1. We consider the problem of using paired comparisons to jointly infer the trajectory of the state $\boldsymbol{w}_{0:T}$ and the dynamics model $f \in \mathcal{F}$ that describes its evolution.

setting, and use these insights in Section 3 to describe our Bayesian particle filter-based method for measurement selection, state tracking, and system identification. We illustrate the operation of our method with an example in Section 4, and evaluate its performance in Section 5.

2. MEASUREMENT SELECTION

At each time step t we take a measurement of the form $\|\boldsymbol{p}_t - \boldsymbol{w}_t\| \gtrless \|\boldsymbol{q}_t - \boldsymbol{w}_t\|$, where $\boldsymbol{p}_t, \boldsymbol{q}_t$ are selected from a known set of landmark points $\mathcal{X} \subset \mathbb{R}^d$. Geometrically, this paired comparison indicates that $\boldsymbol{w}_t$ lies on one side of the hyperplane bisecting points $\boldsymbol{p}_t$ and $\boldsymbol{q}_t$, as illustrated in Figure 1. This hyperplane is defined by normal vector $\boldsymbol{a}_t = \frac{\boldsymbol{p}_t - \boldsymbol{q}_t}{\|\boldsymbol{p}_t - \boldsymbol{q}_t\|}$ and intercept $b_t = \frac{\|\boldsymbol{p}_t\|^2 - \|\boldsymbol{q}_t\|^2}{2\|\boldsymbol{p}_t - \boldsymbol{q}_t\|}$. However, in many practical applications we can access only *noisy* measurements, where comparisons are more likely to be erroneous when $\boldsymbol{w}_t$ is equidistant from the landmark points (i.e., close to the bisecting hyperplane). To represent this type of observation noise, we denote the t^{th} measurement by $Y_t \in \{0, 1\}$, where $Y_t = 1$ indicates that $\boldsymbol{w}_t$ is closer to $\boldsymbol{p}_t$ and $Y_t = 0$ indicates that $\boldsymbol{w}_t$ is closer to $\boldsymbol{q}_t$. We then model Y_t with the logistic likelihood

$$p(Y_t = 1 \,|\, \boldsymbol{w}_t) = \frac{1}{1 + e^{-k(\boldsymbol{a}_t^T \boldsymbol{w}_t - b_t)}}, \tag{2}$$

where k represents the signal-to-noise ratio of the measurements. We assume that Y_t depends only on $\boldsymbol{w}_t$; that is, letting $A \perp\!\!\!\perp B \mid C$ denote that A is conditionally independent of B given C, we have $Y_t \perp\!\!\!\perp \boldsymbol{w}_u \mid \boldsymbol{w}_t$ for $u \neq t$, $Y_t \perp\!\!\!\perp \boldsymbol{v}_{0:T} \mid \boldsymbol{w}_t$, and $Y_t \perp\!\!\!\perp f \mid \boldsymbol{w}_t$. We adopt a Bayesian framework, representing our knowledge of the state $\boldsymbol{x}_t$ and dynamics model f by the posterior densities $p\left(\boldsymbol{x}_t \,|\, y_{0:t-1}\right)$ and $p\left(f_i \,|\, y_{0:t-1}\right)$, where y_t denotes an observed instantiation of Y_t.

A natural question arising in this stochastic measurement model is how to select the measured paired comparison at each time step. In the static setting, it has been shown that some measurements become more informative than others as $\boldsymbol{w}$ is localized, and dramatic improvements in inference are possible by *adaptively* selecting landmark points [17, 14]. In the time-varying setting considered here, at each time step we wish to select the landmark points $(\boldsymbol{p}_t, \boldsymbol{q}_t)$ defining measurement Y_t that provide the most information about the trajectory $\boldsymbol{x}_{0:T}$ and the dynamics model f. We propose a similar approach as [14] by selecting paired comparisons that maximize the *information gain* [18] each measurement provides about both the state trajectory and dynamics model, defined as the mutual information between the measurement Y_t and unknown trajectory $\boldsymbol{x}_{0:T}$ and dynamics f, conditioned on the previous measurements $y_{0:t-1}$:

$$(\boldsymbol{p}_t, \boldsymbol{q}_t) = \underset{\boldsymbol{p}, \boldsymbol{q} \in \mathcal{X}}{\arg\max}\; I\left(Y_t \,;\, \boldsymbol{x}_{0:T}, f \,|\, y_{0:t-1}\right). \tag{3}$$

Intuitively, this quantity represents the amount a paired comparison decreases our uncertainty about the trajectory and dynamics model.

Because we seek to jointly infer the trajectory and dynamics model, it is at first unclear whether one should select measurements that are more informative about the trajectory or about the model. However, the conditional independence of the measurement model in (2) greatly simplifies this design choice: applying the chain rule of mutual information [19] to (3) and simplifying using the conditional independences admitted by our model yields

$$\begin{aligned}
I(Y_t; \boldsymbol{x}_{0:T}, f \,|\, y_{0:t-1}) &= I(Y_t; \boldsymbol{w}_{0:T}, \boldsymbol{v}_{0:T}, f \,|\, y_{0:t-1}) \\
&= I(Y_t; \boldsymbol{w}_{0:T} \,|\, y_{0:t-1}) + I(Y_t; \boldsymbol{v}_{0:T} | \boldsymbol{w}_{0:T}, y_{0:t-1}) \\
&\quad + I(Y_t; f \,|\, \boldsymbol{w}_{0:T}, \boldsymbol{v}_{0:T}, y_{0:t-1}) \\
&= I(Y_t; \boldsymbol{w}_{0:T} \,|\, y_{0:t-1}) \\
&= I(Y_t; \boldsymbol{w}_t \,|\, y_{0:t-1}) + I(Y_t; \boldsymbol{w}_{0:t-1}, \boldsymbol{w}_{t+1:T} | \boldsymbol{w}_t, y_{0:t-1}) \\
&= I(Y_t; \boldsymbol{w}_t \,|\, y_{0:t-1}).
\end{aligned}$$

Therefore, jointly maximizing the information gain with respect to the entire state trajectory and underlying dynamics model is equivalent to simply selecting paired comparisons that maximize the information gain about $\boldsymbol{w}_t$.

3. METHODS

Unfortunately, the measurement likelihood in (2) does not admit a closed-form expression for the information gain $I(Y_t; \boldsymbol{w}_t \,|\, y_{0:t-1})$ of a candidate pair, and approximating it with samples from the posterior $p(\boldsymbol{w}_t | y_{0:t-1})$ is computationally prohibitive when evaluating a large pool of pairs. Instead, we approximate the action of maximizing information gain by using the *mean-cut max-variance* (MCMV) selection strategy of [14], selecting the pair whose bisecting hyperplane cuts

Algorithm 1 MCMV-DF using particle filter

Draw N particles from $p_0(\boldsymbol{x})$ for each candidate dynamics model
for $t = 1, \dots, T$ **do**
  Estimate state $\widehat{\boldsymbol{x}}_t$ using (6)
  Estimate $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ from all particles using (4)-(5)
  $\mathcal{P}_\beta \leftarrow$ downsample set of candidate pairs at rate β
  $(\boldsymbol{p}_t, \boldsymbol{q}_t) \leftarrow \arg\max_{\mathcal{P}_\beta} \sqrt{a_{pq}^T \boldsymbol{\Sigma}_t a_{pq}} - \left| a_{pq}^T \boldsymbol{\mu}_t - b_{pq} \right|$
  $y_t \leftarrow \text{PairedComparison}(\boldsymbol{p}_t, \boldsymbol{q}_t)$, $y_{0:t} \leftarrow y_t \cup y_{0:t-1}$
  for $i = 1, \dots, K$ **do**
    Resample particles using likelihood (2)
    Propagate particles through dynamics (1)
  end for
  Update $p(f_i|y_{0:t})$ using (7)
end for
$f \leftarrow \arg\max_{f_i} p(f_i)$

through the posterior mean in the direction of maximum variance. Specifically, we select measurements by evaluating an acquisition function for each candidate pair in a downsampled pair pool and selecting the maximizing pair, as described in Algorithm 1. This procedure has a computational complexity that scales favorably with the number of candidate pairs, and is a provable approximation to information gain maximization [14].

To select measurement pairs using MCMV, we need to evaluate the posterior mean $\boldsymbol{\mu}_t := \mathbb{E}_{\boldsymbol{w}_t}[\boldsymbol{w}_t|y_{0:t-1}]$ and covariance $\boldsymbol{\Sigma}_t := \mathbb{E}_{\boldsymbol{w}_t}\left[(\boldsymbol{w}_t - \boldsymbol{\mu}_t)^T(\boldsymbol{w}_t - \boldsymbol{\mu}_t)^T|y_{0:t-1}\right]$, which can be computed as

$$\boldsymbol{\mu}_t = \sum_i \mathbb{E}_{\boldsymbol{w}_t}[\boldsymbol{w}_t|f_i, y_{0:t-1}]\, p(f_i|y_{0:t-1}) \tag{4}$$

$$\boldsymbol{\Sigma}_t = \sum_i \mathbb{E}_{\boldsymbol{w}_t}\left[\boldsymbol{w}_t\boldsymbol{w}_t^T|f_i, y_{0:t-1}\right] p(f_i|y_{0:t-1}) - \boldsymbol{\mu}_t\boldsymbol{\mu}_t^T. \tag{5}$$

After taking a measurement, we update the posteriors over $\boldsymbol{x}_t$ and f, estimate the state as the posterior mean

$$\widehat{\boldsymbol{x}}_t := \mathbb{E}[\boldsymbol{x}_t|y_{0:t-1}] = \sum_i \mathbb{E}_{\boldsymbol{x}_t}[\boldsymbol{x}_t|f_i, y_{0:t-1}]\, p(f_i|y_{0:t-1}), \tag{6}$$

and update the dynamics model posterior as

$$\begin{aligned} p(f_i|y_{0:t}) &= \frac{p(y_t|f_i, y_{0:t-1})p(f_i|y_{0:t-1})}{p(y_t|y_{0:t-1})} \\ &= \frac{\mathbb{E}_{\boldsymbol{w}_t}[p(y_t|\boldsymbol{w}_t)|f_i, y_{0:t-1}]\, p(f_i|y_{0:t-1})}{\sum_j \mathbb{E}_{\boldsymbol{w}_t}[p(y_t|\boldsymbol{w}_t)|f_j, y_{0:t-1}]\, p(f_j|y_{0:t-1})}. \end{aligned} \tag{7}$$

We observe that to calculate each of these quantities we can simply track a separate state posterior $p(\boldsymbol{x}_t|y_{0:t-1}, f_i)$ for each candidate dynamical system $i = 1, \dots, K$, from which we can compute the necessary expected values.

Because our pairwise measurement process (2) is nonlinear, we cannot use the closed form updates of the Kalman filter to track each posterior $p(\boldsymbol{x}_t|y_{0:t-1}, f_i)$. Instead, we use the *particle filter*, which allows us to incorporate both the nonlinear likelihood and an arbitrary (potentially nonlinear) candidate dynamics models [20]. In the particle filter framework, the required probability distributions are represented by Monte Carlo particles which can be propagated through the candidate dynamics models f_i. We use N particles to represent the state posterior associated with each candidate dynamics model, resulting in a total of NK tracked particles. We present our algorithm in its entirety, called *mean-cut max-variance dynamic filtering* (MCMV-DF), in Algorithm 1.

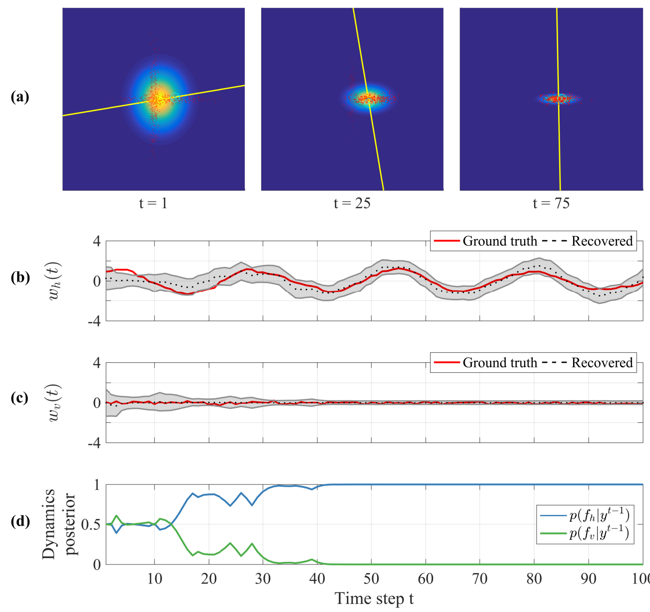

Fig. 2. Stylized demonstration of MCMV-DF. (a) Posterior position distributions at three time instants. Surface plots: state posterior $p(\boldsymbol{w}_t|y_{0:t-1})$; red circles: particles representing $p(\boldsymbol{w}_t|y_{0:t-1}, f_i)\ \forall\ f_i$ with opaqueness representing $p(f_i|y_{0:t-1})$; cyan target: true position $\boldsymbol{w}_t$; yellow line: hyperplane corresponding to selected measurement $(\boldsymbol{p}_t, \boldsymbol{q}_t)$. (b-c) True trajectory and recovered marginal posterior $p(\boldsymbol{w}_t|y_{0:t-1})$ for horizontal and vertical components of position; shaded region corresponds to 95% confidence interval on posterior. (d) Posterior over dynamics models $p(f_h|y_{0:t-1})$ and $p(f_v|y_{0:t-1})$.

4. EXPLANATORY EXAMPLE

Figure 2 illustrates our approach with a stylized numerical example. We track a point $\boldsymbol{w} \in \mathbb{R}^2$ evolving purely along the *horizontal* axis according to a spring-like system, with latent state $\boldsymbol{v} \in \mathbb{R}^4$ representing velocity and acceleration in each dimension. We consider $K = 2$ candidate dynamics models: the true dynamical system f_h, and a similar system f_v that evolves purely along the *vertical* axis. We run MCMV-DF for $T = 100$ time steps and observe the inferred state

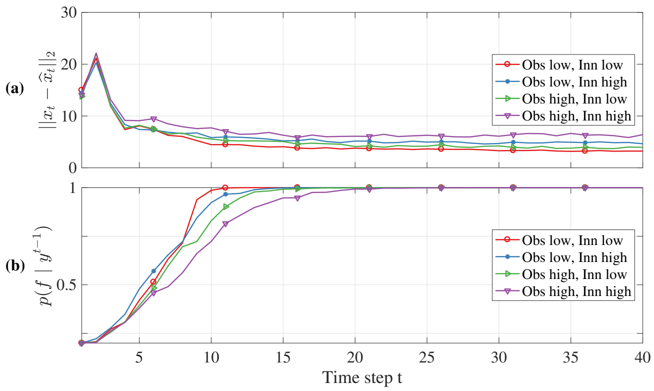

Fig. 3. Tracking performance as observation noise ("obs") and innovation noise ("inn") levels change; each point shows the median over 200 trials. (a) Trajectory reconstruction accuracy, shown as error $\|\boldsymbol{x}_t - \widehat{\boldsymbol{x}}_t\|_2$ at each time step. (b) Dynamical system identification, shown as posterior probability of true system $p(f|y_{0:t-1})$.

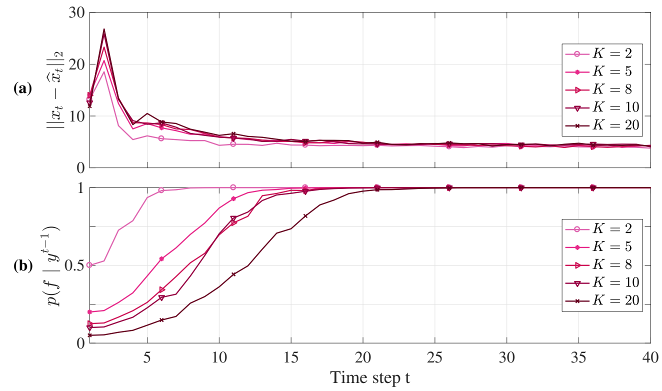

Fig. 4. Tracking performance as the number of candidate dynamics models K increases; each point shows the median over 200 trials. (a) Trajectory reconstruction accuracy, shown as error $\|\boldsymbol{x}_t - \widehat{\boldsymbol{x}}_t\|_2$ at each time step. (b) Dynamical system identification, shown as posterior probability of true system $p(f|y_{0:t-1})$.

and dynamics model posteriors. As MCMV-DF converges to correctly identify the true (horizontal) dynamical system f_h at approximately $t = 40$ (as shown in (d) by the posterior probability $p(f_h|y_{0:t-1})$ approaching unity), the vertical position estimate becomes more accurate. This is reflected in both the tight distribution of probability mass around $\widehat{w}_v(t)$ in (a) and (c) and the closer to vertical orientation of the hyperplanes corresponding to measurements y_{25} and y_{75} in (a). This vertical orientation maximizes variance after the trajectory has been identified as purely along the horizontal axis. As the measurements begin to focus on accurately estimating the horizontal position, the horizontal position estimates also become more accurate, as displayed in (b).

5. NUMERICAL EXPERIMENTS

In this section, we demonstrate the performance of MCMV-DF with simulations on synthetic data, evaluating the effects of observation and innovation noise as well as the number of candidate dynamical systems K on the accuracy of trajectory estimation and system identification. In both experiments, we randomly generate an initial state $\boldsymbol{x}_0 \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$ with dimensionality $d = l = 4$ and compute its trajectory using (1) and f with $\boldsymbol{R} = \sigma^2 \boldsymbol{I}$ for various settings of σ^2. We generate 1500 landmark points, distributed in $\mathbb{R}^4$ as $\mathcal{N}(\boldsymbol{0}, \sigma_p^2 \boldsymbol{I})$ with $\sigma_p^2 = 9$, and use the downsampling rate $\beta = 0.01$ when selecting landmark points for measurements.

In each trial, we generate K random linear dynamics models $f_1, \ldots, f_K$ by placing $d + l$ eigenvalues in complex conjugate pairs on the unit circle (making the resulting systems marginally stable) at angles distributed as $\theta \sim U\left[\frac{\pi}{6}, \frac{\pi}{3}\right]$ (controlling the velocity of the resulting trajectories), with random orthogonal eigenvectors. We arbitrarily select one of the K candidate dynamics models as the true system.

In Figure 3, we evaluate the accuracy of tracking and identification with four different noise levels. To set the observation noise level, we vary the signal-to-noise constant k in (2): "low" observation noise corresponds to $k = 10$ (resulting in approximately 5% incorrect comparisons), and "high" observation noise corresponds to $k = 1$ (resulting in approximately 25% incorrect comparisons). The "low" and "high" values of innovation noise are $\sigma^2 = 10^{-3}$ and $\sigma^2 = 10^{-2}$, respectively. Figure 3 shows the tracking error of the *entire* state $\boldsymbol{x}$, $\|\boldsymbol{x}_t - \widehat{\boldsymbol{x}}_t\|_2$, and posterior probability of the true dynamics model $p(f|y_{t-1})$ over a horizon of $T = 40$ time steps. We observe that MCMV-DF successfully identifies the true dynamics system in a modest number of measurements and quickly achieves low state estimation error. Increasing the observation and innovation noise reduces estimation accuracy and increases the number of time steps required to identify the true dynamical system, but our method still tracks the state and eventually recovers the correct dynamics model.

In Figure 4, we evaluate the effect of the number of candidate dynamics models K on MCMV-DF's performance with fixed noise levels $k = 3$ (resulting in approximately 15% incorrect comparisons) and $\sigma^2 = 5 \times 10^{-3}$. We observe that the higher complexity of the set of candidate systems $\mathcal{F}$ resulting from increasing K makes the problem harder, reducing tracking accuracy and increasing the number of time steps required to identify the true dynamical system; however, system recovery is still possible.

Overall, these results demonstrate MCMV-DF's ability to successfully estimate the state trajectory from intelligently selected paired comparisons and discern between multiple candidate dynamics models. Further study is warranted to evaluate MCMV-DF's performance in real-world systems.

6. REFERENCES

[1] L. Qian, J. Gao, and H.V. Jagadish, "Learning user preferences by adaptive pairwise comparison," *Proc. VLDB Endowment*, vol. 8, no. 11, pp. 1322–1333, 2015.

[2] E. Brochu, N. Freitas, and A. Ghosh, "Active preference learning with discrete choice data," in *Proc. Adv. in Neural Inf. Processing Systems (NeurIPS)*, 2008.

[3] L. Wu, C. Hsieh, and J. Sharpnack, "Large-scale collaborative ranking in near-linear time," in *Proc. ACM Int. Conf on Knowledge Discovery and Data Mining (KDD)*, 2017.

[4] F. Wauthier, M. Jordan, and N. Jojic, "Efficient ranking from pairwise comparisons," in *Int. Conf. on Machine Learning (ICML)*, 2013, pp. 109–117.

[5] Z. Cao, T. Qin, T. Liu, M. Tsai, and H. Li, "Learning to rank: from pairwise approach to listwise approach," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2007, pp. 129–136.

[6] Y. Chen and C. Suh, "Spectral mle: Top-k rank aggregation from pairwise comparisons," in *Int. Conf. on Machine Learning (ICML)*, 2015, pp. 371–380.

[7] B. Eriksson, "Learning to top-k search using pairwise comparisons," in *Proc. Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2013, pp. 265–273.

[8] N. Shah and M. Wainwright, "Simple, robust and optimal ranking from pairwise comparisons.," *Journal of Machine Learning Research*, vol. 18, pp. 199–1, 2017.

[9] C. Coombs, "Psychological scaling without a unit of measurement.," *Psychological review*, vol. 57, no. 3, pp. 145, 1950.

[10] K. Jamieson and R. Nowak, "Active ranking using pairwise comparisons," in *Proc. Adv. in Neural Information Processing Systems (NeurIPS)*, 2011, pp. 2240–2248.

[11] S. Agarwal, J. Wills, L. Cayton, G. Lanckriet, D. Kriegman, and S. Belongie, "Generalized non-metric multidimensional scaling," in *Proc. Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2007, pp. 11–18.

[12] O. Tamuz, C. Liu, S. Belongie, O. Shamir, and A. Kalai, "Adaptively learning the crowd kernel," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2011.

[13] L. Van Der Maaten and K. Weinberger, "Stochastic triplet embedding," in *Proc. IEEE Int. Work. on Machine Learning for Signal Processing (MLSP)*, 2012, pp. 1–6.

[14] G. Canal, A. Massimino, M. Davenport, and C. Rozell, "Active embedding search via noisy paired comparisons," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2019, pp. 902–911.

[15] M. Davenport, "Lost without a compass: Nonmetric triangulation and landmark multidimensional scaling," in *IEEE Int. Work. on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2013.

[16] M. O'Shaughnessy and M. Davenport, "Localizing users and items from paired comparisons," in *Proc. IEEE Int. Work. on Machine Learning for Signal Process. (MLSP)*, 2016.

[17] A. Massimino and M. Davenport, "As you like it: localization via paired comparisons," *arXiv preprint arXiv:1802.10489*, 2018.

[18] D. MacKay, "Information-based objective functions for active data selection," *Neural computation*, vol. 4, no. 4, pp. 590–604, 1992.

[19] T. Cover and J. Thomas, *Elements of information theory*, John Wiley & Sons, 2012.

[20] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," in *IEE proceedings F (radar and signal processing)*, 1993, vol. 140, pp. 107–113.