

The Role of Dynamics in Extracting Information Sparsely Encoded In High Dimensional Data Streams*

M. Sznai¹ O. Camps¹ N. Ozay¹ T. Ding² G. Tadmor¹ D. Brooks¹

¹ ECE Department, Northeastern University, Boston, MA 02115

² Dept. of Elect. Eng., Penn State University, University Park, PA 16802

Summary. A major roadblock in taking full advantage of the recent exponential growth in data collection and actuation capabilities stems from the *curse of dimensionality*. Simply put, existing techniques are ill-equipped to deal with the resulting overwhelming volume of data. The goal of this chapter is to show how the use of simple dynamical systems concepts can lead to tractable, computationally efficient algorithms for extracting information sparsely encoded in multimodal, extremely large data sets. In addition, as shown here, this approach leads to non-entropic information measures, better suited than the the classical, entropy-based information theoretic measure, to problems where the information is by nature dynamic and changes as it propagates through a network where the nodes themselves are dynamical systems.

1.1 Introduction




Fig. 1.1. Examples of sparsely encoded information. (a) Detecting a traffic accident. (b) Tracking a person in a crowd. (c) *gfp*-visualized promoter activation during a diauxic shift experiment in *E. coli* [1]. In all cases fewer than $\mathcal{O}(10^{-3})$ to $\mathcal{O}(10^{-6})$ of the data is relevant.

The recent exponential growth in data collection and actuation capabilities has the potential to profoundly impact society, with benefits ranging from safer, self

* This work was supported in part by NSF grants IIS-0713003 and ECCS-0901433, AFOSR grant FA9550-09-1-0253, and the Alert DHS Center of Excellence.

aware environments, to enhanced image-based therapies. A major road-block to realizing this vision stems from the curse of dimensionality. Simply put, existing techniques are ill-equipped to deal with the resulting overwhelming volume of data.

This chapter discusses the key role that dynamics can play in timely extracting and exploiting actionable information that is very sparsely encoded in high dimensional data streams. Its central theme is the use of dynamical models as information encoding paradigms. Our basic premise is that spatio-temporal dynamic information can be compactly encapsulated in dynamic models, whose rank, a measure of the dimension of useful information, is often far lower than the raw data dimension. This premise amounts to a reasonable localization hypothesis for spatio-temporal correlations, and is a given in mechanical and biological processes. Embedding problems in the conceptual world of dynamical systems makes available a rich, extremely powerful resource base, leading to robust solutions, or, in cases where the underlying problem is intrinsically hard, to computationally tractable approximations with sub-optimality certificates. For instance, in this context, changes in the underlying process can be detected by simply computing the rank of a Hankel matrix constructed from the data and missing information can be recovered by solving a rank minimization problem that can be relaxed to a tractable semi-definite program. A third example is the comparison of data streams in order to establish whether they correspond to time traces of the same phenomenon: it is vastly easier to quantify the difference between two dynamic models (often requiring only a rank comparison) than to search for elusive overlapping time sequences and then compare two such very high dimensional data streams. Finally, the use of dynamic models leads naturally to non-entropic information measures, better suited for problems where the information is by nature dynamic and changes as it propagates through a network where the nodes themselves are dynamical systems. These ideas are illustrated with several examples from different applications, including change detection in video sequences, motion segmentation, and uncovering co-promoted genes.

1.2 Key subproblems arising in the context of dynamic information extraction

The challenges entailed in exploiting dynamic information sparsely encoded in very large data sets are illustrated in Figure 1.1: In all cases, decisions must be taken based on events discernible only in a small fraction of a very large data record: a short video sequence adds up to megabytes, yet the useful information (a change of behavior of a single target), may be encoded in just a portion of a few frames, e.g., less than 10^{-6} of the total data. Similarly, the data from the diauxic shift experiment shown in Figure 1.1(c) consists of 342×10^3 data points from the time traces of 1,920 promoters, (e.g., a total of 19 Mb of data), yet only a few critical time instants and promoter correlations are of interest. Additional challenges arise from the quality of the data, often fragmented and corrupted by noise. Addressing these challenges requires solving the following subproblems:

A: Nonlinear embedding of dynamic data. Finding low dimensional manifold structures in data, a hallmark of machine learning, is a key precursor to both dimensionality reduction and robust information extraction. Existing static techniques ([2] and references therein) provide low dimensional embeddings, but fail to exploit the large gap between data dimension and dynamic rank. As we will show in this chapter, this can be accomplished by employing low rank dynamic models to capture time/parameter dependence on low dimensional manifolds that maximally absorb stationary high dimensions and nonlinearities.

B: Uncovering structures embedded in data. A key step in information extraction is the ability to find structures embedded in the data. For example, when analyzing data generated by an unknown number N_o of sources, it is of interest to identify the number of sources, associated substreams, and the individual dynamics. This is commonly accomplished by searching for statistical correlations or exploiting *a priori* known structural constraints. For example, independently moving objects in a video clip are efficiently detected by factorizing a matrix of feature trajectories [3, 4, 5]. However, methods based on correlations and (application dependent) *a priori* information alone are fragile to missing/corrupted data and have trouble disambiguating structures with overlapping *kinematic* or *statistical properties*. As shown here, these difficulties can be avoided by seeking *dynamically coherent* substreams, e.g., subsets that can be jointly explained by low rank models. Further, this task can be efficiently carried out *without* explicitly finding these models, by estimating ranks of Hankel matrices constructed from time traces. Incorporating *a priori* available information allows for retaining the advantages of existing methods while substantially improving robustness.

C: Dynamic data segmentation. The goal here is to partition the data record into maximal, disjoint sets within which the data satisfies a given predicate. Examples include segmenting a video sequence of a person into its constituent activities, or identifying time periods where a given group of gene promoters is active. While this problem has been the object of considerable research in the past decade, it remains very challenging in cases involving noisy data, where most existing methods lead to computationally demanding problems [6, 7], with poor scaling properties. As we will show in the sequel, the use of dynamics provides a unified, efficient approach to robust segmentation. In its simplest form, the idea is to group data according to the complexity of the model that explains it. Intuitively, models associated with homogeneous data, e.g., a single activity or metabolic stage, have far lower complexity than those jointly explaining multiple datasets. Boundaries are thus characterized by a step increase in model complexity. In turn, these jumps in model complexity can be efficiently detected by examining the singular values of a matrix directly constructed from the data.

D: Dynamic interpolation. Data streams are often fragmented: clinical trial patients may miss appointments, targets may be momentarily occluded. The challenges here are to (i) identify fragments belonging to the same data sets (for instance “tracklets” corresponding to a track of a single target, fragmented due to occlusion), and (ii) interpolate the missing data while preserving relevant dynamical invariants embedded in it. The latter is particularly important in cases where a transition is mediated

by the missing data. An example is detecting an activity change from video data, when the transition point is occluded. Formulating the problem as a minimum order dynamical interpolation one leads to computationally attractive solutions, whereby values for missing data are selected as those that do not increase the complexity - or rank - of the model underlying the data record.

E: Hypothesis testing and distributed information sharing. Examples include determining whether (possibly non-overlapping) data streams correspond to the same process or assessing whether a data set is a realization of a given process. In turn, this entails computing worst-case distances between data and model predictions, a task that can be efficiently accomplished by combining concepts from dynamical systems and information based complexity. Situations involving multiple information sources and users require the ability to (i) maintain consistent data labeling across sources, and (ii) mitigate the communications and computational burdens entailed in sharing very large datasets. Both issues can be efficiently addressed by exploiting the dynamical models underlying the data. Briefly, the idea is to identify a dynamic operator mapping the dynamic evolution of data projections over individual manifolds, amounting to a *dynamical registration* between sources. Sharing/comparing data streams then entails transmitting only the (low order) projections of dynamic variables and running these projections through the interconnecting operator.

In the remainder of this chapter we show how the use of dynamic models that compactly encapsulate relevant spatio-temporal information provides a unified set of tools leading to computationally efficient solutions to problems A –E above. In all cases, these solutions will be illustrated with practical examples.

1.3 Nonlinear embedding of dynamic data.

In the past few years, considerable research has been devoted to the problem of nonlinear dimensionality reduction via manifold embedding. Briefly, the idea is to obtain lower complexity data representations by embedding it into low dimensional nonlinear manifolds while preserving spatial neighborhoods. Commonly used methods include locally linear embeddings (LLE) [8], Isomap [9], Laplacian Eigenmaps [10], Hessian LLE [11], and Semidefinite Embedding [12, 13]. These methods successfully exploit spatial correlations to achieve (often substantial) dimensionality reduction. However, they fail to take advantage of the temporal correlations that are characteristic of dynamic data. As we show next, constraining target manifolds to those spanned by feasible *dynamical* trajectories enables additional (substantial) dimensionality reduction and provides robustness against missing data and outliers.

The starting point is the realization that since projections to/from manifolds can be modeled as memoryless nonlinearities, the problem of jointly identifying the embedding manifold, the dynamics characterizing the evolution of the data on this manifold, and the projection operators can be recast into the Hammerstein/Wiener system identification problem illustrated in Figure 1.2. Here $\Pi_i(\cdot)$ and $\Pi_o(\cdot)$ are memoryless nonlinearities, S is a linear time invariant (LTI) system that describes the temporal evolution of the data on the manifold, and $\mathbf{u} \in R^{n_u}$, $\mathbf{d} \in R^{n_d}$, $\mathbf{u}_m \in R^{n_{u_m}}$ and




Fig. 1.2. Hammerstein/Wiener System Structure. Wiener system.

$\mathbf{y} \in R^{n_y}$, with $n_d \gg n_y$, $n_u \gg n_{u_m}$ represent the respective input (for instance a vector composed of past values of the output and a stochastic driving signal), the raw data, and their projections on the low dimensional manifold. A potential difficulty here stems from the fact that, as recently shown in [14], robust identification of Hammerstein/Wiener systems is generically NP-hard. However, efficient, computationally tractable relaxations that scale polynomially with problem size (both manifold dimension and number of temporal data points) can be obtained by pursuing a *risk-adjusted* approach. The main idea is to identify first the (piecewise) linear dynamics by sampling the set of signals $(\mathbf{u}_m, \mathbf{y})$, and attempting to find, for each sample (typically a subset of a ball in ℓ^2) an LTI operator S (the dynamics on the manifold) compatible with existing *a priori* information and such that $\mathbf{y} = S\mathbf{u}_m$. As shown in [15, 16] both steps reduce to a convex optimization problem via the use of Parrot's Theorem on norm-preserving matrix expansions and standard results on interpolation. The effectiveness of this approach is illustrated in Figure 1.3, where the application of these ideas enabled sustained tracking of multiple subjects in a cluttered outdoor scene. Here, recasting the problem into a non-linear identification form allowed for reducing the problem to an identification/prediction one in a 3-dimensional manifold.

It is worth emphasizing that this approach has the, hitherto unavailable, ability to exploit the synergy between the data embedding and dynamic modeling problems to improve robustness and computational properties. Robustness is improved by *automatically* discarding manifolds incompatible with *a priori* existing information on the dynamics, while computationally attractive models result from maximally absorbing nonlinearities in the manifold structure. Further, the *consistency set* [17] associated with the identification problem provides the means to (in)validate assumptions about the geometry of the manifolds and to quantify the approximation error. Thus, viewing data as a manifestation of hidden dynamics allows a synergy between machine learning (the manifold structure), identification theory (theoretical underpinnings, computational framework) and information based complexity (worst case prediction-error bounds).

1.4 Structure extraction from high dimensional data streams.

Structure extraction methods based on correlations and (application dependent) *a priori* information alone are often fragile to missing/corrupted data and have trouble disambiguating structures with overlapping *kinematic* or *statistical properties*. As

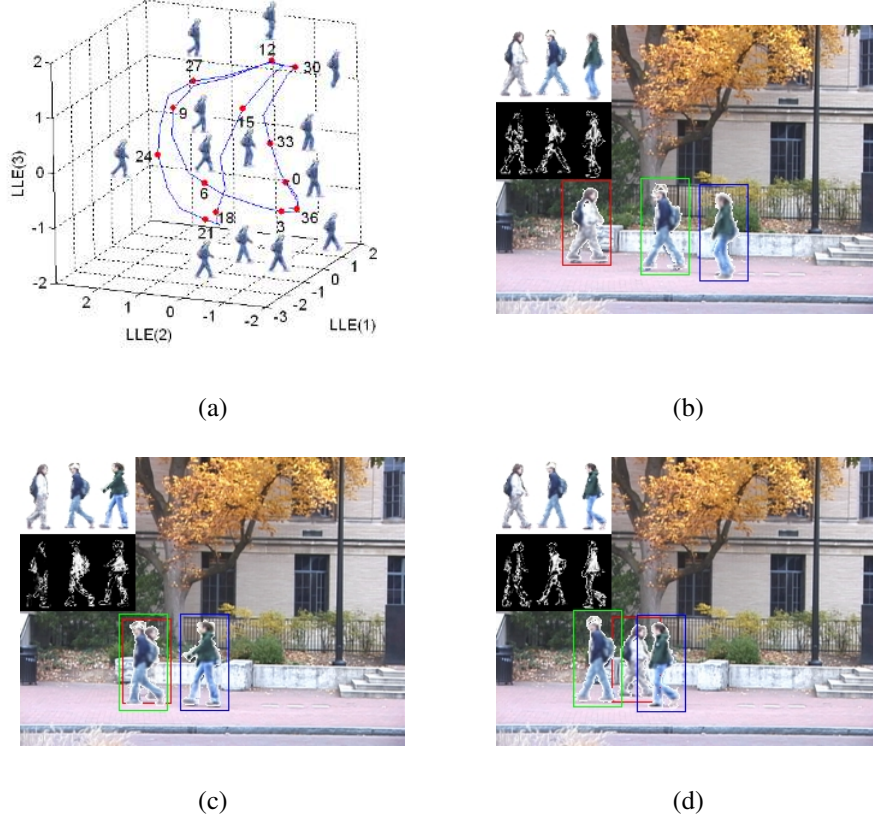


Fig. 1.3. (a): Sample 3 dimensional manifold extracted from a walking sequence. (b)-(d): use of dynamics on this manifold to predict target position and appearance.

an illustrative example, consider time traces $\mathbf{p}_{ti} = (u_{ti}, v_{ti})^T, t = 1 \dots n$ of n_p features $P_i, i = 1, \dots, n_p$, from a single rigid object. Kinematic constraints imply that the rank of the “measurements” matrix $\mathcal{W}_{1:F} \doteq [\mathbf{p}_{ti}] \in \mathbb{R}^{2n \times n_p}$ is at most 4 [18]. The number N_o of *independent* rigid bodies can thus be estimated by factorizing that matrix into rank 4 submatrices. Yet this approach fails to disambiguate objects with partially shared motion, as illustrated in Figure 1.4(a): Here $\text{rank}(\mathcal{W}) = 7$ due to shared propeller rotations; hence any segmentation based solely on factorizing \mathcal{W} will fail to distinguish this case from the case of just two independently moving propellers. The root-cause is that properties that are invariant under row permutations in \mathcal{W} are limited to revealing geometric dependencies but *ignore dynamic constraints*³. As shown next, these ambiguities can be solved through the use of dynamical models that exploit both sets of constraints.

³ Any permutation of the rows of \mathcal{W} satisfies the same geometric constraints, but corresponds to different time trajectories.




Fig. 1.4. (a) Right and left wing propellers move in opposite directions at the same speed. (b) Dynamics based segmentation. (c) Costeira-Kanade segmentation. (d) Zelnik-Manor-Irani segmentation. (e) GPCA segmentation.

The starting point is the realization that for two points $\mathbf{p}_r, \mathbf{p}_s$ belonging to the same source, the time evolution of $y_{r,s}(k) \doteq \mathbf{p}_r(k) - \mathbf{p}_s(k)$ does not carry information about the overall *group* motion of the source. Equivalently, states associated with group motion are *unobservable* from $y_{r,s}$ if \mathbf{p}_r and \mathbf{p}_s belong to the same dynamic cluster. Hence the associated Hankel matrix is rank deficient [17] vis-a-vis the case of points from different sources. This leads to the following simple dynamic clustering algorithm:

- (i) For each pair (r, s) , form the Hankel matrix $\mathbf{H}_{y_{r,s}}$ of pairwise differences $y_{r,s}(k) = \mathbf{p}_r(k) - \mathbf{p}_s(k)$:

$$\mathbf{H}_y = \begin{bmatrix} \mathbf{y}(1) & \mathbf{y}(2) & \cdots & \mathbf{y}(\frac{n}{2}) \\ \mathbf{y}(2) & \mathbf{y}(3) & \cdots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}(\frac{n}{2}) & \cdots & \cdots & \mathbf{y}(n) \end{bmatrix} \quad (1.1)$$

- (ii) Group points according to the minimum value of $\text{rank}[\mathbf{H}_{y_{r,s}}]$.

In this context, robust handling of noisy measurements $\hat{y}(k) = y(k) + \eta(k)$, is accomplished by simply replacing “rank” by the number of singular values above the covariance of the measurement noise⁴, leading to an algorithm computationally no more expensive than a sequence of SVDs. The effectiveness of this approach is illustrated in Figure 1.4 where darker matrix elements indicate higher correlations:

⁴ In this case $\mathbf{H}_{\hat{y}} = \mathbf{H}_y + \mathbf{H}_\eta$, and, under ergodicity assumptions, $\mathbf{H}_\eta^T \mathbf{H}_\eta$ is an estimate of the covariance matrix of the noise.

As shown there, the dynamics based approach achieves perfect segmentation, while methods relying solely on factorizations of \mathcal{W} [5, 19, 20] fail.

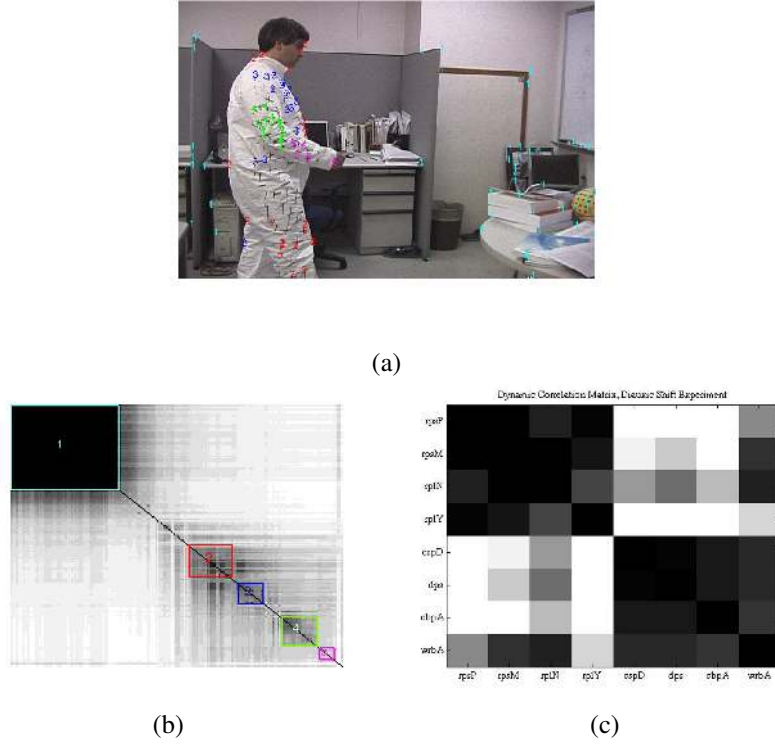


Fig. 1.5. (a) Sample frame. (b) Structures found using dynamic rank (darker color indicates higher dynamic correlation). The hierarchy in the lower right corner corresponds to different portions of the body. (c) Dynamic correlation between genes in the diauxic shift experiment of Figure 1.1(c). The two identified groups correspond to growth related (top left) and stationary (bottom right) genes. The fainter correlation between *wrbA* and (*rpsM*, *rplN*) was unexpected.

An interesting property of the dynamics based approach to segmentation, illustrated in Figure 1.5, is the ability to provide a *hierarchical* segmentation according to the complexity of the joint dynamics. This is key to model the behavior of a target composed by several components acting in a dynamically correlated fashion, e.g., the limbs of a walking person or co-regulated genes. The aggregate behaves as a non-rigid object, whose components share motion modes.

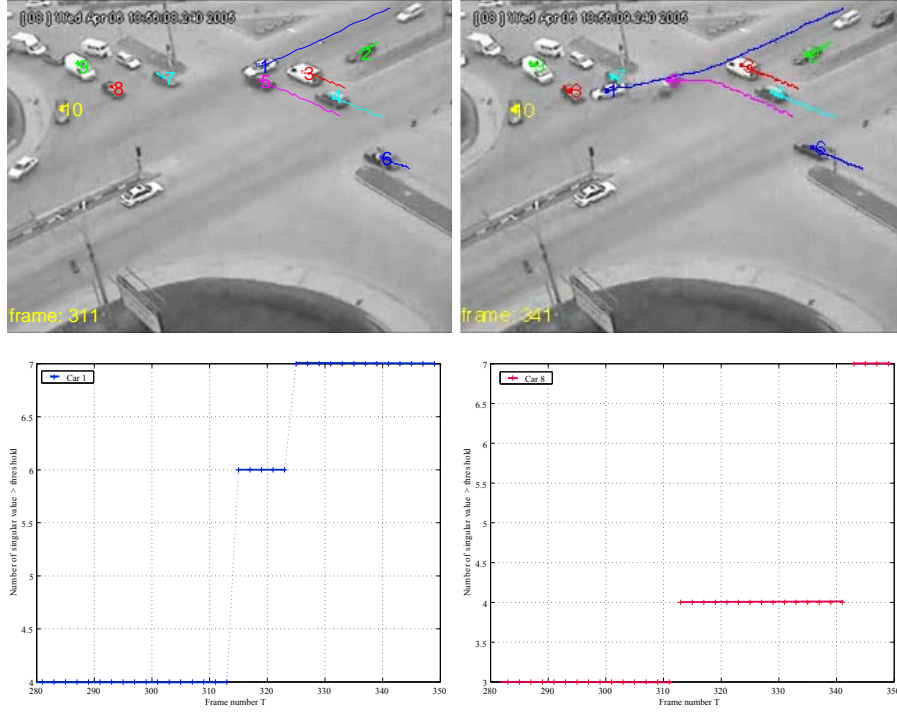


Fig. 1.6. Crash detection. Top: Frames 311 and 341. Bottom: Hankel rank time traces: Car 1 (b) and Car 8 (c).

1.5 Robust Dynamic Data Segmentation.

In principle, changes in the processes underlying a given data record can be detected by a two tiered approach: identification of an underlying set of models (the consistency set) followed by a model (in)validation step to detect points at which new data are inconsistent with all the models in the set. However, the entailed computational complexity is high, roughly n^5 for n data points. A fast, computationally efficient alternative can be obtained by searching for points where the complexity of the underlying model changes. The main idea behind this approach is the fact that models associated with homogeneous data have far lower complexity than those jointly explaining multiple datasets. Further, the complexity of the (unknown) model can be estimated from the experimental data by computing the number $N_{sv,\sigma}(H_y)$ of (significant) singular values of a Hankel matrix similar to H_y in (1.1). Hence the data record can be segmented according to discontinuities in $N_{sv,\sigma}(H_y)$. Figure 1.6 illustrates the effectiveness of this approach in detecting contextually abnormal behavior - an accident - evidenced by a jump in the Hankel rank. An application of this technique to detecting changes in promoter activity in *E. coli* is shown in Figure 1.7.




Fig. 1.7. Detecting transitions in an *E. coli* culture via Hankel rank. Jumps at 20 and 57 correspond to shifts from metabolizing glucose to lactose, to stationary phase, respectively.

The approach outlined above works well for cases where the noise is moderate and adequately characterized as an ℓ^2 bounded signal. Cases where these conditions do not hold (for instance ℓ^∞ noise) can be handled by a modification of this idea (detecting mode changes in piecewise affine models) as follows. The starting point is the assumption that the data record has been generated by a piecewise affine model of the form:

$$\mathcal{H} : f\left(\mathbf{p}_{\sigma(t)}, \{\mathbf{x}(k)\}_{k=t-i}^{t+j}\right) = \eta_f \quad (1.2)$$

where f is an affine function⁵ of the parameter vector $\mathbf{p}_{\sigma(t)}$ which takes values from a finite unknown set according to a piecewise constant function $\sigma(t)$, and η_f denotes an unknown noise signal. Here i and j are positive integers that account for the memory of the model (e.g., $j = 0$ corresponds to a causal model, or $i = j = 0$ corresponds to a memoryless model). Next, consider the sequence of *first order differences* of the parameters $\mathbf{p}_{\sigma(t)}$, given by

$$\mathbf{g}(t) = \mathbf{p}_{\sigma(t)} - \mathbf{p}_{\sigma(t+1)} \quad (1.3)$$

Clearly, a non-zero element of this sequence corresponds to a *change* in the underlying model. Hence, partitioning the data record into maximal homogeneous sequences is equivalent to finding a hybrid model of the form (1.2), consistent with the *a priori* information (e.g., a bound on $\|\eta\|_{\ell^\infty}$) and experimental data, such that the number of non-zero elements of the vector $\mathbf{g}(\cdot)$ is minimized. Formally, defining

⁵ That is: $f\left(\mathbf{p}_{\sigma(t)}, \{\mathbf{x}(k)\}_{k=t-i}^{t+j}\right) = A(\mathbf{x})\mathbf{p}_{\sigma(t)} + \mathbf{b}(\mathbf{x})$

$\delta(t) = \|\mathbf{g}(t)\|_\infty$, the objective is to minimize $\|\delta\|_{\ell^0}$, the number of non-zero elements of δ , subject to (1.2). Using the fact that the *convex envelope* of $\|\cdot\|_{\ell^0}$ in R^N is the ℓ^1 -norm [21], this non-convex problem can be relaxed to:

$$\begin{aligned} & \text{minimize}_{\mathbf{p}(t), \eta(t)} \|\{\mathbf{g}\}\|_{\ell_1} \\ & \text{subject to} \quad f\left(\mathbf{p}(t), \{\mathbf{x}(k)\}_{k=t-i}^{t+j}\right) = \eta(t) \quad \forall t \\ & \quad \|\{\eta\}\|_* \leq \epsilon \end{aligned} \quad (1.4)$$

Since f is an affine function of $\mathbf{p}(t)$, (1.4) has a convex feasibility set \mathcal{F} . Thus, using the ℓ_1 norm leads to a convex, computationally tractable relaxation. The resulting solution can be further improved using the iterative procedure proposed in [22], based on solving, at each iteration, the following weighted ℓ_1 -norm minimization over the convex feasible set \mathcal{F} :

$$\begin{aligned} & \text{minimize}_{z, g, p, \eta} \sum_{t=1}^{T-1} w_t^{(k)} z_t \\ & \text{subject to} \quad \|\mathbf{g}(t)\|_\infty \leq z_t \quad \forall t \\ & \quad f\left(\mathbf{p}(t), \{\mathbf{x}(k)\}_{k=t-i}^{t+j}\right) = \eta(t) \quad \forall t \\ & \quad \|\{\eta\}\|_* \leq \epsilon \end{aligned} \quad (1.5)$$

where $w_i^{(k)} = (z_i^{(k)} + \delta)^{-1}$ are weights with $z_i^{(k)}$ being the arguments of the optimal solution at the k^{th} iteration and $z^{(0)} = [1, 1, \dots, 1]^T$; and where δ is a (small) regularization constant that determines what should be considered zero.

The choice of $*$, the norm characterizing the noise, is application dependent. For instance the ℓ^∞ -norm performs well in finding anomalies, since in this case the change detection algorithm looks for *local* errors, highlighting outliers. On the other hand, when a bound on the ℓ^1 or ℓ^2 -norm of the noise is used, the change detection algorithm is more robust to outliers and it favors the continuity of the segments (i.e., longer subsequences). In addition, when using these norms, the optimization problem automatically adjusts the noise distribution among the segments, better handling the case where the noise level is different in different segments.

1.5.1 Example 1: Video Segmentation.

Segmenting and indexing video sequences have drawn significant attention due to the increasing amounts of data in digital video databases. Systems that are capable of segmenting video and extracting key frames that summarize the video content can substantially simplify browsing these databases over a network and retrieving important content. An analysis of the performances of early shot change detection algorithms is given in [23]. The methods analyzed in [23] can be categorized into two major groups: i) methods based on histogram distances, and ii) methods based on variations of MPEG coefficients. A comprehensive study is given in [24] where a formal framework for evaluation is also developed. Other methods include those where scene segmentation is based on image mosaicking [25, 26] or frames are segmented according to underlying subspace structure [27].

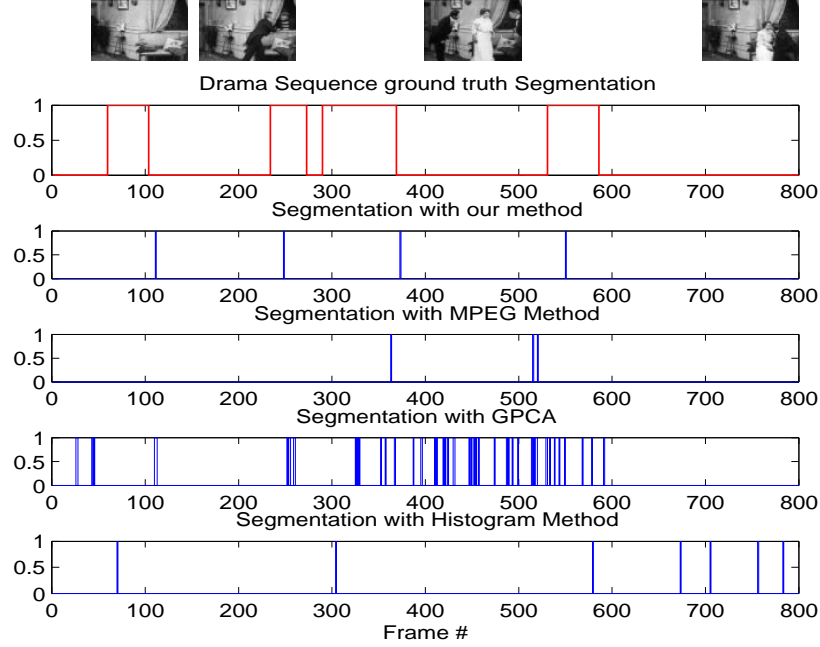


Fig. 1.8. Video segmentation as a hybrid system identification

Given a video sequence of frames $\{\mathcal{I}(t) \in \mathbb{R}^D\}_{t=1}^T$, the video segmentation problem can be solved by first projecting the data into a lower dimensional space, using for instance Principal Component Analysis (PCA), and then applying the sparsification algorithm described above to the projected data (to exploit the fact that the number of pixels D is usually much larger than the dimension of the subspace where the frames are embedded):

$$\mathcal{I}(t) \mapsto \mathbf{x}(t) \in \mathbb{R}^d.$$

Assuming that each $\mathbf{x}(t)$ within the same segment lies on the same hyperplane not passing through the origin⁶ leads to the following hybrid model:

$$\mathcal{H}_1 : f(\mathbf{p}_{\sigma(t)}, \mathbf{x}(t)) = \mathbf{p}_{\sigma(t)}^T \mathbf{x}(t) - 1 = 0 \quad (1.6)$$

Thus, in this context algorithm (1.5) can be directly used to robustly segment the video sequence. It is also worth stressing that as a by-product this method also performs *key frame extraction* by selecting $\mathcal{I}(t)$ corresponding to the minimum $\|\eta(t)\|$

⁶ Note that this always can be assumed without loss of generality due to the presence of noise in the data.

value in a segment (e.g., the frame with the smallest fitting error) as a good representative of the entire segment.

The content of a video sequence usually changes in a variety ways: For instance: the camera can switch between different scenes (e.g., shots); the activity within the scene can change over time; objects or people can enter or exit the scene, etc. There is a hierarchy in the level of segmentation one would require. The noise level ϵ can be used as a tuning knob in this sense.

Figure 1.8 shows the results of applying this approach to a video sequence, `drama.avi`, available from <http://www.open-video.org>. The original mpeg files were decompressed, converted to grayscale and title frames were removed. Each sequence shows a different characteristic on the transition from one shot to the other. The camera is mostly non-stationary, either shaking or moving. For comparison, results using GPCA, a histogram based method and an MPEG method for segmenting the sequences with optimal parameters (found by trial and error) are also shown. Table 1.1 shows the Rand indices [28] corresponding to the clustering results obtained for this sequence and three others from the same database (`roadtrip.avi`, `mountain.avi`, and `family.avi`) using the different methods, providing a quantitative criteria for comparison. Since the Rand index does not handle dual memberships, the frames corresponding to transitions were neglected while calculating the indices. These results show that indeed the sparsity method does well, with the worst relative performance being against MPEG and B2B in the sequence Roadtrip. This is mostly due to the fact that the parameters in both of these methods were adjusted by a lengthy trial and error process to yield optimal performance in this sequence. Indeed, in the case of MPEG based segmentation, the two parameters governing cut detection were adjusted to give optimal performance in the Roadtrip sequence, while the five gradual transition parameters were optimized for the Mountain sequence.

	Roadtrip	Mountain	Drama	Family
Sparsification	0.9373	0.9629	0.9802	0.9638
MPEG	1	0.9816	0.9133	0.9480
GPCA	0.6965	0.9263	0.7968	0.8220
Histogram	0.9615	0.5690	0.8809	0.9078

Table 1.1. Rand indices

1.5.2 Example 2: Segmentation of Dynamic Textures.

Modeling, recognition, synthesis and segmentation of dynamic textures have drawn a significant attention in recent years [29, 30, 31, 32]). In the case of segmentation tasks, the most commonly used models are mixture models, which are consistent with the hybrid model framework.

In the sparsification framework described earlier in this section, the problem of temporal segmentation of dynamic textures reduces to the same mathematical problem as the video segmentation problem, with the difference that now the underlying hybrid model should take the dynamics into account. First, dimensionality reduction is performed via PCA ($\mathcal{I}(t) \mapsto \mathbf{y}(t) \in \mathbb{R}^d$) and then the reduced-order data is assumed to satisfy a simple causal autoregressive model similar to the one in [31]. Specifically, in this case the hybrid model is given by:

$$\mathcal{H}_2 : f\left(\mathbf{p}_{\sigma(t)}, \{\mathbf{y}(k)\}_{k=t-n}^t\right) = \mathbf{p}_{\sigma(t)}^T \begin{bmatrix} \mathbf{y}(t-n) \\ \vdots \\ \mathbf{y}(t) \end{bmatrix} - 1 = 0 \quad (1.7)$$

where n is the regressor order. This model, which can be considered as a step driven autoregressive model, was found to be effective experimentally⁷. The power of this approach is illustrated in Figures 1.9 and 1.10 where two very challenging sequences were segmented. The first sequence consists of a patch of dynamic texture (smoke) appended in time to another patch from the same texture but transposed. Thus, the two subsequences have the same photometric properties but differ in the main motion direction. The second sequence was generated using another dynamic texture (river) by sliding a window both in space and time (by going forward in time in the first half and by going backward in the second), thus reversing the dynamics due to the river flow.

1.6 Constrained interpolation of high dimensional signals:

Consider first the simpler case of interpolating noiseless data, generated by a single LTI system, with McMillan degree bounded by some known n_o . Formally, given a partial sequence $\mathbf{d}_g = \{d_1, \dots, d_q, d_{q+r}, \dots, d_n\}$, the goal is to estimate the missing elements $\mathbf{d}_x = \{d_{q+1}, \dots, d_{q+r-1}\}$ that optimally fit the given data. Intuitively, the best fitting missing elements are those that require adding the least number of modes to the existing model in order to explain the new data. Using the fact that the order of the underlying model is given by the rank of the corresponding Hankel matrix (under the assumption that $n \gg n_o$), this problem can be recast into the following rank minimization form: $\mathbf{d}_{x_o} = \operatorname{argmin}_{\mathbf{d}_x} \operatorname{rank}(\mathbf{H})$ where \mathbf{H} is the Hankel matrix associated with the completed sequence $\mathbf{d} = \{d_i\}$. In this context, noise can be readily handled by simply adding a new variable v such that the measured data $y = d + v$, and a suitable noise description of the form $v \in \mathcal{N}$, a convex, compact set. Finally, since rank minimization is NP-hard [33], using the convex relaxation proposed in [34] leads to the following algorithm:

⁷ The independent term 1 here accounts for an exogenous driving signal. Normalizing the value of this signal to 1, essentially amounts to absorbing its dynamics into the coefficients \mathbf{p} of the model. This allows for detecting both changes in the coefficients of the model and in the statistics of the driving signal.




Fig. 1.9. Comparison of segmentation result for the Smoke sequence concatenated with its transposed dynamics.

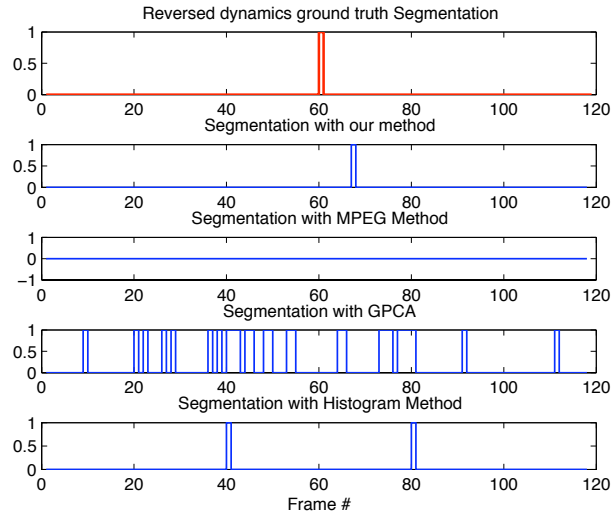


Fig. 1.10. Comparison of segmentation results for the River sequence concatenated with its reversed dynamics.

Algorithm 1: HANKEL RANK MINIMIZATION BASED INTERPOLATION/PREDICTION

Input at time k : N_h : Horizon length; $\mathcal{I}_a \subseteq [k - N_h, k]$: set of indices of available measurements (with $\text{card}(\mathcal{I}_a) \geq n$); $\mathcal{I}_e \subseteq [k - N_h, k + 1]$: set of indices of data to be estimated; with $\mathcal{I}_a \cup \mathcal{I}_e = \mathcal{I}$; available data y_ℓ , $\ell \in \mathcal{I}_a$; set membership description of the measurement noise $v \in \mathcal{N}$.

Output: Estimates $\hat{\zeta}_\ell$ of ζ_ℓ , $\forall \ell \in \mathcal{I}_e \cup \mathcal{I}_a$

1. Let ζ^* denote the following sequence:

$$\zeta_i^* = \begin{cases} y_i - v_i & \text{if } i \in \mathcal{I}_a \\ x_i & \text{if } i \in \mathcal{I}_e \end{cases} \quad \text{where } v, x \text{ are free}$$

variables, and form the matrix $H_\zeta \doteq \begin{bmatrix} \zeta_{i_1}^* & \zeta_{i_2}^* & \cdots & \zeta_{i_{n+n_u}+1}^* \\ \zeta_{i_2}^* & \zeta_{i_3}^* & \cdots & \zeta_{i_{n+n_u}+2}^* \\ \vdots & \vdots & \ddots & \vdots \\ \zeta_{i_{n+1}}^* & \zeta_{i_{n+2}}^* & \cdots & \zeta_{i_{2n+n_u}+1}^* \end{bmatrix}$

2. (Approximately) minimize $\text{rank}[H(x, v)]$ by solving the following convex problem in x, v, R, S :

$$\begin{aligned} & \text{minimize } \text{Tr}(R) + \text{Tr}(S) \\ & \text{subject to } \begin{bmatrix} R & H(x) \\ H(x)^T & S \end{bmatrix} \geq 0, \{v_\ell\} \in \mathcal{N}. \end{aligned}$$

3. Estimate/predict the output ζ_ℓ from the noisy measurements y_ℓ by:

$$\hat{\zeta}_i = \begin{cases} y_i - v_i & \text{if } i \in \mathcal{I}_a \text{ (estimation)} \\ x_i & \text{if } i \in \mathcal{I}_e \text{ (interpolation/prediction)} \end{cases}$$



Fig. 1.11. Data interpolation/association across occlusion. (Note that targets change relative positions while occluded.)

Examples of application of this idea are shown in Figure 1.11, where it was used to establish target identity across occlusion, and in Figure 1.12 where nonlinear embeddings were used first to map the data to a low order manifold where the rank-minimization based interpolation was performed, followed by a remapping of the data to pixel space. Finally, Fig 1.13 shows how a combination of dynamic interpolation and Hankel-rank based segmentation is able to detect occluded events.

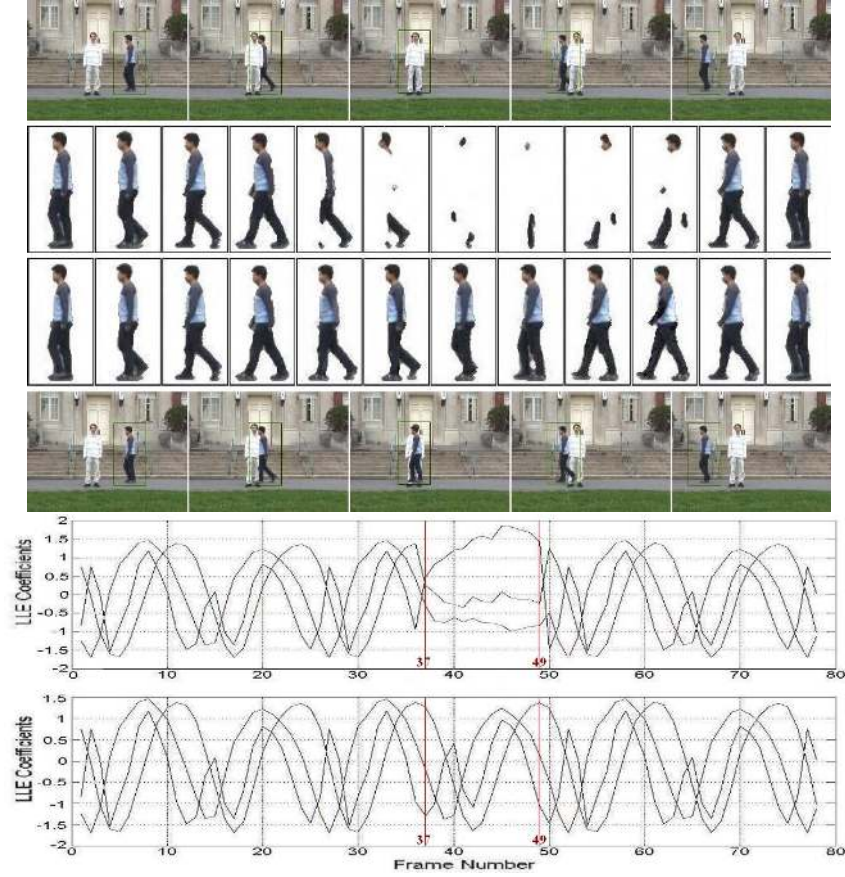


Fig. 1.12. Missing data (second and fifth rows) interpolated by rank minimization on 3D manifolds (third and sixth rows).

It is worth mentioning that the ideas discussed in this section are directly applicable to hybrid models of the form (1.2). In this case, minimizing the rank is roughly equivalent to interpolating the data so that the resulting underlying model exhibits the minimum number of jumps.

1.7 Hypothesis testing and data sharing.

A salient feature of the dynamics-based information extraction framework is its ability to furnish application-relevant *worst-case bounds* on distances between data and model predictions and, significantly, between non-overlapping data streams, in terms of their respective models. These bounds lead to computationally efficient hypothesis testing techniques. Consider a data stream $\{y_k\}_{k=0}^{N-1}$ generated by an underlying




Fig. 1.13. Occluded event detection. Top: Dynamic data interpolation. Bottom: Hankel rank plot showing events.

model of the form:

$$y_{k+1} = \mathcal{F}[y_k, \mathbf{e}_k], \mathbf{y}_k \doteq [y_k, \dots, y_{k-n}], \mathbf{e}_k \doteq [e_k, \dots, e_{k-n}] \quad (1.8)$$

where e is a stochastic input, and \mathcal{F} is the (unknown) evolution operator. Collecting all available *a priori* information about \mathcal{F} and e (e.g., known dynamic modes and noise statistics) into sets \mathcal{S} and \mathcal{N} reduces the problem of finding \mathcal{F} to a finite dimensional optimization via an extended Caratheodory-Fejer interpolation framework [35]. This method is *interpolatory*, hence it generates a model

$$\mathcal{F}_{id} \in \mathcal{T}(\mathbf{y}) \doteq \{\mathcal{F} \in \mathcal{S}: y_{k+1} = \mathcal{F}[\mathbf{y}, \mathbf{e}], \mathbf{e} \in \mathcal{N}\}$$

the set of all models consistent with both the *a priori* information and the experimental data. Since the actual (unknown) model that generated the data must also belong to $\mathcal{T}(\mathbf{y})$, a bound on the worst case prediction error of the identified model \mathcal{F}_{id} is given by:

$$\|\hat{\mathbf{y}} - \mathbf{y}\|_* \leq \sup_{\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{T}(\mathbf{y})} \|\mathcal{F}_1[\mathbf{y}, \mathbf{e}] - \mathcal{F}_2[\mathbf{y}, \mathbf{e}]\|_* = \mathcal{D}[\mathcal{T}(\mathbf{y})] \quad (1.9)$$

where $\|\cdot\|_*$ is a suitable norm and $\mathcal{D}(\cdot)$ denotes the diameter of the set $\mathcal{T}(\mathbf{y})$. When the sets \mathcal{S} and \mathcal{N} are convex, computing this bound reduces to a convex optimization problem [17, Lemma 10.3]. Note that these bounds are computed only once and remain valid as long as the underlying dynamics do not change.

Figure 1.14 compares the actual and upper bound of the error in a human tracking application. In this experiment the measured position in frame 12 was propagated



Frame	13	14	15	16	17	18	19	20
Actual error	8.87	6.14	10.04	13.03	10.31	15.72	19.50	26.04
Worst case bound	13.00	15	17	19	21	23	25	27

Fig. 1.14. Top: Prediction (black cross) versus Ground Truth (white cross). Bottom: Id error.

forward using the identified dynamics and the bounds computed by solving a single Linear Programming problem. If other targets with similar dynamics or photometric properties are present, trackers can safely discard candidates falling outside these bounds.

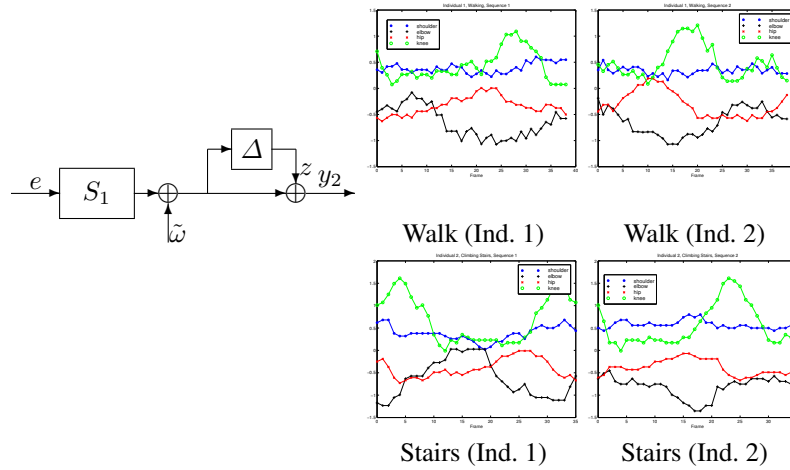


Fig. 1.15. Left: Distance between data streams as a model (in)validation problem. Right: sample joint traces for different activities.

In addition to the **low cost data gating** illustrated above, the worst case bounds provided by $\mathcal{D}[\mathcal{T}(\mathbf{y})]$ can be used to robustly assess the distance between non-overlapping data streams. The idea is to measure this quantity in terms of the distance between the corresponding (model) consistency sets. Intuitively, two partial data streams are considered to be manifestations of the same underlying process if they can be generated by the same dynamic model. The introduction of the consistency set in this context allows for taking into consideration data-quality issues (relatively few observations, corrupted by noise) and *a priori* information. Computing the exact distance between consistency sets is costly, but it can be relaxed to the model (in)validation form shown in Figure 1.15 (a). Given data streams y_1, y_2 , the idea is to identify a nominal model S_1 associated with y_1 and a deformation operator Δ so that the pair (S_1, Δ) generates y_2 . As shown in [36], computing the minimum norm $\gamma_{min} \doteq \min \|\Delta\|_\infty$ over the set of all operators with this property reduces to a convex Linear Matrix Inequality optimization problem. Thus, the value γ_{min} provides a computationally tractable upper bound on the distance between consistency sets.

This idea is illustrated next, using as an example the problem of gait classification. The experimental data listed in Table 1.2 and plotted in Figure 1.15(b), consists of 30 vector sequences, taken from 5 different persons, named A, B, C, D and E. Each sequence contains measurements of the angles of the shoulder, elbow, hip and knee joints of a person walking, running or walking up a staircase. For illustrative sake, these sequences are numbered from 1 to 30 so that the first 10 correspond to walking, the second set of 10 to running and the third set of 10 to walking up a staircase.

Person	Walking	Running	Staircase
A	1, 2	16 to 18	25 to 27
B	3 to 8	11 to 15	21 to 24
C	9, 10	none	28 to 30
D	none	19	none
E	none	20	none

Table 1.2. Experimental Data

Table 1.3 shows the distance from each data set to the dynamic model representing each activity. For each sequence, these nominal models were obtained by first finding a model associated with each of the remaining sequences and then selecting as representative of each class the model closest to its center (e.g., the one solving $\min_i \max_j \|S_i - S_j\|_\infty$). Note that nearest neighbor classification using this metric can successfully recognize 25 out of the 27 sequences under consideration; it only confuses 2 sequences, (y_{26} and y_{29} , belonging to persons A and C walking up a staircase) as walking sequences. The failure is due to the fact that in these instances the experimental data record is too short to disambiguate between activities.

Sequence	S_{walk}	S_{run}	S_{stair}
y_1	0.1743 [†]	0.6758	0.5973
y_2	0.2333 [†]	0.5818	0.2427
y_3	0.0305 [†]	0.6843	0.6866
y_4	0.0410 [†]	0.6217	0.5305
y_5	0.0819 [†]	0.6915	0.6069
y_6	0.0001 [†]	0.6879	0.7688
y_7	0.0900 [†]	0.6892	0.9188
y_8	0.2068 [†]	0.6037	0.7883
y_9	0.0001 [†]	0.6926	0.6028
y_{11}	0.9265	0.3415 [†]	1.0000
y_{12}	0.9676	0.2452 [†]	0.9325
y_{13}	1	0.0002 [†]	0.9323
y_{14}	1	0.0002 [†]	0.9903
y_{15}	1	0.0002 [†]	0.8999
y_{16}	1	0.0005 [†]	0.5707
y_{17}	0.9220	0.0532 [†]	0.5437
y_{18}	1	0.0004 [†]	0.6961
y_{19}	1	0.3545 [†]	0.8374
y_{21}	0.9631	0.5002	0.3174 [†]
y_{22}	0.7952	0.4122	0.0577 [†]
y_{23}	0.7215	0.4089	0.0936 [†]
y_{24}	0.8499	0.4456	0.0805 [†]
y_{25}	0.7252	0.5928	0.3962 [†]
y_{26}	0.6828 [†]	0.7127	0.8827
y_{27}	0.5553	0.5818	0.4682 [†]
y_{28}	0.2650	0.6801	0.1699 [†]
y_{29}	0.0391 [†]	0.6102	0.1470

Table 1.3. Right: distance between data and models (in the \mathcal{H}_∞ norm). In each row [†] denotes the model whose output best matches the given sequence

In the case of **distributed data sources** the high costs (both in bandwidth and computational cost) entailed in sharing information can be avoided by (i) associating to each source a set of intrinsic coordinates on a low dimensional manifold, and (ii) using robust identification techniques [36] to identify dynamic models for the mappings between the projections of the different local data sources (e.g., sensors) onto the respective manifolds (see Figure 1.16). Then, only these low dimensional projections need to be exchanged between nodes, and each node can reconstruct the data observed by other nodes, simply by applying the interconnecting models. Figure 1.16 shows an application of these ideas to the problem of tracking and disambiguating two virtually identical targets.

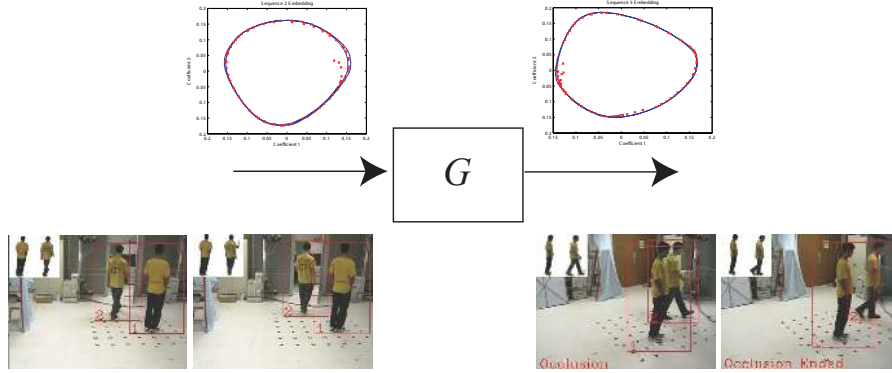


Fig. 1.16. Mapping manifolds between 2 sensors used to recreate an occluded person.

1.8 Conclusions

Arguably, one of the hardest challenges entailed in exploiting actionable information sparsely encoded in high volume data streams is the development of scalable, tractable methods capable of dealing with the overwhelming volume of data [37]. Recent work (manifold embedding [2], compressive sensing, [38, 39, 40]) have led to substantial progress in addressing this issue. However, these methods stop short of fully exploiting the gap between data dimensionality and the rank of the *dynamical* system underlying the data record.

As shown in this chapter, the use *dynamic models as an information encoding paradigm*, can lead to both, substantial dimensionality reduction and computationally attractive algorithms for data extraction/interpretation. Dynamic structures can be tractably discovered from the data in a way which leverages their inherently lower dimensionality. One key feature is the ability of dynamic representations to produce quantifiable measures of uncertainty as *provable error bounds* on the validity of the data interpretation suggested by the model. Another is their relative computational simplicity: in many cases postulating the existence of such a model and associated invariants (e.g., model order) is enough to develop computationally attractive, robust solutions to problems such as segmentation, interpolation and event detection. We believe that these techniques hold the key to render practical several applications, ranging from self-aware environments to automatic discovery of co-regulated genes, that are currently at the proof-of-concept stage, and where the major roadblock is precisely the lack of techniques to robustly handle the extremely high volume of (often relatively low quality) data.

1.9 Acknowledgments

The authors are indebted to Professor Uri Alon, Weizmann Institute, and Dr. Alon Zaslaver, Caltech, for providing the diauxic shift experimental data used in Figures 1.1(c), 1.5(c) and 1.7. The research that generated this data was supported by the Kahn Family Foundation at the Weizmann Institute of Science. The authors are also grateful to Professor Stefano Soatto, UCLA, for providing the data used in the gait classification example shown in Fig. 1.15.

References

1. Zaslaver, A., Bren, A., Ronen, M., Itzkovitz, S., Kikoin, I., Shavit, S., Liebermeister, W., Surette, M.G., Alon, U.: A comprehensive library of fluorescent transcriptional reporters for *Escherichia coli*. *Nature Methods* **3** (2006) 623–628.
2. Lee, J.A., Verleysen, M.: *Nonlinear Dimensionality Reduction*. Springer (2007)
3. Costeira, J., Kanade, T.: A multibody factorization method for independently moving objects. *International Journal of Computer Vision* **29** (1998) 159–179
4. Vidal, R., Ma, Y., Soatto, S., Sastry, S.: Two-view multibody structure from motion. *International Journal of Computer Vision* **68** (2006) 7–25
5. Zelnik-Manor, L., Irani, M.: Degeneracies, dependencies and their implications in multibody and multi-sequence factorization. *Proc. of the 2003 IEEE Conf. on Computer Vision and Pattern Recognition*. (2003) 287–293
6. Ma, Y., Vidal, R.: A closed form solution to the identification of hybrid arx models via the identification of algebraic varieties. *Hybrid Systems Computation and Control*. (2005) 449–465
7. Bemporad, A., Garulli, A., Paoletti, S., Vicino, A.: A bounded-error approach to piecewise affine system identification. *IEEE Trans. on Automatic Control* **50** (2005) 1567–1580
8. Saul, L.K., Roweis, S.T.: Think globally, fit locally: unsupervised learning of low dimensional manifolds. *Journal on Machine Learning Research* **4** (2003) 119–155
9. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290** (2000) 2319–2323
10. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* **15** (2003) 1373–1396
11. Donoho, D.L., Grimes, C.E.: Hessian eigenmaps: locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Arts and Sciences*. Volume 100. (2003) 5591–5596
12. Weinberger, K.Q., Sha, F., Saul, L.K.: Learning a kernel matrix for nonlinear dimensionality reduction. *Proc. of the 2004 International Conference on Machine Learning*, ACM Press (2004)
13. Weinberger, K.Q., Saul, L.K.: Unsupervised learning of image manifolds by semidefinite programming. *Proc. of the 2004 IEEE Conf. on Computer Vision and Pattern Recognition*. (2004) 988–995
14. Ozay, N., Sznajder, M., Lagoa, C., Camps, O.: A sparsification approach to set membership identification of a class of affine hybrid system. *Proc. 47th IEEE Conf. Dec. Control (CDC)*. (2008) 123–130

15. Sznaier, M., Lagoa, C., Mazzaro, M.C.: An algorithm for sampling balls in \mathcal{H}_∞ with applications to risk-adjusted performance analysis and model (in)validation. *IEEE Trans. Autom. Control* **50** (2005) 410–416
16. Sznaier, M., Lagoa, C., Ma, W.: Risk adjusted identification of a class of nonlinear systems. *Proc. 46 IEEE Conf. Dec. Control.* (2007) 5117–5122
17. Sánchez Peña, R., Sznaier, M.: *Robust Systems Theory and Applications*. Wiley & Sons, Inc. (1998)
18. Tomasi, C., Kanade, T.: Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision* **9** (1992) 137–154
19. Xiao, J., Chai, J., Kanade, T.: A closed-form solution to non-rigid shape and motion recovery. *Proc. of the 8th European Conference on Computer Vision (ECCV 2004)*. (2004)
20. Vidal, R., Hartley, R.: Motion segmentation with missing data using powerfactorization and gpca. *Proc. of the 2004 IEEE Conf. on Computer Vision and Pattern Recognition. Volume 2.* (2004) 310–316
21. Fazel, M., Hindi, H., Boyd, S.P.: A rank minimization heuristic with application to minimum order system approximation. *Proceedings of the 2001 American Control Conf.* (2001) Volume 6., AACC (2001) 4734–4739
22. Lobo, M., Fazel, M., Boyd, S.: Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research* **152** (2007) 376–394
23. Gargi, U., Kasturi, R., Strayer, S.H.: Performance characterization of video-shot-change detection methods. *IEEE Transactions on Circuits and Systems for Video Technology* **10** (2000) 1–13
24. Yuan, J., Wang, H., Xiao, L., Zheng, W., Li, J., Lin, F., Zhang, B.: A formal study of shot boundary detection. *IEEE Transactions on Circuits and Systems for Video Technology* **17** (2007) 168–186
25. Osian, M., Van Gool, L.: Video shot characterization. *Machine Vision and Applications* **15** (2004) 172–177
26. Petrovic, N., Ivanovic, A., Jovic, N.: Recursive estimation of generative models of video. *Proc. of the 2006 IEEE Conf. on Computer Vision and Pattern Recognition* (2006) I: 79–86
27. Lu, L., Vidal, R.: Combined central and subspace clustering for computer vision applications. *Proc of the 2006 International Conference on Machine Learning.* (2006) 593–600
28. Rand, W.: Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association* **66** (1971) 846–850
29. Doretto, G., Chiuso, A., Wu, Y., Soatto, S.: Dynamic textures. *International J. Computer Vision* **51** (2003) 91–109
30. Chan, A.B., Vasconcelos, N.: Mixtures of dynamic textures. *Proc. of the 2005 International Conference on Computer Vision (ICCV).* Volume 1. (2005) 641–647
31. Cooper, L., Liu, J., Huang, K.: Spatial segmentation of temporal texture using mixture linear models. *Proceedings of the Dynamical Vision Workshop in the International Conference of Computer Vision* (2005) 142–150
32. Ghoreyshi, A., Vidal, R.: Segmenting dynamic textures with using descriptors, arx models and level sets. *Proceedings of the Dynamical Vision Workshop in the European Conference of Computer Vision* (2006) 127–141
33. Vandenberghe, L., Boyd, S.P.: Semidefinite programming. *SIAM Review* **38**(1) (1996) 49–95
34. Fazel, M., Hindi, H., Boyd, S.P.: Log-det heuristic for matrix rank minimization with applications to hankel and euclidean distance matrices. *Proceedings of the 2003 American Control Conf.* (2003) Volume 3., AACC (2003) 2156–2162

35. Sznaier, M., Mazzaro, C., Camps, O.I.: Open-loop worst-case identification of nonschur plants. *Automatica* **39** (2003) 1019–1025
36. Chen, J., Gu, G.: *Control Oriented System Identification, An \mathcal{H}_∞ Approach*. John Wiley, New York (2000)
37. Donoho, D.: High-dimensional data analysis: The curses and blessings of dimensionality. *AMS Conf. on Math Challenges of the 21st Century*. (2000) Plenary presentation.
38. Candes, E., Romberg, J., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics* **59** (2006) 1207–1223
39. Donoho, D.: Compressed sensing. *IEEE Trans. on Information Theory* **52** (2006) 1289–1306
40. Carin, L., Liu, D., Guo, B.: In situ compressive sensing. *Inverse Problems* **24** (2008) 015023–1–015023–23