Proceedings of the 1999 IEEE
International Symposium on Computer Aided Control System Design
Kohala Coast-Island of Hawai'i, Hawai'i, USA • August 22-27, 1999

A CANONICAL REPRESENTATION FOR UNFALSIFIED CONTROL IN TRUNCATED SPACES

Paul B. Brugarolas[1]
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr, M/S 198-326
Pasadena, CA 91109

Michael G. Safonov[2]
Dept. Electrical Engineering - Systems
University of Southern California
3740 McClintock Ave, EEB-310
Los Angeles, CA 90089-2563

Abstract

In this paper we present a canonical representation for truncated space unfalsified control. We show how it specializes to controller unfalsification and adaptive control problems. Moreover, we propose it to solve system identification problems. In addition, we analyze in detail the core issues of the truncated space unfalsified control theory, give a computational algorithm and practical considerations for its implementation.

1. INTRODUCTION

In the past few years there is been some interest in using falsification ideas for control design and identification [5,7,11]. Unfalsified control emerged as one of these ideas [10]. The problem that unfalsified control solves is the problem of identifying controllers from experimental measurements. Unfalsified control is essentially a variant on the candidate elimination algorithm of Mitchell [6] applied to the control problem. It works by identifying hypothetical controllers that are consistent with past measurement data. In the minds of the researchers who developed this approach was the goal of keeping the formulation as simple and faithful to the scientific method as possible. In doing so they achieved a parsimonious mathematical representation of some essential issues in feedback and learning from data.

In this paper we give a canonical representation for truncated space unfalsified control and show how it is suitable for unfalsified approaches to controller validation and adaptive control problems, which have been already studied in the literature [1-3,4,10,11,13]. In addition, we propose in the same framework a new application of these unfalsified control concepts, application to system identification problems. Truncated space unfalsified control results from working in truncated signals spaces obtained from the application of an observation operator. This subclass of unfalsified control problems is of interest because allows for simple falsification conditions which in turn require a reduced number of computations. So far, all unfalsified control applications up to day fall into this subclass. Moreover, in our development we study in detail this truncated space unfalsified control and present two special cases that allow for direct falsification. In addition, we give a computational algorithm and some practical considerations. We presented some of the ideas here in [8].

2. CANONICAL REPRESENTATION OF UNFALSIFIED CONTROL

In this section we briefly review unfalsified control. The unfalsified control concept is a precise formulation of the controller validation problem in a hypothesis-testing framework [10]. That is given the control *goal* (performance specification) and a set of *hypotheses* (candidate controllers), then it evaluates them against experimental *data*. An important aspect of this approach is that it is data driven. Experimental data is seen as a particular realization of the plant, which often includes actuators and sensors, and their disturbances and noises. For this reason we define an extended plant to be the plant with actuators and sensors. Therefore the experimental data is a realization of the extended plant. Another important aspect is that *a priori* information about the plant is used to define the hypothesis set. Models and any other prior knowledge about the plant are used to design the candidate controllers (hypotheses).

Figure 1 depicts a general representation for an unfalsified control system. It is composed of elements that define relations between signals. The unfalsified control system has an internal structure with three blocks: the controller architecture, the learning processor and the controller. The controller architecture acts as the interface for the extended plant and the controller, and provides the performance signals needed by the learning processor, which in turn evaluates the candidate controllers and selects the best one to be used by the controller. The engine of the unfalsified

[1] Research supported in part by NASA.
[2] Research supported in part by AFOSR Grant F49620-98-1-0026.

control learning process is the evaluation of hypothesis against experimental data. A precise definition follows.

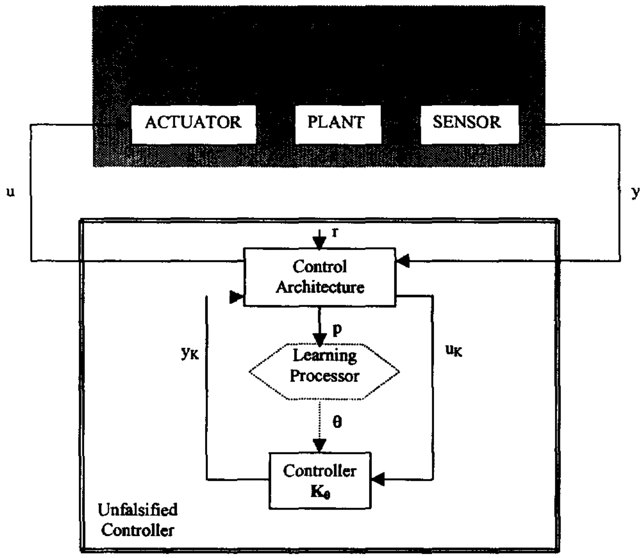

Figure 1: Unfalsified control system

Definition 1 (Unfalsified Controller) (cf. [10]): Given measurement information (*data*), then a controller (*hypothesis*) is said to be falsified if the combined information (*hypothesis* and *data*) is sufficient to deduce that the performance specification (*goals*) would be violated if that controller were in the feedback loop. Otherwise, it is said to be unfalsified.

3. TRUNCATED SPACE PROBLEM FORMULATION

The general unfalsified control problem considered in [2,10] consists of evaluating if the candidate controllers are falsified by the experimental data. In the present paper, we consider only goals that are directly expressible in terms of currently available data. Now we present the elements needed to define and solve the problem of evaluating if a control law is falsified.

We start by defining the signals and then the blocks. We distinguish between two groups of signals, manifest and latent, as in the behavioral approach of Willems [12]. Manifest signals ($Z_{hypotheses} = \cup Z_{hypothesis} \subset P_\tau Z$) as the name expresses are the signals that manifest the observed behavior of the extended plant. These signals are control (u) and measurement (y) signals. Hence the manifest subspace is the input-output space, i.e. $U \times Y$. Latent signals (z_{latent}) are signals internal to the unfalsified controller used to define and evaluate performance and to produce the control signal. These signals are controller (u_K, y_K), command (r) and performance (p) signals, and controller gains (θ). Consider the vector space formed from taking all the signals in one vector $z = (z_{manifest}, z_{latent}) = (u, y, r, p, u_K, y_K, \theta)$, that is, $Z = Z_{manifest} \times Z_{latent}$. In addition, define the observation operator, P_τ, which maps input-output signals to measurement signals. Examples of this operator are the time truncation operator or the sampling operator. Now let's us define a vector space of truncated signals obtained from applying the observation operator to the vector space Z.

Definition 2 (Truncated signal space): Given a space Z, and an observation operator P_τ. The observations truncated signal space is defined as $Z^\tau = P_\tau Z$.

From now on, we are going to work in this truncated signal space for practical reasons. In fact, we will define sets in this truncated signal space; to signify this we will use the superscript τ. Specifically, we consider the following sets in the truncated signal space defined above:

- *Goal set,* $Z^\tau{}_{goal} \subset P_\tau Z$, denotes the performance specification. This could be described in terms of an observation-dependent error functional, $J(P_\tau z)$, which evaluates the error between the system response and the desired one. Without loss of generalty we normalized it to [0,1]. Hence,
$$Z^\tau{}_{goals} = \{P_\tau z | 0 \le J(P_\tau z) \le 1\} \subset P_\tau Z .$$
Observe that since the functional evaluates the error, the value of the cost is a measure of the performance such that smaller cost indicates better performance. Later on in some applications we will use this cost to define a partial ordering of hypotheses.
- *Controller hypothesis,* $Z^\tau{}_{hypothesis}(\theta) \subset P_\tau Z$, denotes the *graph* of a θ–dependent dynamical control law $K_\theta(P_\tau z) = 0$ ($\theta \in \Theta$), viz.,
$$Z^\tau{}_{hypothesis}(\theta) = \{P_\tau z | K_\theta(P_\tau z) = 0\} \subset P_\tau Z .$$
- *Hypotheses set,*
$Z^\tau{}_{hypotheses} : \Theta \to 2^{P_\tau Z}$, $Z^\tau{}_{hypothesis}(\theta) \in Z^\tau{}_{hypotheses}$
denotes the set of candidate controllers parameterized by a vector θ.
- *Data set,* $Z^\tau{}_{data} \subset P_\tau Z$, is defined as
$$Z^\tau{}_{data} = \{P_\tau z | P_\tau (z_{manifest}) = (u, y)_{measurements}\} \subset P_\tau Z$$
where $(u, y)_{measurements}$ represents observed past plant input-ouput data available at time τ.
- *The unfalsification goal:*
$$\left(Z^\tau{}_{data} \cap Z^\tau{}_{hypothesis}(\theta)\right) \subset Z^\tau{}_{goal}$$

These truncated sets will be used in the following section to test falsification for purposes of controller validation, adaptive control and system identification.

Problem 1 (Truncated Space Unfalsified Control Problem): Given a performance specification (*goal set,* $Z^{\tau}{}_{goals} \neq \varnothing$), a set of candidate controllers (*hypotheses,* $Z^{\tau}{}_{hypotheses} = \cup Z^{\tau}{}_{hypothesis}(\theta) \neq \varnothing$), and experimental *data* ($Z^{\tau}{}_{data} \neq \varnothing$) then determine the subset of candidate controllers (*hypothesis*) which is not falsified.

This formulation brings important advantages over past formulations of unfalsified control [10]. One advantage is that every element of the truncated space data set $Z^{\tau}{}_{data}$ is not only consistent with the *observed* data $(u,y)_{measurements} \in P_{Z_{manifest}} Z^{\tau}{}_{data}$ (where $P_{Z_{manifest}}$ denotes the projection $z_{manifest} = P_{Z_{manifest}} z$), but it is also the unique element of the set $P_{Z_{manifest}} Z^{\tau}{}_{data}$. That is, $P_{Z_{manifest}} Z^{\tau}{}_{data} = (u,y)_{measurements}$. This means that in testing unfalsification there is no need to analyze multiplicities of unseen future or intersample behaviors for the manifest signals (u,y). Another advantage is that we have more flexibility in the design since we are not just working in $R \times Y \times U$ but in Z, which may in general be a bigger signal space. This flexibility is very useful in the definition of the performance specification. For example it allows for specifications involving signals derived from a hypothesis-dependent reference model, which is not possible when performance goals must depend only on (r,y,u).

4. MAIN RESULTS

Now we are going to build the background for a theorem that will give a solution to this truncated space unfalsified control problem. We will start by defining the concept of data-hypothesis consistency, which concerns whether a hypothesis (controller) connected to the plant in closed-loop could have produced the measurement data.

Definition 3 (Data-hypothesis consistency): Given a truncated space unfalsified control problem, we say that a *hypothesis* is consistent with the *data* if $P_{Z_{manifest}} Z^{\tau}{}_{data} \subset P_{Z_{manifest}} Z^{\tau}{}_{hypothesis}(\theta)$.

Remark: Note that if a particular controller is not consistent with data, then irrespective of the goal $Z^{\tau}{}_{goal}$ the data cannot falsify this controller.

Now we give a condition for falsification of a candidate controller.

Theorem 1 (Truncated space unfalsified control): For a given truncated space unfalsified control problem, then a candidate control law (*hypothesis,* $Z^{\tau}{}_{hypothesis}(\theta)$) consistent with the experimental *data* is unfalsified by data $P_{Z_{manifest}} Z^{\tau}{}_{data}$ if and only if $Z^{\tau}{}_{data} \cap Z^{\tau}{}_{hypothesis}(\theta) \cap \overline{Z^{\tau}{}_{goal}} = \varnothing$, where $\overline{Z^{\tau}{}_{goal}}$ denotes the complement of the set $Z^{\tau}{}_{goal}$. Otherwise, it is falsified.

Proof: The result follows directly from *Definition* 1 and the equivalence

$$\left\{ \left(Z^{\tau}{}_{data} \cap Z^{\tau}{}_{hypothesis}(\theta) \right) \subset Z^{\tau}{}_{goal} \right\} \Leftrightarrow \left\{ Z^{\tau}{}_{data} \cap Z^{\tau}{}_{hypothesis}(\theta) \cap \overline{Z^{\tau}{}_{goal}} = \varnothing \right\}.$$

Remark: Note that a controller that has not been falsified it has been robust against the uncertainty evidenced in the experimental data [9].

A solution to *Problem 1* is achieved by testing each candidate controller for unfalsification via *Theorem 1.*

In some cases direct falsification is possible, either without data or perhaps without a specifying a controller. These cases arise from incompatibilities between the goal and the data or the hypothesis. To focus attention on these cases we state the following trivial corollary to *Theorem 1.*

Corollary 1 (Special cases for direct falsification):

i. If a *hypothesis* is consistent with *data* but not consistent with the *goal* (i.e. $Z^{\tau}{}_{hypothesis}(\theta) \cap Z^{\tau}{}_{goals} = \varnothing$), then the *hypothesis* (control law) is falsified.

ii. If the *data* is not consistent with the *goal* (i.e. $Z^{\tau}{}_{data} \cap Z^{\tau}{}_{goals} = \varnothing$), then all *hypothesis* (control laws) that are consistent with this *data* are falsified.

Proof: Both results follow directly from *Theorem 1.*

Remarks:

- We start assuming that the goal, hypotheses and data sets were not empty. If any of these sets is empty the problem is trivial.
- *Theorem 1* above is equivalent to [10, Thm. 1], applied to the more specialized truncated space problem formulation of our. That is, *if* we specialize the performance specification of [10] to *Problem 1* we will get our *Theorem 1.*

5. ADAPTIVE TRUNCATED SPACE UNFALSIFIED CONTROL

In this section we are going to apply the above concepts to solve an adaptive control problem [10]. The problem that we would like to solve is the following:

Problem 2 (Adaptive Truncated Space Unfalsified Control Problem): Given a γ-dependent *goal set,* $Z^{\tau}{}_{goals}(\gamma) \neq \varnothing$, ($\gamma \in \mathrm{R}$), a set of candidate controller *hypotheses,*

$Z^{\tau}{}_{hypotheses} = \bigcup Z^{\tau}{}_{hypothesis}(\theta) \neq \varnothing$, and an evolving τ-dependent experimental *data set* ($Z^{\tau}{}_{data} \neq \varnothing$), then at each time τ find the least $\gamma = \gamma_{opt}$ for which the set of unfalsified controllers is non-empty and select a controller $K(\theta_{opt})$ from this set.

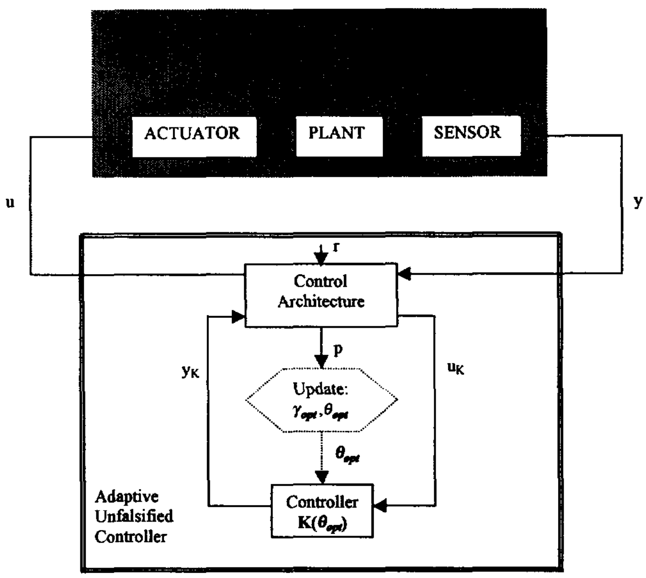

Figure 2 : Adaptive unfalsified control system.

In Figure 2 we specialize the canonical representation given earlier to this problem. In this case the learning processor does two tasks: update the unfalsified candidate controller subset and select from it the best controller and put it in the loop.

Observe that at each time we will have a new measurement so our observation operator will be different, it will have to include the new measurement. In consequence the basic sets will change since they depend on the observations operator. In addition, we have added the parameter γ in the goal set to define a partial ordering of the hypotheses.

$$Z^{\tau}{}_{goals}(\gamma) = \{P_{\tau}z \mid 0 \leq J(P_{\tau}z) \leq \gamma\}$$

The adaptive unfalsified control problem could be posed as an optimization problem as follows:

Theorem 2 (Adaptive Truncated Space Unfalsified Control): Given an adaptive truncated space unfalsified control problem. Then a solution to it is given by the following constraint optimization: At each time τ, find a controller $K_{\theta_{opt}}$ ($Z_{hypothesis}(\theta_{opt})$) that solves:

$$\gamma_{opt} := \left(\begin{array}{l} \underset{Z^{\tau}{}_{hypothesis}(\theta)}{\arg\min} \left(\begin{array}{l} \underset{P_{\tau}z}{\arg\min}\, J(P_{\tau}z) \\ s.t.\;\; P_{\tau}z \in \left(Z^{\tau}{}_{data} \cap Z^{\tau}{}_{hypothesis}(\theta)\right) \end{array} \right) \\ s.t. \;\; \begin{array}{l} Z^{\tau}{}_{data} \cap Z^{\tau}{}_{hypothesis}(\theta) \subseteq Z^{\tau}{}_{goal}(\gamma) \\ Z^{\tau}{}_{hypothesis}(\theta) \in Z^{\tau}{}_{hypotheses} \end{array} \end{array} \right)$$

Remark: Note that in the inner minimization the condition restricts the signals to the possible closed-loops with controller K ($Z_{hypothesis}$), and in the outer minimization the condition restricts the controllers to the unfalsified ones. These optimization constraints could be expressed in terms of inequalities.

Furthermore, it is interesting to mention here that this representation could also be used to represent other adaptive control algorithms like MRAC (Model Reference Adaptive Control) [3].

6. TRUNCATED SPACE UNFALSIFIED SYSTEM IDENTIFICATION

In this section we are going to apply the unfalsified control concepts to build a system identification algorithm. The problem that we would like to solve is the following:

Problem 3 (Closed-Loop Truncated Space Unfalsified System Identification): Given a γ-dependent *goal set,* $Z^{\tau}{}_{goals}(\gamma) \neq \varnothing$, ($\gamma \in R$), a set of candidate model *hypotheses,* $Z^{\tau}{}_{hypotheses} = \bigcup Z^{\tau}{}_{hypothesis}(\theta) \neq \varnothing$, and an evolving τ-dependent experimental *data set* ($Z^{\tau}{}_{data} \neq \varnothing$), then at each time τ find the least $\gamma = \gamma_{opt}$ for which the set of unfalsified models is non-empty and select a model from this set.

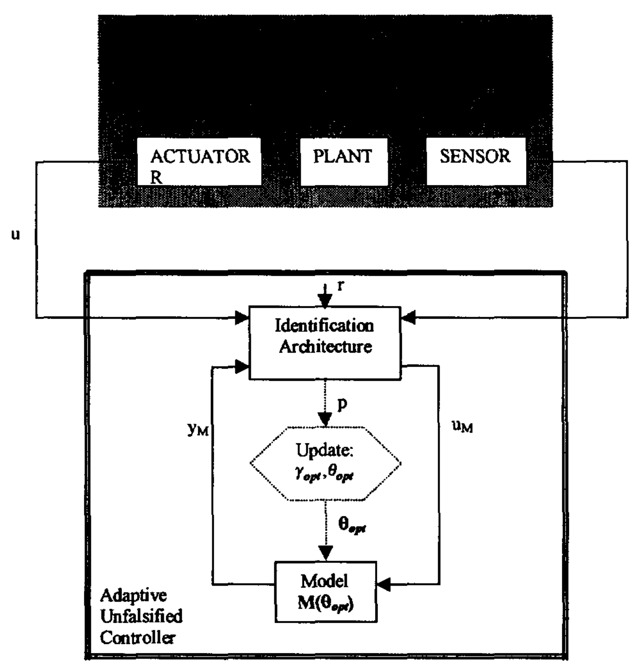

Figure 3: Unfalsified system identification.

In Figure 3 we specialize the canonical representation given earlier to this problem for a closed-loop implementation. Note an open loop implementation is also possible. In the open loop case we will be doing model unfalsification (validation). An application appeared in [13].

In the case presented in Figure 3, the closed-loop case, the learning processor does the job of solving the optimization, which outputs the model parameters for the best model. Here as in the unfalsified adaptive control case at each time τ we have a different observation operator P_τ, and we introduce the parameter γ to define a partial ordering in the goal set.

The closed-loop unfalsified system identification problem can be posed as an optimization problem as follows:

Theorem 3 (Closed-Loop Truncated Space Unfalsified System Identification): Given a closed-loop truncated space unfalsified system identification problem. Then a solution to it is given by the following constraint optimization: At each time τ, find a model $M_{\theta_{opt}}(Z_{hypothesis}(\theta_{opt}))$ that solves:

$$\gamma_{opt} := \begin{pmatrix} \underset{Z^\tau{}_{hypothesis}(\theta)}{\arg\min} \left(\begin{array}{l} \underset{P_\tau z}{\arg\min} J(P_\tau z) \\ s.t.\quad P_\tau z \in \left(Z^\tau{}_{data} \cap Z^\tau{}_{hypothesis}(\theta) \right) \end{array} \right) \\ s.t. \quad \begin{array}{l} Z^\tau{}_{data} \cap Z^\tau{}_{hypothesis}(\theta) \subseteq Z^\tau{}_{goal}(\gamma) \\ Z^\tau{}_{hypothesis}(\theta) \in Z^\tau{}_{hypotheses} \end{array} \end{pmatrix}$$

Remark: Note that this algorithm is analogous to the unfalsified adaptive control one.

7. PRACTICAL CONSIDERATIONS

The three cases studied above, truncated space unfalsified control, adaptive truncated space unfalsified control and truncated space unfalsified system identification, were specializations of the canonical representation. In this section we will give some practical consideration that will be valid for all the problems.

An important first remark is that the unfalsified control could be run just once, or it can be run iteratively on evolving past data as each new datum is acquired. If we run it iteratively on evolving past data then it can be used for real time controller adaptation. Adaptive updates of the current controller can be done either periodically or aperiodically, or by a mixture of both. If we run it periodically, we could run it at any rate. For quickly varying systems we may want to run it at the highest possible rate (sensor sampling rate) but for slowly varying systems we could run it at a slower rate. If we decide to run it aperiodically, it may be convenient to define a function to decide when to run it based on observed performance. A more sophisticated function could be used for the mixed approach.

In Figure 4, we introduce an algorithm for a generic unfalsified learning processor. Every time we call the learning processor, the data available defines the observations operator which in turn will help us define the truncated signal space (*Definition 1*) in which we will work. The block "Choose best" identifies the best hypothesis to be used in the adaptive and identification problems.

As it was explained in *Corollary 1*, there are two cases where direct falsification is possible. If the test associated with the evaluation of the conditions is simple in comparison with the data-hypothesis consistency and falsification tests, then it will be wise to preprocess the hypothesis set trough these tests before the learning processor stage.

Now we present some ideas on how to define these basic sets. These sets are defined in the truncated signal space obtained from the use of the observations operator.

- *Goal*: As expressed earlier without loss of generality the goal, or performance specification, could be defined using a cost function. This cost function should measure in some sense the error between the actual performance and desired one. The measure used could be of many different types. Alternatively, it may have the form of an inequality. For example, [10] used a weighted l_{2e} norm specification of the mixed sensitivity type, [1] used a model reference specification in terms of l_I norms, [4] used an analogous one with l_2 norms. For real-time application with a time varying system we suggest to introduce a forgetting factor to rest importance to the old measurements.

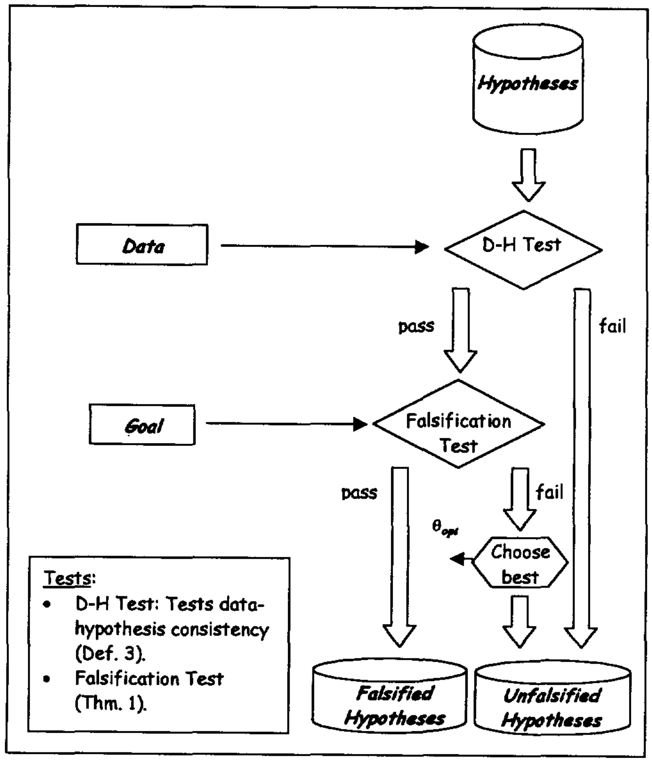

Figure 4: Unfalsified learning processor.

- *Hypotheses*. The controller (or model) hypotheses could be of many different types. However from the computational point of view, it is quite useful if the controllers (or models) are causally left inevitable since then we could use the fictitious reference signal approach, which was introduced in [10] and have been extensible used [1, 3, 9]. In addition, it will be also great advantage to have it parameterized in such a way that the fictitious reference signal could be represented as a product of two vectors, one that depends on the parameters (typically through nonlinear function) and another that contains the dynamics as in [1]. Examples of parameterizations used that meet the causally left invertible condition are the traditional parameterization used in MRAC [10], or parameterizations of the PID type as used in [1]. Other parameterizations that may not satisfy this condition are also possible. For example [4,13] use ARX parameterizations.

- *Data*: The experimental data available defines the observations operator. The data set is the projection of experimental data onto the latent variable subspace of the truncated signal space. Typically, experimental data will be discrete signals of finite length. However, [4] presented a study where signals may asymptotically approach infinite duration.

8. CONCLUSION

In this paper we have presented a tutorial overview of unfalsified control, focusing on the simplifications that result when performance goals are expressed directly in terms of truncated signal spaces. To work in this truncated signal space has the advantage of requiring a reduced number of computations. We have analyzed the core issues in the truncated space unfalsified control theory. We showed how it applies to controller validation and adaptive control, problems that have already been studied from the unfalsified control perspective. In addition, we described how truncated space system identification problems may be studied under the same framework and we showed that the results are analogous to the adaptive control case. We presented some special cases for direct falsification. Finally, linking unfalsified control theory to Mitchell's candidate elimination learning algorithm [6], we have given a computational algorithm and practical considerations for implementation in the context of adaptive feedback control of dynamical systems.

REFERENCES

[1] P. B. Brugarolas, V. Fromion, and M. G. Safonov "Robust switching missile autopilot," in Proc. American Control Conference, Philadelphia, PA, June 22-26, 1998.

[2] T. F. Brozenec, T. C. Tsao, and M. G. Safonov "Controller Invalidation" Proc. IFAC 13th Triennial World Congress pp. 303-308, San Francisco, CA, 1996.

[3] F. B. Cabral and M. G. Safonov "A falsification perspective on Model Reference Adaptive Control" in Proc. IEEE Conf. on Decision and Control, Kobe, Japan, December 11-13, 1996.

[4] R. L. Kosut. "Uncertainty Model Unfalsification for Robust Adaptive Control" Proc. IFAC Workshop on Adaptive Systems in Control & Signal Processing, University of Strathclyde, Glasgow, Scotland, UK August 26-28, 1998.

[5] R.L. Kosut, M. K. Lau, and S. P. Boyd, "Set-membership identification of systems with parametric and nonparametric uncertainty," IEEE Transactions on Automatic Control, 37(7): 929-941, July 1992.

[6] T. M. Mitchell, "Machine Learning". McGraw-Hill, 1997.

[7] K. Poolla, P. Khargonekar, A. Tikku, J. Krause and K. Nagpal, "A time-domain approach to model validation," IEEE Transactions on Automatic Control, 39(5): 951-959, May 1994.

[8] [M. G. Safonov, "Research summary: robust control, feedback and learning," Proc. AFOSR Workshop on Dynamics and Control, Pasadena, CA, May 27-29, 1998.

[9] M. G. Safonov and F. B. Cabral, "Robustness-oriented controller identification", Proc. IFAC Symp. on Robust Control Design ROCOND'97, Budapest, Hungary, June 25-27, 1997.

[10] M. G. Safonov and T. -C. Tsao, "The unfalsified control concept and learning," IEEE Transactions on Automatic Control, 42(6):843-847, June 1997.

[11] R. S. Smith and M. Dahleh, Eds., "The modeling of uncertainty in control system," Proc. of the 1992 Santa Barbara Workshop, Springer-Verlag, New York, 1994.

[12] J. C. Willems, "Paradigms and puzzles in the theory of dynamical systems," IEEE Transactions on Automatic Control, 36(3): 259-294, March 1991.

[13] B. R. Woodley, R. L. Kosut and J. P. How, "Uncertainty model unfalsification with simulation," Proc. American Control Conference 1998, Philadelphia, PA, June 22-26.