

EVALUATING CAUSAL RELATIONSHIPS IN WIRELESS SENSOR/ACTUATOR NETWORKS

Mark Coates

McGill University
Department of Electrical and Computer Engineering
3480 University St, Montreal, Quebec, Canada H3A 2A7.
coates@ece.mcgill.ca

ABSTRACT

This paper addresses the evaluation of causal relationships in large-scale wireless sensor/actuator networks (WSANs). WSANs consist of nodes that can monitor the environment (sensors), actively interact with the environment (actuators), or perform both functions. Such networks, together with the monitored environment, form complex, large-scale dynamic systems. It can be important to identify the causal relationships between variables in these systems; this provides individual actuator nodes with information about the probable consequences of local action on the global monitored environment. Such knowledge is critical to the formation of coordinated action plans, as might be required, for example, in distributed treatment strategies in medical or agricultural applications.

1. INTRODUCTION

Sensor networks have been successfully deployed for passively monitoring environments [1, 2], but there has been relatively little work towards developing networks that make changes to the environment. Wireless sensor and actuator networks (WSANs) represent an important extension, allowing nodes within the network to make autonomous decisions and perform actions (actuation) in response to sensor measurements and shared information. The potential applications of such WSANs are widespread, including agricultural maintenance and localized delivery of medication.

Causal assessment is an important step in the development of WSANs, enabling us to determine whether an actuation has an effect upon the monitored system, and if so, whether it is positive or negative. The formulation of distributed actuation strategies (or treatment plans) in the medical or agricultural context relies on an understanding of how the environment responds. We are interested in the problem of formulating a dynamic actuation strategy for achieving a desired response from the environment. In this paper we restrict our attention to an important first step in

the development of such a strategy: the estimation of the expected response of the system if a given strategy were applied.

The strategy is simply a procedure for deciding at the start of each time-interval t whether or not a given node should actuate; the decision is represented by the binary variable A_t . We consider that the WSAN measures at the start of the time interval a set of environmental state variables V_t . At the end of the time interval, the WSAN measures both the response of the system $Y_t(a)$ (for the actuation decision $A_t = a$) and an additional set of environmental factors W_t . These additional factors may or may not have a causal effect on the outcome Y_t . The WSAN makes these measurements whilst operating under a strategy g , which specifies a probability of actuation conditioned on actuation history and current environmental measurements. Our goal in this paper is to estimate the (counterfactual) expected response of the system if an alternative actuation strategy \tilde{g} had been applied, and hence we conduct our analysis within a counterfactual causality framework [3–5].

The paper is organized as follows. Section 2 specifies the problem we address and clarifies models the assumptions we impose. In Section 3, we describe a semi-parametric estimator that is founded on complexity-penalized (weighted) log-likelihood maximization using dyadic CART principles. Section 4 presents simulation results for an example WSAN system. Section 5 makes some concluding remarks and indicates future research directions.

Related work: Our problem formulation and estimator development is based on the work studying dynamic treatment regimes in [5–8] (the treatment regimes are equivalent to the actuation strategies explored in this paper). We depart from this work in adopting a semi-parametric model for the expected marginal response under a specific strategy. This is a similar approach to the semi-parametric estimation of the (static) treatment-specific mean response described in [9]. Therein a cross-validation approach is adopted for constructing a functional representation from a fixed set of basis functions. In this paper, we use a penalized likeli-

hood formulation and evaluate the estimator using dyadic CART with polynomial regression at the leaf nodes, effectively using a weighted version of the techniques described in [10, 11].

2. PROBLEM FORMULATION

We consider a WSN comprised of joint sensor-actuator nodes labelled $i = 1, \dots, N$. At the beginning of the discrete time interval t , node i measures a set of environmental variables $V_t^{(i)}$, and decides whether to actuate (indicated by the binary random variable $A_t^{(i)}$). In this paper we focus on probabilistic actuation decisions; the probability of actuation for node i at time t is determined by a function $g(V_t^{(i)}, S_t^{(i)})$, where $S_t^{(i)}$ is a set of summary statistics of the actuation history $A_{0:t}^{(i)}$. At the end of the time interval, the WSN node senses the environment to record a local response variable $Y_t^{(i)}$, together with system state variables $W_t^{(i)}$. Measurements are made at the discrete time-intervals $t = 0 : T$. We denote the observed data at the end of the t -th interval as $\{X_{0:t}, Y_{0:t}\}$, where $X_{0:t} = \{A_{0:t}^{(i)}, V_{0:t}^{(i)}, W_{0:t}^{(i)}\}_{i=1}^N$.

We aim to estimate $\Psi_{\tilde{g}}(a_t, L_t, S_t) = E(Y(a_t)|L_t, S_t, \tilde{g})$, a function that maps a subset L_t of the current system state variables V_t , the actuation history summary statistics S_t , and the current actuation decision a_t to an expected marginal response, under the condition that the dynamic actuation strategy were counterfactually specified by the function \tilde{g} . Note that the expectation here is over the probability distribution of being in state (a_t, L_t, S_t) at *any* time throughout the history $t = 0 : T$; this distribution is determined by the probability of the initial conditions $p(L_0)$, the counterfactual actuation decision strategy \tilde{g} and the system dynamics $p(X_t|X_{0:t-1})$. In order to ensure that $\Psi_{\tilde{g}}$ is identifiable and to make the estimation exercise tractable, we make several assumptions on the causal relationships of system variables and on the nature of the actuation and measurement processes.

First, we assume that the function $\Psi_{\tilde{g}}$ is not time-varying. We also assume that both actuation and system state variables have a completely local effect, i.e., conditioned on $V_t^{(i)}$, $W_t^{(i)}$ and $A_t^{(i)}$, the response variable $Y_t^{(i)}$ is independent of the state, actuation decision or response at another sensor j . We assume that the response of the system at the node $Y_t^{(i)}$ is conditionally independent of $\{V_{0:t-1}^{(i)}, W_{0:t-1}^{(i)}\}$ given the current state $\{V_t^{(i)}, W_t^{(i)}\}$ and a set of sufficient statistics $S_t^{(i)}$ that are functions the actuation history $A_{0:t}^{(i)}$. We assume that $V_t^{(i)}$, $S_t^{(i)}$ and $W_t^{(i)}$ are not affected by the current actuation decision $A_t^{(i)}$.

Denote by Z_T the *complete* data, which for each sensor includes all outcomes (all but one counterfactual) $Y_t(a_{0:T})$ for $a_{0:T} \in \mathcal{A}_T$, where \mathcal{A}_T is the set of all possible actuation histories to time T . As in [5, 7], we assume *sequential*

randomization, which specifies that for each $t = 0 : T$, $A_t^{(i)}$ is independent of Z_T given $V_{0:t}^{(i)}, W_{0:t-1}^{(i)}, A_{0:t-1}^{(i)}$. This assumption means that there is no systematic variation of the actual actuation decisions according to potential (counterfactual) past or future responses. In an alternative interpretation, there are no *unmeasured confounders* that affect the decision to actuate and the response of the system. Robins has shown that sequential randomization is a necessary assumption for the development of the G-computation estimator for dynamic treatment regimes [5, 6], which is the foundation of our estimator described in Section 3. We make the further assumption that any actuation history achievable through the adoption of the proposed actuation strategy \tilde{g} has positive probability of occurring under the observed actuation strategy g [6, 7].

3. NON-PARAMETRIC ESTIMATION OF THE MARGINAL EXPECTED RESPONSE

In this section we describe our estimator and the mechanism for its computation. We model a response measurement Y_t to a specific actuation a_t as the expected response plus additive Gaussian noise. Dropping the sensor-specific notation, we have,

$$Y_t(a_t) = \mathbf{E}(Y_t|X_{0:t}) + \epsilon_t, \quad (1)$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is zero-mean Gaussian with variance σ^2 .

We are interested in estimating the marginal mean response to an actuation decision a_t , when the system is in state $\{S_t, L_t\}$, if the system dynamics were being driven by an actuation decision function \tilde{g} . The probability of observing a sample $X_{0:t}^{(i)}$ is equal to the product of the probability of the initial conditions and the probability of the subsequent actuation history and the system conditions it induces. We have:

$$p(X_{0:t}) = p(X_0) \prod_{s=1}^{s=t} p(W_s, V_s, S_s|X_{0:s-1})g(A_s|V_s, S_s) \quad (2)$$

where $p(X_0) = p(V_0, W_0)g(A_0|V_0, S_0)$.

If we were interested in the response of the system under g , then the set of samples $\{X_t^{(i)}, Y_t^{(i)}\}$ for $t = 0 : T$ and $i = 1, \dots, N$ are drawn from the distribution over which we wish to calculate our marginal mean response, so we could form an estimate for Ψ_g by maximizing the likelihood of observing $Y_t^{(i)}$ given the samples $X_{0:t}^{(i)}$. In our formulation this would amount to minimizing a loss function equivalent to the sum of squared error between $Y_t^{(i)}$ and our estimate $\widehat{\Psi}_g(X_t^{(i)})$.

Since we are interested in the counterfactual decision function \tilde{g} , each observed actuation pattern has a different

counterfactual probability of occurrence. Our loss function needs to change to a weighted sum of square errors:

$$L(\widehat{\Psi}_{\tilde{g}}) = \sum_{i=1}^N \sum_{t=0}^T w_t^{(i)} (Y_t^{(i)} - \widehat{\Psi}_{\tilde{g}}(X_t^{(i)}))^2, \quad (3)$$

where

$$w_t^{(i)} = \prod_{s=0}^t \frac{\tilde{g}(A_s|L_s, S_s)}{g(A_s|L_s, S_s)}. \quad (4)$$

This estimator is based on the G-computation formula for mean response to a dynamic treatment regime [5].

We model the expected response $\Psi_{\tilde{g}}(a_t|L_t, V_t)$ as a piecewise smooth function, and use a complexity-penalized dyadic CART procedure to formulate an estimate of the function, based on the approaches described in [10, 11]. In a straightforward implementation of this procedure, we form a dyadic partitioning of the (L_t, V_t) space. The partitioning is achieved through the construction of a tree. The root represents a cell encompassing the entire space. At each layer in the tree we divide the cell in half in one of the dimensions; as we progress through subsequent layers, we cycle through the dimensions in turn. The resultant data structure is a tree of depth L , with a small, possibly empty, subset of the samples $X_t^{(i)}$ residing at any given leaf-node.

At each node in the tree we fit a function to the descendant data points using *weighted* least squares, where the weights are determined by (4). This function may be a wedgelet or a low-order polynomial; for the simulations reported in this paper, we adopt the computationally simpler estimate of a constant (the weighted mean) over each cell in the partition. We associate with each node in the tree a loss (or risk) value which is the sum of squared errors for the chosen fitted function. This is appropriate for the Gaussian model we have adopted.

The formulation of the final estimate is achieved by pruning the tree using a complexity penalty. The leaf nodes of the tree are at layer L . Commencing at layer $L-1$, we check if the loss value of a node in this layer is less than the sum of the loss values of its two children nodes and a constant complexity penalty α . If so, we prune the children from the tree. The process is repeated for each node j at layer $L-2$, but the comparison is now with the $|D_j|$ descendants of j in the pruned tree and the complexity penalty is $\alpha(|D_j| - 1)$. Continuing to the root of the tree, we have constructed the optimal complexity-penalized tree under the specified loss function. The choice of the constant α , which is related to the number of data points and the noise variance σ^2 , is discussed in more detail in [11].

The form of this estimate is reasonably well-suited to distributed implementation. Each WSAN node updates the weight appropriate for its measurements at each time interval. In [12], we describe in detail a distributed algorithm for

evaluation of the estimate, and analyze the computational and communication requirements.

4. AN EXAMPLE: SIMULATION RESULTS

We perform simulations to illustrate the performance of the proposed estimation method for an example scenario. We consider a model of an agricultural system, in which actuation corresponds to the release of a pesticide. Each WSAN node measures the carbon dioxide level L_t at the start of a day t and makes a decision A_t about the release of the pesticide. At the end of the day, it measures the change in weed-density Y_t , estimating it from digital images of the surrounding environment. It also records the change in soil moisture content W_t . The response of the system Y_t is dependent on the variables A_t, S_t, L_t, W_{t-1} and W_t . Here $S_t = \sum_{s=t-K}^{t-1} A_s$ is a statistic indicating the number of actuations in the recent history, and is related to the residual level of pesticide in the soil.

The current WSAN system employs an actuation strategy $g(V_t, S_t)$, where $V_t = \{W_{t-1}, L_t\}$. The actuation strategy is designed with the intent of matching the expected response function (choosing actuation when the expected response is significantly higher if pesticide is released). The lifetime of the sensor network can be increased substantially by turning off the cameras and soil-moisture sensors, so we would like to form an estimate of the expected marginal response of the system if we adopted an alternative actuation strategy $\tilde{g}(L_t, S_t)$ that is only dependent on the carbon dioxide level and the actuation history.

The model of the system is described by the following equations (these are used to generate the measurements in our simulations).

$$Y_t = -2 + 2A_t \cos(\pi S_t)(1 - 2L_t) + L_t^2 + 0.1A_t W_{t-1} + 0.3W_t + 0.2\epsilon_t \quad (5)$$

$$\mathbf{E}(Y_t|V_t, S_t) = -2 + 2A_t \cos(\pi S_t)(1 - 2L_t) + L_t^2 \quad (6)$$

$$g(V_t, S_t) = \cos^2\left(\frac{\pi S_t}{2}\right) e^{-L_t} \text{logit}^{-1}\left(-2 - \frac{W_{t-1}}{2}\right) \quad (7)$$

$$\tilde{g}(L_t, S_t) = \cos^2\left(\frac{\pi S_t}{2}\right) e^{-L_t} \quad (8)$$

Figures 1 and 2 provide an example of the performance of the estimator for a trial of 64 sensors measuring over 256 time instants. In the trial shown in the figures, 6388 data points were available in the estimated response region. Of these, 2224 corresponded to actuation decisions and 4164 to non-actuation decisions. Figure 2 indicates that the selected actuation function \tilde{g} is not a good match to the mean response; the probability of actuation is relatively high in regions where actuation is a poor choice.

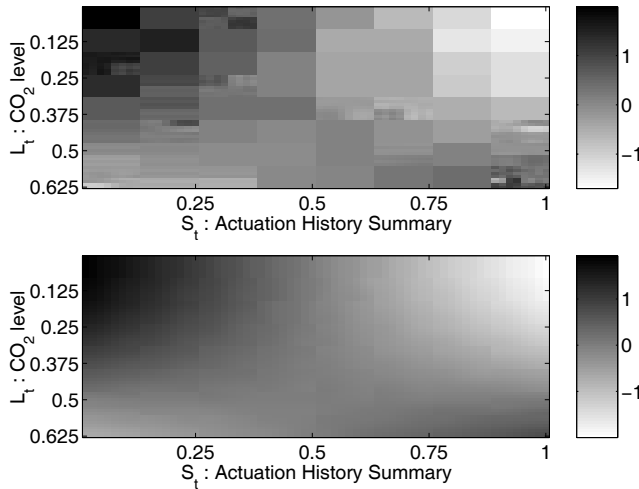


Fig. 1. A comparison between the estimated difference in expected response for the two cases $A_t = 1$ and $A_t = 0$ for an example trial and the true difference. The upper panel is $\hat{\Psi}_{\bar{g}}(S_t, L_t, a_t = 1) - \hat{\Psi}_{\bar{g}}(S_t, L_t, a_t = 0)$. The lower panel is the expected difference as determined from the generating equations of the simulation.

5. CONCLUSIONS

We have addressed the problem of estimating the expected marginal response of a system monitored by a wireless sensor-actuator network if a counterfactual actuation policy were implemented. Our estimator is based heavily on work assessing the response to dynamic treatment strategies [5–8]. The primary contribution of the paper is the identification of an approach for evaluating a specific semi-parametric estimator of the response. We have selected this estimator because its evaluation can be performed in a distributed fashion with relatively minor computational expense and communication overhead, critical aspects for any algorithm in a WSN. In future work, we will focus on developing distributed algorithms addressing the task of identifying optimum actuation policies.

6. REFERENCES

- [1] K.A. Delin, R.P. Harvey, N.A. Chabot, S.P. Jackson, M. Adams, D.W. Johnson, and J.T. Britton, “Sensor web in Antarctica: developing an intelligent, autonomous platform for locating biological flourishes in cryogenic environments,” in *Proc. 34th Lunar and Planetary Science Conf.*, Houston, TX, Mar. 2003.
- [2] D. Myung, B. Duncan, D. Malan, M. Welsh, M. Gaynor, and S. Moulton, “Vital dust: Wireless sensors and a sensor network for real-time patient monitoring,” in *Proc. 8th New England Reg. Trauma Conf.*, Burl., MA, Nov. 2003.
- [3] J. Neyman, “On the application of probability theory to agricultural experiments. (originally appeared in *Essay on prin-*

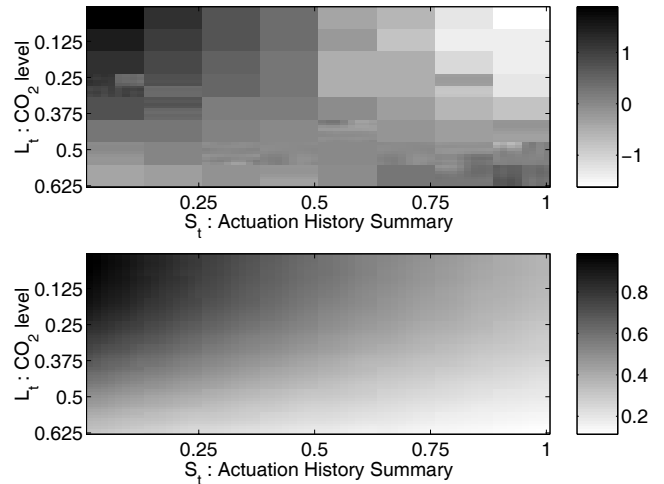


Fig. 2. A comparison between the estimated difference in expected mean response between actuation and no-actuation (upper panel) and the proposed actuation decision function (lower panel).

ciples (section 9), 1923),” *translated in Statistical Science*, vol. 5, pp. 34–58, 1990.

- [4] D.B. Rubin, “Bayesian inference for causal effects: The role of randomization,” *The Annals of Statistics*, vol. 6, pp. 34–58, 1978.
- [5] J.M. Robins, “A new approach to causal inference in mortality studies with a sustained exposure period: Application to control of the healthy worker survivor effect,” *Mathematical Modelling*, vol. 7, pp. 1393–1512, 1986.
- [6] J.M. Robins, *Latent Variable Modelling and Applications to Causality*, vol. 120 of *Lecture Notes in Statistics*, chapter Causal inference from complex longitudinal data, pp. 69–117, Springer-Verlag, New York, 1997.
- [7] S.A. Murphy, M.J. van der Laan, and J.M. Robins, “Marginal mean models for dynamic treatment regimes,” *J. American Statistical Association*, vol. 96, no. 456, pp. 1410–1423, 2001.
- [8] S.A. Murphy, “Optimal dynamic treatment regimes,” *J. Roy. Stat. Soc., Series B*, vol. 65, no. 2, pp. 331–366, 2003.
- [9] Y. Wang and M.J. van der Laan, “Data adaptive estimation of the treatment specific mean,” Tech. Rep., U.C. Berkeley, Division of Biostatistics, 2004.
- [10] C. Scott, R. Willett, and R. Nowak, “Cort: Classification or regression trees,” in *Proc. IEEE ICASSP*, Hong Kong, China, April 2003.
- [11] R. Castro, R. Willett, and R. Nowak, “Coarse-to-fine manifold learning,” in *Proc. IEEE ICASSP*, Montreal, Canada, May 2004.
- [12] Mark Coates, “Estimation of the marginal mean response using wireless sensor-actuator networks,” Tech. Rep., Department of Electrical and Computer Engineering, McGill University, 2005, available at <http://www.tsp.ece.mcgill.ca/Networks/>.