Mobile Robots for Learning Spatio-temporal Interpolation Models in Sensor Networks - The Echo State Map Approach

Erik Schaffernicht[1], Victor Hernandez Bennetts[1], and Achim J. Lilienthal[1]

***Abstract*— Sensor networks have limited capabilities to model complex phenomena occuring between sensing nodes. Mobile robots can be used to close this gap and learn local interpolation models. In this paper, we utilize Echo State Networks in order to learn the calibration and interpolation model between sensor nodes using measurements collected by a mobile robot. The use of Echo State Networks allows to deal with temporal dependencies implicitly, while the spatial mapping with a Gaussian Process estimator exploits the fact that Echo State Networks learn linear combinations of complex temporal dynamics. The resulting Echo State Map elegantly combines spatial and temporal cues into a single representation. We showcase the method in the exposure modeling task of building dust distribution maps for foundries, a challenge which is of great interest to occupational health researchers. Results from simulated data and real world experiments highlight the potential of Echo State Maps. While we focus on particulate matter measurements, the method can be applied for any other environmental variables like temperature or gas concentration.**

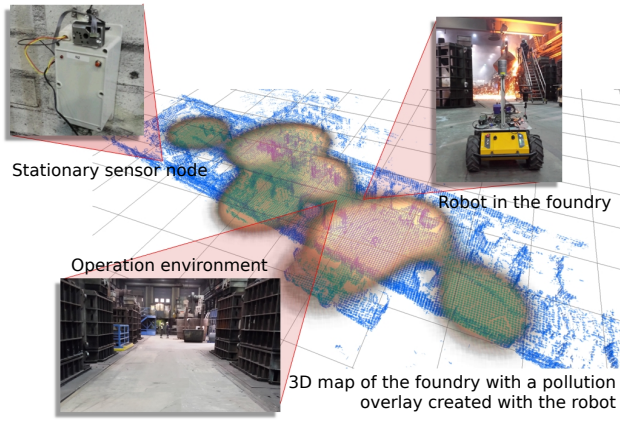

Fig. 1. The RAISE project: Dust measurements from stationary sensor nodes and a mobile robot are combinded to create dust exposure maps in industrial work environments like foundries.

I. Introduction

Wireless Sensor Networks (WSN) are usually installed in order to monitor certain environmental variables over a long time interval and are often subject to severe constraints regarding placement, budget and maintenance frequency. Yet, the resulting measurements and spatial models of the environment should be as accurate as possible. Hence it is necessary to calibrate the sensors properly and to devise an interpolation scheme to generate estimates at places without sensor nodes.

An example application is the creation of dust exposure models for workers in foundries. Excessive exposure to high concentrations of respirable dust causes severe damage to workers' health. Occupational health professionals are highly interested in precise dust distribution models for working environments [1], but high quality measurement devices are expensive and often only available for limited time intervals. It is therefore common practice to perform only point measurements at hand-picked places for short times [2]. It proves to be quite hard to properly generalize those measurements in space and time.

In this work, we propose a novel approach to estimate spatial models (i.e. maps) of environmental variables using Echo State Maps and a hybrid sensing framework, composed of a mobile robot and a WSN. The mobile robot is equipped with sophisticated, high quality sensors, while the sensors in the WSN are of cost effective, low fidelity devices. The contributions that we present in this paper are as follows:

- We introduce an approach that combines data collected with mobile robots and a WSN, using a learning scheme able to handle spatial and temporal interpolation. While the WSN provides non-calibrated measurements at sparse locations for very long periods of time, the robot can provide only time-limited, but calibrated data that is used to learn a model for interpolating between the sensing nodes.
- The comparative evalution of our proposed approach shows that it outperfoms simpler methods that only consider spatial aspects.
- We address a practical application, namely dust pollution monitoring in a foundry hall (see Fig. 1), which is of high interest for occupational health and safety specialists. A more detailed description of this application scenario, its challenges and individual data exploration, i.e. robot-only and sensor network-only, is available in [3].

II. Related Work

The estimation of the spatial distribution of environmental variables (e.g. gas, temperature, etc.) is a problem that has been widely studied. However, most works addresse either sensors on a mobile robot or a WSN, while the fusion of both technologies has yet to be fully explored.

Mobile robots have been widely used to create models of the spatial distribution of, e.g., gaseous emissions [4], [5]. Mobile platforms can only sense one location at a time and have limited operation times. This results in models, often based on kernel estimation methods [6], that are time-limited snapshots of the variable under study.

[1]AASS Research Centre, School of Science and Technology, Örebro University, Fakultetsgatan 1, Örebro, Sweden name.surname@oru.se

In contrast, the predicted models of WSNs are computed with relatively sparse sensing positions but at a high temporal resolution. With a limited number of nodes, research has been focused on optimal sensor placement [7] and local interpolation schemes [8], [9].

A hybrid monitoring system, composed of a mobile robot and a WSN, was presented in [10]. The authors used an aerial robot and a set of sensing nodes to create temperature maps in an ad-hoc testing environment. Under strong a priori assumptions (including time-invariance) regarding the spatial distribution of the ambient temperature, the authors thoroughly evaluated the predictions of their model and concluded that, in order to achieve high spatio-temporal measurement density and prediction accuracy, it is required to combine both, sensor networks and mobile robots.

III. Sensor Characteristics

Two different dust sensors are employed throughout this work. On one hand, we use a low-cost dust sensor as part of our sensing nodes (Fig. 2(a)), which as a reference is also mounted on the robot, and on the other hand, a high-cost dust sensor, usually used by occupational health professionals, which is only mounted on the robot (Fig. 2(b)). Both sensors have a sampling frequency of 1Hz.

The low-cost sensor is a Sharp GP2Y1010AU. It is an optical sensor for fine airborne particles (e.g. dust, smoke), and has a mass concentration resolution of 0.1 mg/m^3. The GP2Y1010AU is directly exposed to the environment and its output is proportional to the particle size and concentration, but it does not differentiate between different particle sizes. A few big particles might give the same reading as a high concentration of smaller particles. In the context of occupational health, the lack of selectivity to the particle size strongly limits the usefulness of the sensor, since the particulate matter size of interest is often very specific based on the type of dust in the environment.

In contrast, the Dusttrak II/DRX is much more sophisticated and expensive. The measurement principle is also based on a light-scattering photometer, but the air sample is pumped into a calibrated measurement chamber that adjusts e.g. for humidity. The dust concentrations are measured for specifically selected particulate matter sizes.

The goal of this work is to estimate and predict the readings of the Dusttrak at an arbitrary position in the area by combining the readings of all stationary low-cost sensor.

First, we investigated how correlated those two types of sensors are. Thereto, the Dusttrak and the Sharp sensor were mounted next to each other on the robot. $S^i = s^i_{t_1}, ..., s^i_{t_n}$ denotes the time series of dust measurements obtained by stationary sensing node i. $D_{(x,y)} = d_{t_1}, ..., d_{t_n}$ denotes the time series of dust measurements obtained by the Dusttrak mounted on the robot at position (x, y). We consider 11 measurement tours conducted in a foundry with changing conditions, each with a duration between 1 and 1.5 hours for a total of 14 hours.

Analyzing the signals of both sensors for those trials (two examples are given in Fig. 3), we observe the following:

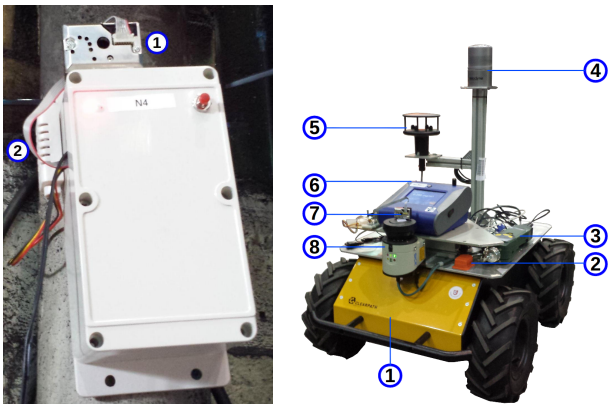

(a) (b)

Fig. 2. Sensors and plattforms. (a) Stationary sensing node (1) Sharp dust sensor, (2) temperature and humidity sensor. The gray box protects a microcontroller and a Wifi connector from the harsh environment conditions in foundries. (b) Dusttrak (6) mounted on a Clearpath Husky plattform (1). Other sensors of the plattform include a GPS/IMU (2), a photo inonization detector (3), LIDAR range sensors (4,8) for mapping and localization, an anemometer (5) and the Sharp dust sensor (7).

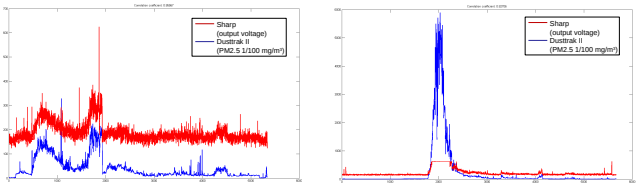

(a) (b)

Fig. 3. Two examples of the data collected (x-axis: time in sec, y-axis: voltage in mV) (a) 90 minutes tour through the foundry with various activities happening. (b) Sand casting creating big amounts of dust particles in close proximity to the robot. Please not that the Sharp sensor quickly goes into saturation.

- The signal-to-noise ratio of the Sharp sensor is much lower.
- The Sharp sensor goes into saturation in high dust concentration scenarios.
- In the non-saturation phases the signals are correlated.

Quantitatively, we calculated Pearson's correlation coefficient between both raw sensor signals[1] S and D as $r = 0.71$.

Using 11-fold cross-validation, we trained Echo State Networks (see next section) with different parameter settings, attempting to predict the Dusttrak signals D with the low-cost sensor data S as input to the network. The correlation coefficient between the network prediction and the real signal improved to $r = 0.87$ on average, indicating that the non-linear processing by the ESN is indeed helpful.

Furthermore, we investigated adding additional input variables (e.g. predicting the Dusttrak data from the low-cost sensor and humidity values) using weighted mutual information as a feature selection method [11]. Out of the investigated inputs (wind speed, wind direction, robot speed, robot pose, humidity and temperature), none but the temperature showed

[1]The problematic periods of saturation are not excluded from the evaluations presented. From an application's point of view, they are of high interest.

a significant improvement in the results of the dust estimation ($r = 0.95$). Examination of these results showed that the temperatures were different for all cases in which the low-cost sensor went into saturation. The network learned to predict the true dust levels based on temperature for these cases, yet occupational health experts and process engineers from the foundry are not aware of a causal relation. Thus, we focus on the dust readings for the remainder of the paper, a proper investigation of additional modalities is future work.

IV. Echo State Maps

In this section we will discuss Echo State Networks (ESN) and how to extend them spatially via Gaussian Processes (GP).

A. Echo State Networks

Introduced by Jaeger [12], [13], ESNs are a well established tool for times series analysis. They are well suited for non-linear signal processing, very quick and easy to train, and can be extended with new output units. Furthermore, the output layer structure is very suitable for spatial interpolation.

This neural network model consists of a single large, fully recurrent hidden layer, the so-called *reservoir*. The weights of this layer, denoted by the matrix $\mathbf{W}$, are not adapted, but instead randomly chosen, subject to the 'Echo State property' (the largest eigenvalue of $\mathbf{W} < 1$). As response to the input data this reservoir creates many temporal patterns that will fade over time (an echo). The internal state $\mathbf{X}$ of the network is computed as

$$\mathbf{X}(t) = f(\mathbf{W} * \mathbf{X}(t-1) + W^{in} * S_t). \tag{1}$$

S_t are the sensor measurements as discussed in the previous section. The function f is in most cases chosen as the hyperbolic tangent, introducing non-linearity to the system, while W^{in} are the randomly chosen weights from the input layer to the reservoir, which are fixed as well. The only learning step is to estimate the weights of the output layer W^{out} as a linear combination of those patterns within the reservoir to match the teaching signal D, which is often executed as a least-squares fitting via the pseudo-inverse:

$$W^{out}_{(x,y)} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}D_{(x,y)}. \tag{2}$$

This simple least-square fitting is not only very fast, but allows to learn new outputs in an already existing network without changing or retraining it.

The output of the network Y is computed as

$$Y_{(x,y)} = W^{out}_{(x,y)} * \mathbf{X}(t). \tag{3}$$

In the context of environmental mapping, the stationary sensor readings are the input time series to the network. We then learn a linear combination of the 'echos' of the input data to match the robot's Dusttrak readings $D_{(x_1,y_1)}$ at position (x_1, y_1), creating a predictor for the dust concentration at that position. Fig. 4(a) visualizes the described use of Echo State Networks for dust concentration estimation. Please note that this predictor implicitly learns the calibration for the sensors as part of the prediction model. Up to this point, there is no spatial information connected to the sensor input, this information will be introduced via the output weights in the next section.

Since the network can be extended easily, robot measurements at a new position (x_2, y_2) can be used to learn a different linear combination of 'echos' existing in the same network. In this way, we iteratively learn the linear combination for each cell in a grid map using a single ESN to create what we call an Echo State Map.

B. Spatial Interpolation of ESN weights

Acquiring the necessary training data for each single cell of the map is very time consuming and often completely unfeasible. For this reason, we propose the use of a Gaussian Process to inter- and extrapolate the output weights $W^{out}_{(x,y)}$ of the ESN to the unobserved cells of the map. This even allows for a continuous representation of the dust exposure, eliminating the explicit need for a grid map.

The general idea of interpolating the network weights between robot measurement positions is visualized in Fig. 4(b).

We follow directly Rasmussen and Williams [14], employing the prediction using noise-free observations. The observations in this case are the estimated output weights of the ESN, which are, according to the network training procedure, the optimal choice in a least-squares sense. Hence, the assumption of noise would change the weights at these known observation positions. This undesired effect is also the reason, why other regression techniques that do not exactly pass through the observations, like e.g. Nadaraya-Watson [15], which is often used in mobile robot olfaction scenarios [6], are unsuitable choices.

The predictive mean of the GP model is defined as:

$$W^{out}_{*} := k(P_*, P)[k(P, P)]^{-1}W^{out}. \tag{4}$$

W^{out} is the vector of a single trained output weight at different positions, while P encodes the 2D position of places with robot measurements, and by extension trained output weights. P_* is the list of query positions where W^{out}_{*} is to be estimated. k is the covariance function, which defines a similarity between the measurement places. We are using the squared exponential covariance function:

$$k(p_i, p_j) = exp\Big(-\frac{1}{2}\frac{|p_i - p_j|^2}{h^2}\Big). \tag{5}$$

h is the length-scale, a free parameter, changing the smoothness of the regression. It can be estimated automatically using log-likelihood maximization. Technically, Eq. 4 has to be evaluated for each output weight from the reservoir. Yet, the computational expensive part $k(P_*, P)[k(P, P)]^{-1}$ has to be computed only once. Therefore, this does not introduce a major computational burden for the method.

We currently do not use the predictive variance estimate provided by the GP framework, but it should be used as

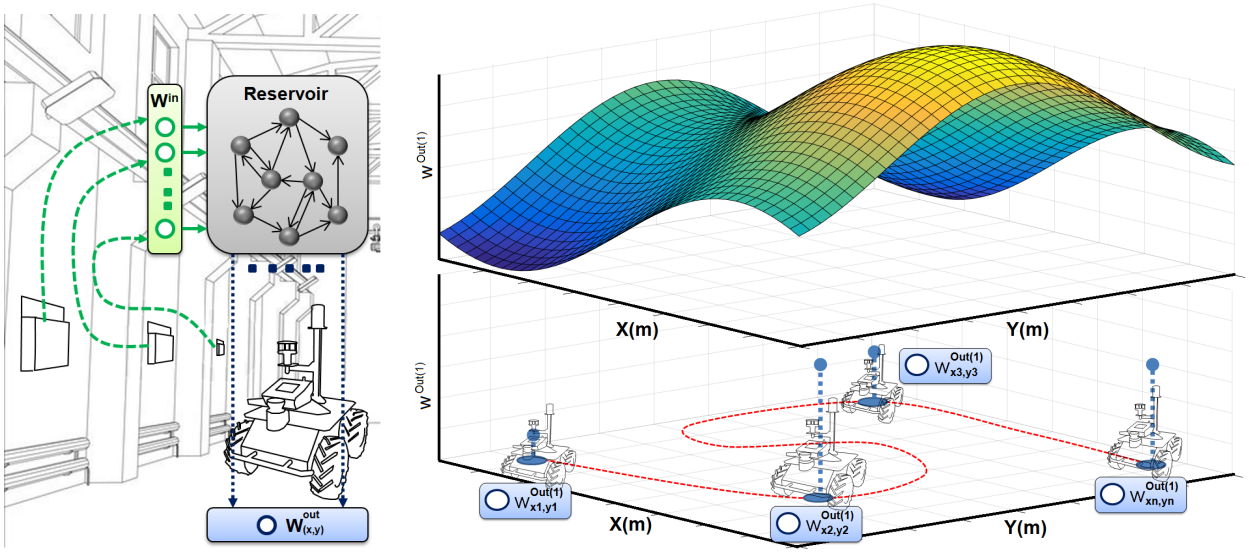

Fig. 4. Core concepts of Echo State Maps. (a) Learning the interpolation model for a robot measurement position. Sensor readings are fed into the ESN and propagated through the reservoir. With the training data provided by robot measurements the output weights are learned. (b) Interpolation of a single output weight of the Echo State Network. For each of the four robot measurement positions an independent estimate for the output weights is learned and then each weight individually is interpolated with a Gaussian Process.

additional information when interpreting Echo State Maps. Furthermore, it could be used as input for a robot exploration strategy, selecting the next measurement position of the robot based on the the predictive variance values.

Due to the random structure and possibly redundant dynamics inside the ESN, it is possible that very similar time series in neighboring positions will yield very different output weights. Since the weights are just a linear combination of the ESN state and the interpolation result is just another linear combination, they are all part of a linear manifold in the weight space. Hence, the transition of the interpolated output values will be smooth (and will change very little despite very different weights). Furthermore, this weight approach allows to interpolate spatial and temporal patterns in the map, compared to a output value interpolation which is purely spatial.

C. Simpler approaches to the problem

The Echo State Map approach will be compared against two other approaches. First, a simple GP interpolation between the sensing nodes (also known as kriging in spatial statistics) disregarding the robot measurements is used. The sensor readings are corrected with a linear calibration model. Here the dust readings are interpolated directly, compared to interpolating the model parameter in Echo State Maps.

Second, we solve a system of linear equations (SEQ) instead of using an ESN to build a model of linear combination of the sensor readings. For each of the robot's measurements at a single location an equation is formed as:

$$d_{t_n} = \alpha s^1_{t_n} + \beta s^2_{t_n} + \cdots + \omega s^i_{t_n} + C \qquad (6)$$

With enough measurements this becomes an overdetermined system of equations and the parameters $\alpha, \beta, \ldots, \omega, C$ are computed with a least squares method. We then interpolate the SEQ parameters between measurement positions with a GP as described in section IV-B. This represents a pure spatial interpolation disregarding temporal patterns.

V. Experimental comparison

First, in order to evaluate the proposed method properly, experiments in simulation with complete ground truth information were conducted. Afterwards we will present results obtained in a real foundry which captures the full complexity of the problem. However, due to the lack of ground truth in the foundries, quantitative evaluation is difficult.

A. Simulation

The used simulation is a very simplified industrial environment shown in Fig 5(a), yet it reflects major challenges encountered in real environments. It consists of two parts: the state/environment model and the sensor model. The environment model includes possible dust sources (representing workplaces that can produce a high dust concentration as well as ambient pollution), dust sinks (representing the place, where the air and dust leaves the area) and a simple airflow model (simulating the effects of the active ventilation used in foundries).

The environment model is organized as a grid map and for each simulation step for each cell the following steps are executed:

1) Introduce new dust particles at dust sources with a predefined activity pattern and high variance (representing work areas) and very low intensity, low variance everywhere else (ambient dust).
2) Remove dust particles at the dust sinks.

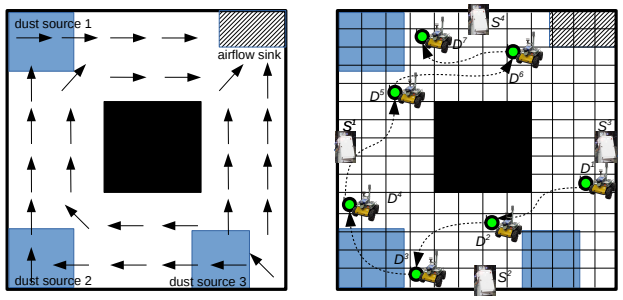

(a) (b)

Fig. 5. An example for the simulated environment. (a) Shows the position of the dust sources, the dust sink and the main airflow as a coarse visualization. The actual airflow model used for the simulation is of three times higher resolution compared to the depicted one. (b) One scenario for acquiring test data, showing the positions of the stationary sensor nodes, the robot measurement positions and grid cells used in the simulation.

3) Propagate particles according to the airflow model. Each cells contains its primary airflow direction, which a particle will follow with a probability of 0.6. In the other cases the particle will stay in its current cell or move to a random adjacent cell.
4) Apply the sensor model to simulate the readings of the stationary sensor nodes.

The sensor model applies all the observations mentioned in Sect. III. The underlying dust concetration value of the environment model is scaled and translated, before white noise is added, using the observed signal-to-noise ratio from Sect. III, and a cutoff is applied, in order to mimic the saturation effect observed with the real sensors at 800mV. For the measurements of the robot the unperturbed values of the simulation where used, reflecting the fact that the expensive sensor is the only ground truth available.

With a predefined pattern of source activity (no dust sources active, single sources active with varying dust intensity, multiple sources active, etc.) a sequence of dust readings for all cells is created, which subsequently were used as training and test data. Within such a scenario, we investigated the performance of the the proposed method, while testing different configurations. Changes included sensor location, robot measurement positions and length of the robot measurement steps. One such test configuration is shown in Fig. 5(b).

B. Results

With the part of data designated as training data set (measurements at specified locations and time intervals imitating robot operation), the output weights of the ESN and the parameters of the SEQ are computed as explained above and then interpolated using the Gaussian Process to build a complete model.

1) Simulator example: We are considering the exact configuration shown in Fig. 5(b). One sensors is placed in the middle of each outer wall, and training data is acquired with the robot at seven measurement positions for 4 minutes each. Afterwards the dust distribution is estimated by the different

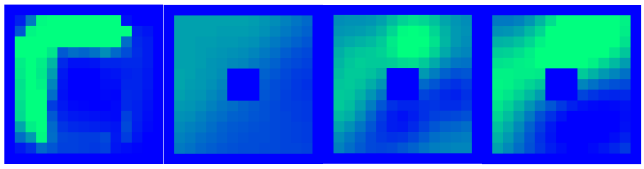

Fig. 6. From left to right: (1) Ground truth of the dust distribution. The situation includes high activity at dust source 2, and low activities at source 1 and 3. The effect of this configuration is a strong dust flow from lower left corner to the upper right corner via the upper half of the map. (2) Result of simple GP interpolation based on four sensors. (3) Estimate via system of equations. (4) Estimate by the Echo State Map.

approaches using only the current sensor readings of the four sensors. A snapshot of a single time step is shown in Fig. 6. In this case the sensor readings are high for S^1 and S^4 and low for S^2 and S^3. The simple GP based interpolation not only fails to capture the structure of the dust distribution, but as a results of the missing implicit calibration and a lower dust concentration along the wall, the estimated values are far too low. Both, the SEQ approach and the Echo State Map show structures more similar to the ground truth. The system of equations shows some false estimates in the lower right corner.

For a numerical evaluation of the model prediction quality at each grid cell of the simulator the quadratic difference between the estimated number of particles and the ground truth is computed and summed for the whole map. We then compute the mean and variance over the whole evaluation part of the data set (40 minutes). The numbers[2] in Table I show the best mean performance by the Echo State Map followed by the SEQ. Interestingly, for the shown example scenario the variance of the error of the SEQ estimate is even higher than for the simple GP interpolation. The second set of numbers is the average over 22 different experiments using the same environment, but different sensing configurations (varying the number of sensors, the number of robot measurements, the position of robot measurements and length of robot measurements).

TABLE I

MEAN AND VARIANCE OF THE SUMMED QUADRATIC DIFFERENCE OVER TIME FOR THE SCENARIO OF FIG. 5(B) AND THE AVERAGE FOR ALL SIMULATED EXPERIMENTS.

Method	Simple GP	SEQ	Echo State Map
mean (example)	9.6	6.7	3.7
variance (example)	8.6	10.9	7.2
mean (all)	12.4	8.7	5.3
variance (all)	11.6	10.3	7.9

The high variances observed are mainly caused by the simulated casting events, where huge numbers of dust particles are introduced. They accumulate a very high error until the dust plume reaches one of the sensors, resulting in a skewed distribtion.

[2]The unit is 10^8 particles, but this arbitrary and subject to the amount of particles used in the simulator.

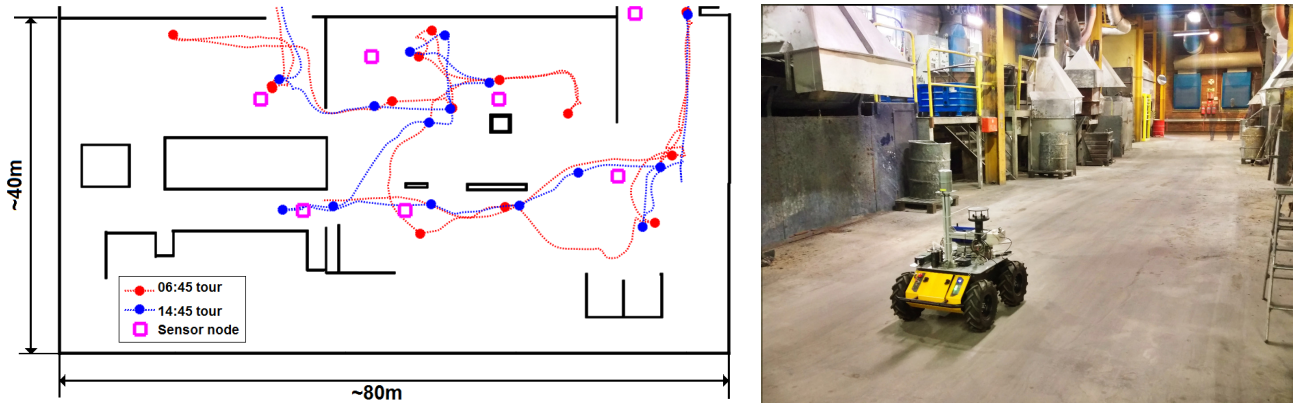

(a) (b)

Fig. 7. Foundry environment. (a) Operation environment in the foundry. Furnaces are situated in the left part of the map, while sand casting and centrifuge casting facilities are situated in the right part of the map. Seven stationary sensor nodes are installed in the area. The robot was performing two measurement tours on the same day, one in the morning (red) and the other in the afternoon (blue). (b) Picture of the environment with the robot at the lower-left position of the blue path.

2) General Observations: Unsuprisingly, the simple distance based interpolation approach is consistently producing the highest deviations from the underlying dust distribution. The method lacks information to differentiate the dust exposure patterns that are created by the airflow.

Most of the time, the ESN approach is outperforming the SEQ approach, especially in situations with a lot of changes in the source activity patterns. As expected, the capabilty to process the data in non-linear fashion as a times series puts the ESN ahead of the SEQ, which just considers each time step individually.

Yet there are scenarios in which both methods produce very similar results. First, with very limited training data per location (often containing only a single source activity change) the SEQ estimate generalizes better, while the ESN overfits to the observed change and deviates in other conditions. Second, the further away from the source the sensors are positioned the smaller the difference between both approaches becomes. Part of the issue is also the applied airflow model, which over distance averages the dust concentration and removes temporal variances. Third, the more robot measurement positions are used the bigger the advantage of the ESN becomes. With very few measurement positions, the only benefit of using an ESN is the local signal processing capability. Adding further measurement positions allows the ESN to propagate temporal patterns through the map.

Regarding internal parameters of the employed ESN, we investigated the influence of the size of the reservoir (from 20 to 2000 neurons) and sparsity of connections therein (from fully connected to 20% connectivity). While there was no clear result regarding sparsity, we observed a distinct response to changes in the size of the reservoir. The mean error was not effected significantly, yet the variance increased with increased size of the reservoir. This behaviour can be reduced by increasing the number of training data (i.e. lengthing the observation time), yet this is undesireable for most applications. In ESN applications it is often favourable to use large reservoirs in order to create rich dynamics, but this leads to another potential issue for very large maps and reservoirs. Since the number of output weights in an Echo State Map is *number of cells* $\times$ *number of neurons*, there might be a limitation concerning available memory. Therefore, we prefer smaller reservoirs.

Since the sampling rate of the employed sensors is rather high (1Hz) compared to the phenomena observed, we kept the spectral radius of the reservoir (highest eigenvalue of the weight matrix $\mathbf{W}$) close to 1. Changing this parameter towards values closer to 0 causes the ESN to perform more like the SEQ, since this parameter controls time constant of the system (i.e. the lifespan of the 'echos' inside the reservoir).

C. Echo State Map from a foundry

In order to test the proposed approach in a real world scenario, we built a dust distribution model in a $300m^2$ area of a foundry shop-floor. A map of the environment with sensor nodes and the robot's exploration paths for two tours are shown in Fig. 7. The length of each tour was slightly over an hour each. The robot stopped for two to three minutes to collect data at several waypoints (denoted by circular markers in Fig 7).

Normal distribution transform (NDT) maps [16] of the environment were built to localize the robot via NDT-MCL [17]. We assume that the uncertainty introduced by errors in the localization system does not significantly influence the results, since they are much smaller compared to uncertainty introduced by turbulent airflow regimes and the chaotic nature of dust emitting processes in a foundry.

The main challenge in the foundry is the absence of ground truth data for evaluation. Instead, we consider the scenario of one robot measurement tour executed to build the exposure model. The data collected during a second measurement tour is used to evaluate the previously built model. It is therefore not possible to evaluate the complete

map, but only the positions the robot stopped at during the second tour. During the first tour mutiple dust producing activities with low-medium intensity where ongoing, while during the second one it was mostly quiet with one casting (medium-high intensity) being performed. The root mean square error of the second measurement tour was $693\mu g/m^3$ for the purely sensor based interpolation with GP, $459\mu g/m^3$ using the SEQ approach and $322\mu g/m^3$ using an Echo State Map.

During the uneventful parts of the second tour all three models could model the low dust conditions. The casting event, in which dust concentrations raise significantly, threw the sensor-based interpolation far off, while the other two approaches gave reasonable results. The observed effects are very similar to those in simulation. One of the issues encountered here was the previously discussed effect with the stationary sensors being further away from the casting event than the robot, and the resulting delay before the sensor measure the changes and the models update accordingly.

VI. Conclusions

We introduced Echo State Maps as a methodology to combine a wireless sensor network with localized robot measurements in order to create a dense interpolation model. Exploiting robot advantages like flexibility in movements, mapping and localization capabilities and combining them with stationary sensor readings allows to build improved environment models. The used Echo State Network additionally provides intrinsic signal processing capabilities.

The initial application of the method in a foundry showed promising results. Currently, we are investigating how to bring these results to, e.g. occupational health studies. This ongoing work requires additional ground truth data, collected during longer periods of time and with additional high quality stationary sensors.

If dependency between the temperature and dust concentration holds (see Sec. III), it would be preferable to integrate this additional variable into the estimation model. This is rather easy to achieve, since it only requires to add additional inputs, one per temperature sensor, to the network of the Echo State Map. Extensions to include further different variables are easy to implement in the presented framework, but have to be balanced against the increased danger for overfitting the model.

Another issue to be explored is the long-term use of Echo State Maps. Here we presented the initial construction of the exposure model, yet the data collected in further robot measurement tours should be used to update the existing Echo State Map. The quick training algorithm makes this certainly feasible, yet this includes interesting research questions regarding knowledge management (i.e. what to forget, what to replace and what to keep or whether to learn a new Echo State Map from scratch), robot exploration strategies in order to identify interesting areas that are volatile and therefore interesting to observe, as well as planning and executing regular autonomous inspection tours with the robot.

Acknowledgements

This work was funded in part by the Swedish Knowledge Foundation (Diarienr. 20130196, RAISE).

References

[1] L. Andersson, I. Bryngelsson, Y. Ngo, C. Ohlson, and H. Westberg, "Exposure assessment and modeling of quartz in swedish iron foundries for a nested case-control study on lung cancer," *Occupational and Environmental Hygiene*, vol. 9, no. 2, pp. 110–119, 2012.

[2] L. Andersson, A. Burdorf, I.-L. Bryngelsson, and H. Westberg, "Estimating trends in quartz exposure in swedish iron foundries predicting past and present exposures," *Annals of Occupational Hygiene*, vol. 56, no. 3, pp. 362–372, 2012.

[3] V. Hernandez, E. Schaffernicht, A. J. Lilienthal, H. Fan, T. P. Kucner, L. Andersson, and A. Johansson, "Towards occupational health improvement in foundries through dense dust and pollution monitoring using a complementary approach with mobile and stationary sensing nodes," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 131–136.

[4] P. Neumann, S. Asadi, A. J. Lilienthal, M. Bartholmai, and J. Schiller, "Micro-Drone for Wind Vector Estimation and Gas Distribution Mapping," *IEEE Robotics and Automation Magazine*, vol. 19, pp. 50–61, 2012.

[5] V. Hernandez, E. Schaffernicht, V. Pomareda, A. Lilienthal, S. Marco, and M. Trincavelli, "Combining non selective gas sensors on a mobile robot for identification and mapping of multiple chemical compounds," *Sensors*, vol. 14, no. 9, pp. 17 331–17 352, 2014. [Online]. Available: http://www.mdpi.com/1424-8220/14/9/17331

[6] A. J. Lilienthal, M. Reggente, M. Trincavelli, J. L. Blanco, and J. Gonzalez, "A statistical approach to gas distribution modelling with mobile robots the kernel dm+v algorithm," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 11 - October 15 2009, pp. 570–576.

[7] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, Jun. 2008.

[8] M. Umer, L. Kulik, and E. Tanin, *Kriging for Localized Spatial Interpolation in Sensor Networks*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 525–532. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-69497-7_34

[9] X. Gao and L. Acar, "Multi-sensor integration to map odor distribution for the detection of chemical sources," *Sensors*, vol. 16, no. 7, p. 1034, 2016. [Online]. Available: http://www.mdpi.com/1424-8220/16/7/1034

[10] W. C. Evans, D. Dias, S. Roelofsen, and A. Martinoli, "Environmental field estimation with hybrid-mobility sensor networks," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 5301–5308.

[11] E. Schaffernicht and H.-M. Gross, "Weighted mutual information for feature selection," in *Int. Conf. on Artifical Neural Networks ICANN (2)*, 2011, pp. 181–188.

[12] H. Jaeger, "Adaptive nonlinear system identification with echo state networks," in *Advances in Neural Information Processing Systems 15*, S. Becker, S. Thrun, and K. Obermayer, Eds. MIT Press, 2003, pp. 609–616.

[13] H. Jaeger and H. Haas, "Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication," vol. 304, no. 5667, pp. 78–80, 2004. [Online]. Available: http://science.sciencemag.org/content/304/5667/78

[14] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2005.

[15] E. Nadaraya, "On estimating regression," *Theory of Probability and Its Applications*, vol. 9, no. 1, pp. 141–142, 1964.

[16] J. Saarinen, H. Andreasson, T. Stoyanov, J. Ala-Luhtala, and A. Lilienthal, "Normal distributions transform occupancy maps: Application to large-scale online 3d mapping," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, May 2013, pp. 2233–2238.

[17] J. Saarinen, H. Andreasson, T. Stoyanov, and A. Lilienthal, "Normal distributions transform monte-carlo localization (ndt-mcl)," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, Nov 2013, pp. 382–389.