

Deep Tracking: Seeing Beyond Seeing Using Recurrent Neural Networks

Peter Ondruška and Ingmar Posner

Mobile Robotics Group, University of Oxford, United Kingdom
{ondruska, ingmar}@robots.ox.ac.uk

Abstract

This paper presents to the best of our knowledge the first end-to-end object tracking approach which directly maps from raw sensor input to object tracks in sensor space without requiring any feature engineering or system identification in the form of plant or sensor models. Specifically, our system accepts a stream of raw sensor data at one end and, in real-time, produces an estimate of the entire environment state at the output including even occluded objects. We achieve this by framing the problem as a deep learning task and exploit sequence models in the form of recurrent neural networks to learn a mapping from sensor measurements to object tracks. In particular, we propose a learning method based on a form of input dropout which allows learning in an *unsupervised* manner, only based on raw, occluded sensor data without access to ground-truth annotations. We demonstrate our approach using a synthetic dataset designed to mimic the task of tracking objects in 2D laser data – as commonly encountered in robotics applications – and show that it learns to track many dynamic objects despite occlusions and the presence of sensor noise.

Introduction

As we demand more from our robots the need arises for them to operate in increasingly complex, dynamic environments where scenes – and objects of interest – are often only partially observable. However, successful decision making typically requires complete situational awareness. Commonly this problem is tackled by a processing pipeline which uses separate stages of object detection and tracking, both of which require considerable hand-engineering. Classical approaches to object tracking in highly dynamic environments require the specification of plant and observation models as well as robust data association.

In recent years neural networks and deep learning approaches have revolutionised how we think about classification and detection in a number of domains (Krizhevsky, Sutskever, and Hinton 2012; Szegedy, Toshev, and Erhan 2013). The often unreasonable effectiveness of such approaches is commonly attributed to both an ability to learn




Figure 1: Often a robot’s sensors provide only partial observations of the surrounding environment. In this work we leverage a recurrent neural network to effectively reveal occluded parts of a scene by learning to track objects from raw sensor data – thereby effectively reversing the sensing process.

relevant representations directly from raw data as well as a vastly increased capacity for function approximation afforded by the depth of the networks (Bengio 2009).

In this work we propose a framework for *deep tracking*, which effectively provides an off-the-shelf solution for learning the dynamics of complex environments directly from raw sensor data and mapping it to an intuitive representation of a complete and unoccluded scene around the robot as illustrated in Figure 1.

To achieve this we leverage Recurrent Neural Networks (RNNs), which were recently demonstrated to be able to effectively handle sequence data (Graves 2013; Sutskever, Hinton, and Taylor 2009). In particular we consider a complete yet uninterpretable hidden state of the world and then train the network end-to-end to update its belief of this state using a sequence of partial observations and map it back into an interpretable unoccluded scene. This gives the network a freedom to optimally define the content of the hidden state

and the considered operations without the need of any hand-engineering.

We demonstrate that such a system can be trained in an entirely *unsupervised* manner based only on raw sensor data and without the need for supervisor annotation. To the best of our knowledge this is the first time such a system has been demonstrated and we believe our work will provide a new paradigm for an end-to-end tracking. In particular, our contributions can be summarized as follows:

- A novel, end-to-end way of filtering partial observations of dynamic scenes using recurrent neural networks to provide a full, unoccluded scene estimation.
- A new method of unsupervised training based on a novel form of dropout encouraging the network to correctly predict objects in unobserved spaces.
- A compelling demonstration and attached video ¹ of the method in a synthesised scenario of a robot tracking the state of the surrounding environment populated by many moving objects. The network directly accepts unprocessed raw sensor inputs and learns to predict positions of all objects in real time even through complete sensor occlusions and noise ².

Deep Tracking

Formally, our goal is to predict the current, un-occluded scene around the robot given the sequence of sensor observations, i.e. to model $P(y_t|x_{1:t})$ where $y_t \in \mathbb{R}^N$ is a discrete-time stochastic process with unknown dynamics modelling the scene around the robot containing other static and dynamic objects and $x_t \subseteq y_t$ are the sensor measurements of directly visible scene parts. We use encoding $x_t = \{v_t, r_t\}$ where $\{v_t^i, z_t^i\} = \{1, y_t^i\}$ if i -th element of y_t is observed and $\{0, 0\}$ otherwise. Moreover, as the scene changes the robot can observe different parts of the scene at different times.

In general, this problem cannot be solved effectively if the process y_t is purely random and unpredictable, as in such a case there is no way to determine the unobserved elements of y_t . In practice, however, the sequences y_t and x_t exhibit structural and temporal regularities which can be exploited for effective prediction. For example, objects tend to follow certain motion patterns and have certain appearance. This allows the behaviour of the objects to be estimated when they are visible, such that this knowledge can be used later for prediction when they become occluded.

Deep tracking leverages recurrent neural networks to effectively model $P(y_t|x_{1:t})$. In the remainder of this section we first describe the probabilistic model assumed to underly the sensing process before detailing how RNNs can be used for tracking.

The Model

Like many approaches to object tracking, our model is inspired by Bayesian filtering (Chen 2003). Note, however,

¹Video is available at: <https://youtu.be/cdeWCpfUGWc>

²The source code of our experiments is available at: http://mrg.robots.ox.ac.uk/mrg_people/peter-ondruska/




Figure 2: The assumed graphical model of the generative process. World dynamics are modelled by the hidden Markov process h_t with appearance y_t which is partially observed by sensor measurements x_t .

the process y_t does not satisfy the Markov property: Even knowledge of y_t at any particular time t provides only partial information about the state of the world such as object positions, but does not contain all necessary information such as their speed or acceleration required for prediction. The latter can be obtained only by relating subsequent measurements. Methods such as the Hidden Markov Model (Rabiner 1989) are therefore not directly applicable.

To handle this problem we assume the generative model in Figure 2 such that alongside y_t there exists another underlying Markov process, h_t , which completely captures the state of the world with the joint probability density

$$P(y_{1:N}, x_{1:N}, h_{1:N}) = \prod_{t=1}^N P(x_t|y_t)P(y_t|h_t)P(h_t|h_{t-1}), \quad (1)$$

where

- $P(h_t|h_{t-1})$ denotes the hidden state transition probability capturing the dynamics of the world;
- $P(y_t|h_t)$ is modelling the instantaneous *unoccluded* sensor space;
- and $P(x_t|y_t)$ describes the actual sensing process.

The task of estimating $P(y_t|x_{1:t})$ can now be framed in the context of recursive Bayesian estimation (Bergman 1999) of the *belief* \mathcal{B}_t which, at any point in time t , corresponds to a distribution $Bel(h_t) = P(h_t|x_{1:t})$. This belief can be computed recursively for the considered model as

$$Bel^-(h_t) = \int_{h_{t-1}} P(h_t|h_{t-1})Bel(h_{t-1}) \quad (2)$$

$$Bel(h_t) \propto \int_{y_t} P(x_t|y_t)P(y_t|h_t)Bel^-(h_t) \quad (3)$$

where $Bel^-(h_t)$ denotes *belief prediction* one time step into the future and $Bel(h_t)$ denotes the *corrected* belief after the latest measurement has become available. $P(y_t|x_{1:t})$ in turn can then be expressed simply using this belief,

$$P(y_t|x_{1:t}) = \int_{h_t} P(y_t|h_t)Bel(h_t). \quad (4)$$




Figure 3: An illustration of the filtering process using a recurrent neural network.

Filtering Using a Recurrent Neural Network

Computation of $P(y_t|x_{1:t})$ is carried easily by iteratively updating the current belief \mathcal{B}_t

$$\mathcal{B}_t = F(\mathcal{B}_{t-1}, x_t) \quad (5)$$

defined by Equations 2,3 and simultaneously providing prediction

$$P(y_t|x_{1:t}) = P(y_t|\mathcal{B}_t) \quad (6)$$

defined by equation 4. Moreover, instead of y_t the same method can be used to predict any state in the future $P(y_{t+n}|x_{1:t})$ by providing *empty* observations of the form $x_{(t+1):(t+n)} = \emptyset$.

For realistic applications this approach is, however, limited by the need for the suitable belief state representation as well as the explicit knowledge of the distributions in Equation 1 modelling the generative process.

The key idea of our solution is to avoid the need to specify this knowledge and instead use a highly expressive neural networks, governed by weights W_F and W_P , to model both $F(\mathcal{B}_{t-1}, x_t)$ and $P(y_t|\mathcal{B}_t)$ allowing to learn their function directly from the data. Specifically, we assume \mathcal{B}_t can be approximated and represented as a vector, $\mathcal{B}_t \in \mathbb{R}^M$. As illustrated in Figure 3, the function F is then simply a recurrent mapping from $\mathbb{R}^M \times \mathbb{R}^{2N} \rightarrow \mathbb{R}^M$ corresponding to an implementation by a Feed-Forward Recurrent Neural Network (Medsker and Jain 2001) where the \mathcal{B}_t acts as the network's memory passed from one time step to the next.

Importantly, we do not impose any restrictions on the content or function of this hidden representation. Provided both networks for $F(\mathcal{B}_t, x_t)$ and $P(y_t|\mathcal{B}_t)$ are differentiable they can be trained together end-to-end as a single recurrent network with a cost function directly corresponding to the data likelihood provided in Equation 4. This allows the neural networks to adapt to each other and to learn an optimal hidden state representation for \mathcal{B}_t together with procedures for belief updates using partial observations x_t and to use this information to decode the final scene state y_t . Training of the network is detailed in the next section while the concrete architecture used for our experiment is discussed as a part of the Experimental Results.

Training

The network can be trained both in a supervised mode, i.e. with both known $y_{1:N}, x_{1:N}$ as well as in an unsupervised mode where only $x_{1:N}$ is known. A common method of supervised training is to minimize the negative log-likelihood of the true scene state given the sensor input

$$\mathcal{L} = - \sum_{t=1}^N \log P(y_t|x_{1:t}) \quad (7)$$

This objective can be optimised by the gradient descent with partial derivatives given by backpropagation through time (Rumelhart, Hinton, and Williams 1985)

$$\frac{\partial \mathcal{L}}{\partial \mathcal{B}_t} = \frac{\partial F(\mathcal{B}_t, x_{t+1})}{\partial \mathcal{B}_t} \cdot \frac{\partial \mathcal{L}}{\partial \mathcal{B}_{t+1}} - \frac{\partial \log P(y_t|\mathcal{B}_t)}{\partial \mathcal{B}_t} \quad (8)$$

$$\frac{\partial \mathcal{L}}{\partial W_F} = \sum_{t=1}^N \frac{\partial F(\mathcal{B}_{t-1}, x_t)}{\partial W_F} \cdot \frac{\partial \mathcal{L}}{\partial \mathcal{B}_t} \quad (9)$$

$$\frac{\partial \mathcal{L}}{\partial W_P} = - \sum_{t=1}^N \frac{\partial \log P(y_t|\mathcal{B}_t)}{\partial W_P} \quad (10)$$

However, a major disadvantage of the supervised training is the requirement of the ground-truth scene state y_t containing parts not visible by the robot sensor. One way of obtaining this information is to place additional sensors in the environment which jointly provide information about every part of the scene. Such an approach, however, can be impractical, costly or even impossible. In the following section we describe a training method which only requires raw, occluded sensor observations alone. This dramatically increases the practicality of the method, as the only necessary learning information is a long enough record of $x_{1:N}$, as obtained by the sensor in the field.

Unsupervised Training

The aim of unsupervised training is to learn $F(\mathcal{B}_{t-1}, x_t)$ and $P(y_t|\mathcal{B}_t)$ using only $x_{1:t}$. This might seem impossible without knowledge of y_t as there would be no way to correct the network predictions. However, *some* values of y_t as $x_t \subseteq y_t$ are known in the directly observed sensor measurements. A naive approach would be to train the network using the objective from Equation 7 of predicting the known values of y_t and not back-propagating from the unobserved ones. Such an approach would fail, however, as the network would only learn to copy the observed elements of x_t to the output and never predict objects in occlusion.

Instead, and crucially, we propose to train the network not to predict the current state $P(y_t|x_{1:t})$ but a state in the future $P(y_{t+n}|x_{1:t})$. In order for the network to correctly predict the visible part of the object after several time steps the network must learn how to represent and simulate object dynamics during the time when the object is not seen using F and then convert this information to predict $P(y_{t+n}|\mathcal{B}_{t+n})$.

This form of learning can be implemented easily by considering the cost function from Equation 7 of predicting visible elements of $y_{t:t+n}$ and *dropping-out* all $x_{t:t+n}$ by setting them to 0 marking the input unobserved. This is similar to

standard node dropout (Srivastava et al. 2014) to prevent network from overfitting. Here, we are dropping out the observations both spatially across the entire scene and temporarily across multiple time steps in order to prevent the network to overfit to the task of simply copying over the elements. This forces the network to learn the correct patterns.

An interesting observation demonstrated in the next section is that even though the network is in theory *trained to predict only the future observations* x_{t+n} , in order to achieve this the network learns to predict y_{t+n} . We hypothesise this is due to the tendency of the regularized network to find the simplest explanation for observations where it is easier to model $P(y_t|\mathcal{B}_t)$ than $P(x_t|\mathcal{B}_t) = \int_{y_t} P(x_t|y_t)P(y_t|\mathcal{B}_t)$.

Experimental Results

In this section we demonstrate the effectiveness of our approach in estimating the full scene state in a simulated scenario representing a robot equipped with a 2D laser scanner surrounded by many dynamic objects. Our inspiration here is a situation of a robot in the middle of a busy street surrounded by a crowd of moving people. Objects are modelled as circles independently moving with constant velocity in a random direction but never colliding with each other or with the robot. The number of objects present in the scene is varied over time between two and 12 as new objects randomly appear and later disappear at a distance from the robot.

Sensor Input

At each time step we simulated the sensor input as commonly provided by a planar laser scanner. We divided the scene y_t around the robot using a 2D grid of 50×50 pixels with the robot placed towards the bottom centre. To model sensor observations x_t we ray-traced the visibility of each pixel and assigned values $x_t^i = \{v_t^i, r_t^i\}$ indicating pixel visibility (by the robot) and presence of an obstacle if the pixel is visible. An example of the input y_t and corresponding observation x_t is depicted in Figure 3. Note that, as in the case of a real sensor, only the surface of the object as seen from particular viewpoint is ever observed. This input was then fed to the network as a stream of 50×50 2-channel binary images for filtering.

Neural Network

To jointly model $F(\mathcal{B}_{t-1}, x_t)$ and $P(y_t|\mathcal{B}_t)$ we used a small feed-forward recurrent network as illustrated in Figure 4. This architecture has four layers and uses convolutional operations followed by a sigmoid nonlinearity as the basic information processing step at every layer. The network has in total 11k parameters, and its hyperparameters, such as number of channels in each layer and size of the kernels, were set by cross-validation. The belief state \mathcal{B}_t is represented by the 3rd layer of size $50 \times 50 \times 16$ which is kept between time steps.

The mapping $F(\mathcal{B}_{t-1}, x_t)$ is composed of the stage of input-preprocessing (the *Encoder*) followed by a stage of hidden state propagation (the *Belief tracker*). The aim of the *Encoder* is to analyse the sensor measurements, to perform operations such as detecting objects directly visible to the

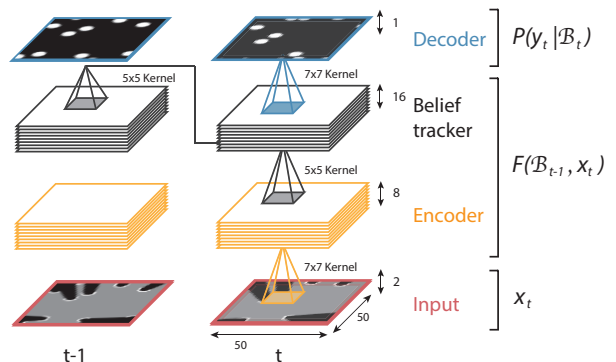


Figure 4: The 4-layer recurrent network architecture used for the experiment. Directly visible objects are detected in *Encoder* and fed into *Belief tracker* to update belief \mathcal{B}_t used for scene deocclusion in *Decoder*.

sensor and to convert this information into a $50 \times 50 \times 8$ embedding as input to the *Belief tracker*. This information is concatenated with the previous belief state \mathcal{B}_{t-1} and in *Belief tracker* combined into a new belief state \mathcal{B}_t .

Finally, $P(y_t|\mathcal{B}_t)$ is modelled by a Bernoulli distribution by interpreting the network final layer output as a *probabilistic occupancy grid* p_t of the corresponding pixels being part of an object giving rise to the total joint probability

$$P(y_t|p_t) = \prod_i (p_t^i)^{y_t^i} (1 - p_t^i)^{(1-y_t^i)} \quad (11)$$

with its logarithm corresponding to the binary cross-entropy. One pass through the network takes 10ms on a standard laptop drawing the method suitable for real-time data filtering.

Training

We generated 10,000 sequences of length 200 time steps and trained the network for a total of 50,000 iterations using stochastic gradient descent with learning rate 0.9. The initial belief \mathcal{B}_t was modelled as a free parameter being jointly optimised with the network weights W_F, W_P .

Both supervised (i.e. when the groundtruth of y_t is known) and unsupervised training were attempted with almost identical results. Figure 5 illustrates the unsupervised training progress. Initially, the network produces random output, then it gradually learns to predict the correct shape of the visible objects and finally it learns to track their position even through complete occlusion. As illustrated in the attached video, the fully trained network is able to confidently and correctly start tracking an object immediately after it has seen even only a part of it and is then able to confidently keep predicting its position even through long periods of occlusion.

Finally, Figure 6 shows example activations of layers from different parts of the network. It can be seen that the *Encoder* module learned to produce in one of its layers a spike at the center of visible objects. Particularly interesting is the learned structure of the *belief state*. Here, the network learned to represent hypotheses of objects having different

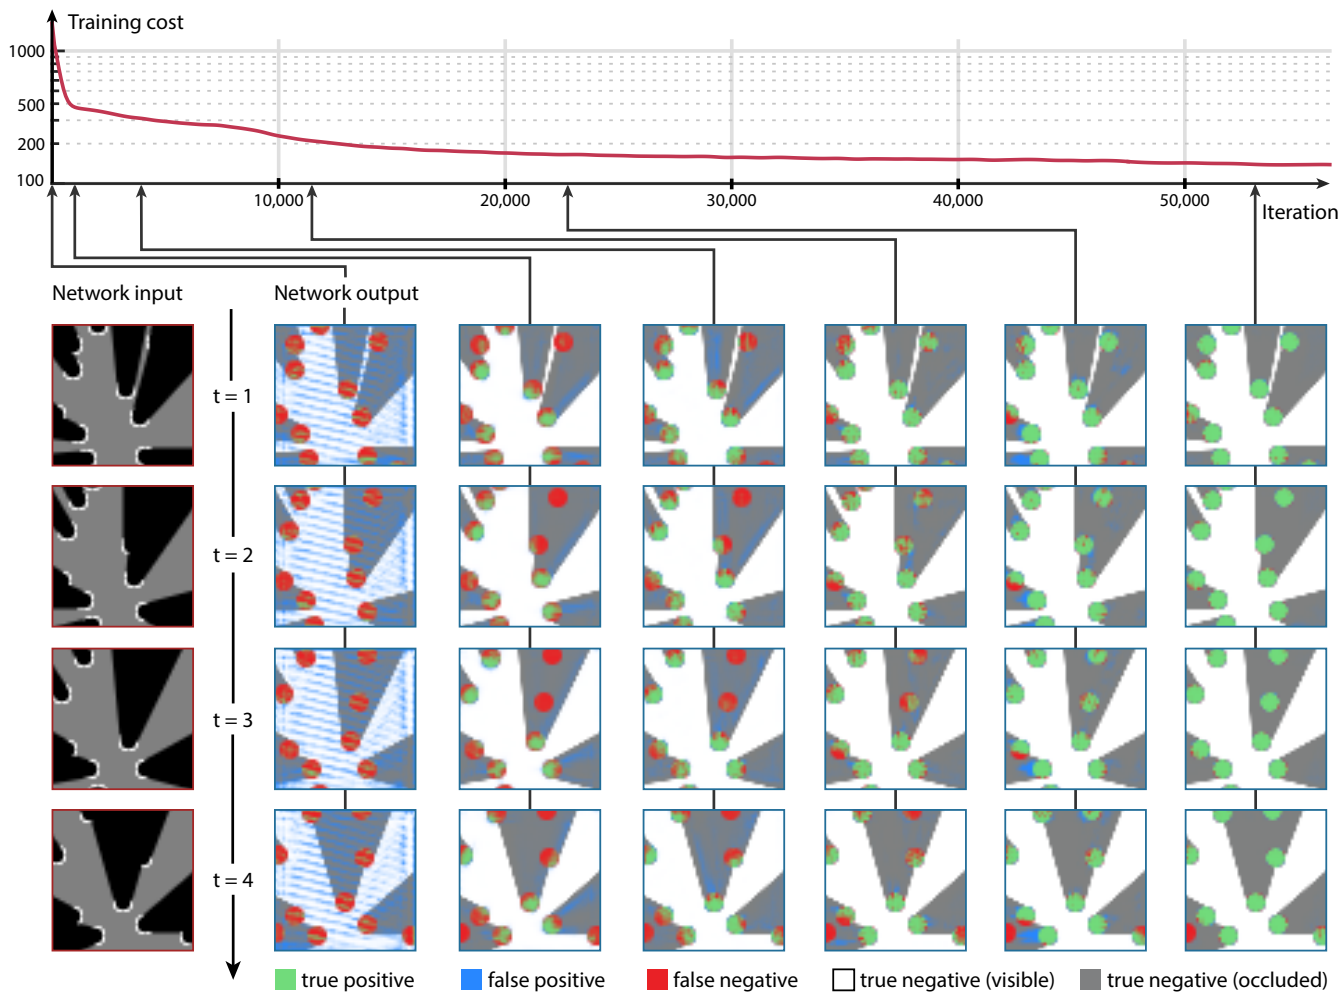


Figure 5: The unsupervised training progress. From left to right, the network first learns to complete objects and then it learns to track the objects through occlusions. Some situations can not be predicted correctly such as the object in occlusion on the fourth frame not seen before. Note, the ground-truth for comparison was used only for evaluation, but not during the training.

motion patterns with different patterns of unit activations and to track this pattern from frame to frame. None of this behaviour is hard-coded and it is a pure result of the network adaptation to the task.

Additionally we performed experiments changing the shape of objects from circles to squares and also simulating sensor noise on 1% of all pixel observations³. In all cases the network correctly learned and predicted the world state, suggesting the approach and the learning procedure is able to perform well in a variety of scenarios. We expect more complex cases will require networks having more intermediate layers or using different units such as LSTM (Hochreiter and Schmidhuber 1997), however.

Related Works

Our work concerns modelling partially-observable stochastic processes with a particular focus on the application

³See the attached video.

of tracking objects. This problem is commonly solved by Bayesian filtering which gives rise to tracking methods for a variety of domains (Yilmaz, Javed, and Shah 2006). Commonly, in these approaches the state representation is designed by hand and the prediction and correction operations are made tractable under a number of assumptions on model distributions or via sampling-based methods. The Kalman filter (Kalman 1960), for example, assumes a multivariate normal distribution to represent the belief over the latent state leading to the well-known prediction/update equations but limiting its expressive power. In contrast, the Particle filter (Thrun, Burgard, and Fox 2005) foregoes any assumptions on belief distributions and employs a sample-based approach. In addition, where multiple objects are to be tracked, correct data association is crucial for the tracking process to succeed.

In contrast to using such hand-designed pipelines we provide a novel end-to-end trainable solution allowing the robot to automatically learn an appropriate belief state rep-

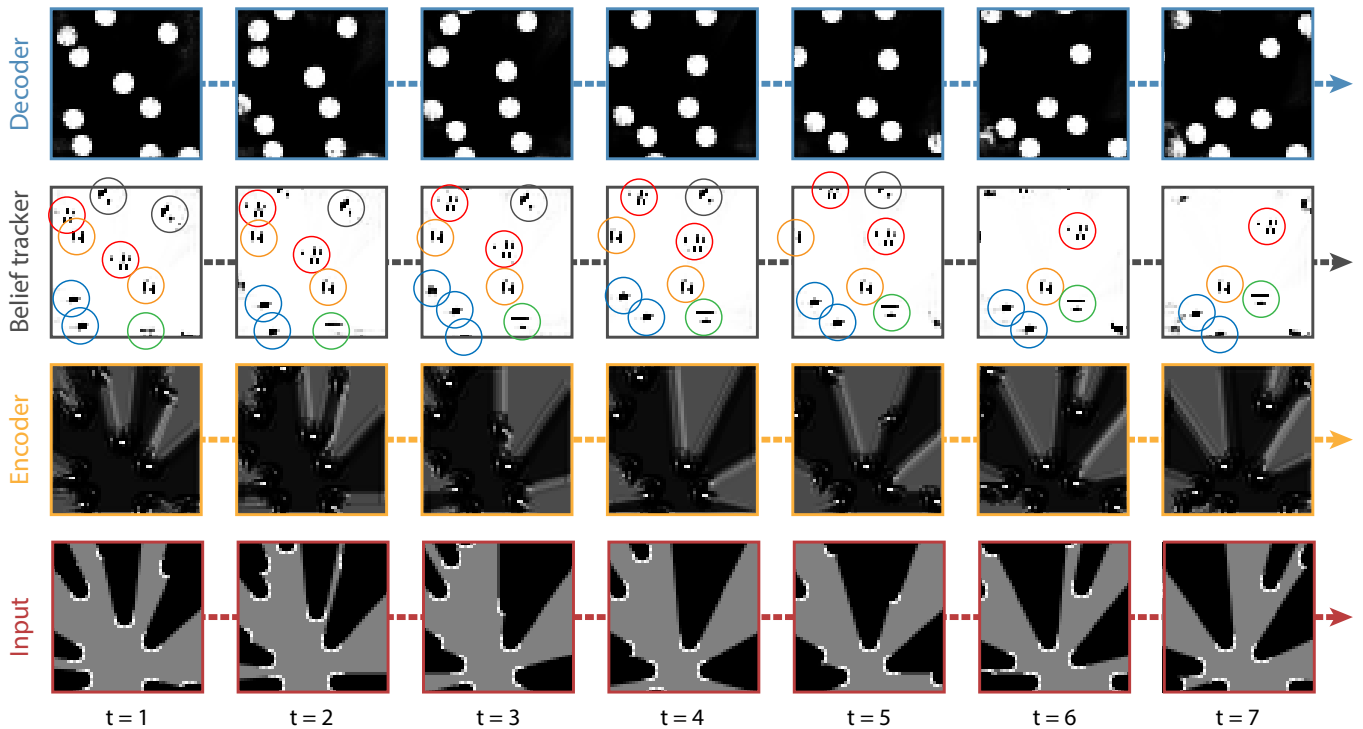


Figure 6: Examples of the filter activations from different parts of the network (one filter per layer). The *Encoder* adapts to produce a spike at the position of directly visible objects. Interesting is the structure of the *belief* \mathcal{B}_t . As highlighted by circles in 2nd row, it adapted to assign different activation patterns to objects having different motion patterns (best viewed in colour).

resentation as well as the corresponding predict and update operations for an environment containing multiple objects with different appearance and potentially very complex behaviour. While we are not the first to leverage the expressive power of neural networks in the context of tracking, prior art in this domain primarily concerned only the detection part of the tracking pipeline (Fan et al. 2010), (Jänen et al. 2010).

A full neural network pipeline requires the ability to learn and simulate the dynamics of the underlying system state in the absence of measurements such as when the tracked object becomes occluded. Modelling complex dynamic scenes however requires modelling distributions of exponentially-large state representations such as real-valued vectors having thousands of elements. A pivotal work to model high-dimensional sequences was based on using Temporal Restricted Boltzman Machines (Sutskever and Hinton 2007; Sutskever, Hinton, and Taylor 2009; Mittelman et al. 2014). This generative model allows modelling the joint distribution of $P(y, x)$. The downside of using an RBM, however, is the need of sampling, making the inference and learning computationally expensive. In our work, we assume an underlying generative model but, in the spirit of Conditional Random Fields (Lafferty, McCallum, and Pereira 2001) we directly model only $P(y|x)$. Instead of RBM our architecture features Feed-forward Recurrent Neural Networks (Medsker and Jain 2001; Graves 2013) making the inference and weight gradients computation exact using a standard back-propagation learning procedure (Rumelhart, Hinton,

and Williams 1988). Moreover, unlike undirected models, a feed-forward network allows the freedom to straightforwardly apply a wide variety of network architectures such as fully-connected layers and LSTM (Hochreiter and Schmidhuber 1997) to design the network. Our used architecture is similar to recent Encoder-Recurrent-Decoder (Fragkiadaki et al. 2015) and similarly to (Shi et al. 2015) which features convolutions for spatial processing.

Conclusions

In this work we presented *Deep Tracking*, an end-to-end approach using recurrent neural networks to map directly from raw sensor data to an interpretable yet hidden sensor space, and employ it to predict the unoccluded state of the entire scene in a simulated 2D sensing application. The method avoids any hand-crafting of plant or sensor models and instead learns the corresponding models directly from raw, occluded sensor data. The approach was demonstrated on a synthetic dataset where it achieved highly faithful reconstructions of the underlying world model.

As future work our aim is to evaluate *Deep Tracking* on real data gathered by robots in a variety of situations such as pedestrianised areas or in the context of autonomous driving in the presence of other traffic participants. In both situations knowledge of the likely unoccluded scene is a pivotal requirement for robust robot decision making. We further intend to extend our approach to different modalities such as 3D point-cloud data and depth cameras.

References

- Bengio, Y. 2009. Learning deep architectures for ai. *Foundations and trends in Machine Learning* 2(1):1–127.
- Bergman, N. 1999. Recursive bayesian estimation. *Department of Electrical Engineering, Linköping University, Linköping Studies in Science and Technology. Doctoral dissertation* 579.
- Chen, Z. 2003. Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics* 182(1):1–69.
- Fan, J.; Xu, W.; Wu, Y.; and Gong, Y. 2010. Human tracking using convolutional neural networks. *Neural Networks, IEEE Transactions on* 21(10):1610–1623.
- Fragkiadaki, K.; Levine, S.; Felsen, P.; and Malik, J. 2015. Recurrent network models for human dynamics. ICCV.
- Graves, A. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Jänen, U.; Paul, C.; Wittke, M.; and Hähner, J. 2010. Multi-object tracking using feed-forward neural networks. In *Soft Computing and Pattern Recognition (SoCPaR), 2010 International Conference of*, 176–181. IEEE.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering* 82(1):35–45.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Lafferty, J.; McCallum, A.; and Pereira, F. C. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data.
- Medsker, L., and Jain, L. 2001. Recurrent neural networks. *Design and Applications*.
- Mittelman, R.; Kuipers, B.; Savarese, S.; and Lee, H. 2014. Structured recurrent temporal restricted boltzmann machines. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, 1647–1655.
- Rabiner, L. R. 1989. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77(2):257–286.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1985. Learning internal representations by error propagation. Technical report, DTIC Document.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1988. Learning representations by back-propagating errors. *Cognitive modeling* 5:3.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-K.; and Woo, W.-c. 2015. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *arXiv preprint arXiv:1506.04214*.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15(1):1929–1958.
- Sutskever, I., and Hinton, G. E. 2007. Learning multilevel distributed representations for high-dimensional sequences. In *International Conference on Artificial Intelligence and Statistics*, 548–555.
- Sutskever, I.; Hinton, G. E.; and Taylor, G. W. 2009. The recurrent temporal restricted boltzmann machine. In *Advances in Neural Information Processing Systems*, 1601–1608.
- Szegedy, C.; Toshev, A.; and Erhan, D. 2013. Deep neural networks for object detection. In *Advances in Neural Information Processing Systems*, 2553–2561.
- Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic robotics*. MIT press.
- Yilmaz, A.; Javed, O.; and Shah, M. 2006. Object tracking: A survey. *Acm computing surveys (CSUR)* 38(4):13.