

# What Makes a Place? Building Bespoke Place Dependent Object Detectors for Robotics

Jeffrey Hawke, Alex Bewley, Ingmar Posner

**Abstract**—This paper is about enabling robots to improve their perceptual performance through repeated use in their operating environment, creating local expert detectors fitted to the places through which a robot moves. We leverage the concept of ‘experiences’ in visual perception for robotics, accounting for bias in the data a robot sees by fitting object detector models to a particular ‘place’. The key question we seek to answer in this paper is simply: how do we define a place? We build bespoke pedestrian detector models for autonomous driving, highlighting the necessary trade off between generalisation and model capacity as we vary the extent of the ‘place’ we fit to. We demonstrate a sizeable performance gain over a current state-of-the-art detector when using computationally lightweight bespoke place-fitted detector models.

## I. INTRODUCTION

Place matters in robotics. Mobile robots frequently traverse the same operating environment over and over, such as an autonomous car performing the same weekday commute, or an autonomous forklift moving goods around the same warehouse. In object detection for robotics, we care about achieving the best possible performance on the data the robot observes in its operating environment. We propose that relaxing the need for detector generalisation in favour of fitting to the operating environment enables better perception performance using a computationally simpler detector.

Prior work in robotics has demonstrated that it is possible to boost the performance of many vision-based systems by using place dependent models. Examples of these place dependent systems include object detectors [1], [2], terrain assessment [3], and visual localisation systems [4], [5]. A key question in considering all these environment specific models is simply: what defines an appropriate ‘place’ for a model if we want to improve a perception system’s performance?

We focus on environment specific models for object detection, which we expect to see vary in performance with appearance change due to viewpoint, lighting, weather, season, and structural change in the environment.

Some of this performance variability will be spatially dependent, with the general scene structure and image texture remaining the same between images taken in the same location. In addition to this spatial variability, there are many temporal sources of environmental variation which will affect image appearance (and thus detector performance). This temporal variation includes appearance change from seasonal change, lighting change, and other dynamic objects operating in the same environment. In this paper we look at

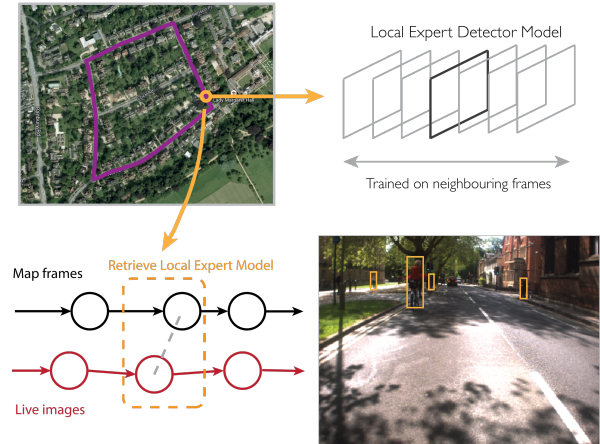


Fig. 1: We can use the knowledge of where our robot is to improve object detection performance using a local expert model fitted to that specific place. Using lightweight detector models with limited model capacity we can boost performance beyond more computationally complex state-of-the-art detectors. We achieve this by limiting the variability of the negative (background) class our model is trained on, mitigating the model capacity limits of lightweight detectors.

the influence of spatial variability on detector performance, and present a framework for improving performance through experience-based methods.

This paper is structured in two parts. We expect image data taken from the same location and time to have similar structure, enabling us to construct bespoke detectors on sequential images. Section III considers a scenario using this spatiotemporal definition of a place. We use an ensemble method: building a suite of local expert models for swathes of a robot’s operating route using a nearest neighbour approach. These models are retrieved at run time using localisation. Section IV extends this bespoke model methodology to use the nearest frames based on image appearance rather than spatial location. We approach all these experiments using pedestrians as the target object class as we have an interest in autonomous driving, though we expect the principles to hold across other object classes.

Our results demonstrate that overcoming limited model capacity (the ability of a classifier to separate the data classes) is the key benefit of using local expert detectors, enabling a lightweight detector to achieve a substantial performance increase over a more complex state-of-the-art pedestrian detector. We are able to do this by limiting the variability of the background (negative) class to local regions of the robot’s operational envelope.

The authors are with the Oxford Robotics Institute, Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, United Kingdom. Email: [jhwake, bewley, ingmar]@robots.ox.ac.uk

## II. LITERATURE REVIEW

The approach we present in this paper can be applied to any object detection problem, but we focussed on pedestrian detection due to our interest in autonomous driving. Pedestrian detection has an extensive body of literature. Much of this work can be summarised as the development of progressively better engineered features [6], [7], different classifiers [8], deformable parts models [9], or adding additional data or context [10], [11], [12], [13].

Recent developments in deep learning in the broader computer vision community have lead to convolutional neural networks being applied to the problem [14], [15], [16]. These approaches are now achieving state-of-the-art performance but tend to be slower at run-time, which makes these ill-suited for use in resource constrained robotic applications.

Zhang *et al.* [17] presented a comprehensive failure analysis of state-of-the-art pedestrian detectors on the Caltech pedestrian dataset [6]. They hypothesised that many of the false negatives (i.e., cases where the detector fails to detect a pedestrian even though the pedestrian is large enough in the image) are due to dataset bias in training, which can be solved by augmenting the training dataset. Their analysis of common sources of false positives highlighted two categories of error: localisation errors where the false detection overlaps with a ground truth bounding box, and background errors where there is zero overlap with ground truth. These background errors were the most common error type, spuriously detecting parts of the background as pedestrians.

In previous work [1] we focussed on addressing these background errors in object detection by applying the concept of robot ‘experiences’ from localisation and navigation [18], [19], [20], [21] to object detection. In an experience-based approach we accept that the world has many different forms of appearance and fit local experts to suit each experience. For perception tasks such as object detection, this means training our detector model to suit the class of data (e.g., season, spatial location) the robot currently experiences. We exploited the fact that mobile robots tend to be constrained to a particular operating environment to boost test performance by fitting models to data captured at a similar time of the year for the same traversed route. Notably, this experience-based detector outperformed a global model which has been trained on data across multiple experiences, highlighting the importance of defining the environment.

Geospatial bias in imagery has been highlighted by Doersch *et al.* [22], who showed that it is possible to automatically find distinctive visual elements for a certain geo-spatial area, comparing data between major cities (for example, Paris and London) and data between districts of a city (Paris). We also know that dataset bias [23] is an influence on the performance of our object detectors, with the training data used being a systematic factor influencing performance across different test datasets [24]. This data bias, however, provides a potential tool to be exploited.

Malisiewicz *et al.* addressed bias in intra-class variability (such as different types of buses) in general purpose image recognition with Exemplar SVMs [25], an ensemble approach which uses a nearest neighbour method to train a

unique model for each positive training datum, fitted to a single positive sample and many thousands of negatives.

Various approaches have been taken to leverage data bias to form scene specific classifiers, largely motivated by visual surveillance tasks. Hattori *et al.* took the concept of scene specific classifiers to one extreme, training a unique detector model for a specific image location in a static scene using synthetic data [26]. Wang *et al.* used a transfer learning framework to adjust a generic pedestrian detector to a scene specific detector [27]. They used visual cues to generate scene specific data in a self-supervised manner, weighting training data samples based on similarity to the observed test data using an affinity graph.

In the domain of mobile robotics, Suzuo *et al.* divided a driven route into a fixed number of scenes using visual bag-of-words, then trained models using negative data gathered from each scene [28]. This has some similarity to our approach here; we build on the ideas from [1] and [28] to understand what influences an appropriate resolution for local expert detectors. Bewley and Upcroft presented similar evidence for the need to adapt object detectors to their operating environment in [29], augmenting a generalised object detector by validating each detection against a background model trained on environment specific data. This approach can also be viewed as exploiting place as context, with some similarities to methods which use context within an image to boost detector performance, e.g. [30].

The idea of experience-based local expert models has since been applied to other perceptual tasks in robotics. Berczi *et al.* [3] created bespoke terrain assessment classifier models for finite patches of a route driven with visual teach and repeat. These models detect when the path in front of a mobile robot has changed from the initial teach pass and previous repeat runs. They demonstrated a marked performance improvement over place-independent methods, comparing their bespoke place dependent models to a place independent learned classifier and a traditional heuristic based method.

Similarly, McManus *et al.* [5] and Linegar *et al.* [4] both proposed using distinct scene specific patches to weakly localise across large appearance change.

These methods all have similar motivation: solving a simpler problem rather than a global one. However, why do these work so well? Why can’t we simply use a generic model? In this paper we present evidence for *why* place dependent models work well for solving these sorts of perception problems in robotics, and offer insight into how we should approach the construction of local expert models.

## III. FITTING TO A KNOWN LOCATION AND TIME

We evaluate how to build bespoke place dependent models by firstly adopting a fine-grained, strict spatiotemporal definition of a place, where we construct a place dependent model for every single frame in a known route. We then immediately repeat the route, localise each image to our mapping run, then retrieve and apply the detector model associated with the localised map frame.

This is the best possible scenario we could hope for when fitting models to similar data. By using a strict spatiotemporal prior we are able to control most sources of variation in scene

content (as the images are from the same spatial location) and appearance (as the images are from effectively the same time). In addition, we use a very fine grained approach, associating a model with every frame, rather than key frames.

In practice, this will form the upper bound in performance for place dependent methods with this classifier and feature type, though we expect to be able to get close to this using experience-based visual localisation [18] with models associated with mapping key frames.

### A. Experimental Approach

We select an urban driving dataset consisting of two consecutive loops of the same route in Oxford [31]. Each loop is 1km in length and we fully annotated these with pedestrian labels. For each image (or model index frame), we fit a model to the closest  $N$  images in time (including the frame itself), illustrated in Figure 1. At run time, we localise each image to the alternate lap and retrieve that model, using the models built on the second lap during the first and vice versa.

### B. Object Detector and Training

Our object detector is a lightweight sliding window implementation using Aggregate Channel Features (ACF) on LUV images [6], with a linear Support Vector Machine (SVM) as the primary classifier. This is very similar to the original work in [6], except we use our own implementation of ACF with LibLinear [32], rather than the original boosted decision trees. This is a lightweight general purpose object detector, able to run at frame rate (20Hz) using a reasonable selection of scales. The ACF feature descriptor consists of ten channels based on colour and gradient, making it suitable for detection tasks such as traffic lights [33] and road signs [34] in addition to the original purpose of pedestrian detection.

We fit detector models to the nearest  $N$  frames in time by performing hard negative mining (HNM) on each image using ground truth labels. This process runs a candidate detector model on the training images, adding the highest scoring false positives to the training data, retraining and iterating until convergence (or a maximum of 20 iterations). The model is initially trained using the INRIA pedestrian dataset [35] for positive samples and an initial seed negative set (10 random patches from each INRIA negative image). We heavily penalise misclassification of training data by setting the SVM optimization parameter to a very large number ( $C = 100$ ), forcing the model to strongly fit to the training data presented. The training data is weighted based on the number of samples to ensure class balance.

### C. Results

To assess the merits of a precise spatiotemporal definition of a place, we compare our place-fitted detector to a generic model for our pedestrian detector trained only on a canonical dataset. This generic model is trained on the same initial training data as the place dependent models, but with HNM performed on the INRIA negative training data rather than data from the robot’s environment. We refer to this baseline trained on generic data as the SVM+ACF model.

Detector Model	Average Precision	Log-Average Miss Rate	Maximum F1 Score
Baseline: SVM+ACF	0.483	0.485	0.473
Baseline: FPDW [6]	0.481	0.440	0.491
Baseline: MSCNN [16]	0.588	0.371	0.579
P.D. Model: N=1	0.622	0.354	0.592
P.D. Model: N=10	<b>0.689</b>	0.292	0.651
P.D. Model: N=100	<b>0.689</b>	<b>0.288</b>	<b>0.672</b>
P.D. Model: N=1000	0.450	0.564	0.521
P.D. Model: N $\sim$ 4100 (full lap)	0.305	0.712	0.405

TABLE I: Performance metrics comparing our place dependent (P.D.) detector models to a number of generic baselines, where  $N$  is the length of image sequence used to train a model. Higher is better for Average Precision and F1 Scores, whereas lower is better for Log-Average Miss Rate.

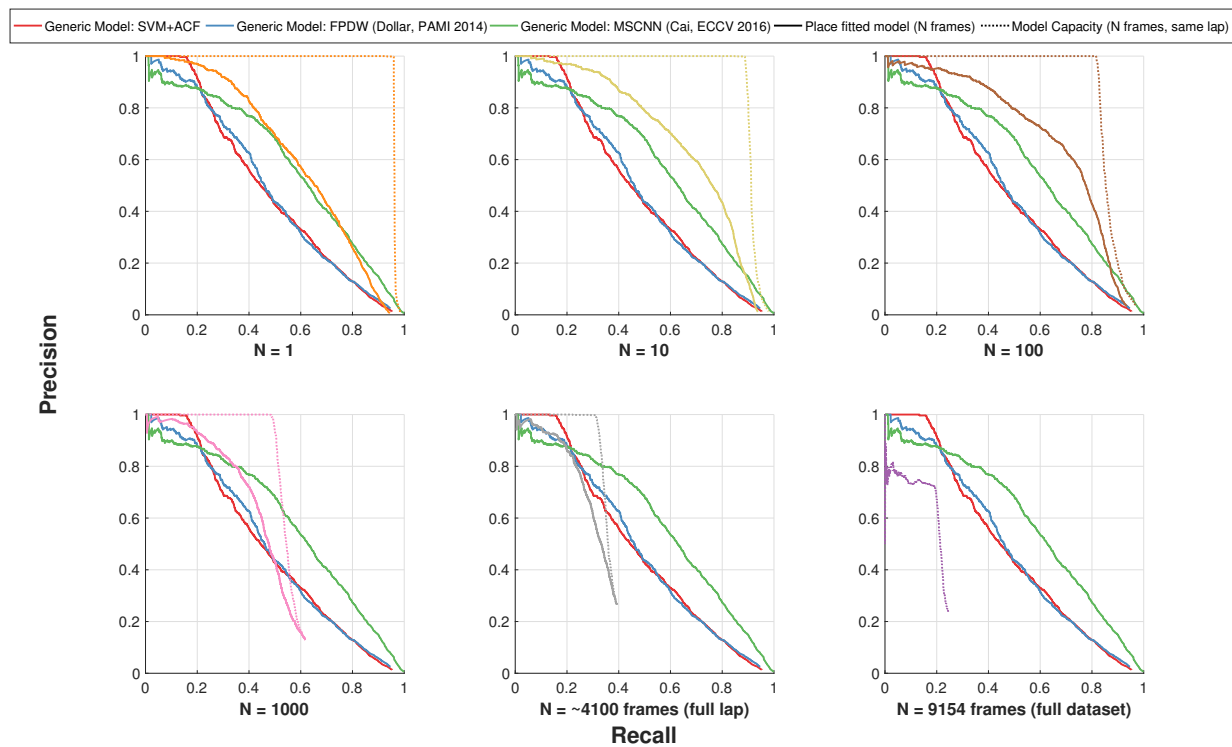
Table I outlines the aggregate performance metrics on our test dataset. We use three common metrics for comparing detector models: average precision, log-average miss rate [6], and maximum F1 score.

1) *Benchmarks*: Results from two publicly available detectors are included to provide a baseline from which to assess potential performance gains. These are FPDW [6], and MSCNN [16]. We use models provided by the respective authors, trained on the Caltech pedestrian dataset [6]. We acknowledge that there will be some bias in performance due to the training dataset, as INRIA has been shown to perform well across a diverse range of benchmark datasets [17]. Furthermore, as INRIA is a cropped dataset, it is only applicable to sliding window approaches making it unsuitable for training detectors like MSCNN.

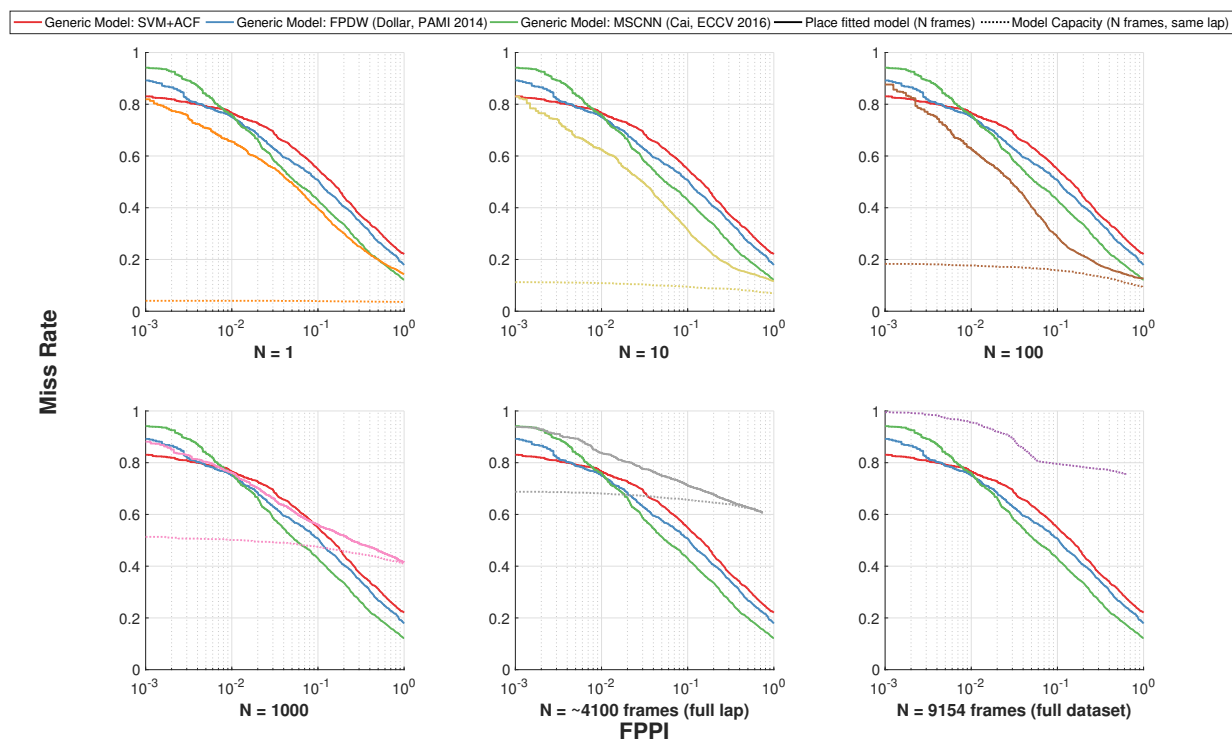
FPDW is very similar to our SVM+ACF generic model, using ACF features but with a boosted decision tree classifier (AdaBoost) rather than an SVM. These two detectors are largely equivalent in performance and speed, able to run at frame rate (20Hz) for a reasonable range of scales on our Bumblebee2 camera using only moderate computational resources. MSCNN, a multi-scale convolutional neural network detector, is currently one of the best performing object detectors, but requires significant GPU based computational power and is only able to run at approximately 2Hz on a top end NVIDIA GTX graphics card. The authors state a run-time performance of 0.4s per frame on KITTI data (1242x375). Many robotic applications require a lower computational budget or higher frame rate.

2) *Detector Performance*: The local expert detector outperforms all the baseline generic models (including the current state-of-the-art), with a twenty percentage point improvement over the reference model of the same type (SVM+ACF). Figure 2 shows the characteristic performance curves for these models as the training sequence length varies. The dotted lines on each plot give us an appreciation for the model capacity, applying the model trained on each image to itself. Model capacity refers to be ability of a classifier to fit (separate) the data correctly without error, which will be determined by the classifier type, feature type, and the classes of data.

3) *Training Swathe Size*: Firstly, we note that using an extremely tightly fitted model (spatially, where the model is fitted to the single closest frame) offers a good improvement,



(a) Precision - Recall (higher is better)



(b) Miss Rate - False Positives Per Image (lower is better)

Fig. 2: Performance of an ideal case place dependent detector compared to detectors using generic models. The place fitted detector uses models fitted on a frame-by-frame basis to a sequence of  $N$  images centred on the closest matching frame from a second lap (the solid coloured lines). The dotted line on each plot represents the model capacity of the place fitted models, i.e., the performance when applying models to the same index frame they were trained on. This highlights the necessary trade-off between the size of negative class we fit to and the point at which we need the model to generalise.

but this is outperformed by models which are fitted to a slightly longer swathe ( $N = 10 - 100$ ). There are a number of factors which will affect this.

Firstly, longer training swathes will inevitably add more training data, which will generally help performance. In addition, these longer swathes will also help to account for the differences in pose between the image and the model index frame. This pose error is small (as per Figure 3), but will limit the efficacy of models fitted to a single frame due to variation in perspective and scene content.




Fig. 3: Variation in pose between localised images on subsequent laps (a median of 0.4m for our dataset) will limit the efficacy of models fitted to a single frame.

On the other hand, using a large sequence of frames comes with a penalty. As we are forcing the model to fit the training data firmly (using a high value for the SVM kernel parameter), we begin to see a drop in performance due to limited model capacity once the negative set is sufficiently large such that we lose class separability.

We start to see the impact of finite model capacity on performance at a swathe size of between 100-1000 frames. At 1000 frames we have lost class separability, in which the linear SVM classifier is unable to separate the positive and negative training data with no errors (illustrated conceptually in Figure 4). As we weight the training data for class balance, this lack of class separability results in the hard negative mining reaching the maximum number of cycles rather than converging.

This performance drop due to model capacity is even clearer in the *global model* case, where we fit the model to the entire dataset. In this case, we only have a model capacity plot as we train the model on every image ( $N = 9154$ ). Even though the global model is trained on the exact test data, it is unable to separate the positive and negative class, requiring the model to generalise.

4) *Concluding remarks:* With lightweight pedestrian detectors performance is boosted by using local expert models which have a very specific place of operation. The extent of a place — measured by a fixed swathe size — should contain multiple frames to avoid sensitivity to the pose difference between training and testing data. Conversely, a narrow definition of a place with minimal variation in the background class is desirable to enhance class separability with finite




Fig. 4: This illustration outlines a simple example of where using a simpler (faster) classifier cannot separate the positive and negative class perfectly due to finite model capacity, indicated by the red dashed decision boundary. Reducing the background class boosts performance by training models for a single environment (e.g., the orange or green images), giving full class separability with a collection of models.

model capacity. In our experiments, models trained with sequences of 10 – 100 frames showed the most significant improvement.

#### IV. LINKING APPEARANCE TO PLACE DEPENDENT MODELS

Spatial location is important, but we also expect to see image appearance change over time: seasons, lighting, and other moving objects. From localisation literature we know that image appearance change is significant, and has a dramatic effect on the performance of perception systems [18].

As we focus on boosting performance by reducing the extent of the background class our model fits to, it logically holds that we should build models on images with similar appearance. In many cases this is influenced by spatial proximity, as higher similarity is correlated with spatial proximity (shown in Figure 7). The previous section (Section III) addresses this. However spatial correlation is not guaranteed, as appearance metrics vary with changes in scene content and perspective. A key example of this is a sequence of images taken driving around a corner: these will be dissimilar by appearance metrics but are spatially close (Figure 6).

##### A. Experimental Approach

We adopt the same experimental scenario as in Section III, where we have two consecutive laps of an urban driving route with ground truth localisation and pedestrian annotations. For every image, we localise to the closest frame in the alternate lap, retrieving the detector model associated with the map frame and apply it. The object detector setup and training is as per Section III-B.

However, to build the detector models we now use the  $N$  nearest frames based on image similarity metrics, rather than  $N$  nearest temporal frames. We use two similarity measures: the  $l^2$ -norm between two GIST global image descriptors [36], and mutual information [37] between two intensity images.

Our results in Section III indicate that our detector model was able to fit well for a place definition of between 1 – 100 frames, with 1000 frame detector models suffering in performance due to model capacity (Figure 5). This result agrees with the data we see in Figure 7, which indicates that the correlation between image appearance and distance




Fig. 5: Aggregate performance metrics for place dependent detector models as a function of  $N$ , the number of frames in the nearest neighbour training swathe. The dotted line represents the model capacity. Note that our generic model (SVM+ACF) and FPDW score almost identically in average precision.

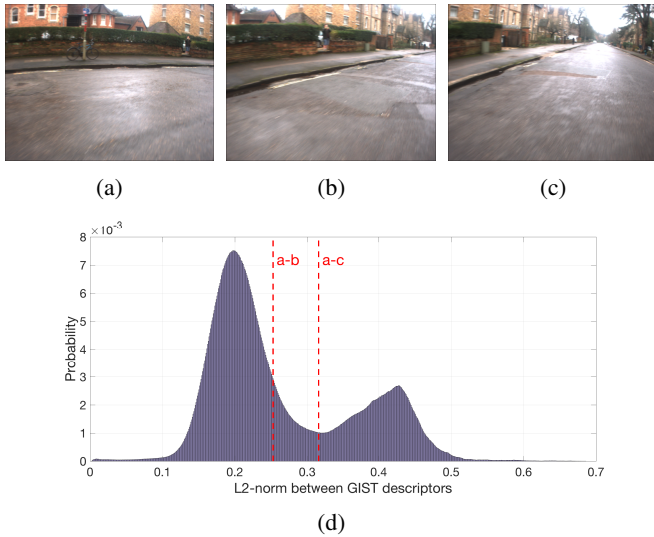



Fig. 6: Images 6a through 6c are taken one second apart (approximately 3 metres), but differ substantially when comparing the GIST descriptors. This is due to the rapidly changing perspective rather than any substantive change in scene content. The histogram in 6d shows the similarity distribution for all image pairs in this route, with the red lines showing the similarity between image pairs 6a-6b and 6a-6c. These fall in the 61st and 71st percentile, highlighting the difference in similarity between spatially close images.

no longer applies beyond 100 frames at the average vehicle speed. We focussed our experiments on models built using the same number of frames as before ( $N = 1 - 100$ ), starting from a single frame and increasing by a decade at each step.


### B. Results

To assess this approach, we use the same experimental methodology as in Section III, using localisation to retrieve the model associated with the map frame closest to the current image.

In Figure 9 and Table II we see a performance boost over the reference baseline detectors. Using the nearest images based on GIST yields more or less equivalent performance.



(a)  $l^2$ -norm between GIST



(b) Mutual Information

Fig. 7: Image similarity metrics correlate with the distance travelled along the robot's trajectory. The mean similarity is shown by the solid line, and the filled polygons represent the  $2\sigma$  and maximum extent. The red dashed lines indicate the distance travelled within  $N = 1, 10, 100$  frames during this test route at the vehicle's average speed of  $4.3 \text{ ms}^{-1}$ .




Fig. 8: These histograms highlight the spatial distance in models’ training data when using different image similarity metrics. For models built using the nearest  $N$  images (spatially), a 100 frame model will on average have been trained on data up to 22 metres from the index frame (as in Figure 7). In contrast, when using the nearest 100 frames by either image similarity metric this can be up to 400 metres.

Model	Average Precision	Log-Average Miss Rate	Maximum F1 Score
Baseline: SVM+ACF	0.483	0.485	0.473
Baseline: FPDW [6]	0.481	0.440	0.491
Baseline: MSCNN [16]	0.588	0.371	0.579
P.D. Spatial: $N=1$	0.622	0.354	0.592
P.D. Spatial: $N=10$	<b>0.689</b>	0.292	0.651
P.D. Spatial: $N=100$	<b>0.689</b>	<b>0.288</b>	<b>0.672</b>
P.D. $l^2$ GIST: $N=10$	0.678	0.303	0.652
P.D. $l^2$ GIST: $N=100$	<b>0.689</b>	0.300	0.666
P.D. Mut. Inf.: $N=10$	0.594	0.382	0.584
P.D. Mut. Inf.: $N=100$	0.656	0.347	0.640

TABLE II: Aggregate performance metrics comparing our bespoke place dependent (P.D.) detector models to a number of generic baselines. A higher number is better for Average Precision and F1 Scores, lower is better for Log-Average Miss Rate. The  $N = 1$  case is the same for all methods as this model is trained on the closest localised frame.

This logically holds: for our dataset we expect the temporally adjacent frames to be very similar in appearance, with the exception of perspective change when turning corners. Using mutual information is less beneficial: this still helps boost performance, but is not as informative as grouping images based on time or GIST. In Figure 7b we note that the correlation between mutual information and distance plateaus at 5 metres, unlike with GIST. The mutual information models are also trained on data which is much further from the source frame than GIST (Figure 8). We expect that mutual information between each colour channel rather than intensity would be better, though this is unlikely to improve significantly over the temporal sequence approach.

It is likely that we could use image appearance to fit better models for larger groups of data, but we are limited by model capacity before this effect is visible. With a higher capacity model (e.g., deep learning approaches or different kernels) we may be able to exploit image appearance to a greater degree where we lack robust localisation.

## V. CONCLUSIONS

In robotics we have the opportunity to exploit knowledge of the world around our robot to improve performance of vision systems. In this paper we demonstrate how prior

knowledge about what a robot will see (by repeated operation along the same route) can be used to boost performance beyond what a more complex state-of-the-art object detector can achieve. However, defining the place of operation is not trivial for mobile robots. We present an assessment of approaches to building bespoke place dependent detector models, investigating which factors influence performance.

Our results indicate that we can heavily leverage a strong spatiotemporal prior, fitting models to a narrow window of operation. Relaxing the need for generality aids performance as we can ensure our models fit a smaller negative background class. This approach avoids the trade off between model capacity and generalisation. When using appearance to build models, we are able to achieve effectively the same performance on groups of similar images.

As model capacity is a major influence on performance of these local models, we anticipate that using higher capacity classifiers (such as deep learning methods or more complex SVM kernels) will enable fitting to larger operating environments. Computational requirements will dictate the choice of classifier used in these local expert detectors. We also only considered expert models fitted to a fixed number of images. It is likely that larger visually uniform sections of a route could be adequately fitted with a single model, whereas more varied routes will need a more specialised model fitted to a shorter section.

From a practical view, we see significant potential for this approach to be coupled to experience-based visual localisation systems [18], [19], [20], [21], using their ability to localise to many parallel ‘experiences’ of the same location. These bespoke detector models are sensitive to the data they are applied to: if the scene content changes substantially from the training environment, it is likely that the detector will generate false positive errors due to the change in background class. This brittleness could be addressed by using models constructed from appearance rather than spatiotemporal sequences, or by using the ability of experience-based localisation to localise across different appearances of the same place, with detector models associated with each experience of the same spatial location.

## ACKNOWLEDGMENT

The authors would like to acknowledge the support from the EPSRC Programme Grant DFR01420 as well as the Advanced Research Computing at the University of Oxford.

## REFERENCES

- [1] J. Hawke, C. Gurau, C. H. Tong, and I. Posner, “Wrong Today, Right Tomorrow: Experience-Based Classification for Robot Perception,” in *Field and Service Robotics (FSR)*, Toronto, ON, Canada, June 2015.
- [2] A. Bewley and B. Upcroft, “From imagenet to mining: Adapting visual object detection with minimal supervision,” in *Field and Service Robotics*. Springer, 2016, pp. 501–514.
- [3] L.-P. Berczi and T. D. Barfoot, “It’s like déjà vu all over again: Learning place-dependent terrain assessment for visual teach and repeat,” in *Intelligent Robots and Systems*. IEEE, 2016.
- [4] C. Linegar and P. Newman, “Made to Measure: Bespoke Landmarks for 24-Hour, All-Weather Localisation with a Camera,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016.
- [5] C. McManus, B. Upcroft, and P. Newman, “Scene Signatures: Localised and Point-less Features for Localisation,” in *Proceedings of Robotics: Science and Systems*, Berkeley, USA, July 2014.

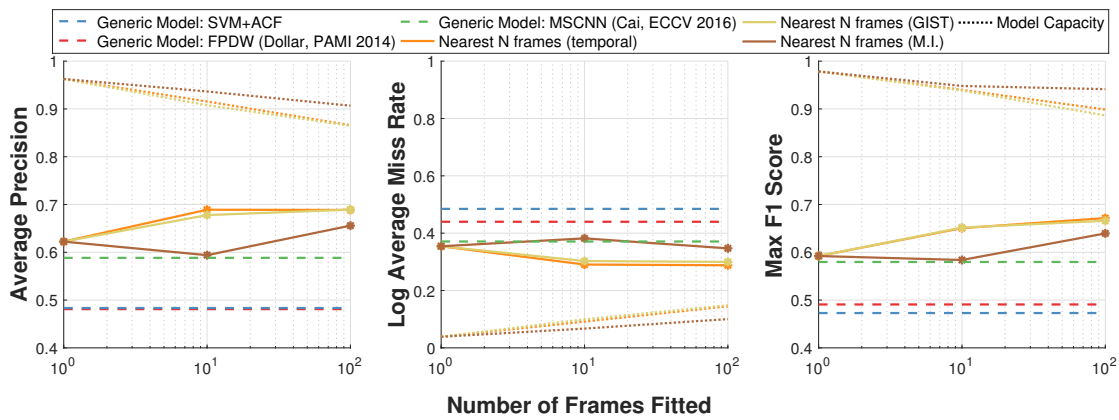


Fig. 9: Aggregate performance metrics for models fitted to the nearest  $N$  frames using image similarity and sequences. The spatiotemporal and GIST based models perform similarly, surpassing the baseline generic models. Two baseline models (SVM+ACF, FPDW) score identically in average precision, and the dotted lines indicate the SVM's limited model capacity.

- [6] P. Dollar, R. Appel, S. Belongie, and P. Perona, "Fast Feature Pyramids for Object Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 8, pp. 1532–1545, Aug. 2014.
- [7] S. Zhang, C. Bauckhage, and A. B. Cremers, "Informed Haar-Like Features Improve Pedestrian Detection," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, June 2014, pp. 947–954.
- [8] C. Wojek and B. Schiele, "A Performance Evaluation of Single and Multi-feature People Detection," in *Pattern Recognition*, ser. Lecture Notes in Computer Science, G. Rigoll, Ed. Springer Berlin Heidelberg, June 2008, no. 5096, pp. 82–91, dOI: 10.1007/978-3-540-69321-5\_9.
- [9] A. Prioletti, A. Mogelmoose, P. Grisleri, M. Trivedi, A. Broggi, and T. Moeslund, "Part-Based Pedestrian Detection and Feature-Based Tracking for Driver Assistance: Real-Time, Robust Algorithms, and Evaluation," *Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1346–1359, Sept. 2013.
- [10] M. Enzweiler, M. Hummel, D. Pfeiffer, and U. Franke, "Efficient Stixel-based object recognition," in *2012 IEEE Intelligent Vehicles Symposium (IV)*, June 2012, pp. 1066–1071.
- [11] D. Geronimo, A. D. Sappa, D. Ponsa, and A. M. Lopez, "2d–3d-based on-board pedestrian detection system," *Computer Vision and Image Understanding*, vol. 114, no. 5, pp. 583–595, May 2010.
- [12] S. Walk, N. Majer, K. Schindler, and B. Schiele, "New features and insights for pedestrian detection," in *Computer Vision and Pattern Recognition*, 2010, pp. 1030–1037.
- [13] W. Ouyang and X. Wang, "Single-Pedestrian Detection Aided by Multi-pedestrian Detection," in *Computer Vision and Pattern Recognition*, 2013, pp. 3198–3205.
- [14] J. Hosang, M. Omran, R. Benenson, and B. Schiele, "Taking a deeper look at pedestrians," in *Computer Vision and Pattern Recognition*, 2015.
- [15] F. Yang, W. Choi, and Y. Lin, "Exploit All the Layers: Fast and Accurate CNN Object Detector With Scale Dependent Pooling and Cascaded Rejection Classifiers," in *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [16] Z. Cai, Q. Fan, R. S. Feris, and N. Vasconcelos, "A unified multi-scale deep convolutional neural network for fast object detection," in *European Conference on Computer Vision*. Springer, 2016.
- [17] S. Zhang, R. Benenson, M. Omran, J. Hosang, and B. Schiele, "How Far Are We From Solving Pedestrian Detection?" in *Computer Vision and Pattern Recognition*, 2016.
- [18] W. Churchill and P. Newman, "Experience-based navigation for long-term localisation," *The International Journal of Robotics Research*, vol. 32, no. 14, pp. 1645–1661, Dec. 2013.
- [19] C. Linegar, W. Churchill, and P. Newman, "Work Smart, Not Hard: Recalling Relevant Experiences for Vast-Scale but Time-Constrained Localisation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Seattle, WA, USA, May 2015.
- [20] M. Gadd and P. Newman, "Checkout My Map: Version Control for Fleetwide Visual Localisation," in *Intelligent Robots and Systems*, 2016.
- [21] M. Paton, K. MacTavish, M. Warren, and T. D. Barfoot, "Bridging the appearance gap: Multi-experience localization for long-term visual teach and repeat," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2016, pp. 1918–1925.
- [22] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. A. Efros, "What Makes Paris Look Like Paris?" *ACM Trans. Graph.*, vol. 31, no. 4, pp. 101:1–101:9, July 2012.
- [23] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *Computer Vision and Pattern Recognition*, 2011, pp. 1521–1528.
- [24] R. Benenson, M. Omran, J. Hosang, and B. Schiele, "Ten years of pedestrian detection, what have we learned?" in *ECCV, CVRSUAD workshop*, 2014.
- [25] T. Malisiewicz, A. Gupta, and A. Efros, "Ensemble of exemplar-SVMs for object detection and beyond," in *2011 IEEE International Conference on Computer Vision (ICCV)*, Nov. 2011, pp. 89–96.
- [26] H. Hattori, V. N. Boddeti, K. Kitani, and T. Kanade, "Learning scene-specific pedestrian detectors without real data," in *Computer Vision and Pattern Recognition*, 2015, pp. 3819–3827.
- [27] X. Wang, M. Wang, and W. Li, "Scene-Specific Pedestrian Detection for Static Video Surveillance," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 2, pp. 361–374, Feb. 2014.
- [28] D. Suzuo, D. Deguchi, I. Ide, H. Murase, H. Ishida, and Y. Kojima, "Environment adaptive pedestrian detection using in-vehicle camera and GPS," in *International Conference on Computer Vision Theory and Applications*, vol. 2, Jan. 2014, pp. 354–361.
- [29] A. Bewley and B. Uppcroft, "Background appearance modeling with applications to visual object detection in an open-pit mine," *Journal of Field Robotics*, 2016.
- [30] G. Chen, Y. Ding, J. Xiao, and T. X. Han, "Detection Evolution with Multi-order Contextual Co-occurrence," in *Computer Vision and Pattern Recognition*, 2013, pp. 1798–1805.
- [31] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017. [Online]. Available: <http://dx.doi.org/10.1177/0278364916679498>
- [32] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "LIBLINEAR: A Library for Large Linear Classification," *Journal of Machine Learning Research*, vol. 9, pp. 1871–1874, 2008.
- [33] M. P. Philipsen, M. B. Jensen, A. Møgelmoose, T. B. Moeslund, and M. M. Trivedi, "Traffic light detection: a learning algorithm and evaluations on challenging dataset," in *Intelligent Transportation Systems*. IEEE, 2015.
- [34] A. Møgelmoose, D. Liu, and M. M. Trivedi, "Detection of us traffic signs," *Intelligent Transportation Systems*, vol. 16, no. 6, pp. 3116–3125, 2015.
- [35] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Computer Vision and Pattern Recognition*, vol. 1, 2005, pp. 886–893 vol. 1.
- [36] M. Douze, H. Jegou, H. Sandhawalia, L. Amsaleg, and C. Schmid, "Evaluation of GIST Descriptors for Web-scale Image Search," in *International Conference on Image and Video Retrieval*. ACM, 2009.
- [37] P. Viola and W. M. Wells III, "Alignment by maximization of mutual information," *International journal of computer vision*, vol. 24, no. 2, pp. 137–154, 1997.