

You can have your ensemble and run it too – Deep Ensembles Spread Over Time

Isak Meding*
Zenseact

Alexander Bodin*
Zenseact

Adam Tonderski
Zenseact
Lund University

Joakim Johnander
Zenseact
Linköping University

Christoffer Petersson
Zenseact
Chalmers University of Technology

Lennart Svensson
Chalmers University of Technology

Abstract

Ensembles of independently trained deep neural networks yield uncertainty estimates that rival Bayesian networks in performance. They also offer sizable improvements in terms of predictive performance over single models. However, deep ensembles are not commonly used in environments with limited computational budget – such as autonomous driving – since the complexity grows linearly with the number of ensemble members. An important observation that can be made for robotics applications, such as autonomous driving, is that data is typically sequential. For instance, when an object is to be recognized, an autonomous vehicle typically observes a sequence of images, rather than a single image. This raises the question, could the deep ensemble be spread over time?

In this work, we propose and analyze Deep Ensembles Spread Over Time (DESOT). The idea is to apply only a single ensemble member to each data point in the sequence, and fuse the predictions over a sequence of data points. We implement and experiment with DESOT for traffic sign classification, where sequences of tracked image patches are to be classified. We find that DESOT obtains the benefits of deep ensembles, in terms of predictive and uncertainty estimation performance, while avoiding the added computational cost. Moreover, DESOT is simple to implement and does not require sequences during training. Finally, we find that DESOT, like deep ensembles, outperform single models for out-of-distribution detection.

1. Introduction

In safety-critical applications, such as autonomous driving (AD), both the predictive performance and the uncertainty quantification performance of neural network models

are essential [1]–[5]. For example, the posterior probabilities produced by a perception model are required to contain all the information needed for the downstream decision-making systems to make safe and efficient driving decisions. Even for a seemingly mundane task such as traffic sign classification one can imagine the problem that could result from a model misidentifying a speed limit sign for a stop sign on the highway, while nevertheless outputting a high confidence. The general problem of overconfidence for certain types of neural networks [3] is particularly troublesome in safety-critical applications.

When it comes to modeling uncertainty, it is often divided into two types of uncertainty – aleatoric uncertainty and epistemic uncertainty [6]. Aleatoric uncertainty is due to inherent randomness in the process that generates the data. It is therefore not possible to decrease this type of uncertainty by improving the model. On the other hand, epistemic uncertainty is that which is caused by a lack of data or knowledge about the underlying process [7]. One example is to predict the outcome of a biased dice throw. It is impossible to predict what side the dice will land on for certain, no matter how accurate the model is. This constitutes irreducible aleatoric uncertainty in the prediction. However, more data can help us identify the bias of the dice and improve the model, thereby decreasing the epistemic uncertainty. A larger number of models reduces epistemic uncertainty, while a larger number of frames reduces aleatoric uncertainty.

Non-Bayesian neural networks often display great predictive performance, but struggle with generating high-quality epistemic uncertainty estimates [3], [8], [9]. Bayesian neural networks on the other hand often yield high-quality epistemic uncertainty estimates but are difficult to train [10]. One way to improve on the epistemic uncertainty estimation of regular neural networks while improv-

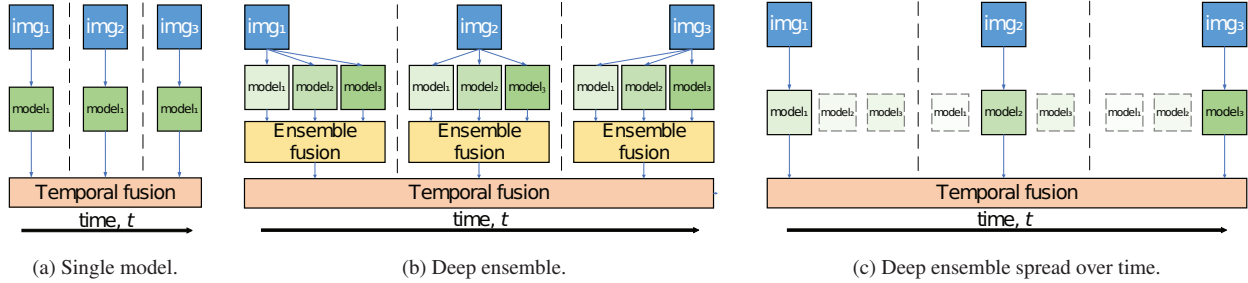


Figure 1: Visualization of the different strategies in an example with three time steps and three ensemble members. The temporal fusion blocks combine predictions produced at each time step using some combination rule. At time step t , $t \in \{1, 2, \dots, T\}$ with $T = 3$ in this example, the temporal fusion block produces a prediction based on the predictions at time steps 1 through t . The ensemble fusion block combines individual ensemble member predictions into a final ensemble prediction at each time step.

ing predictive performance is to use ensembles [8], [9]. It has long been known that ensembles of neural networks can quantify uncertainty in their predictions [11] and increase the predictive performance compared to single models [12], [13]. Lakshminarayanan *et al.* [10] show that neural network ensembles can produce uncertainty estimates that outperform Bayesian models, while also achieving high predictive performance. This allows for a practical and high-performance alternative to Bayesian methods. This type of ensemble is commonly referred to as Deep Ensembles (DEs) [9], [14], [15], and has been verified by other authors to produce high-quality uncertainty estimates [8], [9].

As previously mentioned, ensembles typically also improve predictive performance over single models. This is mathematically proven in the classification setting by Hansen and Salamon [12] under the assumption of independent classification errors between ensemble members. They show that if each model achieves an accuracy greater than 50%, adding more models results in perfect classification performance in the limit. The assumption of independent classification errors does not typically hold in the real world, but the intuition is valid. There are also other potential advantages of using ensembles. Wasay and Idreos [16] show empirically that for a set model parameter budget, using ensembles results in shorter training times and higher accuracy compared to single models.

Because of all the benefits of DEs outlined above, applications of them in safety-critical systems, such as AD, would clearly be desirable. However, an important limiting factor for AD applications is the requirement for low latency, real-time processing on resource-limited embedded hardware – a fact that is typically incompatible with the linear compute increase of DEs in the number of members.

In this paper, we leverage the fact that the sensor data in AD systems are in general sequential in nature and propose Deep Ensembles Spread Over Time (DESOTs). Instead of applying all ensemble members to the sensor data at each time step, as would be the procedure for a conven-

tional DE, *DESOT only applies one member at each time step, but different members at different consecutive time steps*. Hence, DESOT uses the same number of computations as a single model, but the same amount of memory as a DE with an equal number of members. We illustrate the DESOT method by extensively studying the task of traffic sign classification with sequences of tracked patches containing the traffic signs. We choose this task due to its importance in AD applications and its simple formulation, but the method is applicable to a wide range of settings and tasks that use sequence data. We show that DESOTs are competitive with traditional DEs in our setting, both in predictive performance and uncertainty quantification performance. We also show that they outperform both single models and MC-dropout models.

In summary, our contributions are the following:

- (i) We propose Deep Ensembles Spread Over Time (DESOT), an approach applicable to sequences that brings the benefits of Deep Ensembles (DE) without the additional computational cost.
- (ii) We thoroughly analyze the proposed approach on traffic sign classification, where a sequence of tracked patches are fed as input to the model. DESOT obtains the benefit of DEs at the computational cost of a single model.
- (iii) We thoroughly analyze the out-of-distribution detection performance, based on the entropy of the predictions, and find that DESOT, like deep ensembles, substantially outperforms a single model.

2. Related Work

Deep ensembles: Deep ensembles have been shown to be superior to any other ensemble method in uncertainty quantification given a fixed computational budget [9]. Ovadia *et al.* [8] also showed that DEs are some of the best-performing models in uncertainty estimation. The intuition behind why this might be is that ensembling is a variance

reduction technique, and therefore useful for increasing the quality of epistemic uncertainty estimates [17]. Furthermore, they have additional advantages in their simple implementation and high predictive performance.

Though deep ensembles have been used for decades, Lakshminarayanan *et al.* [10] show that they produce state-of-the-art uncertainty estimates. DEs are simple to train, with three basic steps involved: (i) ensure that a proper scoring rule is used as loss function; (ii) optionally use adversarial training to increase robustness; and (iii) train the ensemble using randomized initialization of model parameters to increase variety in the ensemble [10]. Many common loss functions, such as cross-entropy loss, are strictly proper scoring rules and can therefore be used in the deep ensemble framework. In practice, adversarial training is often omitted if improved robustness is not strictly necessary. Lakshminarayanan *et al.* [10] also demonstrate that the model has the attractive property of decreasing its certainty of prediction in out-of-distribution examples, which was demonstrated using an ensemble trained on the MNIST dataset on examples from the NotMNIST dataset which contains letters instead of digits. It has later been verified that DEs are the SotA for UQ on OOD data [8], [9]. For these reasons, and that training time is not critically limited in the same way as computational capacity upon test-time inference, we chose DEs as the framework for generating our ensemble. It should be noted that we are not the first to attempt leveraging ensembles while keeping the test-time inference time down. Wortsman *et al.* [18] propose to average the weights of ensemble members into a single model. Havasi *et al.* [19] propose to divide a given neural network into sub-networks, which essentially constitute an ensemble.

Monte Carlo dropout: Dropout was first introduced by Srivastava *et al.* [20] as a regularization measure during training to limit overfitting and increase the generalizability of the learned representation. With dropout, each neuron is turned off at random during training according to a pre-specified probability, or dropout rate, p . This helps the network not to overfit, and therefore generalize better, as it has to create a more robust representation when any neuron can be dropped at any time. Recognizing that using an ensemble of a set of models is usually beneficial for model performance, Srivastava *et al.* [20] show that using dropout during inference is equivalent to sampling from an exponential set of possible smaller models, which yields higher overall performance. Gal and Ghahramani [21] later showed that performing a number of forward passes through a model with dropout enabled and averaging the results can be seen as a Bayesian approximation. They chose to call this MC-dropout and claimed that it enables superior uncertainty estimation performance in both regression and classification tasks compared to vanilla models. Of note is that since the introduction of MC-dropout, Ovadia *et al.* [8], Ashukha *et*

al. [9], and Lakshminarayanan *et al.* [10], have all claimed that deep ensembles are superior in uncertainty quantification. However, MC-dropout remains widely used due to its simple implementation and general improvement of performance compared to vanilla single models. This makes it a relevant mode of comparison for our proposed model.

Traffic sign recognition: We evaluate the proposed method on Traffic Sign Recognition (TSR), a field with a decades-long history of development [22]. Lately, TSR has become an important part of AD systems, and high and reliable performance is important for safety. This makes it an interesting application for DESOT.

There are two main subproblems of TSR, traffic sign detection and traffic sign classification [23]. This research concerns itself with the latter and assumes that regions of interest in the image have already been identified by another system (in this specific case, human annotators) earlier in the ML pipeline. The domain is characterized by a large number of classes with an imbalanced class distribution. Additionally, variations in illumination, perspective, and occlusions are common [22], making the problem distinctly long-tailed. Furthermore, many of the classes are very similar in shape and color, but with important differences in meaning, such as speed limit signs. Deep learning has recently started revolutionizing this domain, with many models achieving accuracies of over 95% in research settings [24]. This means that any differences in predictive performance between the models tested are likely to be small in absolute terms, and performance benefits might instead lie in the performance of the models on difficult examples such as short sequences or obscured scenes.

3. Deep Ensembles Spread over Time

In this section, we introduce Deep Ensembles Spread Over Time (DESOT), a method to obtain the benefits of a Deep Ensemble (DE) without additional computational cost. In many real-world scenarios, the input data is a sequence of frames. In the literature, it is then common to apply a single model and average the predictions [25]. This strategy is illustrated in Fig. 1a. It is well-known that replacing the single model with a DE improves both predictive- and uncertainty estimation performance (Fig. 1b). Unfortunately, this multiplies the computational cost with a factor M , where M is the number of ensemble members. DESOT (Fig. 1c) instead divides the ensemble over time, running one ensemble member per frame, and thus has the same computational cost as a single model. We hypothesize that DESOT still provides, despite the cheap computational cost, the benefits of DEs.

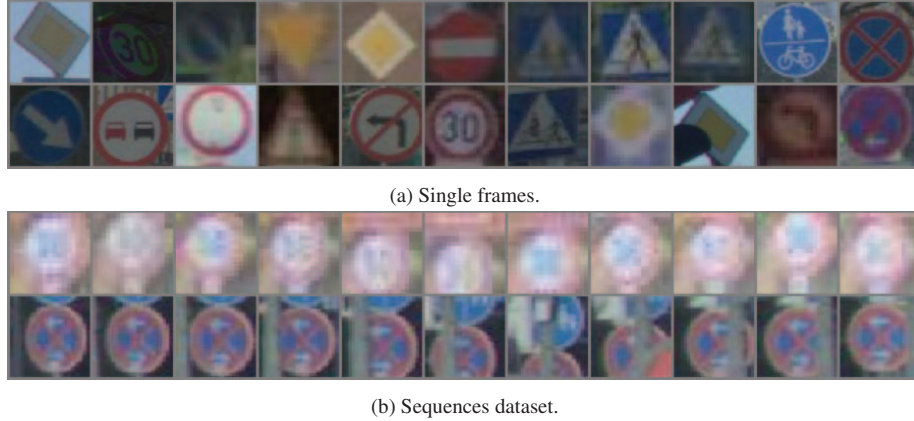


Figure 2: Examples of training and validation data. Training data is cropped from ZOD, and sequences are created by additional tracking across time using adjacent frames.

3.1. Image Classification on Sequences

We will now introduce the application of a single model or DE to *image classification on sequences*¹. First, both the ensemble fusion block and the time fusion block are simple averaging operations. This is how Lakshminarayanan *et al.* [10] implement ensemble fusion for deep ensembles in the classification setting. Now, define a sequence $\mathbf{x} \in \mathbb{R}^{T \times H \times W \times 3}$ as T distinct still images, each with a height of H pixels, a width of W pixels, and three separate color channels. We will consider classification problems where a model should produce a categorical output distribution across C classes for such a sequence \mathbf{x} . Assume there is a set of M different neural networks, each of which can conduct this classification. A single model $m \in \{1, \dots, M\}$ produces a categorical output distribution $\mathbf{p}_m(x^t)$ for each single image x^t at time step $t \in \{1, \dots, T\}$ in the sequence \mathbf{x} . Then, the final output distribution for model m is defined as

$$\mathbf{p}_m(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T \mathbf{p}_m(x^t), \quad (1)$$

which is the element-wise (class-wise) average across the output distributions for each image in the sequence at different time steps. This setup is what we will refer to as a single model (SM).

Now imagine that all M models are used for classifying each image in the sequence, such that the final output distribution for the sequence is

$$\mathbf{p}_{\text{DE}}(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M \mathbf{p}_m(\mathbf{x}) = \frac{1}{MT} \sum_{m=1}^M \sum_{t=1}^T \mathbf{p}_m(x^t), \quad (2)$$

¹Note that image classification on sequences is different from video classification, as the former does not need to model the dynamics of the scene.

which is how we choose to apply deep ensembles (DEs) [10] to sequences – averaged across the images of the sequence \mathbf{x} . DE_M will be used to denote an M -member deep ensemble.

3.2. DESOT

Our proposed method, which we call DESOT, instead uses a single model $m \in \{1, \dots, M\}$ for each image x^t , $t \in \{1, \dots, T\}$, to produce a categorical output distribution, but the models are alternated such that any given model m is used on average T/M times for a certain sequence of T images. Analogously to the notation used for DEs, DESOT_M will denote an M -member deep ensemble spread over time. We let $\sigma(t)$ denote the (repeated) mapping between each time step t and one of the ensemble members $m \in \{1, \dots, M\}$, and write the final output distribution produced by DESOT across time steps, ensemble members, and classes as

$$\mathbf{p}_{\text{DESOT}}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T \mathbf{p}_{\sigma(t)}(x^t), \quad (3)$$

As is clear from this definition, let us again stress that the DESOT method can only be applied to sequences since it fundamentally relies on alternating the ensemble member in use between neighboring frames. If DESOT was to be used on a sequence length of one, the model would be equivalent to a standard single model.

In the real world, one would have a continuous stream of frames from the cameras of the car. In such a case, a set window size might be used such that equations 1-3 each constitute a moving average across some previous images. Due to the shortness of the sequences we use in this paper, which are only 11 frames, we have chosen to use the average across the entire sequence for a model.

4. DESOT for Traffic Sign Recognition

We apply DESOT to traffic sign classification, a challenging image classification problem that is a cornerstone to autonomous driving. While traffic sign classification is oftentimes studied as a single-frame problem, autonomous vehicles will in practice obtain a sequence of frames.

4.1. Data

We use the Zenseact Open Dataset (ZOD) [26], a multi-modal dataset collected by the autonomous driving software company Zenseact. The data has been collected across a number of countries in Europe. For this work, we use the single-frame data in ZOD, which has high-quality annotations for 446k unique signs. Using the annotations, the traffic signs are cropped and saved separately. The majority of these signs are withheld for single-frame training and validation. Due to the lack of temporal sign annotations, 30k randomly selected signs are extended into sequences of crops using an off-the-shelf tracker on the 10 preceding frames, sampled at 15 Hz. These sequences are only used for testing. See Figure 2 for samples from the datasets.

4.2. Architecture and training

We use Resnet18 as the base for all our models tested. Note however that, just like DEs, our approach is architecture agnostic. In addition to the single model and traditional DE, we compare DESOT to an MC-dropout implementation. We use a dropout layer after each non-linearity with a dropout rate of 0.2. The dropout-version is run once on each frame in a sequence and then aggregated by averaging in the same way as for the other models. All models are trained for 30 epochs using the AdamW optimizer [27] with a learning rate of 0.0005. Cosine annealing is used for more stable convergence. A batch size of 256 is used. Classes with fewer than 10 occurrences in the training- and validation datasets are omitted from training and evaluation, as well as crops smaller than 16 pixels along any dimension. In line with the definition of deep ensembles by Lakshminarayanan *et al.* [10], all models are trained independently using a proper scoring rule. We use cross-entropy loss.

For comparisons on uncertainty quantification (UQ), all models are temperature scaled. We use the procedure first introduced by Guo *et al.* [3] where a separate validation set is used to find the temperature that optimizes negative log-likelihood. The validation set in question is a separate subset of single frames from ZOD. For the deep ensemble, a joint temperature scaling scheme is employed where the ensemble is treated as a single model with one temperature parameter. This was found to be superior to creating an ensemble of single models of perfect temperature. This finding agrees with research by Rahaman and Thiery [28]. We adopt the same temperature scaling for DESOT.

5. Results

Our study comprises an analysis of predictive performance (Sec. 5.1) and out-of-distribution uncertainty quantification (Sec. 5.2). The latter is divided into two parts, focusing firstly on real out-of-distribution examples and secondly on systematically increasing the severity of augmentations for in-distribution examples.

5.1. Predictive performance

We assessed the predictive performance of various models on the sequence dataset. Throughout the later epochs, all models achieved high accuracy, yet DESOT₅ and DE₅ particularly outperformed their counterparts in early epochs. Importantly, DESOT matched the performance of traditional DEs, with both models slightly surpassing a single model with temporal fusion (SM). Comprehensive performance details are documented in Table 1. When assessed via F1-score, DESOT maintained parity with DEs, creating a wider gap to SM looking at averages across runs. However, do note that variance also increases compared to accuracy. The inclusion of MC-dropout in the single model significantly diminished its performance. Notably, all models – a single-frame model operating on a randomly selected sequence frame without temporal fusion.

We also evaluated the calibration of each model using Brier reliability and Expected Calibration Error (ECE). Notably, a considerable gap emerged between the single model operating on a single frame and the one fused over the entire sequence. This disparity could be linked to the single-frame temperature scaling procedure used, and the tendency of averaging predictions over multiple images to reduce overall confidence, thereby negatively impacting calibration. However, both DESOT and DE substantially improved sequence calibration, even though they still fell significantly short of the single-frame calibration. These findings suggest the potential value of a sequence-aware fine-tuning step to create better-calibrated models for sequences.

An essential aspect of traffic sign recognition, and by extension, classification, is the long-tailed distribution of classes. The fact that high performance on all classes is important accentuates the significance of performance on rare classes within the aggregated performance metrics. To evaluate the models’ performance on these rare classes, we removed classes with more than 500 samples in the training dataset, leaving us with 625 sequences to evaluate performance on. The predictive performance results for this filtered dataset can be found in Table 2. Due to the smaller dataset (both in training and validation) the results are significantly more noisy. Nevertheless, the results indicate that both ensembling approaches improve performance for these difficult and rare cases.

Overall, our method notably outperforms a single model

Table 1: Performance of each strategy on the ZOD sequence dataset. The results include \pm one standard deviation of performance over 5 runs. All models are temperature scaled. Deep Ensemble (DE) improves over the single model (SM) in terms of predictive- and uncertainty estimation performance, but by multiplying the computational cost with the number of ensemble members. Our DESOT, in contrast, obtains the DE benefits while avoiding the added computational cost.

Model	Accuracy \uparrow	F1-score \uparrow	Brier reliability \downarrow	ECE \downarrow
SM (single frame)	94.85 \pm 0.06	70.42 \pm 1.27	0.0069 \pm 0.0007	0.21 \pm 0.06
SM	97.34 \pm 0.06	81.12 \pm 1.75	0.0124 \pm 0.0003	2.88 \pm 0.06
DESOT ₅	97.60 \pm 0.03	83.26 \pm 0.93	0.0108 \pm 0.0002	2.25 \pm 0.05
DE ₅	97.64 \pm 0.01	82.73 \pm 1.01	0.0108 \pm 0.0002	2.25 \pm 0.05
MC-dropout	97.10 \pm 0.09	76.79 \pm 2.71	0.0165 \pm 0.0004	4.03 \pm 0.07

Table 2: Predictive performance on a minority class version of the ZOD sequence dataset. Without extra computational cost beyond the single-model baseline, DESOT substantially outperform the single model. In contrast to the full dataset, the single-frame-single-model strategy yields competitive performance.

Model	Accuracy \uparrow	F1-score \uparrow
SM (single frame)	87.42 \pm 0.97	46.49 \pm 3.03
SM	87.70 \pm 0.58	43.65 \pm 2.84
DESOT ₅	89.53 \pm 0.71	47.53 \pm 1.58
DE ₅	89.41 \pm 0.38	46.57 \pm 2.19
MC-dropout	85.52 \pm 0.87	36.50 \pm 0.71

with a similar computational footprint, while matching the performance of DE₅, which has a five-fold larger computational footprint. This superior performance could be attributed to the ensemble members collectively offering a more expressive representation of potential traffic signs than a single model. The distinct advantage of DESOTs lies in their capacity to leverage this richer representation while minimizing computational demands.

5.2. Out-of-distribution uncertainty quantification performance

In the real world, ML systems often encounter data absent from the training set, known as Out-of-Distribution (OOD) data. Thus, it’s critical to evaluate how a system performs under such circumstances. Particularly, since machine learning systems can fail silently, confidently misclassifying OOD examples [1], we aim to characterize the behavior of various systems when confronted with OOD data. This is achieved by testing gradual OODness via augmentations, akin to Ovadia *et al.* [8], and by assessing performance on completely unseen data, as per Lakshminarayanan *et al.* [10]. This enables us to simulate common scenarios (e.g., rotated signs or changing lighting) and identify changes in output distribution for OOD detection.

5.2.1 Complete out-of-distribution data

The Zenseact Open Dataset includes a class named NotListed, representing traffic signs not included in any other class, such as destination signs and rare special signs.

This class serves as an unseen set of mixed traffic signs for OOD data testing. Since the models were not trained on this OOD data, previously used metrics like accuracy, F1-score, Brier score, and ECE aren’t applicable. Instead, we use entropy as a metric to quantify if the models ’know what they don’t know’. Higher entropy of the output distribution suggests higher model uncertainty.

The results are illustrated via graphs showcasing the entropy scores of different models. The effect of varying ensemble sizes can be seen in Figure 3, whereas different approaches are more easily compared in Figure 4. As expected, all models exhibit a drastic increase in entropy for OOD data compared to in-distribution data. The single-frame model shows limited adjustment to its uncertainty on OOD data. On the contrary, single-frame DE₅ and DE₁₀ demonstrate a substantial entropy increase on OOD data, which is consistent with previous research [9], [10]. Interestingly, the increase from 5 to 10 ensemble members is relatively minor, suggesting a saturation effect.

When evaluated over a sequence, the entropies of DEs and DESOTs exhibit remarkable similarity, with a slight advantage towards DEs. Surprisingly, the entropy distribution of single models, when aggregated over a sequence, nearly mirrors that of single-frame ensembles. Further, MC-dropout also escalates its entropy in a manner consistent with DEs and DESOTs, contradicting the findings of Lakshminarayanan *et al.* [10]. However, we must emphasize that the adoption of MC-dropout results in degraded predictive performance, as highlighted in our prior results.

Next, we used a simple thresholding strategy on the output entropy to test each model’s OOD detection potential, rewarding methods that clearly separate in- and out-of-distribution entropy distributions. We fitted a threshold to each model using half of the available sequences to detect in- or out-of-distribution samples. The other half of the sequences dataset was then used to evaluate OOD detection performance for each model, with results presented in Table 3, including the entropy threshold value and metrics such as accuracy, precision, recall, and F1-score.

While the optimal threshold value varied widely between models, the OOD detection performance of all models was relatively high. In a standard setting, DESOT and DE per-

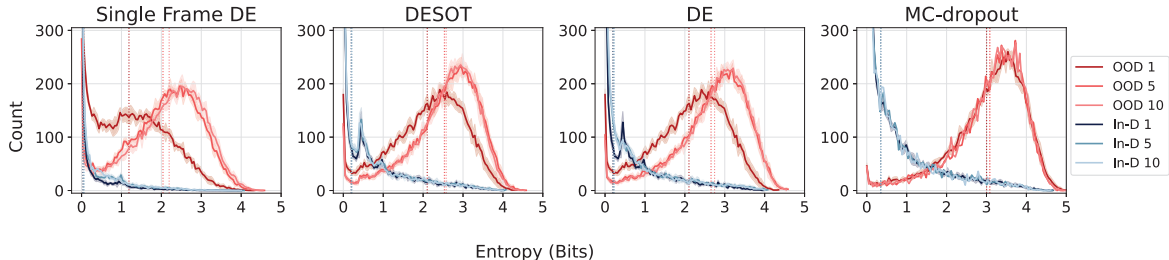


Figure 3: Histogram of the entropy for OOD data (red) and in-distribution data (blue) for the ZOD sequence dataset. We run various ensemble sizes $M \in \{1, 5, 10\}$, which are differentiated by color shade. Again, note that DE_1 and $DESOT_1$ are special cases that are equivalent to a single model. The vertical dashed lines are the mean entropy for the model of the same color. Like standard deep ensembles, DESOT increases its entropy on OOD data as additional ensemble members are added.

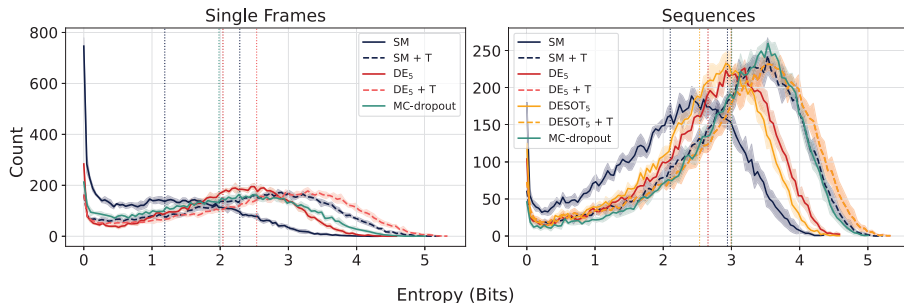


Figure 4: Histogram of the entropy for OOD data for the single-frame test dataset (left) and the sequences dataset (right). The $+T$ suffix denotes temperature scaling. Similar to deep ensembles, DESOT increases the entropy for OOD data, facilitating OOD detection.

form closely and clearly outperform the single model across all metrics. The introduction of temperature scaling, typically used to enhance the calibration of classification models, alters this picture. It marginally impacts the ensembles but dramatically improves the single model’s performance, almost entirely closing the gap with the ensembles. Surprisingly ensembles usually outdo temperature scaling in OOD detection. As our focus is the comparison between DESOTs and DEs, we do not further explore temperature scaling.

We also acknowledge insights from other research [29], highlighting potential pitfalls of this simple experiment, such as the risk of misidentifying ambiguous in-distribution samples as OOD – a problem we also encountered. Despite this, these tests offer a fundamental benchmark for more sophisticated OOD detection strategies, thereby illustrating the effectiveness of even such a straightforward method.

5.2.2 Shifted out-of-distribution data

We generate increasingly out-of-distribution data by escalating the severity of augmentations, inspired by the works of Hendrycks and Dietterich [30] and Ovadia *et al.* [8]. This experiment involved six distinct augmentations, each with a specific intensity range. As we increase the augmentation’s intensity, we anticipate a decrease in model accuracy. Ideally, the models should adjust their certainty to correspond with this accuracy drop, thereby maintaining calibration. Reliable model uncertainty estimates offer actionable

Table 3: Results from applying an entropy threshold for OOD detection on the sequences dataset. Note that DEs and DESOTs perform the best out of all models. When subject to temperature scaling (rows with $+T$ suffix), the difference in performance decreases.

	Threshold	Accuracy	Precision	Recall	F1
SM	0.709	0.895	0.811	0.905	0.855
DESOT ₅	1.048	0.919	0.852	0.922	0.886
DE ₅	1.071	0.919	0.856	0.919	0.886
MC-dropout	1.428	0.910	0.845	0.900	0.872
SM + T	1.463	0.916	0.856	0.906	0.880
DESOT ₅ + T	1.143	0.919	0.850	0.928	0.887
DE ₅ + T	1.178	0.920	0.853	0.925	0.888
MC-dropout + T	1.610	0.914	0.842	0.921	0.880

insights to downstream users.

However, when augmentation intensities are high, calibration may not provide a meaningful measure of uncertainty quantification. Since accurate sequence classification often becomes challenging under intense augmentations, it’s more crucial to identify such samples as out-of-distribution rather than obtaining a well-calibrated class distribution, as the ground truth class loses its relevance. Hence, we have also recorded the mean entropy across the dataset for each model and augmentation severity.

Figure 5 illustrates how various augmentations affect accuracy, Brier reliability [31], [32], and mean entropy of models in sequence settings at different intensity levels. Lower Brier reliability indicates better model calibration, while higher mean entropy suggests increased model un-

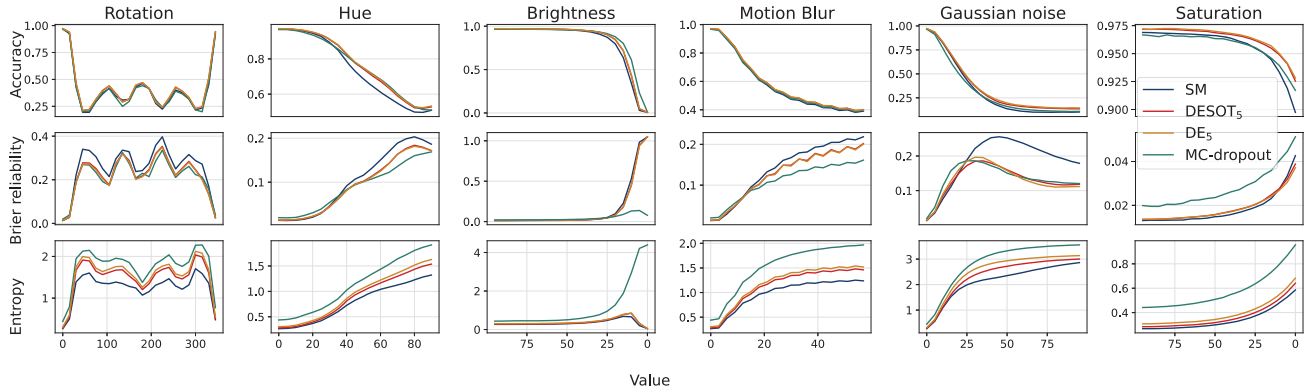


Figure 5: Uncertainty quantification performance for each model on augmented data of increasing intensity. The performance is measured in accuracy, Brier reliability, and mean entropy. Tested on the sequences dataset. All approaches tend to increase their entropy as augmentations become stronger. For some augmentations, the single model without temperature scaling and the MC-dropout increase entropy much more or much less than the other approaches.

certainty. For brevity, we omit results for the single-frame setting.

Observing the graphs, we notice a trend where a decrease in model accuracy corresponds with an increase in the Brier reliability score. This suggests that models become less calibrated as the data turns more out-of-distribution. Furthermore, entropy appears to rise with augmentation intensity, inversely correlating with accuracy. The strange trends in the rotation augmentation graphs are attributed to many signs having less than 360° rotational symmetry.

Examining the results from augmented OOD data, it appears that all models show comparable robustness to the augmentations in terms of predictive performance, deteriorating at similar rates for increased intensities. However, there’s a greater variation in model calibration upon augmentation. Single models struggle to maintain calibration at higher augmentation intensities, particularly for rotations, hue, motion blur, and Gaussian noise, a phenomenon also observed by Ovadia *et al.* [8]. As in the earlier OOD study, MC-dropout delivers convincing robustness, albeit at the expense of predictive performance. It is important to note that our study and that of Ovadia *et al.* [8] are performed on distinct datasets – our more complex sequence data versus their simpler MNIST dataset.

In general, DEs and DESOTs seem to offer similar calibration performances, measured in Brier reliability, for a given augmentation intensity, outperforming single models. Post temperature scaling, the disparity between the ensembling methods and single models is not as pronounced. An exception is Gaussian noise corruption, where single models distinctly fail to maintain calibration compared to others.

6. Conclusions

We have introduced Deep Ensembles Spread Over Time (DESOT), a novel approach that distributes ensemble mem-

bers across a sequence, achieving the predictive power and out-of-distribution robustness of running a full deep ensemble at each time step, with the computational cost of a single model. We extensively demonstrate the viability of this approach on the task of traffic sign classification, a highly relevant task where misinterpretations can lead to catastrophic outcomes and out-of-distribution signs are prevalent.

Looking ahead, we see exciting potential for DESOTs in more complex scenarios, such as 3D object detection where one could use a Kalman Filter for temporal fusion. We believe that our work opens up an avenue for new research on high-performing resource-efficient models, by demonstrating that ensembles can indeed be spread over time.

Limitations: The proposed strategy, DESOT, is applicable only for sequence processing problems. Although, in theory, the computational resources should match those required for a single model, it’s practical to anticipate that at least two members of the ensemble might need concurrent loading into memory. This scenario could potentially double the memory requirements. Moreover, while the improvements of DESOT (and deep ensembles) over SM are significant, we were surprised to see the large benefits of SM over the single-frame-single-model. A similar observation can be made for the out-of-distribution detection, where the effect of temperature scaling almost equals that of DESOT or deep ensembles. Though, combining DESOT or DEs with temperature scaling provides the best performance. We also note that the model calibration is actually better on a single frame than it is for multiple frames. One potential future research direction is thus to investigate how temperature scaling is best applied to the image classification on sequences problem.

Acknowledgements: This work was partially supported by the Wallenberg AI, Autonomous Systems, and Software Programme (WASP).

References

- [1] D. Amodei, C. Olah, J. Steinhardt, P. F. Christiano, J. Schulman, and D. Mané, *Concrete Problems in AI Safety*, 2016. DOI: 10.48550/arXiv.1606.06565.
- [2] D. Feng, A. Harakeh, S. L. Waslander, and K. Dietmayer, “A review and comparative study on probabilistic object detection in autonomous driving,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 9961–9980, 2021.
- [3] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On Calibration of Modern Neural Networks,” *CoRR*, vol. abs/1706.04599, 2017. arXiv: 1706.04599.
- [4] M. Minderer, J. Djolonga, R. Romijnders, *et al.*, “Revisiting the calibration of modern neural networks,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 15 682–15 694, 2021.
- [5] W. Liu, X. Wang, J. Owens, and Y. Li, “Energy-based out-of-distribution detection,” *Advances in neural information processing systems*, vol. 33, pp. 21 464–21 475, 2020.
- [6] A. D. Kiureghian and O. Ditlevsen, “Aleatory or epistemic? Does it matter?” *Structural Safety*, vol. 31, no. 2, pp. 105–112, Mar. 2009, ISSN: 01674730. DOI: 10.1016/j.strusafe.2008.06.020. (visited on 01/26/2023).
- [7] M. Abdar, F. Pourpanah, S. Hussain, *et al.*, “A review of uncertainty quantification in deep learning: Techniques, applications and challenges,” *Information Fusion*, vol. 76, pp. 243–297, 2021, ISSN: 1566-2535. DOI: 10.1016/j.inffus.2021.05.008.
- [8] Y. Ovadia, E. Fertig, J. Ren, *et al.*, *Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift*, 2019. DOI: 10.48550/ARXIV.1906.02530.
- [9] A. Ashukha, A. Lyzhov, D. Molchanov, and D. Vetrov, “Pitfalls of In-Domain Uncertainty Estimation and Ensembling in Deep Learning,” in *International Conference on Learning Representations*, 2020.
- [10] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles,” *International Conference on Neural Information Processing Systems (NIPS)*, vol. 31, pp. 6405–6416, 2017. DOI: 10.48550/arxiv.1612.01474.
- [11] M. Muhlbaiier, A. Topalis, and R. Polikar, “Ensemble Confidence Estimates Posterior Probability,” in *Multiple Classifier Systems*, N. C. Oza, R. Polikar, J. Kittler, and F. Roli, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 326–335, ISBN: 978-3-540-31578-0.
- [12] L. Hansen and P. Salamon, “Neural Network Ensembles,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 12, pp. 993–1001, Nov. 1990. DOI: 10.1109/34.58871.
- [13] T. G. Dietterich, “Ensemble Methods in Machine Learning,” in *Multiple Classifier Systems*, Springer Berlin Heidelberg, 2000, pp. 1–15, ISBN: 978-3-540-45014-6.
- [14] A. Mehrtash, W. Wells, C. Tempany, P. Abolmaesumi, and T. Kapur, “Confidence Calibration and Predictive Uncertainty Estimation for Deep Medical Image Segmentation,” *IEEE Transactions on Medical Imaging*, vol. PP, pp. 1–1, Jul. 2020. DOI: 10.1109/TMI.2020.3006437.
- [15] J. Mukhoti, A. Kirsch, J. van Amersfoort, P. H. S. Torr, and Y. Gal, *Deep Deterministic Uncertainty: A Simple Baseline*, 2021. DOI: 10.48550/ARXIV.2102.11582. (visited on 01/26/2023).
- [16] A. Wasay and S. Idreos, *More or less: When and how to build convolutional neural network ensembles*, Conference Paper, 2021.
- [17] N. Ueda and R. Nakano, “Generalization error of ensemble estimators,” in *Proceedings of International Conference on Neural Networks (ICNN’96)*, vol. 1, Washington, DC, USA: IEEE, 1996, pp. 90–95, ISBN: 978-0-7803-3210-2. DOI: 10.1109/ICNN.1996.548872. (visited on 08/08/2023).
- [18] M. Wortsman, G. Ilharco, S. Y. Gadre, *et al.*, “Model soups: Averaging weights of multiple fine-tuned models improves accuracy without increasing inference time,” in *International Conference on Machine Learning*, PMLR, 2022, pp. 23 965–23 998.
- [19] M. Havasi, R. Jenatton, S. Fort, *et al.*, “Training independent subnetworks for robust prediction,” *arXiv preprint arXiv:2010.06610*, 2020.
- [20] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014, ISSN: 1532-4435.
- [21] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *Proceedings of the 33rd International Conference on Machine Learning - Volume 48*, ser. ICML’16, JMLR.org, 2016, pp. 1050–1059. DOI: 10.48550/arXiv.1506.02142.
- [22] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel, “Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition,” *Selected Papers from IJCNN 2011*, vol. 32, pp. 323–332, Aug. 2012, ISSN: 0893-6080. DOI: 10.1016/j.neunet.2012.02.016.
- [23] M. Mathias, R. Timofte, R. Benenson, and L. Van Gool, “Traffic sign recognition — How far are we from the solution?” In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, 2013, pp. 1–8, ISBN: 978-1-4673-6129-3. DOI: 10.1109/IJCNN.2013.6707049.
- [24] A. F. Magnussen, N. Le, L. Hu, and E. W. Wong, “A Survey of the Inadequacies in Traffic Sign Recognition Systems for Autonomous Vehicles,” *International Journal of Performance Engineering*, vol. 16, no. 10, pp. 1588–1597, 2020. DOI: 10.23940/ijpe.20.10.p10.15881597.

- [25] E. Brissman, J. Johnander, M. Danelljan, and M. Felsberg, “Recurrent Graph Neural Networks for Video Instance Segmentation,” *International Journal of Computer Vision*, vol. 131, no. 2, pp. 471–495, Feb. 2023, ISSN: 0920-5691, 1573-1405. DOI: 10.1007/s11263-022-01703-8. (visited on 08/08/2023).
- [26] M. Alibeigi, W. Ljungbergh, A. Tonderski, *et al.*, “Zenseact Open Dataset: A large-scale and diverse multimodal dataset for autonomous driving,” 2023. DOI: 10.48550/ARXIV.2305.02008. (visited on 05/19/2023).
- [27] I. Loshchilov and F. Hutter, “Decoupled Weight Decay Regularization,” 2017. DOI: 10.48550/ARXIV.1711.05101. (visited on 05/23/2023).
- [28] R. Rahaman and A. H. Thiery, *Uncertainty Quantification and Deep Ensembles*, Nov. 2021. arXiv: 2007.08792 [cs, stat]. (visited on 02/17/2023).
- [29] A. Kirsch, J. Mukhoti, J. van Amersfoort, P. H. S. Torr, and Y. Gal, *On pitfalls in OoD detection: Entropy considered harmful*, 2021.
- [30] D. Hendrycks and T. Dietterich, *Benchmarking Neural Network Robustness to Common Corruptions and Perturbations*, Mar. 2019. arXiv: 1903.12261 [cs, stat]. (visited on 03/23/2023).
- [31] M. H. DeGroot and S. E. Fienberg, “The Comparison and Evaluation of Forecasters,” *Journal of the Royal Statistical Society. Series D (The Statistician)*, vol. 32, no. 1/2, pp. 12–22, 1983, ISSN: 00390526, 14679884. DOI: 10.2307/2987588. JSTOR: 2987588. (visited on 01/20/2023).
- [32] S. Siegert, “Simplifying and generalising Murphy’s Brier score decomposition,” *Quarterly Journal of the Royal Meteorological Society*, vol. 143, no. 703, pp. 1178–1183, 2017, ISSN: 1477-870X. DOI: 10.1002/qj.2985. (visited on 04/18/2023).