

Structured radial basis function network: modelling diversity for multiple hypotheses prediction

Conference or Workshop Item

Accepted Version

Dominguez, A. R., Shahzad, M. ORCID: <https://orcid.org/0009-0002-9394-343X> and Hong, X. ORCID: <https://orcid.org/0000-0002-6832-2298> (2025) Structured radial basis function network: modelling diversity for multiple hypotheses prediction. In: AI-2024 Forty-fourth SGAI International Conference on Artificial Intelligence, 17-19 December 2024, Cambridge, England, pp. 88-101. doi: 10.1007/978-3-031-77915-2_7 Available at <https://centaur.reading.ac.uk/118621/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: http://dx.doi.org/10.1007/978-3-031-77915-2_7

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Structured Radial Basis Function Network: Modelling Diversity for Multiple Hypotheses Prediction

Alejandro Rodriguez Dominguez, Muhammad Shahzad, and Xia Hong

Department of Computer Science, University of Reading, United Kingdom

a.j.rodriguezdominguez@pgr.reading.ac.uk

m.shahzad2@reading.ac.uk

xia.hong@reading.ac.uk

Abstract. Multi-modal problems can be effectively addressed using multiple hypothesis frameworks, but integrating these frameworks into learning models poses significant challenges. This paper introduces a Structured Radial Basis Function Network (s-RBFN) as an ensemble of multiple hypothesis predictors for regression. During the training of the predictors, first the centroidal Voronoi tessellations are formed based on their losses and the true labels, representing geometrically the set of multiple hypotheses. Then, the trained predictors are used to compute a structured dataset with their predictions, including centers and scales for the basis functions. A radial basis function network, with each basis function focused on a particular hypothesis, is subsequently trained using this structured dataset for multiple hypotheses prediction. The s-RBFN is designed to train efficiently while controlling diversity in ensemble learning parametrically. The least-squares approach for training the structured ensemble model provides a closed-form solution for multiple hypotheses and structured predictions. During the formation of the structured dataset, a parameter is employed to avoid mode collapse by controlling tessellation shapes. This parameter provides a mechanism to balance diversity and generalization performance for the s-RBFN. The empirical validation on two multivariate prediction datasets—air quality and energy appliance predictions—demonstrates the superior generalization performance and computational efficiency of the structured ensemble model compared to other models and their single-hypothesis counterparts.

Keywords: diversity · ensemble learning · multiple hypotheses prediction · radial basis functions · Voronoi tessellations.

1 Introduction

Multi-modality focuses on perception with a set of hypotheses instead of a single output to learn processes. Notable existing approaches include Multiple Choice Learning (MCL) [8, 9], Multiple Hypotheses Prediction (MHP) [16], Mixture-Of-Experts [21], Bagging [2], Boosting [7], and Meta-Learning [20]. Among them,

MCL differs in that it uses the output of different models or hypotheses as inputs to a structured ensemble model (or multiple structured prediction/classification task models), which are heterogeneous ensemble predictors that can vary in size, parameters, and architecture [8]. Diverse Multiple Choice Learning (DivMCL), is an extension proposed for diverse multi-output structured prediction by including a diversity encouraging term in the loss function used for training the models [9]. While DivMCL provides diversity, it trains separate networks which makes information exchange between individual predictors harder. To cope with this, the DivMCL ideas are extended by instead of training separate networks for each choice, the individual hypotheses are combined with Voronoi tessellations formed by the predictors' losses in a shared architecture. This allows sharing of information among predictors during training [16]. But, it is not clear how to optimally combine these predictors into an ensemble. To the best of the authors' knowledge, there is no existing method that optimally combines structured predictions from multiple hypotheses prediction with an ensemble learning model that can be trained with a closed-form solution [8, 9, 16].

Another important aspect in enhancing generalization of ensembles is the diversity of individual predictors. Diversity in this context has been extensively researched in literature, e.g., using Bias-Variance-Covariance decomposition [18], ambiguity decomposition [10], and their hybrid extensions [3]. However, there is also not a unifying framework for diversity in ensemble learning. Moreover, there is no clear connection in the literature between geometric properties of loss functions for individual predictors and diversity in ensemble learning [19]. This work focuses on the definition of diversity in ensemble learning as the variety of outputs from base learners that can improve the generalization performance of ensemble models [19].

Building on previous aspects [9, 16, 19], a new approach for multiple hypotheses prediction using a structured ensemble model is presented. In this approach, predictions from a set of base learners or individual predictors are used as inputs for a radial basis function network, with each predictor or hypothesis focusing on a specific basis function. The model is referred to as the Structured Radial Basis Function Network (s-RBFN). During training, the base learners form centroidal Voronoi tessellations (CVT), with each hypothesis or base learner assigned to a particular tessellation. A parametric formula from multiple hypotheses framework [16] is used to weight the updates of the base learner parameters in each iteration of gradient descent, preventing mode collapse and ensuring that all predictions fall within their respective tessellations. In this work, this strategy is applied to control diversity in ensemble learning similar to DivMCL [9] with the mechanism from MHP [16], enhancing generalization performance. The proposed s-RBFN can then be optimized using least squares, providing faster training compared to other existing structured models that rely on gradient descent or non-convex methods [8, 9, 16].

The paper is organized as follows: Section 2 presents a revision of the previous work in structured ensemble learning and diversity; Section 3 presents the

proposed model; Section 4 presents the experimental results and discussions; and finally, Section 5 provides the concluding remarks and outlook.

2 Literature Review

Multiple hypotheses prediction (MHP) methods extend semi-supervised ensembles and other single-loss, single-output systems to multiple outputs providing a piece-wise constant approximation of the conditional output space. They differ from mixture density networks by representing the uncertainty through a discrete set of hypotheses [16]. These models initially employed training techniques from multiple choice learning [1, 9] and later exploited the geometric properties from Voronoi tessellations formed by losses of the individual predictors as multiple hypothesis [16]. These approaches tend to be based on Winner-Takes-it-All (WTA) loss, meaning that the best base learner among all predictors gets updated during their training. A partial solution is a relaxed version of WTA [16] where in addition to the winner predictor, the other predictors also get updated for each iteration. It alleviates the convergence problem of the WTA, but still leads to hypotheses with incorrect modes. Moreover, when optimizing for a mixture distribution, the issues of numerical instabilities and mode collapsing arise. For this purpose, the evolving WTA loss was proposed [13] which addresses these issues by preserving the distribution, yielding regularly distributed hypotheses. Although this somewhat mitigates the issue but still the problem of how to combine the multiple hypothesis efficiently in a structured ensemble model persists.

Another aspect of MHP is the use of diversity which can serve as effective regularization - leading to possibly worse performance on training data, but better generalization on unseen test data [9]. Traditional diversity measures often assess the correlation or discrepancy between predictions of two models and their collective performance [11]. Recent innovations have introduced the Bias-Variance-Diversity decomposition, a nuanced framework that integrates various functional forms for each loss and directly links diversity to the expectation of ensemble ambiguity [19]. This approach goes beyond the traditional Ambiguity and Bias-Variance-Covariance decompositions, limited to squared-loss and arithmetic-mean combiners [10, 18]. In practice, strategies like bagging and boosting facilitate diversity among base learners by manipulating data, thus introducing structural and data diversity. Additional methods quantify diversity through non-maximal predictions and employ metrics such as the logarithm of ensemble diversity (LED) and ensemble entropy [15, 20]. More recently, MCL and DivMCL demonstrate superior test accuracy and better generalization compared to traditional multi-output prediction methods [8, 9]. These approaches emphasize minimizing oracle loss by focusing on specific hypothesis, contrasting with broader Mixture-of-Expert models [15, 20, 9, 8, 12]. Ultimately, the strategic integration of diversity not only serves as an effective regularization mechanism but also critically enhances the predictive accuracy and reliability of ensemble models, especially in managing out-of-distribution data. By optimizing ensemble diversity through sophisticated decomposition models and diverse ensemble

strategies, researchers can effectively balance error components to minimize overall mean-squared error, resulting in significantly improved predictions [18, 4].

3 Proposed Methodology

In this section, first the multiple hypotheses prediction with Voronoi Tessellations is presented. Later, it is explained how this could be scaled to operate in a structured setting for regression applications. This is done by generating the structured dataset using the MHP base learners' predictions. Finally the optimization of the s-RBFN using the structured dataset is efficiently carried out by least squares approach.

In the supervised learning setting, given training instances $\{\mathbf{x}_i\}_{i=1}^N$ and ground-truth labels $\{y_i\}_{i=1}^N$, the multiple hypotheses case involves a set of prediction functions $\{f_{\boldsymbol{\theta}_j}(\mathbf{x})\}_{j=1}^M$ with corresponding model parameters $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_j\}_{j=1}^M$. Assuming the training samples follow the distribution $p(\mathbf{x}, y)$, the expected error for a loss function \mathcal{L} is expressed as:

$$\int_{\mathbf{X}} \sum_{j=1}^M \int_{\mathcal{Y}_j(f_{\boldsymbol{\theta}_j}(\mathbf{x}))} \mathcal{L}(f_{\boldsymbol{\theta}_j}(\mathbf{x}), y) p(\mathbf{x}, y) dy d\mathbf{x} \quad (1)$$

During training, the Voronoi tessellation of the label space is induced by the losses computed from M predictors and given as $\mathcal{Y} = \bigcup_{j=1}^M \mathcal{Y}_j(f_{\boldsymbol{\theta}_j}(\mathbf{x}))$ where $\mathcal{Y}_j(f_{\boldsymbol{\theta}_j}(\mathbf{x}_i))$ represents the j th cell with $f_{\boldsymbol{\theta}_j}(\mathbf{x}_i)$ being the closest of the M predictions to the label data for each training iteration [16]:

$$\mathcal{Y}_j(f_{\boldsymbol{\theta}_j}(\mathbf{x}_i)) = \{y_i \in \mathcal{Y}_j : \mathcal{L}(f_{\boldsymbol{\theta}_j}(\mathbf{x}_i), y_i) < \mathcal{L}(f_{\boldsymbol{\theta}_k}(\mathbf{x}_i), y_i) \forall k \neq j\} \quad (2)$$

While implementing (2), a typical approach adopted to avoid mode collapse is to relax the best-of- M approach by updating all predictors in each iteration [8, 9]. Existing works either focus on multi-output prediction or does not provide an efficient way to combine the base learners or multiple hypotheses, often relying on numerical methods [8, 9, 14]. To this end, the aim of this work is to efficiently combine, in a structured model, the set of hypotheses that form the centroidal Voronoi tessellations. Additionally, the hypothesis that manipulating the shape of the tessellations formed during the training of the predictors, that has direct implications in generalization performance, is validated in the experiments. This is due to the diversity in ensemble learning induced by the predictors.

3.1 Structured Dataset Formation

Two step approach have been taken for structured dataset formation. Firstly, the set of predictors $\{f_{\boldsymbol{\theta}_j}(\mathbf{x})\}_{j=1}^M$ are trained with stochastic gradient descent with randomly initialised weights. Secondly, these learned models are used to generate the predictions that form the structured dataset.

In each i th iteration, using the j th prediction $f_{\theta_j}(\mathbf{x}_i)$ and the true label y_i , the predictors' parameters are updated using the stochastic gradient descent as follows:

$$\boldsymbol{\theta}_j = \boldsymbol{\theta}_j - \eta_j \left(\frac{\partial \mathcal{L}(f_{\theta_j}(\mathbf{x}_i), y_i)}{\partial \boldsymbol{\theta}_j} + \frac{\lambda_p}{N} \boldsymbol{\theta}_j \right) \delta(\mathcal{Y}_j(f_{\theta_j}(\mathbf{x}_i))) \quad (3)$$

where η_j denotes the learning rate for the j th predictor and the norm loss is computed as $\mathcal{L}(f_{\theta_j}(\mathbf{x}_i), y_i) = \|f_{\theta_j}(\mathbf{x}_i) - y_i\|_2^2 + \frac{\lambda_p}{2N} \sum_{j=1}^M \boldsymbol{\theta}_j^2$ with the regularization parameter λ_p . The function $\delta(\mathcal{Y}_j(f_{\theta_j}(\mathbf{x}_i)))$ serves as an indicator with a parameter $0 < \varepsilon < 1$ that can alter the shape of the tessellation during training [16]. This parameter enhances diversity in the structured dataset for ensemble generalization by regulating the extent to which non-top predictors' parameters are updated in each training iteration. It is defined as:

$$\delta(y \in \mathcal{Y}_j(f_{\theta_j}(\mathbf{x}))) = \begin{cases} 1 - \varepsilon & \text{if is true} \\ \frac{\varepsilon}{M-1} & \text{otherwise} \end{cases} \quad (4)$$

When the training of the predictors is completed, the same set of training instances $\{\mathbf{x}_i\}_{i=1}^N$ are used to generate structured dataset. To elaborate, if the prediction obtained after the forward pass for j th predictor on a particular training instance \mathbf{x}_i is denoted as $f_{\theta_j}(\mathbf{x}_i)$, then the resulting predictions for the entire structured dataset can be written in the matrix form as:

$$\mathbf{D}(\varepsilon) = \begin{bmatrix} f_{\theta_1}(\mathbf{x}_1) & \dots & f_{\theta_M}(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ f_{\theta_1}(\mathbf{x}_N) & \dots & f_{\theta_M}(\mathbf{x}_N) \end{bmatrix} \quad (5)$$

with $\mathbf{D}(\varepsilon) \in \mathbb{R}^{N \times M}$ being the matrix of predictions for a particular diversity parameter $0 \leq \varepsilon \leq 1$. Similarly, for any test set with test instances $\{\mathbf{x}'_i\}_{i=1}^n$, the structured test set $\mathbf{D}(\varepsilon)' \in \mathbb{R}^{n \times M}$ is given by the predictions $\{f_{\theta_j}(\mathbf{x}'_i)\}_{i=1}^n$. For any structured test dataset, the predictors use the same set of parameters $\boldsymbol{\Theta}$ obtained after training using stochastic gradient descent.

3.2 s-RBFN Optimisation

The structured dataset is used as input for the radial basis function network, with each j th predictor or hypothesis $f_{\theta_j}(\mathbf{x})$ associated to a particular basis function $\phi(f_{\theta_j}(\mathbf{x}), \mu_j, \sigma_j)$, i.e., a map $\boldsymbol{\Phi}(\mathbf{D}(\varepsilon)) : \mathbb{R}^{N \times M} \rightarrow \mathbb{R}^{N \times M}$ is obtained by applying the basis function $\phi(\cdot)$ to each element of $\mathbf{D}(\varepsilon)$, transforming it into:

$$\boldsymbol{\Phi}(\mathbf{D}(\varepsilon)) = \begin{bmatrix} \phi(f_{\theta_1}(\mathbf{x}_1), \mu_1, \sigma_1) & \dots & \phi(f_{\theta_M}(\mathbf{x}_1), \mu_M, \sigma_M) \\ \vdots & \ddots & \vdots \\ \phi(f_{\theta_1}(\mathbf{x}_N), \mu_1, \sigma_1) & \dots & \phi(f_{\theta_M}(\mathbf{x}_N), \mu_M, \sigma_M) \end{bmatrix}$$

In this work, the Gaussian basis function $\phi(\cdot) = \exp\left(\frac{-1}{2\sigma_j^2} |f_{\theta_j}(\mathbf{x}_i) - \mu_j|^2\right)$ have been used where the centers c_j and scales S_j parameters for the basis functions are computed from each column j of the structured training dataset $\mathbf{D}(\varepsilon)$ and are computed by $\mu_j = \frac{1}{N} \sum_{i=1}^N f_{\theta_j}(\mathbf{x}_i)$, and $\sigma_j = \sqrt{\sum_{i=1}^N \frac{(f_{\theta_j}(\mathbf{x}_i) - \mu_j)^2}{(N-1)}}$.

The s-RBFN formulation can now be expressed in matrix form as follows:

$$\hat{\mathbf{y}} = \Phi(\mathbf{D}(\varepsilon)) \mathbf{w} = \begin{bmatrix} \phi(f_{\theta_1}(\mathbf{x}_1), \mu_1, \sigma_1) & \dots & \phi(f_{\theta_M}(\mathbf{x}_1), \mu_M, \sigma_M) \\ \vdots & \ddots & \vdots \\ \phi(f_{\theta_1}(\mathbf{x}_N), \mu_1, \sigma_1) & \dots & \phi(f_{\theta_M}(\mathbf{x}_N), \mu_M, \sigma_M) \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_M \end{bmatrix} \quad (6)$$

The optimal weights $\{w_i\}_{j=1}^M$ in (6) can now be simply obtained by least-squares with regularization parameter λ_s for the structured model using:

$$\mathbf{w} = \left(\Phi(\mathbf{D}(\varepsilon))^T \Phi(\mathbf{D}(\varepsilon)) + \lambda_s * \mathbf{I}_{(mxm)} \right)^{-1} \Phi(\mathbf{D}(\varepsilon))^T \mathbf{y} \quad (7)$$

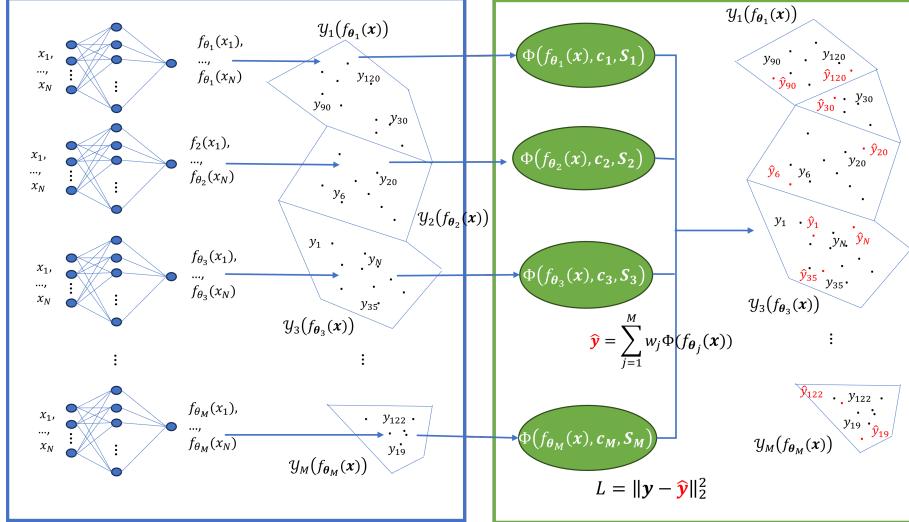


Fig. 1: Model architecture with structured data obtained from neural networks' predictions and the ground-truth labels \mathbf{y} forming centroidal Voronoi tessellations based on the neural networks' losses (Left box). The s-RBFN uses these predictions to estimate the ground-truth labels $\hat{\mathbf{y}}$, with L representing the s-RBFN norm loss (Right Box).

The whole approach presented above has been summarized in the Figure 1 where the model is shown with the structured data obtained in the left box, using neural networks as predictors. The label data is assigned to a particular

Voronoi tessellations depending on how far it is from the predictions of the base learners. This tessellation represents the multiple hypotheses prediction target values, \mathbf{y} (Left box in Figure 1). Once training is completed, the predictions from the trained predictors, using all training instances as input, are used as input data for training a radial basis function network (s-RBFN) via least-squares. The estimates $\hat{\mathbf{y}}$ in Figure 1 of the multiple hypotheses prediction ground-truth labels \mathbf{y} are given by the output from the s-RBFN (Right box in Figure 1). L represents the L_2 norm loss between the ground-truth labels and their estimates.

4 Experiments

4.1 Datasets

An Air Quality dataset [6] and the Appliances Energy Prediction dataset [5] are have been employed in this study. The first dataset consists of 9358 instances of hourly averaged responses from five metal oxide chemical sensors embedded in an air quality chemical multisensor device. The data was recorded from March 2004 to February 2005, and represents the longest freely available recordings of on-field responses from deployed air quality chemical sensor devices [6]. The goal is to predict absolute humidity values with the rest of variables in a multivariate regression problem. The second dataset consist of 10 minutes timestamps for 4.5 months making up over 20 thousand instances from 29 features. The goal is to predict energy appliances in a low energy building [5].

4.2 Models Performance & Comparisons

For the individual predictors, a 2-layer multi-layer perceptron (MLP) have been used with the number of neurons in each layer as κ , learning rates η , multiplicative factor of the initial weights χ , and regularization parameters λ_p . For the s-RBFN model, the number of predictors or hypotheses is given by M , diversity parameter ε , and s-RBFN regularization parameters λ_s . All values used for the hyper-parameters are displayed in Table 1.

For the experiments, the top performing models' versions from the original papers of the two used datasets [6, 17, 5], are replicated for comparison (top competitors). These are the Linear Model (LM), Random Forest (RF), Gradient-Boost (Gboost), and Support Vector Machine Radial Basis Function (SVM-RBF). To elaborate, for the s-RBFN, the experiments are performed with 10 simulations for each combination of hyper-parameters from Table 1. The mean and standard deviations of the RMSE for each of the 10-folds are recorded as performance measures. In total, the experiments have been performed with 80 different model hyper-parameters' configurations (also including the single hypothesis $M = 1$). Additionally, for further comparison, the bench-marking results using the baseline multiple hypothesis prediction (arithmetic combiner) model [16] are also included, in which the ensemble of individual predictors forming Voronoi Tessellations as their arithmetic mean are employed.

Table 1: Sets of values for the s-RBFN hyperparameters

(a) M number of hypotheses, κ number of neurons per layer, η learning rates for the predictors, χ is a multiplicative factor for random initial predictors' weights Θ .

M	κ	η	χ
[2, 5, 10, 20, 35]	[20, 200, 2000]	[0.03, 0.3]	[0.0001, 0.01, 0.1, 1]

(b) ε is the diversity parameter, λ_p is the regularization parameter for the predictors, λ_s is the regularization parameter for the s-RBFN.

ε	λ_p	λ_s
[0, 0.1, 0.35, 0.5]	[0, 0.0001, 0.01, 0.07]	[0, 3, 5]

Absolute Humidity Prediction In Table 2, the 10 cross-folds mean and standard deviation RMSE values for the top performing versions of all models on the test set are presented. The RMSE for the 80 different hyper-parameter configurations are computed and its first and third quartiles are shown in this table. For the rest of the models, 80 different hyper-parameters are applied for comparison.

The best model by generalization performance is the s-RBFN when the hyper-parameters are optimized. The SVM-RBF is the second best performing model. The arithmetic combiner has the lowest standard deviation and consequently has the smallest variation of the mean RMSE for all quartiles. The s-RBFN has a quarter of the 80 different hyper-parameter configurations' mean RMSE values lower than all other models except for the SVM-RBF, due to its higher standard deviation.

Table 2: Absolute humidity prediction: Mean and standard deviation of the 10-fold cross-validation RMSE for the models with the top-performing hyper-parameters configuration in generalization performance. First and third quartiles are shown for all models from 80 different hyper-parameter configurations.

Models	Top Model	std dev	First Quartile	Third Quartile
Linear Model	7692.78	1657.53	8189.84	10488.35
SVM-RBF	29.83	1.99	34.80	37.65
Random Forest	55.66	15.47	69.00	91.55
Gradient Boosting	55.76	38.73	93.92	151.58
Arithmetic Combiner	39.19	0.15	41.75	43.93
s-RBFN	22.46	9.14	38.98	54.71

Table 3: Energy appliance prediction: Mean and standard deviation of the 10-fold cross-validation RMSE for the models with the top-performing hyper-parameters configuration in generalization performance. First and third quartiles are shown for all models from 80 different hyper-parameter configurations.

Models	Top Model	std dev	First Quartile	Third Quartile
Linear Model	281.76	297.69	321.39	803.31
SVM-RBF	104.68	1.27	107.26	109.01
Random Forest	298.46	29.35	328.48	373.81
Gradient Boosting	292.08	67.26	389.10	476.89
Arithmetic Combiner	115.17	0.11	128.54	144.83
s-RBFN	101.12	2.42	102.36	109.96

Energy Appliance Prediction For the energy appliance dataset the same set of experiments are performed as for the air quality dataset. In line with the results displayed in Table 2, in Table 3 it can be seen how the s-RBFN is the best performing model with less standard deviation than in the previous dataset. This makes the model best performer in the first and third quartiles. The arithmetic combiner is the model with lowest standard deviation and the SVM-RBFN is the second best performing model, in line with the air absolute humidity prediction experiments.

Thus, both the dataset, it has been validated empirically that the s-RBFN is the best performing model in terms of generalization performance and for a range of different hyper-parameters.

4.3 Diversity & Generalization Performance

In this section, the hypothesis of the improvement in generalization performance of the s-RBFN for different values of the diversity parameter ε and the number of hypotheses M is verified. Figures 2 and 3 show, for the air quality and energy appliances test sets respectively, the mean RMSE and 90% confidence interval using 10-fold cross-validation for each hyper-parameter configuration, and for different values of the number of hypotheses M and diversity parameter ε . The horizontal axis represents the pairs of hyper-parameters M and ε .

The results in Figure 2 indicate, for the absolute humidity prediction experiments with the air quality test set, that the generalization performance increases with the diversity parameter up to a certain number of hypotheses, but decreases if the number of hypotheses is too large. In this set of experiments the optimal pair for $M = 10$ and $\varepsilon = 0.35$ is well defined. For this pair of hyper-parameters the s-RBFN achieves the best performance, equal to the shown in Table 2. It can be shown that increasing ε for two hypotheses worsen the generalization performance, meaning that a minimum number of hypotheses is needed for diversity to improve generalization capabilities.

In Figure 3, the energy appliances prediction dataset shows the same conclusion with some different results. For relatively large number of hypotheses

($M = 10, 20$) the s-RBFN achieves the best performance in generalization for relatively large $\varepsilon = 0.35$. However, this improvement is not observed for $M = 2$ and $M = 5$, as for five hypotheses the best model has $\varepsilon = 0.1$, with the case of $\varepsilon = 0.35$ being worse than for the case of $\varepsilon = 0$. It is reasonable to believe that for each number of hypotheses there is an optimal level of diversity, or ε , for the s-RBFN model. In the case of two hypotheses ($M = 2$), there is no impact of diversity due to the low number of individual predictors. Moreover, for this case, the performance is very good, suggesting that while diversity can enhance generalization performance for a given number of hypotheses, there may be cases in which the individual predictor alone is good enough for prediction in the test set.

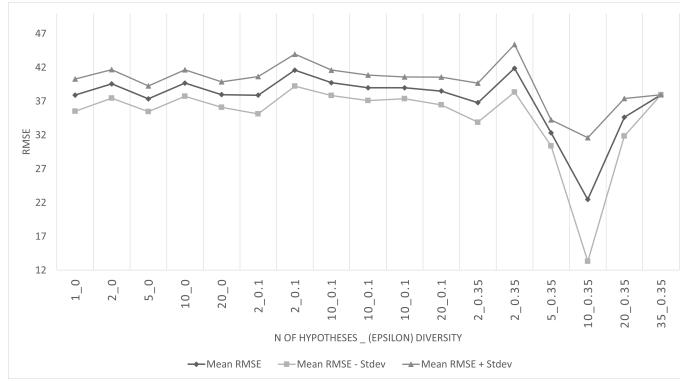


Fig. 2: Air quality test set: Mean RMSE and 90% confidence interval from 10-fold cross-validation for different configurations for hyper-parameters M and ε .

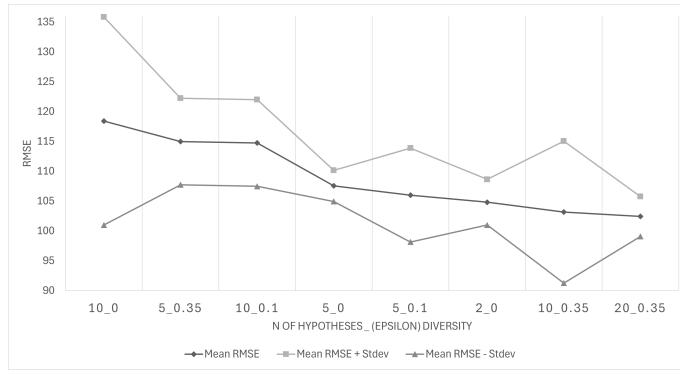


Fig. 3: Energy appliances test set: Mean RMSE and 90% confidence interval from 10-fold cross-validation for different configurations for hyper-parameters M and ε .

4.4 Impact of Regularization

In this section, the purpose is to understand the contribution of the regularization parameter for the s-RBFN in generalization performance. The regularization parameter λ_s has a clear effect in reducing the uncertainty of the hyper-parameters in the prediction of the structured ensemble model. For the air quality test set, in Figure 4a, it can be seen that for greater values of the regularization parameter, the mean RMSE for different hyper-parameter configurations remain more constant. Additionally, the standard deviation is lower for greater values of λ_s , as shown in Figure 4b. The same pattern is observed in the energy appliances test set with Figures 5a and 5b. It can be concluded that the regularization parameter reduces the uncertainty of the s-RBFN hyper-parameters. It also improves the s-RBFN generalization performance, on average, for any value of the hyper-parameters.

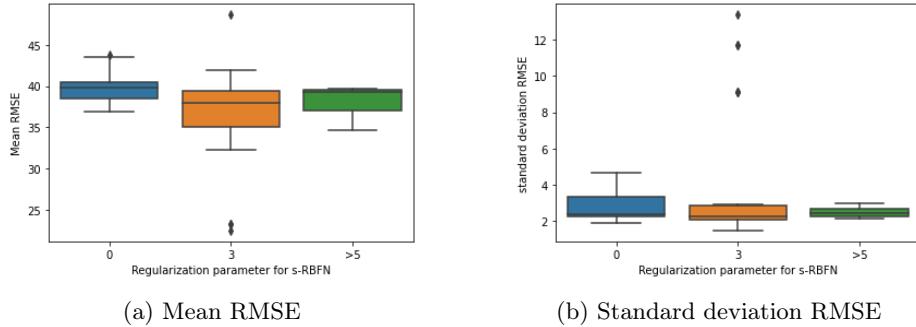


Fig. 4: Air quality test set: (a) Mean and (b) Standard deviation RMSE for 10-Fold cross-validation for different s-RBFN regularization parameters and hyper-parameter configurations.

In summary, the experiments demonstrate that diversity in structured ensemble models, particularly in the s-RBFN, is a distinctive feature of these architectures. There are instances where a single-hypothesis model may perform optimally. This indicates that diversity is not universally beneficial for enhancing generalization performance but rather improves performance contingent on a specific number of hypotheses. The experiments suggest there is indeed an optimal level of diversity, ε , for each number of hypotheses. Conversely, there exists a maximum number of hypotheses beyond which the performance of the ensemble model deteriorates, regardless of the ε level, indicating limits to the benefits of diversification. Similarly, for ε values exceeding 0.35, there is a noticeable decline in overall generalization capabilities.

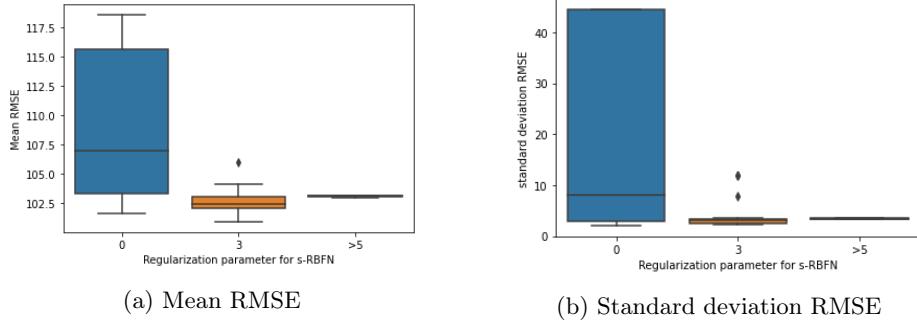


Fig. 5: Energy Appliance test set: (a) Mean and (b) Standard deviation RMSE for 10-Fold cross-validation for different s-RBFN regularization parameters and hyper-parameter configurations.

5 Conclusion

This work introduces a novel structured ensemble model for single-output multiple hypotheses prediction. The presented model incorporates geometric properties of centroidal Voronoi tessellations with the individual predictors' losses during training. By altering the shape of the tessellations through a parametric mechanism, the diversity is introduced to the structured dataset for the s-RBFN model. It has been validated through experiments that the s-RBFN model surpasses other models in generalization performance across a range of hypotheses numbers and diversity parameters. This model is the fastest to train once the structured dataset is prepared using its closed-form expression. Additionally, it facilitates easy control over diversity in structured ensemble learning and multiple hypotheses prediction for single-output regression problems through the diversity parameter. It is crucial to analyze the appropriate number of hypotheses and diversity hyper-parameters for a specific dataset, as these are highly correlated with the generalization performance capabilities of the s-RBFN.

For future work, several areas can be explored to enhance structured ensemble models in multiple hypotheses prediction. For instance, this work uses tabular data for regression with a simple 2-layer network as individual predictors. It would be interesting to employ more datasets from other modalities, e.g., visual or text datasets and use deeper architectures. This would allow to further investigate the relationship between model diversity, complexity and the generalization performance.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Batra, D., Yadollahpour, P., Guzman-Rivera, A., Shakhnarovich, G.: Diverse m-best solutions in markov random fields. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) Computer Vision – ECCV 2012. pp. 1–16. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
2. Breiman, L.: Bagging predictors. *Mach. Learn.* **24**(2), 123–140 (aug 1996). <https://doi.org/10.1023/A:1018054314350>
3. Brown, G.: Diversity in Neural Network Ensembles. Ph.D. thesis, University of Birmingham, United Kingdom (2004), winner, British Computer Society Distinguished Dissertation Award
4. Brown, G., Wyatt, J., Harris, R., Yao, X.: Diversity creation methods: A survey and categorisation. *Information Fusion* **6**, 5–20 (03 2005). <https://doi.org/10.1016/j.inffus.2004.04.004>
5. Candanedo, L.M., Feldheim, V., Deramaix, D.: Data driven prediction models of energy use of appliances in a low-energy house. *Energy and Buildings* **140**, 81–97 (2017). <https://doi.org/https://doi.org/10.1016/j.enbuild.2017.01.083>, <https://www.sciencedirect.com/science/article/pii/S0378778816308970>
6. De Vito, S., Massera, E., Piga, M., Martinotto, L., Di Francia, G.: On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario. *Sensors and Actuators B: Chemical* **129**(2), 750–757 (2008). <https://doi.org/https://doi.org/10.1016/j.snb.2007.09.060>, <https://www.sciencedirect.com/science/article/pii/S0925400507007691>
7. Grandvalet, Y., d'Alché Buc, F., Ambroise, C.: Boosting mixture models for semi-supervised learning. vol. 2130, pp. 41–48 (08 2001)
8. Guzmán-rivera, A., Batra, D., Kohli, P.: Multiple choice learning: Learning to produce multiple structured outputs. In: Pereira, F., Burges, C., Bottou, L., Weinberger, K. (eds.) Advances in Neural Information Processing Systems. vol. 25. Curran Associates, Inc. (2012)
9. Guzman-Rivera, A., Kohli, P., Batra, D., Rutenbar, R.: Efficiently Enforcing Diversity in Multi-Output Structured Prediction. In: Kaski, S., Corander, J. (eds.) Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 33, pp. 284–292. PMLR, Reykjavik, Iceland (22–25 Apr 2014), <https://proceedings.mlr.press/v33/guzman-rivera14.html>
10. Krogh, A., Vedelsby, J.: Neural network ensembles, cross validation, and active learning. In: Tesauro, G., Touretzky, D., Leen, T. (eds.) Advances in Neural Information Processing Systems. vol. 7. MIT Press (1994)
11. Kuncheva, L., Whitaker, C.: Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning* **51**, 181–207 (05 2003). <https://doi.org/10.1023/A:1022859003006>
12. Lee, S., Purushwalkam, S., Cogswell, M., Ranjan, V., Crandall, D., Batra, D.: Stochastic multiple choice learning for training diverse deep ensembles. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. p. 2127–2135. NIPS'16, Curran Associates Inc., Red Hook, NY, USA (2016)
13. Leemann, T., Sackmann, M., Thielecke, J., Hofmann, U.: Distribution preserving multiple hypotheses prediction for uncertainty modeling. *CoRR* **abs/2110.02858** (2021), <https://arxiv.org/abs/2110.02858>

14. Mienye, I.D., Sun, Y.: A survey of ensemble learning: Concepts, algorithms, applications, and prospects. *IEEE Access* **10**, 99129–99149 (2022). <https://doi.org/10.1109/ACCESS.2022.3207287>
15. Pang, T., Xu, K., Du, C., Chen, N., Zhu, J.: Improving adversarial robustness via promoting ensemble diversity. In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 97, pp. 4970–4979. PMLR (09–15 Jun 2019), <https://proceedings.mlr.press/v97/pang19a.html>
16. Rupprecht, C., Laina, I., DiPietro, R., Baust, M., Tombari, F., Navab, N., Hager, G.D.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3591–3600 (2017)
17. Saverio Vito: Air quality (2016), the dataset contains 9358 instances of hourly averaged responses from an array of 5 metal oxide chemical sensors embedded in an Air Quality Chemical Multisensor Device., 10.24432/C59K5F
18. Ueda, N., Nakano, R.: Generalization error of ensemble estimators. *Proceedings of International Conference on Neural Networks (ICNN'96)* **1**, 90–95 vol.1 (1996)
19. Wood, D., Mu, T., Webb, A.M., Reeve, H.W.J., Luján, M., Brown, G.: A unified theory of diversity in ensemble learning. *J. Mach. Learn. Res.* **24**(1) (mar 2024)
20. Yang, Y., Lv, H., Chen, N.: A survey on ensemble learning under the era of deep learning. *Artif. Intell. Rev.* **56**(6), 5545–5589 (nov 2022). <https://doi.org/10.1007/s10462-022-10283-5>, <https://doi.org/10.1007/s10462-022-10283-5>
21. Yuksel, S.E., Wilson, J.N., Gader, P.D.: Twenty years of mixture of experts. *IEEE Transactions on Neural Networks and Learning Systems* **23**(8), 1177–1193 (2012). <https://doi.org/10.1109/TNNLS.2012.2200299>