



HAL
open science

Resilient Multiple Choice Learning: A learned scoring scheme with application to audio scene analysis

Victor Letzelter, Mathieu Fontaine, Mickaël Chen, Patrick Pérez, Slim Essid, Gael Richard

► To cite this version:

Victor Letzelter, Mathieu Fontaine, Mickaël Chen, Patrick Pérez, Slim Essid, et al.. Resilient Multiple Choice Learning: A learned scoring scheme with application to audio scene analysis. *Advances in neural information processing systems*, Dec 2023, New Orleans, United States. hal-04216055

HAL Id: hal-04216055

<https://hal.science/hal-04216055v1>

Submitted on 31 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Resilient Multiple Choice Learning: A learned scoring scheme with application to audio scene analysis

Victor Letzelter^{1,2}

Mathieu Fontaine²

Mickael Chen¹

Patrick Pérez¹

Slim Essid²

Gaël Richard²

¹ valeo.ai, Paris, France

² LTCI, Télécom Paris, Institut Polytechnique de Paris, France

Abstract

We introduce Resilient Multiple Choice Learning (rMCL), an extension of the MCL approach for conditional distribution estimation in regression settings where multiple targets may be sampled for each training input. Multiple Choice Learning is a simple framework to tackle multimodal density estimation, using the Winner-Takes-All (WTA) loss for a set of hypotheses. In regression settings, the existing MCL variants focus on merging the hypotheses, thereby eventually sacrificing the diversity of the predictions. In contrast, our method relies on a novel learned scoring scheme underpinned by a mathematical framework based on Voronoi tessellations of the output space, from which we can derive a probabilistic interpretation. After empirically validating rMCL with experiments on synthetic data, we further assess its merits on the sound source localization task, demonstrating its practical usefulness and the relevance of its interpretation.

1 Introduction

Machine learning models are commonly trained to produce, for any given input, a single prediction. In most cases, for instance, when minimizing the empirical risk with quadratic loss, this prediction can be interpreted as the conditional output expectation given the input. However, there are many tasks for which the conditional output distribution can be multimodal, either by the nature of the task or due to various sources of uncertainty. Temporal tracking and forecasting, for instance, are problems of this type [5, 32, 44]. In such cases, the mean might fall in a low-density region of the conditional probability function, and it would be beneficial to predict multiple hypotheses instead [34, 20]. In this context, Multiple Choice Learning (MCL) has emerged as a simple pragmatic solution [17, 29]. Thanks to a network architecture with multiple heads, MCL models can produce multiple hypotheses, one per head. During supervised training, the gradients are only computed for the head that provides the best prediction given the current input sample. This Winner-Takes-All (WTA) training scheme allows each head to specialize in a region of the output space. Accordingly, Rupprecht et al. [34] have proposed a probabilistic interpretation of MCL based on Voronoi tessellations.

On that account, MCL suffers from two significant issues: *hypotheses collapse* and *overconfidence*. *Hypotheses collapse* occurs during training when a head takes the lead, the others being almost never selected under WTA and thus not updated. This leads to a situation where most heads are not trained effectively and produce meaningless outputs during inference. *Overconfidence* can also be observed at inference time when looking at MCL under the lens of its probabilistic interpretation. Hypotheses that correspond to rare events tend to be over-represented, thus not reflecting the true distribution of

outputs well. Multiple variants of MCL have been developed to alleviate these issues, in particular for classification or segmentation tasks [29, 27, 42, 13].

In this paper, we are interested in MCL in the context of regression settings, which has been mostly overlooked. More specifically, we propose rMCL, for *resilient Multiple Choice Learning*, a MCL variant with a learned scoring scheme to mitigate overconfidence. Unlike prior work, rMCL is designed to handle more general settings where multiple targets may be sampled for each input during training. Moreover, we propose a probabilistic interpretation of rMCL, and show that rMCL’s scoring fixes overconfidence by casting multi-choice regression as conditional distribution estimation. While introduced in the context of the original WTA loss, the proposed scoring approach is also compatible with other variants of WTA. We validate our claims with experiments on synthetic data, and we further evaluate the performance of rMCL on the sound source localization problem. The accompanying code with a rMCL implementation is made available.¹

2 Related Work

Uncertainty estimation. The estimation of machine learning model’s uncertainty can be studied from different perspectives depending on the type of ambiguity being considered. The uncertainty can concern either the model (*epistemic*) or the data (*aleatoric*) [7, 12]. Kendall and Gal [22] showed that those types of uncertainty can be combined in deep learning settings, placing both a distribution on the weights and the outputs of the model. Epistemic uncertainty estimation may, for instance, be achieved by variational approximation of the weights posterior [31, 14]. Independent ensembles (IE) of neural networks [18, 11] is another widespread technique for epistemic uncertainty approximation. Several neural networks, trained from different random initializations, provide samples for approximating the weights posterior distribution [26].

Multiple Choice learning. Originally introduced by Guzman-Rivera et al. [17] and adapted to deep learning settings by Lee et al. [28, 29], MCL is a framework in which multiple output heads propose different possible predictions. Since many datasets contain only a single output realization for each input, the heads are trained with a Winner-Takes-All scheme where only the head that made the best prediction is updated. It can be understood as a specific case of dependent ensembles [8, 46, 4, 47] where the members interact with each other during training. Alternatively, [34] have shown that MCL can be adapted successfully for multi-label classification [23].

The Winner-Takes-All training scheme, however, can cause two main issues known as *hypothesis collapse* and *overconfidence* [34, 6, 10, 19, 10, 27, 42]. Most related previous works tackle either or both of those questions. Hypothesis collapse occurs during training when some heads receive little or no updates due to bad initial values under the WTA loss and other hypotheses taking the lead. To fix this issue, Rupprecht et al. [34] proposed a relaxed Winner-Takes-All (RWTA, or ϵ -WTA) loss that allows the update of non-winner heads, albeit with gradients scaled by a small constant ϵ . While this approach ensures that every head receives gradients updates, it also tends to level the different heads, which counters the initial goal of WTA and thus needs careful tuning. Finally, one can leverage the evolving Winner-Takes-All loss [32], with the top- n heads getting updated (top- n -WTA) instead of only the best one. The authors validate that their approach, with a scheduled decreasing value for n , achieves improved conditioning of the outputs by reducing inconsistent hypotheses in the context of future frames prediction.

On the other hand, overconfidence is an issue that can be observed in inference when evaluating a model as a density estimator. In this case, one would want the different outputs to be distributed across the different modes of the true conditional probability of the outputs. Empirical observations [27, 42] show that this is not usually the case, and rare events tend to be overly represented, rendering such a model inadequate for integration within real-world decision-making systems. In this context, Lee et al. [27] proposed *Confident Multiple Choice Learning* (CMCL) for solving overconfidence in a classification setup. Additionally to the WTA loss, CMCL is trained by maximizing the entropy of the class distributions predicted by the non-selected hypotheses. The classifier’s final prediction is based on a simple average of the discrete class distributions predicted by each of the heads. Although designed for tackling the overconfidence problem, CMCL reduces the diversity of the hypotheses. On that account, *Versatile Multiple Choice Learning* [42] (vMCL) proposed to address the issue by leveraging a choice network aiming at predicting a score for each hypothesis head. This may be seen

¹<https://github.com/Victorletzelter/code-rMCL>

as a variant of Mixture-of-Experts approaches [33], in which the choice network is supervised with explicit targets. The final prediction in vMCL is derived by weighting the class distributions from each hypothesis by their respective scores. These works focus on classification tasks and aggregate the hypotheses, thereby losing individual information from each head. In contrast, we address regression tasks and demonstrate how to benefit from the diversity of predictions *without* aggregation.

We propose to revisit the vMCL approach for regression, single-target or multi-target, and to extend accordingly the mathematical interpretation of the WTA proposed by Rupprecht et al. [34].

3 rMCL regression and its probabilistic interpretation

After reviewing MCL, this section introduces resilient Multiple Choice Learning, dubbed rMCL. This proposed variant handles multi-target settings thanks to its scoring scheme. We then provide a distribution learning interpretation for rMCL that generalizes to regression tasks, thus completing the MCL probabilistic interpretations of prior works [9, 34].

3.1 Fixing the overconfidence issue in Multiple Choice Learning

Let $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{Y} \subseteq \mathbb{R}^q$ be the input and output spaces in a supervised learning setting, such that the training set \mathcal{D} is composed of samples $(\mathbf{x}_s, \mathbf{y}_s)$ from an underlying joint distribution $p(\mathbf{x}, \mathbf{y})$ over $\mathcal{X} \times \mathcal{Y}$.

Multiple choice learning was proposed to address tasks with ambiguous outputs, *i.e.*, for which the ground-truth distribution $p(\mathbf{y} | \mathbf{x})$ is multimodal [17, 29, 10]. Adapted by Lee et al. [29] for deep-learning settings in *Stochastic Multiple Choice Learning* (sMCL), it leverages several models $f_\theta \triangleq (f_\theta^1, \dots, f_\theta^K) \in \mathcal{F}(\mathcal{X}, \mathcal{Y}^K)$, referred to as K hypotheses, trained using the Winner-takes-all (or Oracle) loss. It consists, given an underlying loss function ℓ and for each sample $(\mathbf{x}_s, \mathbf{y}_s)$ in the current batch, of first the computation of

$$\mathcal{L}(f_\theta(\mathbf{x}_s), \mathbf{y}_s) \triangleq \min_{k \in [1, K]} \ell(f_\theta^k(\mathbf{x}_s), \mathbf{y}_s) \quad (1)$$

after each forward pass, followed by backpropagation on the *winner* hypothesis, that is, the minimizing one. Assuming now that a set \mathbf{Y}_s of targets is available in the training set for input \mathbf{x}_s , Firman et al. [10] have shown that (1) can be generalized by updating the best hypothesis per target using the meta-loss

$$\mathcal{L}(f_\theta(\mathbf{x}_s), \mathbf{Y}_s) = \sum_{\mathbf{y} \in \mathbf{Y}_s} \sum_{k=1}^K \mathbf{1}(\mathbf{y} \in \mathcal{Y}^k(\mathbf{x}_s)) \ell(f_\theta^k(\mathbf{x}_s), \mathbf{y}), \quad (2)$$

where

$$\mathcal{Y}^k(\mathbf{x}) \triangleq \{\mathbf{y} \in \mathcal{Y}, \ell(f_\theta^k(\mathbf{x}), \mathbf{y}) < \ell(f_\theta^r(\mathbf{x}), \mathbf{y}), \forall r \neq k\}. \quad (3)$$

This loss can, however, lead to poor conditioning of the output set of predictions (*hypothesis collapse*), with only one or a few hypotheses being exclusively selected for backpropagation. Furthermore, sMCL is subject to the *overconfidence* problem. Following Theorem 1 of [34], a necessary condition for minimizing the risk

$$\int_{\mathcal{X}} \sum_{k=1}^K \int_{\mathcal{Y}^k(\mathbf{x})} \ell(f_\theta^k(\mathbf{x}), \mathbf{y}) p(\mathbf{x}, \mathbf{y}) d\mathbf{y} d\mathbf{x}, \quad (4)$$

is that $f_\theta^k(\mathbf{x})$ amount to the conditional mean of $p(\mathbf{y} | \mathbf{x})$ within the Voronoi cell $\mathcal{Y}^k(\mathbf{x})$ for each non-zero probability cell. We say that the components $\{\mathcal{Y}^k(\mathbf{x})\}_k$ form a *centroidal* Voronoi tessellation [9]. However, this Theorem tells us nothing about the predictions of f_θ^k in very low probability zones; in such regions, the inference-time predictions $f_\theta^k(\mathbf{x})$ will be meaningless. During inference, the sMCL indeed faces a limitation: it cannot solely rely on the predicted hypotheses to identify Voronoi cells in the output space with low probability mass (See Figure 1). This observation leads us to propose *hypothesis-scoring* heads $\gamma_\theta^1, \dots, \gamma_\theta^K \in \mathcal{F}(\mathcal{X}, [0, 1])$. Those aim to predict, for unseen input \mathbf{x} , the probability $\mathbb{P}(Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x}))$ where $Y_{\mathbf{x}} \sim p(\mathbf{y} | \mathbf{x})$ with the aim of mitigating this overconfidence issue. While MCL variants have often been proposed for classification tasks, we are, to the best of our knowledge, the first to propose to solve this issue for multi-target regression, interpreting the problem as a multimodal conditional distribution estimation.

3.2 Resilient Multiple Choice Learning

We consider hereafter a problem of estimation of a multimodal conditional distribution denoted $p(\mathbf{y} | \mathbf{x})$ for each input $\mathbf{x} \in \mathcal{X}$. In a real-world setting, only one or several samples (the *targets*) drawn from $p(\mathbf{y} | \mathbf{x})$ are usually accessible. Sound source localization is a concrete instance of such multimodal prediction for whose $p(\mathbf{y} | \mathbf{x}_s)$ represents the sound source position for an input audio clip $\mathbf{x}_s \in \mathcal{X}$ at a given time t .

In this last example, each target sample in \mathbf{Y}_s may represent the location of a *mode* of the ground-truth multimodal distribution. For such multi-output regression tasks, the MCL training of a randomly initialized multi-hypotheses model with scoring functions, $(f_\theta^1, \dots, f_\theta^K, \gamma_\theta^1, \dots, \gamma_\theta^K)$ can be adapted as follows.

For each training sample $(\mathbf{x}_s, \mathbf{Y}_s)$, let

$$\mathcal{K}_+(\mathbf{x}_s) \triangleq \left\{ k^+ \in \llbracket 1, K \rrbracket : \exists \mathbf{y} \in \mathbf{Y}_s, k^+ \in \underset{k}{\operatorname{argmin}} \ell(f_\theta^k(\mathbf{x}_s), \mathbf{y}) \right\} \quad (5)$$

and $\mathcal{K}_-(\mathbf{x}_s) \triangleq \llbracket 1, K \rrbracket - \mathcal{K}_+(\mathbf{x}_s)$ be the set of *positive* (or *winner*) and *negative* hypotheses respectively. It is then possible to combine the multi-target WTA loss \mathcal{L} in (2) with a hypothesis scoring loss

$$\mathcal{L}_{\text{scoring}}(\theta) \triangleq - \left(\sum_{k^+ \in \mathcal{K}_+(\mathbf{x}_s)} \log \gamma_\theta^{k^+}(\mathbf{x}_s) + \sum_{k^- \in \mathcal{K}_-(\mathbf{x}_s)} \log \left(1 - \gamma_\theta^{k^-}(\mathbf{x}_s) \right) \right), \quad (6)$$

in a compound loss $\mathcal{L} + \beta \mathcal{L}_{\text{scoring}}$. This novel approach differs from previous MCL variants [42] in its ability to predict multimodal distributions in regression settings and by the introduction of separated scoring branches that are updated based on loss (6), such that the target for the scoring branch k is the probability that hypothesis k is among the winners for that sample.

A resource-efficient implementation of rMCL is achieved by deriving the hypotheses and score heads from a shared representation, with distinct parameters at the final stages of the architecture (*e.g.*, through independent fully connected layers). Designed as such, it is also possible to reduce memory cost by updating only a fraction of score heads associated with the negative hypotheses at each training step, typically with distinct samples $k^- \sim \mathcal{U}(\mathcal{K}_-(\mathbf{x}_s))$. This trick could also alleviate the imbalanced binary classification task that the scoring heads face with a large number of hypotheses $|\mathcal{K}_-(\mathbf{x}_s)| \gg |\mathcal{K}_+(\mathbf{x}_s)|$ (typically, the number of target samples at disposal is small relative to the number of hypotheses K). The variants of the WTA (*e.g.*, top- n -WTA, ε -WTA, see Sec. 2) are also compatible with rMCL. Inference with the proposed rMCL model is outlined in Algorithm 1.

As highlighted above, the hypothesis output can be interpreted, in the context of risk minimization, as the conditional mean of the Voronoi cell $\mathcal{Y}^k(\mathbf{x})$ it defines, providing that the cell has non-zero probability. Furthermore, given (6), the output of the score head $\gamma_\theta^k(\mathbf{x})$ can be interpreted as an approximation of the probability of a sample from $p(\mathbf{y} | \mathbf{x})$ to belong to this cell.

Algorithm 1 Inference in the rMCL model

- Input:** Unlabelled input $\mathbf{x} \in \mathcal{X}$. Trained hypotheses and score heads $(f_\theta^1, \dots, f_\theta^K, \gamma_\theta^1, \dots, \gamma_\theta^K) \in \mathcal{F}(\mathcal{X}, \mathcal{Y})^K \times \mathcal{F}(\mathcal{X}, [0, 1])^K$.
- Output:** Prediction of the output conditional distribution $p(\mathbf{y} | \mathbf{x})$.
- 1: Perform a forward pass by computing $f_\theta^1(\mathbf{x}), \dots, f_\theta^K(\mathbf{x}), \gamma_\theta^1(\mathbf{x}), \dots, \gamma_\theta^K(\mathbf{x})$.
 - 2: Construct the associated Voronoi components $\mathcal{Y}^k(\mathbf{x})$ (3) with $\mathcal{Y} = \cup_{k=1}^K \mathcal{Y}^k(\mathbf{x})$.
 - 3: Normalize the predicted scores $\gamma_\theta^k(\mathbf{x}) \leftarrow \gamma_\theta^k(\mathbf{x}) / \sum_{k=1}^K \gamma_\theta^k(\mathbf{x})$.
 - 4: If $Y_{\mathbf{x}} \sim p(\mathbf{y} | \mathbf{x})$, then for k such that $\gamma_\theta^k(\mathbf{x}) > 0$, interpret the predictions as estimations of

$$\begin{cases} \gamma_\theta^k(\mathbf{x}) = \mathbb{P}(Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x})) \\ f_\theta^k(\mathbf{x}) = \mathbb{E}[Y_{\mathbf{x}} | Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x})] \end{cases} \quad (7)$$

3.3 Probabilistic interpretation at inference time

Let us consider a trained rMCL model, such as the one described in the previous section. Following the theoretical interpretation of [34] and the motivation of solving the overconfidence issue, as

explained in Section 3.1, the output objectives of the model can be summarized in (7). Although it is possible, assuming that those properties are verified, to deduce the following law of total expectation

$$\mathbb{E}[Y_{\mathbf{x}}] = \sum_{k=1}^K \gamma_{\theta}^k(\mathbf{x}) f_{\theta}^k(\mathbf{x}), \quad (8)$$

the above quantity may not be informative enough, especially if the law of $Y_{\mathbf{x}} \sim p(\mathbf{y} | \mathbf{x})$ is multimodal. Indeed, to be able to derive a complete probabilistic interpretation, we still need to characterize how rMCL approximates the full conditional density of $Y_{\mathbf{x}}$. For that purpose, it is sufficient to fix a law π_k for $Y_{\mathbf{x}}$ within each Voronoi cell k to accurately represent the predicted distribution in each region of the output space as per the rMCL model, as shown in Proposition 1.

Proposition 1 (Probabilistic interpretation of rMCL). *With the above notations, let us consider a multi-hypothesis model with properties of (7). Let us furthermore assume that*

$$Y_{\mathbf{x}}^k \sim \pi_k(\cdot | \mathbf{x}) \quad (9)$$

i.e., the law $Y_{\mathbf{x}}^k$ of $Y_{\mathbf{x}}$ conditioned by the event $\{Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x})\}$ is described by π_k for each input $\mathbf{x} \in \mathcal{X}$ with constraint $\mathbb{E}[Y_{\mathbf{x}}^k] = f_{\theta}^k(\mathbf{x})$. Then, for each measurable set $A \subseteq \mathcal{Y}$,

$$\mathbb{P}(Y_{\mathbf{x}} \in A) = \int_{\mathbf{y} \in A} \sum_{k=1}^K \gamma_{\theta}^k(\mathbf{x}) \pi_k(\mathbf{y} | \mathbf{x}) d\mathbf{y}. \quad (10)$$

Proof. We have

$$\begin{aligned} \mathbb{P}(Y_{\mathbf{x}} \in A) &= \sum_k \mathbb{P}(Y_{\mathbf{x}} \in A \cap \mathcal{Y}^k(\mathbf{x})) && \text{(by sigma-additivity)} \\ &= \sum_k \mathbb{P}(Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x})) \mathbb{P}(Y_{\mathbf{x}} \in A | Y_{\mathbf{x}} \in \mathcal{Y}^k(\mathbf{x})) && \text{(by Bayes' theorem)} \\ &= \sum_k \gamma_{\theta}^k(\mathbf{x}) \int_{\mathbf{y} \in A} \pi_k(\mathbf{y} | \mathbf{x}) d\mathbf{y} && \text{(by (9)). } \square \end{aligned}$$

Following the principle of maximum entropy [38], the uniform distribution inside each cell $Y_{\mathbf{x}}^k \sim \mathcal{U}(\mathcal{Y}^k(\mathbf{x}))$ is the least-informative distribution, assuming that the output space \mathcal{Y} is with finite volume and the hypotheses lie in the geometric center of their associated cell. With this prior, the predicted distribution can be interpreted as

$$\hat{p}(\mathbf{y} | \mathbf{x}) = \sum_{k=1}^K \gamma_{\theta}^k(\mathbf{x}) \frac{\mathbf{1}(\mathbf{y} \in \mathcal{Y}^k(\mathbf{x}))}{\mathcal{V}(\mathcal{Y}^k(\mathbf{x}))}, \quad (11)$$

where $\mathcal{V}(\mathcal{Y}^k(\mathbf{x})) \triangleq \int_{\mathbf{y} \in \mathcal{Y}^k(\mathbf{x})} d\mathbf{y}$ is the volume of $\mathcal{Y}^k(\mathbf{x})$. Similarly, if we model the output distribution as a mixture of Dirac deltas such that in each cell k , $\gamma_{\theta}^k(\mathbf{x}) > 0 \Rightarrow Y_{\mathbf{x}}^k \sim \delta_{f_{\theta}^k(\mathbf{x})}$, then

$$\hat{p}(\mathbf{y} | \mathbf{x}) = \sum_{k=1}^K \gamma_{\theta}^k(\mathbf{x}) \delta_{f_{\theta}^k(\mathbf{x})}(\mathbf{y}). \quad (12)$$

3.4 Toy example

We build upon the toy example formalized in [34, §4.1] to validate the proposed algorithm and its interpretation. In particular, we seek to illustrate how rMCL handles the overconfidence problem in a regression task. The toy problem involves the prediction of a 2D distribution $p(\mathbf{y} | t)$ based on scalar input $t \in \mathcal{X}$ (the *time*), where $\mathcal{X} = [0, 1]$. The training dataset is constructed by selecting for each random input $t \sim \mathcal{U}([0, 1])$ one of the four sections $S_1 = [-1, 0) \times [-1, 0)$, $S_2 = [-1, 0) \times [0, 1]$, $S_3 = [0, 1] \times [-1, 0)$, $S_4 = [0, 1] \times [0, 1]$, with probabilities $p(S_1) = p(S_4) = \frac{1-t}{2}$, $p(S_2) = p(S_3) = \frac{t}{2}$. Whenever a region is selected, a point is then sampled uniformly in this region. To slightly increase the complexity of the dataset, it has been chosen to enable the sampling of two (instead of one) target training samples with probability $q(t)$ for each input t . Here, q is a piece-wise affine function with $q(0) = 1$, $q(\frac{1}{2}) = 0$, and $q(1) = 1$.

We compared the behavior of our proposed approach rMCL against sMCL and independent ensembles (IE) using a 3-layer perceptron backbone with 256 hidden units and 20 output hypotheses, and ℓ_2 as the underlying distance. For rMCL, we used scoring weight $\beta = 1$. In addition to the full-fledged

rMCL, we assess a variant denoted ‘rMCL*’ where a single negative hypothesis is uniformly selected during training. For independent ensembles, we trained with independent random initialization 20 single-hypothesis models, the predictions of which being combined to be comparable to multi-hypothesis models. The training processes were executed until the convergence of the training loss. We report the checkpoints for which the validation loss was the lowest. All networks were optimized using Adam [24]. We show in Figure 1 (top right) test predictions of the classical sMCL compared to those of the rMCL model. In visualizing the predictions where the ground-truth distribution is known, we observe empirically that the *centroidal* tessellation property, $\mathbb{E}[Y_{\mathbf{x}}^k] = f_{\theta}^k(\mathbf{x})$, is verified with good approximation (see red points Figure 1, bottom right). We otherwise notice the *overconfidence* problem of the sMCL model in low-density zones when the output distribution is multimodal by looking at the output prediction, *e.g.*, for $t = 0.1$ and $t = 0.9$ in Figure 1 where the sMCL predictions fail in low-density zones. In contrast, rMCL solves this issue, assigning to each Voronoi cell a score that approximates the probability mass of the ground-truth distribution in this zone (see Figure 1, bottom). Furthermore, the independent ensembles (triangles) suffer from a collapse issue around the conditional mean of the ground-truth distribution. This behavior is expected as the minimizer of the continuous formulation of the risk (4) on the whole output space is the conditional expectation of the output given the input.




Figure 1: **Comparisons on a 2D regression toy problem.** Comparing sMCL with standard WTA (‘sMCL’, pink diamonds), proposed rMCL (blue circles), a variant of rMCL with a single negative hypothesis chosen uniformly at training (‘rMCL*’, purple crosses), and Independent Ensemble (‘IE’, yellow triangles). (Left) Test EMD for different time t ; (Right) Inference at $t = 0.1, 0.6, 0.9$ and Voronoi tessellations generated by the hypotheses. Green points are samples from the ground-truth distribution at the corresponding time step. Cells conditional mean (red points) matches with the predicted hypotheses. The score predicted in each cell, displayed as color saturation of the blue circles, as per scale on the right, approximates the corresponding proportion of points. At $t = 0.1$ and 0.9 , we can observe how rMCL tackles well the overconfidence issue in the low-density zones.

From a quantitative perspective, as demonstrated in Fig. 1, we can evaluate the effectiveness of rMCL in addressing the *overconfidence* issue compared to the classical sMCL. To ensure a fair comparison with sMCL settings, we have opted for using the Earth Mover’s Distance (EMD) metric, which measures the discrepancy between the ground truth and predicted distributions. Predicted distributions are considered mixtures of Dirac deltas as in (12) for the metric computation, with sMCL predictions assigning uniform mode weights. As depicted in Fig. 1, rMCL outperforms sMCL when the target distribution is multimodal, *i.e.*, for extreme values of t . However, in the unimodal scenario, the performances of rMCL and sMCL are fairly similar. Furthermore, the ‘rMCL*’ variant (depicted by crosses) that is lighter regarding memory consumption exhibits performance comparable to the standard rMCL. Additional assessments, which are detailed in Appendix B.2, have demonstrated the robustness of the rMCL when corrupting the training dataset with outliers in the output space.

4 Experiments with audio data

4.1 Sound source localization as a conditional distribution estimation problem

Sound source localization (SSL) is the task of predicting the successive angular positions of audio sources from a multichannel input audio clip $\mathbf{x} \in \mathcal{X}$ [15]. This problem can be cast either as a regression problem (estimating the continuous positions of the sources, *e.g.*, [37]) or a classification one (segmenting the output space into zones in which we wish to detect a source, *e.g.*, [2]). While not requiring the number of sources to be known in advance, the classification approach suffers from the lack of balanced data and limited spatial resolution. On the other hand, the regression approach enables off-grid localization but typically suffers from the source permutation problem [40]. Assuming that the maximum number of overlapping sources M is known, a solution for handling this multi-target regression problem is to predict, at each time frame t depending on the chosen output resolution, a vector accounting for the sources’ activity $\mathbf{a}_t \in \{0, 1\}^M$, as well as azimuth and elevation angles $\hat{\phi}_t \in \mathbb{R}^M$ and $\hat{\vartheta}_t \in \mathbb{R}^M$. A model can be trained with a permutation invariant training (PIT) approach [1, 37, 45] using an optimization criterion of the form

$$\mathcal{L}(\theta) = \sum_t \min_{\sigma \in \mathcal{S}_M} \ell_{\text{CE}}(\sigma(\hat{\mathbf{a}}_t), \mathbf{a}_t) + \ell_g\left((\sigma(\hat{\phi}_t), \sigma(\hat{\vartheta}_t)), (\phi_t, \vartheta_t)\right), \quad (13)$$

where ℓ_{CE} and ℓ_g correspond respectively to a cross-entropy term and a geometrical loss, the latter being computed only for active sources indexes. \mathcal{S}_M is the set of permutations of M elements and the notation $\sigma(\mathbf{z})$ stands for the M -dim vector \mathbf{z} with its components permuted according to σ . In the following, we will denote \mathbf{Y}_t the set of source positions at time t .

With the distribution learning mindset, this task can be seen as an attempt to estimate, at each time step, the sound source position distribution $p(\mathbf{y} | \mathbf{x})$, which can be viewed as a Dirac mixture, $p(\mathbf{y} | \mathbf{x}) \propto \sum_{\mathbf{y}_t \in \mathbf{Y}_t} \delta_{\mathbf{y}_t}(\mathbf{y})$, if we suppose that the targets are point-wise, with another Dirac mixture representing the predicted active modes at predicted positions. Therefore, a natural way to evaluate such SSL regression models is to solve the linear assignment problem, *e.g.*, using Hungarian matching with spherical distance as an underlying metric. To handle more general distributions, we propose to generalize the metric used to the Earth Mover’s Distance (see Sec. 4.2).

The rMCL framework is well suited to SSL as it allows one to benefit from both the advantages of the regression and classification viewpoints in the same spirit as [41]. There is no need for prior knowledge of the number of sources, and it avoids challenges related to imbalanced spatial positions and the source permutation problem of (13). This method comes at a low computational cost regarding added parameters when opting for a low-level representation shared by the hypotheses and scores heads (see Sec. 3.2). Furthermore, it allows for producing a heat map of the sound sources’ positions with a probabilistic prediction, which could otherwise account for their spatial dispersion depending on the chosen law π_k selected in the Voronoi cells. Modeling the sound sources as point sources, a delta law will be selected following the formulation of Proposition 1.

4.2 Experimental setup

Datasets. We conducted our experiments on increasingly complex SSL datasets originally introduced by Adavanne et al. [1]: i) **ANSYN**, derived from **DCASE** 2016 Sound Event Detection challenge with spatially localized events in anechoic conditions and ii) **RESYN**, a variant of ANSYN in reverberant conditions. Each dataset is divided into three sub-datasets depending on the maximum number of overlapping events (1, 2, or 3, denoted as D1, D2, and D3). These sub-datasets are further divided into three splits (including one for testing). Each training split contains 300 recordings (30 s), of which 60 are reserved for validation. Preprocessing is detailed in Appendix A.2. For each experiment, we used all training splits from D1, D2, and D3. Evaluation was then conducted for each test set based on the overlapping levels. Moreover, we present results from additional SSL datasets, including **REAL** [1] and **DCASE19** [3], in Appendix B.3.

Metrics. To assess the performance of the different methods, we employed the following metrics:

- The ‘Oracle’ error (\downarrow): $\mathcal{O}(\mathbf{x}_n, \mathbf{Y}_n) = \frac{1}{|\mathbf{Y}_n|} \sum_{\mathbf{y}_n \in \mathbf{Y}_n} \min_{k \in [1, K]} d(f_{\theta}^k(\mathbf{x}_n), \mathbf{y}_n)$.
- The *Earth Mover’s Distance* (EMD \downarrow): also known as the 1st Wasserstein distance, it is a distance between the predicted distribution $\hat{p}(\mathbf{y} | \mathbf{x}_n) = \sum_{k=1}^K \gamma_{\theta}^k(\mathbf{x}_n) \delta_{f_{\theta}^k(\mathbf{x}_n)}(\mathbf{y})$

and the ground-truth one $p(\mathbf{y} | \mathbf{x}_n) = \frac{1}{|\mathcal{Y}_n|} \sum_{\mathbf{y}_n \in \mathcal{Y}_n} \delta_{\mathbf{y}_n}(\mathbf{y})$. The EMD solves the optimization problem $W_1(p(\cdot | \mathbf{x}_n), \hat{p}(\cdot | \mathbf{x}_n)) = \min_{\psi \in \Psi} \sum_{k=1}^K \sum_{\mathbf{y}_n \in \mathcal{Y}_n} \psi_{k,n} d(f_{\theta}^k(\mathbf{x}_n), \mathbf{y}_n)$, where $\psi = (\psi(f_{\theta}^k(\mathbf{x}_n), \mathbf{y}_n))_{k,n}$ is a transport plan and Ψ is the set of all valid transport plans [21].

While the oracle error assesses the quality of the *best* hypothesis predicted for each target, the EMD metric looks at *all* hypotheses. It provides insight into the overall consistency of the predicted distribution. The EMD also penalizes the overconfidence issue of the WTA with sMCL as described in Sec. 3.4. In order to fit the directional nature of the sound source localization task, these metrics are computed in a spherical space equipped with distance $d(\hat{\mathbf{y}}, \mathbf{y}) = \arccos[\sin(\hat{\phi}) \sin(\phi) + \cos(\hat{\phi}) \cos(\phi) \cos(\vartheta - \hat{\vartheta})]$, where $(\hat{\phi}, \hat{\vartheta})$ and (ϕ, ϑ) are source positions $\hat{\mathbf{y}}$ and \mathbf{y} , respectively, expressed as azimuth and elevation angles.

Neural network backbone. We employed the CRNN architecture of SeldNet [1] as a backbone, which we modified to enable multi-hypothesis predictions. This adaptation involves adjusting the output format and duplicating the last fully-connected layers for hypothesis and scoring heads. As only the prediction heads are adjusted, the number of parameters added to the architecture is negligible. Refer to Appendix A.2 for additional architecture and training details.

Baselines. We compared the proposed rMCL with several baselines, each utilizing the same feature-extractor backbone as previously described. The baselines include a Permutation Invariant Training variant (‘PIT variant’) proposed in [37] for SSL, the conventional WTA setup (‘WTA, 5 hyp.’) and its single hypothesis variant (‘WTA, 1 hyp.’), its ε -relaxed version with $\varepsilon = 0.5$ (‘ ε -WTA, 5 hyp.’) and its top- n variant with $n = 3$ (‘top- n -WTA, 5 hyp.’), as well as independent ensembles (‘IE’). To ensure a fair comparison, we used the same number of 5 hypotheses for the WTAs, sufficient to cover the maximum of 3 sound sources in the dataset (refer to Sec. 4.5 and the sensitivity study in Table 4). For the single hypothesis WTA, the head is updated by only considering the best target, as this fares better than using one update per target. IE was constructed from five such single hypothesis WTA models trained independently with random initialization.

4.3 Comparative evaluation of rMCL in SSL

Table 1: **Source localization in anechoic conditions.** Average scores (\pm standard deviation) on ANSYN dataset for PIT, IE, various WTA-based methods, and proposed rMCL. The ‘EMD’ evaluates all the hypotheses jointly, while the ‘Oracle’ error only looks at the best hypotheses. D1 corresponds to the single-target setting, while D2 and D3 consider up to 2 and 3 targets, respectively.

Dataset: ANSYN	EMD D1	EMD D2	EMD D3	Oracle D1	Oracle D2	Oracle D3
PIT variant	6.22 \pm 0.80	14.65 \pm 1.22	23.41 \pm 1.39	3.58 \pm 0.46	10.58 \pm 0.91	18.10 \pm 1.04
IE (5 members)	4.05 \pm 0.48	21.64 \pm 2.29	34.34 \pm 2.37	1.24 \pm 0.24	16.91 \pm 2.05	28.82 \pm 2.25
WTA, 1 hyp.	3.97 \pm 0.55	24.69 \pm 2.72	39.66 \pm 2.67	3.97 \pm 0.55	24.69 \pm 2.72	39.66 \pm 2.67
WTA, 5 hyp.	48.22 \pm 1.78	44.41 \pm 1.25	41.83 \pm 0.96	3.56 \pm 0.39	6.53 \pm 0.44	10.44 \pm 0.58
top- n -WTA, 5 hyp.	5.14 \pm 0.80	18.09 \pm 1.32	25.12 \pm 1.30	3.33 \pm 0.46	7.48 \pm 0.79	13.54 \pm 0.96
ε -WTA, 5 hyp.	5.51 \pm 0.66	19.20 \pm 1.78	28.39 \pm 1.68	3.62 \pm 0.57	10.86 \pm 1.26	17.44 \pm 1.22
rMCL, 5 hyp.	7.04 \pm 0.58	13.87 \pm 0.99	20.76 \pm 1.04	3.85 \pm 0.46	7.16 \pm 0.67	11.29 \pm 0.78

Table 2: **Source localization in reverberant conditions.** Results on RESYN dataset, with same table layout as in Table 1.

Dataset: RESYN	EMD D1	EMD D2	EMD D3	Oracle D1	Oracle D2	Oracle D3
PIT variant	10.53 \pm 0.90	25.09 \pm 2.14	35.05 \pm 1.98	4.89 \pm 0.55	15.92 \pm 1.36	24.95 \pm 1.62
IE (5 members)	8.24 \pm 1.02	26.65 \pm 2.49	38.7 \pm 2.62	3.77 \pm 0.61	20.95 \pm 2.27	32.85 \pm 2.48
WTA, 1 hyp.	8.32 \pm 1.28	29.26 \pm 2.85	43.25 \pm 2.94	8.32 \pm 1.28	29.26 \pm 2.85	43.25 \pm 2.94
WTA, 5 hyp.	57.88 \pm 1.71	51.74 \pm 1.36	47.38 \pm 1.26	5.81 \pm 0.58	9.46 \pm 0.71	13.33 \pm 0.69
top- n -WTA, 5 hyp.	42.74 \pm 2.86	37.25 \pm 1.88	36.48 \pm 1.40	6.21 \pm 0.79	11.02 \pm 1.00	17.32 \pm 1.11
ε -WTA, 5 hyp.	8.84 \pm 1.09	27.3 \pm 2.52	38.43 \pm 2.42	6.48 \pm 0.95	20.54 \pm 2.33	30.18 \pm 2.15
rMCL, 5 hyp.	12.14 \pm 1.12	24.45 \pm 1.91	32.28 \pm 1.85	5.74 \pm 0.66	10.5 \pm 0.87	14.6 \pm 0.87

We present in Tables 1 and 2 the results of rMCL and baselines in the anechoic (ANSYN) and reverberant (RESYN) conditions, respectively. The consistency of the results should be noted. For each model, the metrics tend to improve as the number of sources decreases and, under similar conditions, results are marginally lower for more challenging datasets. Unsurprisingly, the single hypothesis approach exhibits strong performances in the unimodal cases (D1), but performs significantly worse in the multimodal cases (D2 and D3). Ensembling (IE) improves its performance but still displays the lack of diversity already exposed in Section 3.4. As for the 5-hypothesis WTA, in its original form, it is ill-suited for multi-target regression due to the overconfidence issue: despite a strong oracle metric, the EMD results are very poor. On the other hand, rMCL surpasses its competitors and shows the best EMD in every multimodal setting. It also consistently obtains the second-best oracle error, only slightly above WTA, while not suffering from overconfidence.

We also present results for REAL and DCASE19 in Appendix B.3, that display similar trends. We did not observe the collapse issue (see Section 3.1) in our settings, neither with WTA nor with the proposed rMCL model. We believe it is solved in practice by the variability of the data samples in the stochastic optimization during training [19]; refer to Appendix B.1 for further discussions.

4.4 Combining rMCL with other WTA variants

Table 3: **Combining proposed rMCL with other WTA variants.** Average scores (\pm standard deviation) for source localization in anechoic (top) and reverberant (bottom) conditions.

Dataset: ANSYN	EMD D1	EMD D2	EMD D3	Oracle D1	Oracle D2	Oracle D3
rMCL, 5 hyp.	7.04 ± 0.58	13.87 ± 0.99	20.76 ± 1.04	3.85 ± 0.46	7.16 ± 0.67	11.29 ± 0.78
top- n -rMCL, 5 hyp.	5.46 ± 0.62	13.88 ± 1.06	21.45 ± 1.10	4.2 ± 0.55	8.04 ± 0.74	13.72 ± 0.89
ε -rMCL, 5 hyp.	5.89 ± 0.69	12.13 ± 0.98	19.95 ± 1.16	3.6 ± 0.54	8.76 ± 0.89	14.47 ± 1.02
Dataset: RESYN	EMD D1	EMD D2	EMD D3	Oracle D1	Oracle D2	Oracle D3
rMCL, 5 hyp.	12.14 ± 1.12	24.45 ± 1.91	32.28 ± 1.85	5.74 ± 0.66	10.5 ± 0.87	14.6 ± 0.87
top- n -rMCL, 5 hyp.	9.3 ± 1.15	23.81 ± 2.22	33.33 ± 2.06	6.99 ± 0.85	13.43 ± 1.40	19.72 ± 1.29
ε -rMCL, 5 hyp.	8.64 ± 1.03	22.82 ± 2.12	32.47 ± 1.97	6.08 ± 0.91	18.39 ± 2.07	26.92 ± 1.94

As noted in Section 3.2, rMCL’s improvements are, in theory, orthogonal to that of other WTA variants. In this section, we combine the rMCL with top- n -WTA and ε -WTA approaches into respectively top- n -rMCL and ε -rMCL. The results are shown in Tables 3. We note that while top- n rMCL does not provide significant improvement to rMCL, ε -rMCL, on the other hand improves EMD scores at the expense of increased oracle error. We observe similar effects when this method is applied to WTA, indicating that there are indeed some additive effects of ε -WTA and rMCL. Also note that with regards to Tables 1 and 2, all the proposed variants still get better EMD than all competitors under multimodal conditions, showing the robustness of our method.

4.5 Effect of the number of hypotheses

The impact of varying the number of hypotheses on the performance of the rMCL model is presented in Table 4 and Figure B.4. First and foremost, we notice the anticipated trend of the oracle metric improving as the number of hypotheses increases. Concerning the EMD metric, the single hypothesis model, which avoids errors from negative hypotheses, is most effective in handling unimodal

Table 4: **Sensitivity analysis.** Effect of the number of hypotheses on the performance of rMCL.

Dataset: ANSYN	EMD D1	EMD D2	EMD D3	Oracle D1	Oracle D2	Oracle D3
PIT variant	6.22 ± 0.80	14.65 ± 1.22	23.41 ± 1.39	3.58 ± 0.46	10.58 ± 0.91	18.10 ± 1.04
WTA, 1 hyp.	3.97 ± 0.55	24.69 ± 2.72	39.66 ± 2.67	3.97 ± 0.55	24.69 ± 2.72	39.66 ± 2.67
rMCL, 3 hyp.	9.89 ± 0.95	14.37 ± 0.91	20.96 ± 1.03	5.65 ± 0.73	8.6 ± 0.63	13.42 ± 0.86
rMCL, 5 hyp.	7.04 ± 0.58	13.87 ± 0.99	20.76 ± 1.04	3.85 ± 0.46	7.16 ± 0.67	11.29 ± 0.78
rMCL, 10 hyp.	9.14 ± 0.76	15.2 ± 0.84	21.28 ± 0.96	2.94 ± 0.35	4.76 ± 0.39	7.54 ± 0.50
rMCL, 20 hyp.	9.13 ± 0.71	16.04 ± 0.84	22.55 ± 0.87	2.06 ± 0.22	3.61 ± 0.30	5.83 ± 0.37

distributions. In the case of the EMD metric applied to multimodal distributions, we observe that multiple hypothesis models improve the results, but excessively increasing the number of hypotheses may marginally degrade performances. A further study could examine this phenomenon in detail, which may be related to the expressive power of the hypothesis heads or the accumulation of errors in score heads predictions.

5 Discussion and limitations

In the audio experiments, the performance of rMCL was found to be affected by the number of hypotheses to tune depending on the complexity of the dataset. Moreover, as with the other variants of WTA, fixing the overconfidence issue with rMCL slightly degrades the performance of the best hypothesis (oracle error) for a reason that is yet to be determined. Otherwise, while ε -WTA behaves as expected with rMCL, trading off the quality of the best hypotheses for overall performance, top- n -WTA does not exhibit the same behavior. This discrepancy warrants further investigation.

The probabilistic interpretation of rMCL, as presented in Algorithm 1 and in Sec. 3.3, states that the different hypotheses would be organized in an optimal partition of the output space, forming a Voronoi tessellation. Ideally, each hypothesis would capture a region of the distribution, and the scores representing how likely this zone would activate in a given context. This interpretation remains theoretically valid whenever train and test examples are sampled from the same joint distribution $p(\mathbf{x}, \mathbf{y})$. However, in realistic settings, the hypotheses might not adhere to meaningful regions, but this could be controlled and evaluated, provided that an external validation dataset is available. In future work, we intend to use calibration techniques [16, 39] to identify and alleviate these issues.

To sum up, this paper proposes a new Multiple Choice Learning variant, suitable for tackling the MCL overconfidence problem in regression settings. Our method is based on a learned scoring scheme that handles situations where a set of targets is available for each input. Furthermore, we propose a probabilistic interpretation of the model, and we illustrate its relevance with an evaluation on synthetic data. Its practical usefulness is also demonstrated in the context of point-wise sound source localization. Further work could include a specific study about the specialization of the predictors, and the validation of the proposed algorithm in increasingly complex real-world datasets.

Acknowledgments

This work was funded by the French Association for Technological Research (ANRT CIFRE contract 2022-1854). We would like to thank the reviewers for their valuable feedback.

References

- [1] Sharath Adavanne, Archontis Politis, Joonas Nikunen, and Tuomas Virtanen. Sound event localization and detection of overlapping sources using convolutional recurrent neural networks. *IEEE Journal of Selected Topics in Signal Processing*, 13(1):34–48, 2018. 7, 8, 14, 15, 16, 17
- [2] Sharath Adavanne, Archontis Politis, and Tuomas Virtanen. Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 1462–1466. IEEE, 2018. 7
- [3] Sharath Adavanne, Archontis Politis, and Tuomas Virtanen. A multi-room reverberant dataset for sound event localization and detection. *arXiv preprint arXiv:1905.08546*, 2019. 7, 17
- [4] Kazi Md Rokibul Alam, Nazmul Siddique, and Hojjat Adeli. A dynamic ensemble learning algorithm for neural networks. *Neural Computing and Applications*, 32:8675–8690, 2020. 2
- [5] Samuel S Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):5–18, 2004. 1
- [6] Mike Brodie, Chris Tensmeyer, Wes Ackerman, and Tony Martinez. Alpha model domination in multiple choice learning. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 879–884. IEEE, 2018. 2

- [7] Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural safety*, 31(2):105–112, 2009. 2
- [8] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000. 2
- [9] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM review*, 41(4):637–676, 1999. 3
- [10] Michael Firman, Neill DF Campbell, Lourdes Agapito, and Gabriel J Brostow. Diversenet: When one right answer is not enough. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5598–5607, 2018. 2, 3
- [11] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019. 2
- [12] Yarin Gal et al. Uncertainty in deep learning. 2016. 2
- [13] Nuno Cruz Garcia, Sarah Adel Bargal, Vitaly Ablavsky, Pietro Morerio, Vittorio Murino, and Stan Sclaroff. Distillation multiple choice learning for multimodal action recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2755–2764, 2021. 2
- [14] Alex Graves. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011. 2
- [15] Pierre-Amaury Grumiaux, Srđan Kitić, Laurent Girin, and Alexandre Guérin. A survey of sound source localization with deep learning methods. *The Journal of the Acoustical Society of America*, 152(1):107–151, 2022. 7
- [16] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. 10
- [17] Abner Guzman-Rivera, Dhruv Batra, and Pushmeet Kohli. Multiple choice learning: Learning to produce multiple structured outputs. *Advances in neural information processing systems*, 25, 2012. 1, 2, 3
- [18] Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE transactions on pattern analysis and machine intelligence*, 12(10):993–1001, 1990. 2
- [19] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 652–667, 2018. 2, 9, 16
- [20] Ehsan Imani and Martha White. Improving regression performance with distributional losses. In *International conference on machine learning*, pages 2157–2166. PMLR, 2018. 1
- [21] Leonid V Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pages 199–201, 1942. 8
- [22] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017. 2
- [23] Youngwook Kim, Jae Myung Kim, Zeynep Akata, and Jungwoo Lee. Large loss matters in weakly supervised multi-label classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14156–14165, 2022. 2
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6, 14
- [25] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 15

- [26] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [27] Kimin Lee, Changho Hwang, KyoungSoo Park, and Jinwoo Shin. Confident multiple choice learning. In *International Conference on Machine Learning*, pages 2014–2023. PMLR, 2017. [2](#)
- [28] Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David Crandall, and Dhruv Batra. Why m heads are better than one: Training a diverse ensemble of deep networks. *arXiv preprint arXiv:1511.06314*, 2015. [2](#)
- [29] Stefan Lee, Senthil Purushwalkam Shiva Prakash, Michael Cogswell, Viresh Ranjan, David Crandall, and Dhruv Batra. Stochastic multiple choice learning for training diverse deep ensembles. *Advances in Neural Information Processing Systems*, 29, 2016. [1](#), [2](#), [3](#)
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. [15](#)
- [31] David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992. [2](#)
- [32] Osama Makansi, Eddy Ilg, Ozgun Cicek, and Thomas Brox. Overcoming limitations of mixture density networks: A sampling and fitting framework for multimodal future prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7144–7153, 2019. [1](#), [2](#), [16](#)
- [33] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *The Artificial Intelligence Review*, 42(2):275, 2014. [3](#)
- [34] Christian Rupprecht, Iro Laina, Robert DiPietro, Maximilian Baust, Federico Tombari, Nassir Navab, and Gregory D Hager. Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In *Proceedings of the IEEE international conference on computer vision*, pages 3591–3600, 2017. [1](#), [2](#), [3](#), [4](#), [5](#), [14](#)
- [35] Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1041–1044, 2014. [17](#)
- [36] Christopher Schymura, Benedikt Bönninghoff, Tsubasa Ochiai, Marc Delcroix, Keisuke Kinoshita, Tomohiro Nakatani, Shoko Araki, and Dorothea Kolossa. PILOT: Introducing transformers for probabilistic sound event localization. In *Interspeech 2021*, pages 2117–2121. ISCA, 2021. [14](#), [15](#), [16](#)
- [37] Christopher Schymura, Tsubasa Ochiai, Marc Delcroix, Keisuke Kinoshita, Tomohiro Nakatani, Shoko Araki, and Dorothea Kolossa. Exploiting attention-based sequence-to-sequence architectures for sound event localization. In *2020 28th European Signal Processing Conference (EUSIPCO)*, pages 231–235. IEEE, 2021. [7](#), [8](#), [14](#), [15](#), [16](#)
- [38] John Shore and Rodney Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on information theory*, 26(1):26–37, 1980. [5](#)
- [39] Hao Song, Tom Diethe, Meelis Kull, and Peter Flach. Distribution calibration for regression. In *International Conference on Machine Learning*, pages 5897–5906. PMLR, 2019. [10](#)
- [40] Aswin Shanmugam Subramanian, Chao Weng, Shinji Watanabe, Meng Yu, and Dong Yu. Deep learning based multi-source localization with source splitting and its effectiveness in multi-talker speech recognition. *Computer Speech & Language*, 75:101360, 2022. [7](#)
- [41] Harshavardhan Sundar, Weiran Wang, Ming Sun, and Chao Wang. Raw waveform based end-to-end deep convolutional network for spatial localization of multiple acoustic sources. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4642–4646. IEEE, 2020. [7](#)

- [42] Kai Tian, Yi Xu, Shuigeng Zhou, and Jihong Guan. Versatile multiple choice learning and its application to vision computing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6349–6357, 2019. [2](#), [4](#)
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [15](#)
- [44] Zehan Wang, Sihong Zhou, Yuyao Huang, and Wei Tian. Dsmcl: Dual-level stochastic multiple choice learning for multi-modal trajectory prediction. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE, 2020. [1](#)
- [45] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 241–245. IEEE, 2017. [7](#)
- [46] Zhi-Hua Zhou, Jianxin Wu, and Wei Tang. Ensembling neural networks: many could be better than all. *Artificial intelligence*, 137(1-2):239–263, 2002. [2](#)
- [47] Xiaoyan Zhu, Jiaxuan Li, Jingtao Ren, Jiayin Wang, and Guangtao Wang. Dynamic ensemble learning for multi-label classification. *Information Sciences*, 623:94–111, 2023. [2](#)

A Experimental setup

A.1 Synthetic data experiments

Dataset. As described in the main paper, the synthetic dataset (§3.4) is based on [34, §4.1]. The complexity of the dataset was increased by allowing two (instead of one) target training samples at each time step. The probability $q(t)$ of sampling two targets instead of a single one at a time t is chosen as the piece-wise affine function such that $q(0) = q(1) = 1$ and $q(\frac{1}{2}) = 0$. This modification aimed to demonstrate how rMCL handles multiple targets sampled for each input during training.

Architectures. We use a three-layer perceptron backbone with 20 output hypotheses and a ReLU activation function for each layer and 256 hidden units. The multi-hypothesis splitting is carried out at the final layer stage. The scoring heads of the rMCL model receive the same representation as the hypothesis heads. The ensemble members have the same backbone but use a single hypothesis at the output stage. The rMCL* model has the same architecture as rMCL but utilizes a different training loss (see Section 3.2).

Training details. For each model, the training was performed on 100,000 training samples and 25,000 validation samples with a batch size of 1024. We employed 20 training epochs with Adam optimizer [24]. The checkpoint retained corresponded to the one with the lowest validation loss. For the *Stochastic Multiple Choice Learning* (sMCL) model, the multi-target version of the Winner-takes-all loss was used. The rMCL and rMCL* models were trained with unit scoring loss weight ($\beta = 1$, see §3.2), but the rMCL* training differed in that it only updated one negative hypothesis (compared to all in standard rMCL). When trained in this manner, the rMCL* model can be considered a more memory-efficient version of the proposed rMCL. The Independent Ensemble (IE) members were trained with a single target update (in the sMCL, the best hypothesis is updated for each target) as it resulted in a better fit to the data. The predictions from the IE members, trained with several initialization instances, were then stacked. The IE results were not plotted in Figure 1 for clarity and comparison purposes, as they were significantly worse (average Earth Mover’s Distance at test time for the IE: 0.62 ± 0.11) compared to the other evaluated models.

Evaluation details. Figure 1, §3.4 (left) displays the Earth Mover’s Distance (EMD) values using ℓ_2 underlying distance calculated for 50 equally spaced input t values in the comparative evaluation. At each time step and for each model, the EMD was computed between 1,000 samples taken from the ground-truth distribution, viewed as a mixture of Dirac deltas and the predicted hypothesis. The centroids in each cell (described in the right part of the figure) were computed using 35,000 samples from the ground-truth distribution for each input.

A.2 Audio data experiments

Datasets preprocessing. As indicated in §4.2, the experimental setup incorporates the ANSYN and RESYN datasets, which feature spatially localized events under anechoic and reverberant conditions respectively [1]. We used the first-order Ambisonics (FOA) format with four input audio channels. The events from 11 possible classes (Clearing throat, Coughing, Door knock, Door slam, Drawer, Human laughter, Keyboard, Keys put on a table, Page turning, Phone ringing and Speech), extracted from the DCASE16 Sound Event Detection dataset, were randomly placed in a spatial grid (see [1]). We adhered to the dataset preprocessing described by Schymura et al. in [37, 36]. The audio signals, with a sampling frequency of 44.1 kHz, were converted into 30-second files. From those files, non-overlapping chunks of 0.5 s were generated to be used as training inputs. Spectrogram computation was performed offline for saving computation, using Hann window with length 0.04 s used for Short Term Fourier Transform (STFT) estimation, with 50% overlapping frames and 2048 Fast Fourier Transform (FFT) bins. The input for the model comprised both amplitude and phase information, stacked channel-wise.

Architecture. We employed the SELDNet backbone [1]. After raw audio preprocessing, it accepts the spectrograms of fixed duration with phase information as input and returns localization output in the chosen output resolution (here, $T = 25$ output time steps were considered for each chunk). The processing includes several feature extraction modules (CNNs, Bi-directional GRUs layers) that generate a representation at each time step in the output resolution. These latent representations are then mapped to the output localization estimates through FC layers. To accommodate the MCL setup, the final FC layers were split into K FC heads, each producing a 2D output at each time step.

The rMCL architecture also includes full confidence heads at the output stage, each producing a scalar output in $[0, 1]$ through a sigmoid function. Therefore, the number of added parameters in the architecture is negligible (~ 4 k when using 5 output hypotheses for an architecture with ~ 1.5 M parameters).

Training details. The training setup is the same as the one used in [37, 36]. Unless specified otherwise, the training details correspond to those in [37], for which the [official code](#) was released, in particular, the implementation of the PIT variant used. The trainings were conducted using the AdamW optimizer [30], with a batch size of 128, an initial learning rate of 0.05, and following the scheduling scheme from [43]. The MCL models were trained using the multi-target version of the Winner-takes-all loss. The rMCL version was trained using confidence weight $\beta = 1$. As for the synthetic data experiments, the training of IE members was performed using different random seeds and single hypothesis loss with a single target update for each.

Evaluation details. As outlined in Section 4.2, the Oracle and EMD metrics with spherical distance were employed for evaluation purposes. The EMD metric is an extension to solve the assignment problem, often tackled using the Hungarian method [25]. Similar to prior studies (*e.g.*, [1, 37, 36]), the localization metric was computed solely for frames (the *active* frames) where at least one active source is present in the sound scene. For each test sample, the metrics were computed and averaged over the active frames among the output T frames. Furthermore, we computed the standard deviation from the average metric for each subsection of the test set (D1, D2, and D3) sample-wise, as presented in the Tables. The standard deviations of the metrics when performing the exact same experiments from different random states are also provided in Appendix B.3. On a separate note, the frame recall metric, which indicates the percentage of time frames in which the number of active sound sources was estimated correctly, is omitted in the results for the sake of conciseness. This is because the EMD already penalizes missing sources in the predictions. In the rMCL model, the number of sources in the sound scene can nevertheless be computed before the normalization described in §3.2 by summing the output scores $\sum_k \gamma_{\theta}^k(\mathbf{x})$.

Visualizations of rMCL outputs for input test samples from the [ANSYN](#) dataset are given in Figure A.1.




Figure A.1: **Qualitative comparisons.** Results for randomly chosen input audio clips (*different rows*) from the D3 test subset of the ANSYN dataset. The columns correspond to the temporal predictions at different time steps t . For each prediction (subplot), the abscissa and ordinates stand for azimuth and elevation angles, respectively (in degrees). We notice the competitive performance of the proposed rMCL model (with *shaded blue circles* for whose the score intensity is displayed in the colorbar) compared with the Permutation Invariant Training (PIT) approach (*orange squares*, baseline used for Tables 1 to 4) for predicting the positions of the targets (*green stars*).




Figure B.1: **Histogram of *winner hypotheses*' indexes at test time for a 20-hypothesis model trained on ANSYN.** The x-axis displays the hypothesis index, and the y-axis the fraction of test samples on which hypothesis k was selected. We observe a high entropy distribution, which shows that the collapse did not occur for this model.

A.3 Computation details

We utilized the **Hydra** library for experimental purposes within the **Pytorch** deep learning framework. Our coding was inspired by [32, 37, 36, 1].

Our experiments were conducted on NVIDIA A100 GPUs. The total computing resources used for this project, including failed experiments, amount to approximately 2,000 GPU hours.

B Further discussions

B.1 About the collapse problem

As highlighted in the main paper, the *collapse* issue in MCL refers to a theoretical situation where one (or a few) of the K hypotheses become dominant, *i.e.*, are almost always selected as a winner and receive the gradient update. In this situation, the other hypothesis heads are not updated, therefore shrinking the diversity of the predictions. As a way to measure the collapse phenomenon for trained model $f_\theta = (f_\theta^1, \dots, f_\theta^K)$ using a validation dataset $\mathcal{D} \triangleq \{(\mathbf{x}_n, \mathbf{Y}_n)\}$, one can compute for each k , the number of samples in \mathcal{D} for which hypothesis k is a *winner*, that is:

$$\mathcal{N}_k(\mathcal{D}) \triangleq |\{(\mathbf{x}_n, \mathbf{Y}_n) \in \mathcal{D} : \exists \mathbf{y} \in \mathbf{Y}_n, \mathbf{y} \in \mathcal{Y}^k(\mathbf{x}_n)\}|.$$

The negative entropy of the histogram values $\{\mathcal{N}_k(\mathcal{D})\}_k$ should therefore inform about the collapse level; if the histogram shows a wide diversity of the selected hypotheses (*i.e.*, a histogram with almost uniform values for each bin), then there is no collapse. On the opposite, if the histogram has only one non-zero bin, the collapse level is maximum.

In our experiments with audio data, we did not observe the collapse problem in practice, neither with WTA nor with the proposed rMCL model. To verify this, one can compute the histogram of the $\{\mathcal{N}_k(\mathcal{D})\}$ values as explained above. See Fig. B.1 for an example of visualization for a 20-hypotheses model trained on ANSYN (corresponding to the results of Table 4 in the paper). As mentioned in [19] (p.8), we believe this issue is, in practice, solved by the variability of the data samples and the training stochasticity.

B.2 Robustness in the presence of target outliers

This section aims to provide insights into how the proposed model can handle the presence of outliers in the output space, which may be critical in real-world datasets. Let's consider a setting where we have outliers in the training dataset, for instance, the toy use-case presented in the paper, where for




Figure B.2: **Robustness of rMCL to outliers.** Evaluation results on the toy example when corrupting the training dataset with outliers, modeled by a bivariate Cauchy distribution. (Left) Sorted values of the unnormalized scores predicted from the hypotheses scoring heads at $t = 0.6$. (Right) Inference at $t = 0.1, 0.6, 0.9$ with zoomed prediction in $[-1, 1]^2$ (top) and Voronoi tessellations in the full plane (bottom) with the same organization as in Figure 1.

each training example sampled, the probability of getting an outlier is $\rho \ll 1$. Then, whenever an outlier is sampled, one hypothesis will be pushed towards it with its associated score heads updated. As the training goes on, some of the hypotheses will manage the outlier samples; let's name them the 'outlier hypotheses'. Thanks to the proposed hypothesis scoring heads, the model will also learn the probability that an outlier hypothesis is chosen for a given training sample. Provided that the outlier likelihood is $\rho \ll 1$, the scoring heads will therefore prevent outlier hypotheses output from deteriorating the quality of the predicted distribution by rMCL. In Fig. B.2 an illustration of this phenomenon is proposed using a Cauchy distribution (we used $\rho = 0.02$). We notice the above-explained phenomenon, where the so-called outlier hypotheses account for the outlier samples, while the other hypotheses lie in the square $[-1, 1]^2$ representing the samples from the ground-truth distribution.

Provided that the probability of sampling an outlier ρ is small enough and the outliers are far enough from the ground-truth distribution to predict, the proposed rMCL model is therefore potentially robust to outliers. In this case, some specific hypotheses, namely the outlier hypotheses, will be assigned to them, preventing the non-outlier hypotheses from being heavily affected. At inference time, it will indeed be possible to set to zero the very low-score hypotheses given an arbitrary threshold so that the outlier hypotheses are not taken into account.

B.3 Further results on audio datasets


We provide in this section further benchmarks of the method on sound source localization (SSL) datasets. Additionally to the results presented in Tables 1-4, where the mean and sample-wise standard deviation of the metrics on ANSYN and RESYN datasets are computed, we provide in Figures B.3 and B.4 further results considering the statistics of the metrics after several runs from different random states, and also including the REAL [1] and DCASE19 [3] datasets. In those datasets, the maximum number of overlapping events is respectively three and two.

Further datasets. In contrast to ANSYN and RESYN datasets which employ simulated Room Impulse Responses (RIRs) for audio spatialization, the REAL and DCASE19 datasets utilize RIRs recorded in real sound scenes. These were captured with a Spherical Microphone Array [1, 3]. Specifically, the recordings took place in various indoor settings inside a university. The REAL and DCASE19 RIRs were convoluted respectively with sound events from the UrbanSound8k [35] and DCASE16 datasets to achieve spatialization. Additionally, ambient noise was collected from the environment of the RIR recordings in the DCASE19 dataset. The same dataset pre-processing and sub-splitting process as presented in ANSYN and RESYN datasets was employed in REAL. For DCASE19, the four development dataset splits were consolidated to form the training set, while the official evaluation dataset served as the test set. The DCASE19 dataset comprises 60-second


recordings. Notably, it has a low percentage of frames with overlapping events. To avoid bias from the dominance of monophonic events in some polyphonic recordings, the unimodal and multimodal splitting for the results in Figures B.3 and B.4 was executed at the prediction level, not the recording level, for this dataset. For each dataset, the WTA variants were trained with $n = 3$ and $\varepsilon = 0.5$.

The outcomes from both the REAL and DCASE19 datasets mirror the patterns identified in preceding datasets in Section 4.3. Within these results, the vanilla WTA approach, represented by the pink diamonds, marginally outperforms our proposed method on the oracle metric when an equivalent number of hypotheses is used. However, a substantial disparity is observed concerning the Earth Mover’s Distance (EMD) metric when multiple hypotheses are predicted. Whenever a single source is present in the scene, the IE with five members (yellow line, triangles) still tend to outperform the other methods both in term of EMD and Oracle. In every multimodal setting within those datasets, we discern competitive results while comparing the EMD of rMCL (blue line, circles) and the other baselines.


Consistently with Section 4.5’s analysis, increasing the number of hypotheses K improves the oracle, but it may also degrade slightly the EMD when K is too large. Furthermore, top- n -WTA shows a disparity of results across datasets. For instance, the EMD results of top- n between ANSYN and other datasets reveal a contrary trend as the number of sources grows. Finally, it is noteworthy that, while oracle results remain stable, greater variability is seen in the rMCL unimodal EMD results in REAL and DCASE19 compared to prior datasets. This fluctuation may be attributed to the stochastic optimization sensibility to the challenging audio conditions, particularly when a single source is active and multiple hypotheses are utilized. Investigation on the explicit evaluation of the uncertainty estimated by rMCL, *e.g.*, due to possible label noise, is left for further work.




(a) Statistics of the metrics on ANSYN.



(b) Statistics of the metrics on RESYN.




(c) Statistics of the metrics on REAL.




(d) Statistics of the metrics on DCASE19.


Figure B.3: **Statistics of the metrics over four datasets.** Mean and standard deviation of EMD (*left*) and Oracle (*right*) over three training runs, on ANSYN, RESYN, REAL and DCASE19 datasets (*a-d*). Details and interpretation of the results are discussed in Sec. B.3.




(a) Sensitivity analysis on ANSYN.



(b) Sensitivity analysis on RESYN.



(c) Sensitivity analysis on REAL.



(d) Sensitivity analysis on DCASE19.

Figure B.4: **Effect of the number of hypotheses on the performance of rMCL.** Mean and standard deviation of EMD (*left*) and Oracle (*right*) over three training runs, on ANSYN, RESYN, REAL and DCASE19 datasets (*a-d*). Details and interpretation of the results are discussed in Sec. B.3.