

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Ambiguity in Sequential Data: Predicting Uncertain Futures with Recurrent Models

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Ambiguity in Sequential Data: Predicting Uncertain Futures with Recurrent Models / Berlati Alessandro.; Scheel Oiver.; Luigi Di Stefano ; Tombari Federico. - In: IEEE ROBOTICS AND AUTOMATION LETTERS. - ISSN 2377-3766. - ELETTRONICO. - 5:2(2020), pp. 9001185.2935-9001185.2942. [10.1109/LRA.2020.2974716]

Availability:

This version is available at: <https://hdl.handle.net/11585/806525> since: 2021-02-25

Published:

DOI: <http://doi.org/10.1109/LRA.2020.2974716>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

A. Berlati, O. Scheel, L. D. Stefano and F. Tombari, "Ambiguity in Sequential Data: Predicting Uncertain Futures With Recurrent Models," in *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2935-2942, April 2020.

The final published version is available online at:
<https://dx.doi.org/10.1109/LRA.2020.2974716>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Ambiguity in Sequential Data: Predicting Uncertain Futures with Recurrent Models

Alessandro Berlati^{1,*}, Oliver Scheel^{2,3,*}, Luigi Di Stefano¹, Federico Tombari^{3,4}

Abstract—Ambiguity is inherently present in many machine learning tasks, but especially for sequential models seldom accounted for, as most only output a single prediction. In this work we propose an extension of the Multiple Hypothesis Prediction (MHP) model to handle ambiguous predictions with sequential data, which is of special importance, as often multiple futures are equally likely. Our approach can be applied to the most common recurrent architectures and can be used with any loss function. Additionally, we introduce a novel metric for ambiguous problems, which is better suited to account for uncertainties and coincides with our intuitive understanding of correctness in the presence of multiple labels. We test our method on several experiments and across diverse tasks dealing with time series data, such as trajectory forecasting and maneuver prediction, achieving promising results.

I. INTRODUCTION

Ambiguity and uncertainty are inherently present in many machine learning tasks, both of sequential and non-sequential nature. A vehicle approaching an intersection might turn left or right, while in text generation (used, e.g., in mobile phones for auto-completion) multiple characters or words might be equally likely to follow the current one. Figure 1 visualizes the multiple trajectories possible when encountering a roundabout.

Several papers have addressed problems stemming from ambiguities and uncertainties, although, mainly, using problem-specific solutions [1], [2], [3]. One general solution are Mixture Density Networks (MDNs) [4]. However, they are defined only for regression problems and have some practical limitations [5]. Another line of work concerns Multiple Choice Learning [6], in which ensembles of M models are used. This uses more parameters and does not share information between predictions. Furthermore, most works in this field focus on non-sequential problems like image classification tasks, and pay little to no attention to time series data.

In this work we propose an extension to recurrent architectures of the Multiple Hypothesis Prediction (MHP) model by Rupprecht et al. [5], originally devised for feed-forward, non-sequential architectures. It is a general multi-purpose ambiguous prediction framework, specifically designed for sequential and recurrent models, and can be used to turn any existing deterministic model into one outputting multiple predictions. Core idea is using an implicit Voronoi tessellation of the label




Fig. 1. When arriving at a roundabout, the cyclist marked in blue has multiple possibilities of continuing its path. At time of recording, a most likely one cannot yet be determined. A good prediction model should account for all possible trajectories at this point in time, and later converge to one, as soon as the cyclist's destination becomes clear.

space, while any arbitrary loss function can be used. By considering a sequence-to-sequence prediction model, an encoder-decoder architecture and a model for generating sequences, we cover the most important recurrent models as of today, showing for each how to incorporate ambiguity. We test our models on multiple problems and find that they outperform standard non-ambiguous models on each of them and compare favorably against related methods.

Additionally, we introduce a new metric tailored for ambiguous problems, as we found previous ones to be unsuited. Indeed, many previous works use some form of oracle metric, which only considers the best hypothesis [1], [2]. Although theoretically sound, this could easily be fooled by guessing a multitude of diverse solutions in hopes of approximating the correct one. Conversely, the proposed metric better captures our understanding of ambiguity. This is achieved by re-labelling data samples, possibly assigning multiple labels to single data points, and requiring predictions to match all of them.

II. RELATED WORK

In last years deep neural networks, most notably Convolutional Neural Networks (CNNs), have proven to be general and flexible function approximators. However, most works focus on learning a one-to-one mapping from input to output, thus optimizing models to predict the single best hypothesis. Although in ambiguous situations this is neither desired nor correct, prediction of multiple hypotheses has been addressed less frequently in research. Rupprecht et al. introduced a general ambiguous prediction framework for feed-forward models dubbed MHP (Multiple Hypothesis Prediction), which is based on a Voronoi tessellation of the label space [5]. Here we extend this principle to recurrent models for sequential data.

a) *General Models: Mixture Density Networks (MDNs)* are a general method for handling ambiguities [4] based on learning the parameters of a Gaussian Mixture Distribution.

* indicates equal contribution

¹Department of Computer Science and Engineering, Università di Bologna, Bologna, Italy

²BMW Group, München, Germany oliver.scheel@bmw.de

³Faculty of Computer Science, Technische Universität München, Garching bei München, Germany

⁴Google, Zürich, Switzerland

However, in contrast to the previously mentioned ambiguous prediction framework, they are limited to regression problems, and model the encountered distributions explicitly instead of implicitly via the Voronoi tessellation. Further, it was shown that MDNs can be difficult to train due to numerical instabilities in high dimensional spaces [5].

Multiple Choice Learning focuses on training an ensemble of M models to output M predictions, backpropagating gradients only to the lowest error predictors [6]. This makes their minimum formulation similar to that of Rupprecht et al. [5] and ours. Indeed, the MHP framework provides mathematical insights on why their formulation works and enables to extend it to regression problems. Further, instead of training M different networks, a joint architecture is applied in [5], thus saving parameters and enabling information exchange. Although Lee et al. show how to address also the sequential problem of image caption generation by deploying LSTMs [6], our approach differs from their work as it is focused on sequential models and describes extensions of different state-of-the-art architectures rather than simply copying the network M times.

(Conditional) Variational Autoencoders (CVAEs) are another common method for modelling uncertainty [7]. This is done by learning a data-conditional latent distribution, s.t. diverse samples can be generated by sampling from this. Gregor et al. introduced sequential VAEs [8]. Although the sequential VAE can be applied in principle to arbitrary problems, still it is a specific type of network model, and existing works focus on problem-specific solutions. In contrast to this, our approach can be used to extend any sequential model to predict multiple hypotheses. Furthermore, in VAEs often a Gaussian distribution is used as prior for the latent distribution, thus limiting the space of possible solutions, whereas our approach is able to model implicitly arbitrary multi-modal distributions.

b) *Image Classification*: Handling multiple predictions has been addressed in image classification [9] as multiple labels may often be assigned to an image. Wang et al. combine a CNN with a Recurrent Neural Network (RNN) to predict multiple labels [3]. Yet, the solutions developed for image classification do address only non-sequential problems.

c) *Sequential Models*: As for sequential problems, ambiguity has been considered less comprehensively. Indeed, most proposals deal with problem-specific approaches, typically leveraging on Generative Adversarial Networks (GANs), VAEs or Reinforcement Learning, or other kinds of sequential generative modelling [10] in order to produce diverse outputs. To predict trajectories, Lee et al. use RNNs combined with CVAEs [2] whilst Gupta et al. GANs and a special diversity loss [1]. Bazzani et al. apply MDNs to the output of RNNs at each timestep (called Recurrent MDNs), to obtain a saliency map of visual attention in videos [11]. Unlike these problem-specific solutions, in this paper we propose a general framework to handle ambiguity in sequential problems.

d) *Ambiguous Labelling*: Kalyan et al. employ multiple labels in order to deal with sparse annotations [12]. The underlying principle stems from the same motivation as our newly proposed metric, namely clustering points w.r.t. the label space and reassigning labels. However, while theirs

is learned in combination with the model, ours is a fixed deterministic mapping intended to be used as a problem and model independent metric. Rhinehart et al. introduce two entropy terms as loss functions for trajectory prediction, encouraging predictions to be diverse whilst simultaneously precise [13]. While again very similar in motivation, this also is no metric. Further we introduce ours in a more general way, e.g. also discussing classification problems, and describing how arbitrary metrics can be extended in our way, s.t. well-known metrics like Precision and Recall can be used and compared in a multi-modal setting.

III. METHODOLOGY

A. Prerequisites

In this section we introduce notation and summarize the MHP Model [5].

1) *Multiple Hypothesis Prediction Model*: Let \mathcal{X} and \mathcal{Y} be vector spaces of input and output variables, or labels, of a problem, and N the number of available data samples. Further, let $p(x, y)$ denote the joint probability distribution over input variables and labels.

Whereas in a classical supervised setting we are interested in training a predictor $f : \mathcal{X} \rightarrow \mathcal{Y}$, in the MHP Model [5] the prediction function is extended with the possibility of outputting M predictions, or hypotheses:

$$f(x) = (f^1(x), \dots, f^M(x)). \quad (1)$$

When computing the loss \mathcal{L} , the best among the M predictors is considered, resulting in the continuous formulation

$$\int_{\mathcal{X}} \sum_{j=1}^M \int_{\mathcal{Y}_j(x)} \mathcal{L}(f^j(x), y) p(x, y) dy dx. \quad (2)$$

Here, $\mathcal{Y} = \cup_{i=1}^M \mathcal{Y}_i$ is the Voronoi tessellation of the label space, which is induced by M generators $g^j(x)$ and the loss \mathcal{L} :

$$\mathcal{Y}_j(x) = \{y \in \mathcal{Y} : \mathcal{L}(g^j(x), y) < \mathcal{L}(g^k(x), y) \forall k \neq j\}. \quad (3)$$

[5] proves that, in order for Equation 2 to be minimal, the generators have to equal the M predictors and that f^j predicts the conditional mean of the Voronoi cell it defines. A full minimization scheme defined from scratch can be found in [5]. Yet, as shown in [5], the method can be implemented efficiently by using a meta-loss on top of any loss \mathcal{L} :

$$M(f(x_i), y_i) = \sum_{j=1}^M \hat{\delta}(y_i \in \mathcal{Y}_j(x_i)) \mathcal{L}(f^j(x_i), y_i). \quad (4)$$

where the Kronecker delta $\hat{\delta}$ is defined by

$$\hat{\delta}(a) = \begin{cases} 1 - \epsilon, & \text{if } a \text{ is true,} \\ \frac{\epsilon}{M-1}, & \text{otherwise,} \end{cases} \quad (5)$$

and used to select the best hypotheses. The ϵ relaxation is used, as a bad initialization could place all hypotheses in a single Voronoi cell. Note that M is a hyper-parameter, which can be chosen freely, and that most models in this field require such a hand-tuned model parameter [6], [4].

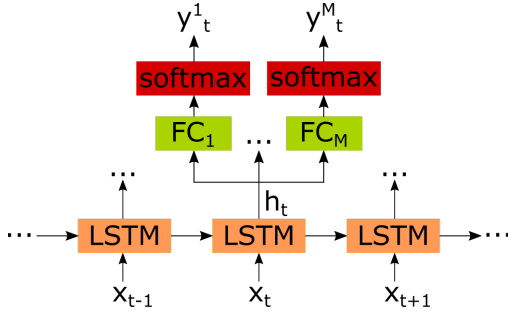


Fig. 2. Visualization of the sequence-to-sequence model. FC_m denotes a fully connected layer with parameters W_m and b_m .

2) *Long Short-Term Memory Cells*: The main goal of this work is to extend the above mentioned prediction method to a variety of different recurrent models. For all of these, we use recurrent neural networks consisting of Long Short-Term Memory Cells (LSTMs) [14]. We refer the reader to [14] for a comprehensive description of LSTMs. Here, for notational purposes, we bundle all internal calculations in a function referred to as `LSTM`, such that one computation step can be written as

$$(\mathbf{h}_t, \tilde{\mathbf{c}}_t) = \text{LSTM}(x_t, \mathbf{h}_{t-1}, \tilde{\mathbf{c}}_{t-1}) \quad (6)$$

with \mathbf{x}_t denoting the input at timestep t , \mathbf{h}_t and $\tilde{\mathbf{c}}_t$ the hidden state and cell state at time t .

B. Sequence-to-Sequence Prediction

Our first proposed model is an extension of the classical sequence-to-sequence prediction architecture. In this, we are given an input sequence and expected to return a prediction in each step, thus also returning a sequence. Thus, each sample (x_i, y_i) is now made out of sequences $x_i = (x_{i,1}, \dots, x_{i,n})$, $y_i = (y_{i,1}, \dots, y_{i,m})$. To obtain an MHP Model, we replace the common fully-connected layer on top of the recurrent network with M copies of it which do not share weights (see Figure 2):

$$\begin{aligned} \mathbf{y}_t^1 &= \text{softmax}(\mathbf{W}_1 \mathbf{h}_t + \mathbf{b}_1) \\ &\dots \\ \mathbf{y}_t^M &= \text{softmax}(\mathbf{W}_M \mathbf{h}_t + \mathbf{b}_M) \end{aligned} \quad (7)$$

Thus, the loss function needs to be applied to sequential data. In our tested problems we use the sequential cross-entropy loss as follows: If $\tilde{\mathbf{y}}_{i,t}^j$ are the predicted probabilities of the ground truth at time t , then for a single sample the standard softmax loss is calculated as $-\log(\tilde{\mathbf{y}}_{i,t}^j)$, hence:

$$\mathcal{L}(f^j(x_i), y_i) = \frac{1}{n} \sum_{t=1}^n -\log(\tilde{\mathbf{y}}_{i,t}^j) \quad (8)$$

C. Encoder-Decoder

Encoder-decoder architectures excel in many sequential problems like trajectory prediction [1], [2]. Encoder-decoder architectures consist of two separate recurrent networks: the encoder processes the input sequence $x_i = (x_{i,1}, \dots, x_{i,n})$ and eventually produces a high-dimensional vector representation enc , which is then fed to the decoder to produce the output

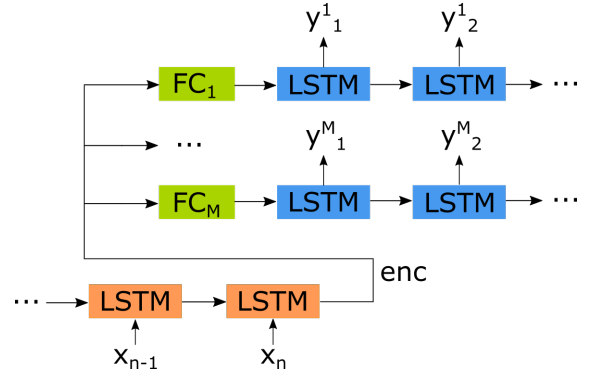


Fig. 3. Visualization of the encoder-decoder model. The encoder is drawn in orange, the decoder in blue. FC_m denotes a fully connected layer with parameters W_m and b_m .

sequence $y_i = (y_{i,1}, \dots, y_{i,m})$. We extend this model to handle ambiguity by introducing a fully connected layer between the encoder and the decoder, so to produce M vectors (see Figure 3):

$$\begin{aligned} \text{enc} &= \text{Encoder}(x_i) \\ \mathbf{y}_i^1 &= \text{Decoder}(\mathbf{W}_1 \text{enc} + \mathbf{b}_1) \\ &\dots \\ \mathbf{y}_i^M &= \text{Decoder}(\mathbf{W}_M \text{enc} + \mathbf{b}_M) \end{aligned} \quad (9)$$

As for the loss, any function suitable to sequences may be used. In our experiments dealing with encoder-decoder models we use the L2-loss $\mathcal{L}(f^j(x_i), y_i) = \frac{1}{m} \sum_{t=1}^m (\mathbf{y}_{i,t}^j - y_{i,t})^2$.

D. Sequence Generation

Although for the problem of generating sequences the samples are still sequence pairs (x_i, y_i) , the focus is to learn conditional probabilities of single frames $P(x_{i,t} | x_{i,t-1}, \dots, x_{i,t-m})$. This way, after training, the network can be initialized with a seed and then run in a closed-loop, thereby creating sequences of arbitrary length. There is a wide range of possible applications, such as text generation [15]. By using a similar model as in Section III-B (duplicating the output layer M times), the training process remains the same, and the same loss functions may be used.

During inference, though, to employ an MHP model, we encounter the problem of exponentially growing possibilities in each step. This is not feasible, not just in terms of practical usage but also in terms of limited resource constraints. Therefore, during inference, our proposed solution makes use of different functions to decide when to split and when to merge predictions. These functions can be chosen in a problem-specific way, but it is also possible to train, e.g., a neural network for this task, implicitly calculating these functions. Note that this procedure is only applied during inference, s.t. these functions do not need to be differentiable. Similar issues have been known before, for related strategies (albeit employing SHP models regarding our definition) we refer to [16].

Inference is started with a single hypothesis H , which can be empty, contain a special starting symbol or an arbitrary number of elements. In each step, inferring for d future steps is simulated. Due to the prediction of M hypothesis

in each of these, a tree of depth d and branching factor M is created. When the predictions in this tree are not yet diverse enough but close together, the first layer of the tree is merged into one prediction, which is appended to H . This merging is done by a function referred to here as `Merge`, while the function to check the tree diversity is denoted as `CheckSplit`. When this encounters a diverse enough tree, function `ChooseTreePaths` finds M hypotheses h_1, \dots, h_M of length d in the tree. H is subsequently split into M different predictions by appending h_1, \dots, h_M , and these hypothesis are consequently followed separately: For each hypothesis path the inference simulation is done, and the resulting M predictions are merged via `Merge`. The pseudocode for this scheme is shown in Algorithm 1, while exemplar implementations of these black box functions are shown in Section V-A.

Algorithm 1 Inference for MHP Sequence Generation, starting with *start* and inferring for a total of l steps. *BuildTree* expects the starting element as well as the desired tree depth as input.

```

procedure INFER(start,  $d$ ,  $l$ )
  DoneSplit = False;
  Predictions = start;
  for  $i$  in  $1 \dots l$  do
    if not DoneSplit then
       $p$  = last item of Predictions
      tree = BuildTree( $p$ ,  $d$ );
      if CheckSplit(tree) then
         $[(p_1^1, \dots, p_d^1), \dots, (p_1^M, \dots, p_d^M)]$  =
        ChooseTreePaths(tree);
        Append ( $[(p_1^1, \dots, p_1^M), \dots, (p_d^1, \dots, p_d^M)]$ ) to Predictions;
        DoneSplit = true;
      else
         $[p^1, \dots, p^M]$  = First layer of tree;
         $p$  = Merge( $[p^1, \dots, p^M]$ );
        Append  $p$  to Predictions;
    else
       $[p_1, \dots, p_M]$  = last item of Prediction;
      for  $m$  in  $1 \dots M$  do
         $p$  = BuildTree( $p_m$ ,  $l$ );
         $p_m$  = Merge( $p$ );
        Append  $[p_1, \dots, p_M]$  to Predictions;
  return Predictions

```

IV. MULTI-MODAL METRIC

In many works concerning multiple predictions, some form of oracle metric is used, which compares the ground truth to the closest prediction among possibly multiple ones [2], [1]. Formally, if l is a metric between prediction set X_i and label y_i , then $l_{\text{oracle}} = (X_i, y_i) = \min_{x \in X_i} l(x, y_i)$. Although this rewards models for predictions close to the “correct” one and allows for diverse predictions, we argue that this metric should not be the gold standard for problems containing ambiguity. Indeed, better scores could be reached by simply guessing hypotheses, and in real-world settings it is not clear what actions should be taken based upon models with good oracle scores. Further, and even more importantly, ambiguous situations sometimes lead to multiple, equally “correct” predictions. On the other hand, when averaging multiple predictions, in most cases the advantage of predicting multiple hypotheses cannot be shown numerically, as the loss is minimized by the (incorrect) mean. Thus we propose a novel, multi-modal (M2) metric specifically designed for

problems containing ambiguities and uncertainties. The core idea is to automatically re-label samples, possibly allowing multiple labels, and extending standard metrics $l(x_i, y_i)$ to a set-based calculation $l_{M2}(X_i, Y_i)$. Note that this principle can be used for basically all prediction tasks and metrics, i.e. also when dealing with non-sequential data. We derive our metric for both the discrete and continuous label cases.

A. Discrete Labels

First step concerns re-labelling of samples. Purposely, for each label class c we calculate a polytope $P(c)$ in the input space, containing all points labeled as c (denoted by $\lambda(x) = c$). Then, for each sample x_i with label y_i , the new label set is defined as $Y_i = \{y_i\} \cup \{c \mid c \in C, x_i \in P(c)\}$, where C is the set of all labels. Such a polytope could be the convex hull. However, this does not scale to high dimensions and it is not reducible in size: especially for high-dimensional and complex data, outliers can inflate the convex hull, which is not desired. Thus, we define the polytope $P(c) = \{x \mid \lambda(x) = c, \min_\tau \leq x \leq \max_\tau\}$, where x , \min_τ and \max_τ are vectors of the d -dimensional sample space. A threshold τ can be defined to control the size of the polytope, denoting that an (approximate) fraction τ of all points labeled c are contained in the polytope. For this, \min_τ^d and \max_τ^d are given by the $([1 - \frac{\tau}{2}] \cdot 100)$ th and $(\frac{\tau}{2} \cdot 100)$ th percentile, respectively, in each dimension d . τ should be chosen such that it best resembles the ambiguity of the problem, e.g. overlaps with our intuitive understanding of multi-modality. Thus it is freely choosable, nevertheless the results obviously are reproducible and comparable given τ . Although this method is applicable for an arbitrary number of dimensions (consider treating MNIST [17] images as 784-dimensional vectors), for high-dimensional spaces it is possible and recommendable to use a lower-dimensional latent space (e.g. intermediate layers of common CNN architectures or autoencoders, similar to [12]) and define the polytope on this. For a better understanding, in Section VI the resulting polytopes of a toy task are shown. Basically, any metric can be extended with this new label set in a natural way. Here we consider the widely-used metrics Precision (Pr) and Recall (Re). The extension is motivated from the context of information retrieval, i.e. a Precision of 1 should only be reached if all predicted hypotheses are correct, analogously a Recall of 1 should mean that all labels are predicted. To satisfy this we define

$$pr(f(x_i), Y_i) = \frac{|f(x_i) \cap Y_i|}{|f(x_i)|} \quad (10)$$

and

$$re(f(x_i), Y_i) = \frac{|f(x_i) \cap Y_i|}{|Y_i|} \quad (11)$$

for each sample i and ambiguous prediction $f(x_i)$ and overall $Pr_{M2} = \frac{1}{N} \sum_{i=1}^N pr(f(x_i), Y_i)$, $Re_{M2} = \frac{1}{N} \sum_{i=1}^N re(f(x_i), Y_i)$.

B. Continuous Labels

Still assuming the set-based interpretation, in the case of continuous labels we require that each prediction is close to




Fig. 4. Visualization of simulated trajectories on a three-way intersection. Each vehicle approaches the intersection from the bottom of the image, and then chooses one of three possible paths.

at least one label, and that for each label there is at least one close hypothesis. However, with continuous labels we cannot directly define the polytope but first apply clustering in the label space to get discrete labels C . Hence, analogously, we define $P(c)$ and $Y_i = \{\mu(c) \mid c \in C, x_i \in P(c)\}$, but here $\mu(c)$ denotes the center of cluster c (see Section VI for visualizations of such clusters). We define the extension of a metric $l(x_i, y_i)$ by

$$l_{M2}(x_i, Y_i) = \frac{\sum_{x \in f(x_i)} \min_{y \in Y_i} l(x, y)}{|f(x_i)| + |Y_i|} + \frac{\sum_{y \in Y_i} \min_{x \in f(x_i)} l(x, y)}{|f(x_i)| + |Y_i|} \quad (12)$$

VI. PROBLEMS AND DATASETS

In this section we introduce several problems coming from different fields, together with the models used to address them. As the general models are introduced in Sections III-B to III-D, here we will just highlight problem specific adaptations (if any).

A. Toy Intersection

To motivate the usage of multiple hypotheses, we start with simple toy problems, which offer clear data. They are based on synthetically created traffic on a three-way intersection (see Figure 4). Each trajectory t is an ordered list of 2D coordinates: $t = ([c_1^1, c_2^1], \dots, [c_1^n, c_2^n])$.

a) *Toy Classification*: Goal of the classification task is to predict each vehicle's destination (left, straight or right) at every time-step. We solve this problem by using the sequence-to-sequence prediction approach from Section III-B. Let l_i be the correct label of sample i , then $x_i = ([c_1^1, c_2^1], \dots, [c_1^n, c_2^n])$ and $y_i = (l_i^1, \dots, l_i^n)$.

b) *Toy Prediction*: For this task future trajectories have to be predicted conditioned on past ones. We showcase the usage of the encoder-decoder and sequence generation models (Sections III-C and III-D). For both models, a first part of each trajectory $x_i = ([c_1^1, c_2^1], \dots, [c_1^m, c_2^m])$, $m < n$, serves as input, while $y_i = ([c_1^{m+1}, c_2^{m+1}], \dots, [c_1^n, c_2^n])$ is the desired output. In the encoder-decoder model x_i and y_i are the input and output of the encoder and decoder, respectively.

As for the sequence generation model, it is trained with full trajectories, learning the conditional probabilities of successive trajectory points. During inference, the model is initialized with x_i and expected to produce y_i , constructing a tree of depth 8 in each step. Function `CheckSplit` checks for the maximum distance of two points in the last layer of the tree and returns true if this exceeds a certain threshold. `Merge` simply returns the mean of the predicted points.

`ChooseTreePaths` sorts all points in the last layer of the tree by their angle to the current point, and then returns those M points that evenly split the range of angles into M parts.

B. Lane Change Prediction

Predicting maneuvers, or more generally action anticipation, is a crucial task in many fields, e.g. autonomous driving, and has sparked great interest in the scientific community [18], [19]. Here we want to predict lane change maneuvers of vehicles as early and accurately as possible. We use the sequence-to-sequence prediction model from Section III-B. At each timestep, five features f_t , like distance to the lane boundaries and lateral velocity, are used as inputs. Each timestep is assigned a label l_t , which can either be L (change left), F (follow lane), R (change right) (see [20] for more details). We use the publicly available NGSIM dataset [21], which offers vehicle trajectories recorded on highways and interstates in the United States.

C. Trajectory Prediction

A multitude of models have been proposed for the problem of predicting trajectories, mainly concerning pedestrians [2], [1]. We employ the encoder-decoder model from Section III-C for this task, using the Stanford Drone Dataset (SDD) [22]. This dataset contains trajectories of, among others, pedestrians, cyclists and skateboarders in eight different scenes. Input to the prediction model is a concatenation of two features: 1) relative 2-D coordinates, consisting for each trajectory of the absolute starting point and subsequently x- and y-offsets to the previous point; and 2) simple semantic maps of the environment, processed by a CNN and centered at the agent's location. These semantic maps are binary images of size 64 x 64, in which each pixel describes membership in the classes obstacle and free space.

VI. RESULTS

In this section we illustrate our extensive experimental evaluation aimed at analyzing the performance of the proposed MHP models. As our goal predominantly is providing a viable general ambiguous prediction framework, for one we compare against non-ambiguous models (SHP), to show the superiority of multiple hypotheses. Further, we provide a competitive comparison of our proposed framework against other well-performing methods yielding multiple hypotheses. Finding best-performing models or outperforming fine-tuned expert ones for specific tasks and datasets is of lesser interest in this work, for the interested reader in the respective sections we compare against some state-of-the art models of this kind nevertheless.

In the field of multi-modal prediction, among the best performing and most commonly used methods are MDNs, MCL and CVAEs, thus we compare against these. In general, we try to use similar models throughout, e.g., an MCL model is identical to the MHP one up to the specific multi-modal part (i.e., using the same input features and exhibiting the same hidden size). As mentioned, other baselines are the SHP models, which are also the same models as presented

in Sections III-B to III-D, except missing the MHP parts (i.e. having no duplicate MHP layers, and just outputting one hypothesis). For a fairer comparison, they are extended with the possibility of predicting multiple hypotheses, if possible.

MDNs are only defined for regression problems, and also cannot be applied in a meaningful way for encoder-decoder models. Thus, these are not considered here, but a comparison is done in [5], yielding favorable results for the MHP framework. Implementation details about the used models are given in the respective sections.

To better understand differences in computational requirements of the tested algorithms, wall clock times needed for training (t_t) and inference (t_i) are listed in the following (in *ms*, per mini-batch)¹. For all our experiments, Adam optimizer is used with a learning rate of 0.001. We split the available data into training, validation and test set with a 60-20-20 ratio, employing early stopping on the validation set. All quantitative and qualitative results in the following sections are reported on the test set.

A. Classification

As simple extensions to a standard SHP model we introduce SHP*, which outputs all classes whose predicted probabilities exceed a certain threshold γ . For all problems we test $\gamma = 0$ and a problem specific γ , which returns best results w.r.t. to the M2 metric. The common metrics Precision (Pr), Recall (Re) and the associated F1 score ($F1 = 2 \cdot \frac{Pr \cdot Re}{Pr + Re}$) are used for evaluation, either using an oracle (Pr_O and Re_O) or as defined in Section IV. An MCL model is further used for comparison: It contains M independent copies of the SHP model, during training gradients are only backpropagated to the best hypothesis (which resembles our MHP training with $\epsilon = 0$). For classification problems, we are not aware of meaningful (C)VAE extensions.

a) Toy Classification: For all models an LSTM network with 512 hidden units is used, the mini-batch size is 32, the M2 metric parameter τ is set to 1, due to the well-behaved structure of the data, and $\epsilon = 0.15$. The synthetic dataset consists of 10000 trajectories with 75 points each.

Quantitative results are shown in Table I. A standard SHP model picks a random prediction when multiple ones are equally likely, but converges to the correct prediction once the situation becomes unambiguous. This is reflected by an oracle Precision and Recall of roughly 2/3. The MHP model fares much better with scores of nearly 1. Setting $\gamma = 0$ reveals a first weakness of the oracle metric, as precision and recall go up to 1: all 3 existing classes are predicted at each timestep, thus the best hypothesis always equals the ground truth. The M2 metric penalizes this behavior with a lower precision. Conversely, due to the re-labeling, the standard SHP model achieves a near perfect M2 precision (since the predicted hypothesis is nearly always correct), but a low recall. Figure 5 visualizes these results. Similar to the SHP model with $\gamma = 0$, the MCL model always outputs 3 different hypotheses. This behaviour is caused by the absence of punishment for




Fig. 5. Correctness of predictions w.r.t. to the M2 metric: If the set of predictions equals the set of new labels, the point is drawn in green, otherwise in red.




Fig. 6. Visualizations of re-labelling step for M2 metric.

(wrongly) diverse predictions due to the principles of MCL ($\epsilon = 0$), and represented by perfect oracle scores, but shows through deductions in the M2 metric. Since the MCL model consists of M full expert-models, the number of parameters as well as time needed for training and inference grows. In contrast to this, the MHP model only contains slightly more parameters than the SHP one, resulting in near-identical running times. Although some categories are won by other models, overall, the MHP model performs best, proving a correct understanding of the ambiguous nature of the problem.

Figure 6a shows the polytope obtained by applying the re-labelling procedure of the M2 metric for this task.

b) Lane Change Prediction: We set $\tau = 0.85$, so that the resulting labelling resembles our human understanding of ambiguity in the problem. For all models an LSTM size of 128 is used, a mini-batch size of 256 and $\epsilon = 0.15$. The results can be found in Table II and are similar to those discussed in the previous paragraph, the MHP model performing best. Note that our used SHP model is the same as in [20], thus a direction comparison to a state-of-the art (non-ambiguous) model is given. To the best of our knowledge, no ambiguous models for this kind of task are employed.


TABLE I
RESULTS OF THE TOY CLASSIFICATION TASK (γ LISTED IN BRACKETS).

Name	Pr_O	Re_O	Pr_{M2}	Re_{M2}	$F1_{M2}$	t_t	t_i
SHP	0.683	0.683	0.999	0.649	0.787	205	105
SHP* (0)	1.0	1.0	0.676	1.0	0.807	205	105
SHP* (0.01)	0.903	0.903	0.750	0.893	0.815	205	105
MCL	1.0	1.0	0.677	1.0	0.807	535	237
MHP	0.999	0.999	0.980	0.984	0.982	210	113

TABLE II
RESULTS ON THE NGSIM DATASET (γ LISTED IN BRACKETS).

Name	Pr_O	Re_O	Pr_{M2}	Re_{M2}	$F1_{M2}$	t_t	t_i
SHP	0.766	0.766	0.797	0.747	0.771	388	263
SHP* (0)	0.952	0.952	0.352	0.952	0.514	388	263
SHP* (0.55)	0.799	0.799	0.572	0.779	0.660	388	263
MCL	0.952	0.952	0.364	0.952	0.527	1005	637
MHP	0.905	0.905	0.732	0.906	0.810	658	506

¹Note that running times vary based on hardware. All experiments are done using a standard laptop containing an NVIDIA Quadro M2000 GPU.



(a) Prediction of the SHP model. (b) Prediction of the MHP model.

Fig. 7. Resulting prediction for a sample trajectory. The input to the decoder is drawn in blue, the ground truth in green, and the predictions in yellow.

B. Prediction

A standard SHP encoder-decoder model cannot easily be extended to output multiple predictions, thus no extension of this is analyzed. We analyze MCL and CVAE extensions though: Again, the MCL model consists of M copies of the encoder-decoder SHP. In the CVAE, the encoder produces a 20-dimensional latent vector z , represented by mean and standard deviation. Via the reparametrization trick the decoder samples from this to generate M hypotheses. Similar to [2], the average error over all hypotheses is used, in addition to the Kullback-Leibler divergence between z 's distribution and a unit normal distribution. To evaluate the outcomes we use the Final Displacement Error (FDE) and Average Displacement Error (ADE) metrics, which denote the metric distance of the prediction to the ground truth considering either only the last timestep or all, respectively. For the clustering used in the M2 metric, we convert trajectories consisting of n timesteps to a $2n$ -dimensional space, and in this employ mean-shift clustering.

a) Toy Prediction: Again, we use the toy intersection from Section V-A, on which we simulate 10000 trajectories containing 60 timepoints each. The first 30 are fed to the encoder, and the decoder is expected to produce the remaining trajectory. Both LSTMs consist of 64 hidden units, $\epsilon = 0.15$, $\tau = 1$ and mini-batch size is 256. The bandwidth used for mean-shift clustering is determined automatically by the algorithm.

The predictions for one sample can be seen in Figure 7, for more we refer to the supplementary video. Whereas the SHP model predicts one random path, the MHP model nicely accounts for all possible outcomes. The corresponding quantitative results over the test dataset are reported in Table III. The SHP model only outputs one mean prediction, resulting in bad overall scores. Also the CVAE model shows problems generating diverse outputs, falling back to the mean as well. This is to be expected, as inputs are very similar, resulting in an average distribution. Both the MCL and MHP model perform much better. Again, the MCL model exhibits a longer training and inference time though, while CVAE and MHP perform similarly timewise. Figure 6b shows the resulting clusters used for the M2 metric.

TABLE III
RESULTS OF THE TOY PREDICTION TASK.

Name	FDE_O	ADE_O	FDE_{M2}	ADE_{M2}	t_t	t_i
SHP	2.82	4.60	2.29	3.64	112	49
MCL	0.15	0.16	0.53	0.53	320	147
CVAE	2.76	4.53	1.99	2.95	218	83
MHP	0.26	0.22	0.51	0.53	210	88



(a) Prediction using 5 hypotheses. (b) Prediction using 3 hypotheses. (c) Prediction of the SHP model.

Fig. 8. Resulting predictions in an ambiguous scene of the SDD.

b) Trajectory Prediction: Using the SDD dataset we experiment with 3 and 5 hypotheses, in addition to the SHP encoder-decoder. For all models we use an LSTM size of 256 and a sequence length of 8 (3.2s) for the encoder and 12 (4.8s) for the decoder, to comply with existing works in this field. Further, $\epsilon = 0.15$, $\tau = 1$ and mini-batch size 256. For the mean-shift clustering we use a bandwidth of 250.

Figure 8 shows the results on the ambiguous scene, which was hinted in Figure 1. The SHP model simply predicts a linear follow-up trajectory, whereas the MHP models cover the full space of possible directions, accurately matching the existing street topology. More hypotheses provide more fine-grained predictions. Again we refer to the supplementary video for more results, also including those of other models. Table IV shows quantitative results over the full test set, proving that the MHP models perform better than the SHP one. The VAE models improve over the toy task given the more diverse input, still sometimes lack diversity and miss hypotheses, as shown by higher oracle scores. Overall, regarding the oracle metric the MHP models perform best, w.r.t. the M2 metric the performances of the ambiguous models are close together - all seem to show an understanding of multimodality when faced with this problem. As before, the MCL models exhibit the longest running times, while VAE and MHP are close together, with MHP scaling slightly better with growing M . Further we compare our methods to SoPhie GAN [23] and the DESIRE framework [2] (results for this obtained from the interpolation in [23]), for a reference to state-of-the art models. For our experiments we use the SDD part of the TrajNet dataset [24], in which the sequences of the SDD have been preformatted to the desired sequence lengths and overly linear trajectories been filtered out. As labels for the test set are not public (yet) and multiple hypotheses not supported in the score calculation, we created our own train / test split from the public training set, which might explain the slightly better scores of our models.

TABLE IV
RESULTS ON THE SDD.

Name	FDE_O	ADE_O	FDE_{M2}	ADE_{M2}	t_t	t_i
SHP	35.94	17.76	49.14	40.36	238	143
MCL 3	27.52	14.46	45.02	41.89	338	187
MCL 5	25.31	13.58	43.81	42.52	450	227
VAE 3	35.07	17.16	44.06	38.86	311	166
VAE 5	34.06	16.62	41.86	38.23	377	193
MHP 3	27.20	14.36	42.66	39.03	312	175
MHP 5	24.89	13.48	41.96	39.11	355	187
DESIRE	34.05	19.25	-	-	-	-
SoPhie	29.38	16.27	-	-	-	-




Fig. 9. Prediction results on one sample trajectory. The ground truth is depicted in green, the predictions in yellow.

C. Sequence Generation

In this section we demonstrate our MHP extension for sequence generation models, predicting trajectories on the toy intersection. We again use the FDE and ADE metrics. The synthetic dataset consists of 10000 trajectories with 25 points each. In sequence generation scenarios, a standard SHP model can only be modified in a very specific way, if at all, to output multiple hypotheses. Experiments did not yield any usable results for SHP extensions. Further, CVAE or MCL models cannot be extended to this application scheme in a trivial way. Thus here we only compare a standard SHP model to the MHP extension. Both LSTMs used have a hidden size of 256, same for mini-batch size, $\epsilon = 0.15$ and $\tau = 1$.

After training, for each test trajectory we initialize the model with 5 points and infer 20. Figure 9 shows the resulting predictions for one sample, in Table V the results for the test set are listed. Again, the MHP model outperforms the SHP model. While training times of both models are similar, inference equals a single forward pass in the SHP model, while in the MHP model a tree is created to generate multiple hypotheses, resulting in higher inference times.

VII. CONCLUSION

In this paper we have proposed an extension of the Multiple Hypothesis Prediction model, resulting in a universally applicable framework for dealing with ambiguity and uncertainty in sequential problems. Experiments with different architectures and in fields as diverse as trajectory analysis and maneuver prediction show the wide applicability and superior performance of ambiguous models over standard non-ambiguous ones in the considered cases. Additionally, as opposed to other multi-modal prediction models, the MHP extension proves to be applicable for all kinds of problems and is among the best performing methods for each of them, while simultaneously exhibiting minimal parameter overhead. Furthermore, to better assess upon problems featuring data with multiple possible labels, we have introduced a novel metric, which we posit be considered whenever dealing with ambiguous problems. Across diverse problems and tasks we showed its applicability and pitfalls of standard oracle metrics.

VIII. ACKNOWLEDGEMENTS

We would like to thank Christian Rupprecht and Iro Laina for valuable discussions and feedback.

TABLE V
RESULTS ON THE TOY SEQUENCE GENERATION TASK.

Name	FDE_O	ADE_O	FDE_{M2}	ADE_{M2}	t_t	t_i
SHP	18.26	5.07	14.07	4.13	59	188
MHP	5.45	1.99	5.48	2.06	67	24433

REFERENCES

- [1] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] N. Lee, W. Choi, P. Vernaza, C. Choy, P. H. S. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [3] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, and W. Xu, "Cnn-rnn: A unified framework for multi-label image classification," in *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [4] C. M. Bishop, "Mixture density networks," Tech. Rep., 1994.
- [5] C. Rupprecht, I. Laina, M. Baust, F. Tombari, G. D. Hager, and N. Navab, "Learning in an uncertain world: Representing ambiguity through multiple hypotheses," *Int. Conf. on Computer Vision (ICCV)*, 2017.
- [6] S. Lee, S. Purushwalkam, M. Cogswell, V. Ranjan, D. J. Crandall, and D. Batra, "Stochastic multiple choice learning for training diverse deep ensembles," in *Int. Conf. on Neural Information Processing Systems (NIPS)*, 2016.
- [7] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *Int. Conf. on Learning Representations (ICLR)*, 2014.
- [8] K. Gregor, I. Danihelka, A. Graves, D. Rezende, and D. Wierstra, "Draw: A recurrent neural network for image generation," *Inf. Conf. on Machine Learning (ICML)*, 2015.
- [9] Y. Gong, Y. Jia, T. Leung, A. Toshev, and S. Ioffe, "Deep convolutional ranking for multilabel image annotation," *Int. Conf. on Learning Representations (ICLR)*, 2014.
- [10] Y. Liu, R. Yu, S. Zheng, E. Zhan, and Y. Yue, "NAOMI: non-autoregressive multiresolution sequence imputation," *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [11] L. Bazzani, H. Larochelle, and L. Torresani, "Recurrent mixture density network for spatiotemporal visual attention," *arXiv preprint arXiv:1603.08199*, 2016.
- [12] A. Kalyan, S. Lee, A. Kannan, and D. Batra, "Learn from your neighbor: Learning multi-modal mappings from sparse annotations," in *Int. Conf. on Machine Learning (ICML)*, 2018.
- [13] N. Rhinehart, K. Kitani, and P. Vernaza, "R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting," *European Conference on Computer Vision (ECCV)*, 2018.
- [14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, 1997.
- [15] A. Graves, "Generating sequences with recurrent neural networks," *arXiv preprint arXiv:1308.0850*, 2013.
- [16] P. Schydlow, M. Rakovi, and J. Santos-Victor, "Anticipation in human-robot cooperation: A recurrent neural network approach for multiple action sequences prediction," *Int. Conf. on Robotics and Automation (ICRA)*, 2018.
- [17] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [18] A. Jain, A. Singh, H. S. Koppula, S. Soh, and A. Saxena, "Brain4cars: Car that knows before you do via sensory-fusion deep learning architecture," *Int. Conf. on Robotics and Automation (ICRA)*, 2016.
- [19] Y. Jiang and A. Saxena, "Modeling high-dimensional humans for activity anticipation using gaussian process latent crfs," in *Robotics: Science and Systems*, 2014.
- [20] O. Scheel, N. N. Shankar, L. Schwarz, N. Navab, and F. Tombari, "Attention-based lane change prediction," *Int. Conf. on Robotics and Automation (ICRA)*, 2018.
- [21] "Ngsim project," <https://ops.fhwa.dot.gov/trafficanalysisstools/ngsim.htm>.
- [22] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," *European Conf. on Computer Vision (ECCV)*, 2016.
- [23] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, and S. Savarese, "Sophie: An attentive gan for predicting paths compliant to social and physical constraints," *Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [24] A. Sadeghian, V. Kosaraju, A. Gupta, S. Savarese, and A. Alahi, "Trajnet: Towards a benchmark for human trajectory prediction," *arXiv preprint*, 2018.