

hypothesis

Authors

Alexander Jung*, Department of Computer Science, Aalto University, Espoo, Finland

*Corresponding author: alex.jung@aalto.fi

Synonyms

predictor

Abstract

A hypothesis is a map $h : \mathcal{X} \rightarrow \mathcal{Y}$ from the feature space $\mathcal{X}$ to the label space $\mathcal{Y}$. Given a data point with features $\mathbf{x}$, the hypothesis returns a prediction $h(\mathbf{x})$ for the label y . A machine learning (ML) method selects a single learned hypothesis $\hat{h}$ from a restricted hypothesis space $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$. This restriction reflects finite computational resources of a practical ML method. As a point in case, empirical risk minimization (ERM)-based methods must be able to search the entire hypothesis space to learn a hypothesis $\hat{h}$ that minimizes the empirical risk on a training set.

Definition

Predicting tomorrow’s maximum daytime temperature from today’s morning temperature is a typical machine learning (ML) task. A hypothesis is the map that performs this prediction. Formally, a hypothesis is a map (also called a function) $h : \mathcal{X} \rightarrow \mathcal{Y}$ from the feature space $\mathcal{X}$ to the label space $\mathcal{Y}$ (Shalev-Shwartz and Ben-David, 2014, Sect. 2.1). Given a data point with features $\mathbf{x}$, the hypothesis h returns an estimate $\hat{y} = h(\mathbf{x})$ of the true label y . Fig. 1 illustrates this for an audio-based application where the hypothesis predicts a Freddie-likeness score from an audio recording.




Figure 1: Hypothesis $h : \mathcal{X} \rightarrow \mathcal{Y}$ mapping the features $\mathbf{x} \in \mathcal{X}$ of a data point to a prediction $h(\mathbf{x}) \in \mathcal{Y}$ of the label. Here, the ML application <https://freddiemeter.withyoutube.com/> uses the samples of an audio recording as features to predict how closely a person’s singing resembles that of Freddie Mercury. The displayed value $h(\mathbf{x}) = 0.82$ indicates an 82% resemblance

An ML method aims to find (or learn) a hypothesis h such that $y \approx h(\mathbf{x})$ for any data point with features $\mathbf{x}$ and label y . Finite computational resources require restricting the search to a subset of

$$\mathcal{Y}^{\mathcal{X}} = \{h : \mathcal{X} \rightarrow \mathcal{Y}\},$$

the set of all functions from $\mathcal{X}$ to $\mathcal{Y}$. This subset is the hypothesis space $\mathcal{H}$ (also called the model) underlying the method. The output of an ML algorithm acting on a training set is a single learned hypothesis $\hat{h} \in \mathcal{H}$ obtained via empirical risk minimization (ERM), i.e., the element of $\mathcal{H}$ that minimizes the empirical risk on the training set.

Different hypothesis spaces admit different representations of their hypothesis maps. Polynomial regression (with a single scalar feature x) represents each $h \in \mathcal{H}$ as an algebraic expression $h(x) = \sum_{j=0}^r w_j x^j$, with coefficients w_j as the model parameters. A hypothesis space of computer programs represents h as executable source code, e.g., a Python function with a prescribed signature. A decision tree or an artificial neural network (ANN) represents h as a signal-flow chart whose nodes apply simple operations (e.g., comparisons, weighted sums, nonlinear activations) and whose edges define the order in which these operations are composed. These representations differ in which functions they can express and in how efficiently the ERM problem can be solved over the corresponding $\mathcal{H}$. Fig. 2 shows the three representations side by side.




Figure 2: Three representations of a hypothesis map. (a) As an algebraic expression, here a degree-2 polynomial in a scalar feature x . (b) As source code, here a Python function. (c) As a signal-flow chart, here a small decision tree that routes a data point to a prediction of 3, 1, or -1 based on the threshold comparisons on the feature x

Cross References

map, function, prediction, predictor, model, hypothesis space, feature vector, label, data point, ERM, training, generalization

References

1. Shalev-Shwartz and Ben-David (2014). Understanding Machine Learning: From Theory to Algorithms. Cambridge Univ. Press, New York, NY, USA.