

Normative View, Induction and Learning

Lecturer: Drew Bagnell

Scribe: Sankalp Arora ¹




Figure 1: Drew Bagnell, with pipe and corduroy jacket, channeling the philosopher David Hume through elaborate chicken analogies.

1 Lecture Abstract

This lecture begins with a small discussion on logic and normative view of the world. Then we will question the nature of induction and its use as a tool for modeling and making decisions about the world. We will explore the limits of the inductive process, and in the process, prove how well different online induction algorithms(*e.g* - weighted majority, randomized weighted majority etc.) perform in relation to each other through the new concepts of *loss* and *regret*.

2 Normative View of the World

Normative view is concerned with identifying the best decision to take or inference to make, assuming an ideal decision maker who is fully informed, is able to compute with perfect accuracy and is fully rational. It assumes probability based inference and general decision making can and should only be defined by pure logic. The normative view suggests that any decision maker or inference engine that doesn't follow the aforementioned rules is wrong or invalid. Hence it does not state the amount of error induced in a system if approximation techniques(*like finite number of particles in*

¹Some content adapted from previous scribes: Mark Desnoyer, Andres Rodriguez, Ben Eckart

a *particle filter*) are used for the lack of adequate computation resources, time or knowledge. Succinctly put the normative view considers “known unknowns”, not the “unknown unknowns”. And this fallacy makes the normative view of little use in real life.

3 Induction

Our standard justification for the use of probability theory is that it seems like the “reasonable” thing to do. Logic is certainly reasonable, and probability naturally and uniquely extends our notion of logic to include degree-of-belief type statements. Thus, if one accepts logic, one must accept probability as a way of reasoning about the world.

In all the probabilistic systems discussed so far (*e.g.* localization and mapping), we can momentarily discard our normative view by looking at their common elements. At the very core, something like mapping, for example, can be seen simply as an induction problem. We take in data about world, and using prior knowledge, we make inferences, and then predictions using the inferences. In fact, Bayes’ rule can be seen as the “canonical” way to do these operations. Induction is a method to predict the future based on observations in the past, and to make decisions from limited information. Thus, we can look generally at induction itself in order to arrive at answers to certain fundamental questions behind many robotic learning systems. Questions such as: “*Why does inductive reasoning work?*”, “*When does it fail?*”, and “*How good can it be?*” will tell us about the nature of learning and lead us to unique notions of optimality with regard to worst-case adversarial conditions.

3.1 Induction from a Chicken’s Perspective

We begin with an illustrative example of how inductive reasoning can fail. Imagine you are a chicken. Every day of your life, the farmer comes to you at a certain time in the day and feeds you. On the 1001th day, the farmer walks up to you. What do you predict? What can you predict other than that he will feed you? Unfortunately for you, this day the farmer has decided to eat you. Thus, induction fails spectacularly. You may object to this example and say that the model of the world that the chicken was using was wrong, but this is precisely the point of the discussion: perhaps we can *never* have a model that perfectly aligns with the real world. Perhaps, you know vaguely of your impending doom, so you assign a low probability of being eaten everyday. You still will (probably) guess incorrectly on that 1001th day because your guess of being fed will dominate your feeling of doom. Furthermore, for any model you bear, an adversarial or random farmer can always thwart it to force false inductive predictions about your fate each day.

3.2 Three views of induction:

1. *The David Hume view.* No matter what, your model will never capture enough of the world to be correct all the time. Furthermore, induction relies heavily on priors made from assumptions, and potentially small changes in a prior can completely change a future belief. If your model is as good as your assumptions or priors, and most assumptions are not likely to be valid at all, then induction is fundamentally broken. We find this view, although true, is not very satisfying or particularly useful.

2. *The Goldilocks/Panglossian/Einstein Position.* It just so happens that our world is set up in a way such that induction works. It is the nature of the universe that events can be effectively modeled through induction to a degree that allows us to make consistently useful predictions about it. In other words, we're lucky. Once again, we do not find this view very satisfying or particularly useful.
3. *The No Regret view.* It is true that you cannot say that induction will work or work well for any given situation, but as far as strategies go, it's provably near-optimal. This view comes from the outgrowth of game theory, computer science, and machine learning. This thinking is uniquely late 20th century.

4 On-line learning Algorithms

Though we can never be *certain* of our prediction (as a chicken, we will always be eaten at some point), we *can* form a strategy that will do nearly as well as we could possibly hope given our current situation. In other words, we can always do something which is the best we can do (and we can prove this is the case).

We begin by studying the problem of “predicting from expert advice.” Our notion of an expert is simply something that makes a prediction. Unfortunately, “expert” is somewhat of a misnomer, given that an expert can give bad, random, or adversarial advice. It is fruitful to think of a collection of experts as a collection of hypotheses about the world.

Lets say each day we have the same n experts that predict from the set {fed, eaten}. Or perhaps a less morbid approach would be the example of predicting whether a given day will be sunny or rainy. In this case, the experts predict from {rainy, sunny}. After they make their prediction, the algorithm makes its own prediction, and then finds out if it actually it rains. Our goal is to perform nearly as well as the best expert so far (being competitive with respect to the best single expert). To quantify these intuitions, we will need the concepts of *loss* and *regret*.

4.1 Loss and Regret

To formalize the No Regret view, let's define a loss function L where:

$$L(wrong) = 1$$

$$L(correct) = 0$$

Therefore, for the optimal play, we want to minimize

$$\sum_t L(p_t) \rightarrow 0$$

Optimal play may be impossible. In this formulation, we cannot guarantee any performance (David Hume view) because the algorithm can be arbitrarily bad if the problem is hard to predict or adversarial.

We'll assume that we have some experts that predict either 0 or 1. Remember, the experts are arbitrary and can be based on a number of things. Examples:

- Always constant
- Markov expert (repeats what it last saw)
- Random (flips a coin and decides the state of the world)
- Something complex (consults the color of the sky)
- Adversarial expert (an expert which somehow knows what state the world is going to be in and predicts the opposite)
- etc..

If we consult N experts, and we want to do nearly as well as the best expert, we can define regret as:

$$Regret = \sum_t [L_t(\text{algorithm}) - L_t(e^*)]$$

where e^* is the best expert at time t

Therefore, our goal is to make regret scale more slowly than time, such that:

$$\lim_{t \rightarrow \infty} \frac{Regret}{t} = 0$$

An algorithm that satisfies the above formula is called “no regret.” Note that even though we cannot minimize loss, we can minimize regret. A side effect of this statement is that a no regret algorithm can still have high loss. This situation occurs when David Hume’s pessimism reigns and no potential model for the world is very good. Minimizing regret in this case is like saying that we can do “as good as the best of the bad things.” However, in situations where no regret results in low loss, Einstein wins out and we end up with a solution that models the world well and can make good predictions.

We will now discuss three attempts to minimize regret: follow the leader, majority vote, and weighted majority.

4.2 Try #1: Follow the Leader

We can also call this algorithm the “best in hindsight” algorithm. We naively pick the expert that has done the best so far and use that as our strategy for timestep t . Unfortunately, blindly picking the leader can lead to overfitting.

Let us assume there are 2 experts :

$e_0 \mapsto$ which always predicts 0

$e_1 \mapsto$ which always predicts 1

Also, let us assume world is such that is alternatively toggles between 0 and 1 with 1 being the initial state. Let the algorithm be designed such that whenever there is the algorithm picks e_0 .

<i>Correct count for e_0</i>	<i>Correct count for e_1</i>	<i>Prediction (Best expert in hindsight)</i>	<i>World State</i>
0	0	e_0	1
0	1	e_1	0
1	1	e_0	1
1	2	e_1	0
.	.	.	.
.	.	.	.
.	.	.	.

From table 1 it can be seen that the algorithm in this case is performing worse than each of the experts. In this example we hope we can somehow take the opinion of all the experts into consideration while making a decision. This is exactly what the next suggested algorithm does.

4.3 Try #2: Majority Vote

Given many experts, this algorithm goes with the majority vote. In a sense, it never gets faked out by observations, because it doesn't pay attention to them. In a statistical sense, we can say that this algorithm is "high bias, low variance," since it is robust but not adaptive, while algorithm #1 is "low bias, high variance," since it is adaptive but not robust. This algorithm can clearly perform much worst than the best expert as it may be the case that majority of the experts are always wrong but one is always correct. To overcome this fallacy we need a method that adapts and decides what experts to respect while predicting, based on the past performance of the experts. To this end we try weighted majority algorithm.

4.4 Try #3: Weighted Majority

Another approach is to notice that at each timestep, each expert is either right or wrong. If it's wrong, we can diminish the voting power of that particular expert when predicting the next majority vote.

More formally, this algorithm maintains a list of weights w_1, \dots, w_n (one for each expert x_1, \dots, x_n), and predicts based on a weighted majority vote, penalizing mistakes by multiplying their weight by half.

Algorithm

1. Set all the weights to 1.
2. Predict 1 (rain) if $\sum_{x_i=1} w_i \geq \sum_{x_i=0} w_i$, and 0 otherwise.

3. Penalize experts that are wrong: for all i s.t. x_i made a mistake, $w_i^{t+1} \leftarrow \frac{1}{2}w_i^t$.
4. Goto 2

Analysis of Algorithm The sum of the weights is $w \leq \left(\frac{3}{4}\right)^m n = \left(\frac{4}{3}\right)^{-m} n$, where m is the number of mistakes. The weight of the best expert $w_i^* = \left(\frac{1}{2}\right)^{m^*} = 2^{-m^*} \leq w$. Therefore,

$$2^{-m^*} \leq w \leq \left(\frac{4}{3}\right)^{-m} n.$$

Taking the \log_2 and solving for m gives,

$$m \leq \frac{m^* + \log_2 n}{\log_2 \frac{4}{3}} = 2.41(m^* + \log_2 n)$$

Thus, the number of mistakes by this algorithm is bounded by m^* plus an amount logarithmic in the number of experts, n .

To get a more intuitive feel for this last inequality, imagine the case when one expert makes *no* mistakes. Due to the binary search nature of the algorithm, it will take $\log_2 n$ iterations to “find” it. Furthermore, if we keep adding good experts, the term m^* will go down, but the term $\log_2 n$ will rise (a fair trade in this case). Adding bad experts will probably not change m^* much and will serve only to add noise, showing up in a higher $\log_2 n$. Thus, we can see the trade-offs involved when adding more experts to a system.

4.5 Randomized Weighted Majority Algorithm

In this algorithm, we predict each outcome with by picking an expert with a probability proportional to its weight. We also penalize each mistake by β instead of by one-half.

Algorithm

1. Set all the weights to 1.
2. Choose expert i in proportion to w_i .
3. Penalize experts that are wrong: for all i s.t. x_i made a mistake, $w_i^{t+1} \leftarrow \beta w_i^t$.
4. Goto 2.

Analysis The bound is

$$E[m] \leq \frac{m^* \ln(1/\beta) + \ln(n)}{1 - \beta}.$$

It is important to note here that the expectation in the expression is caused by probabilistic nature of the method of selecting an expert for prediction. Also note that since we are picking a single

expert, we don't have to worry about figuring out a way to combine multiple experts and can easily work in a space where the experts are predicting abstract things like whether the object seen is a phone, ocean or moon.

Generally, we want β to be small if we only have a few time steps available and vice-versa. Although β can also be time-varying, for example if you start with a large number of experts with a belief that some of the experts are good at predicting the world at the next timestep. In such a case we would want the β to be high initially to reach to the correct experts quickly and then decrease the value of β with time to avoid over-fitting. This algorithm serves to thwart an adversarial world that might know our strategy for picking experts, or any potential collusion between malicious experts and the world.

5 Introduction to Next Lecture

5.1 General Weighted Majority Algorithm

Algorithm

1. Set all the weights to 1.
2. Predict expert i in proportion to w_i .
3. Receive the correct value y_t .
4. Adjust weights s.t. $w_i^{t+1} \rightarrow w_i^t \exp^{-\epsilon l(i, y_t)} \forall i$, where l is the lost function, and ϵ is the penalizer.
5. Goto 2.

Analysis The bound is

$$E[R] \leq \epsilon \sum l(i^*) + \frac{1}{\epsilon} \ln(n),$$

where R is the regret. If ϵ is large, we learn quickly. If ϵ is small, we learn slowly but we'll have little regret.

5.2 Convex Sets and Convex Function

In a convex set, the line between any two points in the set is also in the set. A convex function can be defined s.t.

$$f(ax_1 + bx_2) \leq af(x_1) + bf(x_2),$$

where $a + b = 1$, and $a, b \geq 0$.

6 Suggested Readings

1. Avrim Blum Online Learning survey (<http://www.cs.cmu.edu/~avrim/Papers/survey.ps.gz>)
2. Probability: The Logic of Science, <http://bayes.wustl.edu/etj/prob/book.pdf>, Chapters 1 and 2 (you can skim over the derivations); Probabilistic Robotics, Chapters 1 and 2