Question 1: (6 Marks) If X, X2, ..., Xn Be A Random Sample From Bernoulli (p). 1. Prove That The Pmf

Question 1: (6 Marks) If X, X2, ..., Xn Be A Random Sample From Bernoulli (p). 1. Prove That The Pmf

Understanding the probability mass function (pmf) of a Bernoulli distribution is fundamental in probability theory and statistics, especially when dealing with binary outcomes. In this article, we will explore the derivation of the pmf for a Bernoulli distribution, considering a random sample X_1, X_2, ..., X_n drawn independently from Bernoulli (p). We will systematically prove the pmf and discuss its significance in statistical modeling.

Understanding the Bernoulli Distribution

What Is a Bernoulli Distribution?

The Bernoulli distribution is one of the simplest discrete probability distributions. It models a random experiment with exactly two possible outcomes, often labeled as success (1) and failure (0). The probability of success is denoted by p, where $0 \le p \le 1$, and consequently, the probability of failure is (1 - p).

Mathematically, the Bernoulli distribution is used to describe random variables X that take values in {0, 1} with the following probabilities:

- P(X = 1) = p
- P(X = 0) = 1 - p

Definition of the pmf for Bernoulli

The probability mass function (pmf) for a Bernoulli random variable X is expressed as:

pmf of X:

```
P(X = x) = p^x  (1 - p)^{1 - x},  where x ∈ {0, 1}
```

This compact form encapsulates both outcomes in a single formula, highlighting the distribution's simplicity and elegance.

Sampling from Bernoulli Distribution

Independent and Identically Distributed Samples

Suppose we have a sample of size n, denoted as $X_1, X_2, \ldots, X_n$, where each X_i is an independent Bernoulli(p) random variable. The assumption of independence and identical distribution (iid) is crucial because it simplifies the derivation of the joint pmf and allows us to analyze the collective behavior of the sample.

Joint pmf of the Sample

Given the independence of the samples, the joint pmf is the product of the individual pmfs:

$$P(X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n) = \prod_{i=1}^n P(X_i = x_i)$$

Since each X_i follows the same Bernoulli(p) distribution, this becomes:

$$P(X_1 = x_1, \ldots, X_n = x_n) = \prod_{i=1}^n p^{x_i} (1 - p)^{1 - x_i}$$

Deriving the pmf for the Sum of Bernoulli Random Variables

Sum of Bernoulli Variables as a Binomial Distribution

The sum of n independent Bernoulli(p) variables, denoted as $S = X_1 + X_2 + \ldots + X_n$, follows a Binomial distribution with parameters n and p. This sum counts the total number of successes in n trials.

To find the pmf of S, we consider the probability that exactly k of the n variables are successes (i.e., equal to 1). The pmf of S is then expressed as:

$$P(S = k) = C(n, k) \; p^k \; (1 - p)^{n - k}$$

where $C(n, k) = n! / (k! (n - k)!)$ is the binomial coefficient, representing the number of ways to choose k successes out of n trials.

Connecting Back to the pmf of X_i

Since each X_i is Bernoulli(p), the probability of observing a specific sequence with k successes and (n - k) failures is:

- Number of such sequences: C(n, k)
- Probability of any specific sequence: $p^k (1 - p)^{n - k}$

Multiplying these together, we get the binomial pmf, which sums the probabilities over all sequences with k successes, leading to the conclusion that the sum S follows Binomial(n, p).

Proving the pmf of the Bernoulli Distribution

Step-by-Step Proof

1. **Identify the possible outcomes:** For a Bernoulli random variable X, outcomes are 0 (failure) and 1 (success).
2. **Assign probabilities:** P(X = 1) = p, P(X = 0) = 1 - p.
3. **Express the pmf mathematically:** For $x \in \{0, 1\}$, $P(X = x) = p^x (1 - p)^{1 - x}$.
4. **Verify the pmf properties:** Confirm that for $x \in \{0, 1\}$, $P(X = x) \geq 0$, and that the sum over all possible x equals 1:

```
P(X=0) + P(X=1) = (1 - p) + p = 1
```

Conclusion of the Proof

Since the pmf satisfies the non-negativity condition and sums to 1 over all outcomes, it is a valid pmf. Therefore, the Bernoulli distribution's pmf is confirmed as:

```
P(X = x) = p^x  (1 - p)^{1 - x},  for x ∈ {0, 1}
```

Implications of the pmf in Statistical Modeling

Use Cases of Bernoulli pmf

- Modeling binary outcomes such as success/failure, yes/no responses, or presence/absence data.
- Foundation for more complex models like the Bernoulli process, geometric, and negative binomial distributions.
- Basis for logistic regression and other classification algorithms in machine learning.

Importance in Inferential Statistics

Knowing the pmf allows statisticians to perform hypothesis testing, estimate parameters, and derive confidence intervals for p based on observed data. The Bernoulli pmf's simplicity facilitates analytical calculations and simulations alike.

Summary and Key Takeaways

- The Bernoulli distribution models binary random variables with two outcomes: success (1) and failure (0).
- The pmf of a Bernoulli random variable X with parameter p is $P(X = x) = p^x (1 - p)^{1 - x}$ for $x \in \{0, 1\}$.
- For a sample of n iid Bernoulli(p) variables, the sum follows a Binomial(n, p) distribution with pmf $P(S = k) = C(n, k) p^k (1 - p)^{n - k}$.
- The derivation of the pmf confirms the distribution's validity and provides a foundation for statistical inference and modeling.

In conclusion, the proof of the Bernoulli pmf is a vital step in understanding discrete probability distributions. Its simplicity and versatility make it an essential concept in probability theory, statistics, and various applied fields.

Frequently Asked Questions

What is the probability mass function (pmf) of a Bernoulli(p) random variable?

The pmf of a Bernoulli(p) random variable X is given by P(X=1) = p and P(X=0) = 1 - p, where 0 ≤ p ≤ 1.

How do you derive the pmf for a sum of independent Bernoulli(p) variables?

Since the variables are independent, the sum follows a Binomial distribution with parameters n and p. The pmf is P(S = k) = C(n, k) p^k (1 - p)^(n - k), where S = X_1 + X_2 + ... + X_n.

What is the importance of the binomial coefficient in the pmf of the sum of Bernoulli variables?

The binomial coefficient C(n, k) counts the number of ways to choose k successes out of n trials, reflecting the combinatorial aspect of the distribution.

Can you prove that the sum of n independent Bernoulli(p) variables follows a Binomial(n, p) distribution?

Yes. Since each Bernoulli variable has pmf p^x (1-p)^{1-x}, the sum's pmf is obtained by summing over all combinations of successes, leading to the binomial pmf: P(S = k) = C(n, k) p^k (1 - p)^{n - k}.

What assumptions are necessary for the pmf derivation of the sum of Bernoulli(p) variables?

The variables must be independent and identically distributed (i.i.d.) Bernoulli(p) variables to apply the binomial distribution derivation.

How does the pmf change if the Bernoulli parameters vary across the sample?

If the Bernoulli parameters differ, the sum no longer follows a binomial distribution. Instead, its pmf involves a convolution of individual Bernoulli pmfs, leading to a more complex, non-standard distribution.

Additional Resources

Question 1: (6 Marks) If X, X_2, ..., X_n Be A Random Sample From Bernoulli(p). Prove That The pmf

---

Unraveling the Bernoulli Distribution: A Rigorous Proof of its Probability Mass Function

Understanding the probability mass function (pmf) of the Bernoulli distribution is fundamental in probability theory and statistics. The Bernoulli distribution models binary outcomes—success or failure—making it a cornerstone in fields ranging from economics to machine learning. This article thoroughly explores the derivation of the pmf for a random sample drawn from a Bernoulli distribution with parameter p, providing a detailed, step-by-step proof suitable for academic and professional audiences.

Introduction to the Bernoulli Distribution

The Bernoulli distribution is the simplest discrete probability distribution, characterized by a single parameter p, which represents the probability of success (often coded as 1). Formally, a Bernoulli random variable X takes the value:

$$
X =
\begin{cases}
1, & \text{with probability } p \\[6pt]
0, & \text{with probability } 1 - p
\end{cases}
$$

where $0 < p < 1$.

If we have a random sample $X_1, X_2, \ldots, X_n$ from this distribution, each X_i is an independent and identically distributed (i.i.d.) Bernoulli(p) variable.

Fundamental Concepts and Notation

Before diving into the proof, let's clarify some key concepts:

- Sample Space: For a Bernoulli distribution, the sample space is $\{0, 1\}$.
- Independence: The X_i are independent, meaning the joint probability factors into the product of individual probabilities.
- Identical Distribution: Each X_i has the same distribution, Bernoulli(p).

The goal is to derive the joint pmf of the entire sample $(X_1, X_2, \ldots, X_n)$, which, due to independence, can be expressed as the product of individual pmfs.

Deriving the pmf of a Single Bernoulli Random Variable

The pmf definition

The probability mass function of a Bernoulli(p) variable X is defined as:

$$
P(X = x) = p^x (1 - p)^{1 - x}
$$

for $x \in \{0, 1\}$.

Why this form? Because:

- When $x = 1$, the pmf reduces to $p^1 (1 - p)^0 = p$.
- When $x = 0$, it becomes $p^0 (1 - p)^1 = 1 - p$.

This compact form encapsulates both cases, simplifying calculations and proofs.

Extending to the Sample: Joint pmf of $X_1, X_2, \dots, X_n$

Given the independence assumption, the joint pmf is the product of the individual pmfs:

$$
P(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \prod_{i=1}^n P(X_i = x_i)
$$

Since each X_i follows the same Bernoulli(p) distribution, this simplifies to:

$$
\prod_{i=1}^n p^{x_i} (1 - p)^{1 - x_i}
$$

or equivalently,

$$
\left(\prod_{i=1}^n p^{x_i} \right) \left(\prod_{i=1}^n (1 - p)^{1 - x_i} \right)
$$

Because exponents are multiplicative over products, we can rewrite this as:

$$
p^{\sum_{i=1}^n x_i} (1 - p)^{n - \sum_{i=1}^n x_i}
$$

This expression intuitively reflects that the probability depends only on the total number of successes $\sum_{i=1}^n x_i$.

The Distribution of the Summed Successes: Binomial Connection

The sum $S_n = \sum_{i=1}^n X_i$ follows a Binomial distribution with parameters n and \(p

\).

Why is this important?

Because, in many contexts, the likelihood of observing k successes in n trials—regardless of the specific sequence—is of interest. The pmf of S_n is:

$$
P(S_n = k) = \binom{n}{k} p^k (1 - p)^{n - k}
$$

for $k = 0, 1, 2, ..., n$.

Connecting to the joint pmf

The joint pmf over all sequences with exactly k successes is:

$$
P(\text{sequence with } k \text{ successes}) = p^k (1 - p)^{n - k}
$$

and the number of such sequences is $\binom{n}{k}$.

Thus, the probability of having exactly k successes in the entire sample, regardless of order, is:

$$
P(S_n = k) = \binom{n}{k} p^k (1 - p)^{n - k}
$$

which is the pmf of the Binomial distribution.

Formal Proof of the pmf for the Sample

Given the above, the joint pmf for the sample $(X_1, X_2, ..., X_n)$, considering that each sequence with k successes has the same probability, is:

$$
\boxed{
P(X_1 = x_1, ..., X_n = x_n) = p^{\sum_{i=1}^n x_i} (1 - p)^{n - \sum_{i=1}^n x_i}
}
$$

where each specific sequence corresponds to a particular combination of successes and failures, with the total successes being $\sum_{i=1}^n x_i$.

In conclusion:

- The pmf of the sample vector $(X_1, ..., X_n)$ is multi-point, assigning probability to each sequence based on the number of successes.

- When considering the number of successes k, the pmf simplifies to the binomial pmf:

$$
\boxed{
P(S_n = k) = \binom{n}{k} p^k (1 - p)^{n - k}
}
$$

This is a foundational result underpinning many statistical methods, including hypothesis testing and confidence interval construction for Bernoulli and binomial parameters.

Implications and Applications

The derivation of the pmf for a sample from Bernoulli(p) has critical applications:

- Modeling binary outcomes: In clinical trials, marketing, and survey research.
- Parameter estimation: Estimating p via maximum likelihood, which hinges on the binomial pmf.
- Hypothesis testing: For proportions, using the binomial distribution.
- Machine Learning: Binary classifiers and probabilistic models often depend on Bernoulli and binomial distributions.

Final Remarks

This detailed proof underscores the elegance and interconnectedness of fundamental probability distributions. The Bernoulli distribution's pmf is not just a simple formula; it embodies the essence of binary randomness and serves as a building block for more complex models. Recognizing the joint pmf's structure, especially its binomial nature, allows statisticians and researchers to interpret, analyze, and infer about binary data effectively.

In essence, the derivation exemplifies how individual Bernoulli trials combine to produce the binomial distribution—a cornerstone concept that continues to underpin statistical theory and practice.

[Question 1 6 Marks If X X2 Xn Be A Random Sample From Bernoulli P 1 Prove That The Pmf]

Question 1 6 Marks If X X2 Xn Be A Random Sample From Bernoulli P 1 Prove That The Pmf

Related Articles

- Suppose Someone In A Collaborative Discussion Says That The Theme Of Roll Of Thunder, Hear My Cry Is
- Suppose You Invest In An Asset For Three Years. The Return Earned In Year 1, 2 And 3 Is

20%, -8% And

- One Batch Of Trail Mix Can Be Made By Mixing 3 Cups Of Granola With 0.75 Cup Of Raisins. Hamid Needs

Back to Home