Accurately computing running variance

The most direct way of computing sample variance or standard deviation can have severe numerical problems. Mathematically, sample variance can be computed as follows.

$$\sigma^2 = \frac{1}{n(n-1)}\left(n\sum_{i=1}^{n} x_i^2 - \left(\sum_{i=1}^{n} x_i\right)^2\right)$$

The most obvious way to compute variance then would be to have two sums: one to accumulate the sum of the *x*'s and another to accumulate the sums of the squares of the *x*'s. If the *x*'s are large and the differences between them small, direct evaluation of the equation above would require computing a small number as the difference of two large numbers, a red flag for numerical computing. The loss of precision can be so bad that the expression above evaluates to a *negative* number even though variance is always positive. See Comparing three methods of computing standard deviation for examples of just how bad the above formula can be.

There is a way to compute variance that is more accurate and is guaranteed to always give positive results. Furthermore, the method computes a running variance. That is, the method computes the variance as the *x*'s arrive one at a time. The data do not need to be saved for a second pass.

This better way of computing variance goes back to a 1962 paper by B. P. Welford and is presented in Donald Knuth's Art of Computer Programming, Vol 2, page 232, 3rd edition. Although this solution has been known for decades, not enough people know about it. Most people are probably unaware that computing sample variance can be difficult until the first time they compute a standard deviation and get an exception for taking the square root of a negative number.

It is not obvious that the method is correct even in exact arithmetic. It's even less obvious that the method has superior numerical properties, but it does. The algorithm is as follows.

> Initialize $M_1 = x_1$ and $S_1 = 0$.
>
> For subsequent x's, use the recurrence formulas
>
> $M_k = M_{k-1} + (x_k - M_{k-1})/k$
> $S_k = S_{k-1} + (x_k - M_{k-1})*(x_k - M_k)$.
>
> For $2 \le k \le n$, the k^{th} estimate of the variance is $s^2 = S_k/(k-1)$.

The C++ class `RunningStat` given below uses this method to compute the mean, sample variance, and standard deviation of a stream of data. This code sample shows how to use the class.

```
int main()
    {
        RunningStat rs;

        rs.Push(17.0);
        rs.Push(19.0);
        rs.Push(24.0);

        double mean = rs.Mean();
        double variance = rs.Variance();
        double stdev = rs.StandardDeviation();
    }
```

As new values arrive, say from a simulation, they enter the `RunningStat` class via the `Push` method. To check on the mean, variance, or standard deviation at any time, call the corresponding methods.

The source code for the `RunningStat` class follows:

```
class RunningStat
    {
    public:
        RunningStat() : m_n(0) {}
```

```
    void Clear()
    {
        m_n = 0;
    }

    void Push(double x)
    {
        m_n++;

        // See Knuth TAOCP vol 2, 3rd edition, page 232
        if (m_n == 1)
        {
            m_oldM = m_newM = x;
            m_oldS = 0.0;
        }
        else
        {
            m_newM = m_oldM + (x - m_oldM)/m_n;
            m_newS = m_oldS + (x - m_oldM)*(x - m_newM);

            // set up for next iteration
            m_oldM = m_newM;
            m_oldS = m_newS;
        }
    }

    int NumDataValues() const
    {
        return m_n;
    }

    double Mean() const
    {
        return (m_n > 0) ? m_newM : 0.0;
    }

    double Variance() const
    {
        return ( (m_n > 1) ? m_newS/(m_n - 1) : 0.0 );
    }

    double StandardDeviation() const
    {
        return sqrt( Variance() );
    }

private:
    int m_n;
    double m_oldM, m_newM, m_oldS, m_newS;
};
```

Here are a couple references on computing sample variance.

Chan, Tony F.; Golub, Gene H.; LeVeque, Randall J. (1983). Algorithms for Computing the Sample Variance: Analysis and Recommendations. The American Statistician 37, 242-247.

Ling, Robert F. (1974). Comparison of Several Algorithms for Computing Sample Means and Variances. Journal of the American Statistical Association, Vol. 69, No. 348, 859-866.

Update: See this page for an extension of this code that supports computing skewness and kurtosis as well. Also, the new code supports combining `RunningStat` objects via the + and += operators.

John D. Cook, PhD, President

My colleagues and I have decades of consulting experience helping companies solve complex problems involving math, statistics, and computing.

Email address

Your company's project ...

SEND

Go ahead and send us a note. We look forward to exploring the opportunity to help your company too.

JOHN D. COOK

Search ...

SEARCH