Journal of Universal Computer Science, vol. 1, no. 7 (1995), 504-513
submitted: 15/12/94, accepted: 26/6/95, appeared: 28/7/95©Springer Pub. Co.

Exact Statistics and Continued Fractions

David Lester
dlester@cs.man.ac.uk
Functional Programming Group,
Department of Computer Science, Manchester University,
Oxford Road, Manchester M13 9PL, UK.

Abstract: In this paper we investigate an extension to Vuillemin's work on continued fraction arithmetic [Vuillemin 87, Vuillemin 88, Vuillemin 90], that permits it to evaluate the standard statistical distribution functions. By this we mean: the normal distribution, the χ^2-distribution, the t-distribution, and, in particular, the F-distribution. The underlying representation of non-rational computable real numbers is also as continued fractions, in the style of Vuillemin. This permits arbitrary accuracy over a range of values. The number of terms of a continued fraction that are used by the implementation is dynamically controlled by the accuracy demanded of the final answer.
The use of a modern lazy functional language – Haskell – has considerably eased the programming task. Two features are of note. Firstly, the *type-class* structure allows one to augment the varieties of numbers supported by the language. Secondly, the *laziness* inherent in the Haskell's semantics, makes it very straightforward to dynamically control the accuracy of the intermediate evaluations.

1 Introduction

In this paper we investigate the use of continued fractions to evaluate the standard statistical distribution functions, by which is meant: the normal distribution, the χ^2-distribution, the t-distribution, and the F-distribution. The underlying representation of non-rational computable real numbers themselves is also as continued fractions, in the style of Vuillemin [Vuillemin 87, Vuillemin 88, Vuillemin 90].

The novelty of the work presented in the paper lies in the following areas.

- We use Vuillemin's work in a setting (the calculation of the incomplete beta function) which is highly unstable, and demonstrate that his system is feasible.
- We show that three continued fraction expansions, not discussed by Vuillemin, are nevertheless compatible with his system.
- We show that Gauss' continued fraction may be safely used for a range of values besides those already described by Vuillemin.

The use of Vuillemin's system permits arbitrary accuracy over a range of values. The number of terms of the continued fraction that are used by the implementation is dynamically controlled by the accuracy demanded of the final answer. Here, we exploit the Laziness of mordern functional languages – in this case Haskell [Hudak et al 88].

The use of a modern lazy functional language – Haskell – has considerably reduced the amount of code required; currently amounting to just less than 1000 lines. This includes: an implementation of the rationals, Gosper's continued fraction package, Vuillemin's transcendental function package, and the extra functions for the statistical functions. Two further features are of note. Firstly, the *type-class* structure allows one to augment the varieties of numbers supported by the language. Secondly, the laziness inherent in the Haskell's semantics, makes it very straightforward to dynamically control the accuracy of the intermediate evaluations.

The most interesting problems arise whilst considering the incomplete beta function, which is required to implement Fisher's F-distribution [Fisher 22]. Lackritz [Lackritz 84] has provided a simple method of evaluating this distribution, provided we deal only with small integral degrees of freedom ν_1 and ν_2. The complexity of the evaluation of Lackritz's finite series is $O(N^2)$, where $N = \nu_1 + \nu_2$. We therefore require a method for non-integral ν_1 and ν_2 and for those situations where $\nu_1 + \nu_2$ is large, say 10^5 or more.

As Mardia and Zemroch have observed [Mardia and Zemroch 78], we are mainly concerned with rational-valued degrees of freedom when we consider non-parametric statistical testing. That is, we no longer assume that we are sampling from normally distributed populations. Examples are: the Kruskal-Wallis Test [Kruskal and Wallis 52], and Box's test of Homoscedasticity [Box 49].

We begin by considering the continued fractions that we will be using.

2 Continued Fractions

Continued fractions were much favoured by applied mathematicians of previous centuries, because they provided a relatively straightforward calculating procedure in a pre-computer age (see Gauss and Laplace amongst others) [Gauss 1812, Laplace 1805]. The analytic properties of Gauss' work are well described by Wall [Wall 48]. However, with the development of floating point software and hardware, attention shifted away from them.

A number of researchers are undertaking a reappraisal of algorithms associated with continued fractions. These tend to fall into two categories: firstly those using continued fraction algorithms with floating point implementations of real numbers, and secondly, those using continued fractions to represent real numbers. In this paper I shall be using both techniques to outline some extensions to the Haskell [Hudak et al 88] continued fraction package described in [Lester 92] so that it may perform analysis on standard statistical distributions.

In the field of statistical distribution tabulation – and in particular that of the F-distribution – much use has been made of continued fraction algorithms [Tretter and Walster 79, Tretter and Walster 80]. Unfortunately, it is not possible to make naïve use of these continued fractions for floating point calculations, as cancellation of terms results in a fatal loss of accuracy. To this end, Tretter and Walster use MACSYMA to perform symbolic simplifications on their continued fraction. Brown and Levy [Brown 94] use these simplifications to validate DiDonato and Morris' floating point Algorithm 708 for the incomplete beta function [DiDonato and Morris 92], against a Maple implementation.

The use of continued fractions as a representation of real numbers has been facilitated by two developments. Firstly, Gosper has shown how to manipulate continued fraction represention of real numbers to perform basic arithmetic [Gosper 80, Gosper 81]. His work was subsequently extended to deal with transcendental functions by Vuillemin [Vuillemin 87, Vuillemin 88, Vuillemin 90].

We begin with a review of nineteenth century work on continued fraction algorithms.

3 Gauss' Hypergeometric Function

In the impressively titled paper: "*Disquisitiones Generales Circa Seriem Infinitam*

$$1+\frac{\alpha\beta}{1\,.\,\gamma}x+\frac{\alpha(\alpha+1)\beta(\beta+1)}{1\,.\,2\,.\,\gamma(\gamma+1)}xx+\frac{\alpha(\alpha+1)(\alpha+2)\beta(\beta+1)(\beta+2)}{1\,.\,2\,.\,3\,.\,\gamma(\gamma+1)(\gamma+2)}x^3+\text{ etc.}$$

Pars prior" [Gauss 1812], Gauss outlines the properties of the hypergeometric function. Formally, it can be defined by Equation 1[1].

$$ {}_2F_1(\alpha,\beta,\gamma;z)=1+\sum_{n=1}^{\infty}\left(\prod_{i=1}^{n}\frac{(\alpha+i-1)(\beta+i-1)}{i(\gamma+i-1)}\right)z^n \qquad (1)$$

This has a circle of convergence $|z|=1$; and its behaviour on this circle is:

- divergence whenever $\Re(\gamma-\alpha-\beta)\leq -1$;
- absolute convergence whenever $\Re(\gamma-\alpha-\beta)>0$;
- conditional convergence whenever $-1<\Re(\gamma-\alpha-\beta)\leq 0$, the point $z=1$ is excluded.

One of the properties mentioned by Gauss, occurs in the second section of his paper, where a continued fraction is defined for the ratio of a pair of hypergeometric functions.

$$\frac{{}_2F_1(\alpha,\beta+1,\gamma+1;x)}{{}_2F_1(\alpha,\beta,\gamma;x)}=\left[\frac{1}{1-}\,\frac{ax}{1-}\,\frac{bx}{1-}\,\frac{cx}{1-}\,\frac{dx}{1-}\,\cdots\right] \qquad (2)$$

Where

$$a=\tfrac{\alpha(\gamma-\beta)}{\gamma(\gamma+1)} \qquad b=\tfrac{(\beta+1)(\gamma+1-\alpha)}{(\gamma+1)(\gamma+2)}\ c=\tfrac{(\alpha+1)(\gamma+1-\beta)}{(\gamma+2)(\gamma+3)}$$

$$d=\tfrac{(\beta+2)(\gamma+2-\alpha)}{(\gamma+3)(\gamma+4)}\ e=\tfrac{(\alpha+2)(\gamma+2-\beta)}{(\gamma+4)(\gamma+5)}\ f=\tfrac{(\beta+3)(\gamma+3-\alpha)}{(\gamma+5)(\gamma+6)}$$

and, as Gauss observes: "*cuius lex progressionis obvia est*".

On its own, Equation 2 is not too useful; what we really require is not a continued fraction for a ratio of hypergeometric functions, but one for an individual function. Observe what happens when we substitute $\beta=0$ into Equation 2: the numerator becomes ${}_2F_1(\alpha,0,\gamma;x)$, which is 1. Hence Gauss provides a continued fraction expansion for ${}_2F_1(\alpha,1,\gamma;x)$.

Suitably massaged, this continued fraction can be represented as a potentially infinite sequence of rationals. The sequence will terminate if one of α, γ, or $\gamma-\alpha$ is a negative integer or 0.

$${}_2F_1(\alpha,1,\gamma+1;x)=\left[\frac{1}{1-}\,\frac{ax}{1-}\,\frac{bx}{1-}\,\frac{cx}{1-}\,\frac{dx}{1-}\,\cdots\right] \qquad (3)$$

[1] Here, and throughout, we shall refer to the Hypergeometric function as ${}_2F_1(\alpha,\beta,\gamma;x)$. The leading 2 subscript refers to the two parameters forming the numerator of each coefficient (α and β), and the 1 subscript refers to the one parameter associated with the denominator of each coefficient (γ).

where

$$a = \frac{\alpha\gamma}{\gamma(\gamma+1)} \qquad b = \frac{1(\gamma+1-\alpha)}{(\gamma+1)(\gamma+2)} \; c = \frac{(\alpha+1)(\gamma+1)}{(\gamma+2)(\gamma+3)}$$

$$d = \frac{2(\gamma+2-\alpha)}{(\gamma+3)(\gamma+4)} \; e = \frac{(\alpha+2)(\gamma+2)}{(\gamma+4)(\gamma+5)} \; f = \frac{3(\gamma+3-\alpha)}{(\gamma+5)(\gamma+6)}$$

In Haskell, we can express Gauss' continued fraction (for rational parameters `a` and `c`) as `gauss1CF a c` which is ${}_2F_1(\mathtt{a}, 1, \mathtt{c}; 1)$.

```
> gauss1CF :: QRational -> QRational -> [QRational]
> gauss1CF a c
>  = 0: 1: c/a: a*(c+1)/ca: cf (ca/(a*(a+1)*c)) 0
>    where cf :: QRational -> Integer -> [QRational]
>          cf p n = (c+(2*n+2)%%1)*p:
>                   (c+(2*n+3)%%1)/(p*(ca+n1)*n2):
>                   cf (p*(ca+n1)*n2/((a+n2)*(c+n1))) (n+1)
>                   where n1 = fromInteger (n+1) :: QRational
>                         n2 = fromInteger (n+2) :: QRational
>          ca = c-a
```

To evaluate ${}_2F_1(a, 1, c; x)$ at a particular value of x we simply divide alternate terms of `gauss1CF` by $-x$, evaluate the resultant continued fraction, and multiply this result by $\frac{-1}{x}$. This can be made to work for real valued x not just rational values.

3.1 Kummer's Confluent Hypergeometric Function

If, in the definition of the hypergeometric function, we let $\beta \to \infty$ as $x \to \frac{x}{\beta}$ we obtain Kummer's Function.

$$ {}_1F_1(\alpha, \gamma; z) = \lim_{\beta\to\infty} {}_2F_1(\alpha, \beta, \gamma; \frac{z}{\beta}) = 1 + \sum_{n=1}^{\infty} \left(\prod_{i=1}^{n} \frac{(\alpha + i - 1)}{i(\gamma + i - 1)} \right) z^n \qquad (4)$$

For our purposes, it suffices to consider ${}_1F_1(\alpha, \gamma; z)$, with $\alpha = 1$ and $\gamma \geq 1$. Given the value of γ we may calculate ${}_1F_1(1, \gamma; 1)$ as `phiCF1` γ.

```
> phiCF1 :: QRational -> [QRational]
> phiCF1 c
>  = 0: 1: cf (-1%%1) 0
>    where cf :: QRational -> Integer -> [QRational]
>          cf p n = (c+fromInteger (2*n)) * p :
>                   ((c+fromInteger (2*n+1)) /
>                             (p*fromInteger n')):
>                   cf p' n'
>                   where p' = -p * (1/c')* fromInteger n'
>                         n' = n+1
>                         c' = c + fromInteger n
```

Once more the calculation of ${}_1F_1(1, \gamma; z)$ involves dividing alternate terms by z and dividing the result by z.

4 The Normal Distribution

This, the simplest of the standard statistical distribution functions we shall consider, is defined as the integral in Equation 5.

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{1}{2}x^2} dx \tag{5}$$

To obtain a continued fraction, we need to consider the *Mills' ratio*, of Equation 6.

$$R(x) = \frac{1-\Phi(x)}{\Phi'(x)} = e^{\frac{1}{2}x^2} \int_{x}^{\infty} e^{-\frac{1}{2}x^2} dx \tag{6}$$

Laplace was the first to give a continued fraction for this function in [Laplace 1805]. It occurred as a rejected hypothesis for the thermally induced atmospheric distortion that the astronomers of the day were observing.

$$R(x) = \left[\frac{1}{x+} \; \frac{1}{x+} \; \frac{2}{x+} \; \frac{3}{x+} \; \cdots \; \frac{n}{x+} \; \cdots \right] \tag{7}$$

This continued fraction – Equation 7 – suffers from a low rate of convergence for small x. For this reason we use Shenton's continued fraction – Equation 8 – instead [Shenton 54].

$$\bar{R}(x) = \sqrt{\frac{\pi}{2}} e^{\frac{1}{2}x^2} - R(x) = \left[\frac{x}{1-} \; \frac{x^2}{3+} \; \frac{2x^2}{5-} \; \frac{3x^2}{7+} \; \cdots \right] \tag{8}$$

Once more, this algorithm will work for x represented as a continued fraction as well as rational x.

Oddly enough, the most expensive part of the calculation involves the calculation of $\sqrt{\pi}$ = `exp (log pi/2)` rather than that of the Mills' ratio. Since $\Gamma(\frac{1}{2}) = \sqrt{\pi}$, this leads nicely into the next section, where we obtain a more efficient algorithm to perform this calculation.

5 The Gamma Function

The *Gamma function* $\Gamma(\alpha)$ is Euler's generalization of the factorial function on integers. Although we won't be using it at complex values, this function is analytic over the entire complex plane, except at the points $\alpha = 0, -1, -2, \ldots$. It can be defined by the integral given in Equation 9.

$$\Gamma(\alpha) = \int_{0}^{\infty} x^{\alpha-1} e^{-x} dx \tag{9}$$

Provided that $|\alpha|$ is reasonably small, *i.e.* less than 10^5, the first part of the implementation uses Equation 10 to adjust the value of α to lie within $0 < \alpha < 2$.

$$\Gamma(\alpha+1) = \alpha\Gamma(\alpha) \tag{10}$$

Clearly, if α is one of $0, -1, -2, \ldots$, then we will be using the formula in reverse and will eventually end by dividing by zero.

Associated with the Gamma function are the pair of *incomplete Gamma functions* : $\gamma(\alpha, x)$ and $\Gamma(\alpha, x)$ defined in Equations 11 and 12.

$$\gamma(\alpha, x) = \int_0^x x^{\alpha-1} e^{-x} dx \qquad \text{for } 0 \leq x \tag{11}$$

$$\Gamma(\alpha, x) = \int_x^\infty x^{\alpha-1} e^{-x} dx \qquad \text{for } 0 < x \tag{12}$$

Obviously the following relation holds between the incomplete Gamma functions and the Gamma function proper.

$$\Gamma(\alpha) = \gamma(\alpha, x) + \Gamma(\alpha, x) \qquad \text{for } 0 < x \tag{13}$$

For this to be a viable exact method for the calculation of $\Gamma(x)$ we seek continued fraction expansions for $\gamma(\alpha, x)$ and $\Gamma(\alpha, x)$; these are given in Equations 14 and 15.

$$\gamma(\alpha, x) = \frac{x^\alpha e^{-x}}{\alpha} \, {}_1F_1(1, \alpha + 1, x) \qquad \text{for } 0 < x \tag{14}$$

Recall that ${}_1F_1$ is the confluent hypergeometric function of Kummer.

$$\Gamma(\alpha, x) = x^\alpha e^{-x} \left[\frac{1}{x+} \; \frac{1-\alpha}{1+} \; \frac{1}{x+} \; \frac{2-\alpha}{1+} \; \frac{2}{x+} \; \cdots \right] \qquad \text{for } 0 < x \tag{15}$$

In particular, for computational purposes we select $x = 1$. This eliminates the need to calculate x^α – which given that both α and x might, in general, not be integral – leads to a considerable optimization.

$$\Gamma(\alpha) = e^{-1} \left(\frac{1}{\alpha} \, {}_1F_1(1, \alpha + 1, 1) + \left[\frac{1}{1+} \; \frac{1-\alpha}{1+} \; \frac{1}{1+} \; \frac{2-\alpha}{1+} \; \frac{2}{1+} \; \cdots \right] \right) \tag{16}$$

As a corollary we can compute the *Beta function* $B(p, q)$ which is defined as:

$$B(p, q) = \frac{\Gamma(p)\Gamma(q)}{\Gamma(p+q)} \tag{17}$$

One optimization that is possible in the calculation of $B(p, q)$, is to avoid two of the divisions by e.

The problem with direct use of Equation 17 occurs when p and q are both large – say 10^{50} and 10^{49}. In these circumstances it is no longer practical to use the recurrence relations to simplify the gamma functions. We will explore an alternative at the end of Section 7.

6 The χ^2-Distribution

One direct use that can be made of our work on the Γ functions is that it is now very easy to generate exact values for the χ^2-distribution.

$$Q(\chi^2, \nu) = 1 - \frac{\gamma(\frac{\nu}{2}, \frac{\chi^2}{2})}{\Gamma(\frac{\nu}{2})} \quad \text{where } 0 \leq \chi^2 \text{ and } 0 < \nu \tag{18}$$

Notice that for small positive integral values of $\frac{\nu}{2}$, we may use the alternative definition of Equation 19.

$$Q(\chi^2, \nu) = \frac{\Gamma(\frac{\nu}{2}, \frac{\chi^2}{2})}{(\frac{\nu}{2} - 1)!} \tag{19}$$

Expanding the continued fraction of Equation 15 results in the finite series of Lackritz [Lackritz 83].

7 The F-distribution

There are some interesting features in the calculation of the F-distribution. For example, Mardia and Zemroch [Mardia and Zemroch 78] show that $F = 8.4872 \times 10^{78}$ to obtain $Q(F, 120, 0.1) = 0.0001$. This value of F is, curiously, close to the limit of the maximum floating point number available on their machine. The points at which Brown and Levy [Brown 94] were able to test DiDonato and Morris' algorithm was also artificially restricted, by floating point limitations in their generation routine.

In [Lackritz 84] Lackritz shows how to obtain exact p values for F and t tests, provided that we are dealing with small integral valued degrees of freedom. The following approach gives exact results for rational valued degrees of freedom.

We begin with a definition: the probability of random variable – F-distributed with ν_1 and ν_2 degrees of freedom – being greater than F is $Q(F|\nu_1, \nu_2)$, defined in Equation 20

$$Q(F|\nu_1, \nu_2) = I_x(\frac{\nu_2}{2}, \frac{\nu_1}{2}) \qquad \text{where } x = \frac{\nu_2}{\nu_2 + \nu_1 F} \text{ and } 0 \leq F,\ 0 < \nu_i \tag{20}$$

The normalized incomplete beta function: $I_x(p, q)$, is given in Equation 21.

$$I_x(p, q) = \frac{B_x(p, q)}{B(p, q)} \tag{21}$$

The incomplete beta function is related to Guass' Hypergeometric in Equation 22.

$$B_x(p, q) = \frac{x^p}{p}\ {}_2F_1(p, 1 - q, p + 1; x) \tag{22}$$

However, using one of Gauss' transformations:

$${}_2F_1(\alpha, \beta, \gamma; x) = (1 - x)^{-\beta}\ {}_2F_1(\beta, \gamma - \beta, \gamma; \frac{x}{x - 1})$$

we obtain:

$$B_x(p,q) = \frac{x^p(1-x)^{q-1}}{p} \, {}_2F_1(1-q, 1, p+1; \frac{x}{x-1}) \tag{23}$$

Expanding the hypergeometric function yields Equation 23 which is essentially the continued fraction of Mueller.

The use of this continued fraction is limited to those cases where $-1 \leq \frac{x}{x-1} < 1$ in order that the argument may lie within the radius of convergence of Gauss' Hypergeometric continued fraction. To ensure that this remains the case, whenever $x > \frac{1}{2}$ we use the reflection property of Equation 24.

$$I_x(p,q) = 1 - I_{1-x}(q,p) \tag{24}$$

It is possible – for large values of p and q ($> 10^5$) – for the method of calculating $B(p,q)$ to become too time consuming. For this reason, an alternative exact method for calculating $B(p,q)$ is required.

7.1 Student's *t*-Distribution

This is merely a special case of the F-distribution.

$$A(t|\nu) = 1 - I_x(\frac{\nu}{2}, \frac{1}{2}) \quad \text{where} x = \frac{\nu}{\nu + t^2} \text{ and } \nu > 0 \tag{25}$$

Once more, care needs to be taken to ensure that the value of x lies in the range $0 \leq x \leq \frac{1}{2}$. If it doesn't we must use the reflection formula of Equation 24.

8 Conclusion and Further Work

In this paper we have shown how to extend Vuillemin's work on continued fraction arithmetic [Vuillemin 87, Vuillemin 88, Vuillemin 90], to the evaluation of the standard statistical distribution functions. The number of terms of a continued fraction that are used is dynamically controlled by the accuracy demanded of the final answer.

Two features of Haskell have proved particularly helpful. Firstly, the *type-class* structure allows one to augment the varieties of numbers supported by the language. Secondly, the *laziness* inherent in the Haskell's semantics, makes it very straightforward to dynamically control the accuracy of the intermediate evaluations.

We have shown how to use Tretter and Walsters' continued fraction to evaluate the incomplete beta function [Tretter and Walster 79, Tretter and Walster 80]. DiDonato and Morris incorporate various improvements in their algorithm, but admit that, for small values of p and q their continued fraction requires many terms of the continued fraction to be evaluated [DiDonato and Morris 92]. Our implementation is capable of handling this problem, albeit at reduced speed.

The principle outstanding problem we have is when p and q are large; then the calculation of $B(p,q)$ becomes inordinately expensive. We are able to work with values of p and q less than 10^5. Improving this requires that we are able to give an *exact* value for the log Gamma function $\ln(\Gamma(1+\alpha))$.

Another avenue of investigation lies with the calculation of the probabilities associated with the non-central χ^2- and F-distributions.

References

[Box 49] G.E.P. Box. A general distribution theory for a class of likelihood criteria. *Biometrika*, 36:317–346, 1949.

[Brown 94] B.W. Brown and L.B. Levy. Certification of algorithm 708: Significant digit computation of the incomplete beta function ratios. *ACM Transactions on Mathematical Software*, 20(3):393–397, September 1994.

[DiDonato and Morris 92] A.R. DiDonato and A.H. Morris, Jr. Algorithm 708: Significant digit computation of the incomplete beta function ratios. *ACM Transactions on Mathematical Software*, 18(3):360–373, September 1992.

[Fisher 22] R.A. Fisher. The goodness of fit of regression formulae and distribution of regression coefficients. *Journal of the Royal Statistical Society*, 85:597–612, 1922.

[Gauss 1812] C.F. Gauss. Disquisitiones generales circa serium infinitam $1+\frac{\alpha\beta}{1\,.\,\gamma}x+\frac{\alpha(\alpha+1)\beta(\beta+1)}{1\,.\,2\,.\,\gamma(\gamma+1)}xx+\frac{\alpha(\alpha+1)(\alpha+2)\beta(\beta+1)(\beta+2)}{1\,.\,2\,.\,3\,.\,\gamma(\gamma+1)(\gamma+2)}x^3$ etc. pars prior. In *Werke*, volume 3, pages 123–162. Königlichen Gesellschaft der Wissenschaften, Göttingen, 1812.

[Gosper 80] R.W. Gosper. Continued fraction arithmetic. HAKMEM 101b, MIT, 1980.

[Gosper 81] R.W. Gosper. Continued fraction arithmetic. Unpublished Draft Paper, 1981.

[Gradshteyn and Ryzhik 94] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, Inc. (London) Ltd., 24/28 Oval Road, London NW1 7DX, 5th edition, 1994.

[Hudak et al 88] P. Hudak, P. Wadler, Arvind, B. Boutel, J. Fairbairn, J. Fasel, J. Hughes, T. Johnsson, D. Kieburtz, S. Peyton Jones, R. Nikhil, M. Reeve, D. Wise, and J. Young. Report on the functional programming language Haskell. Draft Proposed Standard, 18th December 1988.

[Kruskal and Wallis 52] W.H. Kruskal and W.A. Wallis. The use of ranks in one-criterion variance analysis. *Journal of the American Statistical Society*, 47:583–621, 1952.

[Lackritz 83] J.R. Lackritz. Exact p-values for Chi-Squared tests. In *Proceedings of the Section on Statistical Education*, pages 130–132. American Statistical Association, 1983.

[Lackritz 84] J.R. Lackritz. Exact p values for f and t tests. *The American Statistician*, 38(4):312–314, November 1984.

[Laplace 1805] P.S. Laplace. *Traité de Mécanique Céleste*, volume 4. Chez Courcier, Imprimeur-Libraire pour les Mathématiques, 71 Quai des Augustins, Paris, 1805.

[Lester 92] D.R. Lester. Vuillemin's exact real arithmetic. In R. Heldal, C.K. Holst, and P. Wadler, editors, *Functional Programming, Glasgow 1991*, pages 225–238, London, 1992. Springer-Verlag.

[Mardia and Zemroch 78] K.V. Mardia and P.J. Zemroch. *Tables of the F- and related Distributions with Algorithms*. Academic Press Inc. (London) Ltd., 24/28 Oval Road, London NW1 7DX, 1978.

[Shenton 54] L.R. Shenton. Inequalities for the normal integral including a new continued fraction. *Biometrika*, 41:177–189, 1954.

[Tretter and Walster 79] M.J. Tretter and G.W. Walster. Continued fractions for the incomplete beta function: Additions and corrections. *The Annals of Statistics*, 7(2):462–465, 1979.

[Tretter and Walster 80] M.J. Tretter and G.W. Walster. Analytic subtraction applied to the incomplete gamma and beta functions. *SIAM Journal on Scientific Statistical Computing*, 1(3):321–326, September 1980.

[Vuillemin 87] J. Vuillemin. Arithmétic réelle exacte par les fractions continues. Technical Report 760, Instituit National de Recherche en Informatique et en Automatique, Domaine de Voluceau, Roquencourt, BP105, 78153 Le Chesnay Cedex, France, November 1987.

[Vuillemin 88] J. Vuillemin. Exact real computer arithmetic with continued fractions. In *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, pages 14–27, Snowbird, Utah, 25–27 July 1988.

[Vuillemin 90] J.E. Vuillemin. Exact Real Computer Arithmetic with Continued Fractions. *IEEE Transactions on Computers*, 39(8):1087–1105, August 1990.

[Wall 48] H.S. Wall. *Analytic Theory of Continued Fractions*. Van Nostrand, Inc., 250 Fourth Avenue, New York 3, 1948.