Kim Boström, July 2007

Statistical Significance

1 Introduction

Given that we have a set $\tilde{x}$ of data and a desired level of significance s, we want a function $\mathrm{ttest}(\tilde{x}, s)$ which can be either zero or one. If the value is one then the function tells us that the sample mean of the dataset $\tilde{x}$ is *significantly different from zero*, at level s. Similarly, if we have two datasets $\tilde{x}_1, \tilde{x}_2$ then we want a function $\mathrm{ttest}(\tilde{x}_1, \tilde{x}_2, s)$ that tells us whether or not the two datasets have a significantly different mean, at level s.

Such functions are *hypothesis test functions*. They aim at testing whether or not some *null hypothesis* can be rejected on the basis of a given data set. The null hypothesis is assumed to be the stronger one: There has to be good evidence against the null hypothesis in order to reject it. Usually, the aim of the researcher is to reject the null hypothesis at a high level of significance, so that the alternative hypothesis, the negation of the null hypothesis, which is the hypothesis of the researcher himself, is to be accepted by the scientific community.

The theory of hypothesis testing is very elaborate and we can only give a very brief and rough sketch here.

In the first case above, the null hypothesis would be that the mean of the underlying random variable is equal to zero. In the second case the null hypothesis would be that the mean of the underlying random variables is equal.

2 Samples

While a random variable is an abstract concept, a *sample* is concrete collection of data. More precisely, a sample $\tilde{x} = (x_1, \ldots, x_N)$ is a finite sequence of N independent realizations of one and the same random variable $\hat{x}$. A real-valued function of a sample is called a *statistic*. The most important statistics of a sample $\tilde{x}$ are the *sample mean*,

$$\mu(\tilde{x}) := \frac{1}{N} \sum_{n=1}^{N} x_n, \tag{1}$$

and the *sample variance*

$$\sigma^2(\tilde{x}) := \frac{1}{N-1} \sum_{n=1}^{N} (x_n - \mu(\tilde{x}))^2. \tag{2}$$

These two statistics are so important because they are good *estimators* of the mean and the variance of the underlying random variable $\hat{x}$, meaning that for large sample size $N \to \infty$,

$$\mu(\tilde{x}) \to \mu(\hat{x}) \tag{3}$$
$$\sigma^2(\tilde{x}) \to \sigma^2(\hat{x}), \tag{4}$$

where $\mu(\hat{x})$ and $\sigma^2(\hat{x})$ are the mean and the variance of the random variable $\hat{x}$. It is crucial to keep in mind the distinction between the sample $\tilde{x}$ and the underlying random variable $\hat{x}$. A random variable $\hat{x}$ is defined by an *apriori* probability distribution $p(x)$, while the corresponding sample $\tilde{x}$ is a set of *aposteriori* data. We can *estimate* the probability $p(x)$ of the random variable $\hat{x}$ taking the value x by

$$p(x) = \frac{n(x)}{N}, \tag{5}$$

where $n(x)$ is the number of occurrences of the value x in the sample $\tilde{x}$ and N is the sample size, that is, the number of elements in $\tilde{x}$. If $\hat{x}$ is a continuous random variable, however, then any particular x in the sample will occur at most once and an estimation is impossible. Thus, in the real world we can only estimate a *discretization* of $\hat{x}$ which is not unique and depends on several free parameters like the binning size and the upper and lower cutoff. Also, estimating the probability distribution from the data requires a very large sample size. On the other hand, estimating particular functions like the mean and the variance can be accomplished already with relatively small sample sizes. This is why mostly one is not interested in the probability distribution of the underlying random variable and rather concentrates on certain statistics, that is, on certain functions of the dataset.

3 Estimators

As we have already seen, the sample mean and the sample variance are good estimators for the mean and the variance of the underlying random variable. But what means a "good" estimator?

Actually, an *estimator* is not a statistic, that is, a real-valued function of a sample, but rather it is a random variable. Say we have a statistic $\theta = \theta(\tilde{x}) = \theta(x_1, \ldots, x_N)$ of a sample $\tilde{x} = (x_1, \ldots x_N)$ of N realizations of the random variable $\hat{x}$. Then the corresponding estimator of θ is defined as

$$\hat{\theta} := \theta(\hat{x}_1, \ldots, \hat{x}_N), \tag{6}$$

where

$$\hat{x}_1 = \ldots = \hat{x}_N = \hat{x}. \tag{7}$$

The motivation is the folllowing. If I have a random variable $\hat{x}$ and I let them realize N times then this sample of N realizations will be different whenever I repeat the same procedure and let $\hat{x}$ realize another N times. Thus, the entire sequence of realizations $\tilde{x} = (x_1, \ldots, x_N)$

can itself be described by a random variable, namely a sequence of N identical copies of the original random variable $\hat{x}$,

$$\hat{\tilde{x}} := (\hat{x}_1, \ldots, \hat{x}_N), \tag{8}$$

where $\hat{x}_1 = \ldots = \hat{x}_N = \hat{x}$. Each time I draw the random variable $\hat{\tilde{x}}$, I obtain another sequence of N realizations of $\hat{x}$.

An estimator is *unbiased* if for any sample size N we have

$$\langle\hat{\theta}\rangle = \theta, \tag{9}$$

otherwise it is *biased*. The bias is defined by

$$B(\hat{\theta}) := \langle\hat{\theta}\rangle - \theta. \tag{10}$$

An estimator is *consistent* if for large sample size N the probability to fail the true value θ goes to zero,

$$\mathcal{P}\{|\hat{\theta} - \theta| > \epsilon\} \to 0 \tag{11}$$

for $N \to \infty$ and arbitrarily small $\epsilon > 0$.

For a given sample $\tilde{x}$ the estimation of some parameter θ will fail by a certain amount. The expected amount of failure, the *mean square error* is related to the overall goodness of the estimator,

$$\mathrm{MSE}(\hat{\theta}) := \langle(\hat{\theta} - \theta)^2\rangle, \tag{12}$$

which is equal to

$$\mathrm{MSE}(\hat{\theta}) = \sigma^2(\hat{\theta}) + B^2(\hat{\theta}). \tag{13}$$

4 Estimating mean and variance

Say, there is a continuous random variable $\hat{x}$ with mean μ and variance σ^2, and we have measured the random variable $\hat{x}$ exactly N times. The set of results $\tilde{x} = \{x_1, \ldots, x_N\}$ then constitutes a *sample* of size N. We now estimate the mean μ of the random variable $\hat{x}$ by the *sample mean* $\mu(\tilde{x})$ defined by

$$\mu(\tilde{x}) := \frac{1}{N}\sum_{n=1}^{N} x_n, \tag{14}$$

as already introduced in (1), so that $\mu(\tilde{x}) \to \mu$ for $N \to \infty$. We know that the larger N, the better the estimate of μ. But how *confident* can we be about an estimate of μ based on a particular sample of size N?

The question is solved by considering the sample mean itself as a random variable, namely the *estimator* of the mean μ,

$$\hat{\mu} := \frac{1}{N}\sum_{i=1}^{N}\hat{x}_i, \tag{15}$$

where all random variables $\hat{x}_i$ are identical copies of the variable of interest, $\hat{x}_i = \hat{x}$. The expectation value of $\hat{\mu}$ is then identical to the mean,

$$\langle\hat{\mu}\rangle = \langle\frac{1}{N}\sum_{i=1}^{N}\hat{x}_i\rangle \tag{16}$$

$$= \frac{1}{N}\sum_{i=1}^{N}\langle\hat{x}_i\rangle \tag{17}$$

$$= \frac{1}{N}N\langle\hat{x}\rangle = \mu, \tag{18}$$

and, as variances simply add up, the variance σ_μ^2 of $\hat{\mu}$ equals

$$\sigma_\mu^2 = \sigma^2(\hat{\mu}) = \sigma^2(\frac{1}{N}\sum_{i=1}^{N}\hat{x}_i) \tag{19}$$

$$= \frac{1}{N^2}\sum_{i=1}^{N}\sigma^2(\hat{x}_i) \tag{20}$$

$$= \frac{1}{N^2}N\sigma^2(\hat{x}) = \frac{\sigma^2}{N}. \tag{21}$$

hence

$$\sigma_\mu = \frac{\sigma}{\sqrt{N}} \tag{22}$$

which goes to zero as N goes to infinity, thus $\hat{\mu}$ is a *consistent* estimator.

Moreover, we estimate the variance σ^2 of the random variable $\hat{x}$ by the statistic

$$\sigma^2(\tilde{x}) := \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \mu(\tilde{x}))^2. \tag{23}$$

The corresponding estimator is given by

$$\hat{\sigma}^2 := \frac{1}{N-1}\sum_{i=1}^{N}(\hat{x}_i - \hat{\mu})^2. \tag{24}$$

One can show that the factor $\frac{1}{N-1}$ causes the estimator to be unbiased, i.e. $\langle\hat{\sigma}^2\rangle = \sigma^2$. If one would take a factor of $\frac{1}{N}$ instead, then there would be a small bias (which vanishes for large N). The estimator (24) is consistent, which means that the realizations of $\hat{\sigma}^2$ approach the true variance σ^2 for large N in the probabilistic sense (11).

5 Student's t distribution

Now assume that the random variable $\hat{x}$ with mean μ and standard deviation σ is Gaussian. Then from (22) we infer that the random variable

$$\hat{z} = \frac{\hat{\mu} - \mu}{\sigma/\sqrt{N}} \tag{25}$$

is *normally distributed*, that is, it has a Gaussian distribution being centered at $\mu_z = 0$ and having a standard deviation of $\sigma_z = 1$. Hence, z indicates the *distance from the mean in standard deviations*, often referred to as the *z-score*. The probability distribution of $\hat{z}$ tells us how good we can estimate the unknown mean μ from the data given that we *apriorily know* the true standard deviation σ. However, we don't know the true standard deviation σ, so we have to estimate it by the estimator $\hat{\sigma}$ given by (24).

Hence, the goodness of our estimation of μ depends on the probability distribution of the random variable

$$\hat{t} = \frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}}. \tag{26}$$

William Gosset, better known under his pseudonym "Student", showed that $\hat{t}$ has a particular distribution which become famous as *Student's t distribution*. For large N this distribution approximates a Gaussian distribution, centered at $\mu_t = 0$ and having standard deviation of $\sigma_t = 1$. Explicitely, the t distribution reads

$$\rho_\nu(t) = \frac{\Gamma((\nu+1)/2)}{\sqrt{\nu\pi}\Gamma(\nu/2)}(1 + t^2/\nu)^{-(\nu+1)/2}, \tag{27}$$

where ν is the *degree of freedom* of the distribution. In our case we have a single random variable which gives the simple value

$$\nu = N - 1. \tag{28}$$

The only important thing to memorize here is that ρ_ν is independent of μ and σ. This means: *The confidence that the estimation of μ is correct does not depend on μ and σ, but it only depends on the sample size N.*

5.1 Confidence interval

Consider the event that the random variable $\hat{t}$ has an outcome greater than a particular positive number $\tau > 0$, then this event can be rewritten as

$$\hat{t} > \tau \tag{29}$$

$$\frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}} > \tau \tag{30}$$

$$\hat{\mu} > \mu + \tau \cdot \frac{\hat{\sigma}}{\sqrt{N}}, \tag{31}$$

which equals the event that we *overestimate* the true mean μ by an amount of

$$u = \tau \cdot \frac{\hat{\sigma}}{\sqrt{N}}. \tag{32}$$

From the above relation we see that the probability for such overestimation reads

$$\begin{aligned} \alpha &:= \mathcal{P}\{\hat{\mu} > \mu + u\} && (33)\\ &= \mathcal{P}\{\hat{t} > \tau\} && (34)\\ &= \int_{\tau}^{\infty} dt\, \rho_\nu(t) && (35)\\ &= 1 - \int_{-\infty}^{\tau} dt\, \rho_\nu(t) && (36)\\ &= 1 - F_\nu(\tau), && (37) \end{aligned}$$

where

$$F_\nu(\tau) = \int_{-\infty}^{\tau} dt\, \rho_\nu(t) \tag{38}$$

is the *cumulative probability distribution* of the Student distribution. Since further above we have assumed $\tau > 0$ and since $\rho_\nu(t)$ is symmetric about $t = 0$, we have

$$\frac{1}{2} < F_\nu(\tau) < 1, \tag{39}$$

and so

$$0 < \alpha < \frac{1}{2}. \tag{40}$$

Similarly, the event that $\hat{t}$ has an outcome smaller than $-\tau$ can be rewritten as

$$\begin{aligned} \hat{t} &< -\tau && (41)\\ \hat{\mu} &< \mu - \tau \cdot \frac{\hat{\sigma}}{\sqrt{N}}, && (42) \end{aligned}$$

which equals the event that we *underestimate* the true mean μ by an amount of u. As can be easily verified, the probability of such underestimation is identical to the probability for the overestimation of μ, due to the symmetry of the Student distribution,

$$\mathcal{P}\{\hat{\mu} < \mu - u\} = \mathcal{P}\{\hat{\mu} > \mu + u\} = \alpha. \tag{43}$$

The complementary probability, that is the probability to either *not overestimate* or *not underestimate* the true mean by more than u is also called the *level of significance*,

$$s := 1 - \alpha, \tag{44}$$

thus we have

$$s = F_\nu(\tau), \tag{45}$$

which admits values $\frac{1}{2} < s < 1$. Evidently, the significance level s is dependent on the value τ and the degree of freedom $\nu = N-1$ of the underlying probability distribution. Therefore, if we fix a desired significance level s and a sample size N we get a value of

$$\tau_{\nu,s} = F_\nu^{-1}(s), \tag{46}$$

where F_ν^{-1} is the *inverse cumulative Student distribution* and $\nu = N - 1$. The above value $\tau_{\nu,s}$ is called the *Student factor* and since it is not quite easy to calculate, the student factor is tabulated in textbooks for certain common values of ν and s.

Now, the probability that the true mean μ lies in the *confidence interval* $\hat{C} = [\hat{\mu} - u, \hat{\mu} + u]$ about the estimated mean $\hat{\mu}$, reads

$$\begin{aligned}
\mathcal{P}\{\mu \in \hat{C}\} &= 1 - \mathcal{P}\{\mu \notin \hat{C}\} && (47)\\
&= 1 - (\mathcal{P}\{\mu < \hat{\mu} - u\} + \mathcal{P}\{\mu > \hat{\mu} + u\}) && (48)\\
&= 1 - (\mathcal{P}\{\hat{\mu} > \mu + u\} + \mathcal{P}\{\hat{\mu} < \mu - u\}) && (49)\\
&= 1 - 2\alpha && (50)\\
&= 2s - 1. && (51)
\end{aligned}$$

So this is why we speak of a *confidence interval*. If we have estimated a mean $\hat{\mu}$ and a standard deviation $\hat{\sigma}$ from a data sample of size N, and we fix a desired significance level s, then for $\nu = N - 1$ we obtain an interval radius of

$$u_{\nu,s} = \tau_{\nu,s} \cdot \frac{s}{\sqrt{N}} \tag{52}$$

so that with probability $P_{\text{conf}} = 2s - 1$ we can be sure that the true mean μ lies within the range $\hat{\mu} \pm u_{\nu,s}$. This corresponds to a *two-tailed t-test* between the estimator $\hat{\mu}$ and the true mean μ, as we will see in the next section. Sometimes, the probability $P_{\text{conf}} = 2s - 1$ is itself to be indicated as the significance level. Then one should either calculate the value $s = (1+P_{\text{conf}})/2$ and look up the Student factor in a standard table for a one-tailed t-test or take P_{conf} as the significance level and look up the Student factor in a table which is explicitely dedicated to a two-tailed t-test. Also, it often occurs that instead of the significance level s the value $\alpha = 1 - s$ is denoted as the level of significance, also often denoted by the letter p in the form $p < 0.001$ or something like that. Lastly, the notions of *two-tailed* or *one-tailed* are sometimes confused, ignored or misused in the literature.

5.2 Example

Given a sample of size $N = 11$ with a sample mean of $\hat{\mu} = 10$ and a sample variance $\hat{\sigma} = 2$. According to (28) we have $\nu = 10$ degrees of freedom. We fix a 95% confidence level, thus $s = 0.95$, and look into Table 1 to get a t-value of $\tau = 1.812$. Using (52) we thus obtain a confidence radius of

$$u = 1.812 \cdot \frac{\sqrt{2}}{\sqrt{11}} \approx 0.426. \tag{53}$$

It follows that the probability that the true mean lies below $\mu + u = 10.426$ reads 95%. Similarly, the probability that the true mean lies above $\mu - u = 9.574$ also reads 95%. Hence, the probability that the true mean lies in the confidence interval $[9.574, 10.426]$ reads $P_{\mathrm{conf}} = 2s - 1 = 90\%$.

6 Estimating two random variables

6.1 Unequal variances

Consider two Gaussian random variables $\hat{x}_1, \hat{x}_2$ with means μ_1, μ_2 and standard deviations σ_1, σ_2, respectively. These variables are each realized N times, which gives the two samples $\tilde{x}_1 \tilde{x}_2$. We want to find out if the underlying random variables $\hat{x}_1$ and $\hat{x}_2$ have the same mean or not. The estimator of the mean of each random variable is given by

$$\hat{\mu}_k = \frac{1}{N_k} \sum_{i=1}^{N_k} \hat{x}_{ki}, \tag{54}$$

where $\hat{x}_{ki} = \hat{x}_k$ for all $i = 1, \ldots N_k$ and $k = 1, 2$. The variance of $\hat{\mu}_k$ is given by

$$\sigma_{\mu_k} = \frac{\sigma_k}{\sqrt{N_k}}. \tag{55}$$

Because variances simply add up, the difference between the two estimators has the variance

$$\sigma^2_{\mu_1 - \mu_2} = \frac{\sigma_1^2}{N_1} + \frac{\sigma_2^2}{N_2}. \tag{56}$$

Therefore, the random variable

$$\hat{z} = \frac{\hat{\mu}_1 - \hat{\mu}_2}{\sqrt{\frac{\sigma_1^2}{N_1} + \frac{\sigma_2^2}{N_2}}} \tag{57}$$

is normally distributed with mean $\mu_z = 0$ and variance $\sigma_z = 1$. Again, we don't know the true variances σ_k but have to estimate them from the data samples. The estimator of the variance of each random variable is given by

$$\hat{\sigma}_k^2 = \frac{1}{N_k - 1} \sum_{i=1}^{N_k} (\hat{x}_{ki} - \hat{\mu}_i)^2, \tag{58}$$

so the random variable

$$\hat{t} = \frac{\hat{\mu}_1 - \hat{\mu}_2}{\sqrt{\frac{\hat{\sigma}_1^2}{N_1} + \frac{\hat{\sigma}_2^2}{N_2}}} \tag{59}$$

is Student distributed with the probability density $\rho_\nu(t)$ given in (27). Unfortunately, the number ν of degrees of freedom is a little bit complicated to calculate. Closer investigations show that the number is approximated by

$$\nu = \frac{\left(\frac{s_1^2}{N_1} + \frac{s_2^2}{N_2}\right)^2}{\left(\frac{s_1^2}{N_1\sqrt{N_1-1}}\right)^2 + \left(\frac{s_2^2}{N_2\sqrt{N_2-1}}\right)^2}, \tag{60}$$

where s_1^2, s_2^2 are the standard deviations of the samples $\tilde{x}_1, \tilde{x}_2$, respectively. The righthand side of the formula above usually returns a noninteger value, where it is conventional to round down to the nearest integer. A quick-and-dirty replacement for the complicated formula is to simply take

$$\nu = \min\{N_1, N_2\}, \tag{61}$$

which underestimates ν and thus gives a stronger bound for significance than necessary: The t-value and thus the distance between the sample means needed for significance becomes greater, so the null hypothesis is strengthened.

Asking whether $\hat{x}_1$ and $\hat{x}_2$ have the same mean or not is equivalent to asking whether the random variable $\hat{t}$ has zero mean or not. Since $\hat{t}$ is Student distributed, we can apply the same strategy as in the last section. We fix a desired significance level s, and for the sample sizes N_1, N_2 given, we look up the Student factor $\tau_{\nu,s}$ in some textbook table to see whether the difference $\hat{\mu}_1 - \hat{\mu}_2$ lies within the confidence interval $C = [-u, u]$ where

$$u = \tau_{\nu,s} \cdot \sqrt{\frac{s_1^2}{N_1} + \frac{s_2^2}{N_2}}. \tag{62}$$

If this is the case, then the null hypothesis cannot be rejected, which means that the two random variables $\hat{x}_1, \hat{x}_2$ are not significantly different at level s. The same care with one-tailed and two-tailed t-tests as in the previous section must also be taken here. By default, the tables list Student factors for the *one-tailed* t-test. This means that the significance level s refers to the null hypothesis of either $\mu_1 < \mu_2$ or $\mu_1 > \mu_2$ being rejected (which both occurs with the same probability). To reject the null hypothesis $\mu_1 = \mu_2$, one has to put up a Student factor from a table that is explicitely dedicated to a *two-tailed* t-test. Alternatively, one calculates $s' = 2s - 1$ and looks for the Student factor corresponding to significance level s' in a standard table dedicated to a one-tailed t-test. It is convenient to indicate the significance not by s but rather by the complementary value $\alpha = 1 - s$ which represents the probability for falsely rejecting the null hypothesis, a so-called "false positive". In this unhappy case, the random variables $\hat{x}_1$ and $\hat{x}_2$ are considered to have different mean although they have not.

6.2 Equal variances

If the variances of $\hat{x}_1, \hat{x}_2$ are assumed to be equal, then one can optimize the test. Instead of considering the separate estimations of each variance one considers the *pooled variance*

$$\hat{\sigma}_p^2 = \frac{(N_1 - 1)\hat{\sigma}_1^2 + (N_2 - 1)\hat{\sigma}_2^2}{N_1 + N_2 - 2}, \tag{63}$$

which is sort of an average between the two individual variances $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$. Now we insert this pooled variance for $\hat{\sigma}_1$ and $\hat{\sigma}_2$ into (59) and obtain

$$\hat{t} = \frac{\hat{\mu}_1 - \hat{\mu}_2}{\sqrt{\frac{(N_1-1)\hat{\sigma}_1^2+(N_2-1)\hat{\sigma}_2^2}{N_1+N_2-2}\left(\frac{1}{N_1}+\frac{1}{N_2}\right)}} \tag{64}$$

Although the above formula is more complicated than the one for unequal variances, we now have an easier time calculating the degrees of freedom,

$$\nu = N_1 + N_2 - 2. \tag{65}$$

6.3 The t-value

We may also use the t-value directly to see the actual level of significance that the two random variables $\hat{x}_1, \hat{x}_2$ differ in their mean. First, we calculate the t-value according to (59) or (64), and then we calculate the degrees of freedom according to (60) or (65), respectively. Then we look up a table of t-values in the row where the number of degrees of freedom equals our own value of ν. The t-value in the table row which is just below our own t-value then reveals the level of significance s.

ν	Significance s							
	0.75	0.875	0.90	0.95	0.975	0.99	0.995	0.999
1	1.000	2.414	3.078	6.314	12.706	31.821	63.657	318.309
2	0.817	1.604	1.886	2.920	4.303	6.965	9.925	22.327
3	0.765	1.423	1.638	2.353	3.182	4.541	5.841	10.215
4	0.741	1.344	1.533	2.132	2.776	3.747	4.604	7.173
5	0.727	1.301	1.476	2.015	2.571	3.365	4.032	5.893
6	0.718	1.273	1.440	1.943	2.447	3.143	3.707	5.208
7	0.711	1.254	1.415	1.895	2.365	2.998	3.499	4.785
8	0.706	1.240	1.397	1.860	2.306	2.896	3.355	4.501
9	0.703	1.230	1.383	1.833	2.262	2.821	3.250	4.297
10	0.700	1.221	1.372	1.812	2.228	2.764	3.169	4.144
11	0.697	1.214	1.363	1.796	2.201	2.718	3.106	4.025
12	0.695	1.209	1.356	1.782	2.179	2.681	3.055	3.930
13	0.694	1.204	1.350	1.771	2.160	2.650	3.012	3.852
14	0.692	1.200	1.345	1.761	2.145	2.624	2.977	3.787
15	0.691	1.197	1.341	1.753	2.131	2.602	2.947	3.733
16	0.690	1.194	1.337	1.746	2.120	2.583	2.921	3.686
17	0.689	1.191	1.333	1.740	2.110	2.567	2.898	3.646
18	0.688	1.189	1.330	1.734	2.101	2.552	2.878	3.611
19	0.688	1.187	1.328	1.729	2.093	2.539	2.861	3.579
20	0.687	1.185	1.325	1.725	2.086	2.528	2.845	3.552
21	0.686	1.183	1.323	1.721	2.080	2.518	2.831	3.527
22	0.686	1.182	1.321	1.717	2.074	2.508	2.819	3.505
23	0.685	1.180	1.319	1.714	2.069	2.500	2.807	3.485
24	0.685	1.179	1.318	1.711	2.064	2.492	2.797	3.467
25	0.684	1.178	1.316	1.708	2.060	2.485	2.787	3.450
26	0.684	1.177	1.315	1.706	2.056	2.479	2.779	3.435
27	0.684	1.176	1.314	1.703	2.052	2.473	2.771	3.421
28	0.683	1.175	1.313	1.701	2.048	2.467	2.763	3.408
29	0.683	1.174	1.311	1.699	2.045	2.462	2.756	3.396
30	0.683	1.173	1.310	1.697	2.042	2.457	2.750	3.385
40	0.681	1.167	1.303	1.684	2.021	2.423	2.704	3.307
50	0.679	1.164	1.299	1.676	2.009	2.403	2.678	3.261
60	0.679	1.162	1.296	1.671	2.000	2.390	2.660	3.232
70	0.678	1.160	1.294	1.667	1.994	2.381	2.648	3.211
80	0.678	1.159	1.292	1.664	1.990	2.374	2.639	3.195
90	0.677	1.158	1.291	1.662	1.987	2.368	2.632	3.183
100	0.677	1.157	1.290	1.660	1.984	2.364	2.626	3.174
200	0.676	1.154	1.286	1.653	1.972	2.345	2.601	3.131
300	–	–	1.284	1.650	1.968	2.339	2.592	3.118
400	–	–	1.284	1.649	1.966	2.336	2.588	3.111
500	0.675	1.152	1.283	1.648	1.965	2.334	2.586	3.107
∞	0.674	1.150	1.282	1.645	1.960	2.326	2.576	3.090

Table 1: Table of Student's t-values for various degrees of freedom ν and significance levels s, for the case of a one-tailed t-test.