
Estimating Distributions and Densities

We have spent a lot of time looking at how to estimate expectations (which is regression). We have also seen how to estimate variances, by turning it into a problem about expectations. We could extend the same methods to looking at higher moments — if you need to find the conditional skewness or kurtosis functions¹, you can tackle that in the same way as finding the conditional variance. But what if we want to look at the whole distribution?

You’ve already seen one solution to this problem in earlier statistics courses: posit a parametric model for the density (Gaussian, Student’s t , exponential, gamma, beta, Pareto, ...) and estimate the parameters. Maximum likelihood estimates are generally consistent and efficient for such problems. None of this changes when the distributions are multivariate. But suppose you don’t have any particular parametric density family in mind, or want to check one — how could we estimate a probability distribution non-parametrically?

14.1 Histograms Revisited

For most of you, making a histogram was probably one of the first things you learned how to do in intro stats (if not before). This is a simple way of estimating a distribution: we split the sample space up into bins, count how many samples fall into each bin, and then divide the counts by the total number of samples. If we hold the bins fixed and take more and more data, then by the law of large numbers we anticipate that the relative frequency for each bin will converge on the bin’s probability.

So far so good. But one of the things you learned in intro stats was also to work with probability *density* functions, not just probability *mass* functions. Where do we get pdfs? Well, one thing we could do is to take our histogram estimate, and then say that the probability density is uniform within each bin. This gives us a piecewise-constant estimate of the density.

Unfortunately, this isn’t going to work — isn’t going to converge on the true pdf — unless we can shrink the bins of the histogram as we get more and more data. To see this, think about estimating the pdf when the data comes from any of the standard distributions, like an exponential or a Gaussian. We can approximate the true pdf $f(x)$ to arbitrary accuracy by a piecewise-constant density (indeed,

¹ When you find out what the kurtosis is good for, be sure to tell the world.

that’s what happens every time we plot it on our screens), but, for a fixed set of bins, we can only come so close to the true, continuous density.

This reminds us of our old friend the bias-variance trade-off, and rightly so. If we use a large number of very small bins, the minimum bias in our estimate of any density becomes small, but the variance in our estimates grows. (Why does variance increase?) To make some use of this insight, though, there are some things we need to establish first.

- Is learning the whole distribution non-parametrically even feasible?
- How can we measure error so deal with the bias-variance trade-off?

14.2 “The Fundamental Theorem of Statistics”

Let’s deal with the first point first. In principle, something even dumber than shrinking histograms will work to learn the whole distribution. Suppose we have one-dimensional samples x_1, x_2, \dots, x_n with a common cumulative distribution function F . The **empirical cumulative distribution function** on n samples, $\tilde{F}_n(a)$ is

$$\tilde{F}_n(a) \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{(-\infty, a]}(x_i) \quad (14.1)$$

In words, this is just the fraction of the samples which are $\leq a$. Then the **Glivenko-Cantelli theorem** says

$$\max_a |\tilde{F}_n(a) - F(a)| \rightarrow 0 \quad (14.2)$$

So the empirical CDF converges to the true CDF everywhere; the maximum gap between the two of them goes to zero. [Pitman \(1979\)](#) calls this the “fundamental theorem of statistics”, because it says we can learn distributions just by collecting enough data.² The same kind of result also holds for the CDFs of higher-dimensional vectors.

If the Glivenko-Cantelli theorem is so great, why aren’t we just content with

² There are some interesting aspects to the theorem which are tangential to what we’ll need, so I will stick them in this footnote. These hinge on the max in the statement of the theorem. For any one, fixed value of a , that $|\tilde{F}_n(a) - F(a)| \rightarrow 0$ is just an application of the law of large numbers. The extra work Glivenko and Cantelli did was to show that this held for *infinitely many* values of a at once, so that even if we focus on the biggest gap between the estimate and the truth, that still shrinks with n . Here’s a sketch, with no details. Fix an $\epsilon > 0$; first show that there is some *finite* set of points on the line, call them $b_1, \dots, b_{m(\epsilon)}$, such that, for any a , $|\tilde{F}_n(a) - \tilde{F}_n(b_i)| < \epsilon$ and $|F(a) - F(b_i)| < \epsilon$ for *some* b_i . Next, show that, for large enough n , $|F(b_i) - \tilde{F}_n(b_i)| < \epsilon$ for all the b_i simultaneously. (This follows from the law of large numbers and the fact that $m(\epsilon)$ is finite.) Finally, use the triangle inequality to conclude that, for large enough n , $\max_a |\tilde{F}_n(a) - F(a)| < 3\epsilon$. Since ϵ can be made arbitrarily small, the Glivenko-Cantelli theorem follows. This general strategy — combining pointwise convergence theorems with approximation arguments — forms the core of what’s called **empirical process theory**, which underlies the consistency of basically all the non-parametric procedures we’ve seen. If this line of thought is at all intriguing, the closest thing to a gentle introduction is [Pollard \(1989\)](#). (If you know enough to object that I should have been writing sup instead of max, you know enough to make the substitution for yourself.)

the empirical CDF? Sometimes we are, but it inconveniently doesn't give us a probability *density*. Suppose that x_1, x_2, \dots, x_n are sorted into increasing order. What probability does the empirical CDF put on the interval (x_i, x_{i+1}) ? Clearly, zero. (Whereas the interval $[x_i, x_{i+1}]$ gets probability $2/n$.) This *could* be right, but we have centuries of experience now with probability distributions, and this tells us that *pretty often* we can expect to find some new samples between our old ones. So we'd like to get a *non-zero* density between our observations.

Using a uniform distribution within each bin of a histogram doesn't have this issue, but it does leave us with the problem of picking where the bins go and how many of them we should use. Of course, there's nothing magic about keeping the bin size the same and letting the number of points in the bins vary; we could equally well pick bins so they had equal counts³. So what should we do?

14.3 Error for Density Estimates

Our first step is to get clear on what we mean by a "good" density estimate. There are three leading ideas:

1. $\int (f(x) - \hat{f}(x))^2 dx$ should be small: the squared deviation from the true density should be small, averaging evenly over all space.
2. $\int |f(x) - \hat{f}(x)| dx$ should be small: minimize the average *absolute*, rather than squared, deviation.
3. $\int f(x) \log \frac{f(x)}{\hat{f}(x)} dx$ should be small: the average log-likelihood ratio should be kept low.

Option (1) is reminiscent of the MSE criterion we've used in regression. Option (2) looks at what's called the L_1 or **total variation** distance between the true and the estimated density. It has the nice property that $\frac{1}{2} \int |f(x) - \hat{f}(x)| dx$ is exactly the maximum error in our estimate of the probability of *any* set. Unfortunately it's a bit tricky to work with, so we'll skip it here. (But see Devroye and Lugosi (2001)). Finally, minimizing the log-likelihood ratio is intimately connected to maximizing the likelihood. We will come back to this (§14.6), but, like most texts on density estimation, we will give more attention to minimizing (1), because it's mathematically tractable.

Notice that

$$\int (f(x) - \hat{f}(x))^2 dx = \int f^2(x) dx - 2 \int \hat{f}(x) f(x) dx + \int \hat{f}^2(x) dx \quad (14.3)$$

³ A specific idea for how to do this is sometimes called a $k-d$ tree. We have d random variables and want a joint density for all of them. Fix an ordering of the variables. Start with the first variable, and find the thresholds which divide it into k parts with equal counts. (Usually but not always $k=2$.) Then sub-divide each part into k equal-count parts on the *second* variable, then sub-divide each of those on the *third* variable, etc. After splitting on the d^{th} variable, go back to splitting on the first, until no further splits are possible. With n data points, it takes about $\log_k n$ splits before coming down to individual data points. Each of these will occupy a cell of some volume. Estimate the density on that cell as one over that volume. Of course it's not strictly necessary to keep refining all the way down to single points.

The first term on the right hand side doesn't depend on the estimate $\hat{f}(x)$ at all, so we can ignore it for purposes of optimization. The third one only involves \hat{f} , and is just an integral, which we can do numerically. That leaves the middle term, which involves both the true and the estimated density; we can approximate it by

$$-\frac{2}{n} \sum_{i=1}^n \hat{f}(x_i) \quad (14.4)$$

The reason we can do this is that, by the Glivenko-Cantelli theorem, integrals over the true density are approximately equal to sums over the empirical distribution.

So our final error measure is

$$-\frac{2}{n} \sum_{i=1}^n \hat{f}(x_i) + \int \hat{f}^2(x) dx \quad (14.5)$$

In fact, this error measure does *not* depend on having one-dimension data; we can use it in any number of dimensions⁴ For purposes of cross-validation (you knew that was coming, right?), we can estimate \hat{f} on the training set, and then restrict the sum to points in the testing set.

14.3.1 Error Analysis for Histogram Density Estimates

We now have the tools to do most of the analysis of histogram density estimation. (We'll do it in one dimension for simplicity.) Choose our favorite location x , which lies in a bin whose boundaries are x_0 and $x_0 + h$. We want to estimate the density at x , and this is

$$\hat{f}_n(x) = \frac{1}{h} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{(x_0, x_0+h]}(x_i) \quad (14.6)$$

Let's call the sum, the number of points in the bin, b . It's a random quantity, $B \sim \text{Binomial}(n, p)$, where p is the true probability of falling into the bin, $p = F(x_0 + h) - F(x_0)$. The mean of B is np , and the variance is $np(1 - p)$, so

$$\mathbb{E}[\hat{f}_n(x)] = \frac{1}{nh} \mathbb{E}[B] \quad (14.7)$$

$$= \frac{n[F(x_0 + h) - F(x_0)]}{nh} \quad (14.8)$$

$$= \frac{F(x_0 + h) - F(x_0)}{h} \quad (14.9)$$

⁴ Admittedly, in high-dimensional spaces, doing the final integral can become numerically challenging.

and the variance is

$$\mathbb{V}[\hat{f}_n(x)] = \frac{1}{n^2 h^2} \mathbb{V}[B] \quad (14.10)$$

$$= \frac{n[F(x_0 + h) - F(x_0)][1 - F(x_0 + h) + F(x_0)]}{n^2 h^2} \quad (14.11)$$

$$= \mathbb{E}[\hat{f}_n(x)] \frac{1 - F(x_0 + h) + F(x_0)}{nh} \quad (14.12)$$

If we let $h \rightarrow 0$ as $n \rightarrow \infty$, then

$$\mathbb{E}[\hat{f}_h(x)] \rightarrow \lim_{h \rightarrow 0} \frac{F(x_0 + h) - F(x_0)}{h} = f(x_0) \quad (14.13)$$

since the pdf is the derivative of the CDF. But since x is between x_0 and $x_0 + h$, $f(x_0) \rightarrow f(x)$. So if we use smaller and smaller bins as we get more data, the histogram density estimate is unbiased. We'd also like its variance to shrink as the same grows. Since $1 - F(x_0 + h) + F(x_0) \rightarrow 1$ as $h \rightarrow 0$, to get the variance to go away we need $nh \rightarrow \infty$.

To put this together, then, our first conclusion is that histogram density estimates will be consistent when $h \rightarrow 0$ but $nh \rightarrow \infty$ as $n \rightarrow \infty$. The bin-width h needs to shrink, but slower than n^{-1} .

At what rate should it shrink? Small h gives us low bias but (as you can verify from the algebra above) high variance, so we want to find the trade-off between the two. One can calculate the bias at x from our formula for $\mathbb{E}[\hat{f}_h(x)]$ through a somewhat lengthy calculus exercise, analogous to what we did for kernel smoothing in Chapter 4⁵; the upshot is that the integrated squared bias is

$$\int (f(x) - \mathbb{E}[\hat{f}_h(x)])^2 dx = \frac{h^2}{12} \int (f'(x))^2 dx + o(h^2) \quad (14.14)$$

We already got the variance at x , and when we integrate that over x we find

$$\int \mathbb{V}[\hat{f}_h(x)] dx = \frac{1}{nh} + o(n^{-1}) \quad (14.15)$$

So the total integrated squared error is

$$\text{ISE} = \frac{h^2}{12} \int (f'(x))^2 dx + \frac{1}{nh} + o(h^2) + o(n^{-1}) \quad (14.16)$$

Differentiating this with respect to h and setting it equal to zero, we get

$$\frac{h_{\text{opt}}}{6} \int (f'(x))^2 dx = \frac{1}{nh_{\text{opt}}^2} \quad (14.17)$$

$$h_{\text{opt}} = \left(\frac{6}{\int (f'(x))^2 dx} \right)^{1/3} n^{-1/3} = O(n^{-1/3}) \quad (14.18)$$

⁵ You need to use the intermediate value theorem multiple times; see for instance Wasserman (2006 sec. 6.8).

So we need narrow bins if the density changes rapidly ($\int (f'(x))^2 dx$ is large), and wide bins if the density is relatively flat. No matter how rough the density, the bin width should shrink like $O(n^{-1/3})$. Plugging that rate back into the equation for the ISE, we see that it is $O(n^{-2/3})$.

It turns out that if we pick h by cross-validation, then we attain this optimal rate in the large-sample limit. By contrast, if we *knew* the correct parametric form and just had to estimate the parameters, we'd typically get an error decay of $O(n^{-1})$. This is substantially faster than histograms, so it would be nice if we could make up some of the gap, without having to rely on parametric assumptions.

14.4 Kernel Density Estimates

It turns out that one can improve the convergence rate, as well as getting smoother estimates, by using kernels. The **kernel density estimate** is

$$\hat{f}_h(x) = \frac{1}{n} \sum_{i=1}^n \frac{1}{h} K\left(\frac{x - x_i}{h}\right) \quad (14.19)$$

where K is a kernel function such as we encountered when looking at kernel regression. (The factor of $1/h$ inside the sum is so that \hat{f}_h will integrate to 1; we could have included it in both the numerator and denominator of the kernel regression formulae, but then it would've just canceled out.) As before, h is the **bandwidth** of the kernel. We've seen typical kernels in things like the Gaussian. One advantage of using them is that they give us a smooth density everywhere, unlike histograms, and in fact we can even use them to estimate the derivatives of the density, should that be necessary.⁶

14.4.1 Analysis of Kernel Density Estimates

How do we know that kernels will in fact work? Well, let's look at the mean and variance of the kernel density estimate at a particular point x , and use Taylor's

⁶ The advantage of histograms is that they're computationally and mathematically simpler.

theorem on the density.

$$\mathbb{E} [\widehat{f}_h(x)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\frac{1}{h} K \left(\frac{x - X_i}{h} \right) \right] \quad (14.20)$$

$$= \mathbb{E} \left[\frac{1}{h} K \left(\frac{x - X}{h} \right) \right] \quad (14.21)$$

$$= \int \frac{1}{h} K \left(\frac{x - t}{h} \right) f(t) dt \quad (14.22)$$

$$= \int K(u) f(x - hu) du \quad (14.23)$$

$$= \int K(u) \left[f(x) - hu f'(x) + \frac{h^2 u^2}{2} f''(x) + o(h^2) \right] du \quad (14.24)$$

$$= f(x) + \frac{h^2 f''(x)}{2} \int K(u) u^2 du + o(h^2) \quad (14.25)$$

$$(14.26)$$

because, by definition, $\int K(u) du = 1$ and $\int u K(u) du = 0$. If we call $\int K(u) u^2 du = \sigma_K^2$, then the bias of the kernel density estimate is

$$\mathbb{E} [\widehat{f}_h(x)] - f(x) = \frac{h^2 \sigma_K^2 f''(x)}{2} + o(h^2) \quad (14.27)$$

So the bias will go to zero if the bandwidth h shrinks to zero. What about the variance? Use Taylor's theorem again:

$$\mathbb{V} [\widehat{f}_h(x)] = \frac{1}{n} \mathbb{V} \left[\frac{1}{h} K \left(\frac{x - X}{h} \right) \right] \quad (14.28)$$

$$= \frac{1}{n} \left[\mathbb{E} \left[\frac{1}{h^2} K^2 \left(\frac{x - X}{h} \right) \right] - \left(\mathbb{E} \left[\frac{1}{h} K \left(\frac{x - X}{h} \right) \right] \right)^2 \right] \quad (14.29)$$

$$= \frac{1}{n} \left[\int \frac{1}{h^2} K^2 \left(\frac{x - t}{h} \right) dt - [f(x) + O(h^2)]^2 \right] \quad (14.30)$$

$$= \frac{1}{n} \left[\int \frac{1}{h} K^2(u) f(x - hu) du - f^2(x) + O(h^2) \right] \quad (14.31)$$

$$= \frac{1}{n} \left[\int \frac{1}{h} K^2(u) (f(x) - hu f'(x)) du - f^2(x) + O(h) \right] \quad (14.32)$$

$$= \frac{f(x)}{hn} \int K^2(u) du + O(1/n) \quad (14.33)$$

This will go to zero if $nh \rightarrow \infty$ as $n \rightarrow \infty$. So the conclusion is the same as for histograms: h has to go to zero, but slower than $1/n$.

Since the expected squared error at x is the bias squared plus the variance,

$$\frac{h^4 \sigma_K^4 (f''(x))^2}{4} + \frac{f(x)}{hn} \int K^2(u) du + \text{small} \quad (14.34)$$

the expected integrated squared error is

$$\text{ISE} \approx \frac{h^4 \sigma_K^4}{4} \int (f''(x))^2 dx + \frac{\int K^2(u) du}{nh} \quad (14.35)$$

Differentiating with respect to h for the optimal bandwidth h_{opt} , we find

$$h_{\text{opt}}^3 \sigma_K^4 \int (f''(x))^2 dx = \frac{\int K^2(u) du}{nh_{\text{opt}}^2} \quad (14.36)$$

$$h_{\text{opt}} = \left(\frac{\int K^2(u) du}{\sigma_K^4 \int (f''(x))^2 dx} \right)^{1/5} n^{-1/5} = O(n^{-1/5}) \quad (14.37)$$

That is, the best bandwidth goes to zero like one over the fifth root of the number of sample points. Plugging this into Eq. 14.35 the best $\text{ISE} = O(n^{-4/5})$. This is better than the $O(n^{-2/3})$ rate of histograms, but still includes a penalty for having to figure out what kind of distribution we're dealing with. Remarkably enough, using cross-validation to pick the bandwidth gives near-optimal results⁷

As an alternative to cross-validation, or at least a starting point, one can use Eq. 14.37 to show that the optimal bandwidth for using a Gaussian kernel to estimate a Gaussian distribution is $1.06\sigma n^{-1/5}$, with σ being the standard deviation of the Gaussian. This is sometimes called the **Gaussian reference rule** or the **rule-of-thumb** bandwidth. When you call `density` in R, this is basically what it does.

Yet another technique is the **plug-in method**. Eq. 14.37 calculates the optimal bandwidth from the second derivative of the true density. This doesn't help if we don't know the density, but it becomes useful if we have an initial density estimate which isn't too bad. In the plug-in method, we start with an initial bandwidth (say from the Gaussian reference rule) and use it to get a preliminary estimate of the density. Taking that crude estimate and "plugging it in" to Eq. 14.37 gives us a new bandwidth, and we re-do the kernel estimate with that new bandwidth. Iterating this a few times is optional but not uncommon.

14.4.2 Joint Density Estimates

The discussion and analysis so far has been focused on estimating the distribution of a one-dimensional variable. Just as kernel regression can be done with multiple input variables (§4.3), we can make kernel density estimates of joint distributions. We simply need a kernel for the vector:

$$\hat{f}(\vec{x}) = \frac{1}{n} \sum_{i=1}^n K(\vec{x} - \vec{x}_i) \quad (14.38)$$

⁷ Substituting Eq. 14.37 into Eq. 14.35 gives a squared error of $1.25n^{-4/5}\sigma_K^{4/5}(\int (f''(x))^2 dx)^{1/5}(\int K^2(u) du)^{4/5}$. The only two parts of this which depend on the kernel are σ_K and $\int K^2(u) du$. This is the source of the (correct) folklore that the choice of kernel is less important than the choice of bandwidth.

One could use any multivariate distribution as the kernel (provided it is centered and has finite covariance). Typically, however, just as in smoothing, one uses a product kernel, i.e., a product of one-dimensional kernels,

$$K(\vec{x} - \vec{x}_i) = K_1(x^1 - x_i^1)K_2(x^2 - x_i^2) \dots K_p(x^p - x_i^p), \quad (14.39)$$

Doing this requires a bandwidth for each coordinate, so the over-all form of the joint PDF estimate is

$$\hat{f}(\vec{x}) = \frac{1}{n \prod_{j=1}^p h_j} \sum_{i=1}^n \prod_{j=1}^d K_j \left(\frac{x^j - x_i^j}{h_j} \right) \quad (14.40)$$

Going through a similar analysis for p -dimensional data shows that the ISE goes to zero like $O(n^{-4/(4+p)})$, and again, if we use cross-validation to pick the bandwidths, asymptotically we attain this rate. Unfortunately, if p is large, this rate becomes very slow — for instance, if $p = 24$, the rate is $O(n^{-1/7})$. There is simply no universally good way to learn *arbitrary* high-dimensional distributions. This is the same “curse of dimensionality” we saw in regression (§8.3). The fundamental problem is that in high dimensions, there are just too different possible distributions which are too hard to tell apart.

Evading the curse of dimensionality for density estimation needs some special assumptions. Parametric models make the very strong assumption that we know exactly what the distribution function looks like, and we just need to fill in a few constants. It’s potentially less drastic to hope the distribution has some sort of special structure we can exploit, and most of the rest of Part II will be about searching for various sorts of useful structure.⁸ If none of these options sound appealing, or plausible, we’ve got little alternative but to accept a very slow convergence of density estimates.

14.4.3 Categorical and Ordered Variables

Estimating probability mass functions with discrete variables can be straightforward: there are only a finite number of values, and so one just counts how often they occur and takes the relative frequency. If one has a discrete variable X and a continuous variable Y and one wants a joint distribution, one could just get a separate density for Y for each value of x , and tabulate the probabilities for x .

In principle, this will work, but it can be practically awkward if the number of levels for the discrete variable is large compared to the number of samples. Moreover, for the joint distribution problem, it has us estimating completely separate distributions for Y for every x , without any sharing of information between them. It would seem more plausible to smooth those distributions towards each others. To do this, we need kernels for discrete variables.

Several sets of such kernels have been proposed. The most straightforward,

⁸ As explained long ago by Wiener (1956), the reason nonparametric estimation doesn’t make scientific theories redundant is that good theories usefully constrain the distributions we’re searching for, and tell us what structures to look for.

however, are the following. If X is a categorical, unordered variable with c possible values, then, for $0 \leq h < 1$,

$$K(x_1, x_2) = \begin{cases} 1 - h & x_1 = x_2 \\ h/(c - 1) & x_1 \neq x_2 \end{cases} \quad (14.41)$$

is a valid kernel. For an ordered x ,

$$K(x_1, x_2) = \binom{c}{|x_1 - x_2|} h^{|x_1 - x_2|} (1 - h)^{c - |x_1 - x_2|} \quad (14.42)$$

where $|x_1 - x_2|$ should be understood as just how many levels apart x_1 and x_2 are. As $h \rightarrow 0$, both of these become indicators, and return us to simple relative frequency counting. Both of these are implemented in `np`.

14.4.4 Practicalities


The standard R function `density` implements one-dimensional kernel density estimation, defaulting to Gaussian kernels with the rule-of-thumb bandwidth. There are some options for doing cleverer bandwidth selection, including a plug-in rule. (See the help file.)

For more sophisticated methods, and especially for more dimensions, you'll need to use other packages. The `np` package estimates joint densities using the `npudens` function. (The `u` is for “unconditional”.) This has the same sort of automatic bandwidth selection as `npreg`, using cross-validation. Other packages which do kernel density estimation include `KernSmooth` and `sm`.

14.4.5 Kernel Density Estimation in R: An Economic Example

The data set `oecdpanel`, in the `np` library, contains information about much the same sort of variables at the Penn World Tables data you worked with in the homework, over much the same countries and years, but with some of the variables pre-transformed, with identifying country information removed, and slightly different data sources. See `help(oecdpanel)` for details.

Here's an example of using `npudens` with variables from the `oecdpanel` data set, from problem set [11](#). We'll look at the joint density of `popgro` (the logarithm of the population growth rate) and `inv` (the logarithm of the investment rate). Figure [14.1](#) illustrates how to call the command, and a useful trick where we get `np`'s plotting function to do our calculations for us, but then pass the results to a different graphics routine. (See `help(npplot)`.) The distribution we get has two big modes, one at a comparatively low population growth rate (≈ -2.9 — remember this is logged so it's not actually a shrinking population) and high investment (≈ -1.5), and the other at a lower rate of investment (≈ -2) and higher population growth (≈ -2.6). There is a third, much smaller mode at high population growth (≈ -2.7) and very low investment (≈ -4).



```

data(oecdpanel)
popinv <- npudens(~popgro + inv, data = oecdpanel)
fhat <- plot(popinv, plot.behavior = "data")$d1
library(lattice)
contourplot(fhat$dens ~ fhat$eval$Var1 * fhat$eval$Var2, cuts = 20, xlab = "popgro",
            ylab = "inv", labels = list(cex = 0.5))

```

Figure 14.1 Gaussian kernel estimate of the joint distribution of logged population growth rate (`popgro`) and investment rate (`inv`). Notice that `npudens` takes a formula, but that there is no dependent variable on the left-hand side of the `~`. With objects produced by the `np` library, one can give the plotting function the argument `plot.behavior` — the default is `plot`, but if it's set to `data` (as here), it calculates all the information needed to plot and returns a separate set of objects, which can be plotted in other functions. (The value `plot-data` does both.) See `help(npplot)` for more.

14.5 Conditional Density Estimation

In addition to estimating marginal and joint densities, we will often want to get conditional densities. The most straightforward way to get the density of Y given X , $f_{Y|X}(y | x)$, is

$$\widehat{f}_{Y|X}(y | x) = \frac{\widehat{f}_{X,Y}(x, y)}{\widehat{f}_X(x)} \quad (14.43)$$

i.e., to estimate the joint and marginal densities and divide one by the other.

To be concrete, let's suppose that we are using a product kernel to estimate the joint density, and that the marginal density is consistent with it:

$$\widehat{f}_{X,Y}(x, y) = \frac{1}{nh_X h_Y} \sum_{i=1}^n K_X\left(\frac{x - x_i}{h_X}\right) K_Y\left(\frac{y - y_i}{h_Y}\right) \quad (14.44)$$

$$\widehat{f}_X(x) = \frac{1}{nh_X} \sum_{i=1}^n K_X\left(\frac{x - x_i}{h_X}\right) \quad (14.45)$$

Thus we need to pick two bandwidths, h_X and h_Y , one for each variable.

This might seem like a solved problem — we just use cross-validation to find h_X and h_Y so as to minimize the integrated squared error for $\widehat{f}_{X,Y}$, and then plug in to Equation 14.43. However, this is a bit hasty, because the optimal bandwidths for the *joint* density are not necessarily the optimal bandwidths for the *conditional* density. An extreme but easy to understand example is when Y is actually independent of X . Since the density of Y given X is just the density of Y , we'd be best off just ignoring X by taking $h_X = \infty$. (In practice, we'd just use a very big bandwidth.) But if we want to find the joint density, we would not want to smooth X away completely like this.

The appropriate integrated squared error measure for the conditional density is

$$\int dx f_X(x) \int dy \left(f_{Y|X}(y | x) - \widehat{f}_{Y|X}(y | x) \right)^2 \quad (14.46)$$

and this is what we want to minimize by picking h_X and h_Y . The cross-validation goes as usual.

One nice, and quite remarkable, property of cross-validation for conditional density estimation is that it can detect and exploit conditional independence. Say that $X = (U, V)$, and that Y is independent of U given V — symbolically, $Y \perp\!\!\!\perp U | V$. Then $f_{Y|U,V}(y | u, v) = f_{Y|V}(y | v)$, and we should just ignore U in our estimation of the conditional density. It turns out that when cross-validation is used to pick bandwidths for conditional density estimation, $\widehat{h}_U \rightarrow \infty$ when $Y \perp\!\!\!\perp U | V$, but not otherwise (Hall *et al.* 2004). In other words, cross-validation will automatically detect which variables are irrelevant, and smooth them away.

14.5.1 Practicalities and a Second Example


The `np` package implements kernel conditional density estimation through the function `npcdens`. The syntax is pretty much exactly like that of `npreg`, and indeed we can think of estimating the conditional density as a sort of regression, where the dependent variable is actually a distribution.

To give a concrete example, let's look at how the distribution of countries' population growth rates has changed over time, using the `oecdpanel` data (Figure 14.2). The selected bandwidth for `year` is 10, while that for `popgro` is 0.048. (Note that `year` is being treated as a continuous variable.)

You can see from the figure that the mode for population growth rates is towards the high end of observed values, but the mode is shrinking and becoming less pronounced over time. The distribution in fact begins as clearly bimodal, but the smaller mode at the lower growth rate turns into a continuous "shoulder". Over time, Figure 14.2 population growth rates tend to shrink, and the dispersion of growth rates narrows.

Let's expand on this point. One of the variables in `oecdpanel` is `oecd`, which is 1 for countries which are members of the Organization for Economic Cooperation and Development, and 0 otherwise. The OECD countries are basically the "developed" ones (stable capitalist democracies). We can include OECD membership as a conditioning variable for population growth (we need to use a categorical-variable kernel), and look at the combined effect of time and development (Figure 14.3).

What the figure shows is that OECD and non-OECD countries both have unimodal distributions of growth rates. The mode for the OECD countries has become sharper, but the value has decreased. The mode for non-OECD countries has also decreased, while the distribution has become more spread out, mostly by having more probability of lower growth rates. (These trends have continued since 1995.) In words, despite the widespread contrary impression, population growth has actually been slowing for decades in both rich and poor countries.




```

pop.cdens <- npcdens(popgro ~ year, data = oecdpanel)
plotting.grid <- expand.grid(year = seq(from = 1965, to = 1995, by = 1), popgro = seq(from = -3.5,
  to = -2.4, length.out = 300))
fhat <- predict(pop.cdens, newdata = plotting.grid)
wireframe(fhat ~ plotting.grid$year * plotting.grid$popgro, scales = list(arrows = FALSE),
  xlab = "year", ylab = "popgro", zlab = "pdf")

```

Figure 14.2 Conditional density of logarithmic population growth rates as a function of time.



```

pop.cdens.o <- npcdens(popgro ~ year + factor(oecd), data = oecdpanel)
oecd.grid <- expand.grid(year = seq(from = 1965, to = 1995, by = 1), popgro = seq(from = -3.4,
  to = -2.4, length.out = 300), oecd = unique(oecdpanel$oecd))
fhat <- predict(pop.cdens.o, newdata = oecd.grid)
wireframe(fhat ~ oecd.grid$year * oecd.grid$popgro | oecd.grid$oecd, scales = list(arrows = FALSE),
  xlab = "year", ylab = "popgro", zlab = "pdf")

```

Figure 14.3 Conditional density of population growth rates given year and OECD membership. The left panel is countries not in the OECD, the right is ones which are.

14.6 More on the Expected Log-Likelihood Ratio

I want to say just a bit more about the expected log-likelihood ratio $\int f(x) \log \frac{f(x)}{\hat{f}(x)} dx$. More formally, this is called the **Kullback-Leibler divergence** or **relative entropy** of \hat{f} from f , and is also written $D(f\|\hat{f})$. Let's expand the log ratio:

$$D(f\|\hat{f}) = - \int f(x) \log \hat{f}(x) dx + \int f(x) \log f(x) dx \quad (14.47)$$

The second term does not involve the density estimate, so it's irrelevant for purposes of optimizing over \hat{f} . (In fact, we're just subtracting off the entropy of the true density.) Just as with the squared error, we could try approximating the integral with a sum:

$$\int f(x) \log \hat{f}(x) dx \approx \frac{1}{n} \sum_{i=1}^n \log \hat{f}(x_i) \quad (14.48)$$

which is just the log-likelihood per observation. Since we know and like maximum likelihood methods, why not just use this?

Well, let's think about what's going to happen if we plug in the kernel density estimate:

$$\frac{1}{n} \sum_{i=1}^n \log \left(\frac{1}{nh} \sum_{j=1}^n K \left(\frac{x_j - x_i}{h} \right) \right) = - \log nh + \frac{1}{n} \sum_{i=1}^n \log \left(\sum_{j=1}^n K \left(\frac{x_j - x_i}{h} \right) \right) \quad (14.49)$$

If we take h to be very small, $K(\frac{x_j - x_i}{h}) \approx 0$ unless $x_j = x_i$, so the over-all likelihood becomes

$$\approx - \log nh + \log K(0) \quad (14.50)$$

which goes to $+\infty$ as $h \rightarrow 0$. So if we want to maximize the likelihood of a kernel density estimate, we *always* want to make the bandwidth as small as possible. In fact, the limit is to say that the density is

$$\tilde{f}(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - x_i) \quad (14.51)$$

where δ is the Dirac delta function⁹. Of course, this is just the same distribution as the empirical CDF.

Why is maximum likelihood failing us here? Well, it's doing exactly what we asked it to: to find the distribution where the observed sample is as probable as possible. Giving any probability to values of x we *didn't* see can only come at

⁹ Recall that the delta function is defined by how it integrates with other functions:

$\int \delta(x) f(x) dx = f(0)$. You can imagine $\delta(x)$ as zero everywhere except at the origin, where it has an infinitely tall, infinitely narrow spike, the area under the spike being one. If you are suspicious that this is really a bona fide function, you're right; strictly speaking it's just a linear operator on functions. We can however approximate it as the limit of well-behaved functions. For instance, take $\delta_h(x) = 1/h$ when $x \in [-h/2, h/2]$ with $\delta_h(x) = 0$ elsewhere, and let h go to zero. But this is where we came in. . .

the expense of the probability of observed values, so Eq. 14.51 really is the unrestricted maximum likelihood estimate of the distribution. Anything else imposes some restrictions or constraints which don't, strictly speaking, come from the data. However, those restrictions are what let us generalize to new data, rather than just memorizing the training sample.

One way out of this is to use the *cross-validated* log-likelihood to pick a bandwidth, i.e., to restrict the sum in Eq. 14.48 to running over the testing set only. This way, very small bandwidths don't get an unfair advantage for concentrating around the training set. (If the test points are in fact all very close to the training points, then small bandwidths get a *fair* advantage.) This is in fact the default procedure in the `np` package, through the `bwmethod` option ("`cv.ml`" vs. "`cv.ls`").

14.7 Simulating from Density Estimates

14.7.1 Simulating from Kernel Density Estimates

There are times when one wants to draw random values from the estimated distribution. This is easy with kernel density estimates, because each kernel is itself a probability density, generally a very tractable one. The pattern goes like so. Suppose the kernel is Gaussian, that we have scalar observations x_1, x_2, \dots, x_n , and the selected bandwidth is h . Then we pick an integer i uniformly at random from 1 to n , and invoke `rnorm(1, x[i], h)`¹⁰ Using a different kernel, we'd just need to use the random number generator function for the corresponding distribution.

To see that this gives the right distribution needs just a little math. A kernel $K(x, x_i, h)$ with bandwidth h and center x_i is a probability density function. The probability the KDE gives to any set A is just an integral:

$$\hat{F}(A) = \int_A \hat{f}(x) dx \quad (14.52)$$

$$= \int_A \frac{1}{n} \sum_{i=1}^n K(x, x_i, h) dx \quad (14.53)$$

$$= \frac{1}{n} \sum_{i=1}^n \int_A K(x, x_i, h) dx \quad (14.54)$$

$$= \frac{1}{n} \sum_{i=1}^n C(A, x_i, h) \quad (14.55)$$

introducing C to stand for the probability distribution corresponding to the kernel. The simulation procedure works if the probability that the simulated value \tilde{X} falls into A matches this. To generate \tilde{X} , we first pick a random data point,

¹⁰ In fact, if we want to draw a sample of size q , `rnorm(q, sample(x, q, replace=TRUE), h)` will work in R — it's important though that sampling be done *with* replacement.

which really means picking a random integer J , uniformly from 1 to n . Then

$$\Pr(\tilde{X} \in A) = \mathbb{E}[\mathbf{1}_A(\tilde{X})] \quad (14.56)$$

$$= \mathbb{E}\left[\mathbb{E}[\mathbf{1}_A(\tilde{X}) \mid J]\right] \quad (14.57)$$

$$= \mathbb{E}[C(A, x_J, h)] \quad (14.58)$$

$$= \frac{1}{n} \sum_{i=1}^n C(A, x_i, h) \quad (14.59)$$

The first step uses the fact that a probability is the expectation of an indicator function; the second uses the law of total expectation; the last steps use the definitions of C and J , and the distribution of J .

14.7.1.1 Sampling from a Joint Density

The procedure given above works with only trivial modification for sampling from a joint, multivariate distribution. If we're using a product kernel, we pick a random data point, and then draw each coordinate independently from the kernel distribution centered on our random point. (See Code Example [29](#) below.) The argument for correctness actually goes exactly as before.

14.7.1.2 Sampling from a Conditional Density

Sampling from a conditional density estimate with product kernels is again straightforward. The one trick is that one needs to do a *weighted* sample of data points. To see why, look at the conditional distribution (not density) function:

$$\hat{F}(Y \in A \mid X = x) \quad (14.60)$$

$$= \int_A \hat{f}_{Y|X}(y \mid x) dy$$

$$= \int_A \frac{\frac{1}{nh_X h_Y} \sum_{i=1}^n K_X\left(\frac{x-x_i}{h_X}\right) K_Y\left(\frac{y-y_i}{h_Y}\right)}{\hat{f}_X(x)} dy \quad (14.61)$$

$$= \frac{1}{nh_X h_Y \hat{f}_X(x)} \int_A \sum_{i=1}^n K_X\left(\frac{x-x_i}{h_X}\right) K_Y\left(\frac{y-y_i}{h_Y}\right) dy \quad (14.62)$$

$$= \frac{1}{nh_X h_Y \hat{f}_X(x)} \sum_{i=1}^n K_X\left(\frac{x-x_i}{h_X}\right) \int_A K_Y\left(\frac{y-y_i}{h_Y}\right) dy \quad (14.63)$$

$$= \frac{1}{nh_X \hat{f}_X(x)} \sum_{i=1}^n K_X\left(\frac{x-x_i}{h_X}\right) C_Y(A, y_i, h_Y) \quad (14.64)$$

If we select the data point i with a weight proportional to $K_X\left(\frac{x-x_i}{h_X}\right)$, and then generate \tilde{Y} from the K_Y distribution centered at y_i , then, \tilde{Y} will follow the appropriate probability density function.

14.7.2 Drawing from Histogram Estimates

Sampling from a histogram estimate is also simple, but in a sense goes in the opposite order from kernel simulation. We first randomly pick a bin by drawing from a multinomial distribution, with weights proportional to the bin counts. Once we have a bin, we draw from a uniform distribution over its range.

14.7.3 Examples of Simulating from Kernel Density Estimates

To make all this more concrete, let's continue working with the `oecdpanel` data. Section 14.4.5 shows the joint pdf estimate for the variables `popgro` and `inv` in that data set. These are the logarithms of the population growth rate and investment rate. Undoing the logarithms and taking the density gives Figure 14.4.

Let's abbreviate the actual (not logged) population growth rate as X and the actual (not logged) investment rate as Y in what follows.

Since this is a joint distribution, it implies a certain expected value for Y/X , the ratio of investment rate to population growth rate¹¹. Extracting this by direct calculation from `popinv2` would not be easy; we'd need to do the integral

$$\int_{x=0}^1 \int_{y=0}^1 \frac{y}{x} \hat{f}_{X,Y}(x,y) dy dx \quad (14.65)$$

To find $\mathbb{E}[Y/X]$ by simulation, however, we just need to generate samples from the joint distribution, say $(\tilde{X}_1, \tilde{Y}_1), (\tilde{X}_2, \tilde{Y}_2), \dots, (\tilde{X}_T, \tilde{Y}_T)$, and average:

$$\frac{1}{T} \sum_{i=1}^T \frac{\tilde{Y}_i}{\tilde{X}_i} = \tilde{g}_T \xrightarrow{T \rightarrow \infty} \mathbb{E} \left[\frac{Y}{X} \right] \quad (14.66)$$

where the convergence happens because that's the law of large numbers. If the number of simulation points T is big, then $\tilde{g}_T \approx \mathbb{E}[Y/X]$. How big do we need to make T ? Use the central limit theorem:


$$\tilde{g}_T \rightsquigarrow \mathcal{N}(\mathbb{E}[Y/X], \mathbb{V}[\tilde{g}_1]/\sqrt{T}) \quad (14.67)$$

How do we find the variance $\mathbb{V}[\tilde{g}_1]$? We approximate it by simulating.

Code Example 29 is a function which draws from the fitted kernel density estimate. First let's check that it works, by giving it something easy to do, namely reproducing the means, which we *can* work out:

```
signif(mean(exp(oecdpanel$popgro)), 3)
## [1] 0.0693
signif(mean(exp(oecdpanel$inv)), 3)
## [1] 0.172
signif(colMeans(rpopinv(200)), 3)
## pop.growth.rate   invest.rate
##           0.0697           0.1660
```

¹¹ Economically, we might want to know this because it would tell us about how quickly the capital stock per person grows.



```
popinv2 <- npudens(~exp(popgro) + exp(inv), data = oecdpanel)
```

Figure 14.4 Gaussian kernel density estimate for the un-logged population growth rate and investment rate. (Plotting code omitted — can you re-make the figure?)

This is pretty satisfactory for only 200 samples, so the simulator seems to be working. Now we just use it:

```
z <- rpopinv(2000)
signif(mean(z[, "invest.rate"]/z[, "pop.growth.rate"], 3)
## [1] 2.61
signif(sd(z[, "invest.rate"]/z[, "pop.growth.rate"])/sqrt(2000), 3)
## [1] 0.0349
```

```

rpopinv <- function(n) {
  n.train <- length(popinv2$dens)
  ndim <- popinv2$ndim
  points <- sample(1:n.train, size = n, replace = TRUE)
  z <- matrix(0, nrow = n, ncol = ndim)
  for (i in 1:ndim) {
    coordinates <- popinv2$eval[points, i]
    z[, i] <- rnorm(n, coordinates, popinv2$bw[i])
  }
  colnames(z) <- c("pop.growth.rate", "invest.rate")
  return(z)
}

```

CODE EXAMPLE 29: *Simulating from the fitted kernel density estimate `popinv2`. Can you see how to modify it to draw from other bivariate density estimates produced by `npudens`? From higher-dimensional distributions? Can you replace the `for` loop with less iterative code?*

This tells us that $\mathbb{E}[Y/X] \approx 2.61 \pm 0.035$.

Suppose we want not the mean of Y/X but the median?

```

signif(median(z[, "invest.rate"]/z[, "pop.growth.rate"]), 3)
## [1] 2.31

```


Getting the whole distribution of Y/X is not much harder (Figure 14.5). Of course complicated things like distributions converge more slowly than simple things like means or medians, so we might want to use more than 2000 simulated values for the distribution. Alternately, we could repeat the simulation many times, and look at how much variation there is from one realization to the next (Figure 14.6).

Of course, if we are going to do multiple simulations, we could just average them together. Say that $\tilde{g}_T^{(1)}, \tilde{g}_T^{(2)}, \dots, \tilde{g}_T^{(s)}$ are estimates of our statistic of interest from s independent realizations of the model, each of size T . We can just combine them into one grand average:

$$\tilde{g}_{s,T} = \frac{1}{s} \sum_{i=1}^s \tilde{g}_T^{(i)} \quad (14.68)$$


As an average of IID quantities, the variance of $\tilde{g}_{s,T}$ is $1/s$ times the variance of $\tilde{g}_T^{(1)}$.

By this point, we are getting the sampling distribution of the density of a nonlinear transformation of the variables in our model, with no more effort than calculating a mean.



```
YoverX <- z[, "invest.rate"]/z[, "pop.growth.rate"]  
plot(density(YoverX), xlab = "Y/X", ylab = "Probability density", main = "")  
rug(YoverX, side = 1)
```

Figure 14.5 Distribution of Y/X implied by the joint density estimate `popinv2`.



```

plot(0, xlab = "Y/X", ylab = "Probability density", type = "n", xlim = c(-1, 10),
     ylim = c(0, 0.3))
one.plot <- function() {
  zprime <- rpopinv(2000)
  YoverXprime <- zprime[, "invest.rate"]/zprime[, "pop.growth.rate"]
  density.prime <- density(YoverXprime)
  lines(density.prime, col = "grey")
}
invisible(replicate(50, one.plot()))

```

Figure 14.6 Showing the sampling variability in the distribution of Y/X by “over-plotting”. Each line is a distribution from an estimated sample of size 2000, as in Figure 14.5; here 50 of them are plotted on top of each other. The thickness of the bands indicates how much variation there is from simulation to simulation at any given value of Y/X . (Setting the `type` of the initial `plot` to `n`, for “null”, creates the plotting window, axes, legends, etc., but doesn’t actually plot anything.)

14.8 Further Reading

Good introductory treatments of density estimation can be found in [Simonoff \(1996\)](#) and [Wasserman \(2006\)](#). My treatment of conditional density estimation is based on [Hall *et al.* \(2004\)](#).

The Glivenko-Cantelli theorem has a more “quantitative” version, the “Dvoretzky-Kiefer-Wolfowitz inequality”, which asserts that with IID samples from a one-dimensional CDF F ,

$$\Pr \left(\sup_x |\hat{F}_n(x) - F(x)| > \epsilon \right) \leq 2e^{-2n\epsilon^2} \quad (14.69)$$

and the constants appearing here are known to be the best that hold over all distributions ([Wasserman 2006](#), §2.2); this can be inverted to get confidence bands for the CDF.

On empirical process theory, see [Pollard \(1989, 1990\)](#); [van de Geer \(2000\)](#); [Pollard \(1989\)](#) is especially good as an introduction. [Devroye and Lugosi \(2001\)](#) applies empirical process theory to density estimation, as well as forcefully advocating measuring error using the L_1 distance, $\int |\hat{f}(x) - f(x)| dx$. In this chapter I have stuck to L_2 , partly out of tradition and partly out of desire to keep the algebra simple (which in turn helps explain the tradition).

Historical notes

I do not know of a good history of the Glivenko-Cantelli theorem (but would like to read one).

Histogram estimates are very old; the word “histogram” was apparently coined by Karl Pearson in the 1890s¹², but as a convenient name for an already-common type of graphic. Kernel density estimation seems to have first been proposed by [Rosenblatt \(1956\)](#) (see especially section 4 of that paper). It was re-introduced, independently, by [Parzen \(1962\)](#), and some of the analysis of the error in KDEs that we saw above goes back to this paper.

Exercises

- 14.1 Reproduce Figure [14.4](#)?
- 14.2 Qualitatively, is this compatible with Figure [14.1](#)?
- 14.3 How could we use `popinv2` to calculate a joint density for `popgro` and `inv` (not `exp(popgro)` and `exp(inv)`)?
- 14.4 Should the density `popinv2` implies for those variables be the same as what we’d get from directly estimating their density with kernels?
- 14.5 You are given a kernel K which satisfies $K(u) \geq 0$, $\int K(u) du = 1$, $\int uK(u) du = 0$, $\int u^2 K(u) du = \sigma_K^2 < \infty$. You are also given a bandwidth $h > 0$, and a collection of n univariate observations x_1, x_2, \dots, x_n . Assume that the data are independent samples from some unknown density f .

¹² See Jeff Miller (ed.), “Earliest Known Uses of Some of the Words of Mathematics”, s.v. “Histogram”, <http://jeff560.tripod.com/h.html>. I have not verified the references cited there, but have found the site to be generally reliable.

1. Give the formula for \hat{f}_h , the kernel density estimate corresponding to these data, this bandwidth, and this kernel.
2. Find the expectation of a random variable whose density is \hat{f}_h , in terms of the sample moments, h , and the properties of the kernel function.
3. Find the variance of a random variable whose density is \hat{f}_h , in terms of the sample moments, h , and the properties of the kernel function.
4. How must h change as n grows to ensure that the expectation and variance of \hat{f}_h will converge on the expectation and variance of f ?

14.6 *The transformation method* Many variables have natural range restrictions, like being non-negative, or lie in some interval. Kernel density estimators don't respect these restrictions, so they can give positive probability density to impossible values. One way around this is the **transformation method** (or "trick"): use an invertible function q to map the limited range of X to the whole real line, find the density of the transformed variable, and then undo the transformation.

In what follows, X is a random variable with pdf f , Y is a random variable with pdf g , and $Y = q(X)$, for a known function q . You may assume that q is continuous, differentiable and monotonically increasing, inverse q^{-1} exists, and is also continuous, differentiable and monotonically increasing.

1. Find $g(y)$ in terms of f and q .
2. Find $f(x)$ in terms of g and q .
3. Suppose X is confined to the unit interval $[0, 1]$ and $q(x) = \log \frac{x}{1-x}$. Find $f(x)$ in terms of g and this particular q .
4. The beta distribution is confined to $[0, 1]$. Draw 1000 random values from the beta distribution with both shape parameters equal to $1/2$. Call this sample \mathbf{x} , and plot its histogram. (Hint: `?rbeta`.)
5. Fit a Gaussian kernel density estimate to \mathbf{x} , using `density`, `npudens`, or any other existing one-dimensional density estimator you like.
6. Find a Gaussian kernel density estimate for `logit(x)`.
7. Using your previous results, convert the KDE for `logit(x)` into a density estimate for \mathbf{x} .
8. Make a plot showing (i) the true beta density, (ii) the "raw" kernel density estimate from [35](#) and (iii) the transformed KDE from [37](#). Make sure that the plotting region shows all three curves adequately, and that the three curves are visually distinct.

[[TODO:
Fix exercise cross-reference numbers here!]]