It's impossible to provide a single "correct" solution to this problem: like so many real-life problems, you can get it wrong, but there's no unique "right answer" to it. You came up with some great approaches to it. I've used several different approaches in the accompanying spreadsheet, which aren't the *best* approaches by any stretch, but I hope them to be demonstrative of several different approaches that make some sense. Here I make some general comments about the questions and explain a few details.

a. This is most easily understood by first considering only the photons in one specific wavelength range: let's consider only photons between 399.5 and 400.5 nm, for starters. According to the AM0 data,[1] the total energy brought into a 1 m^2 area in earth orbit, directly facing the sun, by the photons in this wavelength window is 1.655 W. Those photons are in a tight enough wavelength range that we can treat them as all having the same energy, which is given by

$$E_{photon} = \frac{hc}{\lambda} = \frac{(6.6261\times10^{-34}\,\text{J}\cdot\text{s})(2.9979\times10^{8}\,\text{m/s})}{400.\text{ nm}}\left(\frac{10^9\,\text{nm}}{1\text{ m}}\right) = 4.96_6\times10^{-19}\,\text{J}\left(\frac{1\text{ eV}}{1.602\times10^{-19}\,\text{J}}\right) = 3.10\text{ eV}$$

So we can now calculate how many of these photons are required to supply 1.655 W/m^2 of energy:

$$J_{photon}(\lambda = 400.\text{ nm}) = 1.655\frac{\text{W}}{\text{m}^2\cdot\text{nm}}\left(\frac{\text{J}\cdot\text{s}^{-1}}{\text{W}}\right)\left(\frac{1\text{ photon}}{4.96_6\times10^{-19}\,\text{J}}\right) = 3.33_{27}\times10^{18}\,\frac{\text{photons}}{\text{m}^2\cdot\text{nm}\cdot\text{s}}$$

This same process can be easily applied to each wavelength range in the spectrum if you are using a spreadsheet. The only trick is that the spectral data doesn't always span 1 nm segments, but the power is always given per nm of wavelength; so if you want to use the average λ for each range, that's groovy, but you don't want to divide by the width of the wavelength window. Put more concretely, the AM0 spectrum indicates an incoming solar irradiance of 1.115 $W\cdot m^{-2}\cdot nm^{-1}$ between 811 and 813 nm: the total power provided by photons between 811 and 813 nm is 1.115 $\frac{\text{W}}{\text{m}^2\cdot\text{nm}}(2\text{ nm}) = 2.230\frac{\text{W}}{\text{m}^2}$, 1.115 W/$m^2$ per nm of spectral width. [If you did this wrong, you'd know it, because there would be huge jumps in your $J_{photon}(\lambda)$ spectrum where the data spacing changed.]

Almost as readily, you can modify the $J_{photon}(\lambda)$ function to give spectral irradiance ($W\cdot m^{-2}\cdot nm^{-1}$) by multiplying it by the energy of the photons involved (which is a function of λ and some constants):

$$J_{\text{spectral irradiance}}(\lambda) = \left[J_{photon,blackbody}(\lambda)\right]\left[E_{photon}(\lambda)\right] = \left[\frac{2\pi c}{\lambda^4}\left(\frac{1}{e^{hc/\lambda kT}-1}\right)\right]\left[\frac{hc}{\lambda}\right] = \frac{2\pi hc^2}{\lambda^5}\left(\frac{1}{e^{hc/\lambda kT}-1}\right)$$

b. First, the best data set to use for this purpose is the AM0 data, which describes the solar radiation that hits a 1 m^2 panel in earth orbit, pointed directly at the sun. The AM1.5 data describes the light that arrives at a similar panel pointed at the sun, but located on the earth's surface at a time when the sun isn't directly overhead. Because the atmosphere absorbs a lot of the sun's light, trying to measure the sun's temperature from the AM1.5 data is an appreciably more contorted endeavor; you have to extract out the powerful effect of the atmospheric gases to get back to anything resembling a blackbody spectrum. (What's more, the ozone layer absorbs effectively every UV photon with λ < 300. nm, so you totally lose that portion of the data in the AM1.5 spectrum; so too with several spectral segments in the IR portion of the spectrum that are completely absorbed by terrestrial greenhouse gases.) The AM0 data is much closer to the distribution that's actually emitted by the hot parts of the sun, though it's still not perfect. It is actually the result of blackbody emission from many different-temperature regions of the sun, convoluted with the absorption spectrum of cooler H_2 and He gases in the other layers of the sun and the general scattering of solar photons by the dust particles in space between the earth and the sun (these are what make the short-wavelength portion of the spectrum so jagged; see the description of the sun below for more information.)

The AM0 data provided by NREL is accurate to the number of significant figures they indicate, but because of the factors above, it doesn't look like a perfect blackbody spectrum. It's perfectly legit to de-weight or even throw out data points that don't reflect the blackbody emissive nature of the sun, given that we have a reasonable explanation for their cause and we are trying to answer a question that starts "Assuming the sun is indeed behaving as a blackbody…" I tried several tricks to massage the data and/or pound out a good fit. In each of these cases, I did not attempt to use the intensity of the radiation, and correct for distance attenuation and the surface area of the sun, to determine the temperature of the sun. That's very hard because it is hard to know what surface area of the sun to use (see the descriptions of the sun that appear below), and not as "primary," because you have to rely on other measured data (the radius of the earth's orbit and the surface area of the photosphere, both of which vary somewhat with time). This is not to say that method is wrong or doesn't work, but I don't think it is quite as good as the general wavelength dependence fitting alternative that most of you (and I) used. I tried to find temperatures for which the blackbody emission spectrum had the same spectral distribution as the AM0 data; that is to say, I scaled the intensity vs. wavelength data blindly, and looked for matches in the relative intensity of the two distributions across a wide wavelength range. You can see the details in the accompanying spreadsheet.

First, I tried simple trial-and-error to get a blackbody temperature estimate. That gave me 6200 K. [I reasoned that the best fit would be one that did not under-estimate the spectral data, but consistently overestimated it while staying as close as possible. This is because while we can come up with plenty of reasons for attenuation of photons from the sun, there are no photon sources out there!) Next, I used a straight-up least squares regression applied to all of the AM0 data. That spat back 54_{47} K. Next, unhappy with that fit's match to the peak region of the AM0 spectrum, I employed a weighted regression, multiplying the squared error associated with each point by the square of the AM0 intensity at that point. This has the effect of making the AM0 data of high intensity more important to the curve-fit than the low-intensity data, and improving the match to the peak location at the expense of the fit in the wings of the spectrum. The best temperature for this approach was 56_{76} K, but it still didn't look very good near the peak of the AM0 spectrum. So I upped the ante and weighted the peak region data even more heavily by employing a weighting factor equal to the exponential of the square of the AM0 intensity. That weights most of the low intensity points about equally (1.00 something) but puts a really heavy emphasis on the AM0 points with intensities above 1.0 $W \cdot m^{-2} \cdot nm^{-1}$. This fit came in at 56_{76} K, and got the peak closer to where my eyes thought it should be. Note that these weightings are convoluted with the natural variations in the magnitude of the data and the semi-arbitrary spacing of data points employed by NREL in preparing the AM0 data table.

But there were still some big complaints that could be made, the most pressing being that the data I included in the first set of fits clearly had a lot of intensity cut out of it by absorption due to the cool outer gases of the sun, in the chromosphere and corona. To try to select out the "good" data, I set up a column that looked to see if each point in the AM0 column was smaller than any of the 10 preceding AND simultaneously smaller than any of the 10 subsequent points. If so, I trashed that point as a local low spot. That did a pretty passable job of throwing out the absorption troughs. Modifying the least squares regressions I'd used before to weight any of the excluded points at zero, I repeated my curve fits from before. This time the simple regression netted 56_{43} K, the $AM0^2$ weighted regression spat out 58_{67} K, and the $\exp(AM0^2)$ weighting gave 59_{63} K. These fits were about as good as my trial-and-error result, but I wouldn't say better. I guess that's one of the take-home lessons from all of this…you're smarter than your computer. I knew that near the peak of the AM0 spectrum there's a lot of absorption notches cut out of the spectrum. The undisturbed spectrum has to be at least as high as the highest point in the AM0 spectrum, so a good fit should follow very close to the maximum intensity, erring on the high side, if anything. None of the computerized fits manage that (not that it's really possible to do a great job of it – the AM0 spectrum really isn't *that* close to a blackbody). I'd put the uncertainty in my best temperatures at about ± 200 K; less than that, and it really becomes hard to say that one fit is better than another at the level that really matters, which isn't a square-of-the-errors sum!

I tried one more trick, and was rather surprised at how poorly it worked. As I blathered above, the sun is actually a many-temperature thing. So I tried treating it as two different-temperature blackbodies. I tried this several different ways, in the end resorting to an $\exp(AM0^2)$ weighted fit to the "sifted" data. The solver would constantly set the temperatures of the two blackbodies equal, or reduce the intensity of one of them to the point where it was irrelevant. What it was telling me, and I think it was right about this, is that the AM0 spectrum can't be fit appreciably better with the sum of two blackbody emissions than it can with a single blackbody. That's unusual in the sense that usually the more adjustable parameters you throw at a problem, the better a fit you can get. But a blackbody spectrum is a funny thing, and because it is highly non-linear, with a minimum at both zero and large ordinate values, general rules for simpler systems certainly may not apply to it. At any rate, the temperature that came out of this was effectively the same as that which came out of the single blackbody fit using the same input data and weighting, 59_{64} K.

I want to emphasize that the numbers I give in the work above are written out with a greater number of figures than they deserve, to make it easier to tell which calculation in the spreadsheet I am referring to. As I said, I did not consider any of my fits to be more reliable than ± 200 K, as swings in temperature over that range led to differences in the visible quality of the fit or the value of $\Sigma(\text{weighted error})^2$ that were too small to rely on as meaningful. I also found that small changes in my point selection criteria, equalizing the number of data points per nm of wavelength, and other such changes led to shifts in the best-fit temperatures on the order of 100's of K. The table below indicates the results as I would have reported them as answers to the question posed in the problem set:

Method	Result
trial-and-error visual fit	6200 ± 250 K
simple least squares fit on all data	54_{50} ± 800 K
$[J(\lambda)]^2$ weighted fit on all data	56_{80} ± 400 K
$\exp\{[J(\lambda)]^2\}$ weighted fit on all data	5800 ± 300 K
simple least squares fit on selected data	56_{40} ± 600 K
$[J(\lambda)]^2$ weighted fit on selected data	58_{70} ± 300 K
$\exp\{[J(\lambda)]^2\}$weighted fit on selected data	59_{60} ± 200 K
$\exp\{[J(\lambda)]^2\}$ weighted fit on selected data for two independent blackbody sources	both 59_{60} ± 200 K

c. Let's start with the validity of the question. The ostensible premise is that the sun is a single-temperature blackbody. Well, that's not an *entirely* unreasonable thought. It's darn hot, and it certainly emits more light as a blackbody than as a plasma line emission source. (Some light ***is*** emitted by the electronic relaxation of superheated atoms; the sun is a giant ball of plasma, with a specific atomic composition. As electrons relax on each of those atoms, they emit photons of very particular wavelengths–their distribution certainly doesn't follow the blackbody distribution. But they emit a negligible fraction of the photons that make it to earth orbit.) But single-temperature? The sun is certainly not a single-temperature body. The core is much hotter than the outer layers; the sunspots are local cold spots, and the flares and prominences are exorbitantly hot. Each of these parts of the sun emits radiation characteristic of a blackbody, but that from the core is largely absorbed before it can escape, while the intensity coming from the relatively small number of solar flares is (almost) negligible.

Another factor worthy of consideration is temporal variability. Does the sun remain constant with time? Specifically, are all of its parts at constant temperatures? Actually, no! As you probably know, solar flares occur "seasonally" on the sun, more intense at some times and less intense at others. While the fusion reaction occurring inside the sun is a bit more stable, it too is something of a weather system, with periods of relative activity and inactivity. Most of the temporal variation occurs in the high-energy portion of the spectrum, suggesting it is mostly due to changes in solar flares and related phenomena, and in the changes they cause in the absorption spectrum of the colder outer layers of the sun. The variation has actually been measured! See http://science.nasa.gov/headlines/images/sunbathing/sunspectrum.htm for details.

Finally, understand that the AM0 data we are working with in this problem reflects the effect of scattering by the "space dust" between the sun and the earth. While there is not much of this dust, it consists of very small particles, which interact (scatter) most strongly those photons similar in size to themselves or smaller. This is the reason for the clear fall-off in spectral intensity at short wavelengths, relative to what we would expect for an ideal blackbody. The same phenomenon makes the sky appear blue and the sun appear yellow to red. Shorter wavelength photons are more readily scattered by sub-microscopic particles than are longer wavelength particles, so blue photons are disproportionately scattered and we end up with less of them than we would expect for a blackbody emission spectrum. The better temperature estimates found in the literature are based on spectral data collected from satellites in near-sun orbit, which doesn't suffer this aberration as much, if at all. So those estimates tend to come out a little higher than do ours, but this is the fault of the data we are using, not a fault in our methodology. (Except insofar as we could do better with better input data…the good old "garbage in, garbage out" paradox!)

So what do "the experts" say the temperature of the sun actually is, to the extent that this is a meaningful question? Roy A. Gallant provides a nice description in the National Geographic Picture Atlas of Our Universe.[2] I'll piece together several factoids from this book to give a pretty good overall picture of the sun:

> The core temperature exceeds 15 million Kelvins; at that temperature, mostly cosmic rays are emitted. The temperature gradually falls as you move away from the core, to about 6000 K at the photosphere. Thankfully, most of the radiation emitted by the high-temperature core is absorbed by the cooler outer layers. The bulk of the radiation we see from the sun is emitted by the photosphere. The corona and chromosphere, which are lower-density layers above the photosphere, contain far fewer atoms, so they neither emit nor absorb very much light. Solar flares and prominences can get as hot as 10000 K, but they are rare enough that they don't contribute an appreciable amount of energy to the solar spectrum.

(See http://solar.physics.montana.edu/YPOP/Spotlight/SunInfo/sundiag.html for a diagram of the sun.)

The key statement in the above description is that "the bulk of the radiation we see from the sun is emitted by the photosphere." Our use of the AM0 spectrum amounts to estimating the temperature of the sun based on a photon-weighted average emission: that is to say, we are measuring the temperature of whatever part(s) of the sun provide the bulk of the photons (or, alternately, the spectral irradiance) we are considering. To couch our answer in cautious terms, we should not claim to be actually measuring the temperature of the sun. Rather, the solar emission spectrum agrees most closely with the blackbody emission spectrum of a (insert value here) K blackbody.

The temperature of the surface of the photosphere isn't terribly well-defined, both because it varies somewhat with time and from point to point on the surface, and because the sun doesn't have a distinct surface that you can choose to define as the clear boundary between the photosphere and the chromosphere. But it's widely agreed that this temperature, to the extent it's meaningful, is at or a bit below 6000 K, with specific values given by specific sources appearing in the table below:

Source	Effective Blackbody Temperature of the Sun
NASA National Space Science Data Center, http://nssdc.gsfc.nasa.gov/planetary/factsheet/sunfact.html	5778 K
Dr. Judith Lean at the US Naval Research Laboratory, http://science.nasa.gov/headlines/images/sunbathing/sunspectrum.htm	5770 K
Kristine Spekkens, Astronomy Department, Cornell University http://curious.astro.cornell.edu/question.php?number=126	5880 K
NASA/European Space Agency SOHO Project, http://sohowww.nascom.nasa.gov/explore/sun101.html	5800 K

Here's a quick rundown of how you folks did this, and what you obtained:

Brian, Ghidewon, and Nate used a wavelength-based fit, strongly weighted in favor of wavelengths greater than 600 nm, to come up with a value of 5790 K. Andy of the Wills Clan calculated a theoretical value for the "scaling factor" relating the blackbody equation and the AM0 data, and used this with weightings of the AM0 data that rejected specific absorbances to carry out a least-squares fit. He determined the temperature of the sun's photosphere to be $58_{14} \pm 28$ K. Lawrence, Aistis, and Andy of the Lorenz Clan came up with 5770 K using an approach very similar to that of Sir Wills, but a bit more willing to throw out data points because of specific absorption. Jill employed a similar approach as well, but made use of Mathematica's nonlinear fit algorithm to obtain 5800 K. Anne used both AM0 and AM1.5 data to get a range of results and to keep herself from getting overconfident in her results. She sagely reported a temperature of around 5000 K. Nancy used a very ingenious method for automatically selecting and de-weighting local low points; variations in this approach gave her best fit temperatures of 5668 K, 5673 K, and 5684 K, which are systematically lower than the values in the table above for the very good reason of the preferential scattering of blue photons as sunlight makes its way from the sun to the earth through dusty space. Micah and Andy of the Aults reported a best-fit value of 5634 K, based on a data set carefully de-weighted for specific absorption. Andrew used the segments of the AM0 data that were least impacted by specific absorption, weighting these segments intelligently, to come up with 5600 K. Joey reported 5610 ± 200 K, based on a beautiful data weighting and error analysis job well beyond my own! Devin's 5700 K value was obtained with an effective minimalist strategy and some wise error analysis. You were all in remarkably good agreement with each other, all told!

Several of you found very useful web pages, including this great one:

http://science.nasa.gov/ssl/PAD/SOLAR/default.htm

My only general gripe is that there's a fishy correlation between the values some people reported and the literature value they found. It's as if you trust the literature more than you trust yourselves! Do you really think the people who came up with the literature values were smarter than you, or worked at it more carefully?!? Think again! If anything, they were sloppier, because they know that to a significant extent they are trying to stick a square peg in a round hole in treating the sun as a single-temperature blackbody, and so there's no point in being super-precise with the method used, the question itself involves some inherent inaccuracy. Those few who were trying harder than you were also using a better data set than you were, so you shouldn't expect to get the same results that they did.

Here are a few other useful references I or my students have bumped into over the years:

http://library.wolfram.com/webMathematica/MSP/Explore/Astronomy/Blackbody
http://nssdc.gsfc.nasa.gov/planetary/factsheet/sunfact.html

Bibliography

[1] http://rredc.nrel.gov/solar/standards/am0/#spectrum, accessed on March 21, 2003.
[2] R. A. Gallant, *National Geographic Picture Atlas of Our Universe*, National Geographic Society: Washington, D.C. (1980).