Tips on Teaching General Relativity (with Tensors) to Undergraduates

Thomas A. Moore, Pomona College
AAPT General Relativity Workshop, July 2006

Abstract. This article will present some guiding principles (gleaned from many years of painful experience) for successfully teaching a tensor-based course in general relativity to undergraduates. These principles include (1) simultaneously developing the physics and mathematics, (2) liberally using two-dimensional analogies, (3) building on students' understanding of vectors and vector spaces, (4) designing drills to help students overcome common misconceptions about tensor notation, (5) helping students "own" the derivations, (6) designing a homework-grading scheme that allows students to try hard problems and learn from corrections. I will also describe some tricks and worksheets that I have developed that help students easily evaluate Christoffel symbols and Ricci tensor components for diagonal metrics.

Overview of the Problem

General relativity (GR) has undergone an amazing transformation in the past few decades, moving from being a comparatively inactive and mostly theoretical subject to a topic supporting lively experimental and computational research programs. The topic's current vitality and increasingly important applications in astrophysics (and even engineering!) make the time ripe for GR to become more regular part of our undergraduate offerings. However, the fact that GR is naturally expressed in the abstract and (for most undergraduates) unfamiliar language of tensor calculus makes achieving this goal more difficult.

The first generation of GR textbooks (at least partially) targeted toward U.S. undergraduates (two superb examples are Schutz[1] and Ohanian and Ruffini[2]) more or less followed the order of presentation common to graduate-level GR texts: a thorough discussion of the mathematical concepts of the theory followed by the derivation of solutions to the Einstein equation and discussion of their implications. This (very logical) "math first" approach has serious pedagogical drawbacks, not the least of which is that a long mathematical preamble can be deadly. As a result, many first-generation books tend to rush too quickly through the mathematics for undergraduates to gain the mastery and insight they need.

In reaction, the second generation of GR textbooks now emerging, led by Hartle's excellent textbook[3] and recently followed by Cheng[4], seek to "put the physics first", introducing only the concept of the metric and summation notation before digging into applications. The metrics needed for those applications are asserted, not derived. The full mathematics of curvature and the Einstein equation is only discussed at the end of the book, and both authors are frank about the possibility one might never get to see the full machinery of tensor calculus in a typical undergraduate course.

While this general approach is a necessary corrective step, it also has (in my opinion) a serious pedagogical flaw. GR's deep logic (and graceful beauty) is founded on drawing physical implications from the simple model of curved spacetime. Teaching GR without the tensor language needed to understand this model deeply is like studying Shakespeare's plays after translation into 9th-grade English (or, more prosaically, Newtonian mechanics without calculus): some specific things become easier to understand, but something vital has been lost. This is not simply an idealistic concern: it has very practical pedagogical consequences. When I attempted this approach myself some years back, I found that my students, without the firm moorings provided by the theory's logic, found themselves rudderless in the sea of applications and, in particular, experienced their own understanding of GR to be inadequate. This was *not* the empowering experience I had in mind for them!

The purpose of this paper is to present a third way that minimizes neither the difficulty nor the importance of tensor calculus, but rather faces the challenge squarely and provides the pedagogical tools needed for students to become capable and satisfied users of tensors.

The rest of this will article discuss seven principles for designing undergraduate courses that can successfully teach GR (with tensors) to undergraduates. These principles grow partly out of my experience teaching GR to undergraduates roughly 13 times in the past 24 years, and partly out of my experience in addressing similar challenges while developing my introductory physics text *Six Ideas That Shaped Physics*. They also draw somewhat on the concepts of physics educational research. The principles are:

1. Develop the physics and math simultaneously
2. Use two-dimensional visualizations
3. Keep the math (appropriately) simple
4. Drill students on tensor notation
5. Help students "own" the math
6. Grade homework appropriately
7. Use tricks and tools to avoid tedious calculations

Each remaining section in this article discusses one of these principles in depth.

1 Develop the physics and math simultaneously

Teaching undergraduates tensor calculus takes time, and there are at least two reasons why it is not a good idea to do this all at the beginning of the course. First, delaying the gratification of the juicy physics for many weeks will sap the students' (and even the instructor's) motivation. Secondly, students need time for the strange concepts of GR and tensor calculus to percolate into their none-too porous brains. It is advantageous, therefore, to interweave the math and physics throughout the course.

Table 1 describes a possible 13-week syllabus (with 3 classes / week) that illustrates how one might spread the development of the mathematics over nearly the entire course. (In a 14-week semester, this syllabus allows for in-class midterms and/or review days before take-home exams.) The starred classes are primarily devoted to development of mathematical and/or theoretical concepts.

1	Conceptual Overview	Review of Relativity	Four-Vectors
2	*Index Notation*	*Arbitrary Coordinates*	*Tensor Equations*
3	Maxwell's Equations	*Geodesics*	The Schwarzchild Metric
4	Particle Orbits	Perihelion Precession	Photon Orbits
5	Gravitational Lenses	Event Horizon	Alternative Coordinates
6	BH Thermodynamics	The Kerr Metric	Kerr Particle Orbits
7	Ergoregion and Horizon	Negative Energy Orbits	The Penrose Process
8	*The Absolute Gradient*	*Geodesic Deviation*	*The Riemann Tensor*
9	*Stress Energy Tensor*	*The Einstein Equation*	Interpreting the Equation
10	Schwarzchild Solution	The Observed Universe	A Cosmic Metric
11	Evolution of the Universe	Cosmic Implications	The Early Universe
12	*Linearized Gravity*	*Gauge Freedom*	Gravitational Waves
13	"Energy" in GWs	Generation of GWs	Applications

Table 1 A sample GR course syllabus

The mathematical classes in weeks 2 and 3 provide sufficient basis to understand the metric tensor, how to handle arbitrary coordinates (even in potentially curved spaces), basic tensor concepts (including the the benefit of writing equations in covariant form), and calculating geodesics (using the calculus of variations). The interlude on Maxwell's equations provides a break that gives students a chance to consolidate their understanding tensor notation in the (more or less) familiar context of electricity and magnetism as well as illustrating the power of the notation by "deriving" Maxwell's equations as the natural covariant generalization of Gauss's law. After the derivation of the geodesic equation in Week 3, it is natural to spend time exploring the motion of particles in Schwarzchild and Kerr spacetimes. These metrics are not derived, but once they given, one can study their physical implications in great depth. So far, the syllabus is very similar to the one implied by Hartle's book.

However a second block of mathematical development in weeks 8 and 9 provide the foundation for solving the Einstein equation. Weeks of familiarity with index notation and with practical applications make these difficult concepts easier for students than it would have been if they had been presented at the beginning. They verify the Schwarzchild solution (proving Birkoff's theorem along the way) and then generate the Friedmann-Robertson-Walker metric before exploring its meaning. A third and final blast of mathematics regarding the linearized approximation to the Einstein field equations sets up an exploration of gravitational waves.

The point is that in a class structured this way, students have a chance to ingest the mathematics in more-or-less bite-sized chunks, and thoroughly digest each bit before moving on to the next. This greatly increases the amount of nourishment students can absorb from the mathematics.

Even when appropriately spread out, though, tensor calculus remains difficult. The next four sections describe techniques for making the mathematics more accessible.

2 Use two-dimensional visualizations

Students in my classes practice using *every* new tensor concept (the meaning of the metric, the difference between a coordinate basis and an orthonormal basis, the distinction between covector and vector components, coordinate and tensor transformations, the absolute gradient, Christoffel symbols, geodesic deviation, measures of curvature) in the context of easily visualized two-dimensional flat and curved spaces *before* applying them to four dimensional spacetimes. They explore not only polar coordinates in flat space and longitude/latitude coordinates on a sphere, but stranger coordinate systems for flat space and other types of curved spaces. Practicing the formalism and the use of these ideas in concrete situations that students can easily visualize helps them avoid errors and builds their intuition about and confidence in the mathematics.

This may seem like a simple idea, but I have found it to be both essential and yet underutilized in most textbooks (Hartle's text moves in the right direction, though). Using more two-dimensional visualizations is something that we can all do more of (even in the "physics first" models) to make the mathematics more accessible.

3 Keep the math (appropriately) simple

Having earlier praised the deep beauty of GR and the importance of tensor calculus in appreciating that beauty, let me now issue a warning against too much beauty! Undergraduates *can* be taught tensor calculus, but just barely. This makes it imperative to keep the mathematical formalism as conceptually simple as possible.

In order to teach students effectively, one needs to meet them where they are. Junior and senior undergraduates are generally reasonably competent with vector calculus and using vector components in calculations. Therefore, I have found it helpful to build on this strength by treating tensors almost exclusively as a mathematical quantity represented by a collection of components that transform a certain way, thus making tensors simply one step more complicated than a vector. In keeping with this approach, I deliberately avoid trying to teach students at this level about one-forms, dual spaces and dual bases, and the isomorphism between basis vectors and gradient operators. I also avoid the Levi-Civita tensor, Killing vector fields, wedge products, and so on. These ideas add elegance and power to the mathematics but are beyond what is actually needed to understand the ideas and to do useful calculations.

I realize that ignoring the geometric distinction between vectors and one-forms (and between basis vectors and basis one-forms) is somewhat "retro" and obscures something valuable and conceptually important, but when I have tried to teach these ideas, I have found that they perplex students greatly and (in my opinion) unnecessarily. I honestly do not think that these distinctions are sufficiently helpful to students at this level to spend the time pushing through the confusion they initially create. Indeed, I believe that students who learn about tensors as collections of components will have a firmer foundation for understanding the deeper ideas of differential geometry (and more appreciation for their elegance) when and if they encounter them in a graduate-level course in GR, while those students who are not bound for research in GR (or even graduate school) are not left behind. As I see it, the goal in this course is neither rigor nor elegance but understanding and empowerment.

Because of this, I consider the mathematics in all of the major texts that are oriented (at least partially) to undergraduates is too sophisticated, even in Hartle and Cheng. The concern that those promoting the "physics first" approach have with the difficulty of the math is therefore simultaneously valid and (ironically) self-fulfilling. The alternative to *avoiding* the math is working hard to make it as simple and accessible as possible. I have found experimentally that undergraduates *really can* learn tensor calculus and use it competently if one does this.

While I am on the subject of simplicity, let me talk about a very small thing one can do that has a surprisingly large impact. Instead of using numbers as vector or tensor indices (e.g. p^0, p^1, p^2, p^3), I try whenever possible to use coordinate symbols (e.g. $p^t, p^r, p^\theta, p^\phi$). This is increasingly the common practice (see Hartle, for example), and I want to endorse it as being very helpful: students are less likely to see the superscripts as powers and are less likely to make mistakes. Operationally, I can report that students do not seem to confuse such coordinate-specific indices with abstract indices that can take on any value, even when greek letters such as θ and ϕ are used.

4 Drill students on tensor notation

Tensor notation is new and strange for undergraduates, but to become competent users of tensor calculus, they must learn to interpret the notation both accurately and subconsciously, so that they can focus on what the equations mean. In taking their awkward first steps toward this goal, students make a set of standard errors that can be, recognized, confronted, and resolved. Here are some common errors students make:

- They think that $\delta^\mu_\mu = 1$, not 4.
- They think that $\eta_{\mu\nu} A^\mu B^\nu = \eta_{\mu\mu} A^\mu B^\mu$.
- They think that the equation $g_{\alpha\beta} = g_{\mu\nu}\delta^\mu_\sigma \delta^\nu_\gamma$ looks reasonable.
- They don't understand why $d(p^\mu p_\mu)/d\tau = 2p_\mu(dp^\mu/d\tau)$.

and so on. Ultimately, these errors boil down to deeper problems:

- Failing to recognize implied sums
- Failing to distinguish between free and summed indices
- Failing to understand how index renaming works and what it means

and so on. Observing problems like these in student papers is one of the things that has led people to despair about teaching undergraduates tensor notation.

The alternative to despair is confronting these problems head-on by drilling students on the notation until it becomes second nature. Most published texts have too few problems that ask students to

- Identify free and summed indices
- Write out the implied sums in an equation
- Check that the free indices on both sides of the equation are consistent
- Rename indices in an appropriate way to pull out a common factor
- Identify nonsensical equations

It may seem silly to drill students on such basic things, but my experience is that spending time in such simple drills and addressing the standard errors when first teaching the notation really does improve student competence and confidence.

There is also a mantra that I try to drill (Johnny Cochran-like) into their consciousness: "When in doubt, write it out!" After students write out the implicit sums several times (and recognize that they *could* do this any time), they begin to see the structural meaning of the equations more clearly, and also gain some confidence that they can sort out any interpretation difficulties on their own.

5 Help students "own" the math

One of the most robust results of physics education research is that (at least in the context of introductory physics) students learn the material much more effectively in courses where they are actively engaged in the learning process as opposed to receiving the material in passive lectures or readings[5]. Students need to make the ideas "their own" by actively processing them and holding it up against their own experience. Courses that encourage this process in the design of class activities and homework exercises yield students that perform much better on conceptual tests than those from more traditional classes.

It is certainly plausible that this applies to upper-level courses as well. However, upper-level courses that embrace this philosophy are even rarer than introductory courses that do. Yet the kind of boost promised by active-engagement methods is exactly what we need in a general relativity class, where the concepts are challenging and the mathematics is difficult. Students need not only to *see* the results, they need in a very real sense to *own* them.

A second issue that provides part of the motivation behind the "physics first" movement is the recognition that the math in traditional general relativity texts (and, by extension, in lectures) tends to overpower the physics. Students do have a hard time seeing the flow of the physical argument in the midst of all the derivations and calculations.

I claim that one can address both of these problems by (1) redesigning the textbook and (2) rethinking the purpose of class time. The class notes I have developed over the years (which I am slowly shaping into a textbook) have gradually evolved to a format that emphasizes active learning. At present, each class day's notes begin with a four-page section that summarizes (without derivations) some important mathematical results and physical ideas. The student is then referred to a series of boxes that outline the derivations and other details, with exercises that push the students to fill in intervening steps and/or practice using the ideas (white space for student work is provided in each box). Students are expected to come to class having read the overview and having worked through the exercises in the boxes. We spend most of class time dealing with difficulties the students had with the exercises, answering questions that the overview and exercises have raised, working example problems, and (if there is time) discussing tangential issues of interest. Homework problems are more oriented toward applications than derivations.

The point of this organization is to ensure that students *have* seen the big picture clearly in the overview *and at the same time* have made the derivations their own by working through them personally before even coming to class. Class time is then effectively targeted at the specific difficulties that they are having, difficulties that students have become personally engaged with as a result of their preparation.

I think that this specific approach would be effective in class sizes up to about 15 students, and I can imagine adaptations that would make it work in larger classes as well. For example, one might begin the class by having students coalesce into groups of 5 or so and work for 10 minutes or so to define the difficulties with the reading and exercises that are common to all of them (in the process, helping each other with the simpler problems). The groups could then present their questions to the entire class, and the instructor could address them. I think that this 10-minute pre-processing of questions would save the instructor more than 10 minutes by helping focus the class more effectively on the most important issues. In a large class, one might also want to break up the session at least once in the middle by posing a problem or question that students can address in their groups. This would help keep students from settling back into passive listening.

6 Grade homework wisely

We all know from experience that we really *learn* physics only after doing appropriate homework. However, homework can be a daunting experience in a GR class: the problems are typically quite challenging, and there are many ways that a solution can go wrong. But homework assignments can be structured in a way so that students can face very difficult problems without anxiety, don't spin their wheels uselessly on a problem, and learn from their mistakes. The key is in how the homework is *evaluated*.

Here is an outline of the homework scheme that I use in all my upper-level courses. I grade each student's initial effort on each problem using the following 4-point grading rubric:

	Initial Effort
4	Satisfactory initial effort
3	Missing explanations or steps
2	Major problem parts missing (or didn't finish)
1	Very little coherent effort
0	No initial effort

Note that these initial-effort points have nothing to do with whether that effort is correct: students can earn a full 4 points on this part and be completely wrong. I make no comments at this stage except in the rare case where the student has made an error that I don't think that they will be able to figure out when they look at my solution. I then write something like `4/ /` on each of the students' initial efforts and return the problem sets to the students' department mailboxes with copies of my solutions. (I will talk about the two empty spaces after the 4 in a moment.) It takes literally a few seconds per problem to do this, so I can usually return the initial efforts the same day.

Students then use the printed solutions to correct their work using a different color ink. When they return their corrected homework, I then assess both their original work and the corrections on each problem, grading them according to the following rubrics:

	Correction Quality		**Correction needed**
3	Solution is now completely correct	3	No correction was necessary
2	Minor issues were not corrected	2	Minor corrections were needed
1	Major issues were not corrected	1	Important corrections were needed
0	No correction effort	0	Initial effort needed a complete rewrite

filling these scores into the two blank spaces after the initial-effort score. So, a student who does a problem completely right the first time earns 10 total points; a student who makes minor errors but finds and corrects them earns 9 points, a completely clueless student who nonetheless makes a good effort and fully corrects his or her work can earn up to 7 points, and even a person who submitted no initial effort on the due date but submits a full correction (i.e. a hand-written copy of the printed solution) earns 3 points.

This homework scheme has the following benefits:

1. It strongly rewards effort and care, and does not come down too hard on students that don't solve the problem correctly the first time. This encourages students to face difficult problems without stress and discourages cheating.

2. At the same time, students don't spend hours and hours spinning their wheels. After a reasonable amount of time, they will give up, knowing that they can earn points on the correction.

3. It motivates students to study the printed solutions carefully.

4. It encourages students to become self-critical about their work. Students can learn a lot about the subject and about good problem-solving style by comparing their work to solutions written by an expert.

5. It makes grading easy. Comments are almost never necessary, because students evaluate and correct their own work. Even though the instructor looks at each paper twice, grading goes very rapidly.

While I have found this scheme very useful in smaller classes, in larger classes one might modify the scheme so that a complete initial grade could be given initially, making the second pass optional. One could award something like 4 points for effort, 4 points for correctness, and 2 points for the quality of the initial effort. Students could then optionally earn back "correctness" points by submitting a correction. The disadvantage of such a scheme would be that students would not study the solutions for problems they got right and would not have to assess for themselves the quality of their work. On the other hand, it would greatly reduce the number of submitted corrections. It also avoids the problem of having to harangue the students to turn in corrections: a complete score is on record for each problem even if they only hand in the initial effort.

However one does it, the point is to create a homework scheme that rewards effort, prompts students to think about their work, gives them a chance to study expert solutions, and doesn't make them too anxious. This is valuable in every upper-level class, but perhaps especially so in a GR class.

7 Use tricks and tools to avoid tedious calculations

Students need *some* practice calculating Christoffel symbols and components of the Riemann and Ricci tensors to understand how they work, but once they have done it a few times (preferably in two dimensions), there is not much more to be learned. Solving the Einstein equation can be made much more palatable for students if one uses some powerful and general tools to make such calculations simpler.

If students have access to *Mathematica* or *Maple* at your campus, one can easily set up worksheets that can calculate such components (see Hartle's website[6] for some examples). However, I have for several years used a worksheet (see the appendix) that is almost as good and costs only a few pennies-worth of xeroxing to use. I got the basic idea from Rindler[7], but converted his sketch of a method to a full-blown worksheet. Students can use this worksheet to evaluate Christoffel symbols and Ricci tensor components quickly and accurately for *any* diagonal metric in two, three, or four dimensions by writing down what A, B, C, and D are for the metric in question, writing above each term in the worksheet what that term becomes for that metric, and collecting the nonzero terms.

One can also use computer programs to generate numerical solutions, particularly for displaying particle orbits, graphs, embedding diagrams, and so on. I have written a few such programs (and I understand that Wolfgang Christian of Davidson College is in the process of developing a number of such programs.) Such programs can be terrific tools for helping students visualize the physical implications of the equations.

Such tools are most valuable when students see them as extending abilities that they already have rather than being black boxes. When I use the worksheet, I ask students to check that the worksheet correctly gives the Ricci tensor components for an arbitrary two-dimensional diagonal metric so that they see how it works and begin to trust it. When I use a computer program, I also first have them check that the program reproduces simple cases that we have calculated by hand.

I have also found it useful to re-express the Einstein equation in terms of the Ricci tensor instead of the Einstein tensor[8]:

$$R_{\mu\nu} = 8\pi G(T_{\mu\nu} - \tfrac{1}{2}g_{\mu\nu}T) + \Lambda g_{\mu\nu} \tag{1}$$

The Ricci tensor is simpler than the Einstein tensor to evaluate, and the worksheet makes calculating its components straightforward. Since R_{tt} reduces in the newtonian limit to the laplacian of the newtonian gravitational potential, the Einstein equation expressed this way also has a simple physical interpretation (for a perfect-fluid source at least) that is outlined very nicely in Chapter 19 of Schutz's recent popular book on gravity[9]. (Raising the first index also removes the metric terms from the right side of the equation.) I was greatly surprised at just how much easier calculations became and how much more insight students could get from this form of the equation (particularly in the context of cosmology).

Here is another trick that people might not know. When I first develop the geodesic equation in my course, I use a Lagrangian method to maximize the proper time along the particle's worldline. This nicely connects with methods that students have learned in mechanics, and also reinforces the idea that a geodesic is the curve of extreme "distance" between two points. Midway through the course (after a discussion of the absolute gradient), we also see that a geodesic is the straightest possible line in a given space or spacetime, and give the geodesic equation in terms of Christoffel symbols. Now, comparing the two equations provides a very rapid method of calculating Christoffel symbols for any metric. If one uses the Lagrangian method to evaluate the four components of the geodesic equation, and then compares the result with the geodesic equation expressed in terms of Christoffel symbols, one can immediately read off the Christoffel symbols.

For example, consider the time component of the geodesic equation in Schwarzchild spacetime. The Lagrangian form of the geodesic equation implies that

$$0 = \frac{d}{d\tau}\left(g_{\mu\nu}\frac{dx^\nu}{d\tau}\right) - \frac{1}{2}\frac{\partial g_{\alpha\beta}}{\partial x^\mu}\frac{dx^\alpha}{d\tau}\frac{dx^\beta}{d\tau} \tag{2}$$

If we set $\mu = t$, and recognize that the metric does not depend on t, this becomes

$$0 = \frac{d}{d\tau}\left(g_{tt}\frac{dt}{d\tau}\right) - 0 = \frac{\partial g_{tt}}{\partial r}\frac{dr}{d\tau}\frac{dt}{d\tau} + g_{tt}\frac{d^2t}{d\tau^2} \;\Rightarrow\; 0 = \frac{d^2t}{d\tau^2} + \frac{2GM/r^2}{1-2GM/r}\frac{dr}{d\tau}\frac{dt}{d\tau} \tag{3}$$

after setting $g_{tt} = (1 - 2GM/r)$ and dividing through by g_{tt}. If we compare this to the geodesic equation in the form

$$0 = \frac{d^2x^\mu}{d\tau^2} + \mathit{\Gamma}^t{}_{\alpha\beta}\frac{dx^\alpha}{d\tau}\frac{dx^\beta}{d\tau} \tag{4}$$

we see immediately that $\mathit{\Gamma}^t{}_{rt} = \mathit{\Gamma}^t{}_{tr} = (GM/r^2)(1-2GM/r)^{-1}$ and that all other $\mathit{\Gamma}^t{}_{\alpha\beta} = 0$. Therefore we have calculated 16 Christoffel systems by taking what amounts to a single derivative!

Conclusion

This article has outlined just a few of the things that I have found helpful in teaching a tensor-based GR course over the years. There are also many things that one can do at the detailed level to make certain concepts more transparent and certain derivations easier that I cannot discuss in this short article. Even so, I hope that this article can be helpful to those professors who want to teach a successful undergraduate course in GR that does not abandon tensors.

I have a partial draft of a GR text that embodies the principles described in this paper which I will bring to the conference. It is far from being a complete and polished text, and I have no timetable for its completion, but it does illustrate one concrete approach to teaching an undergraduate GR course with tensors. I used the draft in my GR course this spring and was satisfied that the level and pacing was approximately correct.

The bottom line is that while I think that the mathematics of general relativity is challenging for undergraduates, one *can* successfully teach it if one pays appropriate attention to techniques that simultaneously make the mathematics easier and the pedagogy more efficient and effective. Moreover, I think that there are significant pedagogical advantages to students for doing so: they gain a deeper and more coherent understanding of the material and a valuable sense of empowerment when they are able to master the tensor calculus and solve the Einstein equation on their own power. We need not despair of reaching this goal!

---

[1] Schutz, Bernard F., *An First Course in General Relativity*, Cambridge, 1985.
[2] Ohanian, Hans, and Ruffini, Remo, *Gravitation and Spacetime*, 2nd edition, Norton, 1994.
[3] Hartle, James, *Gravity*, Addison-Wesley, 2003.
[4] Cheng, Ta-Pei, *Relativity, Gravitation and Cosmology*, Oxford, 2005.
[5] For example, see Hake, Richard, "Active engagement vs. traditional methods: A six thousand student study of mechanics test data for introductory physics courses," *Am. J. Phys.* **66** (1), 64-74 (1998).
[6] `http://wps.aw.com/aw_hartle_gravity_1`
[7] Rindler, Wolfgang, *Essential Relativity*, Springer, 1977
[8] I do this only after deriving and fully discussing the Einstein tensor, of course.
[9] Schutz, Bernard F., *Gravity from the ground up*, Cambridge, 2003.