

A PACKING LIST

Before Bob Colwell sets out on his missionary work to explain engineering to the humanities community (At Random, “Frames of Reference,” June 2005, pp. 9-11), he should include the following in his packing list:

- a book on reasoning,
- a history of the 20th century, and
- a subscription to the *Journal of the History of Ideas*.

The New Age mystic and the school board official that Colwell cites as representatives of the humanities are as abhorrent to those of us in the “difficult” sciences as they are to our colleagues in the “hard” sciences.

Colwell appeals to the common prejudice of technologists that the humanities have no standards, hence, anything goes. That appeal is furthered by an inadequate sample that would not pass muster in any technology discussion.

Is scientific illiteracy a problem? Undoubtedly, but then so is humanities illiteracy. What else would explain representatives of the hard sciences writing claptrap like the Bible Code books? It is not good science, much less good humanities work.

But it would be unfair, not to mention inaccurate, to characterize all technologists on the basis of a few carefully chosen examples.

Colwell’s claim in his conclusion that people in the technology business must realize they are “in the service of truth, an absolute truth that must be guarded” is contrary to the history of the 20th century. I will skip the usual list of racial and social outrages perpetrated in the name of science and service of “absolute” truth. No doubt Colwell would reply that those were betrayals of the scientific process, and I would readily agree.

The problem is that when any group—technologists, humanists, or any groups ignorant of one or both—decides that it has a lock on absolute



or even provisional truth, bad consequences follow.

The information overload issues currently reported with such breathless concern in the major news outlets and more than a few scientific circles have been encountered before. For example, the January 2003 issue of the *Journal of the History of Ideas* was entirely devoted to the topic of “Early Modern Information Overload.” Realize that “early modern” in this context refers to the time period 1550-1750.

While the solutions developed at the time did not rely upon the transistor and its now distant cousins, the underlying principles were not all that different from modern information systems. The point being that if we avoid the obviously outlandish on either side, we can learn a great deal by studying both the hard and the difficult sciences.

Colwell starts from the position that “I’m right and you have been too lazy to learn why, so sit down and listen,” which will not excite a lot of interest in his target audience. A better strategy would be to try to understand the other frame of reference before deciding that it is all that different from your own.

Will this strategy mean reading unfamiliar material and having to master new terminology? Yes, but it is also potentially rewarding in terms of new insights into your own discipline.

If the arts/science gap is going to be bridged, it will be by participating in both communities, not by missionaries bringing enlightenment to the savages.
Patrick Durusau
Covington, Ga.
patrick@durusau.net

Bob Colwell responds:

In my June column on “frames of reference,” an idea that it is very difficult to communicate across different

frames, I juxtaposed science and religion and science and the arts as pairs of representative differing frames. To find further evidence for the difficulties of communicating across frames, it isn’t necessary to look any further than Patrick Durusau’s letter.

Within the scientific frame of reference, most of its practitioners (including me) presume that an absolute truth exists. Scientists spend their lives looking for it; engineers spend their lives trying to stay close enough to it that their designs will work as intended. Durusau’s computer works because these engineers succeeded.

The electrical charge on an electron is the same no matter who measures it, no matter where, and no matter what attitude or beliefs the measurer may have. Most of Durusau’s letter seems aimed at the idea that science and the humanities are duals. In many ways they are, and thinking about those similarities may be enlightening to both camps.

But there is a basic way in which they differ, and Durusau has inadvertently highlighted it: Those in the sciences have a different standard for what they consider truth. When people in the humanities say that science has no monopoly on truth, I agree. But in an important sense, we’re talking past each other. We don’t mean the same thing; we’re just using the same word.

Both sides ought to consider that each is using it within its own frame. That was precisely why I bridled at the mystic in the movie—she was using what she meant by “truth” as if it meant the same thing as when a scientist uses it, thus propagating an insidious logical error.

Durusau missed one other major point of the article: It was aimed precisely at *Computer’s* readership, and not at anyone from any other frame of reference. I therefore intentionally used the terms and context that are appropriate to that frame. Were I to have written a similar article for a humanities journal, substantial translation would have been required so as to reach that readership

without giving offense. That was the point of the article.

CAMP A INSIGHT

Regarding Bob Colwell's June column, I think Camp A is actually all right with Camp S. We need a different label than "art" for the other side.

I say this from my experience teaching a new course at the Maryland Institute College of Art titled "Foundations of a Scientific World View." It's a very selective school—these are top-notch kids with great artistic gifts and motivations, and they're no slouches when it comes to traditional academics either.

I also teach a course for graduate engineering students at Johns Hopkins titled "Logical Foundations of Computer Science." The art students have an unfair advantage, being full-time undergraduates versus the part-time graduate engineering students, but I'd have to say the art students are more inquisitive and open to a range of ideas compared to the engineering students. And the art students are very rigorous, too. The art faculty is also very inquisitive, open, and rigorous.

Maybe the opposite camp from Camp S is Camp B, for believers. Camp A is chock full of questioners. They seem to love a radical challenge to their preconceived notions.

*Richard (DJ) L. Waddell Jr.
Laurel, Md.
richard.waddell@jhuapl.edu*

Bob Colwell responds:

I can readily testify that most artists I've known have exhibited as much discipline and rigor as most engineers. But you know how stereotypes go: Artists are creative; engineers are drudges. I don't buy that one, either. Engineers are among the smartest people I've ever met (including all my science friends who are also very bright), and smart people are easily bored. I think people operating at the limits of human intellect and feasibility will always be creative, artist or engineer, and probably to the same extent.

In my June column, I was primarily trying to find different frames of reference that the readers—who are mostly in the scientific camp—would readily understand, without inadvertently extending any stereotypes that should not be propagated. That's why I didn't address creativity, intelligence, inquisitiveness, and so on. I'm not trying to compare and contrast, per se; I'm only trying to show that the idea that differing frames cause differing interpretations of ideas and words is a useful meme.

My motivation is that within the next two to three decades, the probability is quite high that the world will face some fundamental challenges, of types that have never been seen before. Global warming, caused by the wholesale dumping of carbon into the atmosphere, could cause massive, unpredictable changes to the climate worldwide.

To what extent should we make sacrifices today on behalf of future generations, and based on what data and climate models? When the stakes are this high, we had better get this right.

Avoiding such catastrophes will require a reasonably deep understanding of many technical issues. That people not already embedded in the scientific realms must be bright, enthusiastic, and open to ideas is necessary but not sufficient.

The science folks who have studied this issue must find ways to get the information across to their Not Camp-S compatriots, or making communal intelligent decisions will be impossible. That is what I meant when I said the Camp-S folks must bridge the gap.



ARITHMETIC OPTIONS

While my interest in reading "An Open Question to Developers of Numerical

Software" by William Kahan and Dan Zuras (Standards, May 2005, pp. 91-94) kept up till the end, my perturbation grew from nothing to very great.

The authors state that the IEEE 754 Revision Committee's work is directed "into areas of expression evaluation, compiler optimizations, and other language issues that were considered outside the scope of the 1985 standard." Although the article opens with a general description, it is evident that the authors pose the open question because the committee members "wish to get rid of traps and of sNaNs if [they] can." The concluding remarks confirm that this is the primary concern, as does the e-mail address given for sending comments: snans@nonabelian.com.

The authors mention two excellent objectives: "Revisions of substance ... that would most help make modern floating-point programming easier and more reliable, efficient, and portable" and not to "make any change so sweeping it invalidates the operation of most of the world's computers."

What the authors seem to overlook is compatible extension. Defining standards for arithmetic on double (not double-length) numbers as well would best meet their worthy objectives.

There could be four modes of use: extended precision, combined operations on number pairs, operations on intervals, and operations on complex numbers—the last two being of most benefit in computation.

A method of interrogating whether the double arithmetics are available would make transition to the extended standard much easier.

Although the more popular programming languages do not as yet have types available for such extensions, the arithmetic could come into prompt and effective use in standard library code.

The availability of full support for interval arithmetic would benefit technical computation and promote wider use of the proven validated numerics techniques. The availability of full support for complex arithmetic would greatly improve both technical com-

putation and graphical computation, which is of such commercial significance and magnitude that chipmakers might well be persuaded to adopt a new extended standard much earlier than if it merely benefited engineers.

I am not a numerical software developer, and many such developers would be better able to judge the benefits of this kind of extension. If I am wrong, so be it. But if I am right, and if the benefits are as great as I imagine, this is an opportunity for experts to offer the Revision Committee their opinions and their help. If the committee is to extend the standard in this way, they will need the best help they can get.

Neville Holmes

University of Tasmania

Neville.Holmes@utas.edu.au

The authors respond:

Neville Holmes is quite right in stating that the primary goal of our article was to get feedback on the narrow issue of signaling NaNs.

The IEEE 754 Revision Committee has considered many of the extensions Neville mentions, along with others. However, doubled-double generally is an unreliable arithmetic to use for computation.

Suppose a C++ program using float and double variables and constants has been proved to work well, but greater accuracy is desired. One way to achieve that is to promote all double variables and constants to doubled-double (most likely called “long double”), promote all float variables and constants to double, appropriately shrink tolerances controlling iteration convergence, add terms to truncated series, and so on.

Would the promoted program work as well as the old, but twice as accurately? Not necessarily. The promoted program could suffer from mysterious rare malfunctions none of which could afflict the code if quadruple had been used instead of doubled-double.

This is why the committee has chosen to standardize the honest 128-bit quad format already in common use. At least one chip designer has

announced full hardware support for quadruple at full speed.

The IEEE 754r committee has avoided attempting to specify complex arithmetic for three reasons.

First, C99 does a good job. It specifies an imaginary type and then defines the complex type as a formal (unevaluated) sum of a real and an imaginary. Doing so avoids several nasty anomalies in the complex type as Fortran and some packages used in C and C++ specify.

Second, as good as the C99 specification may be, it still suffers from some conundrums. For instance, multiplication that handles complex infinities well is too slow for matrix multiplication.

Correctly rounded complex multiply is probably too much to ask even from systems with a fused multiply-add but not quadruple in hardware. Correctly rounded complex divide still costs too much.

Third, nothing IEEE 754r can say will change entrenched complex arithmetic practices, but attempts to change them might well jeopardize acceptance of the rest of IEEE 754r. Maybe some day.

We have looked into interval arithmetic and a standard set of transcendental functions. Our reluctant conclusion is that the state of the art is not yet up to standardizing these functions in a way that our posterity would surely find satisfactory.

Some things need to be standardized now:

- control of expression evaluation,
- alternate methods of exception handling, and
- aids to debugging.

Our work is not yet finished. The committee’s deliberations are open to anyone who wants to participate (grouper.ieee.org/groups/754/), including those who are not members of the IEEE.

For details about these and other issues see www.cs.berkeley.edu/~wkahan.

William Kahan and Dan Zuras



COMPUTATION AND COMMUNICATION

In “Socially Aware Computation and Communication” (Mar. 2005, pp. 33-40), Alex Pentland discusses some interesting cutting-edge ideas for building machines to understand human behaviors in terms of “quantifying social context in human communication.”

In addition to being the underpinnings of mental partnerships, nonlinguistic signals such as body language, facial expression, and tone of voice are what make humans the most complicated animal. However, how well an expert system could understand and interpret such signals would depend on its capacity to make inferences or draw conclusions. It is not a question of how to build such an expert system, but rather whether the system will work.

In general, an expert system consists of two important components: a knowledge base and reasoning. This base consists of factual knowledge and heuristic knowledge.

Factual knowledge is a vast collection of widely shared and commonly used information. How people use this collection of knowledge depends on multiple factors such as their age, education, culture, gender, and so on. In addition, the same behavior could have a different meaning in a different situation or a different culture.

In some cases, knowledge is incomplete or uncertain. While it’s possible to apply the concept of “fuzzy logic,” assigning an associated confidence factor or weight, these scenarios might excite a human’s imagination but probably not a machine’s.

Heuristic knowledge is useful in testing the efficacy of an intelligent system because it requires making judgments, but this could cause plausible reason-

ing to be ineffective. For example, if people do a good job of pretending, how would forming a line of reasoning using either forward chaining or backward chaining of IF-THEN rules handle interpreting their behavior?

That is what makes simulating the human brain extremely difficult, and I believe no machine can ever replace it.

As a side note, if researchers could build a machine for making important decisions like finding a mate, getting a job, negotiating a salary, or finding a place in a social network, would it become a hindrance to human thinking?

Although we don't know when such a machine might become available, and it seems unlikely that it would be

effective in dealing with the real-life situations, I fully support the spirit of this wonderful pursuit of the concept of building intelligent systems.

Hong-Lok Li
Vancouver, BC
libl@ams.ubc.ca

EMBEDDED ENTERTAINMENT

Readers of my article on smart projectors published in *Computer's* January 2005 issue ("Embedded Entertainment with Smart Projectors," pp. 48-55), might be interested in finding information about a related project. The authors of "Making One Object Look Like Another: Controlling Appearance Using a Projector-Camera System"

(M.D. Grossberg et al., *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, IEEE CS Press, 2004, pp. 452-459) describe a projector-camera system that addresses the problem of radiometric compensation on 3D surfaces with a different technical approach and application orientation.

Oliver Bimber
Bauhaus-University Weimar
bimber@uni-weimar.de

We welcome your letters. Send them to computer@computer.org.



28th International Conference on Software Engineering

20-28 May 2006

<http://www.icse-conferences.org/2006/>

ICSE 2006 will be held in Shanghai, China, a main epicenter of the explosive growth of the software industry in China during these early years of the 21st century. ICSE 2006 will be a signal event in the recognition of this growth, and in the bringing of the Chinese software engineering community into the ICSE mainstream. The web site will detail traveling arrangements as well as the many ICSE activities to be held in Shanghai, a city that blends the dazzlingly ultramodern and the charmingly traditional. A visa is needed for travel—please visit the web site for more information and other conference details. Travel costs less than you think, so join us in Shanghai in May 2006 for an experience that will combine outstanding technical events with a visit to a city and country that will be truly unforgettable!



Major Submission Deadlines

Research Track	9 Sep 2005
Education Track	1 Oct 2005
Workshop and Tutorial Proposals	6 Oct 2005
Experience, Far East Experience, Research Demos, Emerging Results	30 Oct 2005
Doctoral Symposium	5 Dec 2005

