

Sampling Binary Contingency Tables

The binary contingency tables problem has sparked the interest of statisticians and computer scientists alike. Although it's been studied from both theoretical and practical perspectives, a truly usable algorithm remains elusive. The author presents several approaches to the problem.

The traditional motivation for the binary contingency tables problem comes from Charles Darwin. In the early 1830s, Darwin joined a geological expedition to the Galápagos archipelago, where he collected samples and made notes about each of the 17 islands. Back in England, ornithologist John Gould classified the samples Darwin considered to be different bird species as different types of the same bird: a finch. Darwin noticed that not every type of finch lived on every island, so he sketched a so-called *occurrence table* (see Figure 1).¹ In later years, Darwin noted correlations between the types of finches inhabiting the individual islands along with the living conditions on those islands. This observation is often considered to be one of the main sparks that led to his theory of evolution.²

In many modern disciplines, we can arrange data into occurrence tables as well, but is the observed occurrence table the result of pure chance, or does it significantly differ from a random table—in Darwin's case, for example, is the Galá-

pagos occurrence table purely random or a result of evolution? To answer this sort of question, we need to specify which tables constitute the universe for randomness and then set up a test statistics that measures a property of interest on the data. If the value of the statistics for the observed table differs significantly from the expected value across all tables, then the observed table isn't considered random.

Scientists often view the marginal sums of an occurrence table as fixed. For example, the marginal sum in the first row (that is, the sum of all entries in this row) of Darwin's finches table is 14, which tells us how widespread the corresponding type of finch is (in this case, the large ground finch); the marginal sum in the third column is 11, which captures how accommodating the island Isabella is. Now let's rephrase the question: is the observed table just a pure random draw from the set of all occurrence tables with the same marginal sums?

Alan Roberts and Lewis Stone³ designed a possible test statistics that measures competitive pressures between species: for an occurrence table $T = (t_{i,j})$, $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$, where i corresponds to species and j to the environment, they propose computing the function

$$f(T) = \frac{1}{m(m-1)} \sum_{i \neq j} s_{i,j}^2,$$

where

$$s_{i,j} = \sum_{k=1}^n t_{i,k} t_{j,k}$$

If the value of f on the observed table is close to the expected value of f over all tables with the same marginal sums, then the observed table is viewed as random with respect to the test statistics f . (Note that this statistics is just an example—other researchers have designed different statistics that measure other properties of data arranged in occurrence tables.)

After we establish the test statistics, computing the value for the observed table is straightforward. The tricky part is estimating the *expected* value: we can do this via random sampling of all occurrence tables with the appropriate marginal sums.

Formal Statement of the Problem

In the binary contingency tables problem, the task is to generate a uniformly random $m \times n$ binary matrix T that satisfies the row sums r_1, \dots, r_m and the column sums c_1, \dots, c_n . (Every entry of a binary matrix is either zero or one.) Throughout the rest of this article, every matrix T that satisfies these conditions is called a *binary contingency table* (BCT); Ω denotes the sample space—that is, the set of all BCTs with the given row and column sums.

Sequential Importance Sampling

The main idea behind the importance sampling approach⁴ is to sample from a nonuniform (and easier to compute) distribution to estimate the sample space's size. When the size is known, we can then compute the expected value of any function by computing the expected value of a properly adjusted function according to the importance sampling distribution.

More precisely, if σ is a distribution on Ω , we can define a random variable ϕ , which, for an element $\omega \in \Omega$, returns value $1/\sigma(\omega)$. The expected value of ϕ is

$$E[\phi] = \sum_{\omega \in \Omega} \sigma(\omega) \frac{1}{\sigma(\omega)} = |\Omega|,$$

the size of the sample space. Thus, if we keep sampling $\omega_1, \omega_2, \dots$ using the distribution σ , the expression

$$\frac{1}{n} \cdot \sum_{i=1}^n \phi(\omega_i)$$

converges to $E[\phi]$ as n increases. Once we have an estimate for $|\Omega|$, we can then compute $E[f(\omega)]$ for any function f on Ω by computing the expected value of

	Seymour	Baltra	Isabella	Fernandina	Santiago	Rábida	Pinzón	Santa Cruz	Santa Fe	San Cristóbal	Española	Floreana	Genovesa	Marchena	Pinta	Darwin	Wolf
Large ground finch	0	0	1	1	1	1	1	1	1	1	0	1	1	1	1	1	1
Medium ground finch	1	1	1	1	1	1	1	1	1	1	0	1	0	1	1	0	0
Small ground finch	1	1	1	1	1	1	1	1	1	1	1	1	0	1	1	0	0
Sharp-beaked ground finch	0	0	1	1	1	0	0	1	0	1	0	1	1	0	1	1	1
Cactus ground finch	1	1	1	0	1	1	1	1	1	1	0	1	0	1	1	0	0
Large cactus ground finch	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0
Large tree finch	0	0	1	1	1	1	1	1	1	0	0	1	0	1	1	0	0
Medium tree finch	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0
Small tree finch	0	0	1	1	1	1	1	1	1	1	0	1	0	0	1	0	0
Vegetarian finch	0	0	1	1	1	1	1	1	1	1	0	1	0	1	1	0	0
Woodpecker finch	0	0	1	1	1	0	1	1	0	1	0	0	0	0	0	0	0
Mangrove finch	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
Warbler finch	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Figure 1. Occurrence table for Charles Darwin's finches. Rows correspond to the types of finches and columns to the Galápagos islands they inhabit.

$$\frac{f(\omega)}{\sigma(\omega)|\Omega|}$$

according to the distribution σ . In the importance sampling approach, not only do we need to be able to sample from σ and sample every element $\omega \in \Omega$ with nonzero probability, we must also be able to compute the probability $\sigma(\omega)$.

Yuguo Chen and his colleagues¹ first proposed using the importance sampling approach for the BCT problem. Their importance sampling distribution generates a random table in a column-by-column manner, thus they used a sequentially constructed distribution, which is an instance of a so-called *sequential importance sampling* (SIS). The main idea behind their distribution is that if we satisfy all the row sums but only one column sum, the problem becomes significantly easier. A random binary $m \times n$ matrix that satisfies row sums r_1, \dots, r_m and column sum c_1 contains in its first column vector $\mathbf{x} = (x_1, \dots, x_m) \in \{0, 1\}^m$, where $\sum_{i=1}^m x_i = c_1$ with probability proportional to

$$g(\mathbf{x}) := \prod_{i=1}^m \left(\frac{r_i}{n - r_i} \right)^{x_i}.$$

In other words, the probability that the first column is \mathbf{x} is equal to $g(\mathbf{x})/Z$, where the normalizing factor Z equals the sum of $g(\mathbf{y})$ ranging over all vectors $\mathbf{y} = (y_1, \dots, y_m) \in \{0, 1\}^m$ with exactly c_1 ones.

Chen and his colleagues describe their SIS distribution with the following algorithm: first,

choose the first column from the distribution defined by g (even though the normalizing factor Z is a sum ranging over a possibly exponential number of terms, we can sample proportionally to g via a dynamic programming algorithm in polynomial time). Then, once the values in the first column have been chosen, update the row sums so that they represent the required row sums of columns 2 through n (in other words, if $x_i = 0$, row sum r_i stays the same; if $x_i = 1$, then r_i is decreased by one). Choose the second column analogously to the first column. Continue through all columns. (Note that this procedure might end up deadlocked—for example, we might be left with a single column and row sums with numbers higher than one. Chen and his colleagues also propose a more elaborate SIS distribution that gets around this problem, but for the rest of this discussion, the simpler SIS distribution is sufficient.)

We can see that ϕ is an unbiased estimator for the sample space's size (that is, $E[\phi] = |\Omega|$), but it isn't clear how many samples we need for an accurate estimate of $E[\phi]$. This SIS distribution works very well for many practical applications; Jose Blanchet's work recently supported this empirical observation.⁵ He showed that the method converges quickly if the marginal sums are significantly smaller than $\sqrt[4]{S}$, where S is the sum of all marginal sums.

Unfortunately, there are row and column sums for which, with high probability, this SIS distribution requires an exponential number of samples. To give a feeling for what can go wrong, let's look at a simple example of row and column sequences for which the approach fails. Consider the row sums $1, 1, \dots, 1$, with the sum $\lceil \alpha m \rceil$ for some $\alpha > 0$ in the last row and column sums $1, 1, \dots, 1$. (Notice that $n = m - 1 + \lceil \alpha m \rceil$ so that the sum of the row sums equals the sum of the column sums.) This example's simplicity allows us to compare the SIS distribution with the uniform distribution. Let's look more closely at the first column. The SIS distribution (see the definition of the function g) places a one (the single one in this column) in row j with probability proportional to

$$\frac{\lceil \alpha m \rceil}{(m-1 + \lceil \alpha m \rceil) - \lceil \alpha m \rceil} \approx \alpha \quad \text{for } j = m,$$

and with probability proportional to

$$\frac{1}{(m-1 + \lceil \alpha m \rceil) - 1} \approx \frac{1}{(1 + \alpha)m} \quad \text{for } j < m.$$

Therefore,

$$Z \approx \alpha + \frac{1}{1 + \alpha},$$

and the SIS distribution places a one in the last row with probability

$$\approx \frac{\alpha}{Z} \approx \frac{\alpha(1 + \alpha)}{\alpha(1 + \alpha) + 1}.$$

Now let's look at the uniform distribution over all BCTs. It's easy to compute the number of all BCTs for this input: we have

$$\binom{n}{\lceil \alpha m \rceil}$$

choices for the last row and $(m - 1)!$ possibilities for the rest of the table, yielding a total of

$$x := \binom{n}{\lceil \alpha m \rceil} (m - 1)!$$

tables. The same reasoning gives

$$x_1 := \binom{n-1}{\lceil \alpha m \rceil - 1} (m - 1)!$$

tables with a one in the last row of the first column. Therefore, a one should be placed in the last row of the first column with probability $x_1/x \approx \alpha/(1 + \alpha)$. (Notice the difference between this probability and the one arising from the SIS distribution.)

Now let's consider the first βn columns, where β is an appropriately chosen positive constant. It turns out that a similar difference occurs in each of these βn columns. Because all the columns have the same column sum, by symmetry, we expect that approximately a β -fraction of all ones in the last row occurs in these βn columns. However, because of the difference in the probabilities, we'll see significantly fewer ones in the tables sampled by SIS. A concentration bound (namely, the Chernoff bound) implies that the tables with few ones in the first βn columns of the last row form only an exponentially small fraction of all possible BCTs. However, because these are likely to be the only tables sampled by SIS, the technique will estimate only the size of this exponentially smaller set of tables.⁶

Markov Chain Monte Carlo

Markov chain Monte Carlo is a popular sampling technique.⁷ The main idea is to start with an element from the sample space Ω and modify it using a sequence of random (typically local) changes. If we choose the changes from a properly selected set of possibilities and if the sequence of changes is sufficiently long, the resulting element will be an (almost) random sample from Ω . To illustrate the

idea, let's look at a simplified version of card shuffling: we start with a pile of n cards, then choose two cards at random and swap them. If we repeat this move at least $O(n \log n)$ times, the cards will end up in a completely random order.

In mathematical terminology, a Markov chain on a sample space $\Omega = \{\omega_1, \dots, \omega_{|\Omega|}\}$ is given by a transition matrix P of size $|\Omega| \times |\Omega|$, where p_{ij} specifies the probability with which we move from element ω_i to element ω_j . A distribution π on Ω for which $\pi = \pi P$ is called *stationary* (if we choose the initial element from Ω according to π , the next element will also be distributed according to π). If it's possible to get from every element to every other element by a sequence of steps, and if $p_{i,i} > 0$ for every i , then the stationary distribution is unique. Thus, no matter which element we start with, we'll eventually (in the limit) reach the stationary distribution. Finally, if we want to sample from a stationary distribution π , it's sufficient to set up the transition matrix so that it satisfies the so-called detailed balance condition: $\pi(\omega_i)p_{i,j} = \pi(\omega_j)p_{j,i}$ for every $\omega_i, \omega_j \in \Omega$.

Typically, the most challenging part is to decide how many steps of the Markov chain we need to get sufficiently close to the stationary distribution. In computer science, a Markov chain is said to be rapidly mixing if this number of steps is polylogarithmic in $|\Omega|$. For BCTs, this means we want a Markov chain that needs at most a polynomial number of steps in n and m .

Two different Markov chains for the BCT problem have been studied in literature. Persi Diaconis and Anil Gangolli⁸ proposed a Markov chain that, in every step, chooses two rows and two columns of the current table at random, and if the four entries specified by these rows and columns have exactly one 1 (and thus exactly one 0 as well) in each row and each column, the two zeros are replaced with ones and vice versa. Ravi Kannan, Prasad Tetali, and Santosh Vempala⁹ analyzed this Markov chain for the case when all row and column sums are equal and proved that the chain is rapidly mixing (albeit the number of steps is a polynomial of a rather large degree). Unfortunately, there are no known results for this Markov chain in the case when the row and column sums differ significantly (Kannan, Tetali, and Vempala's results can be adapted for marginals that differ only a little).

An alternative Markov chain samples not only from the set of all BCTs, but also from the set of all tables with exactly one row and exactly one column sum decreased by one. To simplify our language, let's call a BCT *perfect* if it has row sums

r_1, \dots, r_m and column sums c_1, \dots, c_n ; we'll call it *near-perfect with deficiencies* at i, j if it has row sums $r_1, \dots, r_{i-1}, r_i - 1, r_{i+1}, \dots, r_m$ and column sums $c_1, \dots, c_{j-1}, c_j - 1, c_{j+1}, \dots, c_n$. The Markov chain starts at a perfect table (it isn't difficult to find one), and the subsequent moves follow these rules:

1. If the current table T is perfect, choose an entry i, j with $t_{i,j} = 1$ at random and set $t_{i,j} = 0$ (creating a deficiency in row i and column j).
2. If the current table T is near-perfect with deficiencies at i, j , choose an entry k, ℓ with $t_{k,\ell} = 1$ or $(k, \ell) = (i, j)$ uniformly at random: if $(k, \ell) = (i, j)$, set $t_{i,j} = 1$ (and T becomes a perfect table); otherwise, flip a coin, and if it comes out heads, do this: if $t_{k,j} = 0$, set $t_{k,j} = 1$ and $t_{k,\ell} = 0$ (move the deficiency from column j to column ℓ); for tails, do this: if $t_{i,\ell} = 0$, set $t_{i,\ell} = 1$ and $t_{k,\ell} = 0$ (move the deficiency from row i to row k).

It follows from the detailed balance condition that the stationary distribution of this Markov chain is uniform over all perfect and near-perfect

Even if the Markov chain converged to the stationary distribution quickly, we would have an extremely small probability of actually sampling the one table we want to see for the marginals just listed.

tables. However, we want to sample only from the perfect tables, so how do we get rid of the near-perfect ones? Consider the marginals $1, 2, 3, \dots, n$ and $1, 2, 3, \dots, n$ for some n . It turns out that a unique BCT satisfies these marginals, but at least $2^{n/4}$ corresponding near-tables also exist. Thus, even if the Markov chain converged to the stationary distribution quickly and we could use it for sampling from all perfect and near-perfect tables, we would have an extremely small probability of actually sampling the one table we want to see for the marginals just listed. Before we sketch a solution to this problem, let's see how a similar issue was overcome in a slightly different setting.

A Different Sampling Problem

Let's start with an $n \times n$ binary matrix A . The goal is to sample from the set of all $n \times n$ binary tables T with all marginals equal to one and, moreover,

if $a_{ij} = 0$, then $t_{ij} = 0$, too. (This problem is known as sampling perfect matchings in a bipartite graph and is directly related to the problem of computing the permanent of a binary matrix.)

If we transfer our perfect and near-perfect table terminology to this setting, we notice that the Markov chain from earlier (with a small modification) gives rise to the uniform stationary distribution for the new problem. Similarly, there exist input matrices A for which the number of perfect tables is exponentially smaller than the number of near-perfect tables.

Through a series of iterations, Mark Jerrum, Alistair Sinclair, and Eric Vigoda¹⁰ came up with a breakthrough idea that tweaked the stationary distribution so that they could sample perfect tables with a probability of roughly $1/(n^2 + 1)$. The idea is to re-weight the distribution so that each of the following $n^2 + 1$ groups of tables is sampled equally likely: perfect tables (denoted with \mathcal{P}) and, for every $i, j \in \{1, \dots, n\}$, near-perfect tables with deficiencies at i and j (denoted with $\mathcal{N}_{i,j}$). In other words, with the uniform distribution, the chance of sampling a table from $\mathcal{N}_{i,j}$ is proportional to $|\mathcal{N}_{i,j}|$; with the re-weighted distribution, we want this chance to be equal to the chance of sampling a perfect table, so we need to sample an element from $\mathcal{N}_{i,j}$ with a probability proportional to $|\mathcal{P}|/|\mathcal{N}_{i,j}|$. The detailed balance condition lets us modify the Markov chain so that the stationary distribution reflects the re-weighted distribution, contingent upon knowing how to compute the ratios $|\mathcal{P}|/|\mathcal{N}_{i,j}|$ (this reweighting technique is known as the Metropolis-Hastings technique). Moreover, the re-weighted Markov chain can be proved to mix rapidly if we have a constant-factor estimate of each of the ratios.

So how do we estimate the ratios? Because computing $|\mathcal{P}|$ is about as hard as being able to sample from \mathcal{P} , we can't compute the ratios directly. However, if the matrix A contained only ones, we would know how to compute $|\mathcal{P}|/|\mathcal{N}_{i,j}| = n!/(n-1)! = n$. Let B represent the $n \times n$ matrix with all entries equal to one and let A be the target matrix for which we want to sample the (perfect) tables. Instead of sampling tables for A directly, we'll sample tables for the matrix $A' = B\lambda + A(1 - \lambda)$, where λ will gradually change from one to (almost) zero. (If you're familiar with the simulated annealing technique, this is an instance of the technique in which $1/\ln\lambda$ corresponds to the current temperature.)

We defined the problem for matrices A with entries zero or one, but the matrix A' might have real-valued entries in $[0, 1]$. In this case, we're interested in

sampling binary tables T with all marginals equal to one; the probability of the table T will be proportional to $\lambda(T) := \prod_{i,j} (a'_{i,j})^{t_{i,j}}$. Notice that $\lambda = 0$ corresponds to the original problem. Similarly, instead of taking the count of all perfect tables, we define the “ λ count” as $\lambda(\mathcal{P}) = \sum_{T \in \mathcal{P}} \lambda(T)$.

We start with $\lambda = 1$ when we know the ratios are equal to n , then we decrease λ slightly so that n is within a factor of 4 from the new ratio $\lambda(\mathcal{P})/\lambda(\mathcal{N}_{i,j})$. Now, we can use the Markov chain to improve our estimate of the new ratio from factor 4 to factor 2 simply by looking at the frequencies of sampling perfect tables versus near-perfect tables with deficiencies at i, j . This, in turn, allows us to decrease λ so that this factor 2 estimate becomes a factor 4 estimate for the decreased λ . We continue decreasing λ in this fashion until λ reaches (almost) zero, at which point we'll have a close estimate of the ratios $|\mathcal{P}|/|\mathcal{N}_{i,j}|$.

Can we use a similar idea for the BCT problem? Earlier, we gradually changed the all-one matrix B into our target matrix A (which restricts some entries of the tables to zeros). However, in the BCT problem, we don't have restrictions on the table entries, so it isn't clear whether we can make analogous gradual transformations of the input. Surprisingly, it turns out that we can. In particular, for any BCT's input, there exists a “restriction matrix” A that we can gradually change into B , our target input (no restrictions). To do this, we include a matrix A' with entries in $[0, 1]$ in the input of the BCT problem—our goal is to sample a BCT T with row sums r_1, \dots, r_m and column sums c_1, \dots, c_n with probability proportional to $\prod_{i,j} (a'_{i,j})^{t_{i,j}}$. As before, we consider matrices A' of the form $B\lambda + A(1 - \lambda)$; again, the tricky part is getting an initial estimate of the ratios $\lambda(\mathcal{P})/\lambda(\mathcal{N}_{i,j})$. Interestingly, for any marginals, we can construct a “restriction matrix” A for which we can estimate the ratios for λ close to zero. Then, we can gradually increase λ until $\lambda = 1$.¹¹

The SIS approach described earlier works well in many real-life situations but provably fails for some inputs. Unfortunately, we can't simply test whether the algorithm converges or not—for the example presented earlier, the algorithm appears to converge to a wrong value.

We've sketched an algorithm for BCTs that runs in a polynomial time for any input, but the polynomial is rather large (for an $m \times n$ input, the running time is upper-bounded by $O^*((nm)^5(n + m))$, where the O^* notation hides constant and

polylogarithmic factors). Thus, a truly usable algorithm for the BCT problem remains elusive. ❧

References

1. Y. Chen et al., "Sequential Monte Carlo Methods for Statistical Analysis of Tables," *J. Am. Statistical Assoc.*, vol. 100, no. 469, 2005, pp. 109–120.
2. F.D. Steinheimer, "Charles Darwin's Bird Collection and Ornithological Knowledge during the Voyage of H.M.S. Beagle, 1831–1836," *J. Ornithology*, vol. 145, no. 4, 2004, pp. 300–320.
3. A. Roberts and L. Stone, "Island-Sharing by Archipelago Species," *Oecologia*, vol. 83, no. 4, 1990, pp. 560–567.
4. I. Beichl, and F. Sullivan, "The Other Monte Carlo Method," *Computing in Science & Eng.*, vol. 8, no. 2, 2006, pp. 42–47.
5. J. Blanchet, "Efficient Importance Sampling for Binary Contingency Tables," to be published in the *Annals of Applied Probability*, 2008.
6. I. Bezáková et al., "Negative Examples for Sequential Importance Sampling of Binary Contingency Tables," preprint available at <http://front.math.ucdavis.edu/math.ST/0606650>.
7. D. Randall, "Rapidly Mixing Markov Chains with Applications in Computer Science and Physics," *Computing in Science & Eng.*, vol. 8, no. 2, 2006, pp. 30–41.
8. P. Diaconis and A. Gangolli, "Rectangular Arrays with Fixed Margins," *Discrete Probability and Algorithms*, D. Aldous et al., eds., Springer-Verlag, 1995, pp. 15–41.
9. R. Kannan, P. Tetali, and S. Vempala, "Simple Markov-Chain Algorithms for Generating Bipartite Graphs and Tournaments," *Random Structures and Algorithms*, vol. 14, no. 4, 1999, pp. 293–308.
10. M. Jerrum, A. Sinclair, and E. Vigoda, "A Polynomial-Time Approximation Algorithm for the Permanent of a Matrix with Non-Negative Entries," *J. ACM*, vol. 51, no. 4, 2004, pp. 671–697.
11. I. Bezáková, N. Bhatnagar, and E. Vigoda, "Sampling Binary Contingency Tables with a Greedy Start," *Random Structures and Algorithms*, vol. 30, nos. 1–2, 2007, pp. 168–205.

Ivona Bezáková is an assistant professor at the Rochester Institute of Technology. Her research interests include randomized and approximation algorithms, discrete random structures, and the Markov chain Monte Carlo technique. Bezáková has a PhD in computer science from the University of Chicago. Contact her at ib@cs.rit.edu.



The American Institute of Physics is a not-for-profit membership corporation chartered in New York State in 1931 for the purpose of promoting the advancement and diffusion of the knowledge of physics and its application to human welfare. Leading societies in the fields of physics, astronomy, and related sciences are its members.

In order to achieve its purpose, AIP serves physics and related fields of science and technology by serving its member societies, individual scientists, educators, students, R&D leaders, and the general public with programs, services, and publications—information that matters.

The Institute publishes its own scientific journals as well as those of its member societies; provides abstracting and indexing services; provides online database services; disseminates reliable information on physics to the public; collects and analyzes statistics on the profession and on physics education; encourages and assists in the documentation and study of the history and philosophy of physics; cooperates with other organizations on educational projects at all levels; and collects and analyzes information on federal programs and budgets.

The scientists represented by the Institute through its member societies number more than 134,000. In addition, approximately 6,000 students in more than 700 colleges and universities are members of the Institute's Society of Physics Students, which includes the honor society Sigma Pi Sigma. Industry is represented through the membership of 37 Corporate Associates.

Governing Board: Mildred S. Dresselhaus† (chair), Lila M. Adair, David E. Aspnes, Anthony Atchley, Slade Cargill, Charles W. Carter Jr., Hilda A. Cerdeira, Timothy A. Cohnt, Lawrence A. Crum, Bruce H. Curren†, Morton M. Denn†, Michael D. Duncan, H. Frederick Dylla† (ex officio), Judith Flippen-Anderson, Judy R. Franz†, Brian J. Fraser, John A. Graham†, Toufic Hakim†, Ken Heller, William Hendee, Judy C. Holoviak, John J. Hopfield, Anthony M. Johnson, Leo Kadanoff, Angela R. Keyser, Timothy L. Killeen, Louis J. Lanzerotti, Harvey Leff, Rudolf Ludeket†, Kevin B. Marvel, Cherry Ann Murray, John A. Orcutt, Elizabeth A. Rogan†, Bahaa E. A. Saleh, Charles E. Schmidt, Joseph Serene, Benjamin B. Snavely† (ex officio), A. F. Spilhaus Jr., Gene Sprouse, Hervey (Peter) Stockman, Quinton L. Williams

†members of the Executive Committee.

Management Committee: H. Frederick Dylla, Executive Director and CEO; Richard Baccante, Treasurer and CFO; Theresa C. Braun, Vice President, Human Resources; James H. Stith, Vice President, Physics Resources; Darlene A. Walters, Senior Vice President, Publishing; Benjamin B. Snavely, Secretary.

www.aip.org