S1 was measured after each x-ray experiment; the K^+-EDTA, Ca^{2+}, and Mg^{2+} ATPase activities were 15 (pH 8.0), 2.6 (pH 8.0), and 0.07 s^{-1} (pH 7.0), respectively, at 25°C. Buffer scattering measurements were performed periodically throughout the measurement sequence. The average count rates were 3,000 to 4,000 counts per second from buffer solution and 6,000 to 11,000 counts per second from sample solutions. After correction with the use of current values measured with the ion chamber, we obtained the net scattering intensity data $I(S)$ by subtracting the buffer scattering from the sample solution scattering. We performed similar measurements with solutions of BSA, assuring the validity of the S1 measurements very near to the beam stop. The data were analyzed on PDP11/34 (DEC) and SPARC Station 330 (SUN) computers.

13. C. R. Bagshaw *et al.*, *Biochem. J.* **141**, 351 (1974).
14. A. Guinier, *Small-Angle Scattering of X-Rays* (Wiley, New York, 1955).
15. O. Glatter, in *Small-Angle X-Ray Scattering*, O. Glatter and O. Kratky, Eds. (Academic Press, New York, 1981), pp. 119–196; L. A. Feigin and D. I. Svergun, *Structure Analysis by Small-Angle X-Ray and Neutron Scattering* (Plenum, New York, 1987).
16. C. C. Goodno, *Methods Enzymol.* **85**, 116 (1982).
17. R. S. Goody and W. Hofmann, *J. Muscle Res. Cell Motil.* 1, 101 (1980); M. A. Geeves and D. R. Trentham, *Biochemistry* **21**, 2782 (1982).
18. J. E. Wells and R. G. Young, *Methods Enzymol.* **85**, 93 (1982).
19. D. A. Winkelmann, T. S. Baker, I. Rayment, *J. Cell Biol.* **114**, 701 (1991).
20. K. M. Kretzschmar, R. A. Mendelson, M. F. Morales, *Biochemistry* **17**, 2314 (1978); R. A. Mendelson and K. M. Kretzschmar, *ibid.* **19**, 4103 (1980).
21. The radius of gyration of R_g must be determined in the range $0 \leq S^2 \leq (1/2\pi R_g)^2$ (the Guinier's criterion) or at least in the range where $S^2 = (1/2\pi R_g)^2$ (estimated with the use of the approximate R_g value of the molecule) is involved (*14, 15*). The size of the region (a Guinier region) over which this approximation holds depends not only on the size but also on the shape of the molecule in a monodisperse solution. Outside the Guinier region the scattering profile is mainly governed by the shape of the molecule. For a head portion (S1) of a myosin molecule, the R_g value has been estimated as ~46 Å from its molecular weight and size (*23*). By referring to this value, we used the region $[S^2_{min} (= 0.0778 \times 10^{-4}$ Å$^{-2}) \leq S^2 \leq S^2_{max} (= 0.2228 \times 10^{-4}$ Å$^{-2})]$ in the determination of R_g to produce a linear plot with low variance as shown in Fig. 1. The excellent linearity of the Guinier plots in the region below S^2_{max} shows that there are no appreciable aggregates such as dimers in any of the samples. The region $< S^2_{min}$ was not used because parasitic scattering around the beam stop influenced the data slightly.
22. R. A. Mendelson, D. Bivin, P. M. G. Curmi, D. K. Schneider, D. B. Stone, *Adv. Biophys.* **27**, 145 (1991).
23. A. Elliott and G. Offer, *J. Mol. Biol.* **123**, 505 (1978).
24. We estimated the molecular mass of S1 [partial specific volume 0.72 cm^3/g (*10*)] by linearly extrapolating the zero angle intensity $I(0)/c$ data of S1 samples and BSA (molecular mass 66.3 kD) to zero protein concentration (*14, 15*).
25. C. Wells and C. R. Bagshaw, *J. Muscle Res. Cell Motil.* **5**, 97 (1984).
26. S. Highsmith and D. Eden, *Biochemistry* **29**, 4087 (1990).
27. R. Aguirre, S.-H. Lin, F. Gonsoulin, C.-K. Wang, H. C. Cheung, *ibid.* **28**, 799 (1989).
28. M. Walker and J. Trinick, *J. Muscle Res. Cell Motil.* **9**, 359 (1988).
29. M. Irving, V. Lombardi, G. Piazzesi, M. A. Ferenczi, *Nature* **357**, 156 (1992).
30. Y. Amemiya *et al.*, *Top. Curr. Chem.* **147**, 121 (1988).
31. K. Wakabayashi and Y. Amemiya, in *Handbook on Synchrotron Radiation*, S. Ebashi, M. Koch, E. Rubenstein, Eds. (North-Holland, Amsterdam, 1991), vol. 4, pp. 597–678. [For the original paper of the present optics, see Y. Amemiya *et al.*, *Nucl. Instrum. Methods* **208**, 471 (1983)].
32. P. R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, New York, 1969).
33. We thank M. Sato for the program used in the calculation of the $p(r)$ function and for assistance; T. Miyanishi for the sequence data of chicken S1; T. Hiratsuka for discussions; and Y. E. Goldman for critical reading of the manuscript and valuable comments.

3 April 1992; accepted 15 July 1992

Neural Computing in Cancer Drug Development: Predicting Mechanism of Action

John N. Weinstein,* Kurt W. Kohn, Michael R. Grever, Vellarkad N. Viswanadhan, Lawrence V. Rubinstein, Anne P. Monks, Dominic A. Scudiero, Lester Welch, Antonis D. Koutsoukos, August J. Chiausa, Kenneth D. Paull

Described here are neural networks capable of predicting a drug's mechanism of action from its pattern of activity against a panel of 60 malignant cell lines in the National Cancer Institute's drug screening program. Given six possible classes of mechanism, the network misses the correct category for only 12 out of 141 agents (8.5 percent), whereas linear discriminant analysis, a standard statistical technique, misses 20 out of 141 (14.2 percent). The success of the neural net indicates several things. (i) The cell line response patterns are rich in information about mechanism. (ii) Appropriately designed neural networks can make effective use of that information. (iii) Trained networks can be used to classify prospectively the more than 10,000 agents per year tested by the screening program. Related networks, in combination with classical statistical tools, will help in a variety of ways to move new anticancer agents through the pipeline from in vitro studies to clinical application.

There are millions of different molecules that should be screened for their activity against cancer. Some are natural products collected from rain forests, oceans, and other habitats; some are products of synthetic organic chemistry. Before 1985, primary screening of new compounds was done in mice bearing the murine leukemia P388. It was not clear, however, that this screening would identify agents effective against solid tumors, including the common human carcinomas (*1*). An alternative possibility, disease-oriented primary screening against panels of in vivo tumors, would not have been feasible for the very large numbers of candidate compounds. These issues motivated development of the current National Cancer Institute (NCI) drug screening, in which compounds are tested in culture for their ability to inhibit the growth of a panel of 60 different human tumor cell lines (*1*). The cell lines in the panel represent a spectrum of histologies and organs of origin. The NCI screening program is currently testing several hundred compounds per week, including both synthetic and natural products.

Implicit in this screening strategy is the premise that agents tested will show reproducible patterns of differential response among these 60 cell lines. Paull *et al.* developed a computer program called COMPARE (*2*) to seek useful information in those patterns. For a designated "seed" compound, COMPARE searches the database of already tested agents for ones with the most similar patterns of activity against the cell panel. The pairwise comparisons are based on Pearson correlation coefficients. Early on it was noticed that compounds matched by pattern often had similar chemical structures. Further study indicated that the matches generally related to in vitro

J. N. Weinstein and V. N. Viswanadhan, Laboratory of Mathematical Biology, Division of Cancer Biology, Diagnosis and Centers, National Cancer Institute, Bethesda, MD 20892.
K. W. Kohn, Laboratory of Molecular Pharmacology, Developmental Therapeutics Program, Division of Cancer Treatment, National Cancer Institute, Bethesda, MD 20892.
M. R. Grever, Developmental Therapeutics Program, Division of Cancer Treatment, National Cancer Institute, Bethesda, MD 20892.
L. V. Rubinstein and A. D. Koutsoukos, Biometric Research Branch, Cancer Therapy Evaluation Program, Division of Cancer Treatment, National Cancer Institute, Bethesda, MD 20892.
A. P. Monks, D. A. Scudiero, L. Welch, Program Resources, Inc./DynCorp., National Cancer Institute/Frederick Cancer Research and Development Center, Frederick, MD 21702.
A. J. Chiausa and K. D. Paull, Information Technology Branch, Developmental Therapeutics Program, Division of Cancer Treatment, National Cancer Institute, Bethesda, MD 20892.

*To whom correspondence should be addressed at Laboratory of Molecular Pharmacology, Developmental Therapeutics Program, Division of Cancer Treatment, National Cancer Institute, Building 37, Room, 5C-25, 9000 Rockville Pike, Bethesda, MD 20892.

biochemical mechanism of action (*3*).

Here, we focus on that issue, the mechanism of action: Can patterns of activity against cell lines in the panel be used to classify new agents by their mechanism of action? If so, then the information could be used to guide further biochemical studies and efforts at rational drug design. It could also be used to aid prospectively in deciding which of the hundreds of compounds tested each week should be selected for more detailed (and expensive) testing both in vitro and in vivo. To address this question, we have developed neural networks in which a drug's pattern of activity in the cell line screen is the input and its mechanism of action is to be predicted.

Neural computing (*4, 5*) is a relatively recent development in the information sciences, an outgrowth of artificial intelligence research in the 1950s and 1960s. Only in the last few years, however, have the algorithms advanced to the point at which they would be useful for our current application. Neural nets are so named because they exhibit certain analogies, at least superficially, to the way in which arrays of neurons probably function in biological learning and memory. They differ from the usual computer programs in that they "learn" from a set of examples rather than being programmed to get the right answer. The information is encoded in the strength of the network's "synaptic" connections.

Neural networks are rapidly making their way now into various areas of the biomedical sciences. They are being applied to pattern recognition or decision making in diverse fields including nucleic acid sequence analysis (6), protein sequence analysis (7), quantitative cytology (8), diagnostic imaging (9), neural organization (*10*), and speech recognition (*11*). They are also being used for clinical diagnosis of conditions such as lower back pain (*12*) and interstitial lung disease (*13*).

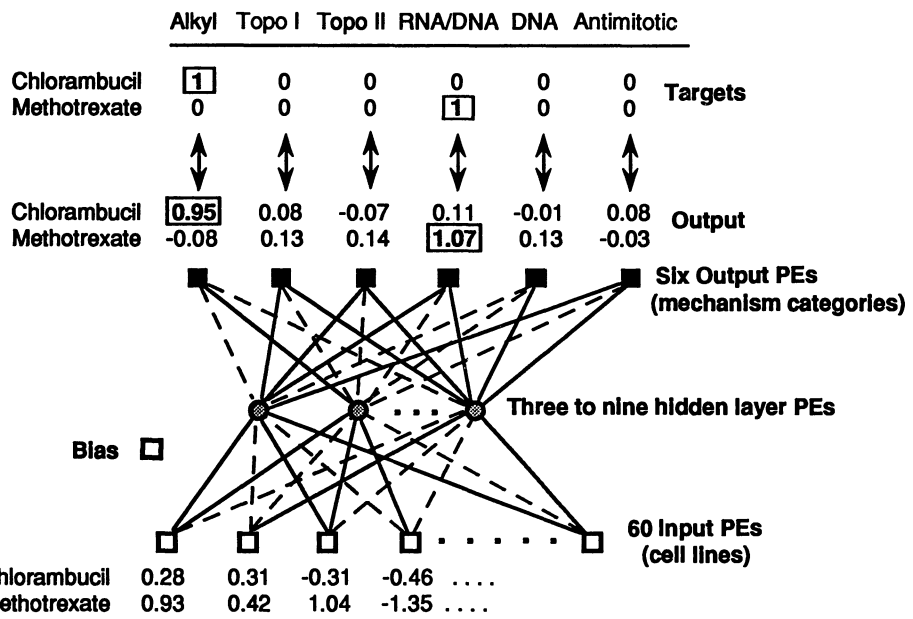

Fig. 1. Neural networks with three to nine hidden layer PEs used to classify agents in the NCI cancer drug screening program according to their mechanism of action. Sample inputs, outputs, and targets for a trained network are shown for two standard agents. Solid and dashed interconnections schematically represent different weights in the trained network. The output patterns closely match the patterns of targets for the two drugs (note values in boxes), which indicates that their mechanisms are correctly predicted. The bias is connected to all hidden layer and output PEs. Topo I and II, topoisomerase I and II inhibitors, respectively; RNA/DNA, RNA/DNA antimetabolites; DNA, DNA antimetabolites; alkyl, alkylating agent.

To predict mechanisms of action in the cancer drug screening program, we developed neural networks such as that shown in Fig. 1. The fundamental building blocks are processing elements (PEs), which can be likened to neurons, and weighted connections, which can be likened to synapses. The network shown has 60 input PEs, one for each cell line, and 6 output PEs, one for each of six mechanistic categories to be considered. The categories selected for this study were alkylating agent, topoisomerase I inhibitor, topoisomerase II inhibitor, RNA/DNA antimetabolite, DNA antimetabolite, and antimitotic agent. A seventh category, represented implicitly, corresponded to none of these.

Between the inputs and outputs of the

Table 1. Anticancer agents included in the NCI database. Numbers in parentheses are NSC numbers.

Alkylating agents (35): Asaley (167780), AZQ (182986), BCNU (409962), busulfan (750), carboxyphthalatoplatinum (271674), CBDCA (241240), CCNU (79037), CHIP (256927), chlorambucil (3088), chlorozotocin (178248), *cis*-platinum (119875), clomesone (338947), Cyanomorpholinodoxorubicin (357704), cyclodisone (348948), dianhydrogalactitol (132313), fluorodopan (73754), hepsulfam (329680), hycanthone (142982), melphalan (8806), methyl CCNU (95441), mitomycin C (26980), mitozolamide (353451), nitrogen mustard (762), PCNU (95466), piperazine alkylator (344007), piperazinedione (135758), pipobroman (25154), porfiromycin (56410), spirohydantoin mustard (172112), teroxirone (296934), tetraplatin (363812), thio-tepa (6396), triethylenemelamine (9706), uracil nitrogen mustard (34462), Yoshi-864 (102627).

Topoisomerase I inhibitors (35): camptothecin (94600), camptothecin Na salt (100880), camptothecin derivatives (95382, 107124, 643833, 629971, 295500, 249910, 606985, 374028, 603071, 176323, 295501, 6016172, 606173, 610458, 618939, 610457, 610459, 606499, 610456, 618939, 364830, 606497, 610458, and 606985, and nine confidential camptothecin derivatives).

Topoisomerase II inhibitors (16): doxorubicin (123127), amonafide (308847), *m*-AMSA (249992), anthrapyrazole derivative (355644), pyrazoloacridine (366140), bisantrene HCl (337766), daunorubicin (82151), deoxydoxorubicin (267469), mitoxantrone (301739), menogaril (269148), morpholinodoxorubicin (354646), *N,N*-dibenzyl daunomycin (268242), oxanthrazole (piroxantrone) (349174), rubidazone (164011), VM-26 (122819), VP-16 (141540).

RNA/DNA antimetabolites (19): L-alanosine (153353), 5-azacytidine (102816), 5-fluorouracil (19893), acivicin (163501), aminopterin (132483), aminopterin derivative (184692), aminopterin derivative (134033), an antifol (633713), an antifol (623017), Baker's soluble antifol (139105), dichloroallyl lawsone (126771), DUP 785 (brequinar) (368390), ftorafur (pro-drug) (148958), 5,6-dihydro-5-azacytidine (264880), methotrexate (740), methotrexate derivative (174121), *N*-(phosphonoacetyl)-L-aspartate (PALA) (224131), pyrazofurin (143095), trimetrexate (352122).

DNA antimetabolites (16): 3-HP (95678), 2′-deoxy-5-fluorouridine (5-FUdR) (27640), 5-HP (107392), alpha-TGDR (71851), aphidicolin glycinate (303812), ara-C (63878), 5-aza-2′-deoxycytidine (127716), beta-TGDR (71261), cyclocytidine (145668), guanazole (1895), hydroxyurea (32065), inosine glycodialdehyde (118994), macbecin II (330500), pyrazoloimidazole (IMPY) (51143), thioguanine (752), thiopurine (755).

Antimitotic agents (13): allocolchicine (406042), an antitubulin (609395), colchicine (757), a colchicine derivative (33410), dolastatin 10 (376128), maytansine (153858), rhizoxin (332598), taxol (125973), taxol derivative (608832), thiocolchicine (361792), trityl cysteine (83265), vinblastine sulfate (49842), vincristine sulfate (67574).

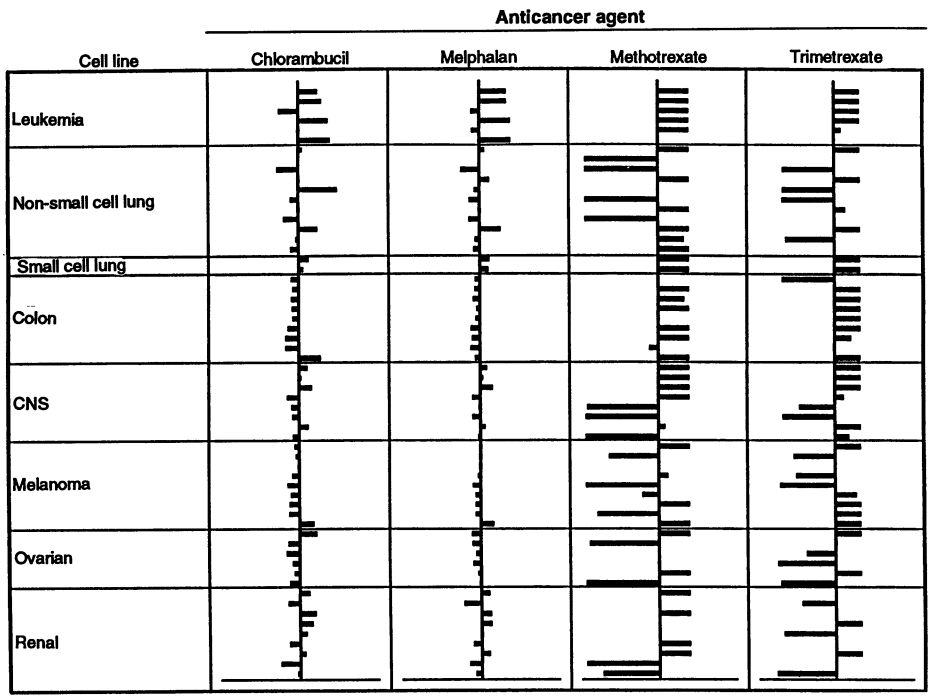

Fig. 2. Patterns of activity in the cancer cell line screen, as illustrated by data for four drugs. The first two, correctly classified by the network as alkylating agents, have quite similar patterns. The second pair, correctly classified as RNA/DNA antimetabolites, is also similar. In some cases, the GI_{50} values were outside of the concentration range tested—hence, the bars were of equal length. CNS, central nervous system.

neural network, we placed one hidden layer whose number of PEs can be varied. As shown in Fig. 1, each hidden layer PE is connected to all inputs and to all outputs; the more hidden layer PEs, the more complex the patterns that can be learned. Using a set of drugs whose mechanisms are (putatively) known, we can train the network to predict mechanism of action. We accomplished this training by iteratively presenting the activity pattern of each agent at the input layer, as shown in Fig. 1 for two standard drugs, methotrexate and chlorambucil. Each output PE corresponds to a category; the target for that output is 1 if the drug belongs to the category and 0 if it does not. Each of the "synaptic" connections between elements has associated with it a weight. During training, the outputs are compared with their corresponding targets, and the error is fed back to update the weights (that is, to reinforce or penalize the connections) with the use of the "back-propagation" algorithm (*14*). By that process, the network outputs come to approximate the target values (1's and 0's for the different classes). The drug is assigned to the category with the largest number as its output. A higher percentage of correct responses might be achieved with the use of a Bayesian optimization (that is, taking into account the fraction of drugs known a priori to be in each category). We chose not to do so, however, because the distribution of category types in the database used here presumably does not reflect that in the much larger database of unknown agents to which the trained neural network is being applied predictively.

The database (Tables 1 and 2) (*15*) was formulated from a list of standard agents, supplemented by additional drugs for which sufficient mechanistic information was available. Putative mechanisms of action were identified with the use of literature sources, structural homologies, and data from experiments on mechanism. The result was a list of 141 dose-range input vectors for a total of 134 agents thought to be in one of the categories noted above. For seven of the drugs, two different dose ranges were kept in the database because it was not clear which best reflected the dose-response characteristics. In most cases, the data represented averages over multiple experiments run at different times. Certain established agents were excluded if their mechanisms of action were unknown, if they fell into categories too small or diffuse to be handled easily, or if they were expected to require biochemical activation not possible in culture (for example, cytoxan, which requires activation by hepatic enzymes). For the same reasons, not all possible mechanisms of action were included.

Specification of mechanism is, of course, uncertain and to some degree arbitrary. A molecule may inhibit cell growth in multiple ways or else in ways dependent on cell type, drug dose, or culture conditions. To the extent that a mechanism could in fact be unambiguously assigned, it was also not clear that the mechanistic category would correlate well with functional pattern. For all of these reasons, and given inevitable experimental noise in the data, we were unsure at the outset whether reasonably predictive results could be obtained. However, the findings have been unexpectedly good.

We used as a cell panel the 60 tumor lines in the current screening program. There were, therefore, 141 input vectors, each consisting of 60 numbers derived from the dose-response curves in the screening. More specifically, we used as input the quantities Δ (*2*), defined as

$$\Delta = \log(1/GI_{50}) - \text{mean}[\log(1/GI_{50})] \quad (1)$$

where GI_{50} is the concentration of a drug required to decrease the growth of a cell line by 50% in the standard assay. Assay conditions and technical details have been described elsewhere (*16*). The value subtracted in Eq. 1 for a given drug relates to the mean of its cytotoxic potencies over the entire panel of cells. Because of this correction, simple differences in relative potency of the agents are nulled out, and their patterns of differential activity are emphasized. Figure 2 shows the cell lines and Δ patterns of four familiar agents: chlorambucil, melphalan, methotrexate, and trimetrexate. It is clear that the first two are closely related in pattern and different from the other pair. In the matrix of 141 × 60 = 8460 values, there were 604 (7.1%) missing values (that is, drug/cell combinations not

Table 2. Cancer cell lines included in the NCI screening program. CNS, central nervous system.

Non–small cell lung cancer (11): NCI-H23, NCI-H522, A549/ATCC, EKVX, NCI-H226, NCI-H322M, NCI-H460, HOP-62, HOP-18, HOP-92, LXFL 529.
Small cell lung cancer (2): DMS 114 and DMS 273.
Colon cancer (9): HT29, HCC-2998, HCT-116, SW-620, COLO 205, DLD-1, HCT-15, KM12, KM20L2.
Ovarian cancer (6): OVCAR-3, OVCAR-4, OVCAR-5, OVCAR-8, IGROV1, SK-OV-3.
Leukemia (6): CCRF-CEM, K-562, MOLT-4, HL-60(TB), RPMI-8226, SR.
Renal cell carcinoma (9): UO-31, SN12C, A498, CAKI-1, RXF-393, RXF-631, 786-O, ACHN, TK-10.
Melanoma (9): LOX IMVI, MALME-3M, SK-MEL-2, SK-MEL-5, M14, SK-MEL-28, M19-MEL, UACC-62, UACC-257.
CNS tumor (8): SNB-19, SNB-75, SNB-78, U251, SF-268, SF-295, SF-539, XF 498.

Table 3. Neural net predictions of mechanism from NCI's cancer drug screening data. The table presents a summary of results as a function of the number of hidden layer PEs and in comparison to the results from an LDA. For comparison with LDA, a two-tailed *P* value was determined from the McNemar test (*19*). The *P* values for seven and nine hidden layer PEs differ because there were a total of ten discrepancies in the case of seven PEs and 12 discrepancies in that of nine PEs.

Hidden layer PEs (*n*)	Incorrect predictions	Correct predictions	Comparison with LDA (*P* value)
3	20 (14.2%)	121	1.00
5	13 (9.2%)	128	0.07
7	12 (8.5%)	129	0.02
9	12 (8.5%)	129	0.04
Linear discriminant analysis	20 (14.2%)	121	—

Table 4. Prediction of mechanism category from cancer drug screening data with the use of a back-propagation neural network with seven hidden layer elements. The left column gives the putative mechanism; numbers in bold are correct predictions. R/D indicates RNA/DNA antimetabolites; DNA indicates DNA antimetabolites. Topo I and II, topoisomerase I and II inhibitors, respectively; Alkyl, alkylating agent. Note that the table reflects results for 141 drug concentration vectors representing 134 different drugs. For seven of the drugs, two different dose-response profiles were included because it was not clear which best indicated the properties of the agent. The total of correct predictions was 129 (91.5%); the total of incorrect predictions was 12 (8.5%).

Actual category	Predicted category						
	Alkyl	Topo I	Topo II	R/D	DNA	Antimitotic	None
Alkylating	**33**	0	2	0	0	0	0
Topo I	0	**35**	0	0	0	0	0
Topo II	1	0	**18**	0	0	0	0
RNA/DNA	1	0	1	**17**	0	1	0
DNA	1	0	1	1	**12**	1	0
Antimitotic	1	1	0	0	0	**14**	0

tested or else eliminated at the time of the experiment for quality control reasons). For each missing value, we inserted the mean value obtained for the cell line over all drugs in the set. Neural networks have a property termed "graceful degradation" (*5*): unlike computer algorithms that fail abruptly when confronted with missing data values or unexpected numbers, neural nets are relatively insensitive to such problems because information is encoded in distributed form throughout the net structure.

In principle, a neural network with enough hidden layer elements can learn to recognize any functional relation between inputs and outputs no matter how complex and nonlinear it is. Therefore, a large enough network could always be trained to recognize particulars of the training data set and thus assign drugs in it to their proper categories with 100% accuracy—even if the net had little or no capacity to predict for agents on which it had not been trained. The aim, then, is to design a network with enough hidden layer elements to capture the rich, nonlinear patterns of interaction without providing so many hidden layer elements (that is, so much plasticity) that the network memorizes the training set and loses the capacity to generalize. We set up our networks (*17*) using the NeuralWorks Professional II/PLUS 386/387 program package (NeuralWare, Inc., Pittsburgh, Pennsylvania).

To be sure that we were considering only the true predictive power of the network, we employed the following statistical cross-validation scheme: The drug vectors were randomly divided into ten approximately equal subsets (14 or 15 in each). The networks were then trained on nine-tenths of the overall data set and tested for predictive power on the other tenth. This procedure was repeated ten times, each time with a different test subset. Thus, each drug's mechanism of action was predicted by a network trained on an entirely independent set. Somewhat better predictions would be expected in the logical limit of this cross-validation scheme, in which only one drug at a time was left out and the train-test cycle was repeated 141 times.

Table 3 shows the percentage of correctly predicted categories as a function of the number of hidden layer elements included in the network. The choice of seven hidden PEs appears at least as good as any other. Table 4 shows in more detail the correct (129) and incorrect (12) choices for a network with seven hidden layer elements. For comparison, the equivalent linear discriminant analysis (with equal prior probabilities) was performed with the use of the SAS program set (*18*) and found to give 20 wrong classifications. As indicated in Table 3, a two-tailed *P* value of 0.02 was calculated by the McNemar test (*19*) for the null hypothesis that the apparent difference in accuracy between neural network and linear discriminant analysis is the result of chance.

Several of the "mistakes" in classification by the neural net involved DNA antimetabolites. When examined by techniques of cluster analysis (*20*), this appeared to be a highly heterogeneous mixture of agents not well correlated with patterns of growth inhibition in the screen. In that sense, the neural network prediction was not "wrong."

Dolastatin 10 and maytansine (antimitotics) were each correctly classified on the basis of experiments in the appropriate dose range but misclassified with the use of Δ values from a second, apparently inappropriate range (whose inclusion in the analysis degraded the results obtained). In the case of guanazole, there was a tie to two decimal places between the output for DNA/RNA antimetabolite (correct) and alkylating agent (incorrect); for the purpose of Tables 3 and 4, it was scored conservatively as incorrect. The other "misses" were morpholinodoxorubicin, hycanthone, macbecin II, carboxyphthalatoplatinum, 2′-deoxy-5-fluorouridine (5-FUdR), 5,6-dihydro-5-azacytidine, ftorafur (a prodrug of 5-fluorouracil), and an antifol (NSC 633713).

These findings substantiate the idea that patterns of differential growth inhibition in the screening panel contain an extraordinary amount of useful information and also that the neural network successfully recognizes that information. Frankly, we did not expect such clear-cut answers. We expected, for instance, that more of these agents would prove to express mixed mechanisms in their patterns of activity. In addition, the subcellular distribution of drugs is influenced by a variety of factors, including transmembrane potentials, transmembrane pH gradients, structure-selective transporters, and P-glycoprotein–mediated multidrug resistance. These factors can be viewed as adding "noise" to the system and may explain some of the incorrect classifications.

Several remaining issues deserve mention. First, categories chosen for this work may subsume more than one mechanism. For example, the alkylating agents include nitrogen mustards and platinum complexes as well as nitrosoureas, which act by different mechanisms and which exhibit distinctive spectra of selective effects against various types of cells and experimental tumors. The neural net analysis of the current cytotoxicity assay, therefore, has the capacity to recognize composite groups of distinct mechanisms if trained to do so. Clearly, the neural networks and databases could be designed in many different ways to answer a variety of specific questions related, but not identical, to those addressed here (*21*). Second, our designation of mechanistic categories should be considered tentative; individual drugs may act by a combination of different mechanisms, and the available evidence is often inconclusive. Third, many close chemical relatives are included

among the drugs used in the current analysis. For example, the topoisomerase I group is comprised solely of camptothecin derivatives, and we do not yet know whether the network will recognize chemically different subclasses of topoisomerase I agents when and if they are identified. Finally, mechanisms operating in the assay may not always be the most prominent ones in vivo or in other cell culture systems.

As often happens, some of the most intriguing clues may come from the data that do not fit. For example, the network with three hidden layer PEs classified mitomycin and porfiromycin as antimitotics instead of alkylating agents. (When five or more hidden layer PEs were used, the classifications were correct.) Interestingly, these compounds react in the minor groove of DNA, whereas all of the other alkylators in the data set react in the major groove. It will be important, therefore, to test other types of minor groove binders, such as tomaymycin, the anthramycins, and the pyrrolo-1,4-benzodiazepines. The classification network thus appears to be a good source of clues as to the fine structure of mechanistic categories, and sometimes those clues come from a comparison of optimal and nonoptimal networks.

We are currently using neural networks in the prospective analysis of new compounds tested by the NCI drug screening program. It appears that neural computing, when combined with other statistical techniques for pattern recognition and decision making, can play a productive role in the development of new agents for the treatment of such diseases as cancer and acquired immunodeficiency syndrome (AIDS).

REFERENCES AND NOTES

1. M. R. Boyd, *Princ. Pract. Oncol. Updates* **3**, 1 (1989).
2. K. D. Paull *et al.*, *J. Natl. Cancer Inst.* **81**, 1088 (1989).
3. R. Bai *et al.*, *J. Biol. Chem.* **266**, 15882 (1991); K. D. Paull *et al.*, *Cancer Res.* **52**, 3892 (1992); H. N. Jayaram *et al.*, *Biochem. Biophys. Res. Commun.* **186**, 1600 (1992); K. D. Paull, E. Hamel, L. Malspeis, in *Cancer Chemotherapeutic Agents*, W. E. Foye, Ed. (American Chemical Society, Washington, DC, in press).
4. T. Khanna, *Foundations of Neural Networks* (Addison-Wesley, New York, 1991); H. H. Szu, *Telemat. Inform.* **7**, 403 (1990).
5. J. Dayhoff, *Neural Network Architectures* (van Nostrand Reinhold, New York, 1990).
6. B. Demeler and G. Zhou, *Nucleic Acids Res.* **19**, 1593 (1991); A. V. Lukashin *et al.*, *J. Biomol. Struct. Dyn.* **6**, 1123 (1989); S. Brunak, J. Engelbrecht, S. Knudsen, *Nucleic Acids Res.* **18**, 4797 (1990); M. C. O'Neill, *ibid.* **19**, 313 (1991).
7. L. H. Holley and M. Karplus, *Proc. Natl. Acad. Sci. U.S.A.* **86**, 152 (1989); D. G. Kneller, F. E. Cohen, R. Langridge, *J. Mol. Biol.* **214**, 171 (1990); M. J. McGregor, T. P. Flores, M. J. E. Sternberg, *Protein Eng.* **2**, 521 (1989).
8. G. L. Wied *et al.*, *Anal. Quant. Cytol. Histol.* **12**, 417 (1990); W. H. Wolberg and O. L. Mangasarian, *ibid.*, p. 314; H. E. Dytch and G. L. Wied, *ibid.*, p. 379.
9. J. M. Boone, V. G. Sigillito, G. S. Shaber, *Med. Phys.* **17**, 234 (1990); J. S. DaPonte and P. Sherman, *Comput. Med. Imag. Graph.* **15**, 3 (1991).
10. M.-M. Mesulam, *Neurol. Prog.* **28**, 597 (1990).
11. S. Kurogi, *Biol. Cybern.* **64**, 243 (1991).
12. N. H. I. Mann and M. D. Brown, *Orthop. Clin. North Am.* **22**, 303 (1991).
13. N. Asada *et al.*, *Radiology* **177**, 857 (1990).
14. D. E. Rumelhart, G. E. Hinton, R. J. Williams, *Nature* **323**, 533 (1986).
15. Insofar as possible, construction of the database was kept independent of prior knowledge of results in the screen. However, screen results were used in a number of cases for the limited purpose of deciding which dose-range vector for a given drug captured the dynamic range of activity. Halichondrin B was initially predicted on the basis of COMPARE analysis of screen results to be an antitubulin agent, and that was later verified experimentally.
16. K. D. Paull *et al.*, *Proc. Am. Assoc. Cancer Res.* **29**, 488 (1988); M. C. Alley *et al.*, *Cancer Res.* **48**, 589 (1988); D. Scudiero *et al.*, *ibid.*, p. 4827; R. H. Shoemaker *et al.*, in *Prediction of Response to Cancer Chemotherapy*, T. Hall, Ed. (Liss, New York, 1988), pp. 265–286; A. Monks *et al.*, *J. Natl. Cancer Inst.* **83**, 757 (1991).
17. Network type: feed-forward, back-propagation; control strategy: normal cumulative delta learning with epoch size 30 and random presentation of input vectors during training. Initial learning coefficients and momentum: 0.3 and 0.45 for the hidden layer and 0.15 and 0.4 for the output layer (all four values decreased by a factor of 2 after 10,000 presentations). The extent of training was 15,000 presentations. (This was determined in preliminary experiments to fall within a broad optimal range with respect to the root-mean-square error of prediction in the ten test sets. This approach to cross-validation has general application.) Classification rule: plurality wins. For each pair of training and test sets, the network was trained and then tested seven times from seven different initial random numbers. The resulting output values were then averaged, a procedure that decreased considerably the number wrongly predicted (for example, from a mean of 15.4 to 12 in the case of seven hidden layer elements).
18. Statistical Analysis System/Statistics (*SAS/STAT*) *User's Guide, Version 6* (SAS Institute, Cary, NC, ed. 4, 1990), vol. 1. Linear discriminant analysis (LDA) is a classical linear statistical technique for classifying vectors of numbers (in this case, the Δ values) according to known examples, where the vectors are assumed to follow a multivariate normal distribution with known or (as in this case) estimated parameters. Thus, LDA can be considered a purely linear equivalent of the network. See, for example, D. J. Hand, *Discrimination and Classification* (Wiley, New York, 1981).
19. The McNemar test is a binomial calculation based on the discrepant predictions of the two methods that asks how likely it is that the difference in the number of incorrect predictions arose by chance [Q. NcNemar, *Psychometrika* **12**, 153 (1947)].
20. Standard techniques of cluster analysis on the 60-line screen patterns (with a Euclidean distance metric and average linkage clustering) supported evidence from the neural network that the category of DNA antimetabolites is functionally highly diverse.
21. As one additional example, calculations were done to assess predictiveness for each category as opposed to any other category. Given a single, overall optimum threshold, a network with nine hidden layer elements made an average of 3.5 mistakes per category out of a possible 141.

22. We thank J. S. Driscoll for help in identifying mechanisms of action; B. Bunow, B. A. Chabner, S. M. Sieber, and J. S. Driscoll for critical reading of the manuscript; and R. D. Everly and G. Ferruzza for technical assistance with the software.

15 May 1992; accepted 1 September 1992

Age and Duration of Weathering by ^{40}K-^{40}Ar and $^{40}Ar/^{39}Ar$ Analysis of Potassium-Manganese Oxides

Paulo M. Vasconcelos, Tim A. Becker, Paul R. Renne, George H. Brimhall

Supergene cryptomelane $[K_{1-2}(Mn^{3+}Mn^{4+})_8O_{16} \cdot xH_2O]$ samples from deeply weathered pegmatites in southeastern Brazil subjected to ^{40}K-^{40}Ar and $^{40}Ar/^{39}Ar$ analysis yielded ^{40}K-^{40}Ar dates ranging from 10.1 ± 0.5 to 5.6 ± 0.2 Ma (million years ago). Laser-probe $^{40}Ar/^{39}Ar$ step-heating of the two most disparate samples yielded plateau dates of 9.94 ± 0.05 and 5.59 ± 0.10 Ma, corresponding, within 2 σ, to the ^{40}K-^{40}Ar dates. The results imply that deep weathering profiles along the eastern Brazilian margin do not reflect present climatic conditions but are the result of a long-term process that was already advanced by the late Miocene. Weathering ages predate pulses of continental sedimentation along the eastern Brazilian margin and suggest that there was a time lag between weathering and erosion processes and sedimentation processes.

Rates of chemical interaction between rocks and the hydrosphere and atmosphere at the Earth's surface have only rarely been constrained by direct dating techniques (*1*). Radiocarbon residence times, uranium series dating (*2*), thermoluminescence techniques, and more recently ^{10}Be (*3*, *4*), ^{36}Cl (*5*, *6*), ^{26}Al (*3*), ^{3}He (*7*, *8*), and oxygen isotopes (*9*) have been used successfully in special cases to constrain weathering or exposure ages. The ages of many weathering surfaces in tectonically stable cratons, however, are beyond the useful limits of some of these techniques. Other techniques either are inapplicable because of the lack of datable overlying volcanic deposits (*9*) or are unreliable because they require assumptions about element immobility after incor-

P. M. Vasconcelos and G. H. Brimhall, Department of Geology and Geophysics, University of California, Berkeley, CA 94720.
T. A. Becker and P. R. Renne, Institute of Human Origins Geochronology Center, 2453 Ridge Road, Berkeley, CA 94709.