Directed Discovery of Novel Drug Cocktails[1]

John H. Miller[2] Ralph Zinner[3] Brittany Barrett[4]

July 20, 2005

[1]Thanks to Juri G. Gelovani for laboratory resources and Andrei Y. Volgin for biological and technical advice on the experiments. Reuban Lotan, Hai T. Tran, Claudio Pisano, Dafna Lotan, John M. Kurie, and Jean P. Issa contributed some of the drugs used in the experiments. Waun K. Hong, Gordon B. Mills, and Li Mao provided a number of useful insights into the project. Miller was supported by core funding from the Santa Fe Institute, and Zinner and Barrett received support from Eli Lilly Oncology and the BESCT Lung Cancer Program at University of Texas, M. D. Anderson Cancer Center.

[2]Carnegie Mellon University and the Santa Fe Institute, *email:* miller@santafe.edu (corresponding co-author).

[3]M. D. Anderson Cancer Center, *email:* rzinner@mdanderson.org (corresponding co-author).

[4]M. D. Anderson Cancer Center, *email:* BriBarret@mdanderson.org.

Abstract

Combinations of drugs can result in effective treatments for certain diseases like HIV/AIDS. Unfortunately, our ability to discover such combinations is quite limited, as drugs often interact in highly nonlinear ways and thus it is difficult to predict *a priori* which cocktails are likely to be effective. Moreover, the brute-force approach of screening all possible combinations fails due to the combinatorial explosion of possible cocktails, even when we consider modest numbers of candidate drugs. As an alternative, here we use a nonlinear search algorithm designed to direct the discovery of novel, effective drug cocktails. We demonstrate this approach by finding chemotherapy cocktails that can inhibit A549 (lung carcinoma) cells using a hybrid, nonlinear-optimization algorithm. We find that directed discovery can be an effective means of automatically deriving novel cocktails using a relatively small number of experiments. The basic idea of directed discovery explored here has a variety of other applications across many fields.

1 Introduction

Certain combinations of drugs result in effective treatments for some diseases like HIV/AIDS. Unfortunately, discovering such drug "cocktails" is often quite difficult. *A priori*, it is difficult to predict the effectiveness of any given cocktail as drugs often interact with one another in highly nonlinear ways. Occasionally, we may have a good understanding of the various molecular pathways that need to be targeted, but such knowledge is typically lacking. An obvious alternative to the above is to search over the set of all possible cocktails. However, the combinatorics implied by such a search quickly become unwieldy, as even modest-sized cocktails result in enormous search spaces.[1] As an alternative, here we employ a non-linear search algorithm to direct the discovery of effective drug cocktails.

To demonstrate this approach, we use a hybrid, nonlinear-optimization algorithm to search for drug cocktails that can effectively kill off A549 cells (a line derived from a lung carcinoma). We search over a space of nineteen possible drugs, and thus there are over half a million possible cocktails. Our pilot results indicate that directed search can indeed locate effective cocktails even with a relatively small sampling of the possible combinations. We find a cocktail that performed 4.18 standard deviations above the mean of a random sample[2] after observing only 370 cocktails.

The general idea of directed discovery has wide applicability. In terms of medical applications, along with developing cancer chemotherapies discussed below, other applications include the development of anti-viral, diagnostic, and other treatment regimes. Such directed search could also be used in other fields, for example, to derive chemical or biological cocktails and processes for industrial applications, such as chemical engineering and environmental clean up, or to conduct other types of experiments needed in business applications such as product development, production, and marketing.

1.1 A Framework for Directed Discovery

The basic framework we consider is one in which we are attempting to experimentally discover good solutions to a problem. We assume that the space of

[1] For example, combinations over a set of twenty potential drugs result in over a million possible cocktails.

[2] Which, if the underlying distribution is normal, would imply the upper 0.0014% of the distribution. The assumption of normality is discussed later in the paper.

possible solutions, S, is large enough that enumerative search is impossible. Each possible solution, $s \in S$, generates a "fitness," with higher values of fitness implying better solutions. If the constituent elements of a solution interact linearly, searching for good solutions is of $O(N)$, where N is the number of elements (and thus enumerative search becomes feasible). However, if the elements of the solution interact nonlinearly, in the worse case scenario search is of $O(2^N)$ (with binary elements). We suspect that typically the world exists between these two extremes, and that while there are nonlinearities, there are also sufficient regularities in the space that a search algorithm designed to seek out high-fitness solutions with a minimal number of experiments can be used to direct the discovery of good solutions.

Below we consider simple drug cocktails as potential solutions to the problem of killing cancerous cells. We use nineteen different drugs, each with a pre-determined, fixed dose, and thus there are 2^{19} or 524,288 possible combinations. Of course, more elaborate search spaces could be considered, for example, drugs could have variable dosing, application order, delivery regimes, and so on.

Each combination of drugs results in a measure of fitness, $F(s)$. Below we use an intentionally simple measure of fitness that combines the effectiveness of the drug cocktail in eliminating cancerous cells with a cost for each drug used in the cocktail. This latter cost is designed to encourage smaller mixes of drugs *ceteris paribus*. The fitness measure results in an objective function for the optimization algorithm. Therefore, the fitness function needs to be carefully chosen to achieve the desired outcome. Much more elaborate fitness functions than the one used below are possible. For example, fitness could be tied to the ability of the cocktail to also preserve normal cells, prevent other side effects, result in low-cost combinations of drugs, exploit nonlinear effects over single-agent doses, and so on.

The final element of the framework is a search algorithm designed to seek out high-fitness combinations with a minimal number of experimental trials. Here we use a hybrid, nonlinear search algorithm. Obviously, there is a large class of such algorithms, and our particular choice, while informed by some simulation experiments, represents just one of many possibilities. We did find that our algorithm was limited in its effectiveness due to a number of experimental confounds identified during the initial series of experiments—even so, the algorithm was able to find some very effective cocktails. Moreover, we were then able to identify even more effective cocktails using a modified search strategy and an improved set of laboratory protocols.

2 A Demonstration of Directed Discovery for Chemotherapy Cocktails

To demonstrate the underlying idea of directed discovery, we consider an application to the development of chemotherapy cocktails for the A549 lung carcinoma cell line (see the Appendix for a fuller description of the methods and materials). The cells were first isolated and concentrated so that the $50\mu l$ of cell suspension added to each well of a 96-well plate[3] resulted in an expected 1500 cells per well. We had access to nineteen different drugs, seventeen of which were targeted agents and two of which are used in standard chemotherapy (see Table 1). Based on the results of prior experiments, we identified drug dilutions so that $10\mu l$ of the single-agent just provoked a detectable inhibition response in a well (the mean inhibition was 8% across the nineteen drugs). Each drug cocktail consisted of the $10\mu l$ increments from each associated drug along with the addition of DMSO and media, so that a total volume of $200\mu l$ of drug/DMSO/media was added to each well. We adjusted the final DMSO concentration of each well to a constant value (set equal to the DMSO concentration of total volume that would result if all nineteen drugs were in a single cocktail, here 0.66%).

Each drug cocktail was replicated in three horizontally contiguous wells of the plate. Three wells in the center of the plate and three on the outside edge received a cocktail with no drugs (but with DMSO and media additions as described above). The plate was then placed in an incubator and 44 hours later $25\mu l$ or 10% of the total well volume of Cell Proliferation Reagent WST-1 (Roche Diagnostics, Indianapolis, IN) was added and the plate was then incubated for another four hours. The WST-1 assay provides a measure of cell survival (via differential absorbance of the reagent). The optical density of the plate was then read on a scanning, multi-well spectrophotometer at 440nm with a 600nm reference, and it was averaged over the three wells used for each cocktail. This mean value was normalized by dividing it by the mean optical density for the six control wells. Our measure of cell survival, $K(s)$, was given by this normalized value.

The algorithm we use here begins with a population of thirty randomly created cocktails. This initial population is known as Generation 0. Each of the nineteen drugs had a 0.33 independent probability of being included in each cocktail. The various cocktails were tested as described above and were

[3]Each well had a capacity of $400\mu l$ with a tissue culture area of $0.33cm^2$.

Drug	Concentration μM	% DMSO at given drug concentration
4HPR	5	0.05000
5-aza-2' deoxycytidine (decitabine)	2	0.02000
anisomycin	0.015	0.00003
ATRA (Vesanoid)	15	0.15000
bortezomib (Velcade)	0.005	0.00000
CD437	0.3	0.00300
cisplatin (Platinol-AQ)	5	0.00000
deguelin	12	0.12000
gemcitabine HC1 (Gemzar)	0.004	0.00000
imatinib mesylate (Gleevec)	3	0.00000
indirubin-3'-oxime	1	0.00200
LY294002 HC1	2	0.02000
MX3350-1	0.5	0.00500
PD 168393	5	0.10000
rapamycin	6	0.10971
SAHA	2	0.02000
SCH66336 (Sarasar)	5	0.05000
SP 600125	5	0.01000
ST1926	0.05	0.00050

Table 1: Drugs used in the experiment. Gemcitabine HC1 and cisplatin are considered to be chemo-therapeutics, the remaining drugs are targeted-therapeutics.

assigned a fitness value equal to

$$-K(s) - 0.1 \times |s|,$$

where $K(s)$ is the measure of cell survival from above and $|s|$ is the number of drugs used in the associated cocktail. Note that the minus sign in front of $K(s)$ implies that cocktails that have lower survival rates (that is, kill more cells) are fitter *ceteris paribus*. The second term of the fitness function imposes a penalty on cocktails that use more drugs, thus, the addition of a new drug in a cocktail needs to result in at least 10% less survival before fitness will increase.

The algorithm then proceeds through a series of iterations or generations. During each generation, a new population of cocktails is formed by first reproducing members of the old population biased by performance. A subset of these reproduced cocktails are then modified via mutation to produce some novel cocktails, and these are then tested using the lab procedures outlined above. The algorithm is then iterated on to the next generation.

In each generation we maintain a population of thirty cocktails. We first select thirty cocktails from the previous generation to serve as the basis for the new population. We use tournament selection, whereby three members of the previous population are randomly chosen (with replacement) and a copy of the one with the highest fitness in this group is added to the new population (in case of a tie, we randomly pick among the best). Note that this selection procedure, while biased in favor of higher fitness cocktails, does not necessarily guarantee that the best member of the population will be retained or the worst eliminated.

Next, we take fifteen randomly chosen members of the new population, and mutate them to generate fifteen novel cocktails. Mutation works by taking an existing cocktail and independently, for each of the nineteen possible drugs in the mix, having a 7% chance of altering that drug (that is, adding the drug to the mix if it was not present or eliminating it if it was used). We continue to mutate a cocktail until it results in a novel cocktail, that is, one that has not been previously observed during the prior series of experiments.

Thus, at the end of the above procedure we have a population of fifteen previously observed (and perhaps repeated) cocktails and fifteen novel cocktails. These latter cocktails are assigned fitness values using the previously described laboratory procedure, and the generation is concluded. Therefore, each new population is a biased (by fitness) selection from the previous pop-

ulation with some local modifications that insure that half of the population is novel.

The algorithm above is a hybrid of some well known non-linear optimization algorithms. Like a genetic algorithm [1], it maintains a population and uses a biased selection mechanism and mutation for modification, however, it does not incorporate any kind of recombination. The algorithm also takes on some characteristics of hill-climbing and simulated annealing [2]. Like these latter two algorithms, the mutation operator performs local search based on a status quo. The key difference, though, is that here the status quo is a population of points (versus a singleton). As the algorithm iterates over time, it is likely that the set of status quo points collapses down to a singleton via selection, and thus the algorithm's behavior progresses from something akin to a genetic algorithm to something more similar to simulated annealing and hill climbing.

2.1 Results

Figure 1 shows the results of the algorithm for a single experiment (PilotA–Experiment 0).[4] The fitness of the initial thirty randomly-generated cocktails is given by the red points on the far left (Generation 0). In Generation 1, the black points are the fitness values for the fifteen selected, but unmodified, cocktails. The red points in this generation, indicate the fitness values of the fifteen novel experiments. Given the selection mechanism, we would expect that the distributions of the black points over any given generation will tend to move upward and tighten over time. (Note that in Generation 5, the best cocktail observed to date was lost due to selection.) Of more interest is the pattern of the red points, as these represent the exploration of novel cocktails. We find that with some exceptions (Generations 6 and 9 are notable), the distribution of newly explored cocktails tends to show increasing fitness.

Table 2 provides some descriptive statistics of the populations during the search. In Generation 0 we randomly created thirty cocktails. In each subsequent generation, fifteen novel cocktails were created via mutation. The statistics in the table are only over the novel cocktails during each generation. The data indicate that the search had mixed success. Clearly in many of the generations the novel populations appeared to be showing clear improvements

[4]Two such experiments were performed. The results of PilotA–Experiment 1 were similar.

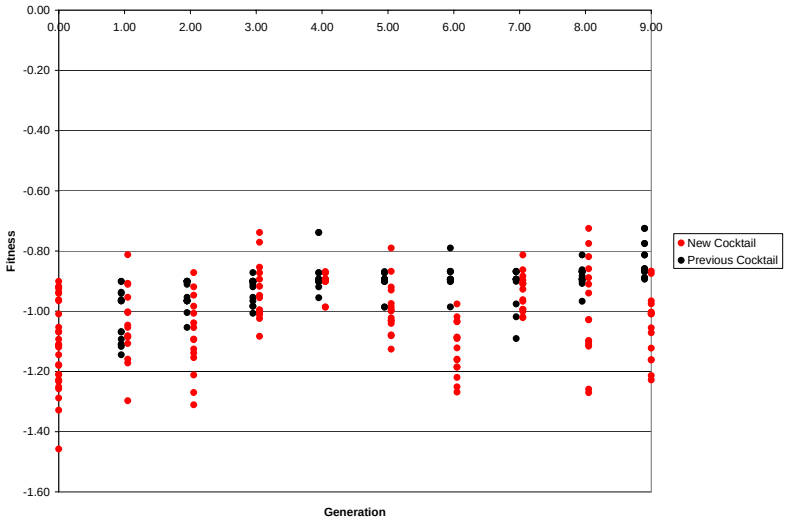

Figure 1: Results of direct search by generation. Red points (RHS) represent newly conducted experiments while black ones (LHS) indicate retained experiments from earlier generations.

Generation	Average	Std. Dev.	Best (New) In Generation	Best To Date
0	-1.11	0.14	-0.90	-0.90
1	-1.05	0.12	-0.81	-0.81
2	-1.08	0.13	-0.87	-0.81
3	-0.94	0.10	-0.74	-0.74
4	-0.99	0.06	-0.87	-0.74
5	-0.99	0.09	-0.79	-0.74
6	-1.13	0.09	-0.98	-0.74
7	-0.94	0.06	-0.81	-0.74
8	-1.00	0.17	-0.72	-0.72
9	-1.05	0.11	-0.87	-0.72

Table 2: Descriptive statistics for novel cocktails during PilotA–Experiment 0 search. The data for Generation 0 are based on thirty, randomly generated cocktails, while those for subsequent generations include only the fifteen novel cocktails created in that generation via mutation. In total, 165 $(30 + 15 \times 9)$ novel cocktails were explored during this experiment out of 524,288 possibilities.

over the initial population; yet, even a casual view of the data does not indicate a clear trend towards higher performing cocktails.

An analysis of the data at the end of Experiment 0 indicated some important (and ultimately confounding) sources of noise in the optical density measure. First, wells with only DMSO and media (but no cells), had a mean optical density of 0.45 with a standard deviation of 0.02 or about 4% of the mean. A similar analysis of (interior) wells containing DMSO, media, and cells, gave a mean optical density of 2.26 with a standard deviation of 0.15 (7% of the mean), and thus the biological activity in these latter wells appears to increase the variance. Another important source of variation is differences caused by the location of the well, in particular, whether or not the well is in the interior of the plate. If we separate out the interior and exterior wells from the previous experiment, we find that the mean optical density is 2.26 ($sd = 0.15$, $n = 60$) for interior wells and 1.31 ($sd = 0.23$, $n = 36$) for exterior wells, and thus exterior wells have an expected assay value of about 58% of the interior wells. The plates used in our experiment had a similar drop off in the six control wells, with the exterior wells having an assay value of about

	Number of Edge Wells		
	0	1	3
Mean Fitness	-1.10	-1.05	-0.96
Std. Dev.	0.13	0.11	0.12
n	121	132	77

Table 3: Number of edge wells versus mean fitness of the cocktails during PilotA–Experiment 0.

67% of the interior ones (2.75 versus 1.84 with $n = 33$ in both cases). Further evidence of an edge effect is apparent in a comparison of the mean *fitness* of the wells with cocktails. As seen in Table 3, the fitness values increase as the number of edges increase which, given that higher fitness is associated with lower optical density, is consistent with the previous observations.

Note that the algorithm we used is sensitive to excessive noise. To economize on the number of experiments, we only assigned fitness to a cocktail once, and did not retest it during the search. Thus, it is possible for an inferior cocktail to be given a high fitness evaluation due to serendipity (say, all of its wells were on the edge and the raw noise from the assay worked in its favor). Similarly, even good cocktails could receive low fitness evaluations if they were unlucky. Given our selection mechanism and the inherent variation in the experiments, it is likely that even low fitness cocktails could have survived for long periods and confounded the search.

To further explore some of the above issues we performed a hill-climbing search using the best cocktail identified by Generation 7 (this cocktail arose in Generation 3). For hill climbing, we searched all of the "one-mutant" neighbors of the cocktail of interest (the "status quo"). One-mutant neighbors of a given cocktail are all those cocktails that differ by either the addition of a single drug that was not in the status quo cocktail or the elimination of a drug that was included. Thus, we explored twenty cocktails (the status quo and its nineteen possible, one-mutant neighbors)—all of these cocktails were placed in the interior of the plates to avoid edge effects. Based on the results of the initial hill climb, we selected a new status quo point (that was given by the cocktail common to the two best mutants observed), and again explored all of the one-mutant neighbors. This latter status quo point proved to be the (perhaps local) optimum. Figure 2 provides a comparison of the original

cocktails (Generation $0r$[5]) and the populations from the first (designated One Mut 1) and second (One Mut 2) hill climbs. The mean fitness of the second one-mutant population was -0.64 ($sd = 0.07$) while that of the Generation $0r$ population was -1.18 ($sd = 0.15$), and the hypothesis that the means of the two distributions were the same is easily rejected ($t = 17.1$). The fitness of the best one-mutant cocktail was -0.54, which is around 4.18 standard deviations away from the mean of the Generation $0r$ distribution (using the higher variance from that distribution). If the underlying distribution is normal, such an outcome would be in the upper 0.0014% of the distribution, and therefore we would expect that given our 370 observations (from the two experiments and two hill climbs), such an outcome would arise about 0.53% of the time by random search. Note that this interpretation of the data is predicated on the assumption of normality, and so caution is urged. As discussed below, we do find that our best cocktails are hitting a lower bound on $K(s)$, and this might cause the tails of the underlying distribution to fatten, implying that it may not be normally distributed. Nonetheless, the search did appear to turn up a very fit cocktail given our objective function.

The results of the one-mutant experiment also illustrate a second issue that may have confounded effective search, namely reaching limits in the assay's ability to differentiate among effective drug combinations. The assay's lower bound appears to have a normalized optical density of around 0.20. Many of the drug cocktails we explored reached this lower bound, that is, they apparently killed off the maximum number of cells we could detect. If such a lower bound is easy to achieve, then there is not much information that can be exploited by the search algorithm, since a large number of perhaps unrelated combinations may lead to high fitness. At some level, this is not a problem, as we are achieving our objective. Moreover, recall that our fitness function incorporates both cell kill and the number of drugs used in the cocktail. Thus, once cocktails hit the lower bound of cells killed, the only way to increase fitness is to seek out cocktails that achieve this outcome with fewer drugs. Such behavior was observed during our hill-climbing search.

The best cocktail we found used three drugs: 4HPR, SAHA, and Velcade. This combination had a $K(s)$ of around 0.18 to 0.22 depending on

[5]The "r" designation is for a replication we did of the original Generation 0 cocktails to explore the repeatability of the experiments with the same cells across 46 days. The mean fitness of the initial data was -1.11 while that of the replication was -1.18. The two sets of data had a correlation coefficient of 0.68 and further analysis indicated that the fitness of the replicated cocktails tended to be slightly worse than that of the original generation.

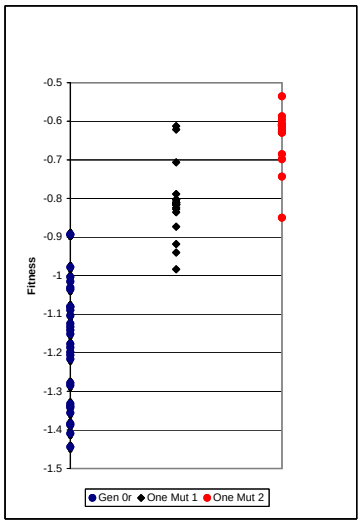

Figure 2: Comparison of Generation 0r (random cocktails) to the one-mutant neighbor populations arising from a hill-climbing search. The first such population (One Mut 1) used as its status quo the best observed cocktail arising in PilotA–Experiment 0 up to Generation 7 (this particular cocktail was first discovered in Generation 3 of that experiment). The second population (One Mut 2) had its status quo given by the cocktail that was formed by the common drugs in the two best cocktails observed in One Mut 1. This latter status quo appears to be the (perhaps local) optimum.

Drug	$K(s)$ as Single Agent
4HPR	0.92
SAHA	0.87
Velcade	0.68

Table 4: Normalized cell survival $K(s)$ given single agent at full dose ($1\times$).

the experiment. During the course of our 370 experiments, we investigated twenty two other triplets (out of 969 possible), and observed $K(s)$ values in the range of 0.43–0.99. These latter triplets were not a random sample (they should be biased toward lower values of $K(s)$), so this provides some support for the notion that the cocktail we found is not a common feature of triplets. Moreover, if we order all of the experimental outcomes from lowest to highest values of $K(s)$, we find that all of the top fifty five cocktails used the three drugs found in our best cocktail (of the top seventy five cocktails, all but four incorporate these drugs). Table 4 lists the single agent effect of the three drugs in question. Both 4HPR and SAHA have relatively small single-agent impacts, with Velcade showing a bit more kill (but within the norms we observed). Table 5 shows the impact on $K(s)$ of the cocktail and of the cocktail minus each individual drug (at both full dose and a 55% dilution designed to move us away from the lower bound of $K(s)$). Please note that the data in both these tables are from a single experiment, so there may be considerable variance underlying the observations. That being said, if each drug just contributed its single-agent impact, we would expect $K(s)$ to equal 0.48 rather than 0.18. If the world were linear, we would also predict that 4HPR and SAHA would have less impact than they do, while Velcade should have more. That is, 4HPR and SAHA seem to work better in combination than you might expect *a priori*. Indeed, this effect is quite dramatic under the 55% dilution regime, where the elimination of either of these two drugs tends to cause a large loss in effectiveness of the remaining drugs. A fuller exploration of the nonlinearities involved in this cocktail, and the underlying mechanisms that account for its behavior, is of interest.

Drug	$K(s)$ at $1\times$	$K(s)$ at $0.55\times$
All 3	0.18	0.35
no 4HPR	0.37	0.92
no SAHA	0.45	0.82
no Velcade	0.37	0.58

Table 5: Normalized cell survival $K(s)$ at full dose ($1\times$) and 55% dilution ($0.55\times$) for the cocktail and the cocktail less the specified single agent

3 Conclusions

The directed discovery of novel drug cocktails, and more broadly, the use of such techniques in a variety of other domains, allows us to break through the usual knowledge and combinatorial bounds that constrain our ability to find good experimental solutions to important problems. Heretofore, searches for, say, novel drug cocktails either required an explicit knowledge of the underlying molecular pathways in the system or an ability to do massive numbers of experiments to explore all possible combinations—both of which are often lacking. As demonstrated above, directed discovery was able to identify some very effective novel drug cocktails by exploring only 370 out of the 524,288 possible cocktails. The ability to efficiently conduct such experimental searches provides a new means by which to make progress on some difficult problems.

The general idea of directed discovery has wide applicability. As shown above, it can be used as a basis to develop novel cancer chemotherapies. Similar techniques could be used to develop anti-viral therapies, diagnostic techniques, and other treatment regimes. Applications are also readily found in a variety of other fields, ranging from chemical and environmental engineering to business.

References

[1] Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor: The U. of Michigan Press.

[2] Kirkpatrick, S., C. D. Gelatt Jr., M. P. Vecchi (1983), "Optimization by Simulated Annealing," *Science*, 220, 4598, 671-680.

Appendix: Materials and Methods

Cell Culture: Human lung carcinoma cells, A549, were obtained from the ATCC (American Type Culture Collection, Manassas, VA). The A549 cells were cultured in a monolayer with a 1:1 (vol/vol) mixture of Gibco Dulbecco's modified Eagle's medium: Nutrient Mixture F-12 from Invitrogen Corporation (Carlsbad, CA). The media was supplemented with Gibco US Qualified Fetal Bovine Serum from Invitrogen Corporation (Carlsbad, CA) and Cellgro Antibiotic-Antimycotic Solution from Mediatech, Inc. (Herndon, VA) at 10% and 1% (vol/vol), respectively. The cells were cultured at 37°C in a humidified atmosphere containing 95% air and 5% CO_2.

Drugs: Anisomycin, indirubin-3'-oxime, LY294002 hydrochloride, and SP 600125 were obtained from Tocris Cookson, Inc. (Ellisville, MO). All were dissolved in dimethyl sulfoxide (DMSO) purchased from Fisher Scientific (Pittsburgh, PA) at concentrations of $50mM$, $50mM$, $10mM$, and $50mM$, respectively. Deguelin and rapamycin were both from Axxora, LLC (San Diego, CA), and dissolved in DMSO to $10mM$ and $5mg/ml$, respectively. Hai Tran's Lab prepared cisplatin (Platinol-AQ) from Bristol-Myers Squibb, Co. (Princeton, NJ) and imatinib mesylate (Gleevec) from Novartis Pharmaceuticals Corporation (East Hanover, NJ). Cisplatin was dissolved in a sodium chloride solution to $1mg/ml$, and imatinib mesylate was dissolved to $12mg/ml$ in sterile water for injection. Both gemcitabine HCl (Gemzar) from Eli Lilly and Company (Indianapolis, IN) and bortezomib (Velcade) from Millennium Pharmaceuticals, Inc. (Cambridge, MA) were dissolved in media to $10mM$ and $1mg/ml$ respectively. 5-aza-2' deoxycytidine was obtained from Sigma-Aldrich, Inc. (St. Louis, MO) and dissolved to $10mM$ in sterile water. PD 168393 from Calbiochem (San Diego, CA) was diluted in DMSO to a concentration of $5mM$. The final seven drugs came from Dr. Reuben Lotan's lab and were prepared by Dafna Lotan. All were prepared at a $10mM$ concentration in DMSO and stored under nitrogen. The seven drugs were as follows: 4-HPR (N-(4-hydroxyphenyl)retinamide) from the National Cancer Institute (Bethesda, MD), ATRA (all trans-retinoic acid) from F. Hoffmann-La Roche (Basel, Switzerland), CD437 (6-[3-(1-adamantyl)-4-hydroxyphenyl]-2-naphthalene carboxylic acid) from CIRD/Galderma R&D (Sophia Antipolis, France), MXC3350-1 from Maxia Pharmaceuticals (La Jolla, CA), SAHA (suberoylanilide hydroxamic acid) from the Midwest Research Institute, SCH66336 (Sarasar) from Schering-Plough (Kenilworth, NJ), and ST1926 (E-3-(4'-hydroxy-3'-adamantylbiphenyl-4-yl) acrylic acid)

from Sigma-Tau (Pomezia, Italy). All drug aliquots were stored in the dark at -20°C, except for 5-aza-2' deoxycytidine, which was stored in the dark at -80°C.

Preparation of Plates: A549 cells were isolated from BD Falcon $75cm^2$ tissue culture treated flasks from BD (Franklin Lakes, NJ) after a Cellgro Phosphate Buffered Saline from Mediatech, Inc. (Herndon, VA) rinse and the addition of Gibco Trypsin-EDTA (0.05% Trypsin, with EDTA 4 Na) from Invitrogen Corporation (Carlsbad, CA). The cells were then plated with $50\mu l$ per well in 96 MicroWell NunclonΔ plates from Nalge Nunc International (Rochester, NY) at a cell concentration of 1,500 cells/well. Plates were then stored at 37°C in a humidified atmosphere containing 95% air and 5% CO_2 until the addition of the drugs.

Drug Preparation: Drug doses were chosen which allowed for a minimal amount of cell kill. Cell proliferation curves under the presence of a given drug were generated to determine the appropriate dose to use for each individual drug. All drugs were further diluted with media from their stock solutions so that the concentration of each drug was such that the addition of $10\mu l$ of a drug into a final well volume of $250\mu l$ resulted in that target concentration across the final well volume.

Combinations: The drugs for a given combination were mixed in a 96-well $2ml$ sterile polypropylene block from Denville Scientific (Metuchen, NJ). All combinations were mixed at four times the needed volume for one well, since each was added in triplicate and the additional amount allowed for pipetting error. Since 19 drugs were used in the experiment, the maximum percent volume of DMSO was calculated based on the assumption that one well could contain all 19 drugs as one combination (0.66% by volume DMSO for the 19 drugs). All combinations in an experiment were normalized to the maximum allowable DMSO concentration of 0.66% DMSO by volume, by adding the necessary amount of DMSO to each combination. Since wells contained varying numbers of drugs and DMSO amounts, media was added to attain a $250\mu l$ final well volume. The cells occupied $50\mu l$, therefore $200\mu l$ of drug/DMSO/media mix was added to the cells in each well. The plates were returned to 37°C in a humidified atmosphere containing 95% air and 5% CO_2.

Determination of Cell Proliferation: After 44 hours of incubation the plates were removed from the incubator and WST-1 Cell Proliferation Reagent from Roche Diagnostics (Indianapolis, IN) was added. The reagent was added at 10% by volume, therefore, $25\mu l$ of WST-1 was added to the $250\mu l$

well volume. The plates were returned to 37°C in a humidified atmosphere containing 95% air and 5% CO_2 for four hours. At 48 hours of total incubation, four hours after the WST-1 addition, the absorbance of the dye solution was measured on a scanning, multi-well spectrophotometer at $440nm$ with a $600nm$ reference.