

DERIVING CHEMOSENSITIVITY FROM CELL LINES: FORENSIC BIOINFORMATICS AND REPRODUCIBLE RESEARCH IN HIGH-THROUGHPUT BIOLOGY

BY KEITH A. BAGGERLY¹ AND KEVIN R. COOMBES²

University of Texas

High-throughput biological assays such as microarrays let us ask very detailed questions about how diseases operate, and promise to let us personalize therapy. Data processing, however, is often not described well enough to allow for exact reproduction of the results, leading to exercises in “forensic bioinformatics” where aspects of raw data and reported results are used to infer what methods must have been employed. Unfortunately, poor documentation can shift from an inconvenience to an active danger when it obscures not just methods but errors. In this report we examine several related papers purporting to use microarray-based signatures of drug sensitivity derived from cell lines to predict patient response. Patients in clinical trials are currently being allocated to treatment arms on the basis of these results. However, we show in five case studies that the results incorporate several simple errors that may be putting patients at risk. One theme that emerges is that the most common errors are simple (e.g., row or column offsets); conversely, it is our experience that the most simple errors are common. We then discuss steps we are taking to avoid such errors in our own investigations.

1. Background.

1.1. *General.* High-throughput biological assays such as microarrays let us ask very detailed questions about how diseases operate, and promise to let us personalize therapy. For example, if we know a priori that only about 70% of a group of cancer patients will respond to standard front-line therapy (as is the case with ovarian cancer), then an array-based test indicating whether a given patient is likely to respond would have great clinical utility.

Data processing in high-throughput studies, however, is often not described well enough to allow for exact reproduction of the results. While various groups have made efforts to construct tools or compendia that should make such reproducibility easier [e.g., Leisch (2002), Ruschhaupt et al. (2004), Gentleman (2005), Gentleman and Temple Lang (2007), Li (2008)], a recent survey [Ioannidis et al. (2009)] of 18 quantitative papers published in *Nature Genetics* in the past two years

Received September 2009.

¹Supported in part by NIH Grant P50 CA083639 09.

²Supported in part by NIH Grant P50 CA070907 08.

Key words and phrases. Microarrays, reproducibility, forensic bioinformatics.

found reproducibility was not achievable even in principle for 10. The time period is relevant, as the journal was then requiring that all raw data be deposited in public repositories so reproducibility would be easier. In some cases, more extensive exercises in “forensic bioinformatics” [e.g., [Stivers et al. \(2003\)](#), [Baggerly, Morris and Coombes \(2004\)](#), [Baggerly et al. \(2004, 2005\)](#), [Coombes, Wang and Baggerly \(2007\)](#)] can use aspects of raw data and reported results to infer the methods that must have been employed, but these are time-consuming.

Unfortunately, poor documentation and irreproducibility can shift from an inconvenience to an active danger when it obscures not just methods but errors. This can lead to scenarios where well-meaning investigators argue in good faith for treating patients with apparently promising drugs that are in fact ineffective or even contraindicated. In some ways, this problem is qualitatively worse for high-throughput assays than it is for simple tests, because the sheer magnitude of the data confounds both our ability to spot simple errors and our intuition about how certain tests “should” behave. For several single-marker assays, investigators may have a well-developed intuition that high values indicate a poor prognosis, but 50-gene signatures require faith that the assembly code has been applied correctly.

In this report we try to illustrate the potential severity of the problem by examining several related papers purporting to use microarray-based signatures of drug sensitivity derived from cell lines to predict patient response. Patients in clinical trials are currently being allocated to treatment arms on the basis of these results. However, we show in five case studies that the results incorporate several simple errors that may be putting patients at risk. The complete code and documentation underlying our results are available as supplementary material (see the Appendix). We then discuss steps we are taking to avoid such errors in our own investigations.

1.2. Chemosensitivity and cell lines. While many microarray studies have clarified aspects of basic cancer biology, there is a definite push to make our studies more “translational” in that our bench or in silico results should be likely to translate into changes in patient care in short order. To that end, studies identifying new targets for drug development are of less import than studies guiding how we use the drugs available to us now. In short, we would like to know whether a given patient is likely to be sensitive (aka be a responder) to a given agent.

In this context, we briefly consider how the activity of new chemotherapeutics is assessed. Shortly after cytotoxic chemotherapy was introduced for the treatment of leukemia, the U.S. government established the Cancer Chemotherapy National Service Center (CCNSC) in part to help test new agents. Drugs were submitted, assigned an NSC number, and used to treat leukemic mice to see if they produced improvement. Both the agency and the method of testing have changed over time. The CCNSC was absorbed into the Developmental Therapeutics Program (DTP), but every new drug submitted is still assigned an NSC number. Tests shifted first from leukemic mice to immunologically compromised (nude) mice onto which various types of human tumors were grafted, and thence to a standard panel of cell

lines (the NCI60) derived from 9 types of human tumors. Sensitivity is assessed in cell lines through dilution assays. Six samples (aliquots) of a given cell line are put into separate wells, a starting dose is applied to well 1, one tenth that dose is applied to well 2, and so on through well 5; well 6 is left untreated. After a fixed period, the amount of growth seen in the treated wells is assessed relative to growth in the untreated well. Summary measures reported are the estimated doses required to achieve 50% growth inhibition (GI50 or IC50), total growth inhibition (TGI) or not only growth inhibition but 50% lethality of the starting cells (LC50). The important thing about this approach is that data on activity against the NCI60 panel is publicly available for every chemotherapeutic now in use, so a method of constructing sensitivity signatures from this panel would have extremely broad applicability.

1.3. *Initial claims.* In late 2006, Potti et al. (2006) introduced a method for combining microarray profiles of the NCI60 with drug sensitivity data to derive “signatures” of sensitivity to specific drugs, which they then used to predict patient response. In theory, the approach is straightforward:

- Using drug sensitivity data for a panel of cell lines, choose those that are most sensitive and most resistant to a drug.
- Using array profiles of the chosen cell lines, select the most differentially expressed genes.
- Using the selected genes, build a model that takes an array profile and returns a classification, and use this model to predict patient response.

They reported success using this approach with several common chemotherapeutic agents, specifically docetaxel, doxorubicin (adriamycin), paclitaxel (taxol), 5-fluorouracil (fluorouracil, FU, or 5-FU), cyclophosphamide (cytoxan), etoposide and topotecan. They also reported some initial success at predicting response to combination therapies. Unsurprisingly, these results generated a lot of attention. The approach was named one of “The Top 6 Genetics Stories of 2006” [Discover (2007)], is the subject of U.S. Patent Application 20090105167, “Predicting responsiveness to cancer therapeutics” [Potti, Nevins and Lancaster (2009)], and had been cited in 212 papers by August 2009 (Google Scholar).

1.4. *Initial questions and response.* Several groups at MD Anderson were among those excited, and we examined the approach in order to help our investigators use this. However, as documented in Coombes, Wang and Baggerly (2007), when we independently reanalyzed the raw data, we noted a number of simple errors affecting the conclusions. Taking doxorubicin as an example, we noted that Potti et al. (2006) predicted response in a test set cohort involving samples from 122 patients; 23 patients were reportedly sensitive to doxorubicin and 99 resistant [Potti et al. (2006), Figure 2C]. However, this data was said to arise from an “independent dataset of samples cultured from adriamycin-treated individuals (GEO

accession nos. GSE650 and GSE651).” These Gene Expression Omnibus (GEO) data sets, which derive from a study on acute lymphocytic leukemia (ALL) by Holleman et al. (2004), give gene profiles for samples from 94 patients sensitive to daunorubicin (in the same chemical family as doxorubicin) and 28 resistant. While the total number of patients (122) is constant, the proportions (23:99 and 94:28) are almost inverted, which led us to suggest that the sensitive/resistant labels might have been reversed. Label reversal implies that an “accurate” signature would preferentially suggest doxorubicin for patients who would *not* benefit.

In reply, Potti and Nevins (2007) disagree, claiming that their approach is “reproducible and robust.” In support of this claim, they comment on “the acute lymphocytic leukemia dataset in which the labels are accurate—full details are provided on our web page.” They note having gotten the approach to work again more recently [e.g., Hsu et al. (2007), Bonnefoi et al. (2007)], and conclude that the Coombes, Wang and Baggerly (2007) analysis is flawed.

1.5. *Subsequent progress.* Progress since 2006 appears substantial. Hsu et al. (2007) used the same approach to develop signatures of response to cisplatin and pemetrexed, extending the range of treatments that could be compared. In examining these signatures, they identified specific components strongly suggesting biological plausibility, noting that

“The cisplatin sensitivity predictor includes DNA repair genes such as ERCC1 and ERCC4, among others, that had altered expression in the list of cisplatin sensitivity predictor genes. Interestingly, one previously described mechanism of resistance to cisplatin therapy results from the increased capacity of cancer cells to repair DNA damage incurred, by activation of DNA repair genes.”

Clinical trial TOP0602 (NCT00509366), now underway, will “assign subjects to either pemetrexed/gemcitabine or cisplatin/gemcitabine therapy using a genomic based platinum predictor to determine chemotherapy sensitivity and predict response to chemotherapy for first-line therapy in advanced non-small cell lung cancer.” [Clinical Trial NCT00509366 (2009)].

Later, Bonnefoi et al. (2007) provided a “validation” of the combination approach, using it to predict patient response to two alternative therapies: taxotere followed by epirubicin and taxotere (TET), and fluorouracil, epirubicin and cyclophosphamide (FEC). This report is a substudy of the European clinical trial EORTC 10994/BIG 00-01, in which breast cancer patients are randomized to receive either TET or FEC. Based on their results, Bonnefoi et al. (2007) suggest that using their predictions might have improved response rates in the target population from 44% to 70%, a huge benefit.

The most recent application was when Augustine et al. (2009) used the approach to construct a signature for temozolomide. The procedure now appears fairly stan-

standardized:

“A signature of gene expression that correlated with resistance to temozolomide was derived from the NCI-60 panel of cancer cell lines (ref. 17; see also Supplementary data). From this panel of 60 cell lines, a smaller subset of 15 was selected that represented two extremes of sensitivity to temozolomide; nine of these cell lines were classified as resistant and six as sensitive. Using the gene expression profiles of these cell lines, we identified 45 genes that showed significantly different expression patterns between the resistant and sensitive cell lines and thus provided a temozolomide sensitivity gene signature (see Supplementary Table S4). The color-coded heatmap of expression of the 45 temozolomide genes across these 15 cell lines (Fig. 4A) shows 8 genes (red) that were more highly expressed in the resistant than in the sensitive cell lines, whereas 37 genes (blue) were more highly expressed in the sensitive than in the resistant cell lines.”

In sum, the procedure apparently gives good predictions in independent test sets, has some biological plausibility, appears to be giving stable results over years of application, and is consequently guiding treatment.

1.6. *Cases we examine here.* While there are other studies, we examine four of the cases listed above in more detail. We chose to examine doxorubicin because it is the drug for which we have the most information about specific predictions. We chose to examine cisplatin and pemetrexed both because the cisplatin signature contains specific genes that have been linked to drug resistance (e.g., ERCC1 and ERCC4) and because these signatures are guiding therapy now. We chose to examine the combination validation because most cancer patients get several drugs, not one. We chose to examine temozolomide because it is the most recent application we have seen.

The goals of our reanalyses differ slightly by study, partially reflecting differences in the types of data available. For doxorubicin, we tried to confirm the accuracy of the test set predictions. For cisplatin and pemetrexed, we tried to confirm the identities of the genes comprising the signature. For combination therapy, we tried to clarify the combination rules and validate predictions for the single best performing drug (cyclophosphamide). For temozolomide, we tried to confirm the association between the drug named and the results provided.

Finally, we provide a broader overview by graphically summarizing the cell lines designated sensitive and resistant to various drugs.

Data sources are listed in Table 1.

2. Case study 1: Doxorubicin. We begin by examining the data for doxorubicin, where the most test set prediction details are available. This data was discussed by Potti et al. (2006), questioned by Coombes, Wang and Baggerly (2007), reaffirmed by Potti and Nevins (2007), and revisited by Potti et al. (2008) as discussed below. In terms of documentation and reproducibility, our goal here is to confirm the test set prediction accuracy.

TABLE 1

Locations of data used in our analyses. Excel files were converted to csv files for loading into R. All data not posted elsewhere is available from our web site

Data sources

Potti et al. (2006) web site, accessed April 4, 2009

<http://data.genome.duke.edu/NatureMedicine.php>

Adria_ALL_data1_n95.doc

Celllines_in_each_predictor1.xls

ChemopredictorsParameters.xls

DescriptionOfPredictorGeneration.doc

GeneLists.zip

MDACC_data.zip

ParametersForImplementingSoftware.xls

Binreg.zip

GEO ids and Sens/Res calls for 95 samples

Sens/Res cell lines for each drug

Binreg parameter settings for 11 drugs

Cell line selection for docetaxel

Probesets splitting Sens/Res cell lines

TFAC Individual and combo predictions

Binreg parameter settings for 7 drugs

Metagene prediction software

Potti et al. (2006) web site, November 6, 2007, no longer posted

Adria_ALL.txt

Numbers, Sens/Res labels for 144 samples,
22 training cell lines, 122 testing samples

Hsu et al. (2007) journal web site, April 4, 2009

10593_Supplementary_Information.doc

10593_Supplementary_Table_1.doc

10593_Supplementary_Table_2.doc

Figure 1

Supplementary Methods

Cisplatin gene list

Pemetrexed gene list

Cisplatin and pemetrexed heatmaps

Gene Expression Omnibus (GEO), April 4, 2009

GSE649

GSE650

GSE651

GSE2351

GSE6861

Holleman et al. (2004) vincristine Res

Holleman et al. (2004) daunorubicin Sens

Holleman et al. (2004) daunorubicin Res

Lugthart et al. (2005) multi-drug Res

Bonnefoi et al. (2007) array CEL files

2.1. *Test data responder/nonresponder counts match those reported.* We first acquired the raw doxorubicin (adriamycin) data (Adria_ALL.txt) posted by Potti and Nevins (2007). This file contains 144 data columns: 22 for the cell lines used as training data, and 122 for test data samples. Sample columns are not named, but sensitive/resistant status is indicated for each. A side comment notes that “Validation data is from GSE4698, GSE649, GSE650, GSE651, and others.” We then checked label counts for the test data: 99 columns are labeled NR (nonresponders) and 23 Resp (responders), matching those reported by Potti et al. (2006).

2.2. *Training data sensitive/resistant labels are reversed.* We next tried to identify the cell lines used in the training set by matching the numbers reported to those in the table of NCI60 quantifications (see Table 1), focusing on those for the 22 cell lines used to produce the initially reported heatmap signature [Potti et al. (2006), Figure 2A]. The posted numbers have been transformed relative to the MAS5 quantifications used earlier. After some experimentation, we

TABLE 1
(Continued)

Data sources	
Holleman et al. (2004) web site, April 4, 2009 http://www.stjuderesearch.org/data/ALL4/key.xls	List of Holleman et al. (2004) files at GEO
Drug sens., April 4, 2009, http://dtp.nci.nih.gov/docs/cancer/cancer_data.html	
GI50_AUG08.BIN	GI50 values, all drugs, August 2008 release
LC50_AUG08.BIN	LC50 values, all drugs, August 2008 release
TGI_AUG08.BIN	TGI values, all drugs, August 2008 release
NCI60 array data, April 4, 2009, http://dtp.nci.nih.gov/mtargets/download.html	
WEB_HOOKS_NOV_GC_ALL.ZIP	MAS5.0 quantifications, U95A arrays
Augustine et al. (2009) journal web site, April 4, 2009	
Figure 4	Temozolomide heatmap
1.pdf	Gene list and metagene scores
Bonnefoi et al. (2007) journal web site, April 4, 2009	
mmc1.pdf	Webappendix with algorithm description
mmc2.pdf	Webfigure of individual drug ROC curves
mmc3.pdf	Webpanel listing cell lines used by drug
mmc4.pdf	Webtable 1 listing X3p Probesets by drug
mmc5.pdf	Webtable 2 Clinical data for samples used
Bonnefoi et al. (2007) files obtained from contact for GSE 6861. Not posted.	
dataBonnefoiPaper.txt	Drug and combination scores
HB131.CEL	CEL file missing from GEO

found that the data were log-transformed, the values for each row (gene) were centered and scaled to have mean zero and unit variance (separately for training and test data), exponentiated to undo the log-transform, and rounded to two decimal places. In order to match the numbers more precisely, we transformed the NCI60 quantifications the same way. The initial training data matrix for doxorubicin has 12558 rows (the 12625 Affymetrix U95Av2 probesets less 67 controls), but Adria_ALL.txt has 8958. This is due to the fact that the test data come from U133Av2 arrays, and the mapping across platforms matched probes based on unique LocusLink IDs. Since the rows are not labeled, we simply tried checking all row-by-row correlations between Adria_ALL.txt and the transformed NCI60 data pertaining to the 22 cell lines in the initial heatmap by brute force. We identified perfect matches for all 8958, establishing that the 10 cell lines labeled “Resistant” are SF-539, SNB-75, MDA-MB-435, NCI-H23, M14, MALME-3M, SK-MEL-2, SK-MEL-28, SK-MEL-5 and UACC-62, and the 12 cell lines labeled “Sensitive” are NCI/ADR-RES, HCT-15, HT29, EKVX, NCI-H322M, IGROV1, OVCAR-3, OVCAR-4, OVCAR-5, OVCAR-8, SK-OV-3 and CAKI-1. These lines (in this order) do produce the heatmap shown in Figure 2A of Potti et al. (2006), and the sensitive/resistant orientation is consistent with their Supplementary Figure 2, which

notes that heatmaps for each predictor have “resistant lines on the left, and sensitive on the right.” However, the labels are reversed relative to those now supplied on the Potti et al. (2006) web site in “Cell lines used in each chemo predictor.” Since the listing above places NCI/ADR-RES (“adriamycin resistant”) in the sensitive group, the training labels in *Adria_ALL.txt* are reversed.

2.3. *Heatmaps show sample duplication in the test data* [Figure 1(a)]. In order to see how well the doxorubicin signature separates test set responders from nonresponders, we extracted the test set expression values for the probesets in the doxorubicin signature (using the identifications from the previous step), log-transformed them, and clustered the samples [Figure 1(a)]. We did not see clear separation of responders from nonresponders, but we did see “blocks” of samples having exactly the same profiles: some test data samples were reused. Two such blocks are marked.

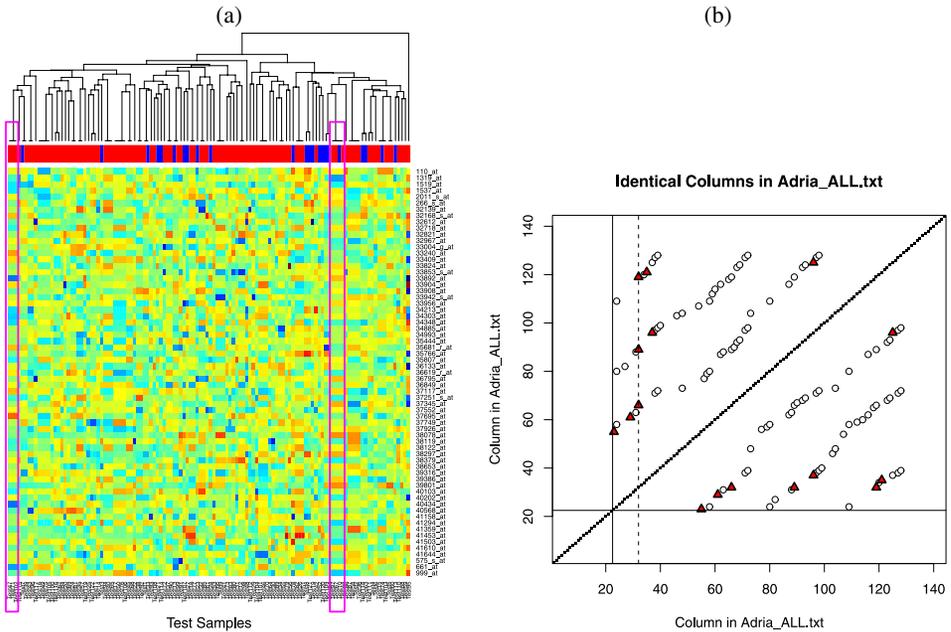


FIG. 1. (a) Heatmap of test samples from *Adria_ALL.txt*, using expression values for genes in the doxorubicin signature. Samples are labeled as “NR” (red) or “Resp” (blue). Responders and nonresponders show no clear separation. However, the clustering shows that some samples are tied (horizontal segments near the dendrogram base). Lines border one group of tied samples at left, and another group of tied samples at right; there are others. Colors for the tied block at right show Resp once and NR three times. (b) Dotplot identifying identical columns in the full *Adria_ALL.txt* matrix. There are no unwanted ties in the training data (lower left corner, separated by solid lines). In the test data (upper right), however, there are tied pairs labeled both consistently (circles) and inconsistently (triangles). One column from the right-hand block from (a) is marked with a dashed line; sample 32 (main diagonal) is labeled “Resp,” but samples 66, 89, and 117 are labeled “NR” (triangles).

2.4. Only 84/122 test samples are distinct; some samples are labeled both sensitive and resistant [Figure 1(b)]. In order to get a more precise idea of the extent of the duplication, we examined all pairwise sample correlations for Adria_ALL.txt. Figure 1(b) highlights pairs with correlations above 0.9999 (duplicates). Circles indicate duplicates with consistent labels; triangles indicate duplicates where the same sample was labeled sensitive in one column and resistant in the other. Only 84 of the 122 test samples are distinct: 60 are present once, 14 twice, 6 three times, and 4 four times. Columns 32 (dashed line), 66, 89 and 117 in Adria_ALL.txt are all the same, but they are labeled Sensitive, Resistant, Resistant and Resistant respectively [this is the set of samples marked at the right in Figure 1(a)].

2.5. *Communication with the journal elicited a second correction.* This correction [Potti et al. (2008)] notes that

“of the 122 samples assayed for sensitivity to daunorubicin for which the authors applied a predictor of adriamycin sensitivity, 27 samples were replicated owing to the fact that the same samples were included in several separate series files in the Gene Expression Omnibus generated in 2004 and 2005, which were the source of the data provided for the study.”

This correction also notes that data were acquired from two other sources in addition to GSE650 and GSE651: GSE2351 [Lugthart et al. (2005)] and GSE649 [Holleman et al. (2004)].

2.6. *The new data also has duplications and samples listed both ways (Table 2).* At the time of the Potti et al. (2008) correction, Adria_ALL.txt was removed from the web site and replaced with Adria_ALL_data1_n95.doc. Adria_ALL_data1_n95.doc gives no quantifications, but rather lists 95 GEO array ids and gives a

TABLE 2
The first 20 rows of Adria_ALL_data1_n95.doc. Visual examination shows rows 3/4, 9/10 and 17/18 (marked) are the same. Rows 17/18 label the same sample both ways

Row	Sample ID	Response	Row	Sample ID	Response
1	GSM44303	RES	11	GSM9694	RES
2	GSM44304	RES	12	GSM9695	RES
3	GSM9653*	RES	13	GSM9696	RES
4	GSM9653*	RES	14	GSM9698	RES
5	GSM9654	RES	15	GSM9699	SEN
6	GSM9655	RES	16	GSM9701	RES
7	GSM9656	RES	17	GSM9708*	RES
8	GSM9657	RES	18	GSM9708*	SEN
9	GSM9658*	SEN	19	GSM9709	RES
10	GSM9658*	SEN	20	GSM9711	RES

TABLE 3

Joint classifications of the 80 distinct samples in Adria_ALL_data1_n95.doc. Holleman et al. (2004) class most samples as sensitive to daunorubicin; Potti et al. (2006) class most as resistant. Adria_ALL_data1_n95.doc has 95 rows with 15 duplicates; 6 duplicates are labeled inconsistently and are classed as “Both”

		Holleman et al. (2004) classifications			
		Sensitive	Intermediate	Resistant	
Potti et al. (2006) classifications	Sensitive	13	0	0	13
	Resistant	29	10	22	61
	Both	6	0	0	6
		48	10	22	

sensitive/resistant label for each. The first twenty data rows from this file are shown in Table 2. Rows 3/4, 9/10 and 17/18 list the same samples. Further, the status labels in rows 17/18 conflict. Looking at all the names shows that only 80 are distinct: 15 are duplicated, and 6 of these show the same sample labeled both RES and SEN.

2.7. *At least 3/8 of the test data is incorrectly labeled resistant (Table 3).* Given the test sample ids, it is also possible to compare the classifications Potti et al. (2008) assigned (not *predicted*, but rather what they are trying to predict) to the 80 unique samples with the classifications assigned using the rules from Holleman et al. (2004), the source of the test data (Table 3). Between 29 and 35 samples Holleman et al. (2004) would call sensitive are classed by Potti et al. (2008) as resistant, with the uncertainty due to duplicate samples with inconsistent calls. There are 10 samples Holleman et al. (2004) would call “intermediate” that are classed by Potti et al. (2008) as resistant. All of the changes add to the number of “resistant” cases. This shift is not achievable by simply moving the LC50 cut-off to redefine the sensitive/resistant boundary, as LC50 values for “sensitive” and “resistant” samples overlap.

2.8. *Summary.* Poor documentation hid both sensitive/resistant label reversal, and the incorrect use of duplicate (and in some cases mislabeled) samples. These problems were hidden well enough to survive two explicit corrections. Code and documentation for this case study are given in Supplementary File 1 [Baggerly and Coombes (2009a)].

3. Case study 2: Cisplatin and pemetrexed. We next examined data for cisplatin and pemetrexed, where (a) specific genes in the cisplatin signature (ERCC1, ERCC4) were identified as providing a plausible rationale for its effectiveness and (b) the signatures are guiding therapy. Most signatures we discuss (including that

for pemetrexed) are assembled from the NCI60, but Hsu et al. (2007) comment that

“The collection of data in the NCI-60 data occasionally does not represent a significant diversity in resistant and sensitive cell lines to any given drug. Thus, if a drug screening experiment did not result in widely variable GI50/IC50 and/or LC50 data, the generation of a genomic predictor is not possible using our methods, as in the case of cisplatin.”

Thus, Hsu et al. (2007) assembled the cisplatin signature from a panel of 30 cell lines profiled by Györfy et al. (2006). Györfy et al. supply both U133A array quantifications and classifications of which cell lines were sensitive, intermediate or resistant in their response to various drugs, including cisplatin (their Figure 2). In terms of documentation and reproducibility, our goal here is to recreate the signature. We acquired the cisplatin and pemetrexed gene lists from Hsu et al. (2007)’s Supplementary Tables 1 and 2, respectively. We acquired the Györfy et al. (2006) expression data from their supporting Table 1. We acquired the `binreg` Matlab scripts used for model fitting and heatmap generation from the Potti et al. (2006) web site. As the analyses are largely parallel, we focus on cisplatin first, turning to pemetrexed only in the next to last subsection.

3.1. *A heatmap using the cisplatin genes shows no separation of the cell lines [Figure 2(a)].* In order to see how clearly the sensitive and resistant cisplatin cell lines were separated, we extracted and clustered the expression submatrix for the signature genes across all 30 cell lines [Figure 2(a)]. The heatmap shows no clear split between sensitive and resistant cell lines.

3.2. *A heatmap using offset cisplatin genes shows clear separation of the cell lines [Figure 2(b)].* Coombes, Wang and Baggerly (2007) noted that several of the gene lists initially reported by Potti et al. (2006) were “off-by-one” due to an indexing error, so we also extracted and clustered the expression submatrix for the “offset” cisplatin signature genes across all 30 cell lines, producing the heatmap shown in Figure 2(b). This offset involves a single row shift: for example, quantifications from row 98 (probeset 200076_at) of the Györfy et al. (2006) table were used instead of those from row 97 (probeset 200075_s_at). This heatmap shows a clear split between sensitive and resistant cell lines.

3.3. *Clustering correlations suggests the cell lines involved [Figure 2(c)].* Györfy et al. (2006) list 20 and 7 lines as resistant and sensitive to cisplatin, respectively (3 are listed as intermediate), but Hsu et al. (2007) show 9 and 6. In order to figure out which cell lines were used, we clustered the correlations between sample columns from the expression submatrix for the offset genes shown in Figure 2(b). This clustering, shown in Figure 2(c), shows two clear groups of cell lines: 10 in the upper right and 7 in the lower left. We used the sensitive and resistant labels from Györfy et al. (2006) to suggest which group was which (upper

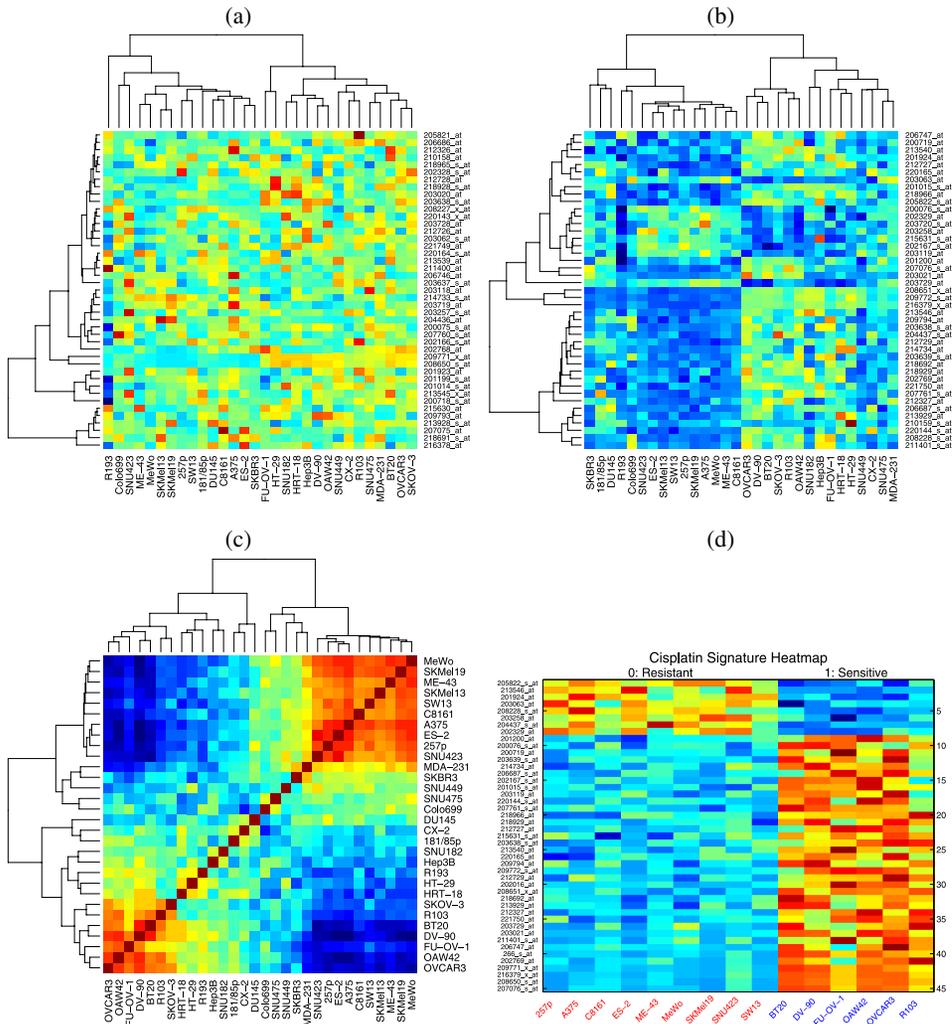


FIG. 2. Images used in reconstructing the Hsu et al. (2007) heatmap for cisplatin. (a) Heatmap of the named cisplatin signature genes across the 30 Györfy et al. (2006) cell lines. There is very little structure visible. (b) Heatmap across the same cell lines using expression values for genes obtained by “offsetting” by one row (e.g., replacing 200075_s_at with 200076_at). There is clear separating structure. (c) Heatmap clustering pairwise sample correlations from the expression submatrix shown in (b). Red clusters (high correlations) in the lower left and upper right suggest initial guesses at the lines that should be treated as “sensitive” and “resistant” respectively. (d) Heatmap produced by binreg using cell lines noted in the text. This is an exact match for the heatmap in Hsu et al. (2007), confirming the cell lines and genes involved.

right is resistant). Checking the labels assigned by Györfy et al. (2006) shows that SKMel13 (Sensitive) and Skov3 (Intermediate) do not fit with the others of their respective groups. Omitting these gives the following 9 “Resistant” cell lines: 257P,

A375, C8161, ES2, me43, MeWo, SKMe119, SNU423 and Sw13; and 6 “Sensitive” cell lines: BT20, DV90, FUOV1, OAW42, OVKAR and R103. We checked this further using a brute force “steepest ascent” method with their `binreg` software as follows. Starting with our initial guess of 10 and 7 lines, we labeled each of the 30 cell lines as resistant, sensitive or not used. We gave our guess a score equal to the number of offset probeset ids matching the `binreg` output (our starting score was 30). We then examined all 60 combinations of cell lines that could be reached by changing the status of one cell line from its current state to resistant, sensitive or not used, and moved to the combination with the highest score (36, obtained by dropping SKOV-3 from the sensitive group), and iterating until a local maximum was reached (41, on the second step, after dropping SKMe1-13 from the resistant group).

3.4. *Applying binreg perfectly reproduces the reported heatmap [Figure 2(d)].* When we apply `binreg` to identify the top 45 genes differentiating the cell lines named above, `binreg` produces the heatmap shown in Figure 2(d). This is an exact match to the cisplatin heatmap reported by Hsu et al. (2007), confirming the cell lines used and identifying the genes involved.

3.5. *The software produces 41/45 offset genes; the others are the ones explicitly mentioned.* We compared the cisplatin gene list `binreg` produced for the matching heatmap with the gene list Hsu et al. (2007) reported. We match 41 of the 45 probesets using the offset gene list. The four probesets we do not match after offsetting (the “outliers”) are 203719_at (ERCC1), 210158_at (ERCC4), 228131_at (ERCC1) and 231971_at (FANCM, associated with DNA Repair), which Hsu et al. (2007) explicitly mention as interesting components of the signature.

3.6. *Two of the “outlier” genes are not on the arrays used.* It is actually vacuous to say we cannot match 228131_at (ERCC1) and 231971_at (FANCM; DNA Repair), as Györfy et al. (2006) report no quantifications for these probesets. Checking the U133 probeset annotation files available from Affymetrix (<http://www.affymetrix.com>) shows that these probesets are on the U133B array platform, not the U133A, so Györfy et al. (2006) did not measure them. [The heatmaps in Figures 2(a) and (b) only have 43 rows.]

3.7. *ERCC1 and ERCC4 have been outliers before.* Coombes, Wang and Baggerly (2007) noted that several gene lists initially reported by Potti et al. (2006) contained genes that `binreg` did not produce. ERCC1 and/or ERCC4 were outliers in the gene lists initially reported for docetaxel, paclitaxel and doxorubicin as well as cisplatin. The other outliers are shown in Table 4.

TABLE 4

“Outlier” probesets in the initially reported gene lists for Potti et al. (2006) (docetaxel, paclitaxel, doxorubicin) and Hsu et al. (2007) (cisplatin). ERCC1 and ERCC4 are common. Probesets for cisplatin marked with an asterisk come from the U133B array platform, not the U133As used. The table does not include 14 outliers in the initial signature for docetaxel that are also found in the list reported by Chang et al. (2003) as being effective separators in the docetaxel test set

U95Av2 PS	Docetaxel	Paclitaxel	Doxorubicin	U133A PS	Cisplatin
114_r_at		MAPT		203719_at	ERCC1
1258_s_at	ERCC4	ERCC4	ERCC4	210158_at	ERCC4
1802_s_at	ERBB2	ERBB2		228131_at	ERCC1*
1847_s_at			BCL2	231971_at	FANCM*
1878_g_at	ERCC1	ERCC1			
1909_at			BCL2		
1910_s_at			BCL2		
2034_s_at			CDKN1B		
33047_at	BCL2L11	BCL2L11			
36519_at		ERCC1			
40567_at	K-ALPHA-1	K-ALPHA-1			

3.8. *Genes are offset, and sensitive/resistant labels are reversed for pemetrexed.* Having matched cisplatin, we then turned to pemetrexed. We used the steepest ascent method described above to identify the cell lines and genes in the pemetrexed signature. After offsetting, we are able to match all 85 genes and to perfectly match the heatmap shown. The 8 “Resistant” cell lines are K-562, MOLT-4, HL-60(TB), MCF7, HCC-2998, HCT-116, NCI-H460 and TK-10, and the 10 “Sensitive” cell lines are SNB-19, HS 578T, MDA-MB-231/ATCC, MDA-MB-435, NCI-H226, M14, MALME-3M, SK-MEL-2, SK-MEL-28 and SN12C. Unfortunately, checking the GI50 data for pemetrexed shows that the sensitive/resistant labels are reversed.

3.9. *Summary.* Poor documentation hid an off-by-one indexing error affecting all genes reported, the inclusion of genes from other sources, including other arrays (the outliers), and a sensitive/resistant label reversal. Our analyses of these signatures are contained in Supplementary File 2 [Baggerly and Coombes (2009b)].

4. Case study 3: Combination therapy. We next examined the approaches used to investigate combination therapy. Potti et al. (2006) mention success in deriving predictions for breast cancer patients treated with a combination of paclitaxel (taxol), 5-fluorouracil, adriamycin and cyclophosphamide (TFAC). Their methods note that the combination predictions were derived from those for the individual drugs using “the theorem for combined probabilities as described by William Feller.” Later, Bonnefoi et al. (2007) provided a “validation” of the combination approach, using it to predict patient response to two alternative therapies:

taxotere (docetaxel) followed by epirubicin (similar to doxorubicin) and taxotere (TET), and fluorouracil, epirubicin and cyclophosphamide (FEC). In terms of documentation and reproducibility, our goals here are to clarify the combination rules and to check predictions for the best single drug. We acquired raw CEL files and array quantifications from GEO (GSE6861; posted quantifications have been revised since we obtained them at the end of 2007). We obtained a missing CEL file (HB131) and a table giving the individual and combination drug predictions from the statistical team in Lausanne which used the predictions to construct ROC curves.

4.1. *Treatment is confounded with run date.* We first checked for gross differences in the array data. Examining high pairwise correlations shows the presence of three clear blocks [Figure 3(a)]. We then extracted run date and lab information from the CEL file headers and plotted the data by run date [Figure 3(b)]; the three clusters correspond to three major run blocks. Different symbols in Figure 3(b) show that the first block contains half of the patients treated with FEC, and that roughly half of the samples run at this time were excluded from the final analysis. The second block contains the second half of the patients treated with FEC. The third block contains all of the patients treated with TET; all of these arrays were run on a different scanner than those from the first two blocks. There is perfect confounding of run date with treatment.

4.2. *Three different combination rules were used.* We then tried to identify the rules used to combine individual drug predictions into a combination score. Letting $P(\cdot)$ indicate probability of sensitivity, the rules used are as follows:

$$P(TFAC) = P(T) + P(F) + P(A) + P(C) - P(T)P(F)P(A)P(C),$$

$$P(TET) = P(ET) = \max[P(E), P(T)],$$

and

$$P(FEC) = \frac{5}{8}[P(F) + P(E) + P(C)] - \frac{1}{4}.$$

Values obtained with the first rule were all greater than 1, so the largest value was set to 1, the smallest to 0, and the others fit by linear interpolation. Values bigger than 1 or less than 0 using the third rule were set to 1 and 0, respectively. These rules are not explicitly stated in the methods; we inferred them either from formulae embedded in Excel files (TFAC) or from exploratory data analysis (TET and FEC). None of these rules are standard. Since all three rules are different, it is not clear what rule was validated.

4.3. *We can't match the accuracy for the best drug/treatment combination [Figure 3(c)].* We also tried to roughly replicate the sensitivity predictions reported for the best performing component: cyclophosphamide in the FEC group. To do this, we first matched the X3P platform probesets Bonnefoi et al. (2007) report

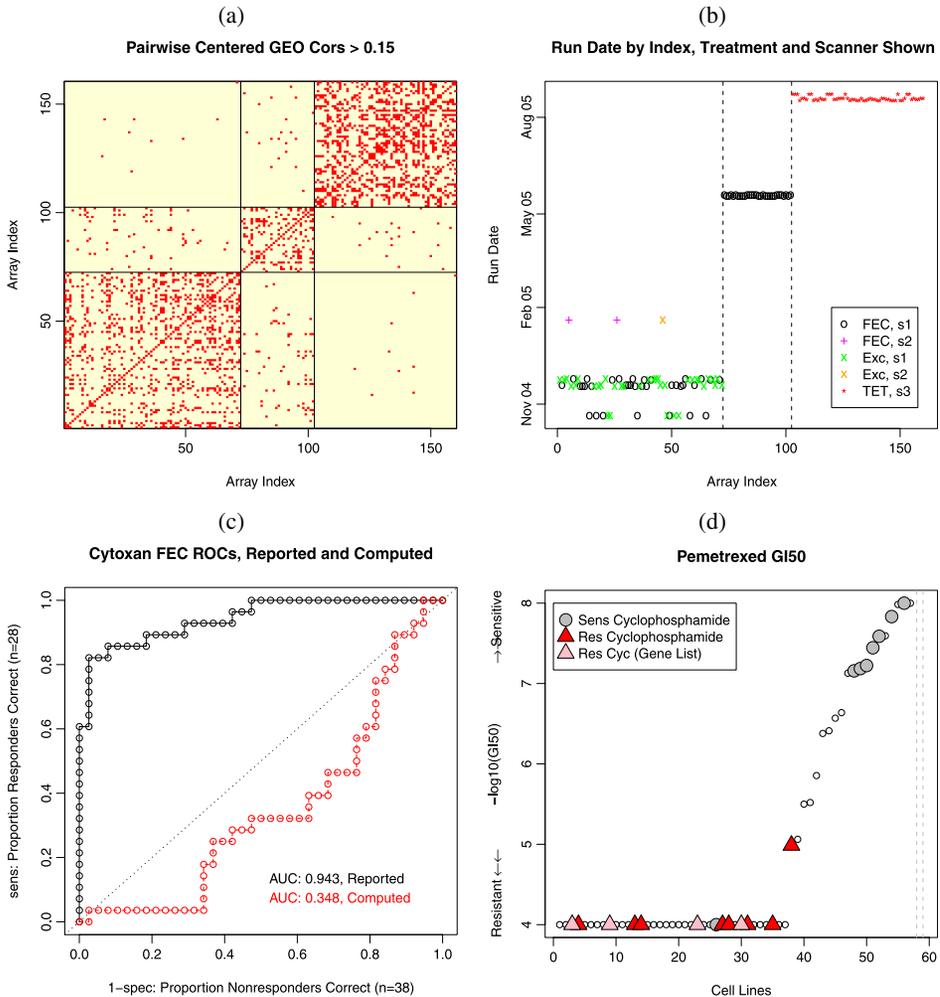


FIG. 3. Aspects of the Bonnefoi et al. (2007) study of combination therapy. (a) Examining array quantifications for high pairwise correlations shows three blocks in the data. (b) Plotting array run date by index shows the same three blocks. The first block contains half of the patients treated with FEC; the other samples in this block were excluded. The second block contains the other half of the patients treated with FEC. The third block contains all of the patients treated with TET, with the arrays run on a different scanner. (c) ROC curves for the single best set of drug predictions: cyclophosphamide for FEC. The reported curve has an AUC of 0.943, indicating extremely good prediction. Our best approximation has an AUC of 0.348, indicating if anything performance worse than chance. (d) Drug sensitivity data for pemetrexed, with the cell lines named for cyclophosphamide indicated. The corresponding plot for cyclophosphamide itself is roughly flat. Pink triangles indicate cell lines that were not reported but are required to generate the gene list reported.

for cyclophosphamide with the U95Av2 platform probesets now available as part of the supplementary information for Potti et al. (2006). This was accomplished first by using “chip comparer” (<http://chipcomparer.genome.duke.edu>), which defined one-to-one mappings for most of the probesets, and using the platform annotation available from either GEO or GeneCards to resolve any ambiguities. We then approximated the `binreg` metagene (SVD) computation in R code, and produced ROC curves using both the reported predictions and those we assembled [Figure 3(c)]. The ROC curve assembled with the reported predictions (AUC of 0.943) perfectly matches that given by Bonnefoi et al. (2007). Our own curve (AUC of 0.348) is qualitatively different.

4.4. *Sensitivity to cyclophosphamide doesn't separate the cell lines used. Sensitivity to pemetrexed does.* We also examined the drug sensitivity data for cyclophosphamide (NSC 26271) to clarify how the cell lines were chosen, but saw no differential activity. This is driven by the fact that cyclophosphamide is a pro-drug (a drug that needs to be processed by the body to produce the active form), and has no direct effect on cell lines. The cell lines used for cyclophosphamide match those found above in the signature for pemetrexed; sensitivity data for pemetrexed is shown in Figure 3(d).

4.5. *Summary.* Design confounding was not mentioned. Poor documentation obscured the fact that different combination rules were used, and leaves both the computation of scores and selection of cell lines for cyclophosphamide unclear. Details of our analyses are given in Supplementary File 3 [Baggerly and Coombes (2009c)].

5. Case study 4: Temozolomide. We next examined the signature for temozolomide reported by Augustine et al. (2009), who discuss how it might be useful for predicting response in melanoma. The initial signature involved 45 genes separating 9 resistant from 6 sensitive cell lines, with cell lines coming from the NCI-60 panel. In terms of documentation and reproducibility, our goal here is to match the results to the drug named. We acquired the heatmap from Augustine et al. (2009)'s Figure 4A and the genes from their supplementary table.

5.1. *The gene list doesn't match the heatmap.* We first tried to confirm the behavior of the individual genes. The gene list indicates that 8 probesets are more strongly expressed in the resistant lines, with the other 37 more strongly expressed in the sensitive lines. However, three genes (RRAGD, SFN and SLC43A3) are listed as higher in both groups (these genes were interrogated by multiple probesets).

5.2. *The heatmap is that published for cisplatin (Figure 4).* We tried to match the heatmap reported following the approach described above for cisplatin, but were unsuccessful. We then visually compared the temozolomide and cisplatin heatmaps (approximated in Figure 4). The heatmaps are identical. Since we inde-

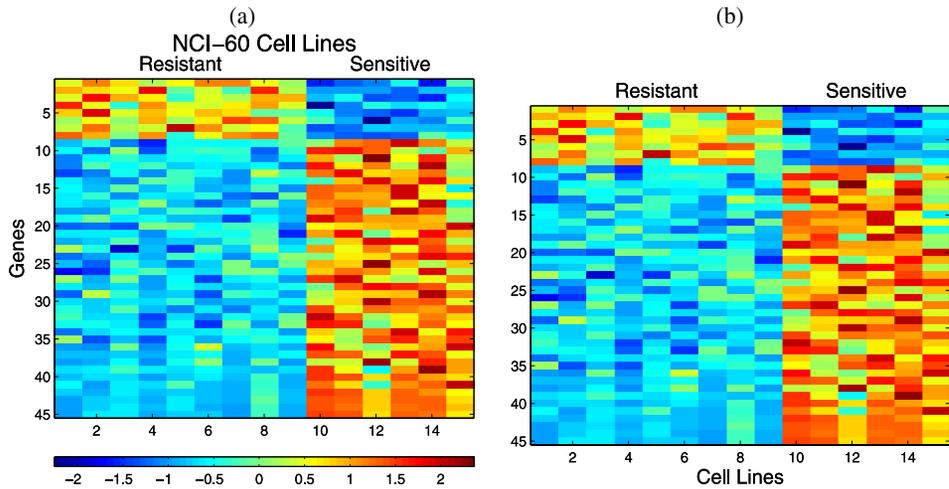


FIG. 4. Approximations to (a) the heatmap initially presented in Figure 4A of Augustine *et al.* (2009) for temozolomide, with lines reportedly chosen from the NCI-60 cell line panel, and (b) the heatmap presented in Figure 1 of Hsu *et al.* (2007) for cisplatin, with cell lines chosen from the 30-line panel of Györfy *et al.* (2006). The heatmaps are the same. We have independently generated the cisplatin heatmap using the Györfy *et al.* (2006) data, but the temozolomide heatmap is neither for temozolomide nor from the NCI-60 panel.

pendently regenerated the cisplatin heatmap using cell lines from Györfy *et al.* (2006), this heatmap does not correspond to temozolomide and was not derived from the NCI-60 cell lines. Since the reported gene list does not match the list for cisplatin (even allowing for offsetting), we presumed that it should correspond to the true list for temozolomide.

5.3. *Journal communication led to a new heatmap with different problems.* In correction, a new figure was supplied [Augustine *et al.* (2009)]. The corrected heatmap involves 150 genes, not 45, so the initial gene list evidently matches neither cisplatin or temozolomide. The revised heatmap shows 5 resistant lines, 5 sensitive lines and 150 probesets, as opposed to 9, 6 and 45 before. The fraction of probesets higher in the resistant group has changed from $8/45$ to about $110/150$. As noted in the Introduction, the initial heatmap was described in detail in the text of Augustine *et al.* (2009), but none of these differences were noted. The new caption names 93 genes as higher in the resistant group. Visual inspection shows about 110.

5.4. *Summary.* Poor documentation led a report on drug A to include a heatmap for drug B and a gene list for drug C. These results are based on simple visual inspection and counting, and are not documented further.

6. Case study 5: Surveying cell lines used. In light of the above issues, we decided to assemble a more extensive overview of which cell lines were treated as

sensitive or resistant for various drugs. We considered 12 sources of information about the signatures for 10 drugs: docetaxel (D), paclitaxel (P), doxorubicin (adriamycin, A), fluorouracil (F), topotecan (T), etoposide (E), cyclophosphamide (C), pemetrexed (Pem), cisplatin (Cis) and temozolomide (Tem). The sources are as follows:

1. Heatmaps from [Potti et al. \(2006\)](#) (D, P, A, F, T, E), November 2006.
2. Heatmaps from [Hsu et al. \(2007\)](#) (Cis, Pem), October 2007.
3. Gene lists from [Potti and Nevins \(2007\)](#) (D, P, A, F, T, E, C), November 2007.*
4. Docetaxel quantifications from the [Potti et al. \(2006\)](#) website (D), November 2007.
5. Doxorubicin quantifications from the [Potti et al. \(2006\)](#) website (A), November 2007.
6. The “list of cell lines used” and the “description of predictor generation” from the [Potti et al. \(2006\)](#) website (D, P, A, F, T, E, C), November 2007.*
7. The “list of cell lines used” supplied as a webpanel supplement to [Bonnefoi et al. \(2007\)](#) (D, A, F, C), December 2007.
8. Numbers of sensitive and resistant lines used by [Salter et al. \(2008\)](#) (P, A, F, C), April 2008.*
9. The “list of cell lines used” posted on the [Potti et al. \(2006\)](#) website (D, P, A, F, T, E, C), August 2008.
10. The “list of cell lines used” by [Bonnefoi et al. \(2007\)](#) after correction (D, A, F, C), September 2008.
11. Numbers of sensitive and resistant lines named by [Riedel et al. \(2008\)](#) (D, P, A, F, T, E, C), October 2008.*
12. Heatmap from [Augustine et al. \(2009\)](#) (Tem), January 2009.

We draw inferences from heatmaps and gene lists when we are able to exactly match the reported results with the `binreg` software used by [Potti et al. \(2006\)](#). This uniquely identifies the two groups being contrasted, but not the direction (which group is sensitive). Direction is inferred from other statements in the relevant papers about what the figures represent. In some cases (*’s above) either the direction or the identity of the cell lines is not precisely specified, so some information must be inferred from other sources. These “other sources” need not be complex; in the case of [Salter et al. \(2008\)](#), for example, we simply examined the heatmaps and counted the numbers of cell lines on the left and right of the major divide.

6.1. *Sensitive/resistant label reversal is common (Figure 5).* The cell line lists and information sources are summarized in Figure 5. Color changes across rows show that there is at least one sensitive/resistant label reversal for every drug checked more than once. The sets of cell lines are different for most drugs, with the exception of cyclophosphamide/pemetrexed noted above. The cell lines reported for cyclophosphamide are a subset of those used for pemetrexed, but running the

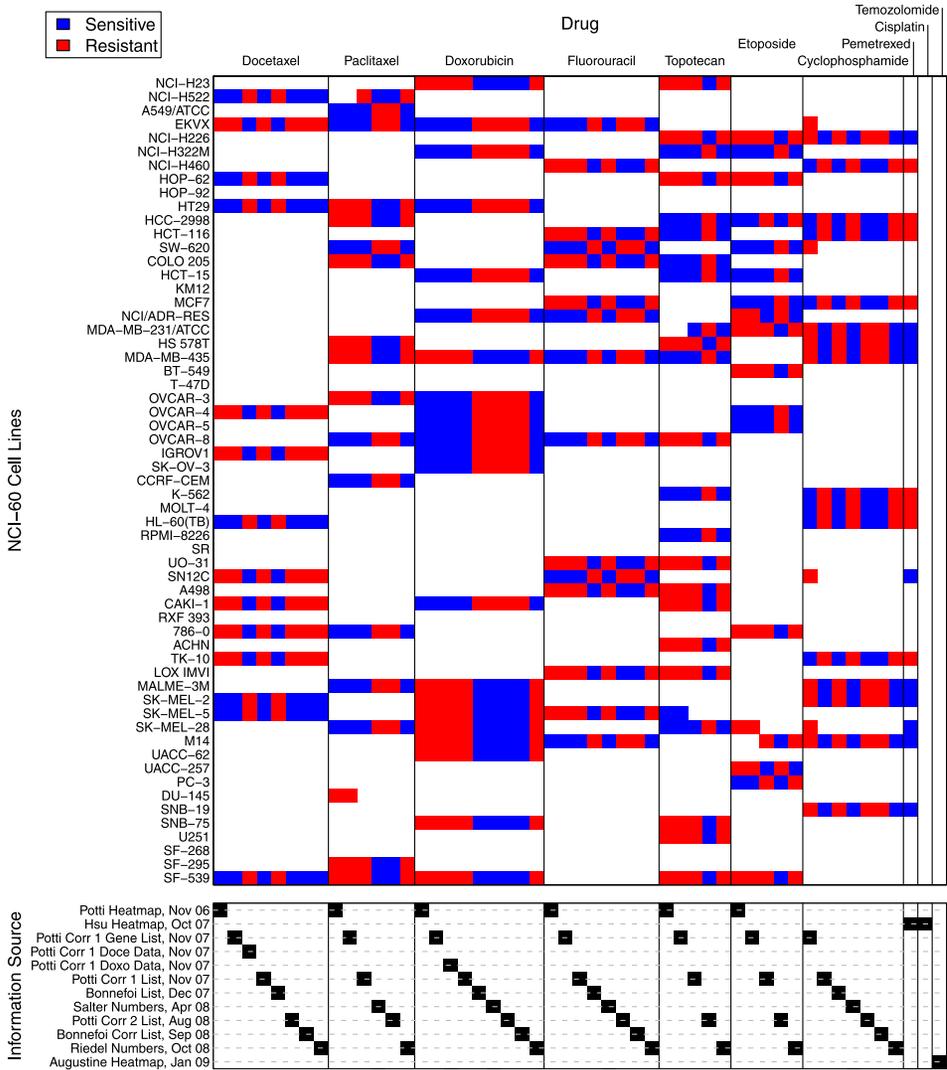


FIG. 5. Mapping of which NCI-60 cell lines were used as sensitive or resistant for which drugs, and the information source used for the inference. For example, there are 8 sources of information (columns) for docetaxel; in column 3, corresponding to the docetaxel quantifications supplied in November 2007, cell line NCI-H522 was listed as resistant. All drug groupings with more than one column show color flips, corresponding to reversal of the sensitive/resistant labeling. Cell lines are different for most drugs save cyclophosphamide/pemetrexed. No cell lines are indicated for cisplatin or temozolomide; the first signature was derived from a different set of cell lines, and in the heatmap reported for temozolomide is actually the heatmap for cisplatin.

cyclophosphamide cell lines through `binreg` does not produce the gene list reported. The cell lines producing the cyclophosphamide gene list are a superset of those used for pemetrexed. The cisplatin signature is based on 30 cell lines assembled by Györfy et al. (2006); the heatmap reported for temozolomide matches that for cisplatin.

Based on the drug sensitivity information, we believe the orientations given in the August 2008 lists of cell lines (Potti Corr 2 list) are correct. Assuming this is the case, Figure 5 shows the sensitive/resistant orientations of the Salter et al. (2008) heatmaps are correct for T and A but incorrect for F and C. Heatmap orientations in Potti et al. (2006) are reversed for T, A and F, and we cannot reproduce the heatmap for C. However, sample predictions shown in both Figure 3A of Potti et al. (2006) and Figure 2B of Salter et al. (2008) suggest results better than even for all four drugs [Potti et al. (2006) p -values: T = 0.002, F = 0.3, A = 0.024, C = 0.003; Salter et al. (2008) p -values: T = 0.07, F = 0.02, A = 0.01, C = 0.02].

6.2. *Summary.* Poor documentation hides the fact that sensitive and resistant labels are being used inconsistently over time, even though this direction determines whether the drug should be offered or withheld. Full details of the figure assembly are given in Supplementary File 4 [Baggerly and Coombes (2009d)].

7. Discussion.

7.1. *On the nature of common errors.* In all of the case studies examined above, forensic reconstruction identifies errors that are hidden by poor documentation. Unfortunately, these case studies are illustrative, not exhaustive; further problems similar to the ones detailed above are described in the supplementary reports. The case studies also share other commonalities. In particular, they illustrate that *the most common problems are simple*: for example, confounding in the experimental design (all TET before all FEC), mixing up the gene labels (off-by-one errors), and mixing up the group labels (sensitive/resistant); most of these mixups involve simple switches or offsets. These mistakes are easy to make, particularly if working with Excel or if working with 0/1 labels instead of names (as with `binreg`). We have encountered these and like problems before. As part of the 2002 Competitive Analysis of Microarray Data (CAMDA) competition, Stivers et al. (2003) identified and corrected a mixup in annotation affecting roughly a third of the data which was driven by a simple one-cell deletion from an Excel file coupled with an inappropriate shifting up of all values in the affected column only. Baggerly, Morris and Coombes (2004), Baggerly et al. (2004) and Baggerly et al. (2005) describe cases of complete confounding leading to optimistic predictions for proteomic experiments. Baggerly, Coombes and Neeley (2008) describe another array study where there was a mixup in attaching sample labels to columns of quantifications, most likely driven by omission of 3 CEL files leading to an off-by-three error affecting most of the names. These experiences and others make

us worry about the dual of the italicized statement above, that *the most simple problems may be common*.

7.2. On the hiding of simplicity. While simple mistakes often allow for simple fixes, incomplete documentation and lack of reproducibility means that this simplicity is often hidden. Further, it means that identifying the problems often requires going across several sources of information, papers and journals, making simple fixes difficult. Identifying the training data cell lines in `Adria_ALL.txt` used the fact that they were in the same order as in the [Potti et al. \(2006\)](#) heatmaps for correlation matching. Identifying cell lines for cisplatin used knowledge of the possibility of an off-by-one error. Identifying the overlap of the cell lines used for cyclophosphamide and pemetrexed required identifying the cell lines and comparing lists across papers. Identifying the initial heatmap for temozolomide required familiarity with the one for cisplatin. This cascade can lead to a certain archaeology with respect to what has been established “as previously shown.” This can also render some conclusions unfalsifiable, for example, with the rejoinder “they get our results when they use our methods” provided without details [[Potti and Nevins \(2007\)](#); see Supplementary File 5, [Baggerly and Coombes \(2009e\)](#)].

7.3. What should be done. So, what can be done? We address this question in two parts: with respect to the specific findings illustrated here, and with respect to reproducibility in general.

7.3.1. In the case of these results. In the case of these results, we don’t think the approach works. We think stronger evidence (including worked examples of how it works) need to be provided before this approach is used to guide patient allocation in clinical trials. In the case of the clinical trial noted above, assuming the general approach works means that sensitive/resistant label reversal for one of the drugs (pemetrexed) may actually put patients at risk by giving guidance at odds with the truth, whereas assuming the general approach doesn’t work means that little will be learned from the trial or even that it may provide misleading support for the approach given the inclusion of genes that shouldn’t be there (ERCC1, ERCC4).

A broader question is whether this approach could work if applied correctly. We don’t think so. We have tried making predictions from the NCI60 cell lines when we step through the process without the errors noted above, and we get results no better than chance. We communicated with the authors extensively in the early phases of our investigation (which we recommend), and shared the fact that we didn’t think it worked before submitting our initial note [[Coombes, Wang and Baggerly \(2007\)](#)]. Empirically, however, we have reached an impasse, as the progression to clinical trials attests.

7.3.2. In the case of reproducibility in general. In the case of reproducibility in general, journals and funding agencies already require that raw data (e.g., CEL files) be made available. We see it as unavoidable that complete scripts will also

eventually be required. General guidelines outlining the types of questions authors should be prepared to answer in detail are discussed in the REporting recommendations for tumor MARKer prognostic studies (REMARK) guidelines [McShane et al. (2005)].

7.4. What we're doing. Partially in response to the examples discussed here, we instituted new operating procedures within our group, mostly simple things having to do with report structure. Reports in our group are typically produced by teams of statistical analysts and faculty members, and issued to our biological collaborators. We now require most of our reports to be written using *Sweave* [Leisch (2002)], a literate programming combination of L^AT_EX source and R [R Development Core Team (2008)] code (SASweave and odfWeave are also available) so that we can rerun the reports as needed and get the same results. Some typical reports are shown in the supplementary material. Most of these reports are written by the statistical analysts, and read over (and in some cases rerun) by the faculty members. All reports include an explicit call to `sessionInfo` to list libraries and versions used in the analysis. The working directory and the location of the raw data are also explicitly specified (in some cases leading to raw data being moved from personal machines to shared drives). We also check for the most common types of errors, which are frequently introduced by some severing of data from its associated annotation [e.g., using 0 or 1 for sensitive or resistant instead of using names (noted above), supplying one matrix of data and another of annotation without an explicit joining feature, calls to order one column of a set]. R's ability to let us use row and column labels which it maintains through various data transformations helps. These steps have improved reproducibility markedly. We have also assumed a fairly regular report structure in order to enhance clarity, including an executive summary (at most two pages of text) detailing Background, Methods, Results and Conclusions in the format most familiar to our biological collaborators. Discussing the Background and Methods with the collaborators ahead of time helps ensure we're working on the problems of actual interest. Finally, we are trying to shift more frequently to the use of standardized templates for common analyses (e.g., two group comparisons for microarray studies). There is definitely a startup cost in time when shifting to this approach, but this is often made up when we have to modify earlier analyses months or years later.

7.5. The bottom line. While it may be difficult, we think the examples above show that this added layer of documentation is required for high-throughput biology.

Acknowledgment. We thank Aron Eklund for comments about the nature of prodrugs.

SUPPLEMENTARY MATERIAL

Supplement A: Examining doxorubicin in detail (DOI: [10.1214/09-AOAS291SUPPA](https://doi.org/10.1214/09-AOAS291SUPPA); .zip). Zipped pdf report describing the identification of ties and sensitive/resistant status for samples checked for doxorubicin.

Supplement B: Cisplatin and pemetrexed (DOI: [10.1214/09-AOAS291SUPPB](https://doi.org/10.1214/09-AOAS291SUPPB); .zip). Zip file containing two pdf reports, matchingCisplatin-Heatmap and matchingPemetrexedHeatmap, describing the matching of the respective heatmaps together with sample and gene identification.

Supplement C: Examining combination therapy (DOI: [10.1214/09-AOAS291SUPPC](https://doi.org/10.1214/09-AOAS291SUPPC); .zip). Zip file containing seven pdf reports: getTestingNumbers, getTestingClinical (for identification of blocks and confounding), checkOldCombinationRule, checkDrugSensitivity (for the combination rules), mapGeneLists and predictCytosanSensitivity (for the ROC curves), and checkingCellLines (for the drug sensitivity values).

Supplement D: Surveying cell lines (DOI: [10.1214/09-AOAS291SUPPD](https://doi.org/10.1214/09-AOAS291SUPPD); .zip). Zip file containing two pdf reports: enumeratingCellLines, describing assembly of the figure, and getTrainingCellLines from the study of combination therapy, for identification of the cell lines needed to produce the gene list for cyclophosphamide.

Supplement E: Examining docetaxel in detail (DOI: [10.1214/09-AOAS291SUPPE](https://doi.org/10.1214/09-AOAS291SUPPE); .zip). Zipped pdf report describing the identification of ties and sensitive/resistant status for samples checked for docetaxel.

All reports are in *Sweave*. More reports, data and code, are at <http://bioinformatics.mdanderson.org/Supplements/ReproRsch-All>. Reports concerning combination therapy are at [ReproRsch-Breast](http://bioinformatics.mdanderson.org/Supplements/ReproRsch-Breast).

REFERENCES

- AUGUSTINE, C. K., YOO, J. S., POTTI, A., YOSHIMOTO, Y., ZIPFEL, P. A., FRIEDMAN, H. S., NEVINS, J. R., ALI-OSMAN, F. and TYLER, D. S. (2009). Genomic and molecular profiling predicts response to temozolomide in melanoma. *Clin. Cancer Res.* **15** 502–510. Correction p. 3240.
- BAGGERLY, K. A., COOMBES, K. R. and NEELEY, E. S. (2008). Run batch effects potentially compromise the usefulness of genomic signatures for ovarian cancer. *J. Clin. Oncol.* **26** 1186–1187.
- BAGGERLY, K. A. and COOMBES, K. R. (2009a). Supplement 1, “Examining doxorubicin in detail,” to “Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology.” DOI: [10.1214/09-AOAS291SUPPA](https://doi.org/10.1214/09-AOAS291SUPPA).
- BAGGERLY, K. A. and COOMBES, K. R. (2009b). Supplement 2, “Cisplatin and pemetrexed,” to “Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology.” DOI: [10.1214/09-AOAS291SUPPB](https://doi.org/10.1214/09-AOAS291SUPPB).
- BAGGERLY, K. A. and COOMBES, K. R. (2009c). Supplement 3, “Examining combination therapy,” to “Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology.” DOI: [10.1214/09-AOAS291SUPPC](https://doi.org/10.1214/09-AOAS291SUPPC).

- BAGGERLY, K. A. and COOMBES, K. R. (2009d). Supplement 4, "Surveying cell lines," to "Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology." DOI: [10.1214/09-AOAS291SUPPD](https://doi.org/10.1214/09-AOAS291SUPPD).
- BAGGERLY, K. A. and COOMBES, K. R. (2009e). Supplement 5, "Examining docetaxel in detail," to "Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology." DOI: [10.1214/09-AOAS291SUPPE](https://doi.org/10.1214/09-AOAS291SUPPE).
- BAGGERLY, K. A., EDMONSON, S. R., MORRIS, J. S. and COOMBES, K. R. (2004). High-resolution serum proteomic patterns for ovarian cancer detection. *Endocr. Relat. Cancer* **11** 583–584.
- BAGGERLY, K. A., MORRIS, J. S. and COOMBES, K. R. (2004). Reproducibility of SELDI-TOF protein patterns in serum: Comparing datasets from different experiments. *Bioinformatics* **20** 777–785.
- BAGGERLY, K. A., MORRIS, J. S., EDMONSON, S. R. and COOMBES, K. R. (2005). Signal in noise: Evaluating reported reproducibility of serum proteomic tests for ovarian cancer. *J. Natl. Cancer Inst.* **97** 307–309.
- BONNEFOI, H., POTTI, A., DELORENZI, M., MAURIAC, L., CAMPONE, M., TUBIANA-HULIN, M., PETIT, T., ROUANET, P., JASSEM, J., BLOT, E., BECETTE, V., FARMER, P., ANDRE, S., ACHARYA, C. R., MUKHERJEE, S., CAMERON, D., BERGH, J., NEVINS, J. R. and IGGO, R. D. (2007). Validation of gene signatures that predict the response of breast cancer to neoadjuvant chemotherapy: A substudy of the EORTC 10994/BIG 00-01 clinical trial. *Lancet Oncol.* **8** 1071–1078.
- CHANG, J. C., WOOTEN, E. C., TSIMELZON, A., HILSENBECK, S. G., GUTIERREZ, M. C., ELLEDGE, R., MOHSIN, S., OSBORNE, C. K., CHAMNESS, G. C., ALLRED, D. C. and O'CONNELL, P. (2003). Gene expression profiling for the prediction of therapeutic response to docetaxel in patients with breast cancer. *Lancet* **362** 362–369.
- CLINICAL TRIAL NCT00509366 (2009). Study using a genomic predictor of platinum resistance to guide therapy in stage IIIB/IV non-small cell lung cancer (TOP0602). Available at <http://clinicaltrials.gov/ct2/show/NCT00509366>.
- COOMBES, K. R., WANG, J. and BAGGERLY, K. A. (2007). Microarrays: Retracing steps. *Nat. Med.* **13** 1276–1277.
- DISCOVER (2007). The top 6 genetics stories of 2006. *Discover* **January**.
- GENTLEMAN, R. (2005). Reproducible research: A bioinformatics case study. *Stat. Appl. Genet. Mol. Biol.* **4** Article 2. [MR2138207](https://doi.org/10.1214/05-STAT)
- GENTLEMAN, R. and TEMPLE LANG, D. (2007). Statistical analyses and reproducible research. *J. Comput. Graph. Statist.* **16** 1–23. [MR2345745](https://doi.org/10.1198/00036810600000001)
- GYÖRFFY, B., SUROWIAK, P., KIESSLICH, O., DENKERT, C., SCHAFER, R., DIETEL, M. and LAGE, H. (2006). Gene expression profiling of 30 cancer cell lines predicts resistance towards 11 anticancer drugs at clinically achieved concentrations. *Int. J. Cancer* **118** 1699–1712.
- HOLLEMAN, A., CHEOK, M. H., DEN BOER, M. L., YANG, W., VEERMAN, A. J., KAZEMIER, K. M., PEI, D., CHENG, C., PUI, C. H., RELING, M. V., JANKA-SCHAUB, G. E., PIETERS, R. and EVANS, W. E. (2004). Gene-expression patterns in drug-resistant acute lymphoblastic leukemia cells and response to treatment. *N. Engl. J. Med.* **351** 533–542.
- HSU, D. S., BALAKUMARAN, B. S., ACHARYA, C. R., VLAHOVIC, V., WALTERS, K. S., GARMAN, K., ANDERS, C., RIEDEL, R. F., LANCASTER, J., HARPOLE, D., DRESSMAN, H. K., NEVINS, J. R., FEBBO, P. G. and POTTI, A. (2007). Pharmacogenomic strategies provide a rational approach to the treatment of cisplatin-resistant patients with advanced cancer. *J. Clin. Oncol.* **25** 4350–4357.
- IOANNIDIS, J. P., ALLISON, D. B., BALL, C. A., COULIBALY, I., CUI, X., CULHANE, A. C., FALCHI, M., FURLANIELLO, C., GAME, L., JURMAN, G., MANGION, J., MEHTA, T., NITZBERG, M., PAGE, G. P., PETRETTO, E. and VAN NOORT, V. (2009). Repeatability of published microarray gene expression analyses. *Nat. Genet.* **41** 149–155.

- LEISCH, F. (2002). Dynamic generation of statistical reports using literate data analysis. In *Compstat 2002—Proceedings in Computational Statistics* (W. Härdle and B. Rönz, eds.) 575–580. Physika Verlag, Heidelberg, Germany.
- LI, C. (2008). Automating dChip: Toward reproducible sharing of microarray data analysis. *BMC Bioinformatics* **9** 231.
- LUGTHART, S., CHEOK, M. H., DEN BOER, M. L., YANG, W., HOLLEMAN, A., CHENG, C., PUI, C. H., RELLING, M. V., JANKA-SCHAUB, G. E., PIETERS, R. and EVANS, W. E. (2005). Identification of genes associated with chemotherapy crossresistance and treatment response in childhood acute lymphoblastic leukemia. *Cancer Cell* **7** 375–386.
- MC SHANE, L. M., ALTMAN, D. G., SAUERBREI, W., TAUBE, S. E., GION, M. and CLARK, G. M. (2005). Reporting recommendations for tumor marker prognostic studies (REMARK). *J. Natl. Cancer Inst.* **97** 1180–1184.
- POTTI, A. and NEVINS, J. (2007). Reply to Microarrays: Retracing steps. *Nat. Med.* **13** 1277–1278.
- POTTI, A., DRESSMAN, H. K., BILD, A., RIEDEL, R. F., CHAN, G., SAYER, R., CRAGUN, J., COTTRILL, H., KELLEY, M. J., PETERSEN, R., HARPOLE, D., MARKS, J., BERCHUCK, A., GINSBURG, G. S., FEBBO, P., LANCASTER, J. and NEVINS, J. R. (2006). Genomic signatures to guide the use of chemotherapeutics. *Nat. Med.* **12** 1294–1300.
- POTTI, A., DRESSMAN, H. K., BILD, A., RIEDEL, R. F., CHAN, G., SAYER, R., CRAGUN, J., COTTRILL, H., KELLEY, M. J., PETERSEN, R., HARPOLE, D., MARKS, J., BERCHUCK, A., GINSBURG, G. S., FEBBO, P., LANCASTER, J. and NEVINS, J. (2008). Corrigendum to “Genomics signatures to guide the use of chemotherapeutics.” *Nat. Med.* **14** 889.
- POTTI, A., NEVINS, J. R. and LANCASTER, J. M. (2009). Predicting responsiveness to cancer therapeutics. U.S. Patent Application 20090105167. Available at <http://www.freepatentsonline.com/y2009/0105167.html>.
- R DEVELOPMENT CORE TEAM (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available at <http://www.R-project.org>.
- RIEDEL, R. F., PORRELLO, A., PONTZER, E., CHENETTE, E. J., HSU, D. S., BALAKUMARAN, B., POTTI, A., NEVINS, J. and FEBBO, P. G. (2008). A genomic approach to identify molecular pathways associated with chemotherapy resistance. *Mol. Cancer Ther.* **7** 3141–3149.
- RUSCHHAUPT, M., HUBER, W., POUSTKA, A. and MANSMANN, U. (2004). A compendium to ensure computational reproducibility in high-dimensional classification tasks. *Stat. Appl. Genet. Mol. Biol.* **3** 37. MR2101484
- SALTER, K. H., ACHARYA, C. R., WALTERS, K. S., REDMAN, R., ANGUIANO, A., GARMAN, K. S., ANDERS, C. K., MUKHERJEE, S., DRESSMAN, H. K., BARRY, W. T., MARCOM, K. P., OLSON, J., NEVINS, J. R. and POTTI, A. (2008). An integrated approach to the prediction of chemotherapeutic response in patients with breast cancer. *PLoS One* **3** e1908.
- STIVERS, D. N., WANG, J., ROSNER, G. L. and COOMBES, K. R. (2003). Organ-specific differences in gene expression and UniGene annotations describing source material. In *Methods of Microarray Data Analysis III* (K. Johnson and S. Lin, eds.) 59–72. Kluwer Academic, Boston.

DEPARTMENT OF BIOINFORMATICS
AND COMPUTATIONAL BIOLOGY
M. D. ANDERSON CANCER CENTER
UNIVERSITY OF TEXAS
1400 PRESSLER
HOUSTON, TEXAS 77030
USA
E-MAIL: kabagg@mdanderson.org
krcoombes@mdanderson.org