

METHOD

Open Access



# A novel independence test for somatic alterations in cancer shows that biology drives mutual exclusivity but chance explains most co-occurrence

Sander Canisius<sup>1</sup>, John W. M. Martens<sup>2</sup> and Lodewyk F. A. Wessels<sup>1,3,4\*</sup>

## Abstract

In cancer, mutually exclusive or co-occurring somatic alterations across genes can suggest functional interactions. Existing tests for such patterns make the unrealistic assumption of identical gene alteration probabilities across tumors. We present Discrete Independence Statistic Controlling for Observations with Varying Event Rates (DISCOVER), a novel test that is more sensitive than other methods and controls its false positive rate. A pan-cancer analysis using DISCOVER finds no evidence for widespread co-occurrence, and most co-occurrences previously detected do not exceed expectation by chance. Many mutual exclusivities are identified involving well-known genes related to cell cycle and growth factor signaling, as well as lesser known regulators of Hedgehog signaling.

**Keywords:** Mutual exclusivity, Co-occurrence, Computational biology

## Background

Tumor development emerges from a gradual accumulation of somatic alterations that together enable malignant growth. As has been revealed by recent genomic profiling efforts, an immense diversity exists in the alterations that tumors acquire [1, 2]. Whether by e.g., copy number aberration, point mutation, or DNA methylation, alterations of many genes may potentially trigger transformation. Often though, the fate of a cell acquiring a certain alteration depends on other alterations already present [3]. Therefore, with an ever-expanding catalog of cancer genes, a need arises to establish how alterations in those genes interact to transform healthy cells to cancer cells. This task can be approached by statistical analyses aiming to uncover more complex, combinatorial patterns in somatic alterations.

Two such patterns are co-occurrence and mutual exclusivity. In the former, alterations of certain combinations of genes tend to co-exist in the same tumor, whereas

in the latter, mostly only one out of a group of genes is altered in a single tumor. Mutual exclusivity is frequently observed in cancer genomics data [4, 5]. Individual alterations targeting similar biological processes are believed to be mutually redundant, with one alteration being sufficient to deregulate the affected process. Identifying mutual exclusivity can therefore help in finding unknown functional interactions. With this in mind, several statistical methods have been proposed to identify significant patterns of mutual exclusivity [6–12].

Just as mutual exclusivity is interpreted as a sign of redundancy, co-occurrence is often held to entail synergy. Alteration of only one of the two genes would be relatively harmless, whereas cells with alterations in both progress to malignancy. If such synergy exists, cancer genomes should be enriched for these co-alterations; i.e., tumors harboring alterations in both genes should be more frequent than expected by chance. Several studies have reported an abundance of co-occurring somatic alterations in various types of cancer [13–19]. For somatic copy number changes, however, it has also been suggested that co-occurring alterations emerge from tumors' overall levels of genomic disruption [20]. Indeed, tumors display a wide diversity in genomic instability, both across

\*Correspondence: l.wessels@nki.nl

<sup>1</sup>Department of Molecular Carcinogenesis, The Netherlands Cancer Institute, Plesmanlaan 121, 1066 CX Amsterdam, The Netherlands

<sup>3</sup>Faculty of EEMCS, Delft University of Technology, Delft, The Netherlands  
Full list of author information is available at the end of the article

and within cancer types. In tumors harboring many alterations, one should not be surprised to see simultaneous alterations in any pair of genes. In contrast, two genes altered in a tumor carrying a small number of alterations might instead have resulted from a purifying selective process. Suggesting synergy as an explanation for observed co-occurrence is only reasonable if a simpler explanation like tumor-specific alteration rates can be rejected.

In this paper, we address the statistical implications of heterogeneous alteration rates across tumors for co-occurrence and mutual exclusivity detection. With extensive analyses of simulated data, we show how commonly used statistical tests are not equipped to deal with the mismatch between what is assumed by the test and what is encountered in the data. In the presence of heterogeneous alteration rates, countless spurious co-occurrences are picked up in data that are controlled not to contain any. At the same time, many instances of true mutual exclusivity are missed. Based on these observations, we introduce DISCOVER, a novel statistical independence test that incorporates the overall alteration rates of tumors to successfully solve the issues encountered with existing tests. We compared the performance of DISCOVER to that of several other published mutual exclusivity tests: *MEMo* [6], *muex* [8], *mutex* [9], *CoMEt* [10], *MEGSA* [11], and *TiMEx* [12]. Across the whole range of significance levels, DISCOVER is more sensitive while controlling the false positive rate at the specified level.

We also applied DISCOVER to a selection of more than 3000 tumors across 12 different cancer types. Only one co-occurrence was detected that is not explained by overall rates of alteration alone. On the other hand, many more cases of mutual exclusivity were detected than would have been possible with traditional tests. The genes targeted by these alterations cover many of the core cancer pathways known to display such exclusivity. However, we also identified exclusivity among less canonical actors in the cell cycle, and among regulators of Hedgehog signaling.

## Results

### Common tests for co-occurrence or mutual exclusivity assume homogeneous alteration rates

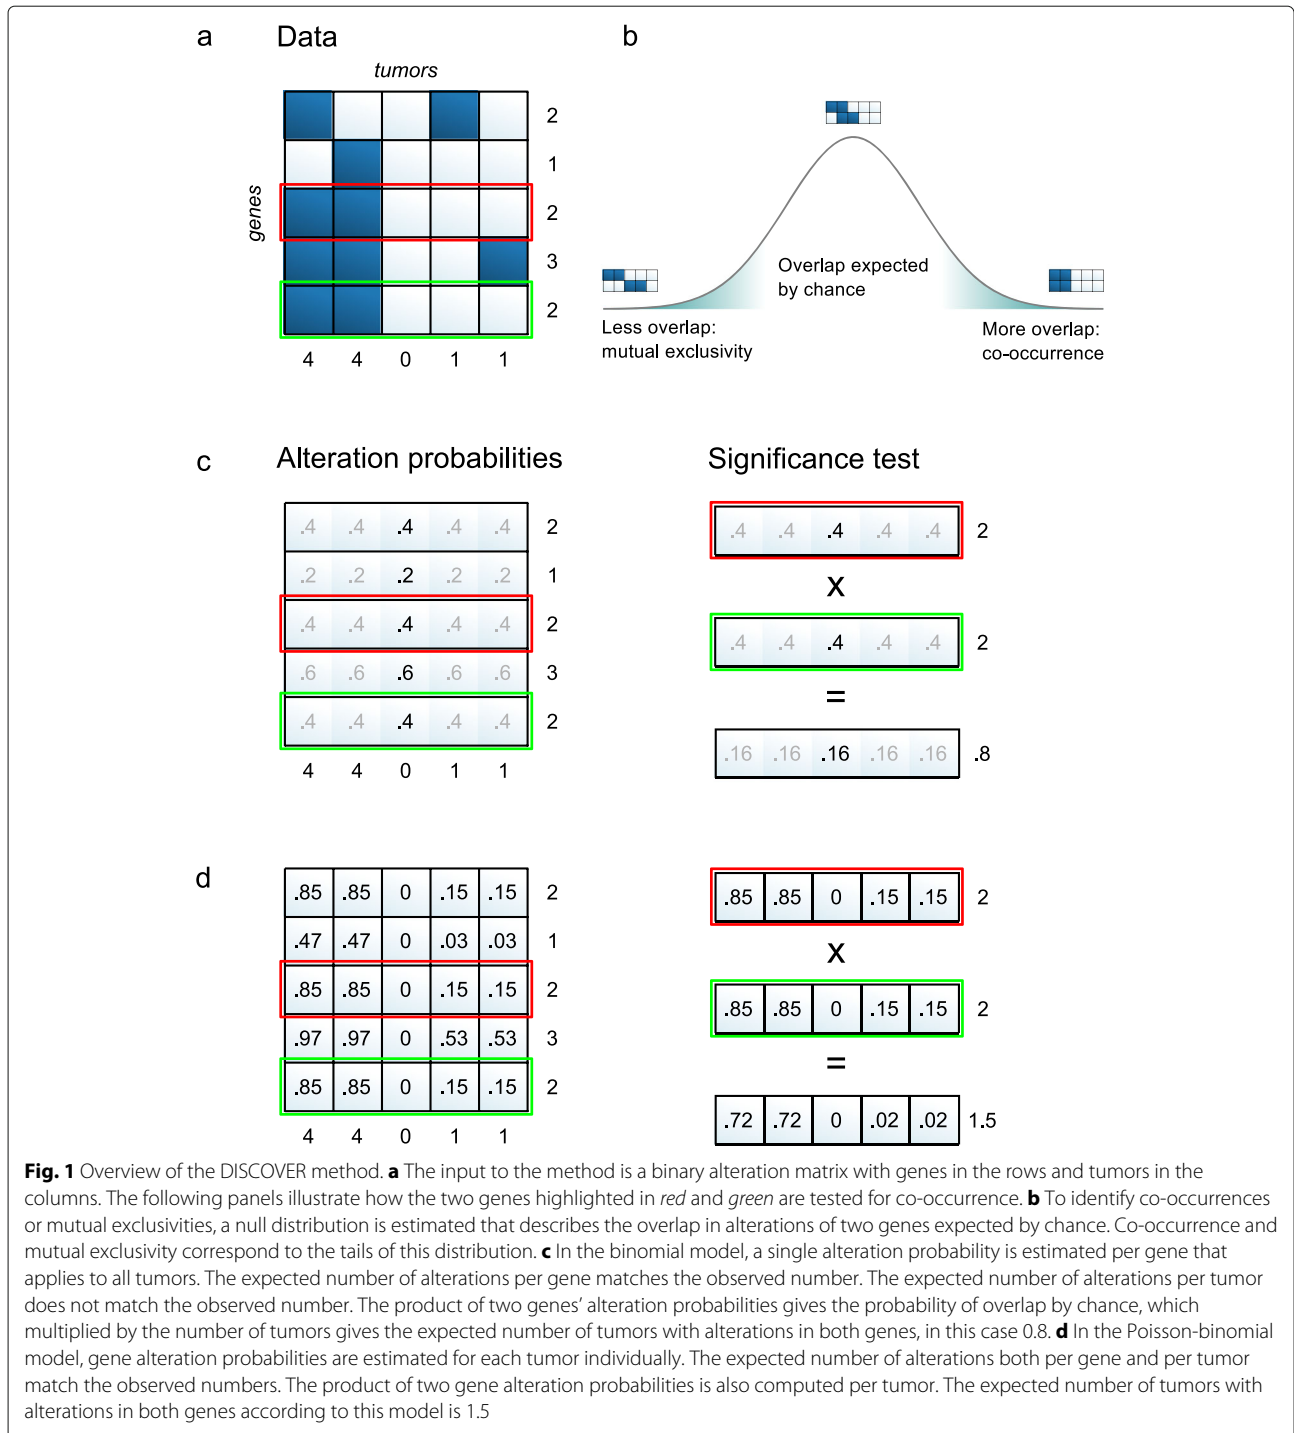
A commonly used test for both co-occurrence and mutual exclusivity is Fisher's exact test applied to a  $2 \times 2$  contingency table [16–18]. The test is used to support co-occurrence when the number of tumors with alterations in both genes is significantly higher than expected by chance. Likewise, it suggests mutual exclusivity when the number of tumors with alterations in both genes is significantly lower. The validity of this test depends on the assumption that genes' alterations across tumors are independent and identically distributed (i.i.d.). Identical distribution implies that the probability of an alteration in a gene is

the same for any given tumor. With cancer's heterogeneity in mind, this assumption may prove problematic. Surely, a gene is more likely found altered in tumors with many somatic alterations overall, than in tumors with only few such changes.

Other tests used for co-occurrence or mutual exclusivity depend on the same i.i.d. assumption as described for Fisher's exact test. This is the case for permutation tests that estimate the expected number of tumors altered in both genes by randomly reassigning gene alterations across tumors [7, 13]. It is also true for a simple binomial test that we will use to illustrate the consequences of violating the i.i.d. assumption. This test is depicted in Fig. 1c. The alteration probability  $p_i$  of a gene is estimated to be the proportion of tumors altered in that gene. For example, gene 3 in Fig. 1a is altered in 2 of the 5 tumors, resulting in  $p_3 = 0.4$  (Fig. 1c). If alterations targeting two genes are independent, the probability of a tumor altered in both genes equals the product  $p_1 \cdot p_2$  of those genes' alteration probabilities. Hence, out of  $m$  tumors,  $m \cdot p_1 p_2$  tumors are expected to harbor alterations in both genes. In the example in Fig. 1a, the probability of alterations in both genes 3 and 5 would be  $p_3 \cdot p_5 = 0.4 \cdot 0.4 = 0.16$ . Therefore, if alterations of genes 3 and 5 were independent, we would expect  $5 \cdot 0.16 = 0.8$  tumors with alterations in both. Observing more such tumors suggests co-occurrence, whereas observing fewer suggests mutual exclusivity (Fig. 1b).


### Assuming homogeneous alteration rates leads to invalid significance estimates

To illustrate the effect of the i.i.d. assumption on the detection of mutual exclusivities and co-occurrences, we performed analyses on simulated data. Genomic alterations were generated such that the alteration frequencies both per gene and per tumor resemble those observed in real tumors, but without any designed relation between the genes' alterations; i.e., genes were simulated to be independent. As these simulated data do not contain co-occurrences or mutual exclusivities, all identified departures from independence are by definition spurious. We can therefore use these data to check the validity of the binomial test. When testing many pairs of independently altered genes, a valid statistical test should produce  $P$  values that approximately follow a uniform distribution. In contrast, when we test for co-occurrence in these data, the  $P$ -value distribution shows a large skew towards extremely low values (Fig. 2a). Even highly conservative significance levels will mark the majority of gene pairs as significant hits. Given that no true co-occurrences exist in the simulated data, all these hits are false positives. If we test for mutual exclusivities instead, we observe a skew towards the high end of the  $P$ -value spectrum (Fig. 2c).



We next evaluated the sensitivity of the binomial test. For this, we tested simulated co-occurrences and mutual exclusivities, which we added to the data. A sensitive test should produce only low *P* values for these positive cases, and so the resulting *P*-value distribution should be heavily skewed towards zero. If we test for co-occurrences, this

is indeed the case (Fig. 2b). Testing for mutual exclusivity, however, reveals a distribution that, although skewed towards lower *P* values, is much more stretched out across the [0, 1] interval (Fig. 2d). Even highly liberal significance levels will only recover a small part of the positive cases.



We conclude that the binomial test is anti-conservative as a co-occurrence test. In contrast, as a mutual exclusivity test, it is conservative. While we used the binomial test for this illustration, we found the same to be true for Fisher's exact test (Additional file 1: Figure S1). To confirm our hypothesis that the i.i.d. assumption is causal to this incorrect behavior, we generated additional simulated data, making sure that the overall alteration rate was similar across the tumors. Using the binomial test to detect co-occurrence and mutual exclusivity of independent genes results in  $P$ -value distributions that are much closer to uniform (Additional file 1: Figure S2). This confirms that statistical tests that rely on the i.i.d. assumption are not suited for co-occurrence analysis, and have reduced sensitivity for mutual exclusivity analysis.

#### A novel statistical test for co-occurrence and mutual exclusivity

Our new method, which we call Discrete Independence Statistic Controlling for Observations with Varying Event Rates (DISCOVER), is a statistical independence test that does not assume identically distributed events. The main ingredients of the method are depicted in Fig. 1d. Unlike the method in the simpler binomial test, we allow different tumors to have different alteration probabilities for the same gene—the alteration probabilities for genes 3 and 5 in Fig. 1d now vary per tumor, in contrast to Fig. 1c. For tumors with many altered genes, this probability is higher than for tumors with only few alterations. To estimate these alteration probabilities, we solve a constrained optimization problem that ensures that the probabilities are consistent with both the observed number of alterations per gene and the observed number of alterations per tumor. The probability of concurrent alterations in two independent genes is then obtained for each tumor individually, by multiplying the tumor-specific gene alteration probabilities, as indicated in the right panel of Fig. 1d. With these probabilities, an analytical test based on the Poisson-binomial distribution can be performed to decide whether the number of tumors altered in both genes deviates from the expectation.

We repeated the simulation study performed for the binomial test, this time applying the DISCOVER test. First, our data only contained independently generated alterations. Testing for co-occurrence (Fig. 2e) and mutual exclusivity (Fig. 2g) resulted in  $P$ -value distributions much closer to uniform, as one would expect. The fact that these distributions are not truly uniform is a property shared by all discrete test statistics [21]; it makes discrete tests slightly more conservative. Most importantly, the anti-conservative bias towards co-occurrence of the binomial test is not present in the DISCOVER test. By testing simulated co-occurrences, we established that the removal

of the anti-conservative bias does not compromise the sensitivity for true co-occurrences (Fig. 2f). Moreover, the sensitivity for mutual exclusivities is improved when compared with the binomial test (Fig. 2h).

#### Extension to a group-based mutual exclusivity test


Mutual exclusivity is not restricted to pairs of genes. Larger groups of genes may also display alteration patterns in which most tumors only have an alteration in one of the genes. We considered three statistics to assess the mutual exclusivity of groups of genes: coverage, exclusivity, and impurity (Fig. 3a). For all three of these statistics, its expectation for groups of independent genes can be described by a Poisson-binomial distribution (see Methods), and thus a statistical test can be formulated for determining significance. Based on simulated data, we established that the impurity-based group test has the best balance between sensitivity and specificity (Additional file 1: Figure S3).

#### Comparison to other mutual exclusivity tests

We compared the performance of the group-based DISCOVER test to that of several other published mutual exclusivity tests: *MEMo* [6], *muex* [8], *mutex* [9], *CoMEt* [10], *MEGSA* [11], and *TIMEx* [12]. In this comparison, we focused on the statistical tests for mutual exclusivity provided by these methods (see Methods). Although the tests differ in the statistical model upon which they are based, all but *MEMo* assume identical alteration probabilities across tumors. Like Fisher's exact test and the binomial test, they are thus examples of tests based on the i.i.d. assumption. *MEMo* does take into account tumor-specific alteration rates by preserving these rates in a permutation scheme. Unlike DISCOVER, it estimates the alteration rate with respect to a small set of recurrently altered genes as opposed to all genes.

The comparison was performed on simulated data. Groups of genes with mutually exclusive alterations of various degrees of impurity served as positive examples (see Methods). For each such group, we also selected groups of independent genes of the same size and matched to have similar alteration frequencies, to serve as negative examples. In total, 10 data sets of 100 positive and 100 negative groups were generated, and evaluation metrics were averaged across these 10 sets. We evaluated the tests for both specificity and sensitivity.

To evaluate specificity, we considered the extent to which a chosen significance level  $\alpha$  predicts the false positive rate obtained when groups with a nominal  $P$  value less than  $\alpha$  are classified as mutually exclusive. By definition of the  $P$  value, rejecting the null hypothesis at a significance level  $\alpha$  should guarantee that the false positive rate (or type I error rate in statistical terminology) is at most  $\alpha$ . Graphically, if the false positive rate is plotted as



**Fig. 3** Extension of the DISCOVER test for mutual exclusivity within groups of genes. **a** Three alternative statistics for measuring the degree of mutual exclusivity within a group of genes. Coverage refers to the number of tumors that have an alteration in at least one of the genes. Exclusivity refers to the number of tumors that have an alteration in exactly one gene. Impurity refers to the number of tumors that have an alteration in more than one gene. **b** *P*-value reliability curves comparing DISCOVER with other mutual exclusivity tests. The false positive rate should not exceed the significance level  $\alpha$ . In such a case, the calibration curve will be below the diagonal. For all tests but muex, this is the case. The curves for CoMet, MEGSA, mutex, and TiMEx are mostly overlapping; their false positive rate stays at 0 until the significance level is almost 1. **c** Sensitivity curves comparing DISCOVER with other mutual exclusivity tests. More sensitive tests will attain higher true positive rates at lower significance levels. Two discontinuities that occur at a significance level of approximately  $1 \times 10^{-16}$  are marked with dotted lines. First, muex compresses all lower *P* values to 0; hence, all lower significance levels have the same true positive rate. Second, this significance level coincides with the change from the slower CoMet exact test to the binomial approximation (see Methods); the two tests seem to behave quite differently

a function of the significance level (Fig. 3b), the resulting curve would ideally follow the diagonal, or it should drop below the diagonal for more conservative tests. With the exception of muex, all methods control their false positive rate below the nominal significance level, but they do so in notably different ways. CoMet, mutex, and TiMEx only yield false positives at extremely high significance levels. Doing so, they are more conservative than required. In contrast, DISCOVER’s curve follows the diagonal more closely. This is another confirmation that tests based on the i.i.d. assumption—like before with the binomial and Fisher’s exact tests—are more conservative than those that model the varying alteration rates. Indeed, MEMo is also less conservative than CoMet, mutex, and TiMEx. It is more conservative than DISCOVER though, which may

be explained by the different strategies for estimating the tumor-specific alteration rates: based on all genes for DISCOVER, or based on frequently altered genes only for MEMo.

To evaluate sensitivity, we compared the increase of the true positive rate as a function of the significance level (Fig. 3c). A sensitive test will already attain high true positive rates at low significance levels. Across the whole range of significance levels, DISCOVER was found to be more sensitive than any of the other tests. It identified more mutually exclusive groups at lower significance levels. Only muex initially shows a higher sensitivity, but it does so at the price of many false positives (Fig. 3b)—we suspect this is partly due to numerical imprecision. At higher significance levels, muex’s sensitivity drops below

that of DISCOVER. MEMo only attains a high sensitivity at higher significance levels: it is affected by the limited resolution of its permutation test. We used 10,000 permutations, which makes the lowest possible  $P$  value  $1 \times 10^{-4}$ . Again contrasting tests based on their underlying assumption, we conclude that the conservatism caused by the i.i.d. assumption is reflected in a lower sensitivity. The majority of mutually exclusive groups are only identified at relatively high significance levels. If correction for multiple testing is applied, this may render many of them insignificant.

### Co-occurrence and mutual exclusivity in pan-cancer somatic alterations

We analyzed a set of 3386 tumors covering the 12 cancer types studied in the TCGA pan-cancer initiative [22]. An alteration matrix was constructed from recurrent copy number changes and high-confidence mutational drivers. Copy number changes were analyzed for 118 genes, of which 40 were gains and 78 were losses. In addition, mutation data were added for 286 genes previously classified as high-confidence driver genes [23]. In total 404 genomic alterations were analyzed covering 374 unique genes, as 30 genes are frequently targeted by both copy number changes and mutations.

We tested for pairwise co-occurrence and mutual exclusivity between pairs of genes not located on the same chromosome. These tests were stratified for cancer type to avoid confounding due to cancer type-specific alteration frequencies. Complementing the pairwise tests, we also employed the DISCOVER group test to detect patterns of mutual exclusivity in larger groups of genes. The groups we tested were selected using two different approaches. In the first approach, we extracted gene sets from the canonical pathway collection of MSigDB [24]. We tested 23 such gene sets based on pathway membership. In the second approach, we aimed to detect de novo gene sets purely based on the data. For this, we applied a clustering algorithm to the pairwise mutual exclusivity results to identify groups of genes showing a high degree of interaction.

### No evidence for widespread co-occurrence

A remarkable outcome of our analysis is that we found no evidence for widespread co-occurrence of somatic alterations. At a maximum false discovery rate (FDR) of 1%, no significant co-occurrences were identified. Relaxing the FDR threshold to 3%, we could recover one co-occurrence, between mutation of *TP53* and amplification of *MYC*. It was recently suggested that *MYC*-amplified tumors show higher levels of *MYC* expression in tumors with a *TP53* mutation than in tumors without [25]. No further, reasonable relaxation of the significance threshold led to additional hits. Certainly, more gene pairs exist that harbor alterations in overlapping sets of tumors. Yet, the

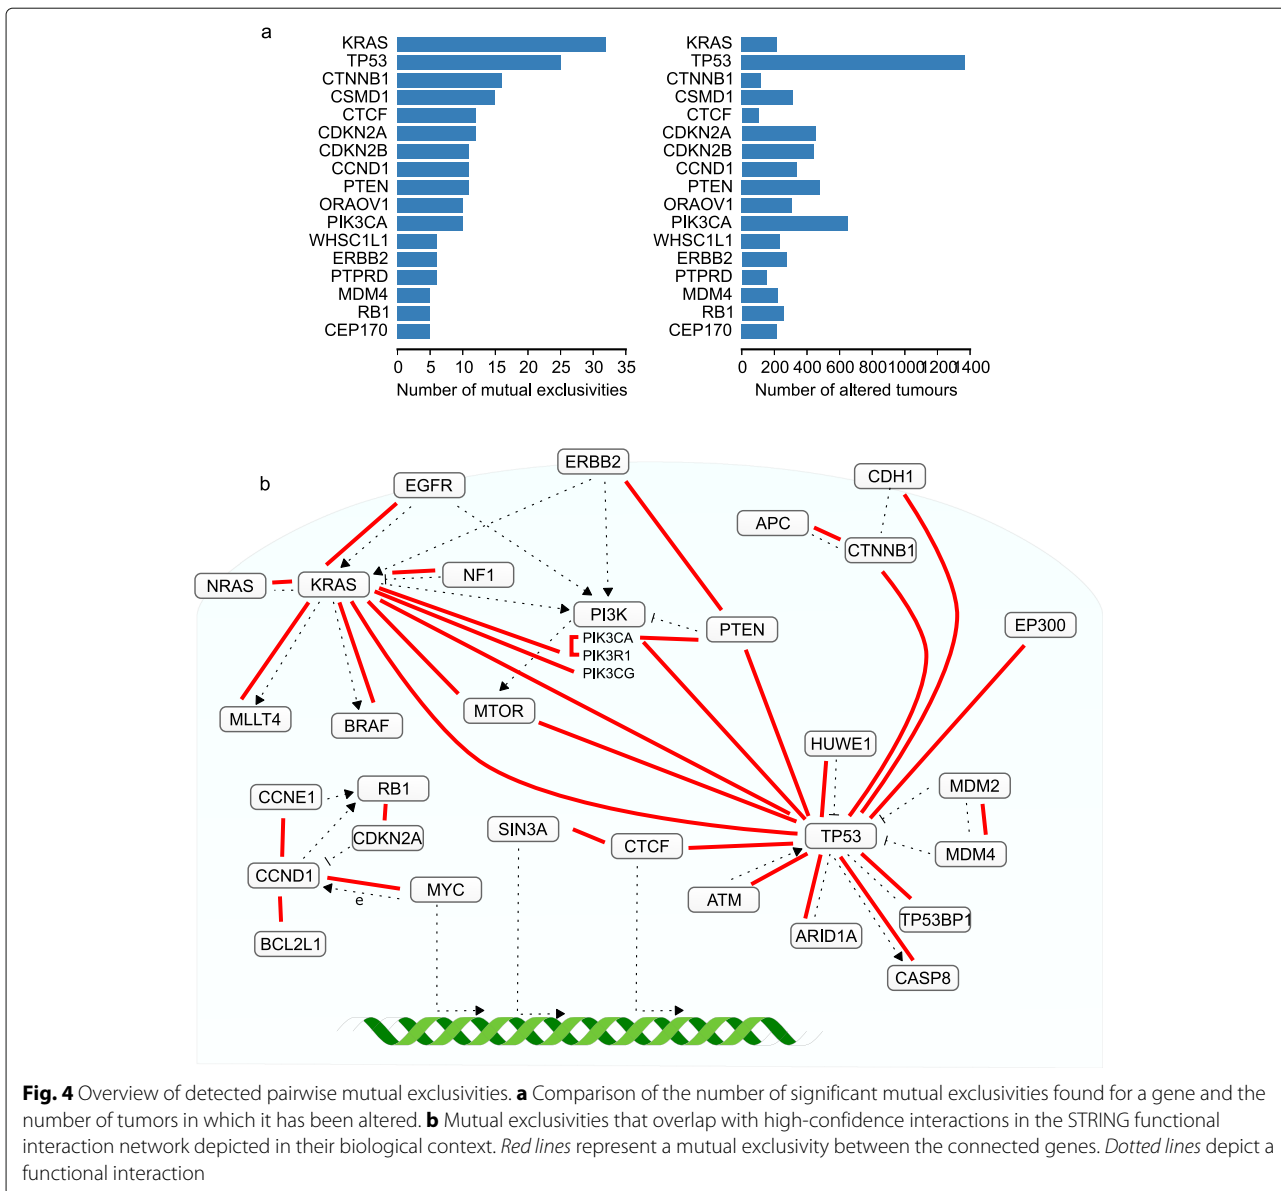
sizes of those overlaps do not exceed what is expected by chance if differences in tumor-specific alteration rates are taken into account. This is in sharp contrast with the significance estimates obtained with the binomial test, which identifies 21,627 significant co-occurrences, almost one-third of all pairs tested.

With the aim of establishing that the DISCOVER test is not overly conservative, we tested for co-occurrence between copy number changes of genes on the same chromosomes. Due to the inherent correlation in copy number of genes situated close to each other, such gene pairs can be considered positive controls. Indeed, all but one of the 112 pairs of tested genes located in the same recurrently altered segment are identified as co-occurring by the DISCOVER test. In addition, 18 pairs of genes situated on the same chromosome arm are detected as co-occurring, as are *DDAH1* on 1p22 and *MCL1* on 1q21. More generally, pairs within the same segment are assigned lower  $P$  values on average than are pairs within the same chromosome arm ( $P = 7 \times 10^{-39}$ , Additional file 1: Figure S4). The same is true, to lesser extents, for pairs within the same chromosome arm compared to pairs within the same chromosome ( $P = 6 \times 10^{-8}$ ) and for pairs within the same chromosome compared to pairs across chromosomes ( $P = 0.0004$ ).

### Mutually exclusive alterations target core cancer pathways

Pairwise mutual exclusivities were found among 181 pairs of genes, at a maximum FDR of 1% (Additional file 2: Table S1). We once more confirmed that detecting mutual exclusivities using the binomial test results in far fewer significant mutual exclusivities—only three pairs were identified. Among the 181 gene pairs, there were 107 unique genes. Many of these are significantly mutually exclusive with only one or a few other genes. For some, reduced statistical power due to low alteration frequency may be the reason for not detecting more associations. However, alteration frequency is not the dominant factor in how often mutual exclusivity is detected (Fig. 4a). For example, mutations of *KRAS* are far less frequent than *TP53* or *PIK3CA* mutations. Yet, *KRAS* was found mutually exclusive with more genes than were the latter two genes.

Since mutual exclusivity is believed often to occur between functionally related genes, we determined the overlap of the identified gene pairs with the STRING functional interaction network [26]. Thirty-one of the identified gene pairs have a high-confidence functional interaction in STRING (Fig. 4b). This overlap is significantly higher than the 5 overlapping pairs expected by chance ( $P < 1 \times 10^{-4}$ ), as determined using a permutation test. Moreover, 121 of the mutually exclusive gene pairs share a common interactor in the STRING network. By chance, this is only expected to be the case for 80 gene



**Fig. 4** Overview of detected pairwise mutual exclusivities. **a** Comparison of the number of significant mutual exclusivities found for a gene and the number of tumours in which it has been altered. **b** Mutual exclusivities that overlap with high-confidence interactions in the STRING functional interaction network depicted in their biological context. Red lines represent a mutual exclusivity between the connected genes. Dotted lines depict a functional interaction

pairs ( $P = 0.003$ ). This suggests that the mutual exclusivities identified are indeed for a large part driven by biological factors. Another confirmation of this is found in the results of the MSigDB gene set tests (Additional file 1: Figure S5). Twelve gene sets representing several cancer-related pathways show significant mutual exclusivity. The mutual exclusivities that overlap with STRING interactions revolve around three commonly deregulated processes in cancer: growth factor signaling, cell cycle control, and p53 signaling.

**Growth factor signaling**

Genes coding for proteins involved in growth factor signaling are frequently altered in cancer. These alterations display a high degree of mutual exclusivity. Mutations

targeting the receptor *EGFR* are mutually exclusive with mutations in its downstream mediator *KRAS*. In turn, *KRAS* mutations are mutually exclusive with mutations in its family member *NRAS*, its negative regulator *NF1*, and its downstream effector *BRAF*. All of these alterations are able to deregulate RAS signaling, and one is sufficient. Mutual exclusivity of mutations in *KRAS* and mutations in both *PIK3R1* and *PIK3CG* may be driven by the known cross-talk between RAS signaling and phosphoinositide 3-kinase (PI3K) signaling [27].

The PI3K signaling cascade itself is also characterized by many mutually exclusive alterations. Mutations in the *PIK3CA* and *PIK3R1* genes—both coding for components of the PI3K complex—are mutually exclusive. Alterations in the *PTEN* gene—a negative regulator of

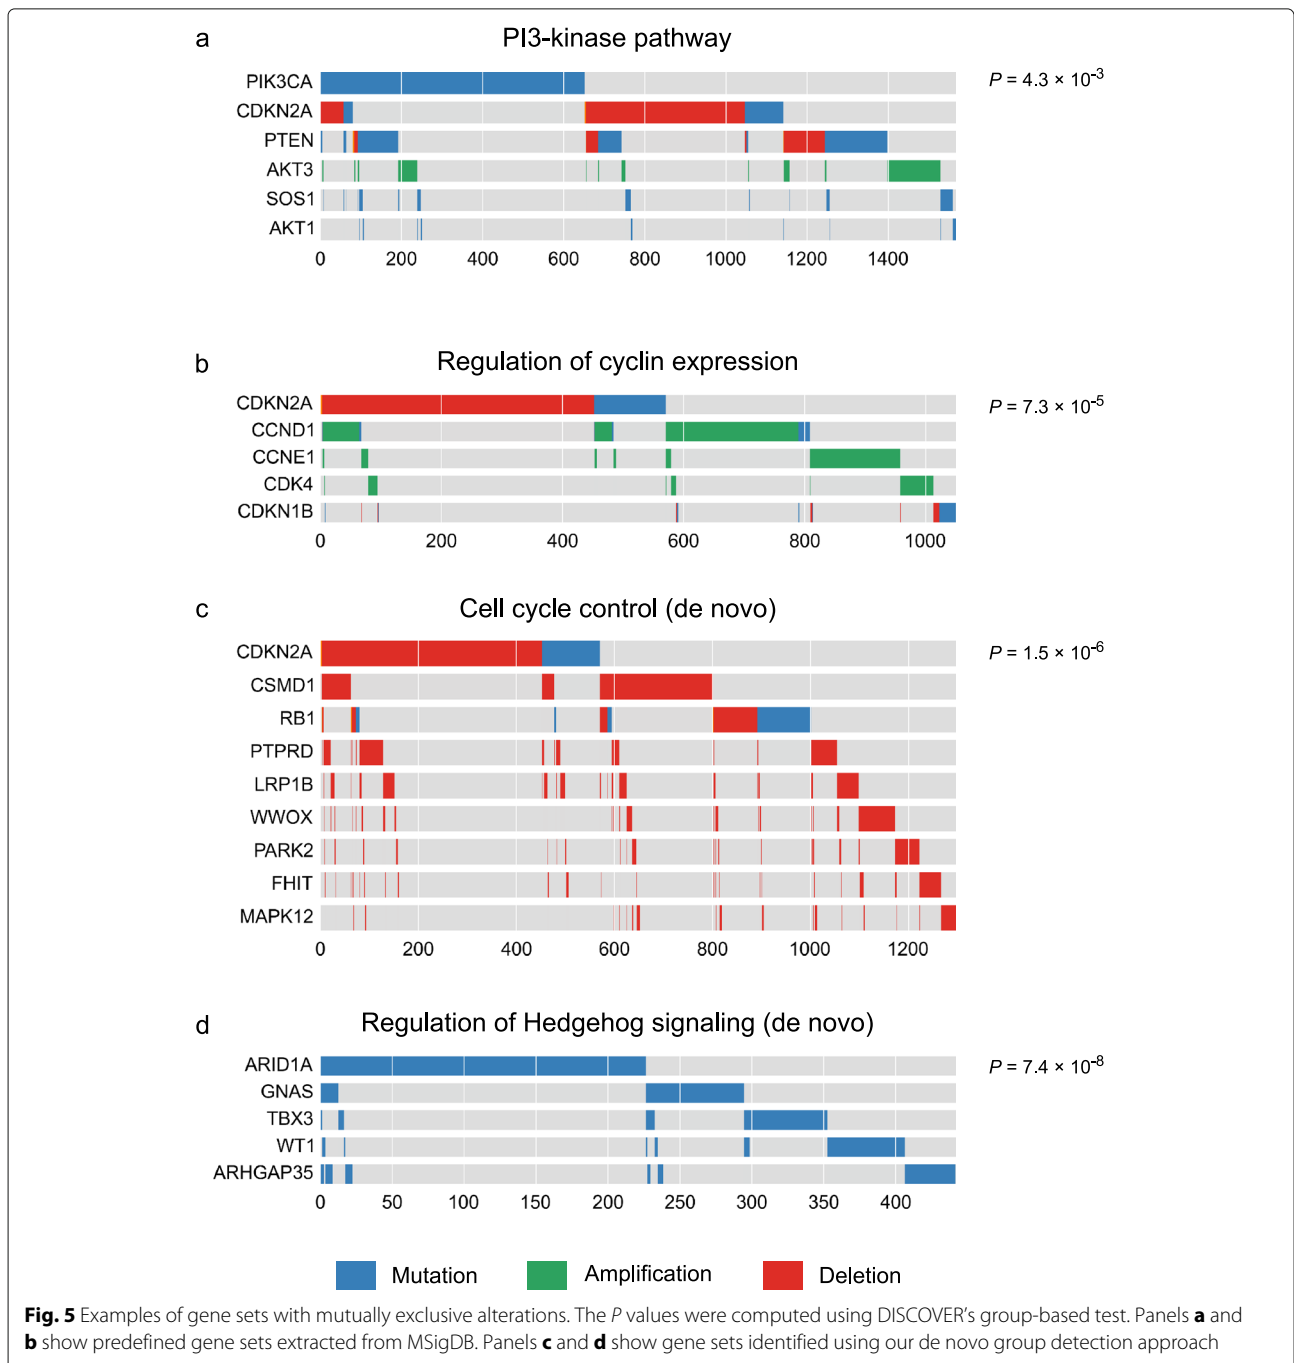


the downstream activation of AKT by PI3K—are mutually exclusive with mutations in *PIK3CA*, but also with alterations in the upstream activator of the cascade *ERBB2*. PI3K signaling is also the central biological process in several of the gene sets found mutually exclusive with the group-based test (Fig. 5a, Additional file 1: Figure S5). Central genes in PI3K signaling such as *SOS1*, *AKT1*, and *AKT3* were not found as mutually exclusive with

other pathway members in the pairwise analysis, yet the groupwise test correctly detects it.

**Cell cycle control**

Many tumors harbor alterations that disable the cell cycle control present in healthy cells. This control arises from a tightly regulated interplay between cell cycle-activating cyclins and CDKs, and CDK inhibitors, linked together by



the master cell cycle regulator RB1. Alterations in these genes are also mutually exclusive. For example, copy number gains in *Cyclins D1* and *E1* are mutually exclusive, as are *CDKN2A* copy number loss and both mutation and copy number loss of *RB1*. The transcriptional activation of *CCND1* by *MYC* is also reflected in the mutual exclusivity between copy number gains in the two genes. Also as a group, cyclins, CDKs, and CDK inhibitors show a clear pattern of mutual exclusivity (Fig. 5b, Additional file 1: Figure S5). *CDK4* and *CDKN1B*, central players in the regulation of the cell cycle, did not appear in the pairwise results, but are highly exclusive with the other genes involved.

### **p53 signaling**

p53 plays a pivotal role in deciding on cell fate after cellular stresses common in cancer development. For this reason, p53 mutations are the most common alterations in cancer. However, not all tumors disable p53 function genetically. Alterations in regulators of p53 provide an alternative way to deregulate p53 function in p53-wild-type tumors, but are likely redundant in tumors that already have a dysfunctional p53 protein. Indeed, we found alterations in several regulators of p53 to be mutually exclusive with *TP53* mutation. For example, mutations in its positive regulator *ATM*, but also mutations in its negative regulator *HUWE1* are mutually exclusive with *TP53* mutations. *MDM2* and *MDM4*, highly similar negative regulators of p53, have a mutually exclusive pattern of copy number gains. Mutations in *CASP8*, a downstream mediator of p53-induced apoptosis, tend also not to overlap with *TP53* mutations.

### **De novo gene set detection**

As a final step in our analysis, we detected de novo gene sets purely based on observed patterns of mutual exclusivity, without input based on recorded biological knowledge. To this end, we applied correlation clustering to a network derived from pairwise mutual exclusivities (see Methods). This identified 120 candidate mutually exclusive gene sets. Testing these gene sets with DISCOVER, 43 were found to be mutually exclusive at a maximum FDR of 1%. The full results are presented in the online Jupyter notebooks (see Availability of data and materials). Below, we discuss two interesting examples.

One of the most significant gene sets includes *RB1* and *CDKN2A*, two pivotal players in cell cycle control (Fig. 5c). *PARK2* [28], *WWOX* [29], *FHIT* [30], *PTPRD* [31, 32], and *MAPK12* [33] have also all been linked to a regulating role in various phases of the cell cycle. They have been found to do so by regulating cyclins, CDKs, or CDK inhibitors. This functional similarity may explain these genes' mutual exclusivity with *RB1* and *CDKN2A*. As of yet, *LRP1B* and *CSMD1* have not been linked to cell cycle control. Their

mutual exclusivity with respect to several regulators of the cell cycle may instigate further study in this direction.

Another group of genes with a high degree of mutual exclusivity ( $P = 7 \times 10^{-8}$ ) consists of genes that have been implicated in the regulation of Hedgehog signaling (Fig. 5d). With the exception of *ARHGAP35*, all genes in this group have experimentally been linked to a regulatory role in Hedgehog signaling. *GNAS* [34, 35], *TBX3* [36], and *WT1* [37] were found to directly regulate the pathway. *ARID1A*, coding for a component of the SWI/SNF complex, is likely to play a similar role, since loss of another component of this complex, Snf5, was found to lead to activation of the Hedgehog pathway [38]. Besides these two examples, several other gene sets were identified that combine known interaction partners with interesting leads for undiscovered interactions.

### **Discussion**

The recent growth in the number of large genomics data sets gives rise to a parallel increase in statistical power to detect ever more complex associations. However, as another consequence of larger sample sizes, poorly matched assumptions will have an increasing impact on the results. A central assumption behind commonly used statistical tests for co-occurrence and mutual exclusivity is that a gene's alteration probability is identical across all tumors. Using simulated data, we have shown that this assumption is not only unjustified, but that it leads to a full reversal of the associations. The binomial test we used for illustration is but a representative of a larger class of independence tests based on the same assumption. This class includes analytical approaches such as Fisher's exact test, CoMet [10], and MEGSA [11], but also permutation tests where gene alterations are uniformly shuffled across the tumors.

We have presented a novel independence test based on assumptions that better match the reality of cancer genomics data. With this new test, we analyzed tumors across 12 different cancer types for the presence of co-occurrence and mutual exclusivity. Only one case of co-occurrence was found, whereas numerous cases of mutual exclusivity were detected. Performing the same analysis with the binomial test led to the detection of many co-occurrences and almost no mutual exclusivity. Many of the mutual exclusivities missed by the binomial test can be related to central processes in cancer biology. We found strong mutual exclusivity between genes involved in growth factor signaling and cell cycle control. Also, lesser known players in the regulation of cell cycle and Hedgehog signaling were identified. Based on the results of our simulation study, we are confident that most of the co-occurrences detected by the binomial test are spurious.

The absence of widespread co-occurrence contradicts what was found in previous genome-wide studies. Besides, it seems counter to our expectation of positive selection for synergy that led us to look for co-occurrence in the first place. It is true that synergy resulting from the alteration of multiple genes has been observed. Co-mutation of genes has been reported to act on a tumor’s response to chemotherapy, or more generally on patient survival [39, 40]. None of these phenotypes, however, has been the subject of the selection from which the original tumor emerged. Only after selective pressure for that particular phenotype has taken place—for example, by treating patients—would enrichment for such co-occurrences be detected. There is no doubt that cancer-driving alterations often act in concert. Yet if statistical results are to serve as support for, or even meant to identify synergy, other possible explanations for the observed co-occurrence should be accounted for. In our pan-cancer analysis, overall alteration rates explained most if not all co-occurrence.

The need to take into account higher level structural features of samples is not unique for co-occurrence and mutual exclusivity analysis. In testing the relationship between high-dimensional gene expression data and phenotypes of interest, latent sources of heterogeneity can have a profound effect on the results. Approaches like surrogate variable analysis [41] have been developed to adjust analyses appropriately. Similarly, genome-wide association studies face the issue of latent population substructure. Again, if ignored, such substructure can drastically alter the findings. Linear mixed models have gained popularity as a method to prevent confounding [42]. Both of these examples have become standard methodologies in many biomedical analyses.

**Conclusions**

Co-occurrence and mutual exclusivity of somatic alterations are helpful concepts for the interpretation of cancer genomics data. For example, hypotheses about functional interactions between genes are often supported by suggested co-occurrence or mutual exclusivity of their alterations. Alarmingly, we have found that the statistical tests most commonly used for this purpose are not appropriate for testing the significance of co-occurrence. Many gene pairs that are believed to be co-altered more often than expected by chance do not exceed this expectation if the confounding effect of tumor-specific alteration rates is taken into account. Hypotheses formulated based on the results of those tests will therefore have limited support from the data. For this reason, we discourage the use of Fisher’s exact test or simple permutation methods for detecting co-occurrence. We have presented DISCOVER as a better alternative. Mutual exclusivity analysis using existing tests does not suffer

from high false positive rates, but the sensitivity is low. DISCOVER identifies more significant mutual exclusivities without increasing the false positive rate. Thus, for both co-occurrence and mutual exclusivity analyses, we expect future cancer genomics studies to benefit from DISCOVER.

**Methods**

**Independence statistic**

We assess both co-occurrence and mutual exclusivity by counting how many tumors have an alteration in both genes and comparing this to the number of tumors expected to have such an overlap by chance if these alterations were independent. Importantly, the overlap expected by chance should factor in the fact that tumors with many alterations have a higher chance of such overlap than tumors with fewer alterations. Our null distribution modeling this overlap therefore takes into account both the alteration rate per gene and the alteration rate per tumor. To this end, let  $p_{ij}$  denote the probability of an alteration in gene  $i$  and tumor  $j$ . We assume that the alteration probability of a gene is higher in tumors with many alterations overall than in tumors with fewer alterations. Therefore,  $p_{ij}$  may be different from  $p_{ik}$  for the same gene  $i$  in two different tumors  $j$  and  $k$ . Then, for two independent genes with alteration probabilities  $p_{1j}$  and  $p_{2j}$ , the probability of an alteration in both genes in tumor  $j$  is  $p_{1j}p_{2j}$ , while for tumor  $k$  it is  $p_{1k}p_{2k}$ . Given such probabilities for a set of tumors, the number of tumors that have an alteration in both genes follows a Poisson-binomial distribution.

The Poisson-binomial distribution [43] describes the sum of independent, non-identically distributed Bernoulli random variables that have success probabilities  $p_1, p_2, \dots, p_n$ . Its probability mass function is defined as follows:

$$P(X = x) = \sum_{A \in \mathcal{F}_x} \left( \prod_{i \in A} p_i \prod_{j \in A^c} (1 - p_j) \right)$$

Here,  $\mathcal{F}_x$  contains all subsets of size  $x$  of  $\{1, 2, \dots, n\}$ , and  $A^c$  denotes the complement of  $A$ .

Based on this distribution, we can estimate the probability of observing a number of tumors with alterations in two genes as extreme—as high for co-occurrence, or as low for mutual exclusivity—as the one observed.

If, for a given gene  $i$ , all probabilities  $p_{ij}$  are equal for every tumor  $j$ , then the Poisson-binomial distribution reduces to a binomial distribution. However, estimating an individual alteration probability for every single tumor ensures that the heterogeneity in alteration rates across tumors is taken into account.

**Estimating gene- and tumor-specific alteration probabilities**

To apply the DISCOVER test, we need estimates of the alteration probabilities  $p_{ij}$  for all genes  $i$  and all tumors  $j$ . Let  $\mathcal{X} \in \{0, 1\}^{n \times m}$  denote the  $n \times m$  binary alteration matrix where an entry  $x_{ij}$  is 1 in case of an alteration in gene  $i$  and tumor  $j$ , and 0 otherwise. We use the notation  $x_{i\bullet}$  and  $x_{\bullet j}$  for the marginal sums of the  $i$ th row and  $j$ th column, respectively. Furthermore, let  $X_{ij}$  denote the random variable for  $x_{ij}$ , and  $X_{i\bullet}$  and  $X_{\bullet j}$  the corresponding marginal sums. If we were to assume that the alteration of a gene is equally likely across all tumors, then the alteration probability only depends on the number of altered tumors  $x_{i\bullet}$  and the total number of tumors  $m$ :

$$p_{ij} = P(X_{ij} = 1 | x_{i\bullet} = k) = \frac{k}{m}, \quad \forall j$$

Estimating the alteration probabilities this way ensures that the expected number of alterations  $E_p(X_{i\bullet}) = \sum_j p_{ij}$  for a gene matches the observed number  $x_{i\bullet}$ . In fact, the familiar expression above is the one that maximizes the likelihood of the observed alterations under the constraint that the expected number of alterations per gene matches the observed number. To make this more explicit, we can reformulate the probability estimation as a constrained optimization problem:

$$\begin{aligned} \max_p \quad & L_p(\mathcal{X}) = \prod_{i=1}^n \prod_{j=1}^m (p_{ij} x_{ij} + (1 - p_{ij})(1 - x_{ij})) \\ \text{s.t.} \quad & p_{i1} = p_{i2} = \dots = p_{im}, \quad 1 \leq i \leq n \\ & \sum_{j=1}^m p_{ij} = \sum_{j=1}^m x_{ij}, \quad 1 \leq i \leq n \\ & 0 \leq p_{ij} \leq 1, \quad 1 \leq i \leq n, 1 \leq j \leq m \end{aligned}$$

All of the above is based on the assumption that alteration probabilities for a gene are equal across tumors. Symptomatic for this assumption are probability estimates such that the expected number of alterations per tumor  $E_p(X_{\bullet j}) = \sum_i p_{ij}$  generally does not match the observed number  $x_{\bullet j}$ . To take into account tumor-specific alteration rates, the above optimization problem can be extended such that this expectation is also matched:

$$\begin{aligned} \max_p \quad & H_p(\mathcal{X}) = - \sum_{i=1}^n \sum_{j=1}^m (p_{ij} \log(p_{ij}) + (1 - p_{ij}) \log(1 - p_{ij})) \\ \text{s.t.} \quad & \sum_{c=1}^m p_{ic} = \sum_{c=1}^m x_{ic}, \quad 1 \leq i \leq n \\ & \sum_{r=1}^n p_{rj} = \sum_{r=1}^n x_{rj}, \quad 1 \leq j \leq m \\ & 0 \leq p_{ij} \leq 1, \quad 1 \leq i \leq n, 1 \leq j \leq m \end{aligned}$$

With this new formulation, the number of parameters to fit is increased by a factor  $m$ . As a consequence, optimizing the likelihood  $L_p(\mathcal{X})$  of the model risks overfitting the data. Therefore, instead of optimizing the likelihood, we choose to optimize the information entropy  $H_p(\mathcal{X})$ . It can be shown that in the optimal solution to this reformulated problem, each alteration probability can be written in terms of two parameters (Additional file 1: Parameter estimation):

$$p_{ij} = \frac{1}{1 + e^{\mu_i + \lambda_j}}$$

Here, each parameter  $\mu_i$  for gene  $i$  is shared by all tumors, and each parameter  $\lambda_j$  for tumor  $j$  is shared by all genes. Because of this, while the original optimization problem aims to estimate  $n \times m$  alteration probabilities, we can obtain the optimal solution by estimating only  $n + m$  parameters. Moreover, all genes with the same number of altered tumors share the same value for  $\mu_i$ . Likewise, all tumors with the same number of altered genes share the same value for  $\lambda_j$ . This sharing of parameters leads to an even larger reduction in the effective dimensionality of the optimization.

Unlike for the binomial case, there is no closed-form solution for estimating the  $\mu_i$  and  $\lambda_j$  parameters. Instead, we use the quasi-Newton numerical optimization algorithm L-BFGS [44].

**Stratified analysis**

When the data consist of clearly separate groups of tumors, such as is the case in the pan-cancer analysis with its different cancer types, it is preferable to stratify the analysis on these groups. For example, in the mutual exclusivity analysis, if group structure is not taken into account, the detected mutual exclusivities may be little more than markers for the underlying cancer types, rather than biologically related genes. An example of this type of confounding is presented in Additional file 1: Stratification in pan-cancer analysis. The DISCOVER test is easily stratified for different groups by solving the constrained optimization problem separately for the tumors of each group. The group-specific background matrices can then be concatenated to construct a single global, but stratified, parameter matrix.

More formally, the binary alteration matrix  $\mathcal{X}$  can be seen as a concatenation of several  $n \times m_c$  submatrices  $\mathcal{X}_c$ , where  $c \in \{1, 2, \dots, C\}$  refers to one of  $C$  possible subgroups—e.g., a cancer type in the pan-cancer analysis—and  $m_c$  is the number of tumors in that group:

$$\mathcal{X} = [\mathcal{X}_1 \ \mathcal{X}_2 \ \dots \ \mathcal{X}_C]$$

To illustrate this, Additional file 1: Figure S6a shows an alteration matrix with tumors of two different subtypes. The parameter estimation procedure described in

the previous section is then applied to each submatrix  $\mathcal{X}_c$  individually, resulting in subgroup-specific probability matrices  $P_c$  (Additional file 1: Figure S6b). The global, stratified probability matrix is obtained by concatenating these matrices:

$$P = [P_1 \ P_2 \ \dots \ P_C]$$

As in the non-stratified case, the expected number of alterations for each gene matches the observed number. However, unlike for the non-stratified probabilities, the expected numbers also match the observed numbers within each subgroup. With this stratified probability matrix, the Poisson-binomial test is applied in the same way as in the non-stratified setting (Additional file 1: Figure S6c).

**False discovery rate control**

Commonly used procedures for multiple testing correction assume that the  $P$  values are distributed uniformly under the null hypothesis. This is the case for, e.g., Bonferroni correction and the Benjamini-Hochberg procedure. However, hypothesis tests that are based on a discrete test statistic, such as our DISCOVER test, are known to lead to non-uniform  $P$ -value distributions under the null hypothesis. In fact, pooling the  $P$  values across tests with a large set of different parameters results in a  $P$ -value distribution that is skewed towards 1.0. This complicates the application of the standard procedures for multiple testing correction. While these procedures would still control the familywise error rate or false discovery rate at the specified threshold, they will be more conservative because of the non-uniformity caused by the discrete test statistic. For the analyses in this paper, we used an adaptation of the Benjamini-Hochberg procedure for discrete test statistics [45]. Further details on this procedure are provided in Additional file 1: False discovery rate control for discrete tests.

**Group-based mutual exclusivity test**

We have defined a family of group-based mutual exclusivity tests. The following statistics can be used to assess groupwise mutual exclusivity. Each of these statistics can be shown to follow a Poisson-binomial distribution, which we make use of to estimate significance.

- *Coverage*: the number of tumors that have an alteration in at least one of the genes. Significance is based on the probability of observing a coverage at least as high in independent genes. The Poisson-binomial parameters for a group of genes  $\{g_i \mid i \in I\}$  can be derived from the individual gene alteration probabilities as follows:

$$p_j = 1 - \prod_{i \in I} (1 - p_{ij}), \quad 1 \leq j \leq m$$

That is, the probability of at least one alteration is one minus the probability of not having any alteration.

- *Exclusivity*: the number of tumors that have an alteration in exactly one of the genes. Significance is based on the probability of observing exclusivity at least as high in independent genes. The Poisson-binomial parameters can be derived from the gene alteration probabilities as follows:

$$p_j = \sum_{i \in I} p_{ij} \prod_{k \in I \setminus \{i\}} (1 - p_{kj}), \quad 1 \leq j \leq m$$

- *Impurity*: the number of tumors that have an alteration in more than one gene. Significance is based on the probability of observing impurity at least as low in independent genes. The Poisson-binomial parameters can be derived from the gene alteration probabilities as follows:

$$p_j = 1 - \prod_{i \in I} (1 - p_{ij}) - \sum_{i \in I} p_{ij} \prod_{k \in I \setminus \{i\}} (1 - p_{kj}), \quad 1 \leq j \leq m$$

That is, the probability of more than one alteration is one minus the probabilities of no alterations and exactly one alteration. As a special case of this, if a group of only two genes is tested, the above expression reduces to  $p_j = p_{1j}p_{2j}$ . This is the same parameterization as was used for the pairwise test.

**Simulation data**

An alteration matrix was constructed such that alteration frequencies across both genes and tumors resembled those of real tumors. For this, we used the copy number data of the TCGA breast cancer study as a reference. Based on the copy number matrix for 24,174 genes and 1044 tumors, we constructed two sequences of marginal counts corresponding to the number of amplifications across genes and across tumors. These two sequences were used as degree sequences to construct a random bipartite graph following the configuration model. The adjacency matrix of this bipartite graph was then used as the alteration matrix for the simulated data analyses. Because of the way this matrix was constructed, the alteration frequencies across both genes and tumors resemble those of the breast cancer tumors used for reference, yet there is no dependence between alterations across genes. For the analyses, only genes with at least 50 alterations were tested.

Mutually exclusive and co-occurring gene pairs, as well as mutually exclusive gene sets, were generated based on two parameters: coverage, the number of tumors altered in at least one of the genes; and impurity or overlap, the proportion of covered tumors altered in more than one of the genes. To generate pairs of mutually exclusive genes, we used quantile regression to relate the coverage of independent gene pairs to their impurity. Simulated

mutually exclusive gene pairs were generated such that their impurity was below the first percentile predicted by the quantile regression model based on their coverage. Likewise, pairs of co-occurring genes were generated such that the number of tumors altered in both genes exceeded the 99th percentile based on the coverage of independent gene pairs.

Mutually exclusive gene sets were generated by first constructing sets of purely mutually exclusive gene alterations and then adding additional, non-exclusive alterations to obtain a prespecified degree of impurity. For the former, the percentage of covered tumors was randomly sampled from a truncated normal distribution with mean 0.4 and standard deviation 0.2, truncated on the interval [0.2, 0.8]. Next, individual gene alteration frequencies were sampled from the empirical distribution of alteration frequencies in the TCGA breast cancer matrix. Gene alteration frequencies were sampled until their sum reached the coverage of the group. The number of genes thus depends on the coverage in a way that is based on realistic cancer data. As some of the mutual exclusivity tests we compared with become intractable with larger numbers of genes, we restricted the maximum number of genes to 6. In addition, we also used a minimum gene set size of 3. Finally, the impurity was sampled from the set {0.02, 0.05, 0.08}. Impure alterations, i.e., additional alterations in an already covered tumor, were assigned to tumors with a probability proportional to the tumor's overall alteration frequency.

For all analyses, the background matrix for the DISCOVER test was estimated on the complete alteration matrix, including genes with fewer than 50 alterations, and including simulated co-occurrences or mutual exclusivities.

#### Comparison to other mutual exclusivity tests

We compared the performance of the group-based DISCOVER test to that of MEMo [6], muex [8], mutex [9], CoMEt [10], MEGSA [11], and TiMEt [12]. Some of these methods do more than just test for mutual exclusivity. They combine a statistical test for mutual exclusivity with an algorithm that identifies groups of genes to test. In our comparison, we were interested in comparing the performance of the statistical tests only. We therefore evaluated the mutual exclusivity tests by applying them to preidentified groups of genes.

For muex, MEGSA, and TiMEt, we used the R implementations provided with their respective publications. For CoMEt, we used a modified version of the official software implementation. Due to the computational complexity of the CoMEt test, it became intractable for some of the gene sets in the comparison. For this reason, the CoMEt publication suggests a set of heuristics to decide between

the exact test and a faster binomial approximation, but we found those to be inadequate in our comparison. Instead, we changed the implementation such that it interrupts the CoMEt exact test after 1 minute and returns the *P* value obtained with the binomial approximation. For the MEMo and mutex tests, we used our own implementations, which we verified to give the same results as their original Java implementations.

#### Pan-cancer alteration data

Preprocessed somatic mutation and copy number data for the 12 cancer types studied in the TCGA pan-cancer initiative [22] were obtained via Firehose (analysis run 2014\_07\_15 at [http://gdac.broadinstitute.org/runs/analyses\\_2014\\_07\\_15/](http://gdac.broadinstitute.org/runs/analyses_2014_07_15/)). Mutations were extracted from the input of the MutSig 2CV analysis. Mutations for genes that have previously been identified as high-confidence mutational drivers [23] were included in the analysis. Discretized copy number changes were extracted from the output of GISTIC2. We considered genes altered if GISTIC2 qualified their copy number change as high level. Pan-cancer recurrently altered regions were obtained via Synapse (syn2203662 at <https://www.synapse.org/#!/Synapse:syn2203662>). For each region, we selected their most likely driver genes for inclusion in the analysis. If a region contained only one gene, this gene was assumed its driver. In the case of more genes, genes were selected if they overlapped with the list of high-confidence mutational driver genes, or with a curated list of cancer genes (<http://www.bushmanlab.org/links/genelists>).

Background matrices for the DISCOVER test were estimated for each type of alteration—mutation, amplification, and deletion—separately, and based on the genome-wide alteration matrices before gene selection. Stratification for the 12 different cancer types was applied as described before. The background matrix used in the analysis was subsequently composed from the relevant rows in the three alteration type-specific background matrices.

#### Overlap with the STRING functional interaction network

Version 10.0 of the STRING network [26] was used to determine overlap of detected mutual exclusivities and functional interactions. We constructed a functional interaction graph by connecting genes with an edge if they had a high-confidence STRING interaction, defined by a combined score greater than 800. A mutual exclusivity graph was constructed by connecting genes with an edge if alterations in these genes were found mutually exclusive at a maximum FDR of 1%. The overlap corresponds to the number of edges appearing in both graphs. To determine the enrichment of this overlap, we estimated a null distribution by randomly shuffling the gene labels of the mutual exclusivity graph 10,000 times and computing the

overlap of these shuffled mutual exclusivity graphs with the unshuffled functional interaction graph.

### De novo gene set detection

Our algorithm for detecting de novo sets of mutually exclusive genes combines two ideas from community detection. Its goal is to detect gene sets with a high likelihood of being mutually exclusive based on the results of a pairwise mutual exclusivity analysis. There are three main steps. First, a mutual exclusivity graph is constructed where genes are connected by an edge if their alterations have been identified as mutually exclusive by the pairwise test. For this step, we used a permissive significance criterion—a maximum FDR of 10%—so as not to exclude potentially interesting gene pairs that may simply not have reached significance due to the limited sample size. Second, groups of genes with a high density of mutual exclusivity edges between them are identified using a graph partitioning algorithm. Finally, these groups are subjected to the groupwise mutual exclusivity test to retain only those groups that are mutually exclusive as a group.

The graph partitioning step is based on overlapping correlation clustering. In correlation clustering, nodes in a graph are clustered such that the combined weight of edges within clusters is maximized and the combined weight of edges between clusters is minimized. The particular algorithm we used [46] allows nodes to be assigned to multiple clusters. Moreover, we modified the original algorithm such that groups of nodes can be designated that should always share the same cluster assignments. We used this for two situations. First, genes in the same copy number segment have highly correlated copy number alterations and, consequently, highly similar mutual exclusivities. Purely based on genomic data, there is no reason to prefer one gene over the other, which is why we always assign all such genes to the same cluster. Second, we assume that copy number alterations and mutations targeting the same gene serve the same function, and therefore add the constraint that these are always assigned to the same cluster.

The edge weights of the mutual exclusivity graph play an important role in the objective function of correlation clustering. A common phenomenon in pairwise associations is that one gene is found mutually exclusive with many other genes, but those genes are not all mutually exclusive with each other. The edges connecting the former gene may therefore not be indicative of gene set membership. They should be assigned a lower weight than edges that more specifically connect genes with a high degree of internal connectivity. To this aim, we selected the edge weights to optimize a modularity objective. In modularity optimization, a graph is compared with random graphs having the same number of nodes, edges, and degree distribution. Edges that are

specific to the graph being partitioned are preferably kept within clusters, whereas edges that also appear in many of the random graphs will often span two clusters. We used a modularity measure based on conditional expected models [47]. This measure ensures that edges connecting sets of nodes with high node degrees receive a lower weight than edges that connect sets of nodes with low node degrees. It also allows for the covariance between the mutual exclusivity tests to be taken into account.

### Additional files

**Additional file 1:** Supplementary methods and **Figures S1–S6.** (PDF 897 kb)

**Additional file 2:** **Table S1.** Significantly mutually exclusive alterations found in the pan-cancer data at a maximum FDR of 1%. (XLSX 16 kb)

### Abbreviations

DISCOVER: Discrete Independence Statistic Controlling for Observations with Varying Event Rates; FDR: False discovery rate; i.i.d.: Independent and identically distributed; TCGA: The Cancer Genome Atlas

### Acknowledgements

We thank Daniel Vis and Evert Bosdriesz for their helpful comments on the manuscript, and John Foekens and Marcel Smid for useful discussions. The results published here are in part based upon data generated by the TCGA Research Network.

### Funding

This study was funded by Cancer Genomics Netherlands.

### Availability of data and materials

Python and R packages implementing DISCOVER are available at <https://github.com/NKI-CCB/DISCOVER>, and are licensed under the Apache License 2.0. The version of the DISCOVER software used for this manuscript has been assigned doi:10.5281/zenodo.167332. Jupyter notebooks containing all the code required to reproduce the figures and results of this manuscript are available at <https://github.com/scanisius/discover-notebooks>. These have been assigned doi:10.5281/zenodo.167469.

### Authors' contributions

SC, JM, and LW conceived the study. SC developed and implemented the method and performed the analyses. LW supervised the study. SC and LW wrote the manuscript. All authors read and approved the final manuscript.

### Competing interests

The authors declare that they have no competing interests.

### Consent for publication

Not applicable.

### Ethics approval and consent to participate

Not applicable.

### Author details

<sup>1</sup>Department of Molecular Carcinogenesis, The Netherlands Cancer Institute, Plesmanlaan 121, 1066 CX Amsterdam, The Netherlands. <sup>2</sup>Department of Medical Oncology, Erasmus University Medical Center, Rotterdam, The Netherlands. <sup>3</sup>Faculty of EEMCS, Delft University of Technology, Delft, The Netherlands. <sup>4</sup>Cancer Genomics Netherlands, Amsterdam, The Netherlands.

Received: 17 October 2016 Accepted: 21 November 2016

Published online: 16 December 2016

## References

- Vogelstein B, Papadopoulos N, Velculescu VE, Zhou S, Diaz Jr LA, Kinzler KW. Cancer genome landscapes. *Science*. 2013;339(6127):1546–58. doi:10.1126/science.1235122.
- Lawrence MS, Stojanov P, Mermel CH, Robinson JT, Garraway LA, Golub TR, Meyerson M, Gabriel SB, Lander ES, Getz G. Discovery and saturation analysis of cancer genes across 21 tumour types. *Nature*. 2014;505(7484):495–501. doi:10.1038/nature12912.
- Ashworth A, Lord CJ, Reis-Filho JS. Genetic interactions in cancer progression and treatment. *Cell*. 2011;145(1):30–8. doi:10.1016/j.cell.2011.03.020.
- Thomas RK, Baker AC, DeBiasi RM, Winckler W, Laframboise T, Lin WM, Wang M, Feng W, Zander T, MacConaill L, Macconnaill LE, Lee JC, Nicoletti R, Hatton C, Goyette M, Girard L, Majmudar K, Ziaugra L, Wong KK, Gabriel S, Beroukhi R, Peyton M, Barretina J, Dutt A, Emery C, Greulich H, Shah K, Sasaki H, Gazdar A, Minna J, Armstrong SA, Mellinghoff IK, Hodi FS, Dranoff G, Mischel PS, Cloughesy TF, Nelson SF, Liaw LM, Mertz K, Rubin MA, Moch H, Loda M, Catalona W, Fletcher J, Signoretti S, Kaye F, Anderson KC, Demetri GD, Dummer R, Wagner S, Herlyn M, Sellers WR, Meyerson M, Garraway LA. High-throughput oncogene mutation profiling in human cancer. *Nat Genet*. 2007;39(3):347–51. doi:10.1038/ng1975.
- Yeang CH, McCormick F, Levine A. Combinatorial patterns of somatic gene mutations in cancer. *FASEB J*. 2008;22(8):2605–22. doi:10.1096/fj.08-108985.
- Ciriello G, Cerami E, Sander C, Schultz N. Mutual exclusivity analysis identifies oncogenic network modules. *Genome Res*. 2012;22(2):398–406. doi:10.1101/gr.125567.111.
- Vandin F, Upfal E, Raphael BJ. De novo discovery of mutated driver pathways in cancer. *Genome Res*. 2012;22(2):375–85. doi:10.1101/gr.120477.111.
- Szczurek E, Beerenwinkel N. Modeling mutual exclusivity of cancer mutations. *PLoS Comput Biol*. 2014;10(3):1003503. doi:10.1371/journal.pcbi.1003503.
- Babur Ö, Gönen M, Aksoy BA, Schultz N, Ciriello G, Sander C, Demir E. Systematic identification of cancer driving signaling pathways based on mutual exclusivity of genomic alterations. *Genome Biol*. 2015;16:45. doi:10.1186/s13059-015-0612-6.
- Leiserson MD, Wu HT, Vandin F, Raphael BJ. CoMET: a statistical approach to identify combinations of mutually exclusive alterations in cancer. *Genome Biol*. 2015;16:160. doi:10.1186/s13059-015-0700-7.
- Hua X, Hyland PL, Huang J, Song L, Zhu B, Caporaso NE, Landi MT, Chatterjee N, Shi J. MEGSA: A powerful and flexible framework for analyzing mutual exclusivity of tumor mutations. *Am J Hum Genet*. 2016;98(3):442–55. doi:10.1016/j.ajhg.2015.12.021.
- Constantinescu S, Szczurek E, Mohammadi P, Rahnenführer J, Beerenwinkel N. TIMEx: a waiting time model for mutually exclusive cancer alterations. *Bioinformatics*. 2016;32(7):968–75. doi:10.1093/bioinformatics/btv400.
- Bredel M, Scholtens DM, Harsh GR, Bredel C, Chandler JP, Renfrow JJ, Yadav AK, Vogel H, Scheck AC, Tibshirani R, Sikic BI. A network model of a cooperative genetic landscape in brain tumors. *JAMA*. 2009;302(3):261–75. doi:10.1001/jama.2009.997.
- Gorringe KL, George J, Anglesio MS, Ramakrishna M, Etemadmoghadam D, Cowin P, Sridhar A, Williams LH, Boyle SE, Yanaihara N, Okamoto A, Urashima M, Smyth GK, Campbell IG, Bowtell DD, Australian Ovarian Cancer Study. Copy number analysis identifies novel interactions between genomic loci in ovarian cancer. *PLoS ONE*. 2010;5(9):. doi:10.1371/journal.pone.0011408.
- Klijn C, Bot J, Adams DJ, Reinders M, Wessels L, Jonkers J. Identification of networks of co-occurring, tumor-related DNA copy number changes using a genome-wide scoring approach. *PLoS Comput Biol*. 2010;6(1):1000631. doi:10.1371/journal.pcbi.1000631.
- Milosevic JD, Puda A, Malcovati L, Berg T, Hofbauer M, Stukalov A, Klampfl T, Harutyunyan AS, Gisslinger H, Gisslinger B, Burjanivova T, Rumi E, Pietra D, Elena C, Vannucchi AM, Doubek M, Dvorakova D, Robesova B, Wieser R, Koller E, Suvajdzic N, Tomin D, Tosic N, Colinge J, Racil Z, Steurer M, Pavlovic S, Cazzola M, Kralovics R. Clinical significance of genetic aberrations in secondary acute myeloid leukemia. *Am J Hematol*. 2012;87(11):1010–6. doi:10.1002/ajh.23309.
- Kandoth C, McLellan MD, Vandin F, Ye K, Niu B, Lu C, Xie M, Zhang Q, McMichael JF, Wyczalkowski MA, Leiserson MD, Miller CA, Welch JS, Walter MJ, Wendl MC, Ley TJ, Wilson RK, Raphael BJ, Ding L. Mutational landscape and significance across 12 major cancer types. *Nature*. 2013;502(7471):333–9. doi:10.1038/nature12634.
- Remy E, Rebouissou S, Chaouiya C, Zinovyev A, Radvanyi F, Calzone L. A modeling approach to explain mutually exclusive and co-occurring genetic alterations in bladder tumorigenesis. *Cancer Res*. 2015;75(19):4042–52. doi:10.1158/0008-5472.CAN-15-0602.
- Pereira B, Chin SF, Rueda OM, Vollan HK, Provenzano E, Bardwell HA, Pugh M, Jones L, Russell R, Sammut SJ, Tsui DW, Liu B, Dawson SJ, Abraham J, Northen H, Peden JF, Mukherjee A, Turashvili G, Green AR, McKinney S, Oloumi A, Shah S, Rosenfeld N, Murphy L, Bentley DR, Ellis IO, Purushotham A, Pinder SE, Borresen-Dale AL, Earl HM, Pharoah PD, Ross MT, Aparicio S, Caldas C. The somatic mutation profiles of 2,433 breast cancers refine their genomic and transcriptomic landscapes. *Nat Commun*. 2016;7:11479. doi:10.1038/ncomms11479.
- Zack TI, Schumacher SE, Carter SL, Cherniack AD, Saksena G, Tabak B, Lawrence MS, Zhang CZ, Wala J, Mermel CH, Sougnez C, Gabriel SB, Hernandez B, Shen H, Laird PW, Getz G, Meyerson M, Beroukhi R. Pan-cancer patterns of somatic copy number alteration. *Nature Genet*. 2013;45(10):1134–40. doi:10.1038/ng.2760.
- Lancaster HO. Significance tests in discrete distributions. *J Am Stat Assoc*. 1961;56(294):223–34. doi:10.2307/2282247.
- The Cancer Genome Atlas Research Network, Weinstein JN, Collisson EA, Mills GB, Shaw KR, Ozenberger BA, Ellrott K, Shmulevich I, Sander C, Stuart JM. The Cancer Genome Atlas Pan-Cancer analysis project. *Nat Genet*. 2013;45(10):1113–20. doi:10.1038/ng.2764.
- Tamborero D, Gonzalez-Perez A, Perez-Llamas C, Deu-Pons J, Kandoth C, Reimand J, Lawrence MS, Getz G, Bader GD, Ding L, Lopez-Bigas N. Comprehensive identification of mutational cancer driver genes across 12 tumor types. *Sci Rep*. 2013;3:2650. doi:10.1038/srep02650.
- Subramanian A, Tamayo P, Mootha VK, Mukherjee S, Ebert BL, Gillette MA, Paulovich A, Pomeroy SL, Golub TR, Lander ES, Mesirov JP. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A*. 2005;102(43):15545–50. doi:10.1073/pnas.0506580102.
- Ulz P, Heitzer E, Speicher MR. Co-occurrence of MYC amplification and TP53 mutations in human cancer. *Nat Genet*. 2016;48(2):104–6. doi:10.1038/ng.3468.
- Szklarczyk D, Franceschini A, Wyder S, Forslund K, Heller D, Huerta-Cepas J, Simonovic M, Roth A, Santos A, Tsafou KP, Kuhn M, Bork P, Jensen LJ, von Mering C. STRING v10: protein-protein interaction networks, integrated over the tree of life. *Nucleic Acids Res*. 2015;43(Database issue):447–52. doi:10.1093/nar/gku1003.
- Rodriguez-Viciana P, Warne PH, Dhand R, Vanhaesebroeck B, Gout I, Fry MJ, Waterfield MD, Downward J. Phosphatidylinositol-3-OH kinase as a direct target of Ras. *Nature*. 1994;370(6490):527–32. doi:10.1038/370527a0.
- Gong Y, Zack TI, Morris LG, Lin K, Hukkelhoven E, Raheja R, Tan IL, Turcan S, Veeriah S, Meng S, Viale A, Schumacher SE, Palmado P, Beroukhi R, Chan TA. Pan-cancer genetic analysis identifies PARK2 as a master regulator of G1/S cyclins. *Nat Genet*. 2014;46(6):588–94. doi:10.1038/ng.2981.
- Yan H, Tong J, Lin X, Han Q, Huang H. Effect of the WWOX gene on the regulation of the cell cycle and apoptosis in human ovarian cancer stem cells. *Mol Med Rep*. 2015;12(2):1783–8. doi:10.3892/mmr.2015.3640.
- Sard L, Accornero P, Torielli S, Delia D, Bunone G, Campiglio M, Colombo MP, Gramegna M, Croce CM, Pierotti MA, Sozzi G. The tumor-suppressor gene FHIT is involved in the regulation of apoptosis and in cell cycle control. *Proc Natl Acad Sci U S A*. 1999;96(15):8489–92. doi:10.1073/pnas.96.15.8489.
- Wang D, Wang L, Zhou J, Pan J, Qian W, Fu J, Zhang G, Zhu Y, Liu C, Wang C, Jin Z, He Z, Wu J, Shi B. Reduced expression of PTPRD correlates with poor prognosis in gastric adenocarcinoma. *PLoS ONE*. 2014;9(11):113754. doi:10.1371/journal.pone.0113754.
- Solomon DA, Kim JS, Cronin JC, Sibenaller Z, Ryken T, Rosenberg SA, Ressler H, Jean W, Bigner D, Yan H, Samuels Y, Waldman T. Mutational inactivation of PTPRD in glioblastoma multiforme and malignant melanoma. *Cancer Res*. 2008;68(24):10300–6. doi:10.1158/0008-5472.CAN-08-3272.
- Zarubin T, Han J. Activation and signaling of the p38 map kinase pathway. *Cell Res*. 2005;15(1):11–8. doi:10.1038/sj.cr.7290257.
- Regard JB, Malhotra D, Gvozdenovic-Jeremic J, Josey M, Chen M, Weinstein LS, Lu J, Shore EM, Kaplan FS, Yang Y. Activation of



- Hedgehog signaling by loss of GNAS causes heterotopic ossification. *Nat Med.* 2013;19(11):1505–12. doi:10.1038/nm.3314.
35. He X, Zhang L, Chen Y, Remke M, Shih D, Lu F, Wang H, Deng Y, Yu Y, Xia Y, Wu X, Ramaswamy V, Hu T, Wang F, Zhou W, Burns DK, Kim SH, Kool M, Pfister SM, Weinstein LS, Pomeroy SL, Gilbertson RJ, Rubin JB, Hou Y, Wechsler-Reya R, Taylor MD, Lu QR. The G protein  $\alpha$  subunit  $G\alpha s$  is a tumor suppressor in Sonic hedgehog-driven medulloblastoma. *Nat Med.* 2014;20(9):1035–42. doi:10.1038/nm.3666.
  36. Takabatake Y, Takabatake T, Sasagawa S, Takeshima K. Conserved expression control and shared activity between cognate T-box genes Tbx2 and Tbx3 in connection with Sonic hedgehog signaling during *Xenopus* eye development. *Dev Growth Differ.* 2002;44(4):257–71. doi:10.1046/j.1440-169X.2002.00640.x.
  37. Kann M, Bae E, Lenz MO, Li L, Trannguyen B, Schumacher VA, Taglienti ME, Bordeianou L, Hartwig S, Rinschen MM, Schermer B, Benzing T, Fan CM, Kreidberg JA. WT1 targets Gas1 to maintain nephron progenitor cells by modulating FGF signals. *Development.* 2015;142(7):1254–66. doi:10.1242/dev.119735.
  38. Jagani Z, Mora-Blanco EL, Sansam CG, McKenna ES, Wilson B, Chen D, Klekota J, Tamayo P, Nguyen PT, Tolstorukov M, Park PJ, Cho YJ, Hsiao K, Buonamici S, Pomeroy SL, Mesirov JP, Ruffner H, Bouwmeester T, Luchansky SJ, Murtie J, Kelleher JF, Warmuth M, Sellers WR, Roberts CW, Dorsch M. Loss of the tumor suppressor Snf5 leads to aberrant activation of the Hedgehog-Gli pathway. *Nat Med.* 2010;16(12):1429–33. doi:10.1038/nm.2251.
  39. Gross AM, Orosco RK, Shen JP, Egloff AM, Carter H, Hofree M, Choueiri M, Coffey CS, Lippman SM, Hayes DN, Cohen EE, Grandis JR, Nguyen QT, Ideker T. Multi-tiered genomic analysis of head and neck cancer ties TP53 mutation to 3p loss. *Nat Genet.* 2014;46(9):939–43. doi:10.1038/ng.3051.
  40. Malinowska-Ozdowy K, Frech C, Schönegger A, Eckert C, Cazzaniga G, Stanulla M, zur Stadt U, Mecklenbräuer A, Schuster M, Kneidinger D, von Stackelberg A, Locatelli F, Schrappe M, Horstmann MA, Attarbaschi A, Bock C, Mann G, Haas OA, Panzer-Grümayer R. KRAS and CREBBP mutations: a relapse-linked malicious liaison in childhood high hyperdiploid acute lymphoblastic leukemia. *Leukemia.* 2015;29(8):1656–67. doi:10.1038/leu.2015.107.
  41. Leek JT, Storey JD. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet.* 2007;3(9):161. doi:10.1371/journal.pgen.0030161.
  42. Yu J, Pressoir G, Briggs WH, Vroh Bi I, Yamasaki M, Doebley JF, McMullen MD, Gaut BS, Nielsen DM, Holland JB, Kresovich S, Buckler ES. A unified mixed-model method for association mapping that accounts for multiple levels of relatedness. *Nat Genet.* 2006;38(2):203–8. doi:10.1038/ng1702.
  43. Wang YH. On the number of successes in independent trials. *Statistica Sinica.* 1993;3(2):295–312.
  44. Liu DC, Nocedal J. On the limited memory BFGS method for large scale optimization. *Math Prog.* 1989;45(1-3):503–28. doi:10.1007/BF01589116.
  45. Carlson JM, Heckerman D, Shani G. Estimating false discovery rates for contingency tables. Technical Report MSR-TR-2009-53. 2009. <http://research.microsoft.com/apps/pubs/default.aspx?id=80571>.
  46. Bonchi F, Gionis A, Ukkonen A. Overlapping correlation clustering. *Knowl Inf Syst.* 2013;35(1):1–32. doi:10.1007/s10115-012-0522-9.
  47. Chang YT, Leahy RM, Pantazis D. Modularity-based graph partitioning using conditional expected models. *Phys Rev E.* 2012;85(1):016109. doi:10.1103/PhysRevE.85.016109.

Submit your next manuscript to BioMed Central and we will help you at every step:

- We accept pre-submission inquiries
- Our selector tool helps you to find the most relevant journal
- We provide round the clock customer support
- Convenient online submission
- Thorough peer review
- Inclusion in PubMed and all major indexing services
- Maximum visibility for your research

Submit your manuscript at  
[www.biomedcentral.com/submit](http://www.biomedcentral.com/submit)

