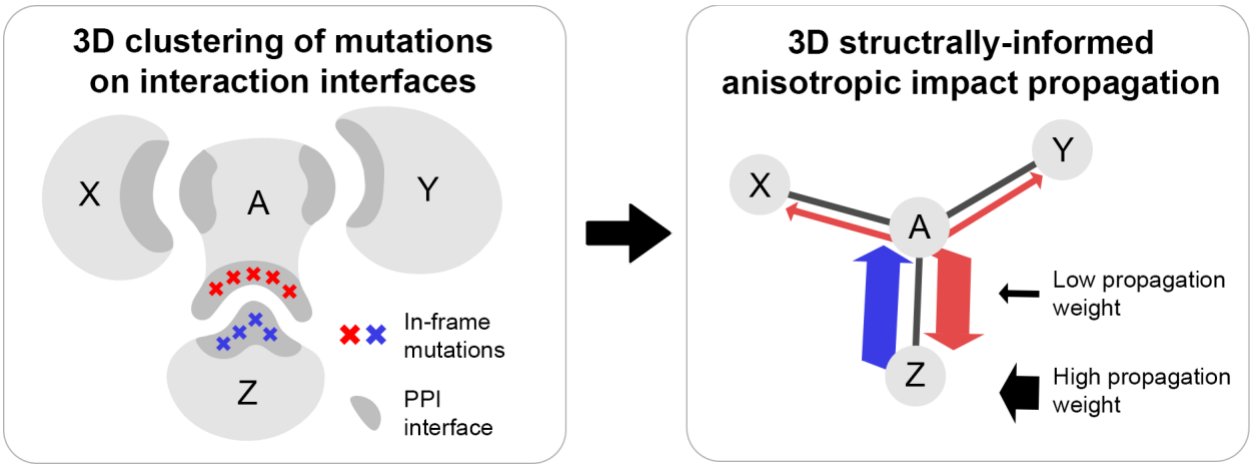

Supplementary Fig. 1. NetFlow3D's weighted propagation strategy. NetFlow3D recognizes that proteins often interact with different partners using distinct three-dimensional (3D) structural interfaces. It, therefore, uses this information to weight the impact of spatially clustered mutations at a specific protein-protein interaction (PPI) interface on different interaction partners differently (anisotropic).

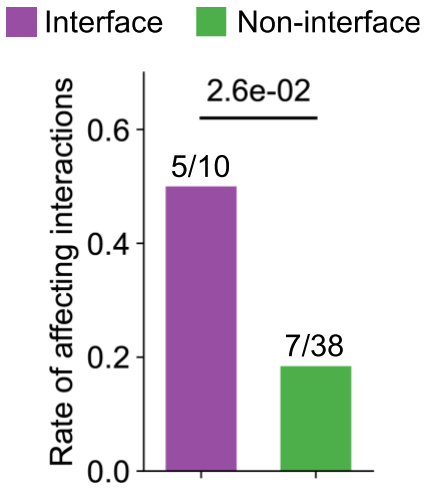

Supplementary Fig. 2. The edgetic effect of somatic cancer driver mutations. To validate the "edgetic effect" of functional somatic in-frame mutations, we experimentally generated mutations within the significant 3D clusters identified by NetFlow3D, and tested their effects on protein-protein interactions via yeast two-hybrid (Y2H) assay (Supplementary Note 5-6). All Y2H experiments were repeated 3 times. As a result, 5 out of 10 measured interactions were disrupted by mutations on their interaction interfaces, and 7 out of 38 were disrupted by non-interface mutations. The results demonstrated that spatially clustered mutations at interaction interfaces exhibited a significantly higher rate of interaction-perturbing effects compared to non-interface mutations ($p = 2.6e\text{-}2$ by a two-sided two-sided two-proportion Z-test), thereby validating their "edgetic effect". Source data are provided as a Source Data file.

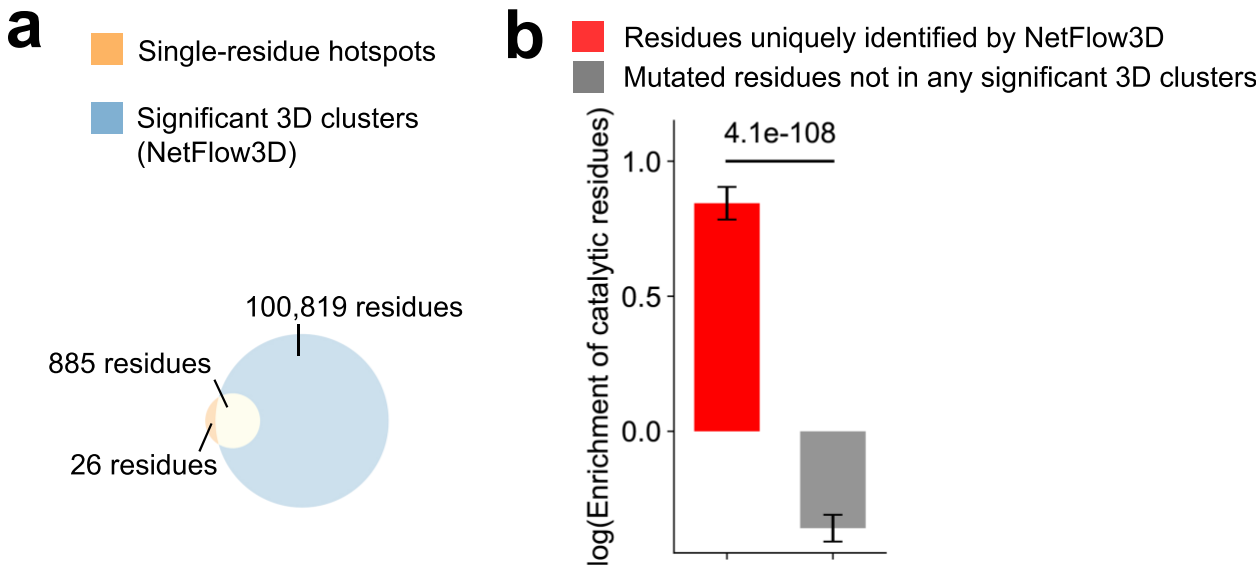

Supplementary Fig. 3. Comparison of 3D clustering results from NetFlow3D with single-residue hotspots detected by the recurrence-based methods. (**a**) Overlap between residues within significant 3D clusters identified by NetFlow3D and those identified as hotspots by the single-residue-based method. For visual clarity, the area proportions in the Venn diagram are scaled based on the square root of the residue count. (**b**) Enrichment analysis of catalytic residues focusing on the residues uniquely identified by NetFlow3D. Enrichment was calculated as the ratio of the observed fraction of catalytic residues among residues under investigation over the fraction of catalytic residues on corresponding proteins (expected fraction). The error bars indicate standard error, calculated using the delta method. P values for each bar were calculated using two-sided Z-tests. Residues uniquely identified by NetFlow3D: n=100,819; statistic (z)=14.0; p=1.4e-44; enrichment=2.3; 95% CI=[2.1, 2.6]. Mutated residues not in any significant 3D clusters: n=682,471; statistic (z)=-7.2; p=4.7e-13; enrichment=0.70; 95% CI=[0.63, 0.77]. The p-value for comparing the observed fractions between the two groups was calculated using a two-sided two-proportion Z-test. Source data are provided as a Source Data file.

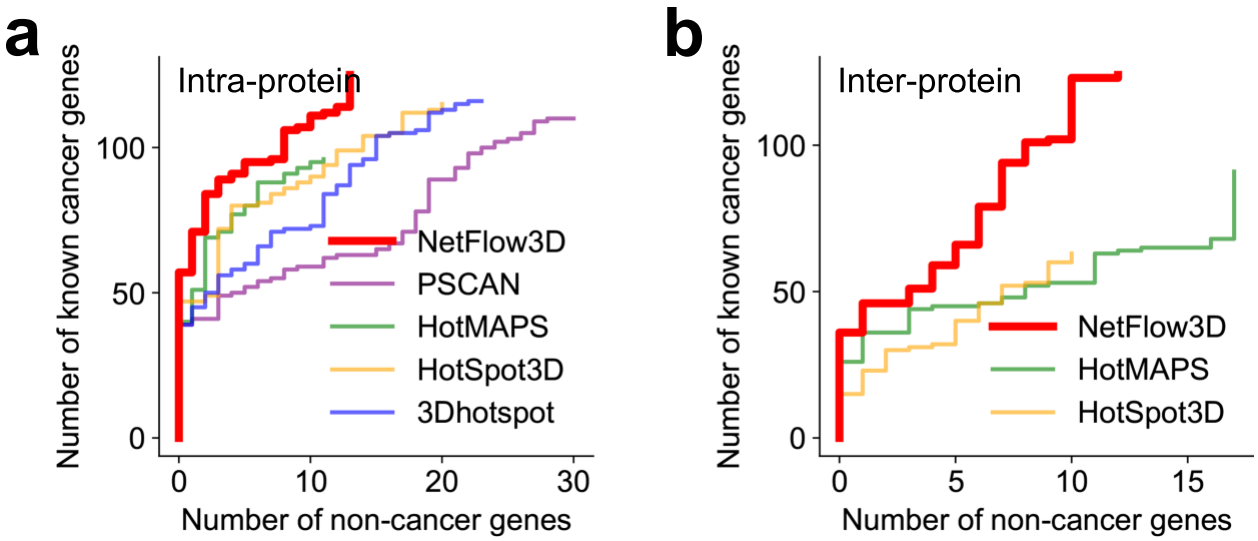

Supplementary Fig. 4. Performance comparison between NetFlow3D and state-of-the-art 3D clustering algorithms on a COSMIC pan-cancer dataset completely independent from TCGA. Performance curves are drawn for the top 1-500 genes, ranked by each algorithm based on the highest scoring 3D cluster on each gene. Source data are provided as a Source Data file. (**a**) Intra-protein 3D clustering results. (**b**) Inter-protein 3D clustering results.

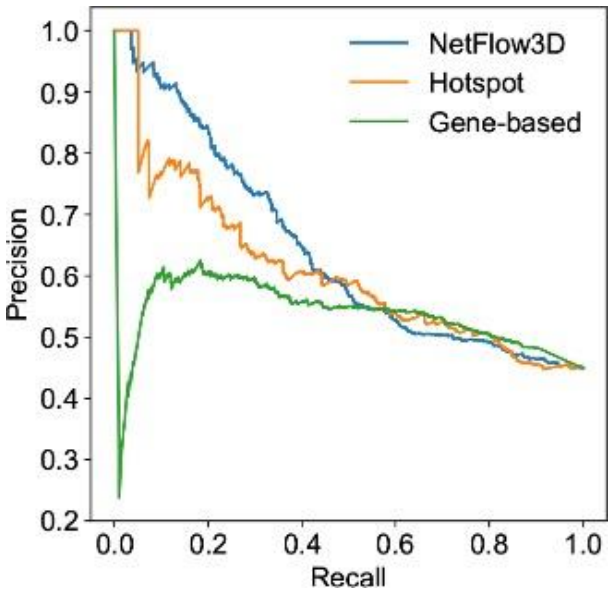

Supplementary Fig. 5. Precision-recall (PR) curves of three methods for identifying cancer driver mutations. Three methods focus on different test units: (i) 3D cluster (NetFlow3D), (ii) single residue (Hotspot), and (iii) whole gene (Gene-based). Gene-level precision-recall curves were generated using known cancer genes from the Cancer Gene Census (CGC) as positive labels and known non-cancer genes curated from the literature as negative labels. Genes with unknown labels were excluded to prevent potential misclassification. For the Hotspot and NetFlow3D methods, the most significant residue or 3D cluster was selected as the representative for each gene. Source data are provided as a Source Data file.

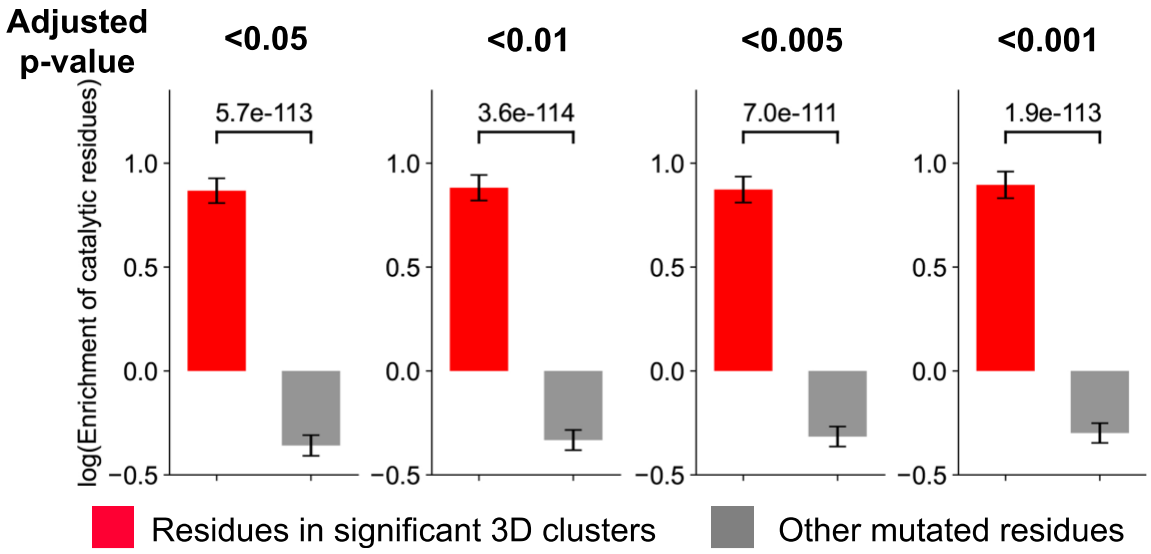

Supplementary Fig. 6. Enrichment analysis of catalytic residues within significant 3D clusters, defined across varying statistical thresholds. Enrichment was calculated as the ratio of the observed fraction of catalytic residues among the residues under investigation over the fraction of catalytic residues on corresponding proteins (expected fraction). The error bars indicate standard error, calculated using the delta method. P values for each bar were calculated using two-sided Z-tests. The p-value for comparing the observed fractions between the two groups was calculated using a two-sided two-proportion Z-test. Adjusted $p<0.05$: n=101,704 residues in significant 3D clusters; n=682,471 other mutated residues. Adjusted $p<0.01$: n=91,199 residues in significant 3D clusters; n=692,976 other mutated residues. Adjusted $p<0.005$: n=87,368 residues in significant 3D clusters; n=696,807 other mutated residues. Adjusted $p<0.001$: n=79,424 residues in significant 3D clusters; n=704,751 other mutated residues. Full statistical details are in the source data. Source data are provided as a Source Data file.

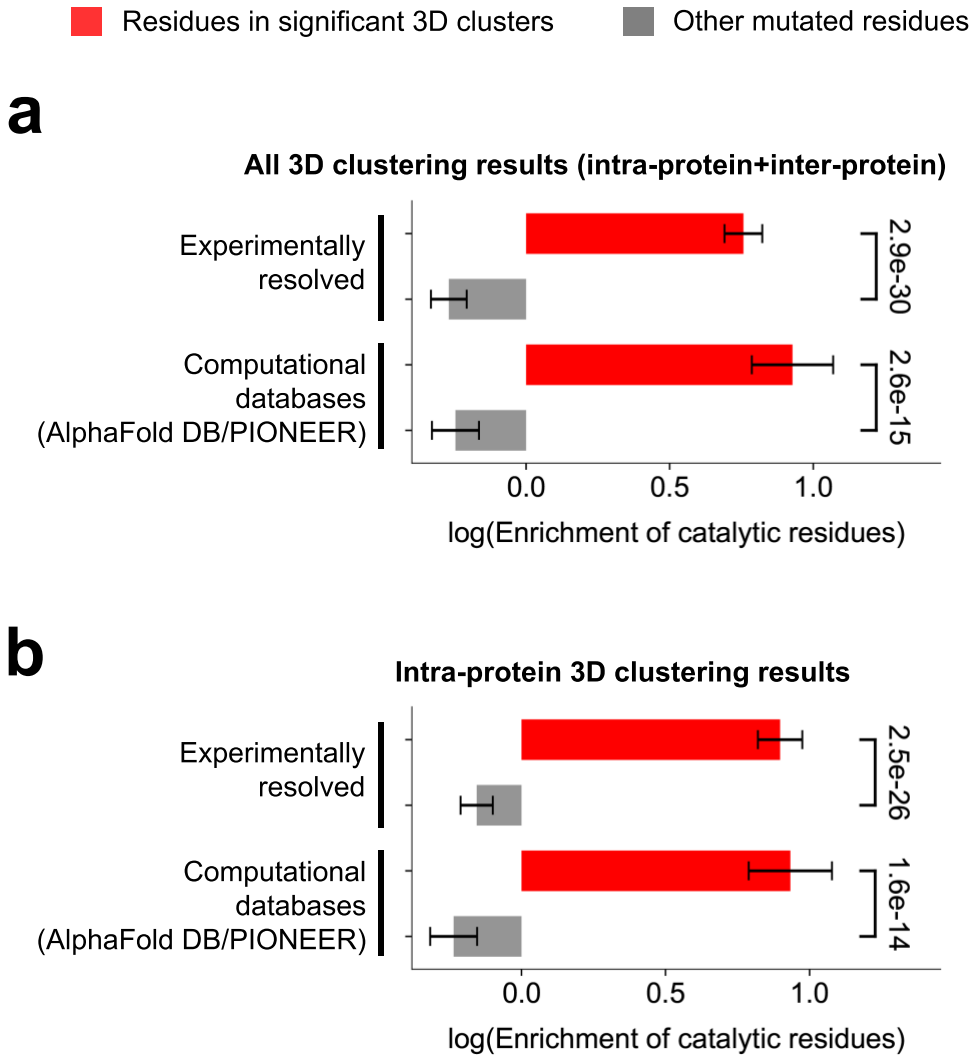

Supplementary Fig. 7. Analysis of catalytic residue enrichment conducted distinctly on 3D clustering results derived from experimentally resolved structures and deep-learning-generated 3D structural data. Enrichment was calculated as the ratio of the observed fraction of catalytic residues among the residues under investigation over the fraction of catalytic residues on corresponding proteins (expected fraction). The error bars indicate standard error, calculated using the delta method. P values for each bar were calculated using two-sided Z-tests. The p-value for comparing the observed fractions between the two groups was calculated using a two-sided two-proportion Z-test. Full statistical details are in the source data. Source data are provided as a Source Data file. (**a**) Results from both intra-protein and inter-protein 3D clustering analyses. Proteins with experimentally resolved structures: n=71,489 residues in significant 3D clusters; n=220,299 other mutated residues. Proteins with only computationally predicted 3D structural data: n=30,215 residues in significant 3D clusters; n=386,898 other mutated residues. (**b**) Results solely from intra-protein 3D clustering analysis. Proteins with experimentally resolved structures: n=45,446 residues in significant intra-protein 3D clusters; n=246,342 other mutated residues. Proteins with only computationally predicted 3D structural data: n=29,363 residues in significant intra-protein 3D clusters; n=387,750 other mutated residues.

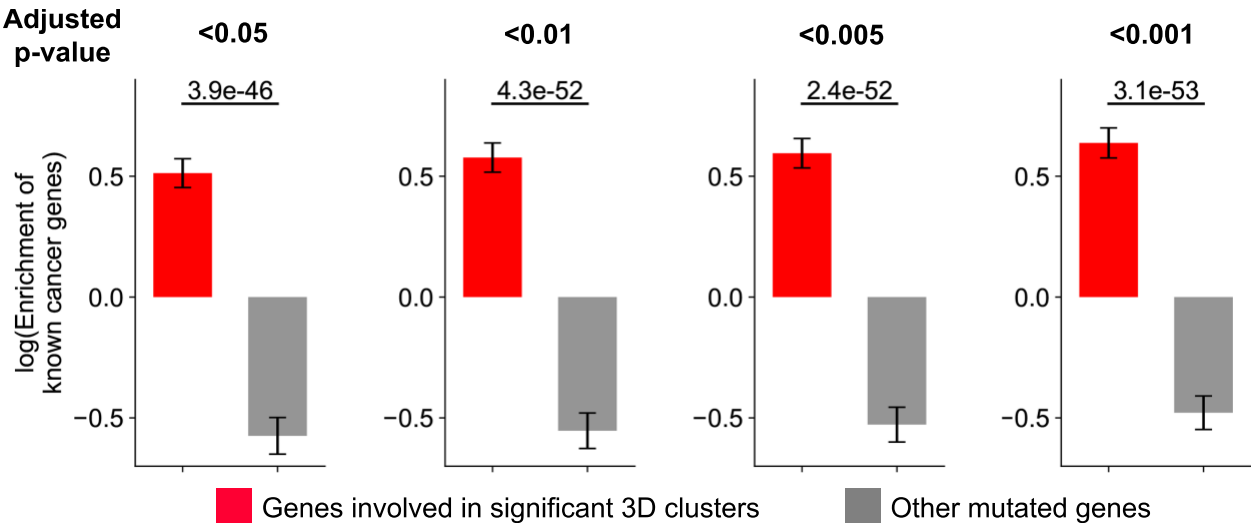

Supplementary Fig. 8. Enrichment analysis of known cancer genes in mutated proteins with or without significant 3D clusters, defined across varying statistical thresholds. Enrichment was calculated as the ratio of the observed fraction of known cancer genes among the genes of interest over the fraction of known cancer genes among all mutated genes (expected fraction). The error bars indicate standard error, calculated using the delta method. P values for each bar were calculated using two-sided Z-tests. The p-value for comparing the observed fractions between the two groups was calculated using a two-sided two-proportion Z-test. Adjusted $p<0.05$: n=5,698 genes involved in significant 3D clusters; n=8,738 other mutated genes. Adjusted $p<0.01$: n=5,085 genes involved in significant 3D clusters; n=9,351 other mutated genes. Adjusted $p<0.005$: n=4,839 genes involved in significant 3D clusters; n=9,597 other mutated genes. Adjusted $p<0.001$: n=4,315 genes involved in significant 3D clusters; n=10,121 other mutated genes. Full statistical details are in the source data. Source data are provided as a Source Data file.

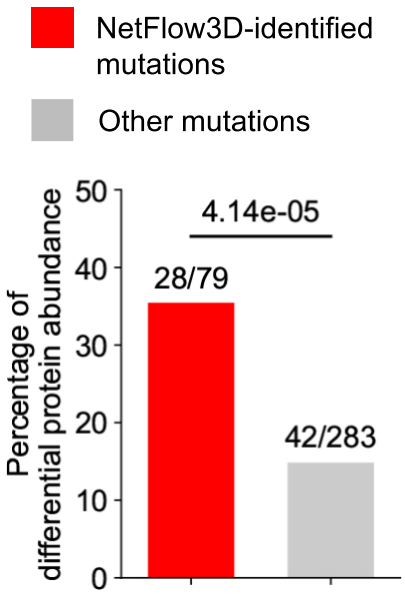

Supplementary Fig. 9. Differential protein abundance analysis per gene and cancer type in scenarios with and without mutations of interest. For each gene in each cancer type, the analysis required both scenarios to have at least 10 samples with available protein expression quantification data. P-values were calculated using two-sided Mann-Whitney U tests, with $p<0.05$ considered indicative of differential protein abundance. The analysis was conducted for NetFlow3D-identified mutations (red) and other mutations (gray). A gene-cancer pair refers to the analysis of a specific gene within a specific cancer type. NetFlow3D-identified mutations led to differential protein abundance in 28 out of 79 gene-cancer pairs, while other mutations affected 42 out of 283 gene-cancer pairs. These results show that NetFlow3D-identified mutations exhibit a significantly higher rate of protein-abundance perturbation compared to other mutations ($p = 4.14e-5$, two-sided two-proportion Z-test). Source data are provided in the Source Data file.

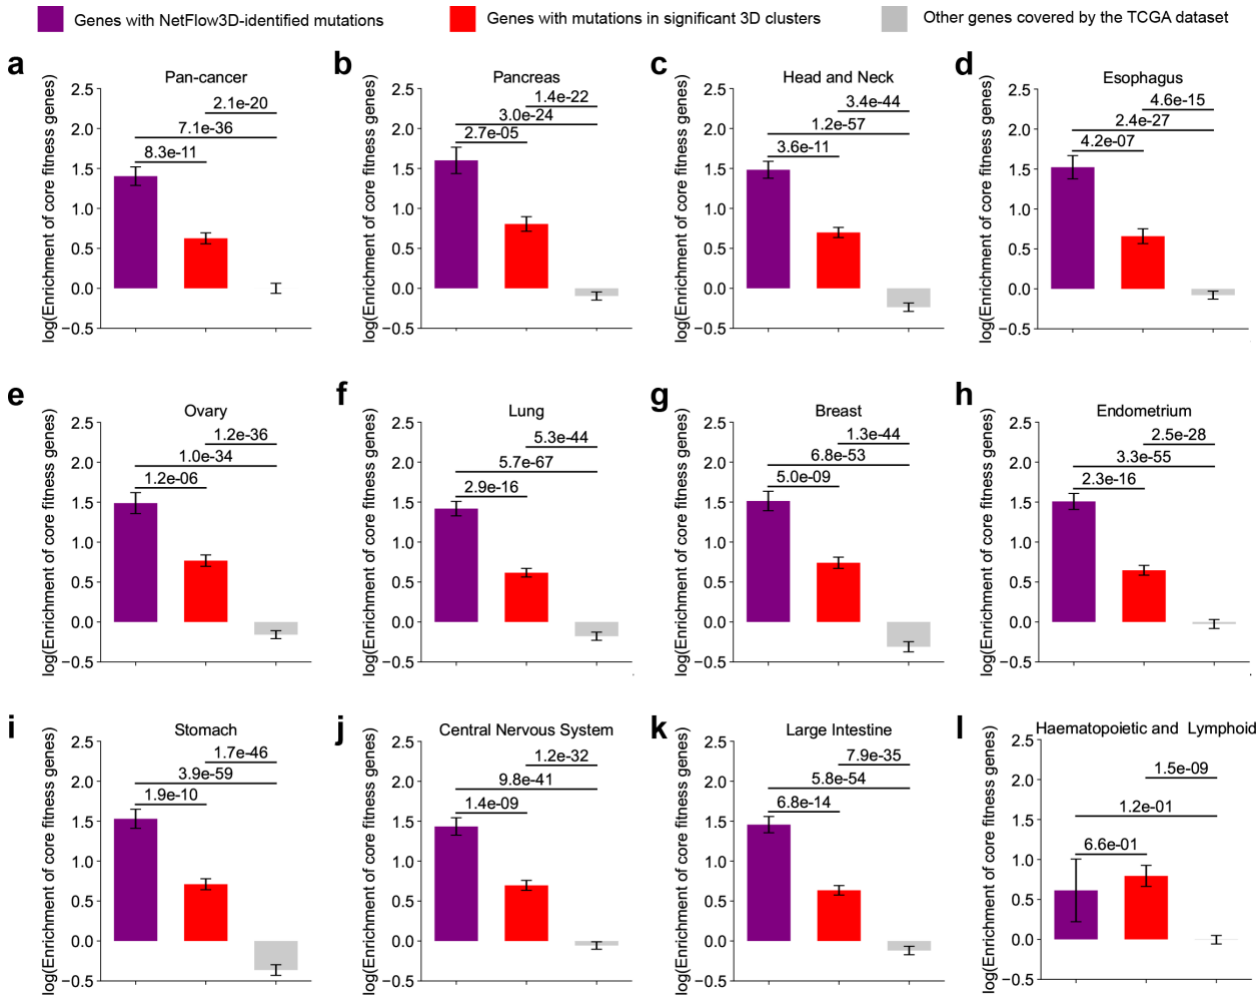

Supplementary Fig. 10. Analysis of core fitness gene enrichment in NetFlow3D results. Pan-cancer and cancer-type-specific core fitness genes were identified by Behan et al.[1] using genome-scale CRISPR-Cas9 screening data in human cancer cell lines. Enrichment was calculated as the ratio of the observed fraction of core fitness genes among the genes under investigation over the fraction of core fitness genes among all genes covered by the TCGA dataset (expected fraction). The error bars indicate standard error, calculated using the delta method. P values for each bar were calculated using two-sided Z-tests. The p-values for comparisons between the observed fractions of core fitness genes in different groups were calculated using two-sided two-proportion Z-tests. Full statistical details are in the source data. Source data are provided as a Source Data file. (**a**) Pan-cancer analysis using pan-cancer core fitness genes from Behan et al. and all mutations identified by NetFlow3D. (**b**-**l**) Cancer-type-specific analysis using cancer-type-specific core fitness genes from Behan et al. and a subset of NetFlow3D-identified mutations that are in the corresponding cancer type.

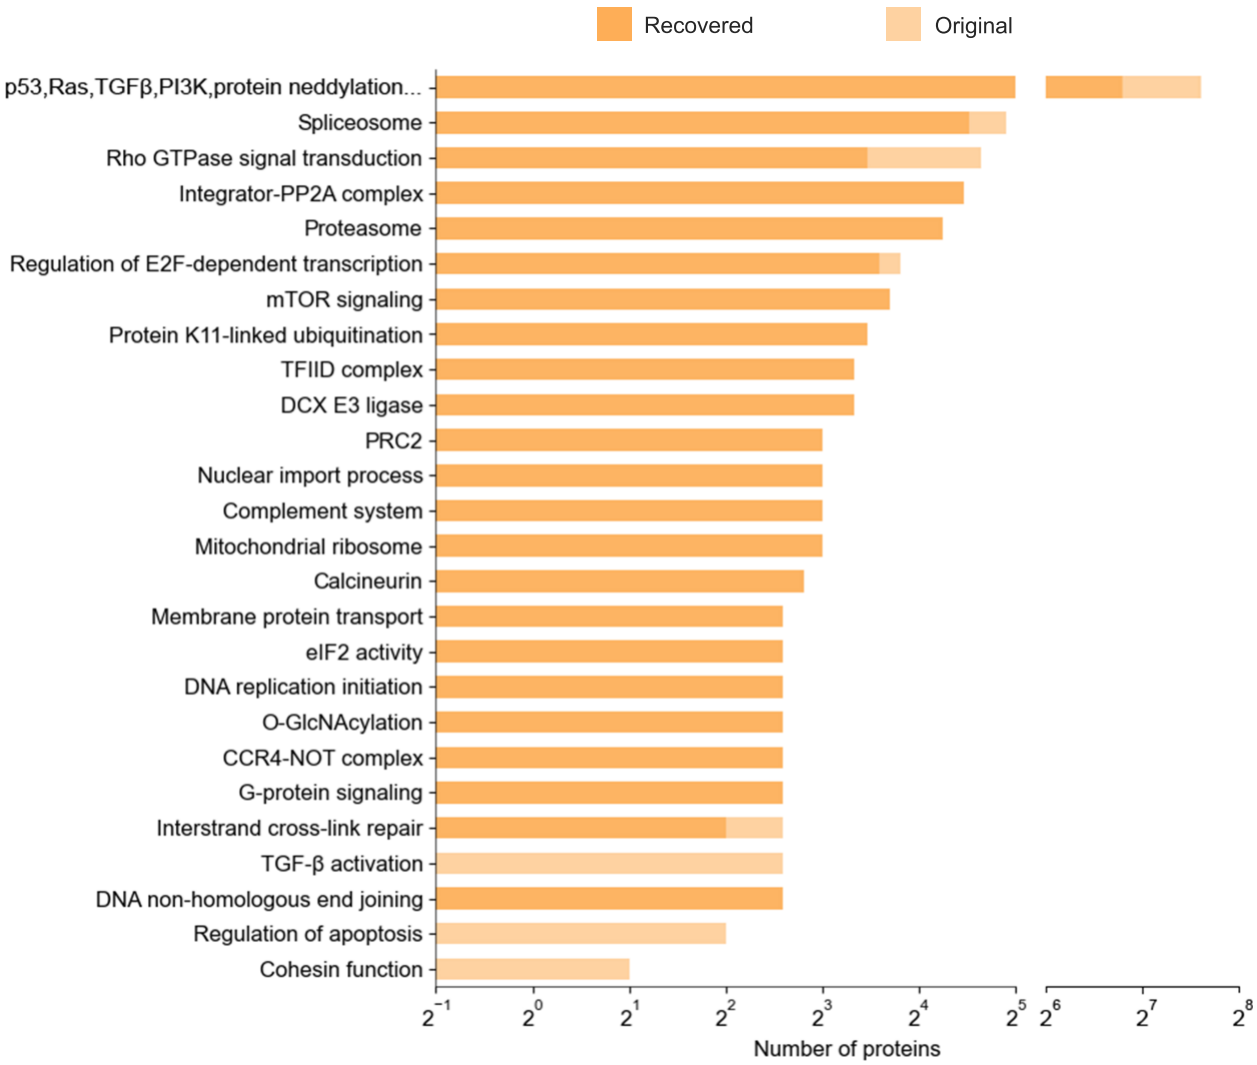

Supplementary Fig. 11. Evaluation of NetFlow3D findings beyond known cancer genes. Mutation signals from known cancer genes were completely removed before re-applying the network propagation model. The light orange bars indicates the number of proteins within each original NetFlow3D module that are not known cancer genes, and the dark orange bars indicate those that were recovered by the newly-identified significantly interconnected modules. Source data are provided as a Source Data file.

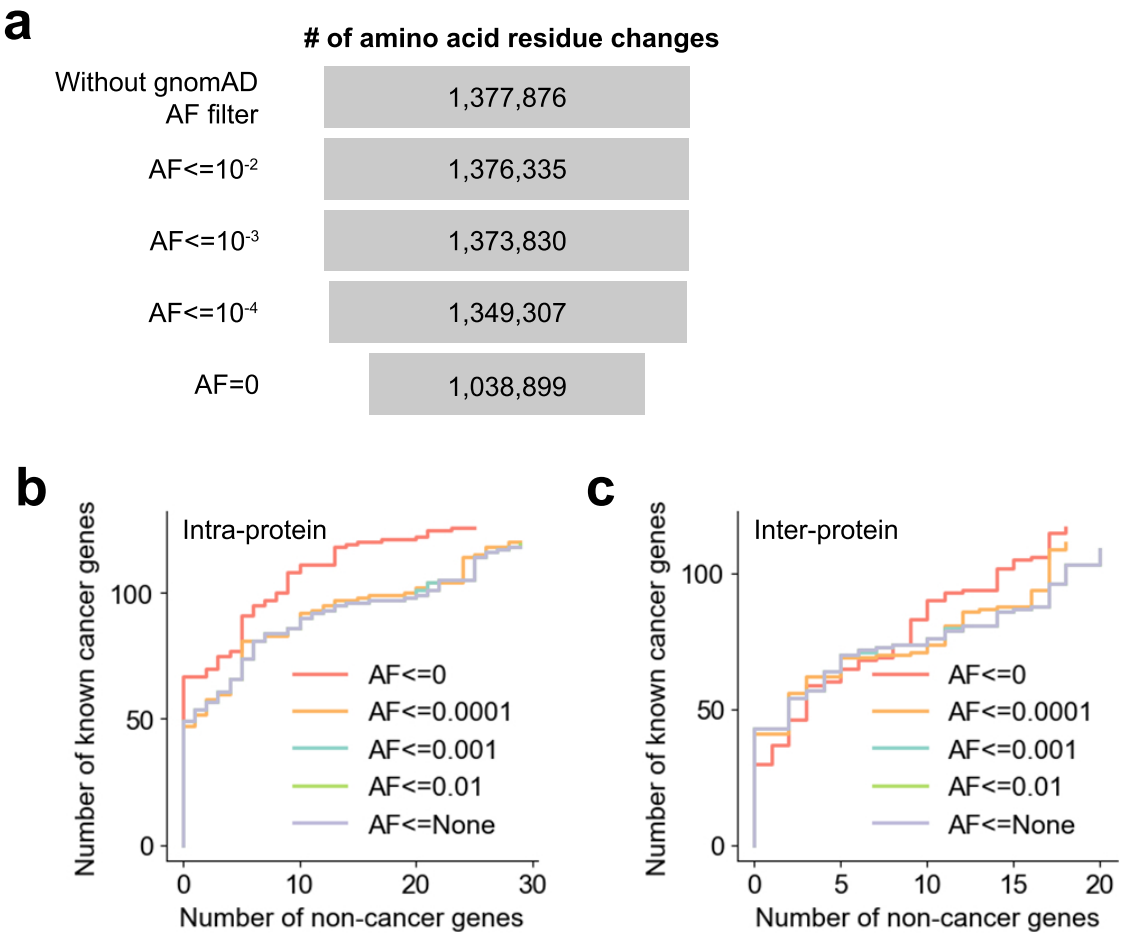

Supplementary Fig. 12. Evaluating the outcomes of applying a range of germline variant filters at different gnomAD allele frequency (AF) thresholds. gnomAD: Genome Aggregation Database. (**a**) Number of preprocessed mutations using each germline variant filter. (**b-c**) After applying NetFlow3D to datasets generated from each germline variant filter, performance curves were plotted for the top 1-500 genes, ranked based on the highest-scoring 3D cluster per gene. Source data are provided as a Source Data file. (**b**) Intra-protein 3D clustering results. (**c**) Inter-protein 3D clustering results.

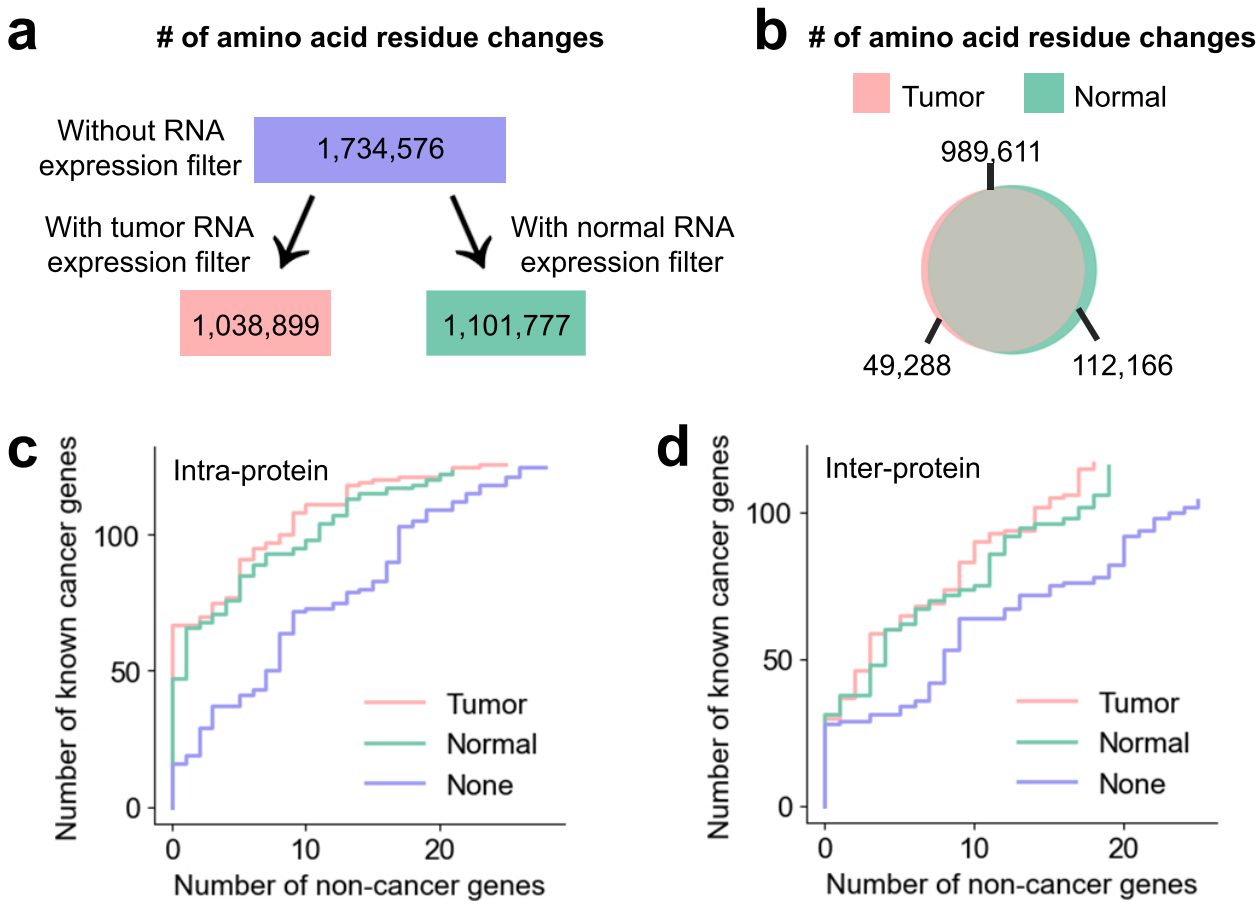

Supplementary Fig. 13. Consequences of three RNA expression filtering approaches. We evaluated the consequences of applying a tumor RNA expression filter (used by NetFlow3D), a normal RNA expression filter (only retained the mutations in those genes with RNA expression levels >=1 FPKM in >=80% of normal tissue samples near the tumors of the corresponding cancer types), and the scenario where no RNA expression filter is applied. (**a**) Number of preprocessed mutations under each filtering scenario. (**b**) A high degree of overlap in the mutations that were retained after applying the tumor and normal RNA expression filter. (**c-d**) After applying NetFlow3D to datasets generated from each RNA filtering approach, performance curves were plotted for the top 1-500 genes, ranked based on the highest-scoring 3D cluster per gene. The results showed that both the tumor and normal RNA expression filters significantly improved performance. Source data are provided as a Source Data file.

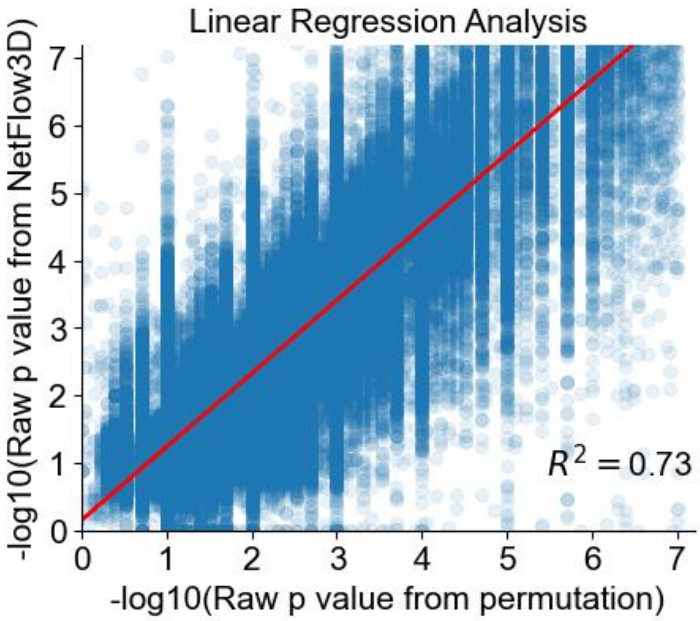

Supplementary Fig. 14. Correlation analysis of 3D cluster p-values derived from permutation tests and NetFlow3D. Each dot corresponds to a 3D cluster. Source data are provided as a Source Data file.

Supplementary Note 1: Linear regression analysis of protein abundance

TCGA proteome profiling data was obtained from Repository in the GDC data portal[2] (https://portal.gdc.cancer.gov/). Protein abundance for each gene in a sample was defined as the maximum quantification value among all its protein peptides. We used a linear regression model to evaluate the statistical association between NetFlow3D-identified mutations and protein abundance. This model controlled for gene-specific and tissue-specific baseline expression levels, as well as clinical covariates including sex, age, tumor stage, and TMB (Model: $\text{protein abundance} \sim \text{C(presence of NetFlow3D identified mutations)} + \text{C(gene ID)} + \text{C(cancer type)} + \text{age} + \text{C(sex)} + \text{TMB} + \text{C(tumor stage)}$. As a negative benchmark, we repeated the analysis by replacing $\text{C(presence of NetFlow3D identified mutations)}$ with $\text{C(presence of other mutations)}$.

Supplementary Note 2: Re-applying network propagation model after excluding the mutation signals from known cancer genes

3D clusters composed exclusively of mutations in known cancer genes, as well as LOF enrichment signals within known cancer genes, were excluded from the analysis. Subsequently, initial heat scores and propagation weights were re-calculated following equations (13)-(16) (Methods), and the network propagation model was re-applied. We then compared the newly-identified significantly interconnected modules to the original NetFlow3D modules to assess the recovery rate of the initial NetFlow3D modules.

Supplementary Note 3: Proteoform diversity

To evaluate the impact of incorporating proteoform diversity, we applied our revised NetFlow3D framework to the same TCGA mutation dataset, this time mapping mutations to all possible isoforms. This thorough mapping approach yielded 3.5 times more amino acid residue changes (2.95M) than the canonical mapping approach (0.85M). We conducted 3D clustering analysis using AlphaFold DB 3D structures, which encompass ~186k human protein isoforms. This analysis was not applicable to PDB or PIONEER, give that 96% (7,195 out of 7,515) of human UniProt entries covered by PDB are canonical isoforms, and 95% (113,949 out of 119,526) of high-quality binary human PPIs in HINT are between canonical isoforms. The 3D clustering results are provided in Supplementary Data 11-12. Our analysis revealed that for 84% of genes (12,167 out of 14,455), the most significant 3D cluster was still identified from their canonical isoforms.

Supplementary Note 4: Permutation tests

To benchmark the p-values of 3D clusters in NetFlow3D via permutation tests, we conducted 10^7 permutations. In each permutation, we randomly shuffled the mutations per patient per protein, and then counted how many patients had mutations within each 3D cluster. These permutations gave us empirical distributions and p-values.

Supplementary Note 5: Mutagenesis experiments

Single-colony-derived mutant clones were constructed using a high-throughput mutagenesis and next-generation sequencing pipeline called Clone-seq[3]. Wild-type clones, sourced from the hORFeome v8.1 collection[4], were used as the template for site-directed mutagenesis conducted by Eurofins Scientific. Mutagenesis was performed at 96-well scales using site-specific mutagenesis primers and full-length human ORF templates. PCR product was digested overnight using DpnI

(NEB) without a ligation step to maximize throughput and then transformed directly into competent cells to isolate single colonies. Then, 4 colonies per mutagenesis reaction were hand-picked and arrayed into 96-well plates. After 21 h incubation at 37 °C, glycerol stocks were generated and then clones were pooled into 4 respective bacterial pools. Maxiprepped DNAs from each of the 4 pools were then combined through multiplexing (NEBNext) and then sequenced in a single 1 × 100 single-end Illumina HiSeq run. Properly mutated clones were then identified by next-generation sequencing analysis and recovered from single-colony glycerol stocks. In total, we generated individual clones for 38 mutations within significant 3D clusters identified by NetFlow3D.

Supplementary Note 6: Y2H assay

We tested the mutation effects on protein-protein interactions via our high-throughput Y2H assay. To perform Y2H, pDEST-AD and pDEST-DB plasmid vectors corresponding to the GAL4-activating domain (AD) and DNA-binding (DB) domain, respectively, were used. Full-length Clone-seq-identified mutant clones were transferred into Y2H-amenable pDEST-DB and pDEST-AD vectors by Gateway LR reactions and then transformed into *MATα* Y8930 and *MAT**a*** Y8800, respectively. All DB-ORF *MATα* transformants, including wild-type ORFs, were then mated against corresponding wild-type and mutant AD-ORF *MAT**a*** transformants in a pairwise orientation on YEPD agar plates. After mating, yeast was replica-plated onto selective SC–Leu–Trp–His+1 mM of 3-amino-1,2,4-triazole (3AT) as well as SC–Leu–Trp–Adenine plates. Interactions were scored after three days of incubation and five days of incubation for SC–Leu–Trp+3AT and SC–Leu–Trp–Ade plates, respectively. To screen out autoactivating DB-ORFs, all DB-ORF *MATα* transformants were also mated pairwise against empty pDEST-AD *MAT**a*** transformants and scored for growth on SC–Leu–Trp+3AT and SC–Leu–Trp–Ade plates. DB-ORFs that trigger reporter activity under this setup were removed from further experiments. In total, we screened 48 mutation-interaction pairs, including 10 pairs where mutations were at the interaction interfaces and 38 pairs where mutations were not.

Reference

1. Behan, F. M. *et al.* Prioritization of cancer therapeutic targets using CRISPR-Cas9 screens. *Nature* **568**, 511–516 (2019).
2. Grossman, R. L. *et al.* Toward a Shared Vision for Cancer Genomic Data. *N. Engl. J. Med.* **375**, 1109–1112 (2016).
3. Wei, X. *et al.* A massively parallel pipeline to clone DNA variants and examine molecular phenotypes of human disease mutations. *PLoS Genet.* **10**, e1004819 (2014).
4. Yang, X. *et al.* A public genome-scale lentiviral expression library of human ORFs. *Nat. Methods* **8**, 659–661 (2011).