INSPIRE: A Network-Based Method for Quantifying Functional Convergence in Cancer Driver Genes

Charles Jacobs[1], George C. Linderman[2,3], Jakob A. Woerner[4], Hyun Jo Kim[5], Gurkan Bebek[5,6,7†]

[1]*Department of Biochemistry, Case Western Reserve University, Cleveland, Ohio, USA.*
[2]*Department of Biomedical Engineering, Case Western Reserve University, Cleveland, Ohio, USA.*
[3] *Applied Mathematics Program, Yale University, New Haven, CT, USA.*
[4]*Genomics and Computational Biology, University of Pennsylvania, Philadelphia, PA, USA.*
[5]*Department of Nutrition, Case Western Reserve University, Cleveland, Ohio, USA.*
[6] *Center for Proteomics and Bioinformatics, Case Western Reserve Univ., Cleveland, Ohio, USA.*
[7]*Department of Computer and Data Sciences, Case Western Reserve Univ., Cleveland, Ohio, USA.*
[†]*E-mail: gurkan.bebek@case.edu*

Recent genomic studies have identified numerous cancer driver genes, many of which occur at low frequencies across patient populations. This long tail of rare driver events suggests that tumorigenesis can arise through alterations in distinct genes that converge on common functional pathways. Characterizing these relationships requires approaches that extend beyond individual gene properties to the network contexts in which driver genes operate. Here, we investigate whether functional similarities among cancer drivers are better captured by comparing the signaling networks they perturb than by conventional gene-centric measures. To address this question, we developed INSPIRE (Integrative Network Similarity PIpEline), a framework for quantifying similarity among within-species biological networks. Applying INSPIRE to the signaling subnetworks of 28 colorectal cancer driver genes mapped topological similarity to downstream functional convergence, clustering driver-gene pathways into distinct, biologically coherent modules based on network architecture. Independent validation using large-scale CRISPR dependency datasets demonstrated that genes assigned to the same INSPIRE cluster exhibited significantly stronger co-essentiality than genes from different clusters ($p - value < 0.038$, r_{rb}=0.567), supporting the functional relevance of the network-derived groupings. Exploratory epistasis analyses identified additional candidate interactions within these clusters. At the same time, aggregation of drivers into INSPIRE-defined functional modules revealed significant network-level mutual exclusivity (FDR = 0.0149) that was not apparent at the individual-gene level. Together, these results demonstrate that signaling-network architecture captures functional convergence among genetically heterogeneous driver genes and provides a complementary framework for studying pathway organization in cancer.

Keywords: Network Similarity, Cancer Driver Genes, Protein Interaction Networks, Functional Convergence, Colorectal Cancer

Availability: Source code and data are available at github.com/bebeklab/inspire.

© 2026 The Authors. Open Access chapter published by World Scientific Publishing Company and distributed under the terms of the Creative Commons Attribution Non-Commercial (CC BY-NC) 4.0 License.

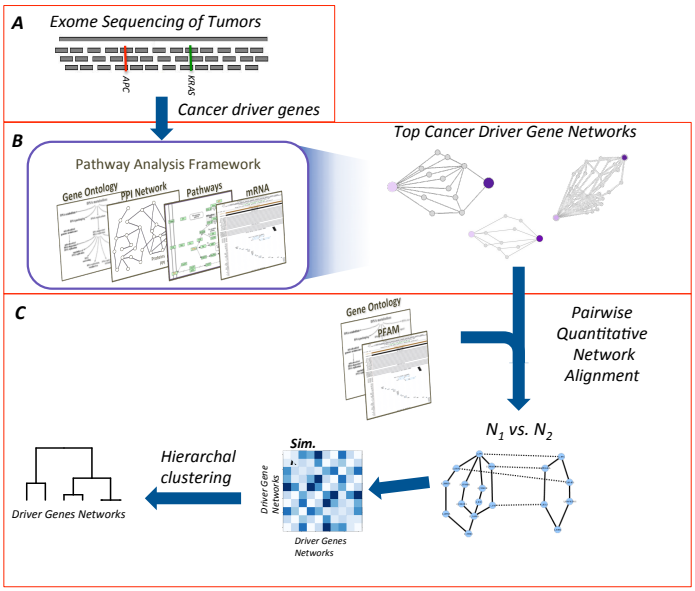


Fig. 1. **Overview of the study design and workflow.** (A) Identification of candidate cancer driver genes exhibiting co-occurrence or mutual exclusivity patterns in patient tumors. (B) Generation of functional signaling subnetworks connecting colorectal cancer (CRC) driver genes to the initiating driver *APC*.[6] (C) Schematic of the `INSPIRE` pipeline: pairwise subnetwork alignments are computed using protein family (Pfam) and Gene Ontology (GO) features to yield similarity scores, which are hierarchically clustered to identify functionally related driver genes based on network topology.

1. Introduction

Large-scale sequencing studies have identified numerous cancer driver genes across many tumor types.[1–4] Despite these advances, understanding how genetically diverse driver genes converge on common biological processes remains a major challenge.

A key challenge in cancer genomics is understanding how driver mutations give rise to malignant phenotypes. Although some driver genes are highly recurrent, many established drivers occur at modest frequencies, suggesting that distinct mutations may converge on common cellular processes. Numerous network-based approaches have therefore been developed to identify disease-relevant subnetworks.[5–10] In colorectal cancer, many well-established drivers beyond *APC* are individually mutated in fewer than 20% of patients, despite their recognized biological importance.[1,11]

The study of mutational patterns has revealed that some driver gene mutations tend to co-occur, suggesting they cooperate in tumor development, while others are mutually exclusive.[12–14] This mutual exclusivity suggests a functional redundancy, where a mutation in either gene is sufficient to dysregulate a common pathway, making a second mutation in the other gene redundant for cancer progression. This supports the notion that functionally similar genes can leverage convergent cellular processes to drive cancer. Understanding and quantifying this functional similarity is central to our work.

In this paper, we introduce INSPIRE , a novel comparative network analysis pipeline to investigate cancer driver genes and elucidate their functional interrelationships (Figure 1). We hypothesize that a pathway-centric approach, which quantifies the similarity between the molecular networks dysregulated by driver genes, will provide a more comprehensive and biologically accurate depiction of their relationships than traditional gene-centric comparisons based solely on sequence or ontology.

A noteworthy contribution of our methodology is our unique approach to quantifying similarity. To measure subnetwork similarity, we introduce a novel network-alignment algorithm

specifically designed to align functionally related PPI subnetworks from the same organism, termed *within-species network alignment*. Crucially, our algorithm seeks to align conserved patterns of interactions rather than just individual proteins, aiming to capture true functional conservation between pathways.

While there is a rich literature on comparing networks across species (cross-species alignment),[15–23] these methods are not directly suited to comparing functionally distinct subnetworks within the same organism. We identify the following as key requirements for our task:

(1) The alignment must be global, yielding a unified similarity score between the two networks.
(2) The mapping must be unambiguous; each protein from one network should map to at most one protein in the other.
(3) Identical proteins shared between two networks (a feature unique to within-species comparison) must be preferentially aligned to each other, reflecting their identical function.

Existing network alignment solutions often fail to meet these requirements. For instance, IsoRank[24] and its derivatives are global alignment tools, but they do not satisfy requirement (3). Because IsoRank's scoring is based on the PageRank algorithm, a protein's alignment score depends on its neighbors' scores. Consequently, two identical proteins may not receive the maximum score if their network neighborhoods differ. The GraphAlignment R package only handles cross-species network alignment and cannot align genes to functional equivalents while also mapping identical genes to each other between networks.[25] Other prominent algorithms focus on local alignments and thus do not satisfy requirement (1).[15,19] While network propagation methods (e.g., HotNet2[26]) identify genes close to each other in network space, and graph embeddings (e.g., node2vec[27]) capture global structural roles, neither explicitly solves the subgraph isomorphism problem to find topologically identical wiring patterns independent of sequence homology. INSPIRE fills this gap by providing a bespoke alignment strategy that quantifies both protein functional similarity and the topological consistency of their interactions.

The primary objective of INSPIRE is to provide a quantitative, pathway-level framework to measure the functional convergence of altered signaling networks. This approach leverages established molecular interactions and functional annotations to determine how distinct driver mutations converge on common mechanisms. By projecting disparate mutational events onto comparative alignment graphs, INSPIRE identifies systems-level functional redundancies that may be missed by gene-centric analyses.

2. Methods

2.1. *Cancer Driver Gene Networks*

To model signaling pathways of colorectal cancer (CRC) driver genes, we constructed protein-protein interaction (PPI) subnetworks connecting each driver to *APC*, a common initiating gene in CRC.[1] Starting from 140 candidate driver genes, putative signaling networks were generated using an adapted PETALS framework[6] on a CRC-specific PPI network derived from IntAct, STRING, and MSS TCGA-COAD expression data.[28–30]

Pathway significance was assessed using GO-derived association rules and a permutation-based empirical null model. Networks were retained at FDR ≤ 0.05, resulting in significant

subnetworks for 28 driver genes. Additional implementation details are provided in Supplementary Methods (Section S1.7).

2.2. *Quantifying Network Similarity*

For two subnetworks to be considered functionally similar, both the proteins and the topological context in which they interact should be similar.[19,31] To incorporate both factors, we formulate network similarity as a network alignment problem.

We represent two subnetworks as graphs $G_A = (V_A, E_A)$ and $G_B = (V_B, E_B)$, where vertices ($v \in V$) correspond to proteins and edges ($e \in E$) to interactions. We define an alignment $\mathbb{A}$ as a one-to-one mapping of a subset of edges from one network to the other. Formally, $\mathbb{A}$ is a function that maps a subset of edges $E_0 \subseteq E_A$ to E_B. The choice of E_0 is part of the optimization process. This edge-based mapping must induce a consistent, one-to-one mapping of the corresponding vertices. That is, if an edge $(u, v) \in E_0$ is mapped to $(x, y) \in E_B$, then no other edge involving u or v can be mapped to an edge involving x or y.

The quality of an alignment is given by an alignment score $W(\mathbb{A}) = \sum_{e \in E_0} w(e, \mathbb{A}(e))$, where $w(e, \mathbb{A}(e))$ is the score for aligning interaction e with its counterpart $\mathbb{A}(e)$. The goal is to find the alignment $\mathbb{A}$ that maximizes this total score.

2.3. *Constructing the Alignment Graph*

To efficiently find the best alignment between G_A and G_B, we adapt the alignment graph concept.[19] The alignment graph $\Psi(G_A, G_B) = (U_\Psi, F_\Psi)$ is constructed as follows:

- Vertices: Each vertex $u \in U_\Psi$ represents a potential alignment of one interaction from each network, i.e., $u = (e_A, e_B)$ where $e_A \in E_A$ and $e_B \in E_B$. Thus, $U_\Psi = E_A \times E_B$.
- Edges: Two vertices in the alignment graph, $u_i = (e_{A_i}, e_{B_i})$ and $u_j = (e_{A_j}, e_{B_j})$, are connected by an edge $f_{ij} \in F_\Psi$ if their corresponding interactions are topologically consistent. We consider three types of connectivities:

(1) Extension: The two pairs of aligned edges share a protein in both graphs (Edge in Graph A matches Edge in Graph B).
(2) Insertion/Deletion (Indel): The pairs share a protein in one graph but are separated by one hop in the other (Edge in Graph A matches a 2-hop path in Graph B).
(3) Jump: The pairs are separated by one hop in both networks (2-hop path in Graph A matches 2-hop path in Graph B).

The edges f_{ij} are weighted based on the scores of the nodes they connect (u_i, u_j) and a penalty for topological inconsistencies (indels and jumps). We then seek a path through this alignment graph that corresponds to a valid, high-scoring alignment.

2.4. *Alignment Scoring*

The scoring scheme is designed to reward both sequence/functional similarity of proteins and topological consistency of interactions.

Scoring Nodes of the Alignment Graph: A node $u = (e_A, e_B)$ in the alignment graph represents an aligned interaction pair. Its weight, $w(e_A, e_B)$, reflects the biological similarity of this pairing. This score is based on the sequence and functional similarity of the corresponding endpoint proteins, calculated using Pfam domains[32] and Gene Ontology (GO) annotations.[33]

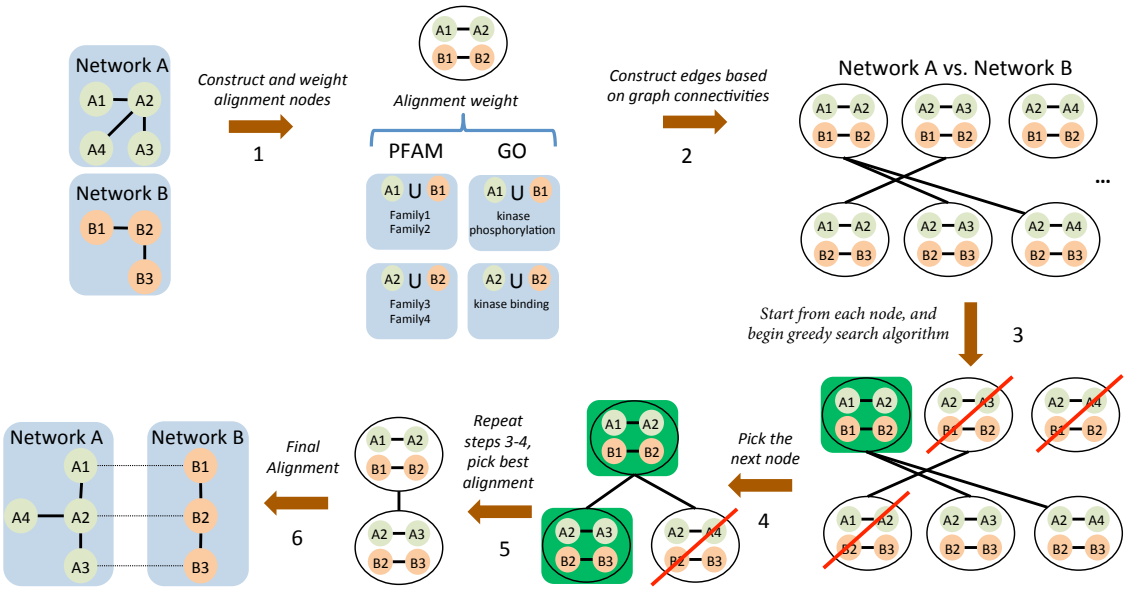


Fig. 2. The INSPIRE alignment algorithm. (1) An alignment graph is constructed in which each node represents a potential correspondence between interactions from networks A and B and is weighted by protein similarity. (2) Nodes are connected according to topological consistency (extensions, indels, and jumps). (3–5) Starting from each node, a greedy MST-inspired search iteratively adds the most favorable consistent correspondence (equivalently, the minimum-cost edge after similarity-to-cost transformation) while removing conflicting alternatives. (6) The procedure is repeated for all starting nodes, and the highest-scoring alignment is returned.

By aligning interactions (edges) rather than just proteins (nodes), we can better capture conserved functional associations.[19]

Scoring Sequence Similarity: We use the Jaccard index to compare the sets of Pfam domains for the proteins in the aligned interactions. For two interactions $e_A = (v_{A_i}, v_{A_j})$ and $e_B = (v_{B_k}, v_{B_l})$, the sequence similarity is:

$$s_{pfam}(e_A, e_B) = \frac{|\mathsf{PF}(v_{A_i}) \cap \mathsf{PF}(v_{B_k})|}{|\mathsf{PF}(v_{A_i}) \cup \mathsf{PF}(v_{B_k})|} \cdot \frac{|\mathsf{PF}(v_{A_j}) \cap \mathsf{PF}(v_{B_l})|}{|\mathsf{PF}(v_{A_j}) \cup \mathsf{PF}(v_{B_l})|} \tag{1}$$

Where $\mathsf{PF}(v)$ is the set of Pfam domains for protein v and $PF(v) = \emptyset$, the score is defined as 0.

Scoring Functional Similarity: We use a semantic similarity measure based on GO Biological Process (BP) annotations to assess functional relationships. The information content (IC) of the shared annotations for two proteins is given by:[34]

$$IC(S_i, S_j) = -\log_2 \frac{|G_{\Lambda(S_i,S_j)}|}{|G_r|} \tag{2}$$

Here, $G_{\Lambda(S_i,S_j)}$ is the set of all genes annotated with the most specific common ancestor of the GO terms for genes S_i and S_j, and G_r is the total set of genes in GO. We use this to compute the functional similarity of two interactions:

$$s_{GO}(e_A, e_B) = IC(v_{A_i}, v_{B_k}) \cdot IC(v_{A_j}, v_{B_l}) \tag{3}$$

This structure mirrors the sequence similarity score, ensuring that corresponding proteins in the two interactions are compared.

Overall Similarity Between Two Interactions: The total weight of aligning two interactions is the sum of these scores:

$$w(e_A, e_B) = s_{pfam}(e_A, e_B) + s_{GO}(e_A, e_B) \tag{4}$$

While semantic similarity packages like GOSemSim offer detailed hierarchy-aware scoring, we employed a Jaccard-based metric here to prioritize computational efficiency for large-scale graph comparisons. We found this provided a sufficient baseline for topological alignment.

2.5. *Scoring the Alignment Graph and Normalization of Scores*

Raw alignment scores were normalized by the maximum attainable score for the smaller network, producing similarities on the interval [0,1]. Pairwise distances were defined as $\hat{s}_{AB}$ and used for downstream clustering (See Section S1.1 for more details).

2.6. *Finding the Optimal Alignment*

Finding the globally optimal alignment is NP-hard.[35] We therefore employ a greedy heuristic search algorithm to find a high-scoring alignment. The challenge is to find a path or tree in the alignment graph Ψ that corresponds to a valid, one-to-one mapping. A Minimum Spanning Tree (MST) algorithm cannot be used because it does not enforce a one-to-one constraint.

Our algorithm is inspired by greedy methods like Prim's algorithm but includes a critical modification: when a node corresponding to the alignment (e_A, e_B) is added to the solution tree, all other nodes in the alignment graph that involve either edge e_A or edge e_B are immediately pruned (removed from consideration). This pruning step enforces the one-to-one mapping constraint and dramatically reduces the search space. The algorithm is initiated from every possible node in the alignment graph. The algorithm then scores the resulting alignments and returns the highest-scoring valid alignment as the final result. To enable MST-style greedy search, the algorithm transforms alignment similarity scores into edge costs using an inverse relationship. Consequently, the algorithm expands the alignment by selecting the minimum-cost candidate at each step.

The algorithm proceeds as follows:

`Initialize:`

$E_{traversed} = \emptyset$

$V_{traversed} = \{x\}$, where $x = (e_{A_x}, e_{B_x})$ is a node from $\mathbb{A}$.

$V_{available} = V_\Psi \backslash [(e_{A_x}, *), (*, e_{A_x}), (e_{B_x}, *), (*, e_{B_x})]$

`Repeat until` $V_{available} = \emptyset$

Select (x, y), where $y = (e_{A_y}, e_{B_y}) \in V_{available}$ with minimum weight in Ψ, $x \in V_{traversed}$, and $y \notin V_{traversed}$

$V_{traversed} = V_{traversed} \cup y$

$E_{traversed} = E_{traversed} \cup (x, y)$

$V_{available} = V_{available} \backslash [(e_{A_y}, *), (*, e_{A_y}), (e_{B_y}, *), (*, e_{B_y})]$

The algorithm is repeated from every node in the alignment graph. The traversal with

the minimum weight corresponds to the best alignment of the graphs. The time complexity of this algorithm depends on the choice of an edge at each step, which requires $O(|V|^2)$ time. However, this can be improved further by using heaps to find the minimum-weight edges in the algorithm's inner loop.

2.7. *Sequence and Gene Ontology-Based Classification of Genes*

As a control (gene-centric) method, we computed a pairwise similarity score w' directly between driver genes, omitting network information. This score combines sequence and GO similarity analogously to our interaction scoring:

$$w'(v_i, v_j) = s'_{pfam}(v_i, v_j) + s'_{GO}(v_i, v_j) = \frac{|\mathsf{PF}(v_i) \cap \mathsf{PF}(v_j)|}{|\mathsf{PF}(v_i) \cup \mathsf{PF}(v_j)|} + IC(v_i, v_j) \tag{5}$$

2.8. *Methodological Benchmarking against MAGNA++ and IsoRank*

To compare with established network alignment approaches, we also performed pairwise alignments using MAGNA++[36] and IsoRank[24] with the same subnetworks as input. Default parameters were used unless otherwise specified. We analyzed network similarity scores from each method using the same clustering framework described below.

2.9. *Clustering*

Stable clustering of the 28 CRC driver genes was performed using ConsensusClusterPlus[37] on the pairwise INSPIRE network similarity matrix. The similarity matrix was symmetrized and converted to a distance matrix as $d_{ij} = 1 - s_{ij}$. Consensus clustering used Partitioning Around Medoids (PAM) with 1,000 bootstrap resamplings of 80% of genes per iteration, testing $k = 2, \ldots, 8$.

2.10. *Co-essentiality Analysis*

CRISPR gene effect (Chronos) scores were obtained from the Cancer Dependency Map.[38] A selective essentiality filter was applied for each gene pair (A, B); only cell lines where at least one gene showed a Chronos score below $\tau = -0.3$ were included in the pairwise Spearman correlation. Common essential ($> 50\%$ of lines) and low-signal genes (< 4 essential lines) were excluded. The primary co-essentiality analysis used the resulting n=24 driver genes. Within-cluster vs. between-cluster co-essentiality was compared using a one-sided Wilcoxon rank-sum test, with effect size reported as the rank-biserial correlation $r_{rb} = W/(n_1 \times n_2)$.

2.11. *Somatic Mutation Profiles and Statistical Evaluation*

Mutational profiles for the 28 driver genes were obtained from TCGA MSS-COAD tumors.[30] We assessed associations between functional similarity and mutational co-occurrence using Mantel tests comparing network-based and mutation-based distances (10,000 permutations).

To overcome the statistical power limitations inherent to the low tumor mutational burden (TMB) in MSS COAD, we implemented a module-level mutual exclusivity framework. Using the Ward.D2 functional dendrogram, genes were partitioned into $K = 6$ clades via tree cutting. For each patient, a clade was designated as mutated if at least one constituent gene contained a somatic mutation. Pairwise mutual exclusivity between clades was evaluated using a one-sided Fisher's Exact Test (alternative = "less"), testing whether co-mutations occurred significantly less frequently than expected by chance. P-values were adjusted for multiple testing using the Benjamini-Hochberg False Discovery Rate (FDR) procedure.

2.12. *Epistasis Proxy Analysis*

To assess functional epistasis between within-cluster gene pairs, we stratified CRISPR dependency scores for a target gene (Gene B) by mutation status of a driver gene (Gene A) across MSS adenocarcinoma cell lines (n=251). Mutation status was derived from the DepMap damaging mutation matrix.[39] Differential dependency was tested using the Wilcoxon rank-sum test. Effect size was quantified using Cliff's delta (δ), a non-parametric dominance statistic ranging from -1 to $+1$, classified using thresholds: negligible ($|\delta| < 0.147$), small ($0.147 - 0.33$), medium (0.33–0.474), and large (≥ 0.474).[40] Given the small mutant group sizes (median $n_{mut} = 5$), results are reported as exploratory. Cliff's delta is the primary metric, and raw p-values are reported without multiple-testing correction.

3. Results

To investigate functional convergence among cancer driver genes, we developed `INSPIRE`, a network-alignment framework that quantifies similarity between protein-protein interaction (PPI) networks. We applied this framework to study driver genes in colorectal cancer (CRC), where a long tail of infrequently mutated genes suggests that distinct genetic alterations may converge on common pathways.

3.1. *Driver-Gene Networks*

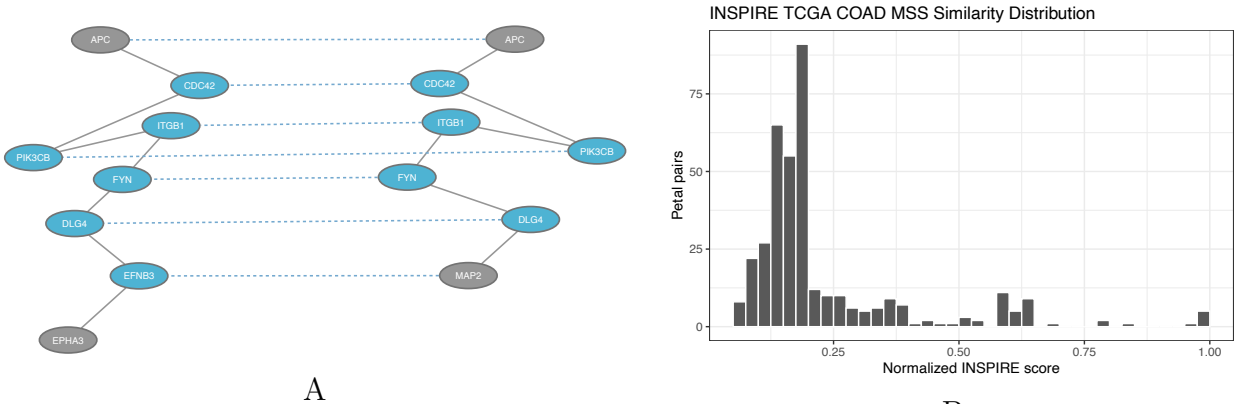
Our analysis began with 140 candidate CRC driver genes. For each gene, we constructed a signaling subnetwork connecting it to *APC*, a key initiator of colorectal tumorigenesis.[41,42] We derived these networks from a CRC-specific PPI graph built from curated databases and TCGA gene-expression data (see Methods 2.1). We identified 28 genes with high-confidence subnetworks (FDR ≤ 0.05), (Supp. Table S1). The resulting subnetworks ranged in size from 4 to 29 proteins.

3.2. *Network similarity reveals functional relationships among driver genes*

We applied INSPIRE to all pairwise comparisons among the 28 CRC driver-gene networks. Each alignment produced a normalized network similarity score reflecting similarity of signaling architecture rather than sequence homology.

The distribution of pairwise similarity scores revealed substantial heterogeneity among the driver networks, with most pairs exhibiting low similarity and a smaller subset displaying strong similarity. Using the resulting distance matrix, hierarchical clustering identified biologically coherent groups of driver genes. Members of known functional modules, including the TGF-β pathway genes *SMAD2*, *SMAD3*, and *SMAD4*, clustered closely together. Similarly, genes with established pathway-level relationships but limited sequence similarity frequently appeared within the same network-derived clusters.

To obtain a stable cluster assignment, we applied consensus clustering using Partitioning Around Medoids (PAM) with 1,000 bootstrap resamplings. Evaluation of solutions ranging from $k = 2$ to $k = 8$ identified five robust clusters (Supp. Figures S6 and S7), selected based on consensus heatmaps and biological interpretability. The final clustering exhibited positive separation in the original INSPIRE similarity space, with an average silhouette width of 0.132, indicating non-random structure among the inferred pathway relationships.



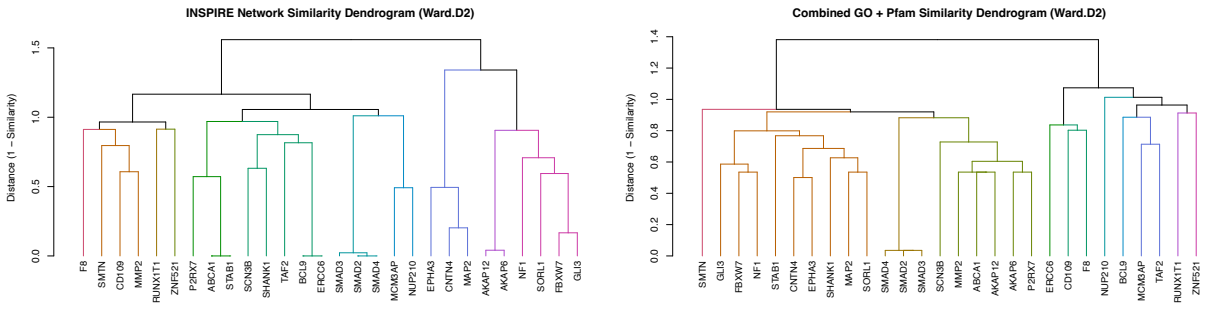
A B

Fig. 3. A. Example alignment of two CAN-gene networks (*EPHA3* (left) vs. *MAP2* (right)). CAN-genes are colored grey, and dashed (blue) edges show aligned nodes (hence edges) that returned the best score. This alignment has a normalized score of 0.69735. The alignment includes EFNB3 and MAP2. Despite distinct sequences, their network topology indicates that they occupy equivalent positions in the signaling cascade, illustrating functional convergence. B. Histogram of normalized alignment scores ($\hat{s}_{AB}$). All 28 networks are compared pairwise. While most pairs show low similarity, a subset ($\sim 10\%$) achieves high similarity scores ($\hat{s}_{AB} \geq 0.5$), suggesting functional convergence.

3.3. *Network-based clustering captures functional relationships beyond gene-level similarity*

To assess whether network alignment provides information beyond gene-level annotations, we compared INSPIRE clustering against a control approach based on Gene Ontology semantic similarity and Pfam domain similarity (Figure 4). Whereas the control approach evaluates genes individually using shared annotations and domain content, INSPIRE additionally incorporates the signaling-network context surrounding each driver gene.

Both approaches recovered the canonical TGF-β module (SMAD2, SMAD3, SMAD4). However, INSPIRE additionally identified several relationships not apparent from gene-level annotations. Most notably, MCM3AP and NUP210 formed a distinct two-gene cluster, and AKAP12, AKAP6, NF1, SORL1, FBXW7, and GLI3 were grouped within a common higher-



A B

Fig. 4. A. Hierarchical clustering of CAN-gene networks using distance scores from our INSPIRE algorithm. The clustering reveals functional relationships based on network similarity. B. Hierarchical clustering of the same CAN-genes using a gene-centric distance measure that combines GO semantic similarity[43] and sequence similarity based on Pfam domains.[32]

order clade. These findings suggest that signaling-network context captures functional relationships beyond sequence and annotation similarity.

Importantly, the biological relevance of the INSPIRE-derived clusters was supported by independent validation analyses. Genes assigned to the same network-defined clusters exhibited significantly greater co-essentiality than those in different clusters, and several candidate epistatic interactions were identified within cluster boundaries. Together, these results indicate that network topology captures biologically meaningful relationships that complement information obtained from sequence- and annotation-based similarity measures.

3.4. *Functional validation using co-essentiality analysis*

To independently evaluate whether INSPIRE clusters represent biologically coherent functional modules, we analyzed CRISPR dependency data from the Cancer Dependency Map. Co-essentiality was quantified using Chronos gene effect scores. Because CRC driver genes are enriched for tumor suppressors that are nonessential in most cell lines, we applied a selective-essentiality filtering strategy prior to analysis. Two genes (*MCM3AP* and *TAF2*) were classified as common essentials and removed, while *F8* and *SHANK1* were excluded due to insufficient signal. The final analysis therefore included 24 genes. Gene pairs assigned to the same INSPIRE cluster exhibited significantly stronger co-essentiality than pairs belonging to different clusters (one-sided Wilcoxon rank-sum test, $p = 0.038$, rank-biserial correlation $r_{rb} = 0.567$). This result indicates that genes grouped by network alignment display more similar dependency profiles across cancer cell lines, consistent with convergence on common biological processes. The same directional trend remained detectable when all 28 genes were included ($p = 0.055$), whereas the standard unfiltered analysis showed no signal ($p = 0.91$). Although cluster separation was evident in INSPIRE similarity space (average silhouette width $= 0.132$), separation in co-essentiality space was nearly absent (average silhouette width $= -0.025$).

3.5. *Epistasis proxy analysis identifies candidate functional interactions*

We next investigated whether genes grouped within the same INSPIRE cluster also exhibited evidence of functional interaction. Using 251 microsatellite-stable adenocarcinoma cell lines, we tested whether damaging mutations in one cluster member altered CRISPR dependency on another member.

Among 134 within-cluster directional gene pairs, nine showed nominal significance ($p < 0.05$), although none remained significant after multiple-testing correction. Several interactions nevertheless exhibited large effect sizes, including *SMAD2* $\rightarrow$ *SMAD4* ($\delta = -0.69$, $p = 0.004$), *EPHA3* $\rightarrow$ *CNTN4* ($\delta = -0.72$, $p = 0.014$), and *SMTN* $\rightarrow$ *MMP2* ($\delta = 0.76$, $p = 0.025$) . All three large-effect interactions occurred within INSPIRE-defined clusters (Figure 5).

3.6. *Functional Network Architecture is Independent of Global Co-occurrence but Exhibits Module-Level Mutual Exclusivity*

To test whether functional network similarity correlates with patient-level co-mutation, we performed a global Mantel test between the functional distance matrix and the mutational distance matrix (Figure S4). The test revealed no significant global correlation ($p = 0.50$). This result demonstrates that global network distance does not linearly track pairwise mutation

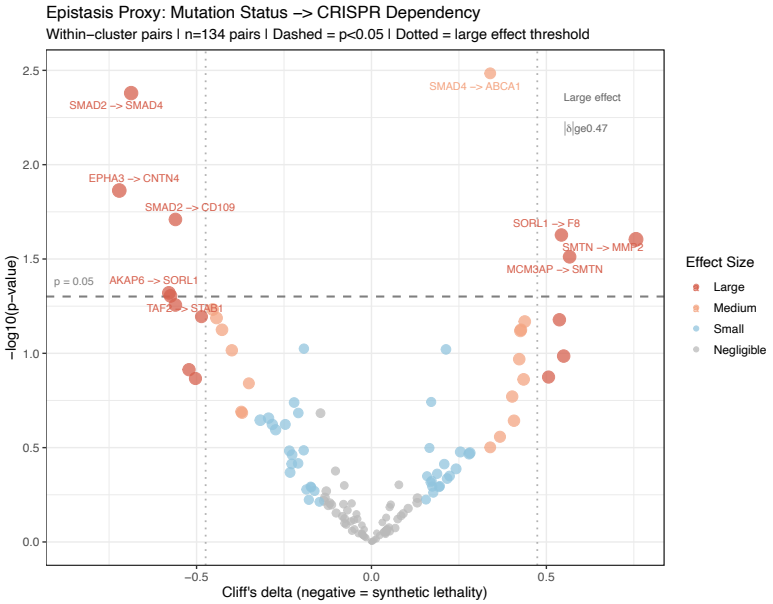


Fig. 5. Epistasis proxy analysis of within-cluster INSPIRE driver gene pairs. Volcano plot of Cliff's delta (effect size) versus -log10(p-value) for 134 directional within-cluster gene pairs tested across MSS adenocarcinoma cell lines. Points are colored by effect size magnitude.[40] A negative delta indicates synthetic lethality; a positive delta indicates epistatic buffering. Dashed line: p = 0.05. Results are exploratory given small mutant group sizes (median n=5 mutant lines per gene).

co-occurrence across the entire panel.

Pairwise single-gene mutual exclusivity tests between adjacent dendrogram leaves were statistically underpowered due to the characteristic low mutation frequencies of MSS COAD driver genes (typically 3%–20% prevalence). For instance, testing the adjacent TGF-β pathway components *SMAD2* (27 mutations) and *SMAD4* (67 mutations) yielded 4 observed co-mutated patients against an expected baseline of 5.71 patients (Odds Ratio = 0.627, $p = 0.286$, FDR = 1.00). While functionally convergent, single-gene sample sizes in MSS COAD lack the statistical power required to differentiate minor co-occurrence drops from stochastic background noise. To resolve this limitation, we aggregated individual low-frequency genes into $K = 6$ higher-order functional clades based on dendrogram topology (Supp. Table S3). Aggregating genes into functional modules elevated module mutation rates (e.g., 20%–50% prevalence), restoring statistical power. Pairwise Fisher's exact tests across clades identified significant pathway-level mutual exclusivity. Notably, Clade 2: *AKAP12, AKAP6, FBXW7, GLI3, NF1, SORL1* and Clade 6: *SMAD2, SMAD3, SMAD4* exhibited strong mutual exclusivity (Odds Ratio = 0.460, $p = 9.96 \times 10^{-4}$, FDR = 0.0149) (Table S2).

3.7. *Comparison with Existing Network Alignment Methods*

To evaluate whether standard alignment frameworks could resolve functional convergence among CRC driver networks, we benchmarked our approach against two widely used global network alignment tools: MAGNA++[36] and IsoRank[24] (Supplementary Figure S8). Both methods proved ill-suited for intra-species, pathway-level comparison. Hierarchical clustering of pairwise alignments from MAGNA++ and IsoRank produced highly degenerate, unstructured dendrograms characterized by chaining artifacts and low modular resolution. Notably, neither method recovered canonical pathway relationships—such as grouping members of the TGF-β/SMAD signaling cascade (*SMAD2*, *SMAD3*, *SMAD4*)—and frequently failed to align identical proteins shared across subnetworks. These shortcomings stem from their foundational design for cross-species orthology mapping, which assumes sequence divergence and prioritizes conserved topological isomorphism across species. In contrast, by formulating align-

ment around edge-cost transformations that explicitly enforce identity mapping and reward functional conservation, `INSPIRE` resolves distinct, biologically cohesive clusters that reflect genuine pathway-level convergence.

4. Discussion

A central premise of this study is that cancer driver genes can be understood through the signaling networks they perturb rather than solely through their individual molecular characteristics. The results support this hypothesis. INSPIRE consistently identified functionally coherent groups of CRC driver genes using network topology alone and produced clusters that were more biologically meaningful than those obtained from sequence similarity or ontology-based annotations.

The strongest independent validation came from the co-essentiality analysis. Genes assigned to the same INSPIRE clusters exhibited significantly greater similarity in CRISPR dependency profiles than genes belonging to different clusters. Importantly, this pattern emerged only after applying a selective-essentiality filter tailored to tumor suppressor genes. Conventional co-essentiality analyses did not recover a comparable structure. These findings suggest that network-derived relationships capture functional convergence reflected in cellular dependency phenotypes.

The exploratory epistasis analysis further supported the biological relevance of the network-derived clusters. Although statistical power was limited by the rarity of many driver mutations, nine within-cluster gene pairs showed nominally significant interactions, several with substantial effect sizes. While these observations require further validation, the enrichment of candidate interactions within INSPIRE-defined clusters suggests that network similarity may help prioritize functionally relevant genetic relationships.

Functional similarity was not associated with global mutational co-occurrence (Mantel test, $p = 0.50$), likely reflecting the low prevalence of MSS colorectal cancer driver mutations and the limited power of pairwise co-mutation analyses. Consistent with this interpretation, mutual exclusivity emerged only after aggregating genes into INSPIRE-derived functional clades. Whereas single-gene analyses failed to identify significant associations, module-level testing revealed robust mutually exclusive interactions, suggesting that functional redundancy operates at the level of network-defined pathways. These results demonstrate that network-based aggregation can expose mutational patterns that gene-centric analyses miss.

Several limitations should be acknowledged. First, the epistasis analysis was underpowered because most driver mutations were present in only a few cell lines. Consequently, these results should be viewed as hypothesis-generating rather than definitive evidence of genetic interaction. Second, the co-essentiality validation involved a relatively small set of driver genes following filtering, which may have reduced statistical sensitivity. Finally, the inferred signaling networks depend on available protein interaction resources and expression measurements, although restricting network construction to curated interactions increases biological confidence.

Despite these limitations, the network, co-essentiality, epistasis, and mutational analyses converge on a common conclusion that the colorectal cancer driver genes frequently partici-

pate in shared functional programs despite genetic heterogeneity and mutual exclusivity. By quantifying these relationships at the level of signaling architecture, INSPIRE provides a robust framework for uncovering pathway-level organization that conventional gene-centric approaches overlook.

While this initial validation leveraged well-characterized pathways to establish baseline accuracy, this comparative engine is naturally positioned for exploratory annotation. Future applications of INSPIRE can systematically align established reference networks against those of poorly characterized genes, providing a computational pipeline to generate targeted, testable hypotheses to guide downstream experimental design.

Acknowledgement

Research reported in this publication was supported by the National Center for Research Resources (NCRR) of the National Institutes of Health under Award Number KL2RR024990. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

1. L. D. Wood, D. W. Parsons, S. Jones, J. Lin, T. Sjoblom, R. J. Leary, D. Shen, S. M. Boca, T. Barber, J. Ptak, N. Silliman, S. Szabo, Z. Dezso, V. Ustyanksky, T. Nikolskaya, Y. Nikolsky, R. Karchin, P. A. Wilson, J. S. Kaminker, Z. Zhang, R. Croshaw, J. Willis, D. Dawson, M. Shipitsin, J. K. V. Willson, S. Sukumar, K. Polyak, B. H. Park, C. L. Pethiyagoda, P. V. K. Pant, D. G. Ballinger, A. B. Sparks, J. Hartigan, D. R. Smith, E. Suh, N. Papadopoulos, P. Buckhaults, S. D. Markowitz, G. Parmigiani, K. W. Kinzler, V. E. Velculescu and B. Vogelstein, The genomic landscapes of human breast and colorectal cancers, *Science* **318**, 1108 (nov 2007).
2. D. W. Parsons, S. Jones, X. Zhang, J. C. Lin, R. J. Leary, P. Angenendt, P. Mankoo, H. Carter, I. M. Siu, G. L. Gallia, A. Olivi, R. McLendon, B. A. Rasheed, S. Keir, T. Nikolskaya, Y. Nikolsky, D. A. Busam, H. Tekleab, J. Diaz, L. A., J. Hartigan, D. R. Smith, R. L. Strausberg, S. K. Marie, S. M. Shinjo, H. Yan, G. J. Riggins, D. D. Bigner, R. Karchin, N. Papadopoulos, G. Parmigiani, B. Vogelstein, V. E. Velculescu and K. W. Kinzler, An integrated genomic analysis of human glioblastoma multiforme, *Science* **321**, 1807 (2008).
3. M. H. Bailey, C. Tokheim, E. Porta-Pardo, S. Sengupta, D. Bertrand *et al.*, Comprehensive characterization of cancer driver genes and mutations, *Cell* **173**, 371 (apr 2018).
4. Z. Sondka, S. Bamford, C. G. Cole, S. A. Ward, I. Dunham and S. A. Forbes, The COSMIC cancer gene census: describing genetic dysfunction across all human cancers, *Nature Reviews Cancer* **18**, 696 (oct 2018).
5. A. Torkamani and N. J. Schork, Identification of rare cancer driver mutations by network reconstruction, *Genome Res* **19**, 1570 (2009).
6. G. Bebek, V. Patel and M. R. Chance, Petals: Proteomic evaluation and topological analysis of a mutated locus' signaling, *BMC Bioinformatics* **11**, p. 596 (2010).
7. N. Yosef, E. Zalckvar, A. D. Rubinstein, M. Homilius, N. Atias, L. Vardi, I. Berman, H. Zur, A. Kimchi, E. Ruppin and R. Sharan, Anat: a tool for constructing and analyzing functional protein networks, *Sci Signal* **4**, p. pl1 (2011).
8. H. Sun, B. Han, X. Cheng and K. Ma, Transcriptional regulatory network and protein–protein interaction to reveal the mechanism of pancreatic cancer, *Molecular Biology Reports* **41**, 387 (dec 2013).
9. G. Wu, X. Feng and L. Stein, A human functional protein interaction network and its application to cancer data analysis, *Genome Biology* **11**, p. R53 (2010).

10. Y. Xue, G. Cooper, C. Cai, S. Lu, B. Hu, X. Ma and X. Lu, Tumour-specific causal inference discovers distinct disease mechanisms underlying cancer subtypes, *Scientific Reports* **9** (sep 2019).
11. S. D. Markowitz and M. M. Bertagnolli, Molecular origins of cancer: Molecular basis of colorectal cancer, *N Engl J Med* **361**, 2449 (Dec 2009).
12. S. Park and B. Lehner, Cancer type-dependent genetic interactions between cancer driver alterations indicate plasticity of epistasis across cell types, *Mol Syst Biol* **11**, p. 824 (Jul 2015).
13. P. Priestley, J. Baber, M. P. Lolkema, N. Steeghs, E. de Bruijn, C. Shale, K. Duyvesteyn, S. Haidari, A. van Hoeck, W. Onstenk, P. Roepman, M. Voda, H. J. Bloemendal, V. C. G. Tjan-Heijnen, C. M. L. van Herpen, M. Labots, P. O. Witteveen, E. F. Smit, S. Sleijfer, E. E. Voest and E. Cuppen, Pan-cancer whole-genome analyses of metastatic solid tumours, *Nature* **575**, 210 (oct 2019).
14. F. Sanchez-Vega, M. Mina, J. Armenia, W. K. Chatila, A. Luna *et al.*, Oncogenic signaling pathways in the cancer genome atlas, *Cell* **173**, 321 (apr 2018).
15. M. Koyuturk, A. Grama and W. Szpankowski, An efficient algorithm for detecting frequent subgraphs in biological networks, *Bioinformatics* **20 Suppl 1**, i200 (2004).
16. M. Koyuturk, Y. Kim, S. Subramaniam, W. Szpankowski and A. Grama, Detecting conserved interaction patterns in biological networks, *J Comput Biol* **13**, 1299 (2006), Koyuturk, Mehmet Kim, Yohan Subramaniam, Shankar Szpankowski, Wojciech Grama, Ananth R01 GM068959-01/GM/NIGMS NIH HHS/ J Comput Biol. 2006 Sep;13(7):1299-322.
17. J. Flannick, A. Novak, B. S. Srinivasan, H. H. McAdams and S. Batzoglou, Graemlin: general and robust alignment of multiple large interaction networks, *Genome Res* **16**, 1169 (2006).
18. B. Dost, T. Shlomi, N. Gupta, E. Ruppin, V. Bafna and R. Sharan, Qnet: a tool for querying protein interaction networks, *J Comput Biol* **15**, 913 (2008).
19. X. Guo and A. J. Hartemink, Domain-oriented edge-based alignment of protein interaction networks, *Bioinformatics* **25**, i240 (2009).
20. C. S. Liao, K. Lu, M. Baym, R. Singh and B. Berger, Isorankn: spectral methods for global alignment of multiple protein networks, *Bioinformatics* **25**, i253 (2009).
21. U. Pieper, N. Eswar, B. M. Webb, D. Eramian, L. Kelly, D. T. Barkan, H. Carter, P. Mankoo, R. Karchin, M. A. Marti-Renom, F. P. Davis and A. Sali, Modbase, a database of annotated comparative protein structure models and associated resources, *Nucleic Acids Res* **37**, D347 (2009).
22. X. Qian and B. J. Yoon, Effective identification of conserved pathways in biological networks using hidden markov models, *PLoS One* **4**, p. e8070 (2009).
23. F. Ay, M. Kellis and T. Kahveci, Submap: aligning metabolic pathways with subnetwork mappings, *J Comput Biol* **18**, 219 (2011).
24. R. Singh, J. Xu and B. Berger, Global alignment of multiple protein interaction networks with application to functional orthology detection, *Proc Natl Acad Sci U S A* **105**, 12763 (Sep 2008).
25. J. Berg and M. Laessig, Cross-species analysis of biological networks by bayesian alignment, *Proceedings of the National Academy of Sciences* **103**, 10967 (2006).
26. M. D. M. Leiserson, F. Vandin, H. T. Wu, J. R. Dobson, J. V. Eldridge, J. L. Thomas, A. Papoutsaki, Y. Kim, B. Niu, M. McLellan, M. S. Lawrence, A. G. Perez, D. Tamborero, Y. Cheng, G. A. Ryslik, N. Lopez-Bigas, G. Getz, L. Ding and B. J. Raphael, Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes, *Nature Genetics* **47**, 106 (feb 2015).
27. A. Grover and J. Leskovec, node2vec: Scalable feature learning for networks, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (Association for Computing Machinery, New York, NY, USA, 2016).
28. D. Szklarczyk, R. Kirsch, M. Koutrouli, K. Nastou, F. Mehryary, R. Hachilif, A. L. Gable, T. Fang, N. T. Doncheva, S. Pyysalo, P. Bork, L. J. Jensen and C. von Mering, The string

database in 2023: protein–protein association networks and functional enrichment analyses for any sequenced genome of interest, *Nucleic Acids Research* **51**, D638 (2023).
29. S. Orchard, M. Ammari, B. Aranda, L. Breuza, L. Briganti, F. Broackes-Carter, N. H. Campbell, G. Chavali, C. Chen, N. Del-Toro, M. Duesbury, M. Dumousseau, E. Galeota, U. Hinz, M. Iannuccelli, S. Jagannathan, R. Jimenez, J. Khadake, A. Lagreid, L. Licata, R. C. Lovering, B. Meldal, A. N. Melidoni, M. Milagros, D. Peluso, L. Perfetto, P. Porras, A. Raghunath, S. Ricard-Blum, B. Roechert, A. Stutz, M. Tognolli, K. van Roey, G. Cesareni and H. Hermjakob, The MIntAct project—IntAct as a common curation platform for 11 molecular interaction databases, *Nucleic Acids Research* **47**, D529 (2019).
30. The Cancer Genome Atlas Network, Comprehensive molecular characterization of human colon and rectal cancer, *Nature* **487**, 330 (07 2012).
31. G. Bebek and J. Yang, Pathfinder: mining signal transduction pathway segments from protein-protein interaction networks, *BMC Bioinformatics* **8**, 335 (2007).
32. M. Punta, P. C. Coggill, R. Y. Eberhardt, J. Mistry, J. Tate, C. Boursnell, N. Pang, K. Forslund, G. Ceric, J. Clements, A. Heger, L. Holm, E. L. Sonnhammer, S. R. Eddy, A. Bateman and R. D. Finn, The pfam protein families database, *Nucleic Acids Res* **40**, D290 (2012).
33. M. Ashburner, C. A. Ball, J. A. Blake, D. Botstein, H. Butler, J. M. Cherry, A. P. Davis, K. Dolinski, S. S. Dwight, J. T. Eppig, M. A. Harris, D. P. Hill, L. Issel-Tarver, A. Kasarskis, S. Lewis, J. C. Matese, J. E. Richardson, M. Ringwald, G. M. Rubin and G. Sherlock, Gene ontology: tool for the unification of biology, *Nature Genetics* **25**, 25 (may 2000).
34. J. Pandey, M. Koyutürk, S. Subramaniam and A. Grama, Functional coherence in domain interaction networks, *Bioinformatics* **24**, i28 (Aug 2008).
35. C. Y. Ma and C. S. Liao, A review of protein-protein interaction network alignment: From pathway comparison to global alignment, *Computational and Structural Biotechnology Journal* **18**, 2647 (2020).
36. V. Vijayan, V. Saraph and T. Milenkovic, Magna++: Maximizing accuracy in global network alignment via both node and edge conservation, *Bioinformatics* **31**, 2409 (07 2015).
37. S. Monti, P. Tamayo, J. Mesirov and T. Golub, Consensus clustering: A resampling-based method for class discovery and visualization of gene expression microarray data, *Machine Learning* **52**, 91 (2003).
38. A. Tsherniak, F. Vazquez, G. Montgomery, D. A. Weir, G. Kryukov, G. S. Cowley, S. Gill, W. F. Harrington, S. Pantel, R. Jménez *et al.*, Defining a cancer dependency map, *Cell* **170**, 564 (2017).
39. A. Walther, E. Johnstone, C. Swanton, R. Midgley, I. Tomlinson and D. Kerr, Genetic prognostic and predictive markers in colorectal cancer, *Nature Reviews Cancer* **9**, 489 (jun 2009).
40. MetricGate, Cliff's Delta Effect Size Calculator `https://metricgate.com/docs/cliffs-delta/`, (2024), Online statistical calculator.
41. B. Vogelstein, E. R. Fearon, S. R. Hamilton, S. E. Kern, A. C. Preisinger, M. Leppert, Y. Nakamura, R. White, A. M. Smits and J. L. Bos, Genetic alterations during colorectal-tumor development, *N Engl J Med* **319**, 525 (1988).
42. K. W. Kinzler and B. Vogelstein, Lessons from hereditary colorectal cancer, *Cell* **87**, 159 (1996).
43. J. Pandey, M. Koyutürk, Y. Kim, W. Szpankowski, S. Subramaniam and A. Grama, Functional annotation of regulatory pathways, *Bioinformatics* **23**, i377 (Jul 2007).
44. E. Banks, E. Nabieva, R. Peterson and M. Singh, Netgrep: fast network schema searches in interactomes, *Genome Biol* **9**, p. R138 (2008).