

A Computational Workflow for Analysis of Missense Mutations in Precision Oncology

Rayyan Tariq Khan^{1,3}, Petra Pokorna^{5,7}, Jan Stourac^{1,2,3}, Simeon Borko^{2,3,4}, Ihor Arefiev^{1,2}, Joan Planas-Iglesias^{1,2,3}, Adam Dobias^{1,2}, Gaspar Pinto^{1,2,3}, Veronika Szotkowska^{1,2}, Jaroslav Sterba⁶, Ondrej Slaby^{5,7}, Jiri Damborsky^{1,2,3}, Stanislav Mazurenko^{1,2,3,*}, David Bednar^{1,2,3,*}

¹ Loschmidt Laboratories, Department of Experimental Biology, Faculty of Science, Masaryk University, Brno, Czech Republic; ² Loschmidt Laboratories, RECETOX, Faculty of Science, Masaryk University, Brno, Czech Republic; ³ International Clinical Research Center, St. Anne's University Hospital Brno, Brno, Czech Republic; ⁴ IT4Innovations Centre of Excellence, Faculty of Information Technology, Brno University of Technology, Brno, Czech Republic; ⁵ Central European Institute of Technology, Masaryk University, Brno, Czech Republic; ⁶ Department of Paediatric Oncology, University Hospital Brno and Faculty of Medicine, Masaryk University, Brno, Czech Republic; ⁷ Department of Biology, Faculty of Medicine, Masaryk University, Brno, Czech Republic

* Corresponding authors: SM: mazurenko@mail.muni.cz; DB: davidbednar1208@gmail.com

Keywords: Bioinformatics; Cancer; Function; High-Performance Computing; Machine learning; Molecular Modelling; Oncology; Personalised Medicine; Single Nucleotide Polymorphism; Stability; Treatment

Abstract

Every year, more than 19 million cancer cases are diagnosed, and this number continues to increase annually. Since standard treatment options have varying success rates for different types of cancer, understanding the biology of an individual's tumour becomes crucial, especially for cases that are difficult to treat. Personalised high-throughput profiling, using next-generation sequencing, allows for a comprehensive examination of biopsy specimens. Furthermore, the widespread use of this technology has generated a wealth of information on cancer-specific gene alterations. However, there exists a significant gap between identified alterations and their proven impact on protein function. Here, we present a bioinformatics pipeline that enables fast analysis of a missense mutation's effect on stability and function in known oncogenic proteins. This pipeline is coupled with a predictor that summarises the outputs of different tools used throughout the pipeline, providing a single probability score, achieving the balanced accuracy above 86%. The pipeline incorporates a virtual screening method to suggest potential FDA/EMA-approved drugs to be considered for treatment. We showcase three case studies to demonstrate the timely utility of this pipeline. To facilitate access and analysis of cancer-related mutations, we have packaged the pipeline as a web server, which is freely available at <https://loschmidt.chemi.muni.cz/predictonco/>.

Introduction

More than 19 million cancer cases were diagnosed in 2020 (Cancer Today, 2020) with a projected load of 28.4 million cases in 2040 (Sung et al., 2021). The three traditionally used approaches to treat cancer, namely chemotherapy, surgery, and radiotherapy, generally result in higher mortality rates compared to the less adopted precision medicine-based techniques (Krzyszczuk et al. 2018). Next-generation sequencing technologies form the basis of precision oncology and can help generate a large amount of transcriptomic and genomic

data. On the other hand, these technologies often do not provide clinically actionable data. This leads to a divide between generation of the said data and their utility, as mutants with unknown effects are often found during clinical testing (Buzdin et al., 2021).

There are not many tools that can help bridge the gap between data generation and creation of actionable insights. Swiss-PO, an online tool, allows for mapping experimentally determined mutations on a curated list of 50 genes and their various associated 3D structures. It also allows users to visualise multiple molecular interactions; however, it leaves it to the user to intuitively assess the structural implications of mutations that have not been experimentally determined (Krebs et al., 2021) and it can also not predict patient survival outcomes. PSnpBind, a database, catalogues changes to binding affinities of ligands due to binding site single-nucleotide polymorphisms (SNPs), however this database is limited to 26 human proteins and is limited to interactions between ligands and binding site residues (Ammar et al., 2022). We sought to overcome some of these limitations by creating a robust pipeline that can predict the effects of missense mutations, even for ones which are not experimentally determined, on cancer-related proteins.

The pipeline relies on advances in fast protein modelling, such as AlphaFold (Jumper et al., 2021), prediction of the effect of missense mutations on a protein structure (Bendl et al., 2014), and protein stability prediction (Blanco et al., 2018; Kellogg et al., 2010). This allows for a lot more information to be harvested from mutations identified by exome sequencing and can then be used for actionable decision making. Additionally, coupling fast ligand docking in proteins (Trott & Olson, 2009) and with the availability of multiple drug libraries online, such as ZINC (Irwin et al., 2012), it is possible to screen novel potential inhibitors for the mutated proteins.

As the interpretation of large-scale genomic and transcriptomic data is limited due to the need to utilise multiple computational tools, performing the aforementioned analysis on exome sequences can take time if done manually. After a cancer diagnosis, treatment is generally a race against time, and with the variable success rates of conventional “one size fits all” therapies, fast and accurate interpretation of molecular findings and assessment of their actionability are of vital importance, especially in difficult-to-treat cases. This is where an automated precision oncology approach will be most useful as it can optimise treatment strategies, improve outcomes and increase the quality of life for many patients (Lassen et al., 2021).

Here we introduce a bioinformatics pipeline for the analysis of the effect of mutations on stability and function in cancer-related proteins. The pipeline applies *in silico* methods of molecular modelling, structural bioinformatics and machine learning, and outputs actionable data which can be used for decision making. The coupled predictor produces a decision on the oncogenicity of the protein mutation by utilising the outputs derived at various stages of the pipeline. Moreover, we show the application of the pipeline on three use cases and highlight the importance of advanced bioinformatics in precision oncology.

Methodology

Manual Curation, Structure Repairs and Geometry Optimization

A list of 44 cancer-related proteins (including 1 isoform of a selected protein) were chosen as targets for the manual curation. The selection was based on the importance of the respective proteins for cancer diagnostics and, notably, in cancer treatment. The vast majority of curated proteins are either direct targets of therapeutic agents or, despite not being targets themselves, represent established predictive biomarkers for administering targeted treatments aimed at downstream members of the same pathway. Additionally, we included proteins that are frequently altered across various cancer types and are relevant to both diagnostics and cancer research (e.g., p53). The proteins with their various annotations are listed in the Supplementary material [SI 1](#).

The 44 protein sequences and their annotations were fetched from the UniProt database (The UniProt Consortium, 2023). In the case of KRAS, two isoforms are provided, including the canonical isoform and an isoform that is commonly utilized across clinical databases of genetic variants. The essential residues were re-confirmed in the literature as well as in the Mechanism and Catalytic Site Atlas (M-CSA, Ribeiro et al., 2017) and the SWISS-PROT (Boeckmann, 2003) databases. For the purposes of this study, in the case of multi-domain proteins, only the catalytic cytoplasmic domains of the proteins were considered. The best available structure from the wwPDB database (wwPDB consortium, 2019), the ideal biological assembly, as well as the relevant chain (in multimeric structures) were selected based on resolution and missing parts. Canonical co-factors for structures were established using the UniProt database, these were retained in the structure, and all other ligands, ions, and water molecules were removed from the structure ([SI 1](#)). The residue indexes were mapped using the SIFTS database (Dana et al., 2018). After visual inspection of each target protein, these four key problematic regions/positions were identified: (i) missing regions, i.e., low resolution regions in the crystal structure; (ii) long, missing and/or intrinsically disordered regions not influencing the catalytic site of the protein; (iii) missing atoms in the side chain; (iv) amino acids requiring identity correction, i.e., the sequence in the 3D structure did not correspond to that recorded in UniProt.

Each protein structure which required any of these structural improvements (for the aforementioned problematic regions/positions i, iii, or iv), was modelled using MODELLER version 9.24, 2020/04/06, r11614 (Eswar et al., 2008). The modelling was guided by the UniProt-PDB alignment provided by SIFTS. Regions identified as intrinsically disordered (repair ii) were omitted from the modelling. Custom extensions of three MODELLER python classes (Environment, Model, and AutoModel) were developed to ensure that the produced models: (i) incorporated any relevant co-factor from the template, (ii) were not optimised on the regions that did not require repairs, and (iii) structures containing multiple chains could be modelled and minimised at once. If no experimental structure was available, the AlphaFold database (Jumper et al., 2021) was searched. The mutant structure was generated by introducing the desired mutation in the target wild type structure by MODELLER, and it was guided by a trivial alignment between the wild type and the mutant sequences.

For each protein structure, inconsistent torsion angles, total energy, or Van der Waals clashes were reduced using RepairPDB feature of FoldX 4.0 (Blanco et al., 2018). Then

minimization of structures was performed in Rosetta 3.11-static (Kellogg et al., 2010) with constraints using the Talaris2014 force field (O'Meara et al., 2015). The wild type and mutant structures were then aligned using DeepAlign 1.135-2-foss-2018b (Jiménez-Moreno et al., 2021) to ensure that their coordinates match for further analysis.

Protein Stability Prediction

The impact of the missense mutation on the stability of the protein structure was calculated using Rosetta and FoldX. For FoldX the PssmStability command was used, water molecules were only taken from the 'crystal', pH was set to 7 and the number of runs was set to 5. Rosetta calculations were made on the minimised structures using the ddg_monomer command, following protocol 3 (Kellogg et al., 2010), for which the extent of sidechain repacking is set to within 8Å while using the soft-rep energy function and the Talaris2014 force field.

Protein Function Prediction, Phylogenetic Analysis and Consensus Classification

Additionally, PropKa 3.4.0 (Rostkowski et al., 2011) was used to predict the impact of the mutation on the pKA values of the proteins, using the propka3 command. Homologous sequences with sufficient identity (more than 50%) and coverage (\pm 20% of the query sequence), i.e., UniRef50 sequences, were downloaded from the UniRef database (Suzek et al., 2014) and multiple sequence alignment were generated using Clustal-Omega (Sievers et al., 2011) tool from the EMBL-EBI web server (Madeira et al., 2022). This was used for conservation analysis using Jensen-Shannon Divergence algorithm (Capra & Singh, 2007) and transformed to mutability grades by using HotSpot Wizard (Sumbalova et al., 2018) thresholding. The mutations were also submitted to the HOPE (Venselaar et al., 2010) web server to collect information from a multitude of information sources; including calculations on the 3D coordinates of the protein, sequence annotations from the UniProt database, and predictions by DAS (Distributed Annotation System) services (Prlić et al., 2007). Furthermore, PredictSNP (Bendl et al., 2014), was used to predict the effect of the amino acid substitution on the target protein function through consensus classification.

Pocket Analysis and Virtual Screening

Potential binding pockets within the structures of the analysed proteins were calculated using the prank predict command in P2Rank 2.3 (Krivák & Hoksza, 2018), the resulting pockets were visually analysed and manually optimised to cover the entire binding sites. Selected pockets were listed in [SI 2](#) according to their colour codes.

Virtual screening was performed on both the wild type and the mutant protein structure. A set of 4380 small molecules that were approved by the Food and Drug Administration and European Medicines Agency was taken from the ZINC database (Irwin et al., 2012). AutoDock Vina 1.1.2 (Trott & Olson, 2009) was run using the standard vina command, within a parameterized grid within each protein. The grid coordinates ([SI 1](#)) were created by visually placing the grid on the protein structure in PyMOL using the ADPlugin (Seeliger & de Groot, 2010) and ensuring that the binding pockets with essential residues were completely within the grid. The values for the binding energy of each small molecule to

a wild type structure as well as its mutant structure were used to calculate the impact of the mutation on the binding energy.

Machine Learning Predictor Development

The predictive part of the pipeline is a machine-learning based tool that was trained on 1073 single-point mutants whose effect was classified as Oncogenic or Benign. The variants for the Benign class were selected from the gnomAD and ClinVar (Landrum et al., 2017) databases. Variants with >1% population frequency in gnomAD, variants annotated as „benign“ or „likely benign“ in the ClinVar database, and variants without ClinVar annotation, for which the classification as „benign“ or „likely benign“ is at the same time supported by applicable ACMG criteria (Richards et al., 2015) were utilised. The variants for the Oncogenic class were collected in expert-curated precision oncology knowledgebases, mainly, but not limited to, precision oncology knowledge base OncoKB by Memorial Sloan Kettering Cancer Center (Chakravarty et al., 2017), as well as The JAX Clinical Knowledgebase by The Jackson Laboratory (Patterson et al., 2016), Personalized Cancer Therapy Knowledge Base by MD Anderson Cancer Center (Kurnit et al., 2017), cBioPortal (Gao et al., 2013), and the DoCM database (Ainscough et al., 2016). Variants with conflicting interpretations across multiple sources were not included in the list. Both subsets were manually filtered for any possible overlaps with the datasets used in the PredictSNP consensus predictor and its constituents.

The entire dataset (SEQ: 509 oncogenic and 564 benign data points) was further annotated by the pipeline of PredictONCO. The following six features were calculated regardless of the structural information available: essentiality of the mutated residue (yes -1 /no - 0), the conservation of the position (the conservation grade and MSA score), the domain where the mutation is located (“cytoplasmic”, “extracellular”, “transmembrane”, “other” - one-hot encoded), the PredictSNP score, and the number of essential residues in the protein. For approximately half of the data (STR: 377 oncogenic and 76 benign data points), the structural information was available, and six more features were calculated: FoldX and Rosetta ddg_monomer scores, whether the residue is in the ligand-binding pocket obtained from P2Rank (yes -1 /no - 0), and the pKa changes of essential residues obtained from PROPKA3. The dataset is available at <https://zenodo.org/records/10013764>.

For the training protocol, 20% of the data in each of the two sets was kept aside for testing, chosen randomly but grouped by positions to ensure that no specific position in a protein from the test set appears in the training set. The following types of predictors were tested: the support vector machine (SVM), decision tree (DT), and XGBoost classifier (XGB), taken as they are implemented in the scikit-learn 1.2.0 and xgboost 1.7.3 libraries for Python 3.8.15. We also used the PredictSNP score alone as a baseline. For each method, we tested a set of hyperparameters based on 5-fold cross-validation implemented on the training data and receiver operating characteristic (ROC) area under the curve (AUC) as the metric (Table S1 in [SI 3](#)).

The final evaluation consisted of constructing the ROC and Precision-Recall curves. Furthermore, a round of 100 random-state re-initialisations with different random seeds was performed to evaluate the robustness of the final models. For the area under the ROC curve and the average precision values, we reported the average and standard deviation from the

100 runs. Since any change to the predictor or data split results in a different set of x-axis coordinates in the ROC and Precision-Recall curves, we used a fixed grid of 30 points and applied 1D linear interpolation to obtain the y-axis value for each iteration. These values were then plotted as 10% and 90% quantiles.

The version numbers for softwares and python packages that were used are noted in [SI 4](#).

Results

Development of a Fully Automated Computational Workflow

We created a bioinformatics pipeline for structure and sequence based analysis of the effects of missense mutations on cancer-related proteins (Figure S1). Since the pipeline requires curated protein structures, a method for curation was developed and applied to a list of 44 proteins ([SI 1](#)) which were then tested to ensure they can be handled in the pipeline. The pipeline was assembled using multiple bioinformatics tools, databases, and techniques. **Figure 1** represents a schematic outline of the pipeline, the output of which ultimately feeds into the machine learning predictor. The predictor gives a binary decision on the effect of mutation with confidence score which is helpful in the summation and comprehension of results. Three cases of oncological interest were then studied using the developed method.

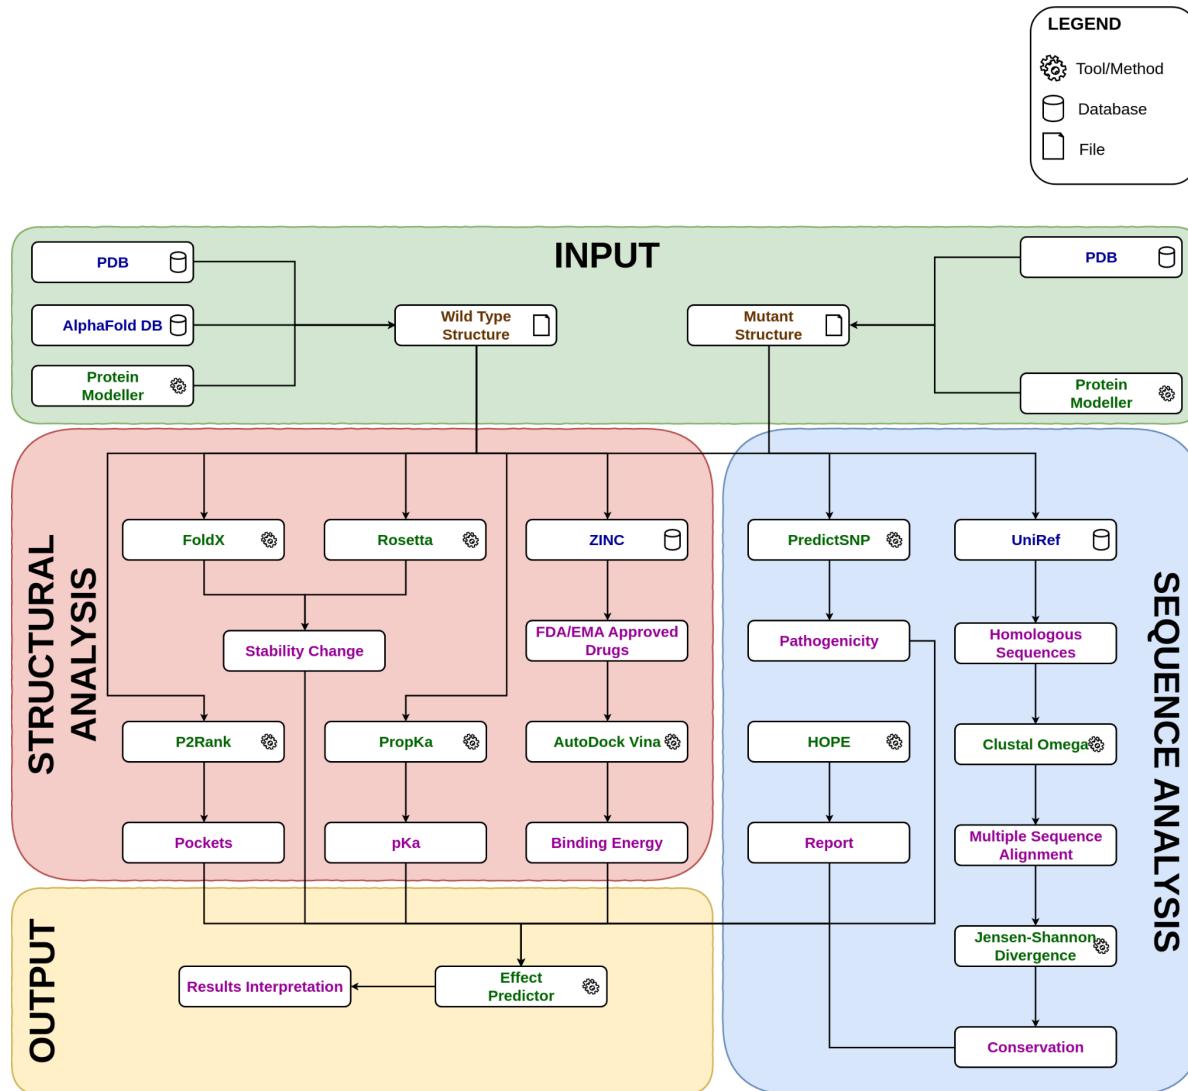


Figure 1. A schematic representation of the bioinformatic pipeline used to predict the effect of a missense mutation on the oncogenicity of the protein.

Training of Sequence-Based and Structure-Based Machine Learning Predictors

Initially, we trained three different types of predictors, covering different trade-offs between explainability and flexibility, and compared their performance with the baseline model using the PredictSNP score alone. After optimising the hyperparameters (Table S1 in [SI 3](#)), we evaluated the performance on the held-out 20% of the dataset split by position in a protein. The support vector machines and XGBoost classifiers showed superior yet similar performance based on the area under the ROC curve and the average precision from the Precision-Recall curve ([Figure 2](#)), also confirmed statistically ([Figure S2 in SI 3](#)). We selected the XGBoost predictor for the final model due to the interpretability of its scores: the SVM model evaluation is based on the signed distance to the separating hyperplane, without intuitive interpretation. On the other hand, the XGBoost classifier directly returns the probability that a particular mutation is oncogenic. The final XGBoost predictor is made up of 15 and 9 decision trees of the depth of 1 for structure and sequence data sets, respectively. The feature importance scores revealed that the PredictSNP score and conservation had the highest information gains ([Figure S3 in SI 3](#)). We also tested if using the train/test split by

proteins would compromise the performance and saw only a marginal decrease (Figure S4 in [SI 3](#)), indicating the significant potential of the pipeline for other protein targets. The balanced accuracy for the sequence-based XGBoost predictor is 87%, and for the structure-based XGBoost predictor is 90%.

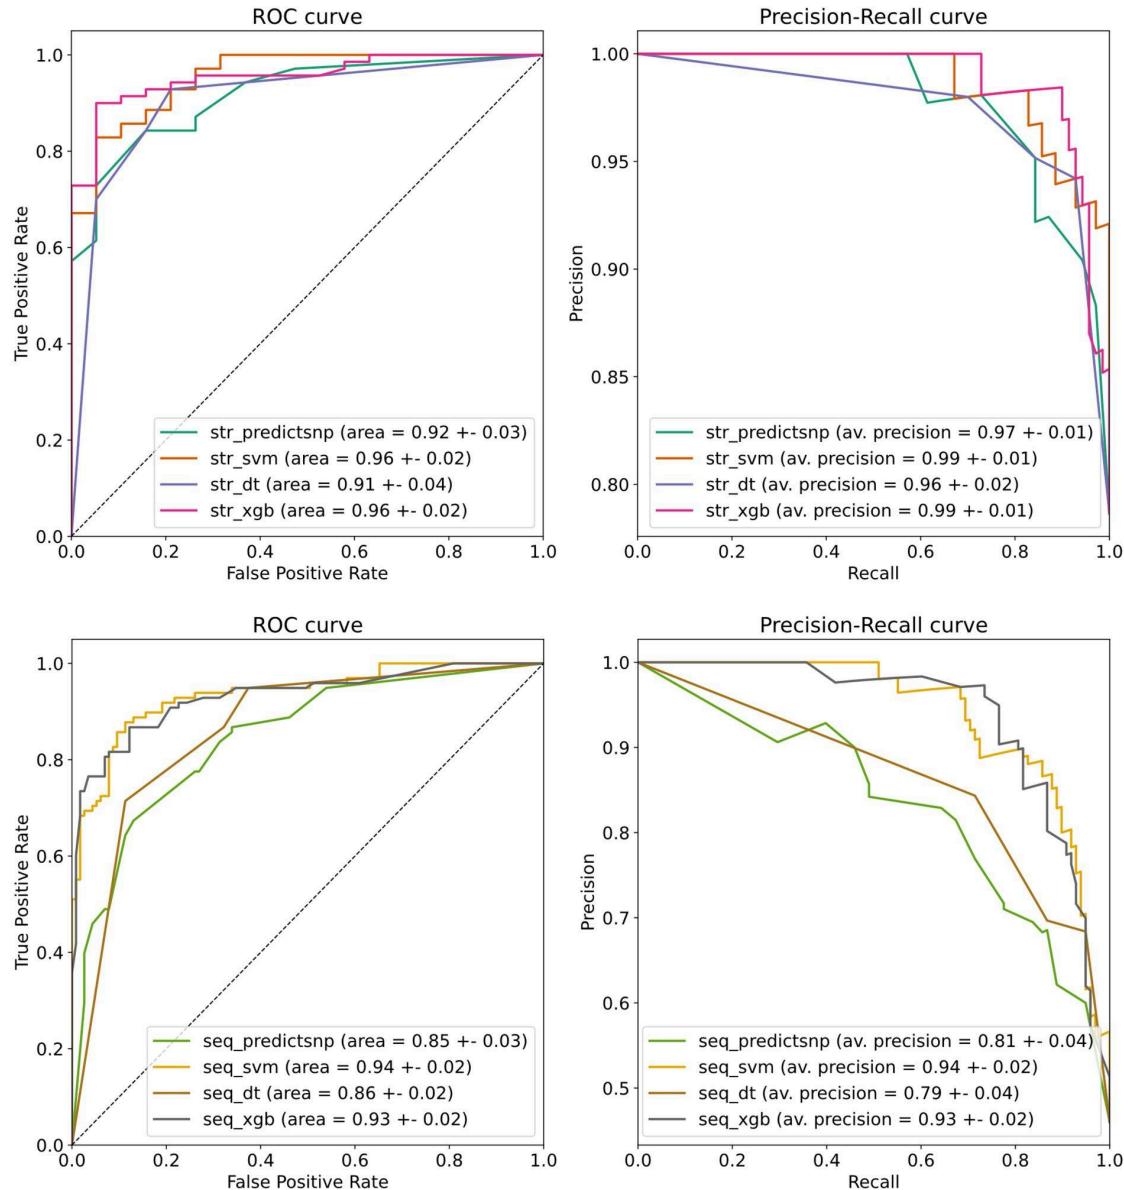


Figure 2. The Receiver Operating Characteristic and Precision-Recall curves based on held-out test sets. Top: classifiers trained on the dataset with the structural features available (STR). Bottom: classifiers trained on the dataset with the sequence-only features (SEQ). Both the support vector machine (SVM) and XGBoost (XGB) showed comparable performance superior to the baseline model and decision tree (DT). The reported errors are standard deviations obtained by bootstrapping (N=1000). The PredictSNP score was used as the baseline.

We also compared the performance of our predictor on the test set against several other models (Table 1). We evaluated the following individual scores as baselines: conservation, predictSNP, FoldX, and Rosetta. In addition, we evaluated the performance of the ESM variants model, a recently published workflow based on the 650-million-parameter protein language ESM1b, which was used to score all possible missense variant effects in the

human genome (Brandes et al., 2023). In both settings (SEQ and STR), PredictONCO showed superior performance.

Table 1. Comparison of PredictONCO with other models on the test set. The models selected for comparison were individual features and the ESM variants predictor. The reported errors are standard deviations obtained by bootstrapping (N=1000).

	Predictor	ROC AUC↑	Avg. Precision↑
SEQ	PredictONCO	0.932 ± 0.018	0.934 ± 0.018
	conservation	0.872 ± 0.026	0.802 ± 0.042
	predictSNP	0.845 ± 0.027	0.808 ± 0.041
	ESM variants	0.923 ± 0.018	0.911 ± 0.023
STR	PredictONCO	0.955 ± 0.020	0.988 ± 0.006
	FoldX	0.575 ± 0.064	0.867 ± 0.037
	Rosetta	0.628 ± 0.064	0.876 ± 0.039
	conservation	0.937 ± 0.037	0.970 ± 0.020
	predictSNP	0.918 ± 0.030	0.973 ± 0.011
	ESM variants	0.929 ± 0.027	0.981 ± 0.009

Case Studies with Selected Cancer-associated Proteins

The following case studies demonstrate scenarios in which the tool has helped to facilitate further clinical decision-making. The respective variants featured in the case studies were identified across research projects utilizing high-throughput DNA sequencing techniques, which were conducted by the co-authors of this manuscript.

Case study 1 - Platelet Derived Growth Factor Receptor Beta PDGFRB N666T

In a patient with myofibroma, sequencing analysis revealed an N666T variant of the PDGFRB protein (UniProt ID:P09619). Even though some mutations of the N666 residue, including N666K (Iwamura et al. 2023), N666H (Pond et al. 2018), or N666S (Ortiz et al. 2020), have already been documented in myofibroma patients, N666T, in particular, lacks published functional evidence and was reported in a total of one patient in combination with another mutation. Therefore, a comprehensive assessment of its effect would provide further confirmatory evidence on the variant's pathogenicity, which is substantial, given the therapeutic implications of receptor tyrosine kinase inhibition. Conservation status showed high evolutionary conservation, which is typically damaging to the protein. For amino acid 826, one of the essential catalytic residues, a large increase in dissociation constant was predicted, suggesting a significant functional impact. Both stability predictors suggested a

deleterious effect, which is also in agreement with the deleterious effect on protein function predicted by PredictSNP. Given all this data, the oncogenic effect was predicted by the XGBoost classifier with 100% confidence. Furthermore, in virtual screening, Sunitinib showed a slightly better increase in binding affinity compared to Imatinib, which was used as a drug of choice in different myofibroma preclinical studies, making Sunitinib a suitable alternative option for therapeutic planning. The full report can be accessed at - https://loschmidt.chemi.muni.cz/predictonco/job/pdgfrb_N666T

Case study 2 - Angiopoietin-1 receptor TIE2 G1036D

In a patient with a vascular tumour, sequencing analysis revealed a G1036D variant in the TIE2 (UniProt ID:Q02763) gene. The G1036D variant represents a previously undescribed alteration, which has not been documented in the literature, clinical, or population databases of genetic variants. Given the rapidly evolving field of vascular tumour genetics and the possibility of targeted therapeutics administration, identifying novel potentially activating alterations is vastly important. Although the residue is non-essential, moderately evolutionarily conserved, and only moderate changes were predicted for the catalytic residues, the overall impact was evaluated by the XGBoost classifier as oncogenic with a 99% confidence score and was based on a deleterious prediction by both the PredictSNP algorithm and stability predictors FoldX and Rosetta. This could be approached as a basis to facilitate further functional tests to measure mutant receptor phosphorylation and, if proven as activating, introduce a considerable therapeutic opportunity (by potentially using one of the suggested inhibitive compounds such as Ecteinascidin, Ponatinib, etc., or other inhibitors of downstream signalling cascade) as well as an addition to the knowledge on disease pathogenesis. The full report can be accessed at - https://loschmidt.chemi.muni.cz/predictonco/job/tie2_G1036D

Case study 3 - Cellular Tumor Antigen p53 K101Q

In next-generation sequencing screening for cancer predispositions, the K101Q variant of p53 (UniProt ID:P04637) was identified in an individual with a negative family history of cancer. p53 represents the most commonly altered gene in all cancers, and p53 variants predispose to cancer development when of germline origin. Therefore, a careful assessment must be performed for further genetic counselling. The respective variant has not been documented in the literature or functionally characterised. With lacking evidence from literature and databases of genetic variants, typically only prediction algorithms that employ sequence-based information without structural data are available. Therefore, combining both structural and sequence related perspectives might yield a more accurate prediction. The XGBoost classifier predicted the mutation as neutral with an 81% confidence score, supported by both the PredictSNP prediction and the stability predictors. Information on evolutionary conservation showed that the wild-type residue is not conserved at this position, which may suggest that the variant is not damaging to the protein. Based on these results and no family history of cancer, the variant should not influence subsequent clinical management. Given the importance of p53 variants in both somatic and germline contexts and their same functional impact, this case study exemplifies the utility of the tool in the assessment of hereditary cancer predisposition. The full report can be accessed at the following link - https://loschmidt.chemi.muni.cz/predictonco/job/p53_K101Q

Discussion

Prediction of the effect of missense mutations on cancer-related protein structures is a complicated task. This paper presents our pipeline for tackling this problem thus allowing clinical bioinformaticians to easily run multiple cancer-related analyses for their target mutations on a curated list of proteins.

A major part of the pipeline capitalises on structural bioinformatics and it requires the presence of good quality protein structures for accurate analysis. However, a good amount of cancer-associated structures are transmembrane channels and thus only have fragmented domain level structures. Some of them can be multimeric, and thus modelling proves a challenge. Despite AlphaFold (Jumper et al., 2021) being touted as a major groundbreaker in the field of protein structure modelling, it proves inefficient in modelling large multi-subunit, multimeric proteins as quaternary domain level interactions are difficult to model. Thus the structural bioinformatics part of the pipeline is limited to working with high-quality structures at the domain level. AlphaFold-Multimer (Evans et al., 2021) can be used to predict the multimeric conformation in 70% of heteromeric cases and 72% of homomeric cases to limit this problem; it is unclear whether this accuracy of predictions is viable for working with oncogenic or tumour suppressor proteins, especially when the final prediction will likely be used in a medical context.

Currently, the web server provides predictions for 44 target proteins, which were selected based on their relevance to the field of oncology. Appropriate processing of a new structure to be used in the pipeline requires expert level knowledge of multiple bioinformatic tools. Curation in this field is a recognized bottleneck, especially in the case of the interpretation of results (Bungartz et al., 2018.) The addition of new target proteins to the internal database connected to the PredictONCO web server is possible and it is offered to the user community based on direct requests. Once a protein is curated, all mutations on its structure can be easily analysed. Moreover, the pipeline can also work with sequence only data and the trained XGBoost classifier can also reliably predict using only the sequence based features with only a 4% drop in average precision.

The pipeline has no standard run time as it mostly depends on whether structural analysis needs to be performed along with sequence-based analysis or not. The structural analysis increases the computational load and the complexity of the structure can further increase the run time. However, the calculations generally do not take more than two days to complete. It is unclear whether this time frame is long or short as run time benchmarking would require the existence of other similar tools, techniques or pipelines for comparative purposes, and specialised methodologies that deal with the same case do not exist. However, this time window meets the initial requirements for the use of the web server in clinical practice as well as for research and educational purposes. Furthermore, it helps assist in making the result interpretation step easier as interpretation itself is a recognized bottleneck (Bungartz et al. 2018.)

Comparison to other similar tools is difficult as, as of this writing, we did not come across a pipeline integrating multiple approaches to predict the effect of a missense mutation on a cancer-related protein. However several databases and online data integrating tools do exist. The two most prominent of these databases would be the International Cancer Genome Consortium (ICGC) (The International Cancer Genome Consortium, 2010) and The

Cancer Genome Atlas (TCGA) (Weinstein et al. 2013). Furthermore, survival analysis tools also exist and are primarily based on 4 types of data: (i) mRNA data, such as PRECOG (Gentles et al., 2015), (ii) ncRNA data, such as OncoLnc (Anaya, 2016), (iii) DNA methylation and mutation data, such as cBioPortal (Gao et al., 2013), and (iv) Protein data, such as TCPA (Li et al., 2013). Additionally, the Swiss-PO web tool for mapping gene mutations on the 3D structure can be used, but it only allows for intuitive and qualitative analysis of mutations that have already been experimentally determined (Krebs et al., 2021). Comparison to the aforementioned database PSnpBind is also difficult as it only catalogues changes to binding affinities of ligands due to binding site single-nucleotide polymorphisms (SNPs) (Ammar et al., 2022).

Our pipeline currently only supports missense mutations, as it is unable to handle insertions, deletions, or fusions of oncogenic proteins because individual tools in the pipeline are not able to analyse them. However, substitutions do make up a large number of cancer associated mutations as a large number of genes associated with various cancer types contain single nucleotide variants (Deng et al., 2017). For common solid tumours, 95% of cancer driver mutations in humans are single-base substitutions. Approximately, 90.7% of these result in the amino acid being substituted for another, non-synonymous one (Darbyshire et al. 2019). Thus, even though insertions, deletions, and fusions cannot be analysed using the pipeline, it still provides predictions for a significant majority of cancer-related alterations. The tool is freely accessible to the community of bioinformaticians and medical doctors and will provide fast and useful analysis of data from sequencing of patient samples.

Funding

The authors would like to express their thanks to the Czech Ministry of Education [TEAMING -CZ.02.1.01/0.0/0.0/17_043/0009632; ESFRI CZECRIN LM2023049; ESFRI eINFRA LM2018140]; the Technology Agency of the Czech Republic [TREND FW03010208; PERMED TN02000109]; European Union [TEAMING 857560]; Brno University of Technology [FIT-S-23-8209]; Ministry of Health [NU20-03-00240]. The research was further supported by the project National Institute for Oncology Research [Programme EXCELES, ID Project No. LX22NPO5102 funded by the European Union — Next Generation EU].

Availability of data and materials

- The pipeline is available as a web server, at <https://loschmidt.chemi.muni.cz/predictonco/>.
- The list of proteins, definition of binding pockets, and ML model validation are attached as supplementary files.
- The training and testing dataset are available at <https://zenodo.org/records/10013764>.

Conflict of interest statement. None declared.

References

Ainscough, Benjamin J., et al. "DoCM: A Database of Curated Mutations in Cancer." *Nature Methods*, vol. 13, no. 10, Sept. 2016, pp. 806–7, doi:<https://doi.org/10.1038/nmeth.4000>.

Ammar, Ammar, et al. "PSnpBind: a database of mutated binding site protein–ligand complexes constructed using a multithreaded virtual screening workflow." *Journal of Cheminformatics* 14.1 February 2022, 8, doi:<https://doi.org/10.1186/s13321-021-00573-5>

Anaya, Jordan. "OncoLnc: Linking TCGA Survival Data to MRNAs, MiRNAs, and LncRNAs." *PeerJ Computer Science*, vol. 2, 13 June 2016, p. e67, <https://doi.org/10.7717/peerj-cs.67>.

Bendl, Jaroslav, et al. "PredictSNP: Robust and Accurate Consensus Classifier for Prediction of Disease-Related Mutations." *PLoS Computational Biology*, edited by Paul P. Gardner, vol. 10, no. 1, Jan. 2014, p. e1003440, doi:<https://doi.org/10.1371/journal.pcbi.1003440>.

Blanco, Javier Delgado, et al. "FoldX Accurate Structural Protein–DNA Binding Prediction Using PADA1 (Protein Assisted DNA Assembly 1)." *Nucleic Acids Research*, vol. 46, no. 8, Mar. 2018, pp. 3852–63, doi:<https://doi.org/10.1093/nar/gky228>.

Boeckmann, B. "The SWISS-PROT Protein Knowledgebase and Its Supplement TrEMBL in 2003." *Nucleic Acids Research*, vol. 31, no. 1, Jan. 2003, pp. 365–70, doi:<https://doi.org/10.1093/nar/gkg095>.

Bungartz, Kathryn D, et al. "Making the Right Calls in Precision Oncology." *Nature Biotechnology*, vol. 36, no. 8, 6 Aug. 2018, pp. 692–696, doi:<https://doi.org/10.1038/nbt.4214>.

Brandes, Nadav, et al. "Genome-wide prediction of disease variant effects with a deep protein language model." *Nature Genetics* 55.9, Sep. 2023 1512-1522, doi:<https://doi.org/10.1038/s41588-023-01465-0>

Buzdin, Anton, et al. "Editorial: Next Generation Sequencing Based Diagnostic Approaches in Clinical Oncology." *Frontiers in Oncology*, vol. 10, Jan. 2021, doi:<https://doi.org/10.3389/fonc.2020.635555>.

"Cancer Today." *Iarc.fr*, 2020, gco.iarc.fr/today/home.

Capra, John A., and Mona Singh. "Predicting Functionally Important Residues from Sequence Conservation." *Bioinformatics*, vol. 23, no. 15, May 2007, pp. 1875–82, doi:<https://doi.org/10.1093/bioinformatics/btm270>.

Chakravarty, Debyani, et al. "OncoKB: A Precision Oncology Knowledge Base." *JCO Precision Oncology*, no. 1, Nov. 2017, pp. 1–16, doi:<https://doi.org/10.1200/po.17.00011>.

Dana, Jose M., et al. "SIFTS: Updated Structure Integration with Function, Taxonomy and Sequences Resource Allows 40-Fold Increase in Coverage of Structure-Based

Annotations for Proteins." *Nucleic Acids Research*, vol. 47, no. D1, Nov. 2018, pp. D482–89, doi:<https://doi.org/10.1093/nar/gky1114>.

Darbyshire, Madeleine, et al. "Estimating the frequency of single point driver mutations across common solid tumours." *Scientific Reports*, 9.1, Sep. 2019, 13452, doi: <https://doi.org/10.1038/s41598-019-48765-2>

Deng, Na, et al. "Single Nucleotide Polymorphisms and Cancer Susceptibility." *Oncotarget*, vol. 8, no. 66, 10 Nov. 2017, <https://doi.org/10.18632/oncotarget.22372>. Accessed 17 Apr. 2019.

Eswar, Narayanan, et al. "Protein Structure Modeling with MODELLER." *Methods in Molecular Biology*, Jan. 2008, pp. 145–59, doi:https://doi.org/10.1007/978-1-60327-058-8_8.

Evans, Richard, et al. "Protein complex prediction with AlphaFold-Multimer." *BioRxiv*, Oct. 2021, doi:<https://doi.org/10.1101/2021.10.04.463034>.

Gao, J., et al. "Integrative Analysis of Complex Cancer Genomics and Clinical Profiles Using the CBioPortal." *Science Signaling*, vol. 6, no. 269, Mar. 2013, pp. pl1–0, doi:<https://doi.org/10.1126/scisignal.2004088>.

Gentles, Andrew J, et al. "The Prognostic Landscape of Genes and Infiltrating Immune Cells across Human Cancers." *Nature Medicine*, vol. 21, no. 8, 20 July 2015, pp. 938–945, <https://doi.org/10.1038/nm.3909>.

Irwin, John J., et al. "ZINC: A Free Tool to Discover Chemistry for Biology." *Journal of Chemical Information and Modeling*, vol. 52, no. 7, June 2012, pp. 1757–68, doi:<https://doi.org/10.1021/ci3001277>.

Iwamura, Ryuji, et al. "PDGFRB and NOTCH3 Mutations Are Detectable in a Wider Range of Pericytic Tumors, Including Myopericytomas, Angioleiomyomas, Glomus Tumors, and Their Combined Tumors." *Modern Pathology*, vol. 36, no. 3, 1 Mar. 2023, p. 100070, doi:<https://doi.org/10.1016/j.modpat.2022.100070>.

Jiménez-Moreno, A., et al. "DeepAlign, a 3D Alignment Method Based on Regionalized Deep Learning for Cryo-EM." *Journal of Structural Biology*, vol. 213, no. 2, June 2021, p. 107712, doi:<https://doi.org/10.1016/j.jsb.2021.107712>.

Johnston, W. T., et al. "Childhood Cancer: Estimating Regional and Global Incidence." *Cancer Epidemiology*, vol. 71, Apr. 2021, p. 101662, doi:<https://doi.org/10.1016/j.canep.2019.101662>.

Jumper, John, et al. "Highly Accurate Protein Structure Prediction with AlphaFold." *Nature*, vol. 596, no. 7873, July 2021, pp. 583–89, doi:<https://doi.org/10.1038/s41586-021-03819-2>.

Kellogg, Elizabeth H., et al. "Role of Conformational Sampling in Computing Mutation-Induced Changes in Protein Structure and Stability." *Proteins: Structure*,

Function, and Bioinformatics, vol. 79, no. 3, Dec. 2010, pp. 830–38, doi:<https://doi.org/10.1002/prot.22921>.

Krebs, Fanny S., et al. "Swiss-PO: a new tool to analyze the impact of mutations on protein three-dimensional structures for precision oncology." *NPJ precision oncology*, 5.1, 2021, 19, doi:<https://doi.org/10.1038/s41698-021-00156-5>

Krivák, Radoslav, and David Hoksza. "P2Rank: Machine Learning Based Tool for Rapid and Accurate Prediction of Ligand Binding Sites from Protein Structure." *Journal of Cheminformatics*, vol. 10, no. 1, Aug. 2018, doi:<https://doi.org/10.1186/s13321-018-0285-8>.

Krzyszczuk, Paulina, et al. "The Growing Role of Precision and Personalized Medicine for Cancer Treatment." *Technology*, no. 03n04, Sept. 2018, pp. 79–100, doi:<https://doi.org/10.1142/s2339547818300020>.

Kurnit, Katherine C., et al. "Personalized Cancer Therapy": A Publicly Available Precision Oncology Resource." *Cancer Research*, vol. 77, no. 21, Oct. 2017, pp. e123–26, doi:<https://doi.org/10.1158/0008-5472.can-17-0341>.

Lam, Catherine G., et al. "Science and Health for All Children with Cancer." *Science*, vol. 363, no. 6432, Mar. 2019, pp. 1182–86, doi:<https://doi.org/10.1126/science.aaw4892>.

Landrum, Melissa J., et al. "ClinVar: Improving Access to Variant Interpretations and Supporting Evidence." *Nucleic Acids Research*, vol. 46, no. D1, Nov. 2017, pp. D1062–67, doi:<https://doi.org/10.1093/nar/gkx1153>.

Lassen, Ulrik N., et al. "Precision oncology: a clinical and patient perspective." *Future Oncology*, 17.30, Jul 2021, 3995–4009, doi:<https://doi.org/10.2217/fon-2021-0688>

Li, Jun, et al. "TCPA: A Resource for Cancer Functional Proteomics Data." *Nature Methods*, vol. 10, no. 11, 15 Sept. 2013, pp. 1046–1047, doi:<https://doi.org/10.1038/nmeth.2650>.

Madeira, Fábio, et al. "Search and Sequence Analysis Tools Services from EMBL-EBI in 2022." *Nucleic Acids Research*, vol. 50, no. W1, Apr. 2022, pp. W276–79, doi:<https://doi.org/10.1093/nar/gkac240>.

O'Meara, Matthew J., et al. "Combined Covalent-Electrostatic Model of Hydrogen Bonding Improves Structure Prediction with Rosetta." *Journal of Chemical Theory and Computation*, vol. 11, no. 2, Jan. 2015, pp. 609–22, doi:<https://doi.org/10.1021/ct500864r>.

Ortiz, Erica, et al. "Invasive Myofibromatosis with Visceral Involvement in a Term Newborn: A Case Report." *American Journal of Pediatrics*, vol. 6, no. 2, 30 Apr. 2020, pp. 173–173, doi:<https://doi.org/10.11648/j.ajp.20200602.30>.

Patterson, Sara E., et al. "The Clinical Trial Landscape in Oncology and Connectivity of Somatic Mutational Profiles to Targeted Therapies." *Human Genomics*, vol. 10, no. 1, Jan. 2016, doi:<https://doi.org/10.1186/s40246-016-0061-7>.

Pond, Dinel, et al. "A Patient with Germ-Line Gain-of-Function PDGFRB P.N666H Mutation and Marked Clinical Response to Imatinib." *Genetics in Medicine*, vol. 20, no. 1, Jan. 2018, pp. 142–150, doi:<https://doi.org/10.1038/gim.2017.104>.

Prlić, Andreas, et al. "Integrating Sequence and Structural Biology with DAS." *BMC Bioinformatics*, vol. 8, no. 1, Sept. 2007, doi:<https://doi.org/10.1186/1471-2105-8-333>.

Ribeiro, António J M, et al. "Mechanism and Catalytic Site Atlas (M-CSA): A Database of Enzyme Reaction Mechanisms and Active Sites." *Nucleic Acids Research*, vol. 46, no. D1, Nov. 2017, pp. D618–23, doi:<https://doi.org/10.1093/nar/gkx1012>.

Richards, Sue, et al. "Standards and Guidelines for the Interpretation of Sequence Variants: A Joint Consensus Recommendation of the American College of Medical Genetics and Genomics and the Association for Molecular Pathology." *Genetics in Medicine*, vol. 17, no. 5, May 2015, pp. 405–24, doi:<https://doi.org/10.1038/gim.2015.30>.

Rostkowski, Michał, et al. "Graphical Analysis of PH-Dependent Properties of Proteins Predicted Using PROPKA." *BMC Structural Biology*, vol. 11, no. 1, Jan. 2011, doi:<https://doi.org/10.1186/1472-6807-11-6>.

Seeliger, Daniel, and Bert L. de Groot. "Ligand Docking and Binding Site Analysis with PyMOL and Autodock/Vina." *Journal of Computer-Aided Molecular Design*, vol. 24, no. 5, Apr. 2010, pp. 417–22, doi:<https://doi.org/10.1007/s10822-010-9352-6>.

Sievers, Fabian, et al. "Fast, Scalable Generation of High-Quality Protein Multiple Sequence Alignments Using Clustal Omega." *Molecular Systems Biology*, vol. 7, no. 1, Jan. 2011, p. 539, doi:<https://doi.org/10.1038/msb.2011.75>.

Steliarova-Foucher, Eva, et al. "International Incidence of Childhood Cancer, 2001–10: A Population-Based Registry Study." *The Lancet Oncology*, vol. 18, no. 6, June 2017, pp. 719–31, doi:[https://doi.org/10.1016/s1470-2045\(17\)30186-9](https://doi.org/10.1016/s1470-2045(17)30186-9).

Sumbalova, Lenka, et al. "HotSpot Wizard 3.0: Web Server for Automated Design of Mutations and Smart Libraries Based on Sequence Input Information." *Nucleic Acids Research*, vol. 46, no. W1, May 2018, pp. W356–62, doi:<https://doi.org/10.1093/nar/gky417>.

Sung, Hyuna, et al. "Global Cancer Statistics 2020: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers in 185 Countries." *CA: A Cancer Journal for Clinicians*, vol. 71, no. 3, Feb. 2021, pp. 209–49, doi:<https://doi.org/10.3322/caac.21660>.

Suzek, Baris E., et al. "UniRef Clusters: A Comprehensive and Scalable Alternative for Improving Sequence Similarity Searches." *Bioinformatics*, vol. 31, no. 6, Nov. 2014, pp. 926–32, doi:<https://doi.org/10.1093/bioinformatics/btu739>.

Sweet-Cordero, E. Alejandro, and Jaclyn A. Biegel. "The Genomic Landscape of Pediatric Cancers: Implications for Diagnosis and Treatment." *Science*, vol. 363, no. 6432, Mar. 2019, pp. 1170–75, doi:<https://doi.org/10.1126/science.aaw3535>.

The International Cancer Genome Consortium. "International Network of Cancer Genome Projects." *Nature*, vol. 464, no. 7291, Apr. 2010, pp. 993–998, doi:<https://doi.org/10.1038/nature08987>.

The UniProt Consortium. "UniProt: The Universal Protein Knowledgebase in 2023." *Nucleic Acids Research*, vol. 51, no. D1, Nov. 2022, pp. D523–531, doi:<https://doi.org/10.1093/nar/gkac1052>.

Trott, Oleg, and Arthur J. Olson. "AutoDock Vina: Improving the Speed and Accuracy of Docking with a New Scoring Function, Efficient Optimization, and Multithreading." *Journal of Computational Chemistry*, 2009, pp.455-461, doi:<https://doi.org/10.1002/jcc.21334>.

Venselaar, Hanka, et al. "Protein Structure Analysis of Mutations Causing Inheritable Diseases. An E-Science Approach with Life Scientist Friendly Interfaces." *BMC Bioinformatics*, vol. 11, no. 1, Nov. 2010, doi:<https://doi.org/10.1186/1471-2105-11-548>.

Ward, Zachary J., et al. "Estimating the Total Incidence of Global Childhood Cancer: A Simulation-Based Analysis." *The Lancet Oncology*, vol. 20, no. 4, Apr. 2019, pp. 483–93, doi:[https://doi.org/10.1016/s1470-2045\(18\)30909-4](https://doi.org/10.1016/s1470-2045(18)30909-4).

Weinstein, John N, et al. "The Cancer Genome Atlas Pan-Cancer Analysis Project." *Nature Genetics*, vol. 45, no. 10, 26 Sept. 2013, pp. 1113–1120, doi:<https://doi.org/10.1038/ng.2764>.

wwPDB Consortium. "Protein Data Bank: The Single Global Archive for 3D Macromolecular Structure Data." <Https://Academic.oup.com/Nar/Article/47/D1/D520/5144142>, vol. 47, no. D1, Oct. 2018, pp. D520–28, doi:<https://doi.org/10.1093/nar/gky949>.