PUNCH: An Interactive Web Server for Predicting Intrinsically Disordered Regions in Protein Sequences ☆

Di Meng * **and Gianluca Pollastri**

School of Computer Science, *University College Dublin, Ireland*

Correspondence to Di Meng:*School of Computer Science, University College Dublin, Ireland. *di.meng@ucdconnect.ie (D. Meng)*

https://doi.org/10.1016/j.jmb.2025.169018

Editor: Rita Casadio

Abstract

PUNCH is a freely accessible web server designed for the rapid and accurate prediction of intrinsically disordered regions (IDRs) in protein sequences. Built on a high-performance computational framework, PUNCH web server which built on PUNCH2-Light predictor, combines speed with predictive accuracy, offering users a streamlined interface for generating predictions from sequence input. Validated against the CAID2 benchmarking datasets, PUNCH web server demonstrates competitive performance in detecting IDRs across diverse protein sequences. Notably, it excels in the Disorder_PDB dataset and provides reliable results for the Disorder_NOX dataset, addressing the challenges of predicting disordered regions with low sequence similarity. The server is available at https://alienlabs.ucd.ie/punch2/, with extensive documentation and downloadable example datasets to support researchers in structural biology and bioinformatics.

Introduction

Proteins with intrinsically disordered regions (IDRs) lack stable three-dimensional structures, yet they play crucial roles in cellular processes such as signaling, regulation, and molecular recognition.[1,2] Despite their biological significance, predicting IDRs remains a challenging task due to their structural flexibility and dynamic nature. To address these challenges, our group recently developed two IDR prediction tools, PUNCH2 and its computationally efficient variant, PUNCH2-Light.[3] Unlike PUNCH2, which incorporates multiple sequence alignment (MSA) as part of the feature extraction process, PUNCH2-Light eliminates the need for MSA, significantly reducing computational demands while maintaining high prediction accuracy.

The Critical Assessment of Intrinsic Disorder prediction (CAID)[2,4] initiative provides a rigorous benchmarking framework for evaluating IDR prediction methods. By offering standardized datasets and evaluation metrics, CAID enables an objective comparison of prediction tools, driving progress in the field. Our predictors were primarily evaluated using datasets from the second round of CAID (CAID2) and were also participants in the recently conducted third round (CAID3),[5] which was held in conjunction with the CASP-16[6] conference. The CAID3 results, now publicly available, provide valuable insights into the evolving landscape of IDR predictors and highlight the increasing demand for methods that offer a balance between accuracy, computational efficiency, and usability. In CAID3, both PUNCH2 and PUNCH2-Light achieved exceptional performance, securing the top two ranks on

☆ This article is part of a special issue entitled: 'Computation Resources (2025)' published in Journal of Molecular Biology.

Please cite tts article as: D. Meng and G. Pollastri, PUNCH: An Interactive Web Server for Predicting Intrinsically Disordered Regions in Protein SequencesJournal of Molecular Biology, https://doi.org/10.1016/j.jmb.2025.169018

the *Disorder_PDB* benchmark and ranking 12th and 13th on the *Disorder_NOX* dataset. These results establish PUNCH2 and PUNCH2-Light as valuable resources for studying protein structure and function.

Building on these advancements, the PUNCH web server provides a freely accessible platform for IDR prediction, combining the computational efficiency of PUNCH2-Light with a user-friendly interface. Designed to accommodate users at all levels, the web server enables rapid and accurate predictions for both large-scale datasets and single-sequence analyses. This manuscript presents the design, implementation, and performance of the PUNCH web server and discusses its potential applications in IDR research. The PUNCH web server is freely accessible at https://alienlabs.ucd.ie/punch2/.

Materials and Methods

Overview of prediction algorithm

PUNCH2 and PUNCH2-Light leverage a lightweight yet powerful convolutional neural network (CNN) architecture,[7] specifically optimized for the identification of intrinsically disordered regions (IDRs) in protein sequences. The core model integrates multiple embedding techniques to effectively capture essential sequence features and biochemical properties. Initially, three primary embedding strategies were evaluated for their effectiveness in representing IDR-relevant sequence information: One-Hot Encoding, Multiple Sequence Alignment (MSA)-based embeddings, and Protein Language Models (PLMs). The evaluation results revealed that combining One-Hot Encoding with ProtTrans embeddings achieved comparable performance to utilizing all three methods. This finding indicated that the computationally expensive MSA-based embeddings provided minimal additional value when One-Hot and ProtTrans embeddings were already in use. As a result, the web server implementation, known as PUNCH2-Light,[3] incorporates only the One-Hot and ProtTrans embeddings, striking a balance between speed and memory efficiency.

Among these embeddings, ProtTrans embeddings provide deep contextualized representations that enhance PUNCH2-Light's ability to recognize disordered sequence patterns without relying on traditional MSAs. These embeddings are processed through a CNN framework featuring a "narrow but deep" architecture that efficiently captures IDR-associated patterns. This structure eliminates the need for MSA-based inputs, significantly reducing computational overhead, and allows the model to process variable-length protein sequences efficiently. The purely convolutional architecture also supports parallelized computations, further enhancing processing speed. While convolutional layers excel at capturing local sequence dependencies, the depth of the model allows it to effectively detect long-range dependencies. Despite its depth, PUNCH2-Light remains highly efficient, with approximately 200,000 parameters, offering a favorable balance of predictive accuracy and minimal resource consumption.

The PUNCH2-Light predictor is fully integrated into the PUNCH web server, ensuring both accuracy and computational efficiency for large-scale IDR analyses in an accessible web-based format. For the purposes of this manuscript, the term "PUNCH web server" refers to the online platform, while "PUNCH2-Light" denotes the underlying prediction model.

Data sources and benchmarking

The training and validation datasets for PUNCH2-Light were carefully curated from reputable sources, including the Protein Data Bank (PDB)[8] for structured sequences and DisProt[3] for fully disordered proteins. To construct the PDB dataset, sequences were retrieved using the query: [Structure Determination Methodology = "experimental" AND (Experimental Method = "X-RAY DIFFRACTION" AND Polymer Entity Type = "Protein")]. Missing residues within these sequences were labeled as disordered regions. After filtering out sequences with more than 30% similarity to the test set, a final dataset comprising 72,958 sequences, referred to as "PDB_missing," was compiled. These sequences represent proteins that are either partially disordered or fully structured. However, this dataset lacked fully disordered proteins necessary for comprehensive model training.

DisProt,[9] recognized as a comprehensive resource for disordered protein data, was selectively utilized to address the annotation gaps in partially disordered proteins. While the disordered regions annotated in DisProt are highly reliable, unannotated regions may correspond to either disordered regions lacking experimental evidence or structured regions. To ensure the inclusion of fully annotated disordered sequences, a subset of 158 fully disordered proteins (termed DisProt_FD) was extracted from DisProt. This dataset provided unique sequence characteristics essential for effective model training.[3]

To rigorously benchmark PUNCH2-Light, evaluations were conducted using the CAID2 challenge datasets,[4,5] which serve as standardized benchmarks for IDR predictors. The evaluation encompasses two subsets: **Disorder_NOX** and **Disorder_PDB**. Disorder_NOX classifies all available disordered annotations as positive and treats the remaining sequence regions as negative. In contrast, Disorder_PDB labels disordered regions as positive, structured regions (as annotated by the PDB) as negative, and disregards regions without sufficient annotation. Due to the inherent incompleteness in DisProt annotations,[9] Disorder_PDB

was chosen as the primary evaluation dataset, offering a more reliable classification of structured and disordered regions. However, to account for potential annotation inconsistencies, including possible false-negative classifications in Disorder_PDB, additional evaluations were conducted on Disorder_NOX to provide a comprehensive performance assessment across both scenarios.

To ensure robust evaluation, the training set was rigorously filtered to exclude sequences sharing more than 30% similarity with those in both Disorder_PDB and Disorder_NOX. This filtering process allowed us to retain the full benchmarking datasets, which include 348 sequences for Disorder_PDB and 210 sequences for Disorder_NOX. By eliminating potential overlaps, this approach minimizes bias and provides an accurate assessment of the model's generalization capability on previously unseen data.

Performance

PUNCH2-Light was benchmarked against the top 10 IDR predictors (on Disorder_PDB dataset) from CAID2, demonstrating competitive accuracy and rapid processing speeds. The model excelled in detecting subtle IDR patterns across diverse protein sequences. Table 1 presents a comparative performance summary of PUNCH2-Light against leading predictors on both the Disorder_PDB and Disorder_NOX datasets, emphasizing its robustness in both accuracy and efficiency. On the Disorder_PDB dataset, PUNCH2-Light achieved an AUC of 0.950, an APS of 0.912, and a Max F1 of 0.845, performing comparably to the top-performing predictor, SPOT-Disorder2.[10] In the Disorder_NOX dataset, PUNCH2-Light attained an AUC of 0.779, an APS of 0.374, and a Max F1 of 0.516, remaining competitive with the top predictors on Disorder_PDB. The best-performing predictor on the CAID2 Disorder_NOX dataset, Dispredict3,[11] achieved an AUC of 0.838, an APS of 0.560, and a Max F1 of 0.548, with corresponding performance on Disorder_PDB of 0.895, 0.777, and 0.731, respectively, as shown in Table 1 (row 1a). This performance discrepancy highlights the annotation differences between the two datasets and underscores the challenges posed by varying annotation standards.

To assess the statistical significance of PUNCH2-Light's performance across the Disorder_PDB and Disorder_NOX datasets, we performed a bootstrap resampling analysis with 1000 iterations to estimate the variability of AUC, APS, and Max F1 scores. The results demonstrate that PUNCH2-Light achieves consistent performance within each dataset, particularly for Disorder_PDB, with mean AUC, APS, and Max F1 values of 0.94 8 ± 0.003, 0.910 ± 0.009, and 0.842 ± 0.011, respectively. For Disorder_NOX, the corresponding values were 0.781 ± 0.012, 0.376 ± 0.023, and 0.518 ± 0.022, indicating greater variability. A statistically significant difference was observed between the two datasets, with PUNCH2-Light showing significantly different performance levels, as supported by low *p*-values (≈**0.001**) obtained from paired *t*-tests and 95% confidence intervals. These findings suggest that while PUNCH2-Light delivers stable predictions within each dataset, its performance is influenced by the differing annotation standards between Disorder_PDB and Disorder_NOX.

Moreover, the recently released CAID3 results[5] demonstrate that PUNCH2-Light achieved an AUC of 0.950 (0.830), an APS of 0.925 (0.563), and a Max F1 of 0.862 (0.646) on the Disorder_PDB (Disorder_NOX) dataset. Notably, PUNCH2-Light ranked as the second-best predictor on Disorder_PDB, surpassed only by our PUNCH2 predictor, which obtained an AUC of 0.956 (0.832), an APS of 0.927 (0.573), and a Max F1 of 0.865 (0.646). On the Disorder_NOX dataset, PUNCH2-Light ranked 13th. These results highlight PUNCH2-Light's ability to deliver reliable predictions across diverse evaluation criteria, further

Table 1 Benchmark performance comparison of PUNCH2-Light with the Top 10 CAID2 Predictors on *Disprot_PDB* (*Disprot_NOX*).

Top	*Predictor*	*AUC* *Disorder_PDB* (*Disorder_NOX*)	*APS* *Disorder_PDB* (*Disorder_NOX*)	*Max F1* *Disorder_PDB* (*Disorder_NOX*)
1	*SPOT-Disorder2*	*0.949 (0.780)*	*0.928 (0.558)*	*0.860 (0.633)*
2	*AlphaFold-rsa*	*0.944 (0.747)*	*0.916 (0.407)*	*0.849 (0.525)*
3	*PredIDR-long*	*0.934 (0.741)*	*0.870 (0.341)*	*0.800 (0.464)*
4	*IDP-Funsion*	*0.933 (0.818)*	*0.878 (0.475)*	*0.822 (0.539)*
5	*SPOT-Disorder*	*0.931 (0.772)*	*0.889 (0.364)*	*0.824 (0.514)*
6	*SETH-0*	*0.930 (0.772)*	*0.893 (0.374)*	*0.830 (0.516)*
7	*AlphaFold-pLDDT*	*0.929 (0.695)*	*0.881 (0.335)*	*0.821 (0.506)*
8	*PredIDR-short*	*0.927 (0.737)*	*0.859 (0.340)*	*0.790 (0.459)*
9	*metapredict*	*0.923 (0.756)*	*0.834 (0.321)*	*0.819 (0.516)*
10	*DeepIDR-2L*	*0.922 (0.800)*	*0.858 (0.460)*	*0.794 (0.513)*
1a	*Dispredict3*	*0.895 (0.838)*	*0.777 (0.560)*	*0.731 (0.548)*
	PUNCH2-light	***0.950 (0.779)***	***0.912 (0.374)***	***0.845 (0.516)***

underscoring its potential as a robust tool for IDR prediction.

The numbers from *1* to *10* represent the performance rankings of various predictors on the CAID2[4] challenge *Disorder_PDB* dataset. *1a* represents the best predictor on the CAID2 challenge *Disprot_NOX* dataset. The last predictor, PUNCH2-Light, is our newly developed predictor.

Server Description

The framework of the website

PUNCH web server is an intuitive web server developed to predict intrinsically disordered regions (IDRs) in protein sequences, designed for ease of use and efficiency. Targeted at both bioinformatics and structural biology researchers, it provides a straightforward, interactive experience for a wide range of users.

Web pages

The website framework includes four primary web pages to facilitate user interaction and access to PUNCH web server's functionalities (web pages section, left part in Figure 1):

1. **Home Page**: This introductory page provides multiple ways for users to submit a query in FASTA format, check task status, and view basic information about PUNCH web server. It also includes a link to the Download page for users interested in local installation options.
2. **Task Result Page**: This page displays task information, including a unique task ID, and shows whether a submitted task is still in progress or completed. Once completed, the user can download the prediction results. The page also presents a table with the list of query sequences; each sequence has an individual link that leads to a detailed view of the prediction results for that sequence (prediction result page).
3. **Prediction Result Page**: Displaying results for individual protein sequences, this page includes a colour-coded visualization, with disordered amino acids highlighted in pink. A prediction plot provides a per-amino-acid breakdown of predicted values, and an interactive threshold slider allows users to adjust the threshold dynamically, with the plot and highlights updating in real-time.
4. **Download Page**: This page contains all information for users who wish to install PUNCH locally, detailing installation instructions for two different methods. It links to the software on GitHub and Docker for easy access.

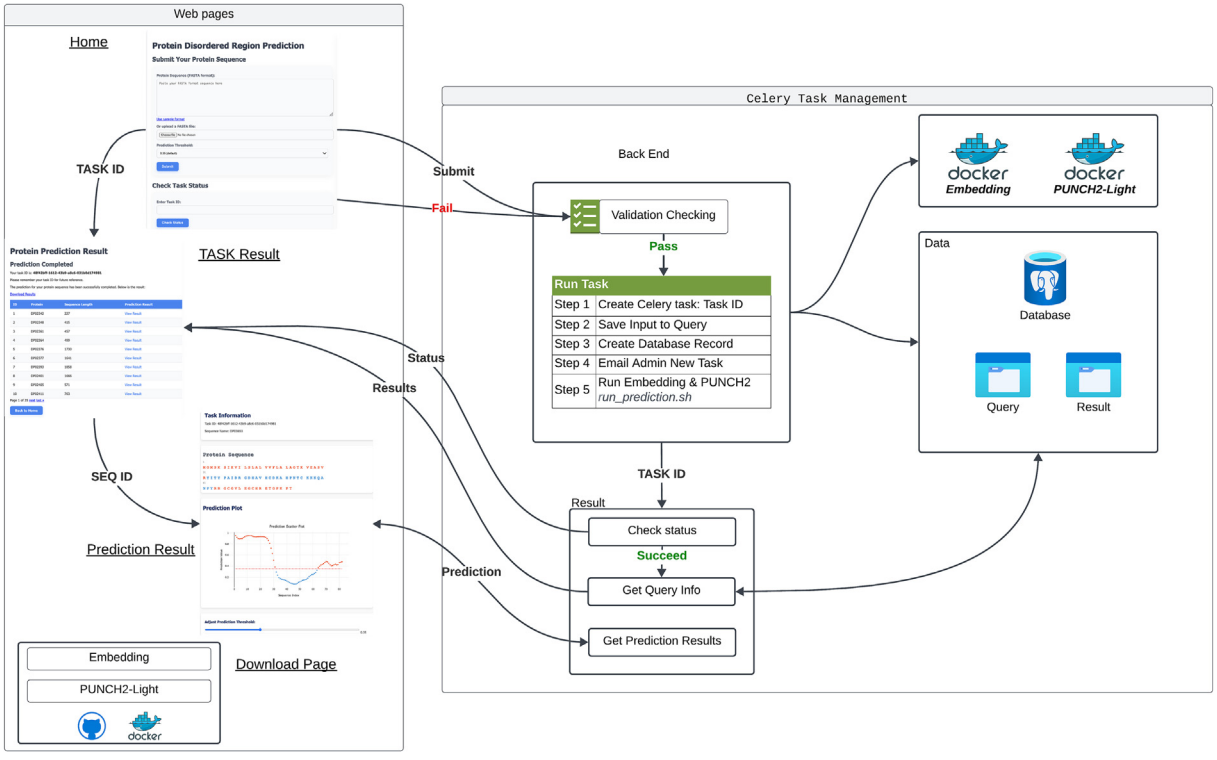

Figure 1. Framework of the PUNCH Web Server. The diagram illustrates the website structure, with the left section showing the primary web pages: (1) Home page for task submission, status checks, basic information, and download links; (2) Task Result page for viewing task details and sequence-specific prediction results with download options; (3) Prediction Result page displaying interactive predictions, including a threshold-adjustable plot; and (4) Download page with installation instructions for local use. The right section outlines the back-end framework, encompassing page interactions, Celery task management, data handling, and prediction task execution.

This structured, user-friendly setup ensures that users can easily navigate through submissions, review results, and explore download options, making PUNCH web server accessible for users of various expertise levels.

Back-end

The PUNCH web server is built on the Django[12] framework, with Python handling back-end processes and HTML and Javascript supporting the front end. Each user request is managed as a "Task," handled asynchronously by integrating Celery[13] to process predictions independently of the main application thread. This structure enables users to check the status of tasks later using their task IDs, without needing to keep the page open. Prediction results are stored for a week, allowing users ample time to retrieve their data.

Since the PUNCH2-Light predictor does not require MSA results, and both One-Hot encoding and ProtTrans embeddings operate with minimal resource demands, the computational load remains low for each sequence batch (~50 sequences). As tasks queue in Celery, the server processes one task at a time following a first-come, first-served rule, ensuring an efficient workflow without resource strain. This setup minimizes wait times by ensuring that each prediction completes quickly.

To manage data effectively, PUNCH web server uses a hybrid storage solution (Figure 1, right side):

- **Relational Database** (PostgreSQL)[14]: Stores essential metadata for each Task, such as the task ID, creation timestamp, file paths to input and output files, and other relevant task information. This structure minimizes database load by storing only critical metadata, improving server performance.
- **File System**: Large files, including the actual sequence data and prediction results, are saved on the server's disk. This approach avoids the need to store extensive data within the database itself, reducing database size and complexity while ensuring that users can download large results files directly from the file system.

Together, this structure optimizes the website's ability to handle concurrent users and sizable query data while keeping operational costs low.

Implementation and server runtime

PUNCH web server is designed with a lightweight, high-performance architecture that efficiently handles high-traffic loads. Asynchronous task management is handled by Celery, ensuring smooth processing of user requests and predictions even during peak usage. The server infrastructure operates on a Linux system with 24 CPU cores, 16 GB of RAM, an NVIDIA GeForce RTX 4080 GPU (16 GB memory), and 1 TB of storage. Using the Disorder_PDB dataset and applying the same method as the CAID2 benchmarking for runtime calculation—measuring the average prediction time for a sequence of 1,000 residues—PUNCH web server achieves an average runtime of 6 s per sequence, completing predictions for all 348 sequences in *Disorder_PDB* in approximately 1,660 s. Furthermore, when tested on the UniProt Reviewed Swiss-Prot[15] human proteome dataset (20,421 sequences, totaling 13.7 MB), PUNCH web server generated predictions for all sequences in roughly 1,600 s. With support for input files up to 40 MB, PUNCH web server is well-suited for large datasets and high-throughput applications.

User guide and an example

As illustrated in Figure 1, users can submit a query from the homepage by entering one or more protein sequences in the text area or uploading a FASTA file (.fasta). Both submission methods require the input to be in FASTA format. For example, submitting a file named p49913.fasta, containing one Uniprot sequence (P49913), will initiate the task. The prediction threshold can be adjusted, with a default value set at 0.35. Once a query is submitted, a task is created immediately, and the user is directed to the result page, where the task ID (e.g., bb880a50-6955-4d17-aafe-85b7bfaa8644) is displayed. Users may refresh the result page or save the task ID to check the task status from the homepage at a later time.

Upon completion, a download link appears, allowing users to obtain the prediction results as CSV files, with each sequence saved in a separate file. Additionally, users can view the prediction results directly on the website: the result page lists all submitted sequences in a table, with links for each sequence leading to a detailed prediction result page.

In the Prediction Result page (illustrated in Figure 2 for sequence P49913), detailed predictions are displayed as both colour-coded amino acids and a scatter plot. Disordered amino acids are highlighted in pink, while structured amino acids appear in blue. Users can adjust the threshold slider, and the scatter plot and amino acid colours will dynamically update to reflect the changes.

Results for each task are available to download and review for one week, after which task information is automatically deleted.

Conclusion

PUNCH web server offers a powerful and accessible platform for predicting intrinsically disordered regions (IDRs) in protein sequences. Through its efficient, high-performance

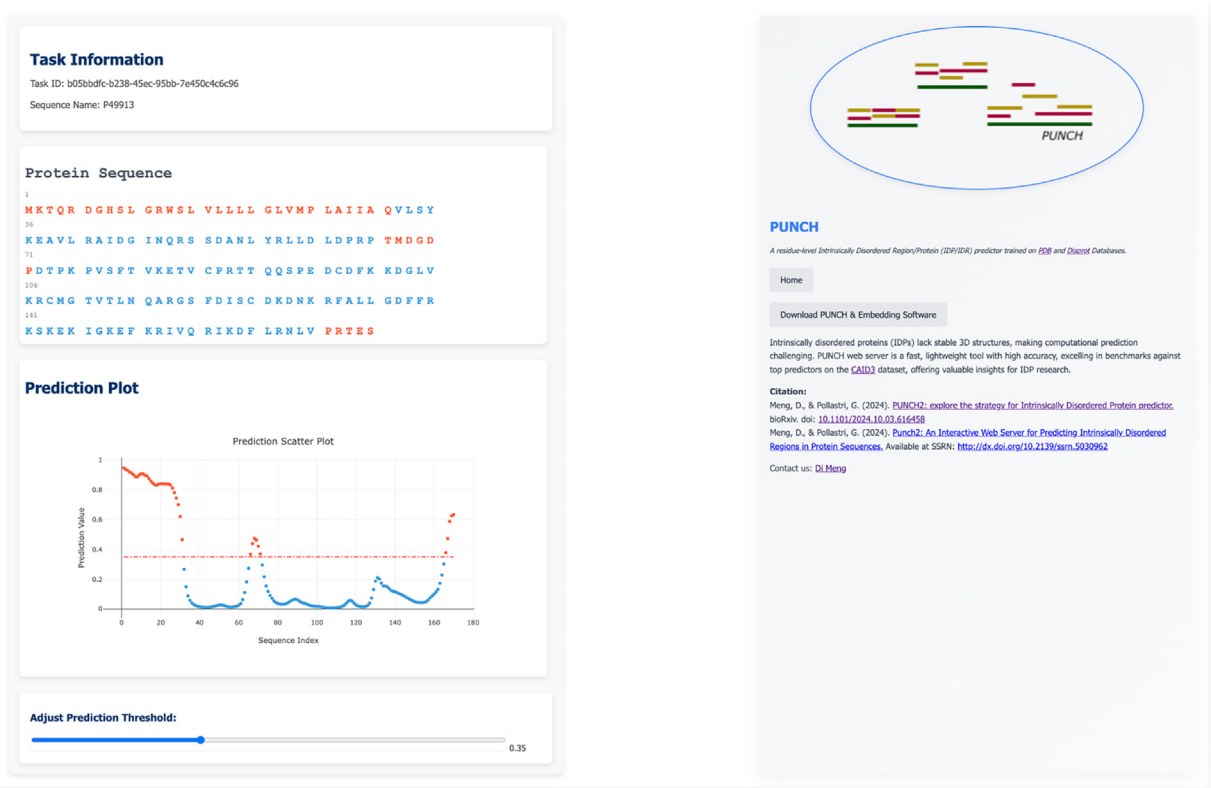

Figure 2. Prediction Result page. This screenshot shows the prediction result of Uniprot sequence P49913.

architecture, the server allows users to generate accurate predictions quickly, supporting both high-throughput analysis and detailed individual investigations. PUNCH2-Light, the core predictor, combines One-Hot encoding and ProtTrans embeddings in a deep convolutional neural network, effectively capturing IDR-related features without the computational cost of traditional alignment-based methods. Comprehensive benchmarking against CAID2 standards confirms PUNCH web server's competitive accuracy and processing speed, making it a valuable resource for the bioinformatics community.

The server's user-friendly interface includes well-structured pages for query submission, task status monitoring, and result visualization, with dynamic elements that allow users to adjust prediction thresholds in real time. Backed by a robust data management system, PUNCH web server securely handles large datasets and stores results temporarily, providing users ample time to retrieve and analyze their data. With freely available access, clear documentation, and downloadable resources, PUNCH stands as a versatile tool for researchers across bioinformatics and structural biology. The server's architecture and adaptability position it as a sustainable solution for IDR prediction, supporting ongoing research efforts and future advancements in the study of protein disorders.

Availability

Website: https://alienlabs.ucd.ie/punch2/.
Dataset: https://huggingface.co/datasets/deeeeeeeeee/PUNCH2_data.
GitHub:

- Embedding: https://github.com/deemeng/embedding
- PUNCH2-Light: https://github.com/deemeng/punch2_light

Docker:

- Embedding: https://hub.docker.com/r/dimeng851/embedding
- PUNCH2-Light: https://hub.docker.com/r/dimeng851/punch2_light

CRediT authorship contribution statement

Di Meng: Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Formal analysis, Data curation, Conceptualization.

Gianluca Pollastri: Writing – review & editing, Supervision, Project administration.

DATA AVAILABILITY

The data is from open access databases and the description is in the paper.The code is shared through a Github link.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was funded by Science Foundation Ireland (SFI) under awards #1636933 and #1920920. A big thanks to Andrew Fulcher for his creative input on the logo design.

Received 22 November 2024;
Accepted 16 February 2025;
Available online xxxx

Keywords:
intrinsically disordered regions;
protein prediction;
web server;
convolutional neural network;
bioinformatics

References

1. van der Lee, R., Buljan, M., Lang, B., Weatheritt, R.J., Daughdrill, G.W., Dunker, A.K., Fuxreiter, M., Gough, J., Gsponer, J., Jones, D.T., Kim, P.M., Kriwacki, R.W., Oldfield, C.J., Pappu, R.V., Tompa, P., Uversky, V.N., Wright, P.E., Babu, M.M., (2014). Classification of intrinsically disordered regions and proteins. *Chem. Rev.* **114** (13), 6589–6631. https://doi.org/10.1021/cr400525m. Epub 2014 Apr 29. PMID: 24773235; PMCID: PMC4095912.
2. Necci, M., Piovesan, D., CAID Predictors, et al., (2021). Critical assessment of protein intrinsic disorder prediction. *Nature Methods* **18**, 472–481. https://doi.org/10.1038/s41592-021-01117-3.
3. Meng, D., Pollastri, G., (2024). PUNCH2: explore the strategy for Intrinsically Disordered Protein predictor. *bioRxiv*10.03.616458. https://doi.org/10.1101/2024.10.03.616458.
4. Conte, A.D., Mehdiabadi, M., Bouhraoua, A., Monzon, A. M., Tosatto, S.C.E., Piovesan, D., (2023). Critical assessment of protein intrinsic disorder prediction (CAID) – results of round 2. *Proteins* **91** (12), 1925–1934. https://doi.org/10.1002/prot.26582.
5. CAID challenge results. Critical assessment of intrinsic disorder (CAID). Available from: https://caid.idpcentral.org/challenge/results. Accessed 19th December 2024.
6. CASP16 conference. Critical assessment of protein structure prediction. Available online: https://predictioncenter.org/casp16/. Accessed 27th December 2024.
7. LeCun, Y., Bengio, Y., Hinton, G., (2015). Deep learning. *Nature* **521** (7553), 436–444. https://doi.org/10.1038/nature14539. PMID: 26017442.
8. Berman, H.M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., Bourne, P.E., (2000). The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242. https://doi.org/10.1093/nar/28.1.235.
9. Aspromonte, M.C., Nugnes, M.V., Quaglia, F., Bouharoua, A., DisProt Consortium, Tosatto, S.C.E., Piovesan, D., (2023). DisProt in 2024: improving function annotation of intrinsically disordered proteins. *Nucleic Acids Res., Database Issue*. https://doi.org/10.1093/nar/gkad928. PubMed:37904585.
10. Hanson, J., Paliwal, K.K., Litfin, T., Zhou, Y., (2019). SPOT-Disorder2: improved protein intrinsic disorder prediction by ensembled deep learning. *Genomics Proteomics Bioinf.* **17** (6), 645–656. https://doi.org/10.1016/j.gpb.2019.01.004.
11. Kabir, Md Wasi Ul, Hoque, Md Tamjidul, (2024). DisPredict3.0: Prediction of intrinsically disordered regions/proteins using protein language model. *Appl. Math. Comput.* **472(C)**
12. Django Software Foundation. Django: High-level Python web framework (Version 5.1.1) [Software]. Available from https://www.djangoproject.com/. Solem, A., Celery Project contributors.
13. Celery: Distributed task queue (Version 5.4.0) [Software]. Available from https://docs.celeryproject.org/. PostgreSQL Global Development Group.
14. PostgreSQL: Open source relational database (Version 14.13) [Software]. Available from https://www.postgresql.org/.
15. The UniProt Consortium, (2023). UniProt: the universal protein knowledgebase in 2023. *Nucleic Acids Res.* **51** (D1), D523–D531. https://doi.org/10.1093/nar/gkac1052.