The structural coverage of the human proteome before and after AlphaFold

Eduard Porta-Pardo[1,2,], Victoria Ruiz-Serra[1,2] and Alfonso Valencia[1,3,*]*

1 - Barcelona Supercomputing Center (BSC)
2 - Josep Carreras Leukaemia Research Institute (IJC), Badalona, Spain
3 - Institució Catalana de Recerca Avançada (ICREA)
* To whom correspondence should be addressed: eporta@carrerasresearch.org / alfonso.valencia@bsc.es

Abstract

The protein structure field is experiencing a revolution. From the increased throughput of techniques to determine experimental structures, to developments such as cryo-EM that allow us to find the structures of large protein complexes or, more recently, the development of artificial intelligence tools, such as AlphaFold, that can predict with high accuracy the folding of proteins for which the availability of homology templates is limited.

Here we quantify the effect of the recently released AlphaFold database of protein structural models in our knowledge on human proteins. Our results indicate that our current baseline for structural coverage of 47%, considering experimentally-derived or template-based homology models, elevates up to 75% when including AlphaFold's predictions, reducing the fraction of "dark proteome" from 22% to just 7% and the number of proteins without structural information from 4.832 to just 29.

Furthermore, although the coverage of disease-associated genes and mutations was near complete before AlphaFold release (70% of ClinVar's pathogenic mutations and 74% of oncogenic mutations), AlphaFold models still provide an additional coverage of 2% to 14% of these critically important sets of biomedical genes and mutations. We also provide several examples of disease-associated proteins where AlphaFold provides critical new insights. Overall, our results show that the sequence-structure gap of human proteins has almost disappeared, an outstanding success of direct consequences for the knowledge on the human genome and the derived medical applications.

Introduction

Ever since the first protein structure was published in 1958[1] it was clear that structure information is essential to understand the biological functions of proteins. In that sense, the last 63 years have witnessed an outstanding process thanks to several milestones. First and foremost, the Protein Data Bank[2] has emerged, since its creation in 1971, as the key biological database to organize and standardize experimental protein structural data. Also, several groups in the late 1980s and early 1990s observed that the protein structure was much more conserved than its sequence[3], which led to the creation of the first computational tools to predict protein folding[4–8]. These tools are extremely important, as the growth of protein sequence data has far outpaced that of experimentally-determined protein structures. In order to systematically assess the performance of all these tools and monitor the progress of the protein folding prediction field, the Critical Assessment of protein Structure Prediction (CASP) experiments were established in 1994[9]. These experiments have held a high-standard in the field and in recent years have witnessed the massive progress that protein structure prediction has made thanks to, among others, the use of artificial intelligence approaches[10].

Protein structures have a wide-array of applications in biotechnology and biomedicine. Understanding the structure of a protein can help in identifying which mutations in enzymes can make them acquire new properties that make them interesting from a biotechnological or environmental perspective[11]. Also, in the case of biomedicine, they can be used to perform *in silico* experiments to compute the affinity of tens of thousands of small molecules or antibodies for proteins associated with different diseases[12]. Further biomedical application of protein structures involves the study of the consequences of genetic variants, both acquired[13] or inherited[14]. For instance, computational tools based on protein structures are often more accurate [15,16] than those based on linear features[17] when prioritizing relevant protein-coding variations in patients[18], a much needed strategy when experimental characterization of millions of mutations is not realistic.

Over the last two editions of CASP, the computational tool AlphaFold[19], developed by DeepMind, has achieved a performance that suggests that the prediction of individual protein domains has virtually been solved[20]. Moreover, in conjunction with the EMBL-EBI, the team behind AlphaFold has recently released a database of predicted protein structures for the whole human proteome[21]. These two outstanding accomplishments, together, have generated a great excitement in the scientific community in general and amongst structural biologists in particular, as it opens a wide-array of research possibilities.

Given the profound transformations that structural biology has recently experienced, the importance of protein structures in biomedicine, and the potential transformation that the field might experience in light of the release of AlphaFold's code and models for human proteins, it is important to quantify the contribution of these models in general and of high-quality models in particular. The latter are particularly relevant as they might be able to provide insights into the biochemistry of proteins that had not yet been amenable to experimental structure determination or the effects of mutations in such proteins.

Here we quantified the consequences of the addition of the AF models to the accumulated knowledge on protein structures and models. Our results show that AlphaFold predictions

have significantly increased our coverage for the entire human proteome, particularly in proteins integral to membrane, enzymes involved in lipid metabolism and olfactory receptors. The direct contribution of AlphaFold's protein models to proteins with biomedical interest is more limited, as these proteins were already better characterised. Overall, AlphaFold increases the structural coverage from 47% to 75% of all human protein residues and, importantly, the coverage with high-quality models (sequence identity with a PDB chain is ≥ 50%) from 31% to 50%. Moreover, the number of proteins with no structural coverage at all has dramatically decreased from 4.832 proteins to just 29. Together, the combined data from experimental structures, template-based homology modelling with high sequence identity and high-accuracy models from AlphaFold will open a new era in structural biology in particular and human biomedicine in general.

Results

The structural coverage of the human proteome before AlphaFold

Over the last 25 years, the protein structure field has made significant advances. This is evidenced by the fact that the Protein Data Bank (PDB)[2], the main database for protein structure coordinates, in 1995 had 4.455 protein coordinate files, whereas by 2020 the number had increased to 177.806 (**Figure 1a**). The relative contribution of human proteins (sequence identity of human proteins with PDB chains is >95%) has also increased over these 25 years, now making 27% of all PDB structures, compared to only 17% in 1995 (**Figure 1b**). These advances, importantly, have translated into a much higher coverage of the human proteome. We quantified it by aligning all the principal protein isoforms of the human protein coding genes from ENSEMBL[22] according to APPRIS[23] (ENSEMBL v103, n = 20.312) against all the PDB chains (n = 596.542) using BLAST[24]. We considered only hits with an e-value below 1e-7 and sequence identity ≥ 20%, the limit thresholds for template-based homology modelling[3]. With this definition our results show that the current structural coverage of the human proteome is 47% of all human protein residues (**Figure 1c**). This is significant progress compared to 1995, when the coverage was 7 times less (6.3%). If we are more stringent and consider only those regions of the proteome with a sequence identity >95% with a PDB protein chain, then the current coverage decreases to 16%, but even so this is 33 times higher than in 1995 (0.47%).

Not all the human proteome has a structure (such as in the case of intrinsically disordered proteins[25]) or its structure has been more difficult to experimentally determine (as in the case of membrane proteins). To further explore this issue, we predicted the presence of two types of linear regions: PFAM domains[26], which are evolutionary conserved functional regions, and intrinsically disordered regions (IDRs), which are protein regions predicted not to have any tertiary structure in their native state. Moreover, by exclusion, we also determined whether a protein region was neither structured or predicted to be disordered, the so-called dark proteome[27] (**Figure 1d**). Results show that for 13% of all residues of the human proteome there is a PFAM domain assigned but no template available to model a structure, 18% of all residues are predicted to be disordered and we have no features at all for the additional 22% of all residues.

In summary, in terms of three-dimensional coverage, 17% of all human protein residues currently have experimental structures (sequence identity with a PDB chain is >95%), 14% can be modelled with good accuracy (sequence identity between 95% and 50%) and 16% can be modelled with reasonable accuracy (sequence identity between 50% and 20%). As for the rest of the proteome, without considering IDRs, 35% of residues are missing structural coverage even though 13% of them belong to a known PFAM domain (**Figure 1d**).

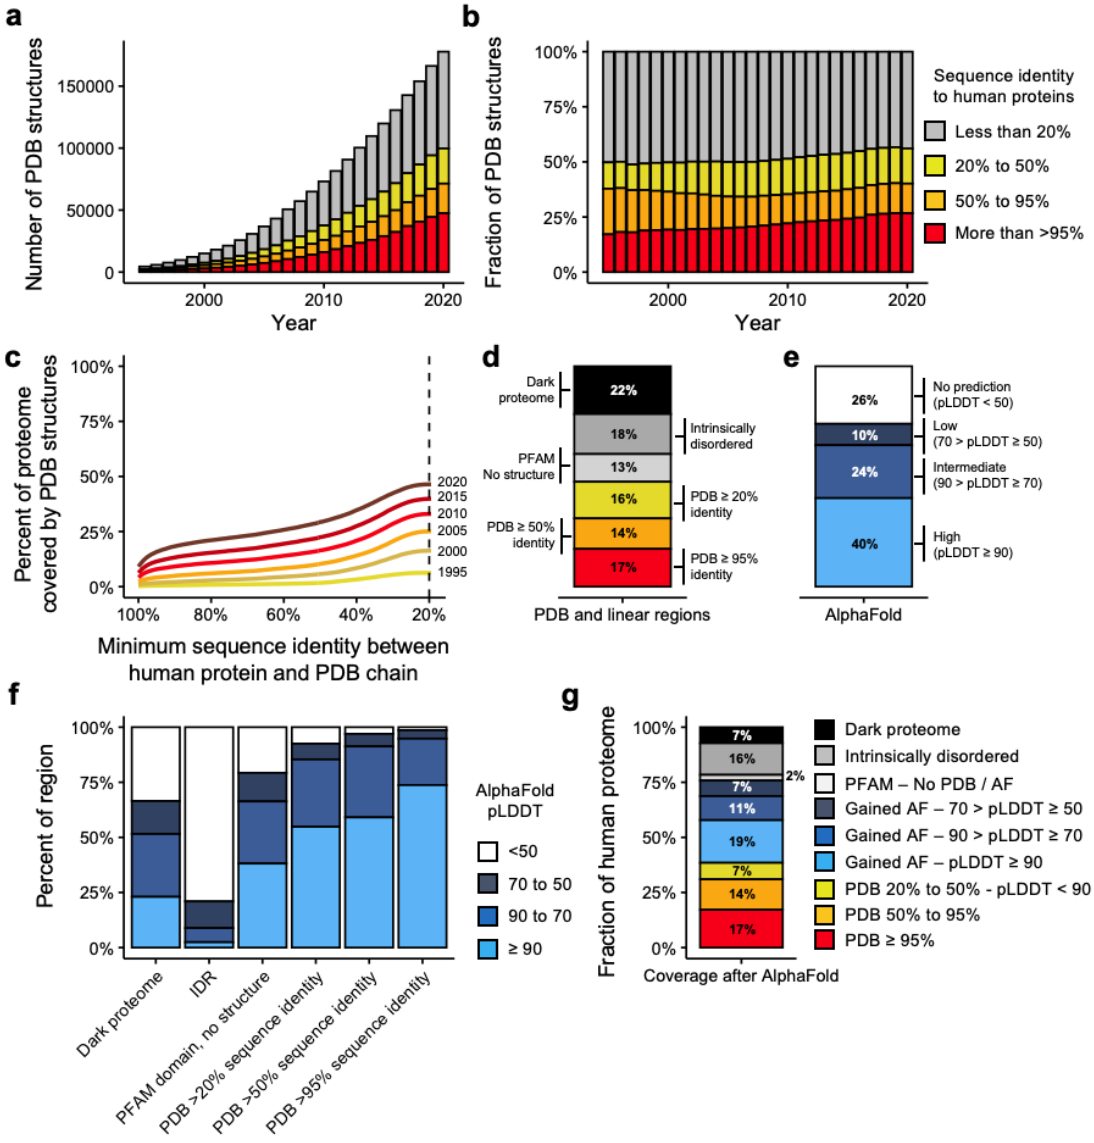

Figure 1 - Current coverage of the human proteome. a-b) Barplot showing the absolute (**a**) or relative (**b**) number of PDB coordinate files mapping to human proteomes at different thresholds of sequence identity. Legends in barplots **a** and **b** are the same. **c)** Evolution of the coverage of the human proteome by three-dimensional coordinate files in the Protein Data Bank (y-axis) according to the minimum percent identity of the BLAST hits (x-axis). Each line represents the coverage using only the coordinate files available in PDB in a given year. **d)** Barplot showing the coverage of the human proteome by different types of structural features, both linear (PFAM domains and IDRs) and three-dimensional (PDB) (y-axis is the same as in **c**). **e)** Coverage of the proteome by different

AlphaFold pLDDT score thresholds (y-axis is the same as in **c**). **f)** Coverage (y-axis) of different types of regions (x-axis) depending on AlphaFold confidence levels. **g)** Current coverage (y-axis) of the human proteome.

The structural coverage of the human proteome after AlphaFold

We next analyzed the structural predictions of AlphaFold for the human proteome[21]. AlphaFold has, along with the actual three-dimensional coordinates, a measure of how confident one can be about the prediction itself: predicted local distance difference test (pLDDT). According to its developers, a pLDDT above 90 is of high confidence, between 90 and 70 of intermediate confidence, between 70 and 50 of low confidence, and the coordinates of any residue with a pLDDT below 50 should be disregarded[21]. At the residue level, AlphaFold models have high confidence predictions for 40% of the human proteome, intermediate-confidence predictions for 24%, low-confidence predictions for 10% and no prediction (i.e. pLDDT < 50) for 26% (**Figure 1e**).

Then, we calculated the overlap between the distribution of linear and template-based homology regions (**Figure 1d**) with the AlphaFold prediction confidence scores (**Figure 1e**). As expected, there was a correspondence between the sequence identity of template proteins with target PDB chains and AlphaFold's pLDDT score: 74% of all residues with over 95% sequence identity with a PDB chain had a pLDDT above 90, while the percent decreased to 55% when the sequence identity of all residues ranges between 20% and 50% (**Figure 1f**). Consistently, we observe that 75% of all residues that were predicted to be in IDRs (i.e. disordered) had a pLDDT below 50. Overall these results show consistency between the foldability of a protein region and AlphaFold's pLDDT score.

Regarding the portion of the human proteome that lacks structure coverage, i.e. a number of PFAM domains without known experimental structure or template to build homology models and the dark proteome, AlphaFold makes a significant contribution. High-confidence structures (pLDDT >90) are predicted for 38% and 23% of the residues corresponding to PFAM domains without structure or the dark proteome respectively (**Figure 1f**). Still, the 21% and 33% of the residues of PFAM domains without structure and the dark proteome respectively are associated with a pLDDT < 50 (**Figure 1f**), which could be indicative of truly disordered regions.

The integration of our previous structural knowledge through experimental structures, template-based homology modelling, PFAM domains, IDRs together with AlphaFold predictions, shows that the current overall coverage of the human proteome at the residue levels is excellent (**Figure 1g**). For the rest of the paper, we consider a region as having a highly accurate protein structure model or prediction if it either has a sequence identity above 50% with a PDB chain or it has an AlphaFold pLDDT score above 90.

With this definition, excellent structures or models are available for 50% of the residues, 31% coming from PDB and 19% from AlphaFold. An additional 25% of the human proteome can be covered with other structural models, either through template based modelling of PDB chains with a sequence identity with the template between 20% and 50% (7%) or through AlphaFold models with intermediate (11%) or low confidence (7%). Approximately 16% of the proteome is predicted as disordered and only 2% of all residues have a PFAM domain

without structure nor model. Finally, the fraction of the proteome without any annotation has decreased to only 7%, albeit given the high correlation between regions with pLDDT below 50 and intrinsic disorder, these are likely to also be IDRs, which would put the fraction of disordered proteome to 23%.

Changes in coverage at the protein level

In this section, we analyse how AlphaFold complements our knowledge at the protein level. When comparing the number of human proteins without structural coverage before and after considering AlphaFold predictions, we observe that the number drops from 4.832 to just 29 proteins (Figure 2a). This means that we currently have either experimental or predicted structural data for some region of 99,8% of all the human proteins we considered (17.006 proteins, **Methods**).

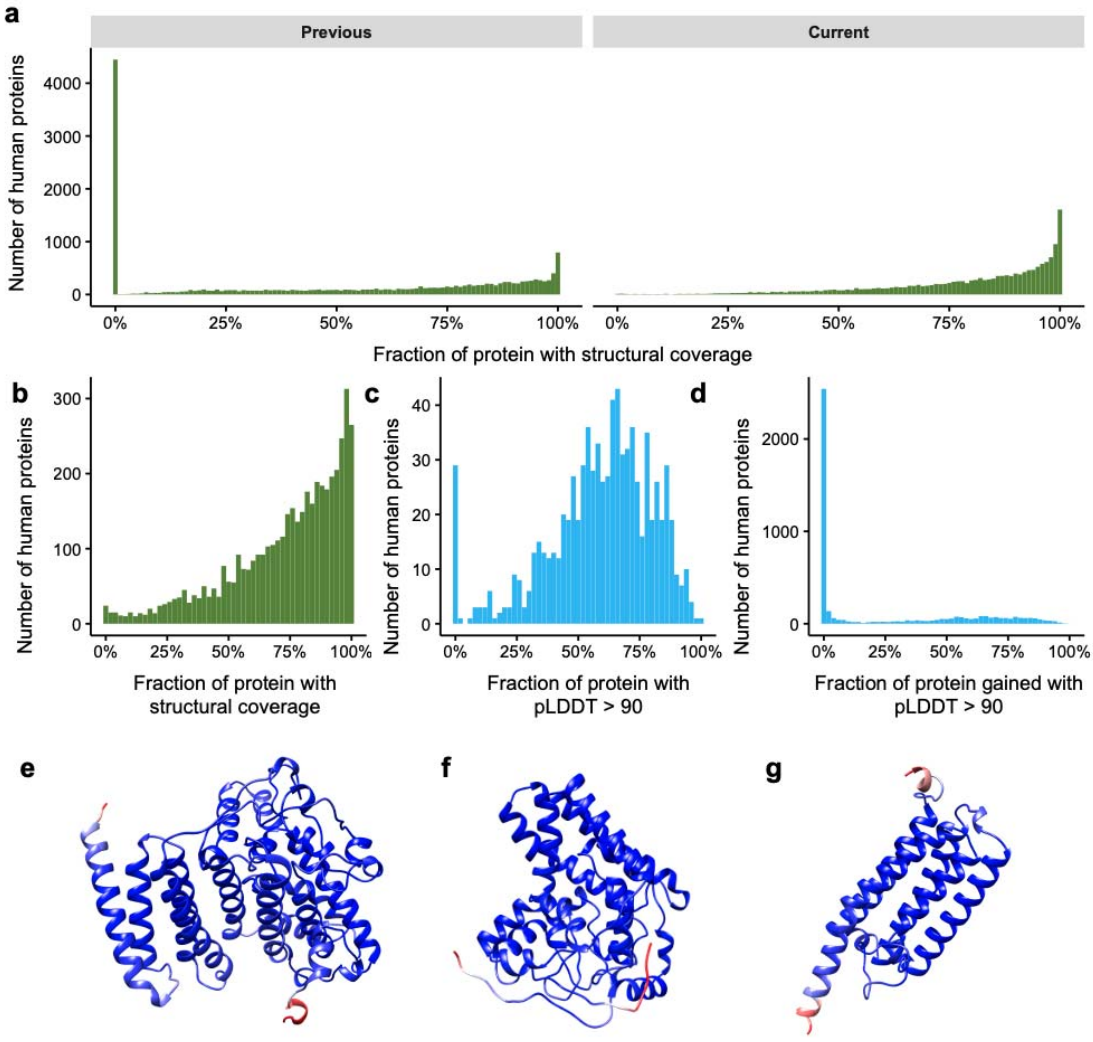

Figure 2 - Changes in the structural coverage at the protein level after AlphaFold. a) Histogram showing the number of proteins (y-axis) according to their structural coverage (x-axis) before (left) and after (right) the release of AlphaFold models. **b)** Histogram showing the number of proteins for

which we previously had less than 1% of structural coverage (y-axis) according to their current structural coverage after AlphaFold. **c)** Same as **b** but now including only high-confidence (pLDDT > 90) AlphaFold predictions (x-axis). **d)** Histogram showing how much AlphaFold high-confidence predictions contribute (x-axis) to our coverage of proteins with >95% structural coverage. **e-g)** AlphaFold models for AGMO, DEGS1 and PEMT. Models are colored in blue-red scale showing the pLDDT score for the residue, with red representing low pLDDT and blue high pLDDT.

We studied in more detail the proteins for which we previously had no structural data, i.e. the aforementioned 4.832 proteins. A GO term enrichment analysis[28] on these proteins reveals a strong enrichment in olfactory receptors (269 genes, OR > 2,5, p < 1e-30), lipid metabolism (66 genes, OR > 2, p < 1e-10), genes integral to membrane (1700 genes, OR 1,5, p < 1e-50) and genes without any previous GO annotation (1056 genes, OR > 1.8, p < 1e-70). Notably, AlphaFold models provide structural coverage for over 50% of the protein's length for 4.063 of these 4.832 proteins (**Figure 2b**). While this number boils down to 1.316 proteins if we just refer to high-quality predictions (pLDDT > 90 covering more than half the protein's length, **Figure 2c**), this is still a significant improvement considering that before AlphaFold we had no structural data for these proteins.

Lastly, we quantified how much AlphaFold contributes to the proteins for which our current structural coverage is above 95% of the protein length (n = 5.162, **Figure 2d**). This revealed that, while in most cases we already had a significant amount of structural data, either experimental or through template-based homology, there are 1.413 proteins with very high structural coverage where AlphaFold contributes over 50% of the structural data. These include relevant proteins in the biomedical context such as AGMO, an enzyme that has been linked to numerous diseases such as cancer and diabetes[29], DEGS1, another enzyme with a critical role in lipid metabolism[30], or PEMT, another lipid metabolism enzyme and was amongst the first genes associated with non-alcoholic fatty liver disease[31] (**Figure 2e-g** respectively). In all three cases no structural information was known nor any template with enough sequence identity to build a model. However, AlphaFold predicted with high confidence the structures of most of all these proteins (90%, 95% and 91%, respectively).

Structural coverage of regions with biomedical interest

Next, we focused on different subsets of genes of biomedical interest, as their protein structures can be useful in the design of new treatments[32] or to understand the consequences of genetic variants[33,34]. For this purpose, we focused on the protein structural coverage of disease related genes and mutations extracted from DisGeNet[35], OncoKB[36], TCGA[15] and ClinVar[37] databases (**Methods**).

Based on PDB structures or template-based homology modelling solely, the structural coverage of disease genes from DisGeNet (50%), cancer driver genes from TCGA (54%) and OncoKB (50%) or genes related to neurological (58%) or autoimmune (61%) diseases is always higher than the average coverage for the human proteome (47%, **Figure 3a**), likely due to an increased effort in determining the structure of disease-associated proteins which actually count with more experimental structures (sequence identity >95%) compared to non disease-associated proteins. Therefore, it is not surprising that for these genes, the contribution of AlphaFold is not as high as for the rest of the proteome (**Figure 3a**). Particularly, DisGeNet genes gain between 14% to 28% of structural coverage, depending if

we consider only high confidence results (pLDDT >90) or all AlphaFold predictions (pLDDT >50), reaching a total coverage up to 78%. In the case of cancer driver genes, the contribution is lower, as TCGA and OncoKB genes gain 5% and 7% of high confidence structural coverage. Overall, the contribution from AlphaFold with some degree of confidence together with PDB previous knowledge, sets the current high-quality structural coverage for cancer driver genes (TCGA and OncoKB) to 52% and the overall coverage to 70%.

In the context of mutations, we analyzed the structural coverage of Clinvar (**Figure 3b**) considering their annotated pathogenicity. In the case of mutations annotated as “benign”, before AlphaFold we had high-quality annotations for approximately 31% and AlphaFold added roughly 12% of high quality annotations, putting our current high-quality structural coverage at 43%. For pathogenic mutations, instead, we already had a better high-quality structural coverage (71%), so when adding the high-quality AlphaFold predictions (12%) we now reach a high quality coverage for 83% of all pathogenic mutations in Clinvar.

In the case of TCGA (**Figure 3c**), mutations were classified according to whether they were predicted to be either oncogenic, located in cancer driver genes or located in any other protein (likely passengers). Results resemble those observed in the ClinVar dataset. The coverage of likely passenger mutations follows the same distribution as for the rest of the proteome, rising from 50% to 77% when considering AlphaFold predictions. In the case of driver genes, however, we started from a much higher high-quality structural coverage (**Figure 3c**). We had high-quality structural coverage from PDB for 51% of the non-oncogenic mutations in these genes, and AlphaFold added an additional 4%, putting the total at 55%. For mutations predicted to be oncogenic[15], we already had high-quality models for 92% of them, and AlphaFold only added 2% more, putting the total at 94%. We observed the same trend for another recent set of predicted oncogenic mutations by *in silico* mutagenesis of cancer driver genes[38], with high-quality structural data from PDB covering 74% of all oncogenic mutations and AlphaFold only adding 2%, putting the total at 76% (**Figure 3d**). It should be noted, though, that the algorithms predicting the oncogenicity of both sets of somatic mutations use in part structural information, so the results are likely biased towards regions with pre-existing structural data.

In summary, although the additional contribution of AlphaFold models to pathogenic and oncogenic mutations is apparently low, it still provides important structural context for some mutations. This is the case of B3GALT6, a gene without structural information and that, when mutated, can cause a Mendelian disease. Thanks to the high-confidence structural model provided by AlphaFold, it is now possible to appreciate a 3D clustering of the disease-causing mutations in this gene (**Figure 3e**). Similarly, oncogenic mutations in the cancer driver gene MED12 previously had no structural data at all, but are now covered by a high-confidence model. In fact, the mapping of the three oncogenic somatic mutations in MED12 to the AlphaFold model reveals that, despite being over 600 aminoacids apart (positions 521, 879 and 1138), they form a cluster together according to mutation3D[39] ($p < 0.002$, 10.000 bootstrapping iterations), suggesting that they potentially share an oncogenic mechanism (**Figure 3f**).

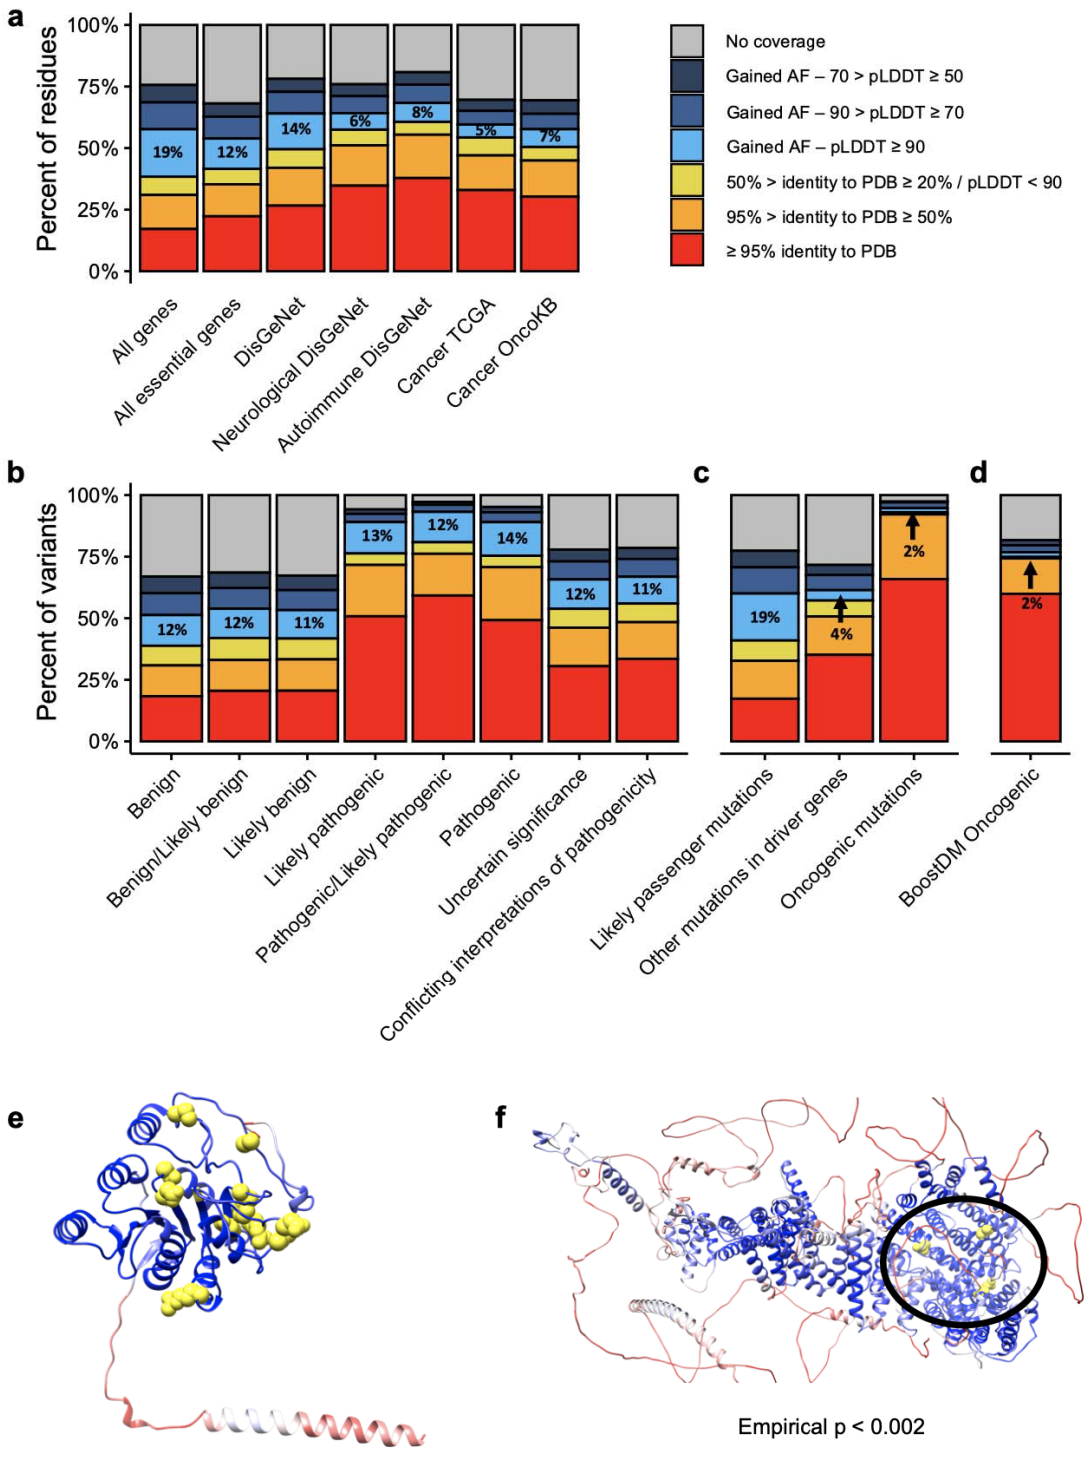

Figure 3 - Changes in structural coverage of biomedical proteins due to AlphaFold models. a) Current structural coverage (y-axis) of different subsets of proteins (x-axis). Bars are colored according to the source of the structural coverage. **b)** Same as **a** but focusing on Clinvar mutations classified by their pathogenicity (x-axis). **c)** Same as **a** but focusing on somatic mutations from TCGA, classified by their likely oncogenicity (x-axis). **d)** Same as **a** but focusing on oncogenic mutations from BoostDM. **e)** AlphaFold model for B3GALT6. Residues are colored according to their

pLDDT from red (lower values) to blue (higher values). Pathogenic mutations from Clinvar are highlighted in yellow. **e)** AlphaFold model for MED12. Coloring is the same as for **d**, but yellow residues indicate oncogenic mutations.

Discussion

In this study we quantified the contribution of AlphaFold protein models to our already considerable knowledge of the human proteome. Before AlphaFold models were released we had high-quality structural data for a significant fraction of the human proteome (31%), either through experimental structures (17% of all residues) or through homology models built with templates of high sequence identity (14% of all residues). However, there was still a large fraction of the proteome for which, despite enough evidence about its foldable potential (i.e. proteins with good representation in terms of the corresponding PFAM), we had no structural data.

The models from AlphaFold have increased our high-quality coverage by an additional 20%, putting it at half of all human protein residues. Moreover, they have virtually closed the gap of structureless foldable regions, as we now have some structural coverage (either through AlphaFold or PDB low sequence identity models) for 75% of all human residues, including almost all PFAM domains. Regarding the latter, another important dataset that has been recently published are the structural models predicted with RoseTTAFold, another artificial intelligence tool, for all PFAM domains[40]. In the future, it will be interesting to compare the predictions of both tools for different PFAM domains, especially for those for which we previously had no experimental structures. Finally, approximately 23% of all human residues are likely to be disordered, which puts the fraction of the human proteome that can potentially have a structure but for which we have no models or templates at only 2%, a remarkable achievement.

In terms of individual proteins, while before AlphaFold we had no structural coverage at all for 4.832 proteins, this number has been reduced to only 1.336 if one considers only proteins with high-quality AlphaFold models (pLDDT > 90 in more than 10% of the protein) and to 29 if one includes all AlphaFold predictions. Importantly, these previously structureless proteins for which we now have high-accuracy models include several important genes in human disease such as DEGS1, ARID2 or B3GALT6.

Our results also suggest that most of the genuinely new contributions of AlphaFold to the structural coverage are in genes that are olfactory receptors, proteins integral to the membrane or, most importantly, genes that have no Gene Ontology annotations. However, one must also bear in mind that some of these predictions will not be correct, either due to actual errors of the software or, in other cases, due to missannotations of public databases such as in the case of CRIPAK[41]. On the positive side, we also expect AlphaFold predictions to enlighten the function of many of the proteins for which we previously had no GO annotations. Since we now have high-confidence predictions for, at least, some parts of many of these proteins, it should be easier to identify protein domains and begin to assign functions and biological processes to these proteins.

In the specific case of disease-associated genes, our high-quality coverage before AlphaFold models were released was higher (between 50% and 60% of all residues) than

the average of the human proteome (38%). This could be, among other reasons, because disease-associated genes are more studied than the rest, as evidenced by their increase in their relative representation in PDB: in 1995 only 2% of all PDB structures were from human cancer genes, whereas by 2020 the fraction has increased to 4%. Therefore, the contribution of AlphaFold to this group of proteins is more limited, between 5% and 14% if one considers high quality models, compared to 19% for all residues in the human proteome.

In the case of mutations the bias is even higher, not only because disease-associated mutations are more studied, but also because they tend to be located in protein regions that form structures. For this reason AlphaFold models only provide high-confidence structural details for an additional 12% and 2% for Clinvar pathogenic and TCGA oncogenic mutations respectively. That being said, AlphaFold models can be used to generate hypotheses behind the pathogenic mechanisms of some of these mutations where we previously had no structure data at all, as we have exemplified for MED12.

In conclusion, we now have either experimental data, template-based homology models or artificial intelligence-derived models for virtually all foldable protein regions in the human proteome (75% of the proteome). Moreover, 50% of the human proteome is covered with experimental or high-quality structural data (more than 50% sequence identity to PDB or more than 90 pLDDT for AlphaFold models). While at the end of 2020 we were already in a very good position (30% of the proteome), AlphaFold has greatly contributed in closing the gap of protein structure coverage, providing details in protein families that were not amenable to other structure determination approaches. Overall, we have practically completed our coverage of the human foldable proteome, an outstanding achievement that will open the door to more challenging problems in structural biology in general and in biomedicine in particular, such as the prediction of protein-protein interaction complexes[42,43], the relative positioning of protein domains in multi-domain proteins[44], the identification of immunogenic peptides[45] or the prediction of the consequences of different types of mutations.

Acknowledgements

E.P-P and V.R-S are supported by the La Caixa Junior Leader Fellowship LCF/BQ/PI18/11630003. A.V. is supported by Institució Catalana de Recerca Avançada (ICREA). The authors would like to thank the DeepMind team for sharing their models of human proteins.

Methods

Homology search

A BLAST[24] search was performed against 596.542 PDB chains extracted from a total of 253.444 protein bioassemblies deposited in PDB by the date of download (04/03/2021) using the human proteome as a query (Ensembl v103) [24]. We only used one protein isoform per gene, either the one defined as the principal isoform by APPRIS[23] or, when two or more isoforms were principal we randomly chose one of them. In the end we used 20.503 protein sequences. Significant hits between ENSEMBL protein sequences and PDB chains were those with e-value lower than 1e-7 and sequence identity percent greater or equal to 20%.

PFAM domains

We used Pfamscan[46] to identify PFAM domains[26] in the 20.503 ENSEMBL sequences. The database of PFAM-A models was downloaded on June 29th 2021 and created on March 19th 2021. We only kept those PFAM domains identified with an e-value below 1e-8.

IUPRed2

To identify intrinsically disordered regions we used IUPRed2[47] (downloaded on April 17th 2021). We used the “long” disorder option and considered as disordered all those residues with a score above 0.5.

AlphaFold models

AlphaFold models for human proteins were downloaded on July 23rd 2021 from https://alphafold.ebi.ac.uk. We extracted the sequences and compared them to the ENSEMBL protein sequences used for the PDB analysis. For comparison purposes all the analyses and results presented here are based on the subset of 17.006 protein sequences for which the ENSEMBL and AlphaFold protein sequences were identical. We also extracted pLDDT values for each residue from the AlphaFold models, as these are stored as if they were the B-factor of the protein coordinates file[19,21].

Gene lists

The list of “All essential genes” includes all those with a pLI score[48] above 0.9. Disease-associated gene annotations come from DisGeNet[35]. The list of “DisGeNet” genes from Figure 3a consists of all DisGeNet genes with a score above 0.45 and where “diseaseType” is “disease”.

The list of “Neurological DisGeNet” genes includes all genes from DisGeNet where the score is above 0.45 and the diseaseName is in the following list: "Alzheimer_s Disease", "Amyotrophic Lateral Sclerosis", "Bipolar Disorder", "Schizophrenia", "Depressive disorder", "Mental Depression", "Autistic Disorder", "Major Depressive Disorder" and "Parkinson Disease".

The list of “Autoimmune DisGeNet” genes includes all genes from DisGeNet where the score is above 0.45 and the diseaseName is in the following list: "Asthma", "Diabetes Mellitus,

Non-Insulin-Dependent", "Lupus Erythematosus, Systemic", "Rheumatoid Arthritis", "Crohn Disease", "diabetes Mellitus, Insulin-Dependent", "Ulcerative Colitis" and "Psoriasis".

The list of “Cancer TCGA” genes includes all those found as drivers in the PanCancerAtlas analysis[15]. “Cancer OncoKB” includes all genes annotated as cancer genes in OncoKB[36] (downloaded on July 26th 2021).

Variant datasets

Variant files from TCGA and ClinVar datasets were retrieved from public repositories (https://api.gdc.cancer.gov/data/1c8cfe5f-e52d-41ba-94da-f15ea1337efc and https://ftp.ncbi.nlm.nih.gov/pub/clinvar/vcf_GRCh38/clinvar.vcf.gz). The protein location of all variants was predicted using the Ensembl Variant Effect Predictor (VEP[49]; version 98.3). Annotations regarding the clinical significance of Clinvar variants were extracted from the Clinvar file. Oncogenic mutations from TCGA include all those predicted to be oncogenic in [15]. The list of driver genes used in Figure 3d is from OncoKB. Mutations from BoostDM[38] were obtained from the IntoGen website (www.intogen.org).

Software

All statistical analyses were done using R 4.0.2. Graphical plots were create with the packages “ggplot2”[50], “pacthwork” and “reshape2”. Molecular graphics and analyses performed with UCSF Chimera[51], developed by the Resource for Biocomputing, Visualization, and Informatics at the University of California, San Francisco, with support from NIH P41-GM10331.

References

1. Kendrew, J. C. *et al.* A three-dimensional model of the myoglobin molecule obtained by x-ray analysis. *Nature* **181**, 662–666 (1958).
2. Berman, H. M. *et al.* The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242 (2000).
3. Chothia, C. & Lesk, A. M. The relation between the divergence of sequence and structure in proteins. *EMBO J.* **5**, 823–826 (1986).
4. Sali, A. & Blundell, T. L. Comparative protein modelling by satisfaction of spatial restraints. *J. Mol. Biol.* **234**, 779–815 (1993).
5. Godzik, A., Kolinski, A. & Skolnick, J. Topology fingerprint approach to the inverse protein folding problem. *J. Mol. Biol.* **227**, 227–238 (1992).
6. Göbel, U., Sander, C., Schneider, R. & Valencia, A. Correlated mutations and residue contacts in proteins. *Proteins* **18**, 309–317 (1994).
7. Bowie, J. U., Lüthy, R. & Eisenberg, D. A method to identify protein sequences that fold into a known three-dimensional structure. *Science* **253**, 164–170 (1991).
8. Jones, D. T., Taylor, W. R. & Thornton, J. M. A new approach to protein fold recognition. *Nature* **358**, 86–89 (1992).
9. Moult, J., Pedersen, J. T., Judson, R. & Fidelis, K. A large-scale experiment to assess protein structure prediction methods. *Proteins* **23**, ii–v (1995).
10. Kryshtafovych, A., Schwede, T., Topf, M., Fidelis, K. & Moult, J. Critical assessment of methods of protein structure prediction (CASP)-Round XIII. *Proteins* **87**, 1011–1020 (2019).
11. Korkegian, A., Black, M. E., Baker, D. & Stoddard, B. L. Computational thermostabilization of an enzyme. *Science* **308**, 857–860 (2005).
12. Meng, X.-Y., Zhang, H.-X., Mezei, M. & Cui, M. Molecular docking: a powerful approach for structure-based drug discovery. *Curr. Comput. Aided Drug Des.* **7**, 146–157 (2011).
13. Tokheim, C. *et al.* Exome-Scale Discovery of Hotspot Mutation Regions in Human Cancer Using 3D Protein Structure. *Cancer Res.* **76**, 3719–3731 (2016).
14. Mosca, R. *et al.* dSysMap: exploring the edgetic role of disease mutations. *Nat.*

Methods **12**, 167–168 (2015).

15. Bailey, M. H. *et al.* Comprehensive Characterization of Cancer Driver Genes and Mutations. *Cell* **174**, 1034–1035 (2018).
16. Tokheim, C. & Karchin, R. CHASMplus reveals the scope of somatic missense mutations driving human cancers. doi:10.1101/313296.
17. Adzhubei, I., Jordan, D. M. & Sunyaev, S. R. Predicting functional effect of human missense mutations using PolyPhen-2. *Curr. Protoc. Hum. Genet.* **Chapter 7**, Unit7.20 (2013).
18. Chen, H. *et al.* Comprehensive assessment of computational algorithms in predicting cancer driver mutations. *Genome Biol.* **21**, 43 (2020).
19. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* (2021) doi:10.1038/s41586-021-03819-2.
20. Callaway, E. 'It will change everything': DeepMind's AI makes gigantic leap in solving protein structures. *Nature* **588**, 203–204 (2020).
21. Tunyasuvunakool, K. *et al.* Highly accurate protein structure prediction for the human proteome. *Nature* (2021) doi:10.1038/s41586-021-03828-1.
22. Howe, K. L. *et al.* Ensembl 2021. *Nucleic Acids Res.* **49**, D884–D891 (2021).
23. Rodriguez, J. M. *et al.* APPRIS 2017: principal isoforms for multiple gene sets. *Nucleic Acids Res.* **46**, D213–D217 (2018).
24. Altschul, S. F., Gish, W., Miller, W., Myers, E. W. & Lipman, D. J. Basic local alignment search tool. *Journal of Molecular Biology* vol. 215 403–410 (1990).
25. Dunker, A. K. *et al.* Intrinsically disordered protein. *J. Mol. Graph. Model.* **19**, 26–59 (2001).
26. Mistry, J. *et al.* Pfam: The protein families database in 2021. *Nucleic Acids Res.* **49**, D412–D419 (2021).
27. Perdigão, N. *et al.* Unexpected features of the dark proteome. *Proc. Natl. Acad. Sci. U. S. A.* **112**, 15898–15903 (2015).
28. Mi, H., Muruganujan, A., Casagrande, J. T. & Thomas, P. D. Large-scale gene function

analysis with the PANTHER classification system. *Nat. Protoc.* **8**, 1551–1566 (2013).

29. Sailer, S., Keller, M. A., Werner, E. R. & Watschinger, K. The Emerging Physiological Role of AGMO 10 Years after Its Gene Identification. *Life* **11**, (2021).
30. Karsai, G. *et al.* DEGS1-associated aberrant sphingolipid metabolism impairs nervous system function in humans. *J. Clin. Invest.* **129**, 1229–1239 (2019).
31. Song, J. *et al.* Polymorphism of the PEMT gene and susceptibility to nonalcoholic fatty liver disease (NAFLD). *FASEB J.* **19**, 1266–1271 (2005).
32. Somody, J. C., MacKinnon, S. S. & Windemuth, A. Structural coverage of the proteome for pharmaceutical applications. *Drug Discov. Today* **22**, 1792–1799 (2017).
33. Wang, Z. & Moult, J. SNPs, protein structure, and disease. *Hum. Mutat.* **17**, 263–270 (2001).
34. Raimondi, F. *et al.* Insights into cancer severity from biomolecular interaction mechanisms. *Sci. Rep.* **6**, 34490 (2016).
35. Piñero, J. *et al.* The DisGeNET knowledge platform for disease genomics: 2019 update. *Nucleic Acids Res.* **48**, D845–D855 (2020).
36. Chakravarty, D. *et al.* OncoKB: A Precision Oncology Knowledge Base. *JCO Precis Oncol* **2017**, (2017).
37. Landrum, M. J. *et al.* ClinVar: improvements to accessing data. *Nucleic Acids Res.* **48**, D835–D844 (2020).
38. Muiños, F., Martínez-Jiménez, F., Pich, O., Gonzalez-Perez, A. & Lopez-Bigas, N. In silico saturation mutagenesis of cancer genes. *Nature* (2021) doi:10.1038/s41586-021-03771-1.
39. Meyer, M. J. *et al.* mutation3D: Cancer Gene Prediction Through Atomic Clustering of Coding Variants in the Structural Proteome. *Hum. Mutat.* **37**, 447–456 (2016).
40. Baek, M. *et al.* Accurate prediction of protein structures and interactions using a three-track neural network. *Science* (2021) doi:10.1126/science.abj8754.
41. Abascal, F. *et al.* Loose ends: almost one in five human genes still have unresolved coding status. *Nucleic Acids Res.* **46**, 7070–7084 (2018).

42. Porta-Pardo, E., Garcia-Alonso, L., Hrabe, T., Dopazo, J. & Godzik, A. A Pan-Cancer Catalogue of Cancer Driver Protein Interaction Interfaces. *PLoS Comput. Biol.* **11**, e1004518 (2015).
43. Wang, X. *et al.* Three-dimensional reconstruction of protein networks provides insight into human genetic disease. *Nat. Biotechnol.* **30**, 159–164 (2012).
44. Xu, D., Jaroszewski, L., Li, Z. & Godzik, A. AIDA: ab initio domain assembly for automated multi-domain protein structure prediction and domain-domain interaction prediction. *Bioinformatics* **31**, 2098–2105 (2015).
45. De Mattos-Arruda, L. *et al.* Neoantigen prediction and computational perspectives towards clinical benefit: recommendations from the ESMO Precision Medicine Working Group. *Ann. Oncol.* **31**, 978–990 (2020).
46. Mistry, J., Bateman, A. & Finn, R. D. Predicting active site residue annotations in the Pfam database. *BMC Bioinformatics* **8**, 298 (2007).
47. Mészáros, B., Erdos, G. & Dosztányi, Z. IUPred2A: context-dependent prediction of protein disorder as a function of redox state and protein binding. *Nucleic Acids Res.* **46**, W329–W337 (2018).
48. Lek, M. *et al.* Analysis of protein-coding genetic variation in 60,706 humans. *Nature* **536**, 285–291 (2016).
49. McLaren, W. *et al.* The Ensembl Variant Effect Predictor. *Genome Biol.* **17**, 122 (2016).
50. Wilkinson, L. ggplot2: Elegant Graphics for Data Analysis by WICKHAM, H. *Biometrics* vol. 67 678–679 (2011).
51. Pettersen, E. F. *et al.* UCSF Chimera--a visualization system for exploratory research and analysis. *J. Comput. Chem.* **25**, 1605–1612 (2004).