Atypical Structural Tendencies Among Low-Complexity Domains in the Protein Data Bank Proteome

Sean M. Cascarina*, Mikaela R. Elder, and Eric D. Ross*

Department of Biochemistry and Molecular Biology, Colorado State University, Fort Collins, CO 80523, USA

*Co-corresponding authors

Email: Sean.Cascarina@colostate.edu

Email: Eric.Ross@colostate.edu

Abstract

A variety of studies have suggested that low-complexity domains (LCDs) tend to be intrinsically disordered and are relatively rare within structured proteins in the protein data bank (PDB). Although LCDs are often treated as a single class, we previously found that LCDs enriched in different amino acids tend to exhibit substantial differences in protein metabolism and function. Therefore, we wondered whether the structural conformations of LCDs are likewise dependent on which specific amino acids are enriched within each LCD. Here, we directly examined relationships between enrichment of individual amino acids and secondary structure preferences across the entire PDB proteome. Secondary structure preferences varied as a function of the identity of the amino acid enriched and its degree of enrichment. Furthermore, divergence in secondary structure profiles often occurred for LCDs enriched in physicochemically similar amino acids (e.g. valine vs. leucine), indicating that LCDs composed of related amino acids can have distinct secondary structure preferences. Comparison of LCD secondary structure preferences with numerous pre-existing secondary structure propensity scales resulted in relatively poor correlations for certain types of LCDs, indicating that these scales may not capture secondary structure preferences as sequence complexity decreases. Collectively, these observations provide a highly resolved view of structural preferences among LCDs parsed by the nature and magnitude of single amino acid enrichment.

Author Summary

The structures that proteins adopt are directly related to their amino acid sequences. Low-complexity domains (LCDs) in protein sequences are unusual regions made up of only a few different types of amino acids. Although this is the key feature that classifies sequences as LCDs, the physical properties of LCDs will differ based on the types of amino acids that are found in each domain. For example, the sequences "AAAAAAAAAA", "EEEEEEEEEE", and

"EEKRKEEEKE" will have very different properties, even though they would all be classified as LCDs by traditional methods. In a previous study, we developed a new method to further divide LCDs into categories that more closely reflect the differences in their physical properties. In this study, we apply that approach to examine the structures of LCDs when sorted into different categories based on their amino acids. This allowed us to define relationships between the types of amino acids in the LCDs and their corresponding structures. Since protein structure is closely related to protein function, this has important implications for understanding the basic functions and properties of LCDs in a variety of proteins.

Keywords

Low-complexity domain; α-helix propensity; β-sheet propensity; intrinsic disorder; protein structure

Introduction

Protein sequence, in combination with surrounding physicochemical environment, fundamentally determines protein form and function. Protein sequences vary with respect to sequence "complexity", which is effectively a measure of the amino acid frequencies within a region of defined length. Most naturally-occurring proteins contain a rich mixture of residues drawn from the 20 canonical amino acids. However, a surprising number of proteins possess relatively simple regions, termed low-complexity domains (LCDs), which are characteristically comprised of only a subset of amino acids.

LCDs are found across a broad array of organisms [1], but are particularly abundant in eukaryotes and are involved in a variety of normal and pathological processes [2,3,12–19,4–11]. Consequently, a number of studies have examined the functions of LCDs and LCD-containing proteins. Proteins with LCDs exhibit differences in metabolic regulation, molecular interaction

partners, subcellular localization, and molecular functions [10,20,21], all of which hinge upon the primary amino acid of which the LCDs are comprised [10]. Importantly, even LCDs enriched in highly similar amino acids can be associated with specialized molecular functions and distinct trends in metabolic regulation. For example, while both N-rich and Q-rich LCDs are associated with transcription, N-rich domains are also associated with P-bodies, whereas Q-rich domains are associated with functions related to endocytosis [10]. Similarly, while enrichment of most hydrophobic amino acids is associated with low protein abundance, low protein half-life, and low translation efficiency, progressive enrichment of alanine or valine is associated with complete opposite trends in protein metabolism.

While the functions of LCDs have been explored in a variety of organisms, the structural characteristics of LCDs are relatively poorly understood. Early studies of LCDs indicated that they are underrepresented among structured proteins in the Protein Data Bank (PDB) [22,23], leading to the pervasive perception that LCDs are generally intrinsically disordered. Indeed the most abundant types of LCDs in eukaryotic proteomes are those enriched in polar and charged amino acids [6,7,10]. Naturally, these types of LCDs tend to adopt an ensemble of disordered conformations, which may be expanded or collapsed depending on the composition and distribution of constituent amino acids [24]. However, under certain conditions, even LCDs that favor disorder may adopt ordered conformations. For example, changes in intracellular environment (e.g. pH) might induce folding of a disordered domain, or disordered regions may adopt stable conformations when docked to folded intramolecular domains or when bound to intermolecular binding partners. Additionally, LCDs with less-extreme polar/charge content may inherently favor folded conformations. Therefore, it is not unreasonable to expect a variety of LCDs to adopt stable structures under particular conditions or in particular sequence contexts, which is supported by recent analyses of LCDs (treated as a single class) in structured proteins [3].

A variety of methods have been developed to identify LCDs [25–30]. Most of these methods employ pre-defined statistical thresholds to classify domains as low or high complexity and demarcate LCD boundaries. While these methods are extremely valuable and have contributed to a number of important discoveries, they also inherently possess a number of limitations [10]. In essence, all LCD detection methods are fundamentally attempting to define regions in which a small subset of the possible amino acids are present at high density. While unified by this defining feature, LCDs are really a diverse set of domains with distinct physical properties; these properties depend upon the amino acids predominantly enriched in each LCD. Furthermore, most methods for defining LCDs employ a statistically-defined threshold (which may not reflect biologically-relevant criteria [10]) and do not distinguish between LCDs of moderate compositional enrichment (e.g. 40% alanine content) and LCDs of extreme compositional enrichment (e.g. 100% alanine). Thus, parsing protein domains by both the predominant amino acid enriched in each domain and its degree of enrichment may reveal trends that are obscured when LCDs are grouped without taking these compositional features into consideration.

Here, we directly examine the relationship between local amino acid composition and structural tendencies, with a particular focus on regions at compositional extremes (i.e. LCDs). Although structured proteins are strongly biased toward sequences of high complexity, LCDs are nevertheless abundant among structured proteins. Secondary structure preferences among LCDs diverge as sequences become less complex, and the secondary structure profiles depend on both the type of amino acid enriched and its degree of enrichment. Despite the common assertion that LCDs tend to be disordered, some types of LCDs exhibit a strong preference to be ordered. Additionally, structural preferences diverge even for related amino acids, suggesting that physicochemically-related LCDs are not always functionally synonymous. Finally, LCD-containing proteins are associated with distinct classes of folds, dependent upon the type of LCD (i.e. the amino acid most enriched within the LCD). These observations suggest that

grouping all LCDs into a single class grossly oversimplifies the structural tendencies exhibited by each type of LCD.

Results

A systematic survey of the PDB proteome reveals an abundance of residue-specific LCDs

Although structured proteins present in the PDB are generally biased against LCDs [22,23], this does not preclude the possibilities that: 1) certain types of LCDs (e.g. F-rich LCDs) may tend to be highly structured, and 2) LCDs may be abundant and important within structured proteins, despite the tendency to be less common than regions of higher complexity.

To systematically examine relationships between local amino acid composition and abundance within protein sequences in the PDB, we first obtained a set of non-redundant PDB protein sequences with < 90% homology, with subsequent filtering to remove chains with perfect sequence identity (see Methods). Additionally, membrane-embedded proteins were excluded to account for differences in folding environment, which could confound downstream analyses (see File S1 for a list of all PDB chain ID's included in the final dataset). Each PDB chain was exhaustively scanned with window sizes ranging from 10 to 100 amino acids, increasing the window size by 10 with each iteration (similar to the method implemented in [10]). As mentioned previously, the essential feature underlying all methods of LCD detection is enrichment in one (or a select few) of the 20 canonical amino acids. Therefore, for each window, the percent composition for each amino acid was calculated. For any given type of amino acid, the majority of protein chains exhibit low compositional enrichment (Fig 1), indicating that most proteins are reasonably complex, which is consistent with previous studies and *a priori* expectations. However, for most amino acids, a substantial number of protein chains contain subsequences that achieve high compositions, particularly for smaller scanning

window sizes (Fig 1). For example, for most amino acids there are hundreds to tens of thousands of proteins with 10-amino acid subsequences containing at least 40% composition (as a point of comparison, LCDs identified by the SEG algorithm typically have a minimum of ~25-45% composition of the most enriched amino acid within a 12aa segment). In the most extreme case, >18k protein chains possess a region containing ≥40% leucine residues in a 10-amino acid window. Therefore, LCDs are quite abundant in PDB sequences despite the overall bias toward sequences of higher complexity.

To directly compare the frequencies in Fig 1 with frequencies obtained using traditional methods of identifying LCDs, we employed a modified version of the SEG algorithm. Briefly, PDB sequences were scanned using the default 12aa window size, and all segments scoring ≤2.2 bits were classified as LCD segments. By default, the SEG algorithm does not distinguish between LCDs of different types (i.e. LCDs strongly enriched in different amino acids). Therefore, we modified the SEG algorithm to parse LCDs into amino acid-specific categories based on the most frequent amino acid within each LCD, as indicated in [10] (also described in the Methods section). This modification allows us to dissect specific relationships between LCD types and their corresponding frequencies.

Distinct classes of LCDs exhibit markedly different frequencies within the PDB and within eukaryotic proteomes (Fig 2). While A-rich and L-rich LCDs tend to be the most common individual LCD types among PDB proteins (Fig 2A,B), LCDs enriched in charged or polar residues occupy most of the subsequent top-ranked positions with respect to frequency. When LCDs are combined based on general physicochemical categorization, LCDs enriched in a hydrophilic amino acid are the most abundant class both in terms of the number of PDB proteins containing a hydrophilic LCD (albeit very slightly; Fig 2A, *inset*) and the total number of regions classified as hydrophilic LCDs (Fig 2B, *inset*). To compare LCD frequencies in the PDB to LCD frequencies among whole proteomes, identical analyses were performed for the yeast and human proteomes (Fig 2C,D). In general, for nearly every type of LCD, the number of proteins

with LCDs is higher among the yeast and human proteomes compared with the PDB for nearly every type of LCD, indicating that the PDB is selectively biased against sequences with LCDs both globally and for nearly all types of LCDs individually (which is consistent with previous analyses [3]). Proteins with hydrophilic LCDs are far more common than proteins with hydrophobic LCDs in the yeast proteome, both in terms of individual amino acids (Fig 2C) and when grouped (Fig 2C, *inset*), and this margin is much larger in the yeast proteome than in the PDB. With respect to LCD frequencies, the PDB more closely resembles the human proteome (Fig 2D), although minor differences are also apparent. For example, the margin between proteins with hydrophilic LCDs and proteins with hydrophobic LCDs is slightly larger in the human proteome compared to the PDB, although this margin is not as dramatic as that observed for the yeast proteome. Proteins with G-rich LCDs are similarly abundant in the PDB and human proteomes, while proteins with P-rich LCDs are substantially less common in the PDB compared to the human proteome.

Overall, these data highlight the abundances of LCDs when distinguished on the basis of the predominant amino acid enriched in each LCD within the PDB and certain eukaryotic proteomes. If LCDs are treated as a single class, broad trends in structural preferences will be heavily biased by how common each type of LCD is in PDB sequences, even though LCDs of different types may exhibit dramatically different structural preferences and behavior.

Stratification of LCDs by degree of amino acid enrichment reveals dose-dependent relationships between composition and structural tendencies

We were particularly interested in secondary structure preferences corresponding to short regions of high density of one or a few amino acid types (i.e. LCDs). The SEG algorithm (the most frequently implemented method to detect LCDs) uses a 12-amino acid scanning window by default. Therefore, for consistency with subsequent analyses, we exhaustively re-scanned all PDB sequences with a 12-amino acid scanning window and binned each

subsequence based on the percent composition of each amino acid within each subsequence (Fig 3). For each subsequence, the fraction of each secondary structure type was also calculated and binned with the corresponding subsequence. The fraction of each secondary structure type was then averaged across all subsequences within each percent composition bin (Fig 3) and plotted (Fig 4; see Methods for additional description). For most amino acids, subsequences at extreme percent compositions are overwhelmingly disordered (light grey line), consistent with previous observations (Fig 4). However, for a number of amino acids (particularly hydrophobic and aromatic residues), subsequences maintain a substantial mean fraction of alternative secondary structures. For example, alanine, phenylalanine, isoleucine, leucine, methionine, valine, tryptophan, tyrosine are predominantly associated with α-helix or β-sheet preference at their respective extremes of enrichment. While the degree of enrichment achieved by many of these residue types are often not as extreme as the degree of enrichment for polar/charged residues, these subsequences are nevertheless classified as LCDs.

A detailed inspection of the composition-to-structure preferences reveals more nuanced differences even between amino acids of very similar types. In order to facilitate a direct visual comparison between related amino acids, the mean fraction of each type of secondary structure at increasing composition levels for each amino acid (from Fig 4) were linearly interpolated and indicated as dynamic stacked bar charts in Movie 1. For ease of comparison, amino acids were clustered based on physicochemical properties.

For regions entirely lacking the amino acid of interest ("0 Residues of Interest" frame), secondary structure preferences do not differ dramatically, and each category roughly approximates the whole-PDB preferences. Regions lacking P or G have the highest proportions of α-helix and β-sheet, and these proportions drop to among the lowest within the first shift (domains with only one P or G), consistent with the known role of these residues in disfavoring or precluding these secondary structure types (see, e.g., [31] and [32] for comparisons of multiple α-helix and β-sheet propensity scales, respectively). Mean α-helix and β-sheet content

continues to decline precipitously as P content increases, while the disordered and irregularly-structured fractions increase. Additional dose-response relationships are apparent between composition and secondary structure for many other amino acids. With respect to physicochemically-related groups of amino acids, differences in secondary structure preferences begin to emerge with as few as 3 residues of interest and often become more pronounced as composition continues to increase. For example, among hydrophobic residues, V-rich regions tend to have a higher proportion of β-sheet content relative to regions enriched in other hydrophobic residues, and this disparity progressively increases from 3-6 residues of interest. Regions with higher enrichment of these types of residues are rare or absent within the PDB, resulting in spuriously large apparent changes in secondary structure preferences. Similar trends are apparent for other groups of amino acids. Among aromatic residues, mean α-helix content is highest in F-rich regions, moderate in W-rich regions and low in Y-rich regions, with a concomitant inverse relationship with respect to β-sheet content. Q-rich regions tend to maintain a higher fraction of α-helix content and lower fraction of β-sheet content compared to other uncharged polar residues, and this divergence progressively increases at higher compositions. Among the charged residues, H-rich and D-rich regions tend to be associated with a higher mean fraction of disorder and lower fraction of α-helix content relative to E-rich, R-rich, and K-rich regions; this trend begins to develop at milder levels of enrichment but is most apparent between 5-7 residues of interest.

From these observations we conclude that, while LCDs as a single class are often associated with intrinsic disorder, certain types of LCDs exhibit clear preferences for particular secondary structures, and these structural preferences are dependent upon the degree of enrichment for each residue type.

Direct comparison of compositional enrichment to traditionally-defined LCDs

To directly compare compositional enrichment with a set of LCDs defined using traditional methods, we again implemented our modified version of the SEG algorithm, which parses LCDs into types based on the most strongly-enriched amino acid. To gain an initial point of reference, mean secondary structure content for the entire set of non-redundant PDB chain sequences was first calculated. On average, α-helix content constitutes the highest fraction of structured proteins, followed by β-sheet content and non-disordered "loop or other irregular structure" content (Fig 5). Notably, the mean fraction of intrinsic disorder constitutes ~5-10% of all secondary structure annotations. In contrast, LCDs as a single class are associated with a greater degree of intrinsic disorder and concomitant compaction of the mean fraction of every other secondary structure type except α-helix content (Fig 5). However, more importantly, secondary structure profiles differ considerably when LCDs are parsed into sub-classes based on the predominant enriched amino acid. These differences closely mirror those observed when using our unbiased, composition-centric approach. For example, Q-rich LCDs almost exclusively prefer α-helical or disordered conformations, whereas N-rich LCDs exhibit lower α-helix content, higher β-sheet content, and a higher fraction of residues adopting non-disordered "loop or irregular structure". A remarkably similar trend is observed among their negatively charged counterparts (E-rich and D-rich LCDs, respectively), suggesting that the slightly longer side chains of Q and E (rather than their terminal functional groups) explain these differences. Strikingly, all three main aromatic-type LCDs exhibit remarkably low intrinsic disorder content, roughly equal to that of structured proteins in general ("PDB" group), further emphasizing that the relationship between low sequence complexity and intrinsic disorder is heavily dependent on LCD type.

While these results are largely in agreement with the data in Fig 4 and Movie 1, they highlight an important limitation of grouping LCDs with no regard for composition. Parsing LCDs based on the most-enriched amino acid resolves the issue of combining physically disparate domains, but does not specifically relate secondary structure preferences to the degree of

enrichment for the indicated amino acid. For example, the class of N-rich LCDs include domains

that contain ~25-35% N up to 100% N. As indicated in Fig 1, sequences with lower percent

composition tend to be much more abundant within the PDB, indicating that these sequences

would disproportionately weight the secondary structure preferences within each LCD category.

Furthermore, using a strict, arbitrarily defined complexity threshold does not indicate the degree

of enrichment required to observe the initial development of secondary structure preference, or

how these preferences change at different levels of enrichment. Therefore, although traditional

methods for defining LCDs may still be sufficient to yield important structural insights, further

resolving LCDs based on degree of enrichment for the predominant amino acid allows for

elucidation of dose-dependent relationships between the amino acid composition of LCDs and

their corresponding structural preferences.

Secondary structure preferences among highly enriched LCDs systematically deviate
from secondary structure propensity scales

A variety of statistical and experimental approaches have been used to develop α-helix

and β-sheet propensity scales for the 20 canonical amino acids. However, these scales are

nearly always developed using a set of structured proteins which, as mentioned previously, are

skewed toward high-complexity sequences, yet LCDs often exhibit unusual biophysical behavior

[24]. A simplified interpretation of the secondary structure propensity scales might suggest that

as LCDs approach 100% composition of a single type of amino acid (i.e. homopolymers), the α-

helix and/or β-sheet propensity of that domain as a whole should approach the exact propensity

value for the predominant amino acid within the LCD. However, at present, it is unclear whether

these scales extrapolate well to LCDs.

To directly assess the applicability of current secondary structure propensity scales to

LCDs in structured proteins, we first obtained a set of highly enriched LCDs for each amino acid

type by combining all previously-identified 12aa LCDs with ≥50% of the composition attributable

to that amino acid. Therefore, in principle, the predominant amino acid in each type of LCD should exert the greatest influence on the overall secondary structure propensity of the LCD. The mean fraction of secondary structure content for each type of LCD was then plotted against each of seven α-helix propensity scales [31,33–36] and nine β-sheet propensity scales [32–34,37–40] in a pairwise fashion (Figs S1-S2).

In general, a moderate correlation is observed between LCD-specific mean fraction of α-helix or β-sheet content and established α-helix or β-sheet scales (Figs S1-S2). Correlation coefficients for the α-helix scales range from 0.57 to 0.82, while correlation coefficients for β-sheet scales span a much broader range of 0.31 to 0.75 (Fig S3). However, for both α-helix content and β-sheet content, correlation coefficients tend to be weaker between the LCD-specific mean fraction content and established scales.

For most classes of LCDs, the mean fraction of α-helix and β-sheet content is fairly consistent with propensity values derived from α-helix and β-sheet scales respectively. However, a detailed examination of the correlation plots indicates that LCDs enriched in a few specific amino acids diverge in a consistent manner from each of the secondary structure propensity scales (Figs S1-S2). In order to quantitatively assess these tendencies, the vertical displacement from the line of best fit (i.e. the residual) was calculated for each amino acid within each regression plot, and averaged across all α-helix or β-sheet propensity scales. The residual essentially represents both the direction and magnitude by which structural preferences among LCDs tend to deviate from predicted values according to established scales. With respect to α-helices, LCDs highly enriched in F, L, I, A, or G tend to contain a high fraction of α-helix compared to their predicted values based on regression with established scales (Fig 6A). Notably, in all cases, the residual values for these classes of highly enriched LCDs are positive for all α-helix scales examined, indicating that the observed α-helix content is unanimously higher than expected across all α-helix scales. Conversely, LCDs enriched in H, K, S, Y, E, D, R, T, and N tend to have a lower α-helix content relative to their expected values. For the β-

sheet propensity scales, LCDs enriched in V, Y, G, and T tend to have high β-sheet content relative to their expected fraction of β-sheet content, while a cluster of certain LCDs (including F, Q, H, R, K, L, and E type LCDs) typically have low β-sheet content relative to predicted β-sheet propensity (Fig 6B). For many of the LCD classes we observe a dose-dependent relationship between the amino acid composition and the mean residual value across all α-helix and β-sheet propensity scales (see Figs S4-S5 and File S2 for extensive analysis and associated supplemental discussion). Additionally, complete compositional analysis of the highly-enriched LCDs indicates that the structural preferences exhibited by each class of LCDs are not substantially influenced by a latent preference for another amino acid as the second most common residue (Fig S6).

Overall, these data suggest that while secondary structure propensity scales perform reasonably well when assessing the conformations of some types of LCDs in structured proteins, the structural preferences among specific classes of LCDs consistently deviate from existing α-helix and β-sheet propensity scales.

Specialized LCDs preferentially occur within proteins that adopt particular folds

While LCDs are more often intrinsically disordered relative to sequences of higher complexity, it is clear that many of the specialized LCD types are capable of (if not prefer) adopting stable secondary structures. It is possible that structured LCDs are favored within particular 3D protein fold families, effectively occupying a protein structure niche. To examine this possibility, Pfam fold annotations mapped to PDB entries were obtained from the PDBfam database [41]. For each category of specialized LCD type (as defined by the modified SEG algorithm), the frequency of each protein fold annotation associated with the specialized LCDs was compared to the frequency of the same annotation within all PDB entries included in our nonredundant PDB proteome using Fisher's exact test with Holm-Sidak correction for multiple hypothesis testing (see Methods section for additional details). For each type of LCD, the top 10

significantly enriched Pfam annotations are indicated in Table 1 (complete statistical analyses are provided in File S3). For some well-characterized types of LCDs, enriched Pfam annotations correspond to expected types of folds. For example, C residues are often involved in coordination of Fe-S clusters, and C-rich domains are associated with proteins containing Fe-S binding domains. L-rich LCDs are associated with a number of distinct leucine-rich repeat domain (LLRD) proteins. Collagen proteins are textbook examples of P-rich proteins, and proteins with P-rich LCDs identified in our analyses are associated with proteins containing collagen triple helix repeats.

However, aside from these well-known associations, our data implicate additional relationships between LCDs and tertiary structure families. For example, both G-rich and V-rich LCDs are significantly associated with zinc-binding dehydrogenases and alcohol dehydrogenase GroES-like domains. N-rich and I-rich domains are associated with clostridium neurotoxin proteins. W-rich domains are associated with multiple glycosyl hydrolase families, as well as reverse transcriptase and RNaseH families. S-rich and T-rich LCDs are associated with immunoglobin V-set domains and immunoglobulin C1-set domains; interestingly, immunoglobulin V-set and immunoglobulin C1-set domains tended to be the most strongly underrepresented among nearly all other types of LCDs (File S3), indicating that proteins containing these folds avoid most types of LCDs yet specifically allow for S- and T-rich LCDs. In addition to the shared associations (either between or within a LCD type), many of the enriched Pfam classes are uniquely related to a particular type of LCD. Collectively, these results suggest that each type of LCD may fulfill a specialized structural role in particular tertiary structure families.

Discussion

LCDs have gained recent attention for their importance in a number of molecular processes. While LCDs in protein sequences have been studied for nearly two decades [22,25,42], a complete understanding of the forms and functions of LCDs is still lacking. It is often asserted that LCDs tend to be intrinsically disordered. On a statistical level this is true, since the majority of LCDs are enriched in polar and charged amino acids (see Fig 2 and [6,7,10]) which, at high local density, tend to promote disordered conformations [24]. However, this is an oversimplification on two levels: 1) LCDs are, in reality, a highly heterogeneous class of domains with distinct physicochemical properties based on LCD sequence and amino acid composition, which results in diverging structural preferences, and 2) traditional LCD methods do not distinguish between LCDs with only mild compositional enrichment (~25-35% of a single type of amino acid) and those with extreme compositional enrichment (e.g. true homopolymers). Therefore, both the nature and magnitude of amino acid enrichment within LCDs are effectively lost when LCDs are treated as a single category.

In the present work, we advance the view that LCDs ought to be parsed into sub-types based on the predominant amino acid enriched and its degree of enrichment within each LCD. Analyses by Saqi [43] and Kumari *et al.* [3] indicate that LCDs, when treated as a single class, typically adopt α-helical or coiled conformations in structured proteins. Furthermore, both studies showed that the amino acid compositions of these LCDs differ from the compositions of LCDs in naturally occurring proteins in general protein sequence databases. These observations provided important broad statistics with respect to LCDs as a single class and highlighted the fact that LCDs are not universally disordered. However, they did not specifically dissect the relationships between the compositions of LCDs and their corresponding structural preferences. Here, we observe that although many types of LCDs are indeed predominantly disordered, a number of LCD sub-classes exhibit substantial tendencies to form structured conformations that are strongly dependent the predominant amino acid comprising the LCD. Furthermore, at milder levels of enrichment, even LCD classes typically associated with disorder

may retain at least a minor preference for ordered conformations (e.g. Q-rich, T-rich, E-rich, R-rich, and K-rich LCDs). For all LCD types, structural preferences vary as a function of amino acid composition, and often change in a dose-dependent manner upon progressive enrichment.

With respect to tertiary structure, certain types of LCDs were preferentially associated with particular protein fold families, suggesting that LCDs may fulfill specific structural roles in some protein folds. Although we were able to map Pfam annotations to each PDB chain, Pfam annotations could not be mapped to precise locations within PDB chain sequences. Therefore, Pfam associations with specific types of LCDs can be interpreted in multiple ways. For example, it is possible that LCDs tend to lie directly within associated Pfam-annotated regions, or that LCDs lying outside of Pfam-annotated regions interact with or dock on the domain.

When compared with traditional secondary structure propensity scales, the average secondary structure content among LCD subtypes exhibited only a moderate or weak correlation. Most secondary structure scales are based predominantly on experiments or statistics that are heavily or exclusively reliant upon observations among high-complexity sequences. Our observations suggest that traditional scales may only weakly capture structural tendencies at the more "unusual" compositional extremes that define LCDs. Perhaps the broader principle (which we are certainly not the first to appreciate) is that secondary structure scales are not static: secondary structure propensities for each of the amino acids, though often represented by a single value, likely change based on local amino acid sequence and composition. By definition, LCDs provide an unusual local environment and are, therefore, no exception to this principle.

It is important to note that, although LCDs are abundant within structured proteins, these LCDs likely constitute only a small fraction of all LCDs found in native proteins (refer back to Figs 1 and 2). The secondary structure content of LCDs observed in structured proteins might necessarily be skewed toward structured conformations, and may not reflect the preferences of each type of LCD globally. Furthermore, it is conceivable that structural studies are biased at

the experimental level: 1) LCDs might in some cases be eliminated *a priori* from protein sequences due to the pervasive perception that they tend to be disordered, 2) solution conditions used in structural studies might favor or disfavor ordered conformations for certain types of LCDs, and 3) many types of LCDs may adopt structures contingent upon interaction with specific binding partners. Therefore, at present, it is difficult to evaluate whether incorporating composition-dependence into secondary structure propensity scales would improve *de novo* predictions of secondary structure based on sequence alone.

In conclusion, our results highlight the abundance of LCDs within structured proteins, the unique structural tendencies of LCDs, and the differences in structural tendencies across LCDs of distinct types within the PDB proteome.

Methods

Data acquisition and processing

Protein sequences as well as corresponding secondary structure and intrinsic disorder annotations for all protein chains found in the Protein Data Bank (PDB) were downloaded from https://www.rcsb.org/pdb/static.do?p=download/http/index.html on 6/19/2019. To reduce bias due to redundancy in the PDB, a set of PDB IDs clustered based on a 90% homology threshold were obtained from http://www.rcsb.org/pages/download/ftp on 6/19/2019. For each cluster, the first PDB ID (which is pre-sorted based on the quality of the structure deposited in the PDB) was retained in the "non-redundant" set. While the non-redundant set obtained from the PDB eliminates redundancy *across* PDB structure files, each structure file may contain multiple chains with identical amino acid sequences (e.g. a homotetrameric protein complex). Therefore, the non-redundant set was further filtered by randomly selecting a single chain (and corresponding secondary structure and disorder annotation) from all chains with perfect sequence identity). Finally, to further eliminate LCD biases due to N- or C-terminal tags (e.g.

HA-tags, polyHis tags, etc.), tags were trimmed from PDB chain sequences if the tag sequence overlapped with the first 20 or final 20 residues within the sequence (see File S4 for full list of tags and corresponding tag sequences included in search criteria). The human and yeast reference proteomes (Uniprot ID's UP000005640 and UP000002311, respectively) were downloaded from https://www.uniprot.org/ on 10/5/2019.

Calculation of Shannon entropy

Shannon entropy, which is implemented in the SEG algorithm to identify LCDs [25], was calculated as indicated in [10]. Briefly, the Shannon entropy corresponding to each 12aa subsequence was calculated as:

$$SE = - \sum_{i=1}^{N=20} \frac{n_i}{L} \left(log_2 \frac{n_i}{L} \right) \#(1)$$

where N represents the size of the residue alphabet (N = 20, for the canonical amino acids), n_i is the number of occurrences of the ith residue within the given subsequence of length L=12. LCDs (defined as having SE ≤ 2.2 bits) were then assigned to a LCD category based on the most frequent amino acid within the subsequence (i.e. A-rich LCDs, C-rich LCDs, etc.).

Local amino acid composition algorithm

To examine the distribution of protein segments as a function of composition and window size for the entire non-redundant PDB, all PDB chain protein sequences were exhaustively scanned using window sizes ranging from 10 amino acids to 100 amino acids. For each amino acid, the percent composition for each subsequence was calculated and binned based on the composition for the amino acid of interest. All subsequences containing a non-canonical amino acid were excluded from analyses.

To examine relationships between amino acid composition and secondary structure preferences, a 12-amino acid scanning window (which corresponds to the default window size

implemented in the SEG algorithm [25,42]) was used to exhaustively scan each PDB chain

sequence. The mean fraction of each secondary structure type was calculated independently for

each combination of amino acid and percent composition. Within each dataset corresponding to

one amino acid, these mean fraction values for each secondary structure type were linearly

interpolated and plotted as a dynamic bar chart to generate Movie 1. Additionally, to compare

with LCDs defined using traditional methods, the mean fraction of each secondary structure

type was calculated for each class of LCDs defined by Shannon entropy, as indicated above.

Regression analyses comparing observed secondary structure content to established secondary structure propensity scales

For each set of LCDs analyzed, the mean fraction of α-helix or β-sheet content was

calculated across the set of LCDs, and those values were plotted in a pairwise fashion against

each of the α-helix or β-sheet propensity scales, respectively. Ordinary least squares

regressions were performed using the Python statsmodels package. Residual values were

calculated from the resulting regression lines and the associated mean fraction of α-helix or β-

sheet content for each LCD class. The absolute value of the correlation coefficients for all

pairwise comparisons between secondary structure propensity scales, and between LCD

secondary structures and each secondary structure propensity scale (Fig S3) were calculated

using the Pandas (Python package) correlation function. Correlation matrices in Fig S3A,B were

sorted based on average correlation coefficient.

Statistical analysis of Pfam annotations associated with LCDs

Pfam protein fold annotations were obtained from the PDBfam database ([41];

http://dunbrack2.fccc.edu/protcid/pdbfam/Download.aspx) on 6/25/2019. For each LCD type,

Pfam annotations were mapped to each PDB chain in the LCD category. The frequency of each

Pfam annotation found in each LCD category was then compared to the overall frequency of the

Pfam annotation within the nonredundant PDB proteome using a two-sided Fisher's exact test. Adjusted *p*-values were calculated within each LCD category using the Holm-Šidák method (implemented in the statsmodels package available in Python) to account for multiple hypothesis testing. Additionally, to evaluate the degree of enrichment or depletion of each Pfam annotation associated with LCDs, the odds ratio was calculated as:

$$OR_A = \left[\frac{f_{LCD}}{1 - f_{LCD}}\right] \Big/ \left[\frac{f_{PDB}}{1 - f_{PDB}}\right] \tag{2}$$

where f_{LCD} represents the frequency of the Pfam annotation among a given class of LCDs, and f_{PDB} represents the frequency of the Pfam annotation among all chains in the nonredundant PDB proteome.

Acknowledgements

We would like to thank Dr. Olve Peersen for critical suggestions regarding analytical approach.

Figure Legends

Fig 1. LCDs are abundant in structured proteins. Each subplot indicates the frequency distributions for all peptide subsequences within PDB proteins as a function of scanning window size and amino acid composition for each amino acid indicated in the subplot titles.

Fig 2. Abundances of classically-defined LCDs within the PDB and eukaryotic proteomes. The PDB proteome was exhaustively scanned using a 12aa window size, and the Shannon entropy was calculated for each segment. Segments with Shannon entropy ≤2.2 bits were classified as LCDs. For each amino acid, LCDs were assigned to that category if the frequency of the amino acid was ≥ the maximum frequency for all amino acids within the LCD sequence. The total number of PDB chain sequences with each type of LCD (A) and the total number of

regions scoring as LCDs for each LCD category (B) are indicated. Analogous frequencies for the yeast and human proteomes are indicated in (C) and (D) respectively. For all plots, the LCDs for which the indicated amino acid was clearly the predominant amino acid in the LCD sequence represent "unambiguous" LCDs, while LCDs for which another amino acid was equally abundant within the LCD sequence represent "ambiguous" LCDs. Proteins were preferentially assigned to the "unambiguous" category if they contained at least one region that could be unambiguously identified as an LCD of a given type. Insets in the upper-right corner of each plot indicate the frequencies of LCDs when grouped by physicochemical categories: hydrophobic (A, I, L, M, and V); charged (D, E, H, K, and R); polar (C, N, Q, S, and T); aromatic (F, W, Y); and hydrophilic (combination of charged and polar classes). Total values corresponding to broad physicochemical categories in the insets do not represent the sum of the frequencies of individual types of LCDs, since some proteins contain multiple distinct types of LCDs that fall into the same physicochemical category.

Fig 3. Depiction of computational strategy for relating local amino acid composition to secondary structure annotations across the PDB proteome. For each amino acid, PDB sequences were scanned with a 12aa window size. For each peptide subsequence, the corresponding DSSP (i.e. secondary structure) annotations were sorted into bins based on the frequency of the amino acid of interest (e.g. serine, in the depicted example). After an exhaustive scan of the PDB proteome, the mean fraction of each secondary structure type was calculated within each residue count bin.

Fig 4. Conformational preferences are highly dependent on both amino acid type and degree of enrichment. The PDB proteome was exhaustively scanned using a 12-amino acid window size as depicted in Fig 3. Secondary structure tendencies across a range of enrichment for each amino acid are depicted in separate subplots. Within each subplot, scatter points

represent the mean fraction of each secondary structure type across all peptide sequences within each indicated "residue count" bin.

Movie 1. Progressive development of secondary structure preferences as a function of compositional enrichment. For each amino acid, the mean fraction values corresponding to each type of secondary structure were linearly interpolated across all composition bins. Each stacked bar represents the mean secondary structure content for each secondary structure type corresponding to all peptides with the indicated amino acid composition. For example, in the first frame (titled "0 Residues of Interest"), the "A" stacked bar represents the mean secondary structure content corresponding to all peptides with 0 alanine residues (note that at this stage, all stacked bars a quite similar and effectively represent rough approximations of the mean secondary structure profile for the entire PDB proteome). As the movie progresses, the stacked bars change to reflect the secondary structure content as the composition increases for each of the amino acids on the x-axis. To continue the example, "1 Residue of Interest" indicates the secondary structure profile corresponding to all peptides with 1 alanine for the "A" stacked bar, "2 Residues of Interest" represent the secondary structure profile corresponding to all peptides with 2 alanines, etc. Short interframe intervals are indicative of interpolated values, whereas longer pauses occur as each integer value for "X Residues of Interest" is reached.

Fig 5. LCD classes parsed by predominant amino acid exhibit unique structural tendency profiles. LCDs with Shannon entropy ≤ 2.2 bits were parsed into amino acid categories based on the most frequent amino acid within each peptide subsequence. Bars indicate the mean fraction of each secondary structure type for all subsequences within each amino acid specific LCD bin. For comparison, the mean fraction of each type of secondary structure for all subsequences across all proteins ("PDB" group) and for all peptides qualifying as LCDs combined into a single category ("LCDs" group) are also shown.

Fig 6. Observed secondary structure content for some LCDs systematically deviates from secondary structure propensity scales. The mean fraction of α-helix and β-sheet content corresponding to highly enriched LCDs (defined as those for which a single amino acid comprises at least 50% of the overall composition) were plotted against all secondary structure propensity scales in a pairwise fashion (see Figs S1-S2). For each amino acid, the residuals were calculated from the regression line and indicated in the box plots for α-helix propensity scales (A) and β-sheet propensity scales(B).

Table 1. Top 10 significantly enriched Pfam annotations associated with each LCD category. For each LCD class, up to 10 significantly enriched (lnOR > 0.0) Pfam annotations are indicated in ascending order according to Holm-Sidak corrected *p*-value.

Supplemental File Legends

Fig S1. Observed mean fraction of α-helix among highly enriched LCDs exhibits a moderate correlation with α-helix propensity scales. Scatter plots indicate all pairwise comparisons between the mean fraction of α-helix among each type of LCD and values from established α-helix propensity scales. Each shaded band indicates the 95% confidence interval around the regression line.

Fig S2. Observed mean fraction of β-sheet among highly enriched LCDs exhibits a moderate-to-weak correlation with β-sheet propensity scales. Scatter plots indicate all pairwise comparisons between the mean fraction of β-sheet among each type of LCD and values from established β-sheet propensity scales. Each shaded band indicates the 95% confidence interval around the regression line.

Fig S3. α-helix and β-sheet content among highly enriched LCDs diverges from secondary structure propensity scales. All pairwise comparisons between secondary structure propensity scales, as well as mean fraction of each secondary structure type among highly enriched LCDs, were performed for α-helix (A) and β-sheet (B) propensity scales (see also Figs S1-S2 for each pairwise correlation plot). Heatmap intensities and corresponding values indicate the absolute value of the Pearson's correlation coefficient for each comparison. To indicate general ranking with respect to the overall degree of correlation with all other α-helix or β-sheet propensity scales, the heatmaps were pre-sorted based on the average absolute correlation coefficient in descending order (top-to-bottom, and left-to-right).

Fig S4. Dose-dependent relationships between amino acid composition and divergence from α-helix propensity scales. For each residue count bin, the average fraction of α-helix content was calculated separately for each LCD class. Pairwise regression analyses were performed between each α-helix propensity scale and the mean fraction of α-helix content among LCDs in the given residue count bin. From the regression analyses, the residual values for each LCD class (i.e. each amino acid) were averaged across all α-helix propensity scales. This process was repeated independently for each residue count bin, and the resulting mean residual values are indicated in the figure above. Additionally, the size of each point (as well as the opacity of each point and the opacity of the preceding line segment) reflects the sample size for each residue count bin for each LCD class (i.e. the number of protein regions that were parsed into each bin based on the composition of the amino acid indicated in each subplot title).

Fig S5. Dose-dependent relationships between amino acid composition and divergence from β-sheet propensity scales. Mean residual values were calculated from pairwise regression analyses between observed fraction of β-sheet content among LCDs and each of the

β-sheet propensity scales are indicated above. Regression analyses, calculation of the residuals, and plotting were performed as indicated in the Fig S4 legend.

Fig S6. Highly enriched LCDs do not exhibit strong subsidiary preferences for other amino acids. Heatmap indicating the average frequency of all residues calculated from the highly enriched LCD sequences within each LCD category.

File S1. PDB chain ID's included in the dataset analyzed in this work. The original set of non-redundant PDB sequences (at the 90% sequence identity level) was further filtered to exclude all PDB chains with perfect identity, as well as all PDB chains associated with membrane proteins (as indicated in the Methods section). File S1 contains the final set of PDB chain ID's used for all analyses.

File S2. Supplemental discussion for Figs S4-S5.

File S3. Statistical analyses of Pfam annotations associated with each type of LCD. All Pfam annotations associated with each LCD category (separate tabs) are indicated, along with their frequencies among the LCD category and PDB proteome, their degrees of enrichment/depletion (lnOR), and their associated probability statistics (p-value and Holm-Sidak Corrected p-value).

File S4. Common protein expression/purification tags trimmed from protein termini prior to analyses.

References

1. Haerty W, Golding GB. Low-complexity sequences and single amino acid repeats: not just "junk" peptide sequences. Genome. 2010;53: 753–762. doi:10.1139/G10-063

2. Harrison PM, Gerstein M. A method to assess compositional bias in biological sequences and its application to prion-like glutamine/asparagine-rich domains in eukaryotic proteomes. Genome Biol. 2003;4: R40.

3. Kumari B, Kumar R, Kumar M. Low complexity and disordered regions of proteins have different structural and amino acid preferences. Mol Biosyst. 2015;11: 585–594. doi:10.1039/c4mb00425f

4. Michelitsch MD, Weissman JS. A census of glutamine/asparagine-rich regions: Implications for their conserved function and the prediction of novel prions. Proc Natl Acad Sci. 2000;97: 11910–11915. doi:10.1073/pnas.97.22.11910

5. DePristo MA, Zilversmit MM, Hartl DL. On the abundance, amino acid composition, and evolutionary dynamics of low-complexity regions in proteins. Gene. 2006;378: 19–30. doi:10.1016/j.gene.2006.03.023

6. Harrison PM. Exhaustive assignment of compositional bias reveals universally prevalent biased regions: analysis of functional associations in human and Drosophila. BMC Bioinformatics. 2006;7: 441. doi:10.1186/1471-2105-7-441

7. Radó-Trilla N, Albà M. Dissecting the role of low-complexity regions in the evolution of vertebrate proteins. BMC Evol Biol. 2012;12: 155. doi:10.1186/1471-2148-12-155

8. Coletta A, Pinney JW, Solís DYW, Marsh J, Pettifer SR, Attwood TK. Low-complexity regions within protein sequences have position-dependent roles. BMC Syst Biol. 2010;4: 43. doi:10.1186/1752-0509-4-43

9. Radó-Trilla N, Arató K, Pegueroles C, Raya A, de la Luna S, Albà MM. Key Role of Amino Acid Repeat Expansions in the Functional Diversification of Duplicated Transcription Factors. Mol Biol Evol. 2015;32: 2263–72. doi:10.1093/molbev/msv103

10. Cascarina SM, Ross ED. Proteome-scale relationships between local amino acid

composition and protein fates and functions. Schlessinger A, editor. PLOS Comput Biol. Public Library of Science; 2018;14: e1006256. doi:10.1371/journal.pcbi.1006256

11.   Gomes E, Shorter J. The molecular language of membraneless organelles. Journal of Biological Chemistry. 2019. pp. 7115–7127. doi:10.1074/jbc.TM118.001192

12.   Karlin S, Brocchieri L, Bergman A, Mrazek J, Gentles AJ. Amino acid runs in eukaryotic proteomes and disease associations. Proc Natl Acad Sci U S A. 2002;99: 333–338. doi:10.1073/pnas.012608599

13.   Marcotte EM, Pellegrini M, Yeates TO, Eisenberg D. A census of protein repeats. J Mol Biol. 1999;293: 151–160. doi:10.1006/jmbi.1999.3136

14.   Sim KL, Creamer TP. Abundance and Distributions of Eukaryote Protein Simple Sequences. Mol Cell Proteomics. 2002;1: 983–995. doi:10.1074/mcp.M200032-MCP200

15.   Albà MM, Guigó R. Comparative analysis of amino acid repeats in rodents and humans. Genome Res. 2004;14: 549–554. doi:10.1101/gr.1925704

16.   Faux NG, Bottomley SP, Lesk AM, Irving JA, Morrison JR, De La Banda MG, et al. Functional insights from the distribution and role of homopeptide repeat-containing proteins. Genome Res. 2005;15: 537–551. doi:10.1101/gr.3096505

17.   Simon M, Hancock JM. Tandem and cryptic amino acid repeats accumulate in disordered regions of proteins. Genome Biol. 2009;10. doi:10.1186/gb-2009-10-6-r59

18.   Lobanov MY, Klus P, Sokolovsky IV, Tartaglia GG, Galzitskaya OV. Non-random distribution of homo-repeats: Links with biological functions and human diseases. Sci Rep. 2016;6. doi:10.1038/srep26941

19.   Chong S, Dugast-Darzacq C, Liu Z, Dong P, Dailey GM, Cattoglio C, et al. Imaging dynamic and selective low-complexity domain interactions that control gene transcription. Science (80-). 2018;361. doi:10.1126/science.aar2555

20.   Chavali S, Chavali PL, Chalancon G, de Groot NS, Gemayel R, Latysheva NS, et al. Constraints and consequences of the emergence of amino acid repeats in eukaryotic

proteins. Nat Struct Mol Biol. 2017; doi:10.1038/nsmb.3441

21. Pelassa I, Fiumara F. Differential occurrence of interactions and interaction domains in proteins containing homopolymeric amino acid repeats. Front Genet. 2015;6. doi:10.3389/fgene.2015.00345

22. Wootton JC. Non-globular domains in protein sequences: Automated segmentation using complexity measures. Comput Chem. 1994;18: 269–285. doi:10.1016/0097-8485(94)85023-2

23. Romero P, Obradovic Z, Li X, Garner EC, Brown CJ, Dunker AK. Sequence complexity of disordered protein. Proteins Struct Funct Genet. 2001;42: 38–48. doi:10.1002/1097-0134(20010101)42:1<38::AID-PROT50>3.0.CO;2-3

24. Das RK, Ruff KM, Pappu R V. Relating sequence encoded information to form and function of intrinsically disordered proteins. Curr Opin Struct Biol. 2015;32: 102–112. doi:10.1016/j.sbi.2015.03.008

25. Wootton JC, Federhen S. Statistics of local complexity in amino acid sequences and sequence databases. Comput Chem. 1993;17: 149–163. doi:10.1016/0097-8485(93)85006-X

26. Shin SW, Kim SM. A new algorithm for detecting low-complexity regions in protein sequences. Bioinformatics. 2005;21: 160–170. doi:10.1093/bioinformatics/bth497

27. Li X, Kahveci T. A novel algorithm for identifying low-complexity regions in a protein sequence. Bioinformatics. 2006;22: 2980–2987. doi:10.1093/bioinformatics/btl495

28. Promponas VJ, Enright AJ, Tsoka S, Kreil DP, Leroy C, Hamodrakas S, et al. CAST: an iterative algorithm for the complexity analysis of sequence tracts. Bioinformatics. 2000;16: 915–922. doi:10.1093/bioinformatics/16.10.915

29. Nandi T, Dash D, Ghai R, C BR, Kannan K, Brahmachari SK, et al. A novel complexity measure for comparative analysis of protein sequences from complete genomes. J Biomol Struct Dyn. 2003;20: 657–668. doi:d=3013&c=4104&p=11631&do=detail

[pii]\n10.1080/07391102.2003.10506882 [doi]

30. Claverie JM, States DJ. Information enhancement methods for large scale sequence analysis. Comput Chem. Pergamon; 1993;17: 191–201. doi:10.1016/0097-8485(93)85010-A

31. Pace CN, Scholtz JM. A helix propensity scale based on experimental studies of peptides and proteins. Biophys J. 1998;75: 422–7. Available: http://www.ncbi.nlm.nih.gov/pubmed/9649402%0Ahttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1299714

32. Smith CK, Regan L, Withka JM. A Thermodynamic Scale for the β-Sheet Forming Tendencies of the Amino Acids. Biochemistry. 1994;33: 5510–5517. doi:10.1021/bi00184a020

33. Koehl P, Levitt M. Structure-based conformational preferences of amino acids. Proc Natl Acad Sci U S A. 1999;96: 12524–9.

34. Fujiwara K, Toda H, Ikeguchi M. Dependence of α-helical and β-sheet amino acid propensities on the overall protein fold type. BMC Struct Biol. 2012;12. doi:10.1186/1472-6807-12-18

35. Muñoz V, Serrano L. Elucidating the folding problem of helical peptides using empirical parameters. II†. Helix macrodipole effects and rational modification of the helical content of natural peptides. J Mol Biol. 1995;245: 275–296. doi:10.1006/jmbi.1994.0023

36. Luque I, Mayorga OL, Freire E. Structure-based thermodynamic scale of α-helix propensities in amino acids. Biochemistry. 1996;35: 13681–13688. doi:10.1021/bi961319s

37. Minor DL, Kim PS. Context is a major determinant of β-sheet propensity. Nature. 1994;371: 264–267. doi:10.1038/371264a0

38. Minor DL, Kim PS. Measurement of the β-sheet-forming propensities of amino acids. Nature. 1994;367: 660–663. doi:10.1038/367660a0

39. Street AG, Mayo SL. Intrinsic -sheet propensities result from van der Waals interactions between side chains and the local backbone. Proc Natl Acad Sci. 2002;96: 9074–9076. doi:10.1073/pnas.96.16.9074

40. Kim CA, Berg JM. Thermodynamic β -sheet propensities measured using a zinc-finger host peptide. Nature. 1993;362: 267–270. doi:10.1038/362267a0

41. Xu Q, Dunbrack RL. Assignment of protein sequences to existing domain and family classification systems: Pfam and the PDB. Bioinformatics. 2012;28: 2763–2772. doi:10.1093/bioinformatics/bts533

42. Wootton JC, Federhen S. Analysis of Compositionally Biased Regions in Sequence Databases. Methods Enzymol. 1996;266: 554–571.

43. Saqi M. An analysis of structural instances of low complexity sequence segments. Protein Eng Des Sel. 1995;8: 1069–1073. doi:10.1093/protein/8.11.1069

Figure 1

Figure 3

Figure 4

Figure 6

Figure S1

Figure S2

Figure S4

Figure S5

Figure S6

Figure 2

Figure S3

Figure 5