

Traits and tangles: An analysis of the Big Five paradigm by tangle-based clustering

Hanno von Bergen and Reinhard Diestel

27.11.2024

Abstract

Using the recently developed mathematical theory of tangles, we re-assess the mathematical foundations for applications of the five factor model in personality tests by a new, mathematically rigorous, quantitative method. Our findings broadly confirm the validity of current tests, but also show that more detailed information can be extracted from existing data.

We found that the big five traits appear at different levels of scrutiny. Some already emerge at a coarse resolution of our tools at which others cannot yet be discerned, while at a resolution where these *can* be discerned, and distinguished, some of the former traits are no longer visible but have split into more refined traits or disintegrated altogether.

We also identified traits other than the five targeted in those tests. These include more general traits combining two or more of the big five, as well as more specific traits refining some of them.

All our analysis is structural and quantitative, and thus rigorous in explicitly defined mathematical terms. Since tangles, once computed, can be described concisely in terms of very few explicit statements referring only to the test questions used, our findings are also directly open to interpretation by experts in psychology.

Tangle analysis can be applied similarly to other topics in psychology. Our paper is intended to serve as a first indication of what may be possible.

This paper can be read at two levels.

- At the methodological level it showcases a new mathematical clustering method suited particularly to the social sciences [1, 9]: indirect clustering by *tangles*.
- At the content level it offers additional, non-statistical, soundness validation (which we shall argue is needed) and new structural analysis of the Five Factor Model in personality testing.

Tangles have their origin in structural graph theory. Although they do offer a new, indirect, approach to generic clustering, which we apply here, their particular strength lies in direct applications to problems typical for the social sciences, as outlined in [9].

1 Introduction

The ‘big five’ paradigm, also known as the *Five Factor Model* (FFM) [14], holds that the essence of an individual’s personality can be captured by evaluating it in terms of the five ‘OCEAN’ traits: *openness to experience*, *conscientiousness*, *extraversion*, *agreeableness* and *neuroticism*. These five traits have been advanced as interpretations of the five principal factors originally found by Tupes and Christal [16] in data of Cattell [5] and Fiske [10], and also of five factors found by others in different data; see [13]. In a typical personality assessment based on this paradigm, such as [6, 11, 15], the person to be evaluated has to answer, for each of these five traits, a set of questions designed to test to what degree an individual has the trait in question.

There are a number of fundamental assumptions underlying this approach, which are the topic of this paper.

The first of these stems from the fact that the above five traits are not themselves the factors, in the technical factor-analytical sense, but are interpretations of such factors. In personality tests such as [15], however, they are targeted directly. The multitude of traits studied, e.g., in [5] and [10], has thus been compressed into five factors; these were interpreted; and their interpretations were fanned out again into the many aspects displayed in [12] and assessed in [15].

Compression followed by de-compression always bears the risk of information loss or shift. Hence the following question arises:

- (1) To what extent are current personality tests that are designed to target the traits named as OCEAN related, in a reproducible quantitative sense, to the original Five Factors from which they derive their justification?

We believe that settling this question requires data of a kind we have been unable to find in the literature; see our discussion of this problem in Section 1.1.

Alternatively, one can interpret the reference which these tests make to the original Five Factors as inspiration rather than justification, and seek to validate these tests directly. In particular, do the several questions into which each of the five OCEAN traits is fanned out really test the same thing? For example, we might, in our own use of the term ‘conscientiousness’, count both orderliness and a sense of responsibility for others as aspects of this ‘trait’. But do these two phenomena really occur together in individuals more often than not?

Note that this does not follow even if we assume that ‘conscientiousness’ is a suitable name for one of the five principal factors found in those original studies. What this assumption and the definition of factor analysis tell us is that several of the traits originally tested for are highly correlated with *some* aspect of conscientiousness. It does not tell us that all its aspects we might think of when designing tests such as [15] must be highly correlated with each other.

Ideally, we would like to start with a positive answer to the following question:

- (2) Are the five OCEAN traits extensionally verifiable *as traits*, in some objective sense that does not depend on our own, culturally determined, notions associated with those five words?

In other words, can we ground our interpretative notions of the five groups of personality aspects named as OCEAN in quantitatively verifiable data that

confirms them as genuine traits, whether or not these correspond exactly to the original Five Factors as stipulated in (1)?

Note that no experiment can be expected to prove assertions such as ‘these ten questions are well suited to test for openness’ (say). This is because the only way to define a comprehensive trait such as openness rigorously at the data level is by what mathematicians would call a *class*: to name some explicitly chosen set of ‘markers’ known to be highly correlated with each other, and decree that whatever they have in common (a postulated entity thought to ‘explain’ their high correlation) shall be called ‘openness’. But then this set of markers could be tested for directly – which would make its justification tautological and offer no added value.

However, one *can* develop criteria for the following generalisation of (2), which can also be applied directly to any concrete personality test:

- (3) To what extent does a given set of questions form markers of *some* trait?

Such criteria are suggested in Section 1.2. Their validation involves clustering.

Buchanan, Johnson and Goldberg [4] offer evidence of a positive answer to (3) for a test similar to [15]. This evidence is based on factor analysis run on the answers they obtained for questions designed to test for the OCEAN traits. While factor analysis can, in principle, be used for clustering, this is not its brief or even its forte. We shall argue in Section 1.3 that clustering by factor analysis is too narrowly defined to meet our criteria for (3) from Section 1.2. If one accepts our criticism, a consequence will be that the results from [4] do not sufficiently support a positive answer to (3).

Using a new clustering method developed specifically for ‘fuzzy’ datasets as are common in the social sciences, clustering by *tangles* [9], we offer here an analysis more in line with our criteria for (3) of the test results obtained in [2, 3] for the IPIP questions [12] developed to test the FFM markers of Goldberg [11]. We give a preview of this in Section 2.

Our analysis broadly confirms that the 50 questions asked in [12] correspond to five traits, in a sense that meets our criteria from Section 1.2. Interestingly, though, these traits are not visible at the same level of ‘resolution’ of our clustering tool of tangles: some already emerge at a coarse resolution at which others cannot yet be distinguished, while at a resolution where these *can* be distinguished some of the former traits are no longer visible as clusters, but have disintegrated into single questions. *Our analysis, thus, finds more structure* than just five clusters: structure between the five expected traits, but also between them and their subtraits and more comprehensive traits.

We have found this structure to be robust across different, indeed disjoint, sets of participants whose answers to the questions in [12] we evaluated. But our results are limited, so far, to the data published in [2, 3]. It would be interesting to see if they can be replicated with results obtained in other personality tests developed for the ‘big five’ paradigm.

Tangles have their origin in structural discrete mathematics, in particular in graph theory. The use of tangles in data analysis, including for the social sciences, is new and treated in depth in [9]. An easy-to-read introduction can be found in [7], and we outline what we need here in Section 3. Tangle software is freely available via [1].

1.1 From factors to traits: information loss or shift through interpretation in pursuit of challenge (1)

The FFM has its origin in the fact that several studies found five main factors in personality tests each testing for some larger set of traits, different sets for these original studies. Formally, this precludes any notion that these are ‘the same’ five factors. However, they are said [13] to have similar *interpretations*: those summarised, e.g., as ‘OCEAN’. Moreover, only this shift from the data level to that of interpretation enables us to talk about (five) *traits* at all, rather than about intuitively inaccessible mathematical entities constructed to formally ‘explain’ (linearly combine to) various other traits.

When FFM-based personality tests such as [15] are developed, however, these interpretations necessarily take on a life of their own: the questions designed to test for the ‘big five traits’ target whatever our intuitive ideas for these traits encompass. In particular, they do not target the entirely abstract five factors that Tupes and Christal [16] found in Cattell’s test for 35 other traits.

While this is entirely legitimate, and in practical terms unavoidable, it does involve a double jump: one from the original factors to their interpretations, and another from those interpretations to the inventories designed to test them. Any claim that these inventories test for the original Five Factors, therefore, is a hypothesis that needs validation.

In the case of Tupes and Christal’s five factors this would require a study in which, for example, the 50-question inventory from Goldberg [11] was tested alongside Cattell’s original 35 traits. The set of 10 questions targeting a particular one of the five traits, t say, would then have to be shown to be highly correlated to one of the principal five factors found for Cattell’s 35 traits in *this* study, which in turn would have to be shown to be similar in terms of the Cattell questions to the factor found by Tupes and Christal which t interprets.

However, this is not what happened. Instead, Goldberg [11] subjected the returns for his 50 questions to another run of factor analysis. As expected, he found five factors: one for each set of 10 questions targeting a trait, and highly correlated to these 10 questions. This does not amount to a positive answer to our question (1) – only, at best, to a positive answer to our question (2) or (3).

We look into this further in Section 1.3. We shall argue that factor analysis, while being the well-established tool of choice for *finding* hitherto untargeted fundamental traits in terms of which others can be explained, is not the most powerful tool available for evaluating whether inventories of questions designed to test for these traits live up to their brief. Our aim in this paper is to offer an additional tool for this purpose.

1.2 What is a trait? Two extensional criteria for (2) and (3)

Our question (2) asks whether the five FFM traits are ‘extensionally verifiable’: whether the groupings of personality aspects implicit in our semantic notion of each of the five terms defining these traits are borne out in reality.

For example, there may be reasons for us to subsume certain distinct personality aspects in our notion of agreeableness that do not, in fact, occur more often together in an individual than other aspects. In such a case ‘agreeableness’, while being a convenient term for us to refer to such a *combination* of personality

aspects, would not name a ‘trait’ – at least none fundamental enough to serve as one of five pillars on which to build comprehensive personality assessments.

The work reported in this paper is based on the assumption that there is a consensus that a *trait* should satisfy the two formal requirements outlined below. In Section 1.3 we describe a method which, we shall argue, can validate these two criteria better than what appears to be the current standard method.

Consider a term t in our language for which we wish to decide whether or not we should call it a ‘trait’. To decide this, we might devise a questionnaire Q_t each of whose questions targets one aspect of t included in our informal notion for t , so that all these aspects are represented in Q_t .

Our first requirement on t for qualifying as a trait is what we might call *extensional cohesion*: that the answers received from any sufficiently diverse population for any such set Q_t of questions are highly correlated,¹ statistically or in some other relevant sense.² Only if they are can the answers to the questions in Q_t given by a concrete person be treated as measuring the same thing – the ‘trait’ t . However, such cohesion alone is not enough to warrant this.

Our second requirement on t is what we might call *extensional completeness*. This is defined only when Q_t is one of several sets of questions in some comprehensive questionnaire Q , such as the ‘big five’ inventories in [11] and [12]. It stipulates that the answers received for Q_t shall *not* be highly correlated to those received for any set $Q' \subseteq Q \setminus Q_t$, such as a similar set $Q_{t'}$ devised to test for a trait t' other than t . For if they were, and both Q_t and $Q_{t'}$ (say) had already been shown to be extensionally cohesive, then t and t' would merely be two manifestations of some more comprehensive trait, testable by $Q_t \cup Q_{t'}$ but perhaps not yet understood in interpretation terms. This would raise legitimate concerns about any view of t and t' as different personality traits: it might still be convenient for us, but would not be borne out extensionally as envisaged in (2).

In this paper we report on our analysis of the data in [2, 3], where Q is the set of questions from [12]. This is composed of five subsets Q_t , one for each of the five terms t in OCEAN. We found that each of these Q_t is both extensionally cohesive and extensionally complete in Q at *some* level of ‘resolution’ of our tools (see Section 2.2). But we also found that this does not happen at the same level: there is no level of resolution for our tangle-based validation tool at which all the five sets Q_t simultaneously satisfy the above two requirements on traits.

1.3 Traits as clusters: a generic paradigm for (3)

Our two requirements from Section 1.2 for sets of variables, in our case questions, to qualify for being regarded as capturing some common ‘trait’ as envisaged in (3) are best described in clustering terms.

Let us view the set Q of all the questions included in a given FFM test as a ‘space’,³ each question $q \in Q$ being a point in this space. If we interpret

¹The correlation may be positive or negative. Since positive and negative correlations between two questions can be converted into each other by inverting one of these questions, and positive correlation is transitive if it is high enough, we can make all correlations in a group of pairwise highly correlated questions positive simply by rephrasing some of the questions.

²Tangles in datasets are defined in terms of set partitions rather than pairs of data points. This enables us to apply more comprehensive notions of correlation than statistical correlation of pairs of questions, such as that of ‘mutual information’ from discrete entropy.

³A metric space, for example – we do not need a vector space. Formally, all we need is a set with a positive real-valued function on its pairs.

high correlation between two questions as close proximity in this space, the extensional cohesion and completeness required of each of the five subsets Q_t of Q , one for each trait t , translate to the requirement that these five sets Q_t *should form distinguishable clusters in the space Q* . For example, a subset of Q found to be extensionally incohesive could not be a cluster: it might be a union of several clusters not hanging together, or even of proper subsets of such clusters. An extensionally cohesive but incomplete subset of Q might form only a partial cluster, or there might be no clusters in Q at all.

Buchanan, Johnson and Goldberg [4] and Goldberg [11], offer validations of their FFM test inventories, which are similar to those we investigated [2, 3]. These can be described in clustering terms, too. In their studies, as well as in [2] and [3], there is an inventory Q consisting of five sets Q_t of ten questions, one for each of the five terms t in OCEAN. In order to validate their test, they ran factor analysis on the returns they received. They found five main factors, one factor f_t for each of the five sets Q_t , and note that the questions in Q_t are highly correlated to f_t but not to any other $f_{t'}$. They interpret this as validation of both their test and, presumably, of the FFM itself in the sense of our question (1).

This should then imply a positive answer also to our weaker question (3), and indeed to an extent it does. Expressed in clustering terms, their validation shows that the answers received for their tests form clusters around the five factors they found, viewed as virtual centre points: clusters with respect to correlation interpreted as proximity, as earlier.

However, this confirmation of (3) is not the strongest possible which their data might offer. This is because factor analysis, when used for clustering as here, produces clusters subject to constraints that are not needed: it usually requires the centre points of the clusters to be orthogonal when viewed as vectors in the appropriate inner product space, and the centre points of clusters are chosen so as to maximise proximity not just to the points in their corresponding cluster, but to *all* the data points (subject to being orthogonal). This is mitigated somewhat by an application of *varimax* to the five factors originally found, but remains an unnecessary constraint from a pure clustering perspective.




Figure 1: Factor analysis, used for clustering of variables with respect to correlation as proximity, might find the black dot as a cluster centre.

Figure 1 shows an extreme example to illustrate the latter point. The small dots indicate the variables to be clustered. Two dots are close in the figure if they are highly correlated as variables. The central black dot indicates a combined variable that factor analysis might find: it is computed so as to maximise its

average correlation to *all* the variables to be clustered, so the black dot is placed so as to minimise its average distance from the small dots. Standard clustering would instead find the three coloured dots as cluster centres: they each maximise the average correlation only to the group of variables in their respective cluster.

More generally, recasting our two criteria from Section 1.2 in clustering terms indicates that statistical correlation – even without the two unnecessary constraints just mentioned – is not the only possible basis for their quantification: it is one way to measure the similarity of variables (questions), but not the only one. Even more generally, there may be – and indeed are – ways of defining and detecting clusters that are not just based on the similarity of *pairs* of points.

Tangles can detect, and quantify, clusters defined in more general ways than by similarity of pairs of points, let alone just by statistical correlation. Moreover, they unify all standard clustering methods by extracting something like their common structural essence.

1.4 Testing for our two criteria: from clustering to tangles

In Section 1.3 we argued that generic clustering may offer a way of validating FFM tests that is not subject to some constraints inherent in using factor analysis for this purpose. In this section we make the case that, among the generic clustering methods available, tangle-based clustering may be a particularly good choice to detect clusters meeting our criteria from Section 1.2 in particularly ‘fuzzy’ environments such as personality data.

Clustering, traditionally, seeks to divide a set of data points into subsets of points that are particularly tightly knit together. Sometimes, in *soft clustering*, points are allowed to spread their ‘membership’ over more than one cluster. But there is no generally accepted definition of ‘cluster’. This appears to be unavoidable, because the task to decide for every single data point which, if any, cluster it should be assigned to, is simply too much to ask of real-world data.

Tangles are designed to overcome this problem. They deliberately refrain from attempting to define a cluster as a set of points. Instead, they merely point to roughly where in the data space clusters can be found, and say something about their relative structure – such as whether some large but loose cluster includes one or more smaller dense ones. The gain from ignoring the precise (but error-prone and partly random) information about individual points given in the data, which yields only fuzzy information about cluster-like sets of such points, is that the coarser information retained is precise, robust, and can be examined mathematically. This is what tangles do: *tangles are a precise way of capturing intrinsically fuzzy data* [9]. This applies in particular when the data in question consists of answers to an FFM questionnaire.

The precise but coarser information from the data of [2, 3] that is encoded in tangles will be enough to answer our question (3): we do not need to come up with concrete subsets of Q , possibly but not necessarily the sets Q_t , to prove the extensional cohesion and completeness of the five corresponding traits. And moreover, we obtain additional information on how these traits are related: how they emerge from more comprehensive traits, how they split into subtraits, and all this with quantitative information about the level of precision, or ‘resolution’,⁴ at which this happens.

⁴This will be encoded as the *order* of the corresponding tangles.

2 Data, objectives, and setup

Our aim was to validate, and possibly refine, the Five Factor Model (FFM) [14] by analysing the returns received in [2, 3] for the IPIP questions [12] developed to test the FFM markers developed by Goldberg [11]. These questions were designed specifically to test for the five traits commonly abbreviated as OCEAN, ten questions for each of these traits. We applied tangle-based clustering to the set of these questions based only on the answers received, that is, ignoring the information of which trait each question was designed to test.

2.1 Clusters and traits: the expected outcome

The anticipated result was that, at the most basic level, we would recover five clusters in this set of questions: the clusters corresponding to, in traditional clustering terms, the five subsets of the questions designed to test for the five traits.

However, tangle-based clustering can reveal more detailed information than just detecting a certain number of clusters. It can relate these to each other in structural ways which, when we interpret them here, may reveal how those five traits are related: how some of them can be distinguished less clearly than others, and hence form more comprehensive traits at a coarser level, and perhaps split into more detailed traits at a finer level. These levels come on an absolute scale, the *order* of the tangles found, so that these developments can be compared across the various traits and subtraits as they emerge when the data is analysed at increasing ‘resolution’. In particular, the five target traits, if recovered, would have their place in any such evolutionary hierarchy of traits that may be factually observable in [2] and [3].




Figure 2: Four clusters and three bottlenecks

2.2 Clustering by tangles

Consider a dataset as depicted in Figure 2. Intuitively, there are four clusters: a central one, and three around it. However as soon as we try to come up with a rigorous explicit definition of ‘cluster’ that underpins this intuition, we get into trouble. For example, if we say that a *cluster* is a set of pairwise close points, or a maximal such set, we can easily find pairs of points that invalidate this attempt: points that are close to each other but which lie in different point clouds that form our intuitive clusters, such as the two red dots in the picture, or points that lie in the same intuitive cluster but far from each other, such as the two green dots.

Clustering by tangles sidesteps these issues by taking a radically different, indirect approach. On the right in the picture we can see some natural ways to split the dataset in two, splits that do not cut right through any of the visible clusters but skirt around them. We call such natural partitions the *bottlenecks* of the dataset. For every bottleneck, every visible cluster lies mostly on one of its two sides: otherwise it would not be a bottleneck. In this way, each cluster *orients* every bottleneck towards the side that contains most of it. The figure shows this for the central cluster by little arrows on the three bottlenecks.

This information, the simultaneous orientation of each bottleneck towards one of its sides, is called a *tangle*; it is all that we wish to remember about each of those four clusters. Note that this information is robust in a fuzzy environment such as Figure 2: the way in which each of the four clusters orients those three bottlenecks will not change if we draw the bottlenecks slightly differently or move a few of our data points.

Note, however, that we are still far from a general definition of tangles that can replace the more traditional attempts at defining clusters. This is because our definition of the four tangles in the figure depended on a pre-conceived, if intuitive, traditional notion of cluster: we defined ‘bottlenecks’ as splits of our dataset that leave its clusters largely intact, and we defined ‘tangles’ as simultaneous orientations of all the bottlenecks towards some cluster.

The crucial starting observation in applied tangle theory is that one can define both ‘bottlenecks’ and ‘tangles’ axiomatically, without reference to clusters. One can then either extract point clusters from those tangles – or choose not to, if the information about the dataset sought from determining its clusters can be derived from the tangles directly, as will be the case in our application here.

Although tangles contain less detailed information than traditional point clusters, they offer additional insight. One is that they come at different levels: if we allow only few, particularly narrow, bottlenecks we get few tangles that hang together loosely. If we orient wider bottlenecks too, we get more and denser tangles. These will *refine* the earlier loose ones in that they orient the same bottlenecks as those do, but some wider ones in addition.




Figure 3: A hierarchy of tangles of increasing order $k = 1, 2, 3$

Figure 3, from [9], shows tangles of increasing *order*, or bottleneck width. There are four tangles, ρ_1, \dots, ρ_4 , of order $k = 1$. Adding slightly wider bottlenecks to find the tangles of order 2, we see that the tangle ρ_1 of order 1 remains

unchanged, the tangles ρ_2 and ρ_4 disintegrate, and the tangle ρ_3 splits into the four tangles $\sigma_2, \dots, \sigma_5$. When we allow even wider bottlenecks for tangles of order 3, we find that the tangle $\sigma_1 = \rho_1$ disintegrates, the tangles σ_2 and σ_5 remain unchanged, and the tangles σ_3 and σ_4 of order 2 each split into two tangles of order 3, into τ_2, τ_3 and τ_4, τ_5 respectively.

The lowest order at which a certain tangle first appears, where it arises as one of several tangles of that order which refine the same tangle of lower order, is known as its *complexity*; the highest order up to which this tangle remains unchanged is its *cohesion*, and the difference between the two (plus 1) is its *visibility*. In the example of Figure 3, the tangle σ_1 of order 2 has complexity 1, because it equals (is the unique extension to an order-two tangle of) the tangle ρ_1 of order 1 which refines no tangle of lower order still; it has cohesion 2, because it is still a tangle of order 2 but no longer one of order 3; and so it has visibility $2 = (2 - 1) + 1$. The tangle σ_4 of order 2 has complexity 2, because it is not the only order-two tangle that refines the tangle ρ_3 of order 1 (as $\sigma_2, \sigma_3, \sigma_5$ do too); it also has cohesion only 2, since τ_4 and τ_5 are two tangles of order 3 that both refine it; and so it has visibility $(2 - 2) + 1 = 1$. See [9, Ch.7.4, 14.6] for details.

Another advantage over traditional clustering is that the latter often requires us to specify in advance how many clusters we seek to detect. Tangles reflect whatever clusters the dataset contains. We can set a number of parameters, such as the tangles' required agreement value, which will influence this in general. But once that is done, the number of tangles of any given order found will depend on the data only, not on any choice of ours.

Most importantly, tangles are robust when the data is inherently fuzzy, such as survey data in psychology, sociology, or economics.

2.3 Tangles and traits

When the points in our dataset of Section 2.2 are questions about personality, narrow bottlenecks will divide questions whose answer sets are only loosely correlated, in some well-defined sense of our choice.⁵ As the width of the bottlenecks considered increases, these bottlenecks (and hence their tangles) can distinguish also between more highly correlated questions. A tangle will 'disintegrate' when the bottlenecks allowed become so wide that they cease to be true 'bottlenecks' but dissect the set of questions into singleton subsets (as indicated in Figure 3): when a question is considered as 'highly correlated' only to itself.

In our context it will be possible, in principle,⁶ to interpret each of the tangles of our set Q of 50 questions found in the data of [2, 3] as related to one particular personality aspect. These different aspects will be more fundamentally different when their corresponding tangles have low order, because they will be distinguished by splits that cannot separate highly correlated questions. They will be more specific when their corresponding tangles have higher order.

We can then formalise traits as equivalence classes of tangles of different order, where a tangle τ of order k is *equivalent* to a tangle τ' of order $k' \geq k$

⁵As mentioned earlier, the similarity measure between different questions we shall use will mostly be that of mutual information in the sense of information theory, not statistical correlation.

⁶Such interpretations need not be readily expressible in traditional terms. Indeed, one of the main advantages of tangles is that they enable us to detect hitherto unknown patterns: patterns we do not have to first guess and then confirm, the standard process in the social sciences. But tangles, once found, are immediately accessible to interpretation by experts; see below.

if, essentially, τ' is the unique extension of τ to a tangle of order k' .⁷ A *trait*, then, formalised as such an equivalence class of tangles, can be thought of as a single tangle that ‘is born’ at some order k and ‘dies’ after some maximum order $k' \geq k$.

In Section 2.2 we called the minimum and the maximum k for which ‘a given tangle’ is a k -tangle its ‘complexity’ and ‘cohesion’. We can now define this more precisely: the *complexity* and the *cohesion* of a trait – an equivalence class of tangles – are the smallest and the largest k for which this trait (class) contains a tangle of order k . As before, the *visibility* of a trait is the difference between its cohesion and its complexity plus 1.

Inasmuch as (equivalence classes of) tangles of partitions of a set Q of questions formalise the informal notion of clusters in Q , our formal notions of complexity and cohesion for traits formalise what in Section 1.2 we called the desired external completeness and cohesion of clusters in Q . They do so in a precise, quantitative, way which can be computed from the data of answer sets received for Q . And they give rise to the new formal notion of visibility, which may be an interesting new parameter for traits found in personality tests.

3 Methods 1: Tangle basics

For this section let Q be an arbitrary finite set.

3.1 Partitions

Definition 1 (partition and sides)

Let $A, B \subseteq Q$ be two subsets of Q . If A and B are disjoint and $A \cup B = Q$, then $\{A, B\}$ is called a partition of Q . The sets A and B are called the sides of the partition.

We denote partitions either as a set, like $\{A, B\}$, or by single letters such as r or s . We denote the two sides of a partition s by \vec{s} and \bar{s} . The directions of these arrows are not fixed in advance. For example, if $s = \{A, B\}$ then A could be denoted by either \bar{s} or \vec{s} . But once we have chosen A to be denoted as \vec{s} , say, the other side B of this partition will implicitly have been denoted as \bar{s} .

Given a set S of partitions of Q , we denote the set of all their sides as \vec{S} . Note that for every $\vec{s} \in \vec{S}$ also \bar{s} lies in \vec{S} , since it is also a side of $s \in S$. If s is a partition of Q and $q \in Q$, we write $q(s)$ for the side of s that contains q .

Two sets are commonly called *nested* if one contains the other. We call two partitions *nested* if they have nested sides.

3.2 Order functions and similarity functions

Definition 2 (order of partition)

Given a set S of partitions of Q and an $S \rightarrow \mathbb{R}$ function $s \mapsto |s|$ that we have called an ‘order function’, we call $|s|$ the order of s , for any $s \in S$.

⁷The precise definition is a little narrower: a tangle τ of order k is *equivalent* to a tangle τ' of order $k' \geq k$ if every tangle of order ℓ with $k \leq \ell \leq k'$ that induces τ is in turn induced by τ' ; see [9, Ch.14.6].

All order functions considered in this paper will take non-negative values. They will also all be injective: not by definition, but because we shall tweak the data so as to make them injective. This makes our computations easier, because it spares us having to make arbitrary choices, but it has no material consequences.

Our aim in choosing order functions will be to assign low order to particularly natural partitions of Q , partitions we called ‘bottlenecks’ in Section 2.

One way of defining order functions without reference to any informal notion of cluster is to assign low order to partitions that split few pairs of points we regard as ‘similar’. Functions $Q^2 \rightarrow \mathbb{R}$ intended to express this are called *similarity functions* on Q . One simple choice of an order function on the partitions of Q , then, is to sum up the similarities of the pairs of points split by a partition:

Definition 3 (cut weight order)

Let $w: Q^2 \rightarrow \mathbb{R}$ be a similarity function on Q , and let $\{A, B\}$ be a partition of Q . The cut weight of $\{A, B\}$ is defined as

$$\text{cut}(A, B) := \sum_{a \in A, b \in B} w(a, b).$$

While the cut weight order function is natural, it tends to assign higher order to more balanced partitions, because these split more pairs. For example, if the sides A, B of a partition of Q are about equally large, then the number of pairs (a, b) whose weights $w(a, b)$ are summed in the definition of *cut* is quadratic in $|Q|$, about $(|Q|/2)^2$, while if the partition is unbalanced in that it divides only some k points of Q from the rest, this number is linear in $|Q|$, about $k \cdot |Q|$. Thus, no matter what the weights $w(a, b)$ are, the value of $\text{cut}(A, B)$ tends to be larger when $\{A, B\}$ is more balanced.

This is a problem, since we do not want our particularly natural, and hence low-order, partitions to be cluttered up by unbalanced but unnatural partitions. We therefore try to correct the bias of the order function *cut* towards unbalanced partitions. For example, we often take as the order of $\{A, B\}$ not the sum of the weights $w(a, b)$, as above, but something more like their average:

Definition 4 (ratio cut weight order)

Given a similarity function on Q , the ratio cut weight of a partition $\{A, B\}$ of Q is defined as

$$\text{Rcut}(A, B) := \frac{\text{cut}(A, B)}{|A||B|}|Q| = \text{cut}(A, B) \left(\frac{1}{|A|} + \frac{1}{|B|} \right).$$

There are many other commonly used order functions; see [9]. Our order function of choice will be that of ratio cut weight, based primarily on a similarity function called ‘mutual information’ borrowed from information theory. But we tested also the ‘cosine’ similarity function. Both these are defined in Section 4.2.

3.3 Tangles: basic notion, types, and hierarchy

Let S be a set of partitions of Q . The few graph-theoretic terms used in this section are explained in [8].

Definition 5 (orientation of set of partitions)

A set $\sigma \subseteq \vec{S}$ which, for each $s \in S$, contains exactly one side of s is called an orientation of S . The side of s that lies in σ is denoted by $\sigma(s)$.

Figure 2 indicates an orientation, by arrows, of the set of three ‘bottleneck’ partitions shown in blue. As in that example, orientations of partitions can be defined by subsets of Q that lie mostly on the same side of any such partition. We say that an orientation σ of S is *guided* by any subset X of Q that satisfies

$$|X \cap \vec{s}| > |X \cap \vec{s}|$$

for every $\vec{s} \in \sigma$. An orientation σ of S can have more than one guiding set, and it can also have none.

Example 1 For every $q \in Q$ there is a unique orientation of S that is guided by $\{q\}$. We call it the principal orientation of S defined by q .

As explained in Section 2.2, the idea behind tangles is that while clusters, conceived of as subsets X of Q , are notoriously difficult to pin down precisely, they will – whatever their precise definition if any – guide the same orientations of the ‘bottleneck’ partitions of Q if these are narrow enough. For example, if more than two thirds of X lie on the same side of any $s \in S$, the orientation σ of S guided by X will not change if we modify X a little, perhaps by using a competing definition of point cluster, or because our data is not perfectly reliable.

In this example, the intersection of any three sets in σ contains an element of X , since each of them misses less than a third of X . This motivates the following formal definition of tangles:

Definition 6 (tangle)

An orientation τ of S is a tangle of S if $\vec{r} \cap \vec{s} \cap \vec{t} \neq \emptyset$ for all triples $\vec{r}, \vec{s}, \vec{t}$ of elements of τ .

Note that this definition, while motivated by the properties of a subset X of Q that we might think of as a point cluster, no longer depends on such an informal notion of ‘cluster’: the point it requires to lie in the intersection of any three⁸ elements of a tangle can be any point in Q and does not have to come from any previously defined ‘cluster’ X .

Due to this difference, a set of partitions of a dataset Q can have tangles that are not guided by any obvious subsets of Q : tangles are more general than orientations of set partitions guided by ‘clusters’ in that set, however defined.

In order to reduce the number of tangles of S , we often require that its triples contain not just one but some specified number of elements of Q :

Definition 7 (agreement value)

The agreement value $a(\tau)$ of a tangle τ of S is the minimum number of elements of Q that lie in the intersection of any three sets in τ .

Thus, if every three sides of partitions of S that lie in a given tangle τ have at least n elements of Q in common, the agreement value of τ is at least n .

⁸There is some magic in this number, which is discussed at length in [9, Section 7.3].

Recall that every $q \in Q$ defines a map $s \mapsto q(s)$ from S to \vec{S} , just as every tangle τ of S defines an $S \rightarrow \vec{S}$ map $s \mapsto \tau(s)$. We think of q as being similar to τ if these two functions agree on many $s \in S$:

Definition 8 (point-tangle similarity)

Given a tangle τ of S and any $q \in Q$, we call the number

$$\sigma(q, \tau) := |\{s \in S : q(s) = \tau(s)\}|$$

the similarity between q and τ .

Definition 9 (S_k ; k -tangles of S ; tangles in \vec{S} ; order of a tangle; $\mathcal{T}(S)$)

Let $s \mapsto |s| \in \mathbb{R}$ be an order function on S . For every integer k we write

$$S_k := \{s \in S : |s| < k\}$$

and call the tangles of S_k the k -tangles of S or the tangles of order k in \vec{S} . We write $\mathcal{T}(S)$ for the set of all tangles in \vec{S} , irrespective of their order.

Note that the k -tangles of S are not tangles of S (unless k is very large), since they only orient its subset S_k . For $k < k'$, every k' -tangle τ' of S defines (or *induces*, or *refines*) a k -tangle of S , the tangle $\tau := \tau' \cap \vec{S}_k$ of S_k . Note that $\tau \subseteq \tau'$ in this case. Conversely, given any two tangles $\tau \subseteq \tau'$ in \vec{S} , there will be integers $k \leq k'$ such that τ' is a k' -tangle of S and $\tau = \tau' \cap \vec{S}_k$ a k -tangle.

Definition 10 (tangle equivalence)

A k -tangle τ of S is called equivalent to a k' -tangle τ' of S with $k \leq k'$ if every ℓ -tangle of S with $k \leq \ell \leq k'$ that induces τ is in turn induced by τ' .

It is not hard to prove, but also not entirely obvious, that this is indeed an equivalence relation on the set of tangles in \vec{S} . Note that refinement is well defined on equivalence classes: if a tangle τ refines a tangle σ not equivalent to τ , then any tangle τ' equivalent to τ refines any tangle σ' equivalent to σ .

Definition 11 (complexity, cohesion, visibility)

The complexity and the cohesion of an equivalence class of tangles in \vec{S} are the smallest and the largest order of a tangle in that class, respectively. The visibility of the class is its cohesion minus its complexity plus 1.

Definition 12 (tangle search tree)

The tangle search tree of S is the graph with node set $\mathcal{T}(S)$ in which two nodes ρ, τ are adjacent if $\rho \subsetneq \tau$ and there is no tangle σ in \vec{S} such that $\rho \subsetneq \sigma \subsetneq \tau$. We often write $\mathcal{T}(S)$ also for this tree, not just for its set of nodes.

The tangle search tree is indeed a tree. Since $S_k = \emptyset$ for sufficiently small k , and \emptyset is a k -tangle for this k , the empty set \emptyset is a node of $\mathcal{T}(S)$, which we take as its root. If our order function is injective, as we shall make it throughout, then every child of a tangle in the tangle search tree orients exactly one more partition in S than its parent, so every node has at most two children. The tangle search tree has a non-root level for every integer k that occurs as an order of a partition in S . If these are all the integers between 1 and some K ,

the tangles at level $k = 1, \dots, K$ are the $(k + 1)$ -tangles of S , those that orient every partition in S of order up to k .

Note that two tangles in $\mathcal{T}(S)$ are equivalent if and only if the path that links them in the tangle search tree is vertical⁹ and has no inner vertices at which the tree branches: in graph-theoretical terms, its inner vertices all have degree 2. Contracting the maximal such paths in the tree $\mathcal{T}(S)$ to single edges between their ends, and deleting the root \emptyset if it has only one child, turns $\mathcal{T}(S)$ into a *binary* tree $\hat{\mathcal{T}}(S)$, one whose non-leaf nodes¹⁰ all have exactly two children. These nodes correspond to the equivalence classes of tangles defined earlier.

Definition 13 (tree of traits)

In contexts when the equivalence classes of tangles in \vec{S} (other than possibly $\{\emptyset\}$) are called ‘traits’, the tree on $\hat{\mathcal{T}}(S)$ will be called the tree of traits. We denote this tree by $\hat{\mathcal{T}}(S)$ too.

For every pair of tangles in \vec{S} of which neither induces (i.e., is a superset of) the other there exists an $s \in S$ which these two tangles orient differently. We say that s *distinguishes* these two tangles. If s has lowest order among all the partitions in S that distinguish the two tangles, it distinguishes them *efficiently*.

Definition 14 (efficient distinguisher of incomparable tangles)

If the order function on S is injective, then any $s \in S$ that distinguishes two tangles in \vec{S} efficiently is unique. We call it their efficient distinguisher.

4 Methods 2: Tangle particulars for our analysis

Our aim in this paper is to determine, and analyse, tangles of partitions of the set Q of questions used in the FFM test in [2, 3]. We shall always denote as I the set of individuals on whose answers to these questions we base our tangles. This set I will vary as we compute and test our tangles, as will the set S of partitions of Q whose tangles we compute. But Q will always be the same. Thus, we shall study the tangles in \vec{S} for a suitable set S of partitions of Q , based on the answers to Q given by various sets I of people.

Before we describe more specifically how we chose the parameters giving rise to these tangles, let us formalise the notion of ‘trait’ as envisaged in Section 2.3, now rigorously on the basis of Definition 10:

Definition 15 (tangle trait)

A (tangle) trait is an equivalence class of tangles in \vec{S} other than $\{\emptyset\}$.

Note that $\{\emptyset\}$ may or may not form an equivalence class in $\mathcal{T}(S)$: it does so if and only if the empty tangle \emptyset has two descendents in the tangle search tree.

In order for the term of ‘tangles in \vec{S} ’ to be defined, we had to choose an order function on S . We used the order function of ‘ratio cut weight’ introduced in Section 3.2. This order function requires us to first choose a similarity function on Q . We used two different options for this, in order to test our results for stability. These similarity functions are defined in Section 4.2.

⁹Formally: ascending or descending in the tree order on $\mathcal{T}(S)$ associated with the root \emptyset .

¹⁰The root of a tree never counts as a leaf [8].

Whichever similarity function on Q we worked with, a key element of our work was that these similarities should be extensionally determined: they should be computable from how the individuals in I answered those questions, regardless of any interpretation, or intention for why those questions were included in Q , such as to test for a particular trait. Moreover, our aim was that every individual in I should carry the same weight in determining the similarities between questions in Q . Since different individuals have different habits in how they answer *any* question, such as their different levels of assertiveness or positivity, our aim of giving them equal weight required us to factor out those individual habits. We describe how we did that in Section 4.1. Note that such normalisation must be done with care, since the personal habits we normalise are themselves personality traits like those that Q intends to measure. We also normalised the answers for each particular question to give them mean zero, so that questions suggesting traits generally seen as positive were treated no differently from other questions.

In Section 4.3 we explain how we chose S . In Section 4.4 we present our tangle search algorithm, slightly simplified; the precise version is described in [1].

4.1 Normalising the answers received for Q

In our data, the answers solicited by the questions in Q were on scale from 1 to 5, with 1 indicating strong disagreement and 5 strong agreement. This makes Q into a set of functions $q: I \rightarrow \mathbb{R}$, where $q(i) \in \{1, \dots, 5\}$ records the answer given to q by i . Dually, every $i \in I$ defines a function $i: Q \rightarrow \mathbb{R}$ that sends each question $q \in Q$ to its answer $i(q)$ given by i . Thus, $i(q) = q(i)$ for all $i \in I$ and $q \in Q$ whenever some I has been fixed.

Some individuals tend to agree more than others with questions as they are phrased. This can be seen formally by the fact that the functions $i: Q \rightarrow \mathbb{R}$ have different means μ_i for different $i \in I$. To compensate for this we normalised, for every $i \in I$ separately, their answers to give them mean 0, by subtracting μ_i from their given answers.

Next, we compensated for different levels of assertiveness amongst different people. The assertiveness of person $i \in I$ can be measured by the standard deviation σ_i of the function $i: Q \rightarrow \mathbb{R}$. We normalised these by dividing, for each $i \in I$ separately, their answers (after normalising their mean to 0) by σ_i , to give them all the standard deviation of 1.

In a third step, we did some normalisation between the questions $q: I \rightarrow \mathbb{R}$. Some questions have more people agree with them than other questions, perhaps because they are testing for a trait that is generally seen as positive. We compensated for this by subtracting, separately for each question $q \in Q$, from the answers received for q the median of all these answers.

After this last normalisation, every $q \in Q$ had equally many answers $i(q) \leq 0$ as answers $i(q) \geq 0$, counted over all $i \in I$. However, this last normalisation between the questions may have undone some of our earlier normalisation between individuals. We therefore repeated all three steps cyclically until this process converged to functions that satisfied all three of our normalisation goals: to functions $i: Q \rightarrow \mathbb{R}$ and $q: I \rightarrow \mathbb{R}$ with $i(q) = q(i)$ such that every i has mean 0 and standard deviation 1, and every q has median 0. On the data we investigated, this convergence happened quickly, and only a few rounds of iteration were needed.

From now on we shall use our notation of $q: I \rightarrow \mathbb{R}$ for $q \in Q$, and $i: Q \rightarrow \mathbb{R}$ for $i \in I$ to refer to these normalised functions computed from their original versions in the data of [2, 3].

4.2 Similarity functions

Our first similarity function is based on the notion of discrete entropy in information theory, defined below. More background, examples, and a discussion of how this captures similarities can be found in [9, Ch. 9.5].

Definition 16 (entropy and mutual information)

The (discrete) entropy of a function $f: I \rightarrow X$, where X is any finite set, is the number

$$\mathbb{H}(f) := \sum_{x \in X} (|f^{-1}(x)|/|I|) \cdot \log(|I|/|f^{-1}(x)|).$$

The product $(f \times g): I \rightarrow X \times Y$ of two functions $f: I \rightarrow X$ and $g: I \rightarrow Y$ is defined by setting $(f \times g)(i) := (f(i), g(i))$ for all $i \in I$. The mutual information of $f: I \rightarrow X$ and $g: I \rightarrow Y$, then, is defined as

$$\mathbb{I}(f, g) := \mathbb{H}(f) + \mathbb{H}(g) - \mathbb{H}(f \times g).$$

Definition 17 (entropy similarity)

The entropy similarity of two questions $p, q \in Q$ is their mutual information $\mathbb{I}(p, q)$ as $I \rightarrow \mathbb{R}$ functions.

If we view the set \mathbb{R}^I of all $I \rightarrow \mathbb{R}$ functions as a copy of the vector space \mathbb{R}^n with $n = |I|$, we can compare entropy similarity to other similarity functions between pairs of points of \mathbb{R}^n . We shall do this with the following similarity measure on \mathbb{R}^n , which is based on the cosine of the angle between two vectors:

Definition 18 (cosine similarity)

The cosine similarity of two points $p, q \in \mathbb{R}^n$ is defined as

$$c(p, q) := \frac{p \cdot q}{\|p\| \|q\|}.$$

Here \cdot denotes the dot product in \mathbb{R}^n , while $\|p\|$ and $\|q\|$ denote the lengths of p and q , respectively. The absolute cosine similarity of p and q is defined as

$$ac(p, q) := |c(p, q)|.$$

4.3 Choosing partitions

As $|Q| = 50$, there are about 2^{50} partitions of Q , so we could not include them all in our set S of partitions whose tangles we computed. This set S should contain a diverse selection of low-order partitions, so that our results approximated the tangles of the set of all partitions of Q .

We constructed such a set S in three steps, described below. Performing these steps entailed two choices: of one of two matrices, L or J (see below); and of one of the two similarity measures from Section 4.2. As our order function $s \mapsto |s|$ we always used ‘ratio cut weight’ from Definition 4.

In the first step we computed an orthonormal set of eigenvectors of a certain matrix, which we then converted into partitions of Q . This yielded a set of at most $|Q| = 50$ partitions, typically fewer. We did this for one of two $m \times m$ matrices, L and J , where $m := |Q| = 50$. Let the rows and columns of these matrices be indexed not by numbers but by the elements of Q directly.

The matrix L is commonly known as the ‘combinatorial Laplacian’ of the complete graph on Q with ‘edge weights’ given by a similarity function σ on Q . In row p and column q the matrix L has the entry $-\sigma(p, q)$ if $p \neq q$, and entry $\sum_{q' \in Q \setminus \{q\}} \sigma(q, q')$ if $p = q$. See [9, Ch. 9.2, 10.3] for more about L and the role of its eigenvectors. The similarity functions we used are those from Section 4.2.

L has m real eigenvalues $0 = \lambda_0 \leq \dots \leq \lambda_{m-1}$, counted with multiplicities. If Q had a partition $\{A, B\}$ such that $\sigma(a, b) = 0$ for all $a \in A$ and $b \in B$, we would add it to S straight away and start again with Q replaced by A and, separately, by B . This did not in fact happen. The first eigenvalue $\lambda_0 = 0$ of L therefore always had multiplicity 1 (as is well known), so $0 < \lambda_1 \leq \dots \leq \lambda_{m-1}$. We then picked eigenvectors v_1, v_2, \dots for the eigenvalues $\lambda_1, \lambda_2, \dots$ one by one, all orthogonal to both the eigenvector $\mathbb{1}$ for λ_0 and to each other.

The matrix J has entries $-(p \cdot q)$ in row p and column q , for all $p, q \in Q$, where ‘ \cdot ’ once more denotes the dot product in \mathbb{R}^n , as in Definition 18. The matrix J has m real eigenvalues $\lambda_1 \leq \dots \leq \lambda_m \leq 0$, counted with multiplicities, whose corresponding orthogonal eigenvectors v_1, v_2, \dots are the *principal components* of the ‘variables’ in Q for the data given by $I \subseteq \mathbb{R}^Q$. Note, however, that we are using these principal components not as discussed, and criticised, in Section 1.3 (see Figure 1), but in the dual (and opposite) role: rather than proposing them as centre points of clusters in Q , we use them to define natural partitions of Q .

The eigenvectors v_1, v_2, \dots of either L or J were then turned into partitions of Q . Given $v_i : Q \rightarrow \mathbb{R}$, we chose the partition

$$s_i := \{\{q \in Q : v_i(q) < 0\}, \{q \in Q : v_i(q) \geq 0\}\}.$$

Let us call the set S of these partitions S^0 .

In the second step we iteratively enlarged our current set S of partitions by including, for all the pairs of partitions $\{A, B\}, \{C, D\}$ already in S , also their four *corners* in \vec{S} . These are the sets

$$A \cap C, B \cap C, A \cap D, B \cap D.$$

The partitions of Q which these sets form with their complements in Q are the four *corner partitions* of $\{A, B\}$ and $\{C, D\}$. Tangle algorithms, see [1], work better when corners of partitions in S are also in \vec{S} if they have low order. Let us call the resulting set S of partitions $S^1 \supseteq S^0$.

We then performed the third step on the partitions in S^1 . For each $s \in S^1$ we tried to find a single $q \in Q$ such that moving q across the partition s would reduce the order of s . If we were able to find such elements q , we moved the q across s for which this decreased the order of s the most, adding the modified partition to S . We then repeated this with the modified partition, iteratively until no such q could be found. We added all these modified partitions to S^1 . We took the resulting set of partitions of Q as the set S whose tangles we would finally compute.

Let us take a moment to motivate the above three steps; more background can be found in [9]. Partitions obtained from orthonormal eigenvectors of L

and J are particularly useful starting partitions for our construction: they are more balanced than other partitions of low order, in that they divide Q into sides of roughly equal size, and among such balanced partitions they have particularly low order.

Adding corners helps with weeding out ‘fake’ tangles: tangles of only some of the partitions of Q that do not extend to tangles of all the partitions – those we are trying to approximate.

Partitions whose orders are local minima with respect to moving single elements across include our target partitions, the efficient distinguishers of tangles of agreement value at least 2 of *all* the partitions of Q . Since tangles of agreement 2 cannot orient two partitions differently if these differ by a single $q \in Q$ (ie, both towards the moving q or both away from it), the way in which our tangles orient the local minima partitions determines by how they orient our earlier set S^1 . Thus, only tangles of S^1 that extend to tangles of the increased final set S were considered further. This weeds out tangles of S^1 that satisfy the tangle axioms ‘for the wrong reason’ that S^1 did not include enough low-order partitions to put the tangle axioms to a sufficiently rigorous test.

4.4 Finding the tangles

Having constructed one of the sets S of partitions of Q described in Section 4.3, we constructed all the tangles in \vec{S} of agreement value at least $a = 2$, as follows.

We began by sorting the partitions in S by order, as $|s_0| < |s_1| < \dots$ say. Then for $k = 0, 1, \dots$ in turn we built the set \mathcal{T}_k of all tangles of S_k , where

$$S_k := \{s_i \mid i < k\}.$$

Note that this coincides with our Definition 9 of S_k if $|s_i| = i$ for all $i < k$, which we shall assume to simplify our exposition.¹¹ Note that $\mathcal{T}_0 = \{\emptyset\}$, since the empty tangle is the unique tangle of $S_0 = \emptyset$.

Having constructed \mathcal{T}_k for some $k \geq 0$, we let \mathcal{T}_{k+1} consist of all the tangles of the form $\tau \cup \{\vec{s}_k\}$ or $\tau \cup \{\hat{s}_k\}$ with $\tau \in \mathcal{T}_k$. For any given $\tau \in \mathcal{T}_k$, sometimes both these extensions of τ were tangles, or just one of them, or neither: this depended on whether \vec{s}_k or \hat{s}_k formed a triple with two elements of τ whose intersection consisted of fewer than a elements of Q , in which case adding it to τ would not yield a tangle of S_{k+1} of agreement value at least a .

This process ended only once the last partition in S , its element of largest order, was examined. Then $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2 \cup \dots$ was our set of all tangles in \vec{S} .

5 Results

In this section we present our findings. In Section 5.1 we specify the datasets in which we looked for tangles. We state which parameters we used to compute these tangles (see Sections 4.2 and 4.3), and note how many traits they yielded (Definition 15). We finally define what we mean when we use the ‘big five’ labels of O, C, E, A, N to name some of these traits.

¹¹The actual value of an order never matters for tangles: all that matters is the linear ordering which the order function used imposes on S .

In Section 5.2 we present the trees of traits we found (Definitions 12 and 13), analyse their structure, and indicate what this means from an interpretation point of view.

Interpretation of our results is addressed more comprehensively in Section 5.3. While we have to leave interpretation as such to experts in psychology, we say how exactly such experts can read off the facts they need from our structural analysis in Section 5.2 and supplementary data listed in the Appendix. One of the fortes of tangle clustering is that it condenses the most crucial information into a very small set of facts that can be stated explicitly at the interpretation level. This is done in the Appendix, in Section 7.3.

In Section 5.4 we list, for each of the traits found, which of the 50 questions in Q represent that trait best. In particular, we do this for the *OCEAN* traits, but also for the other traits we found. This also has a direct impact on interpretation, which is discussed in Section 5.4 too.

Finally, we address the question of how reliable our results are. This has two aspects, which we call ‘robustness of methods’ and ‘stability of findings’. By *robustness* of our methods we mean invariance of our findings under small changes of the input, while we keep those methods unchanged. By *stability* we mean the converse: the invariance of our findings under small changes of the methods, performed on the same input. These are discussed in Sections 5.5 and 5.6, respectively.

5.1 The tangle traits we found

We computed, separately for the two datasets in [2] and [3], the tangles in \vec{S} of agreement at least 2. We focus here on reporting our findings for the larger of those two datasets, the data from [3]. We also report more briefly what tangles we found in [2], and discuss the differences (Section 5.2) and their significance. Both studies were conducted with the same questionnaire Q of 50 questions. When computing their tangles we used identical parameters, specified below.

Let us remark right away that there will indeed be differences between our findings for the two studies of [2] and [3]. This motivated us to investigate to what extent these differences were due to substantial differences in the data, or perhaps to a lack of robustness of our methods. To answer this we re-applied our algorithm, with the same parameters as before, to disjoint subsets of the participants in the larger study of [3]. As we shall see in Section 5.5, the tangles we found on these subsets were very similar to those we found on the entire dataset of [3]. Hence they all differed, and in similar ways, from the tangles we found in [2].

We have thus been able not only to find reproducible tangle-based traits in the studies of [2] and [3]. We can also say with some confidence which aspects of the relationship between those traits, in each of the two datasets, are stable across those two studies (which were four years apart and conducted with different target sets of participants), and which were specific to one study or the other.

The tangles analysed below are based on mutual information as the similarity function (Definition 17), and on its Laplacian L as the matrix (Section 4.3). In Section 5.6 we shall compare these tangles with the tangles we found for the cosine similarity function (Definition 18), again with the matrix L . Since these two similarity functions measure different things, we expected them to return

different tangles, which was indeed the case. As both mutual information and cosine similarity are accessible to interpretation, the (small but interesting) difference between the tangles found can also be interpreted. We must leave this to the experts in psychology, but would like to point out this possibility.

We further ran tests to see whether our choice of the matrix L rather than J affected the tangles we found. Any difference here would be harder to interpret. However, we found no such difference; see Section 5.6.

In each of the datasets from [2] and [3] we found several thousand tangles. In the case of [2] these combined to eleven traits, in the case of [3] to fifteen (see Definition 15). We denoted some of these traits by the letters O, C, E, A, N , according to the following rule.

Let us first denote by the letters O, C, E, A, N the five subsets of Q designed to test for the trait indicated by the respective letter. Consider any one of these subsets, O say. If it so happened that O , but none of the other four sets $C, E, A, N \subseteq Q$, guides the lowest-order tangle in some trait T , and if no trait refined by T has this property too, we denoted T as O . Similarly, we used each of the letters C, E, A, N to denote the set of questions in Q it indicates, and then by implication also the trait corresponding to that set, as defined above for O .

The tree of the fifteen traits we found in [3] is depicted in Figure 4. The tree of the eleven traits we found in [2] is depicted in Figure 5.

5.2 Structure: how different traits are related

The trees of traits, shown in Figure 4 for [3] and in Figure 5 for [2], display a hierarchy between those traits. This has two aspects. The first aspect is the structure of the tree: this tells us how the traits are related to each other, i.e., which of them refines which. The second aspect is at which level the various traits occur: this tells us their complexity, cohesion, and visibility (Definition 11).




Figure 4: The tree of traits found in the larger study of [3]

Let us look at the first of these aspects. Traits sitting above other traits in the figure can be viewed as a common generalisation of those. The trait T_8 in Figure 4, for example, is shown above the traits A and E , so it is something like a common generalisation of the two. Indeed, the partitions in S oriented by the tangles in T_8 are oriented in the same way by the tangles in A and by the tangles in E . Those latter tangles, however, may disagree on their orientation of partitions whose order is too large for them to be oriented by the tangles in T_8 . In particular, they disagree on their unique efficient distinguisher, the partition s of lowest order which they orient differently. Since A and E both refine T_8 , the order of s is at least that of the largest tangle in T_8 .

By contrast, the trait T_2 is an even broader common generalisation of A and E , one that also encompasses O – but neither C nor N . The traits A and E are distinguished from O by the same efficient distinguisher s , whose order is at least that of the largest tangle in T_2 . This s is also the efficient distinguisher of T_8 and O , but it does not distinguish A from E .

Let us now turn to the second aspect of Figure 4, the levels at which the various traits are shown. The vertical scale in the figure indicates order: of the tangles representing the traits depicted, and of the partitions of Q in S that distinguish them.¹²

The trait T_6 , for example, is entered in the diagram at level 59. This means that the largest order of a tangle in this trait is 59. The lowest order of a tangle in T_6 can also be read out of the picture: it is the lowest order at which the highest-order tangle in the parent of T_6 , trait O , splits into tangles belonging to its two children, the traits T_4 and T_6 . As O is entered at level 43, this order of the efficient distinguisher between the tangles in T_4 and those in T_6 is at least 43.¹³ If it is exactly 43, the trait T_6 has complexity 43, cohesion 59, and visibility $59 - 43 + 1 = 17$. Its visibility is indicated roughly by the length of the line that joins T_6 to O in the figure. This is higher than the visibility of the trait T_5 , say, but lower than that of the traits C and N .

A particularly interesting aspect of the vertical distribution of the traits in Figure 4 is that some of the *OCEAN* traits are born only when others have already died: there is no order of separations in S at which they overlap. Traits A and E , for example, are born just below level 940, when their parent trait T_8 splits. At this time the traits C and O are already dead. Trait E then outlasts trait A by a considerable margin, as does trait N , which was born even before A and E .

In summary, while the ‘big five traits’ targeted by Q do show up as tangle traits, they do so at very different times, or levels of ‘resolution’ of the lens – order – through which we view our tangles and the traits to which they combine.

Let us now turn to the smaller study of [2]. Its tree of traits is depicted in Figure 5. For any comparison with the tree of traits for the larger study [3], note that while the respective traits O, C, E, A, N were named according to the same rule, applied independently to the two datasets, the enumerations T_1, T_2, \dots and T'_1, T'_2, \dots of the other traits are arbitrary and not related.

¹²But note that Figure 4 is not drawn to scale. The lengths of the lines, which we shall see indicate the visibility of traits, thus have to be compared with caution: lines further up span a smaller vertical difference than their lengths appear to indicate.

¹³More precisely, it is the lowest order of a separation s not oriented by the largest tangle in O . Since that tangle has order 43, it orients precisely the separations in S that have order less than 43. So s is the element of S of the smallest order that is at least 43.




Figure 5: The tree of traits found in the smaller study of [2]

The most important result on which the two studies agree is that both feature traits that merit being called O, C, E, A, N according to our rule in Section 5.1: their lowest-order tangle is guided by the subset of Q specified by the letter that names them, but by no other such subset, and this applies to no smaller trait.

One structural similarity between [2] and [3] is that the traits A and E are direct siblings: in both studies they have a common parent (the trait T_8 in [2] and the trait T'_4 in [2]). And in both datasets the trait A then splits into two more traits, while E does not have distinct refinements.

A closer look also reveals some interesting differences between [2] and [3].

A particularly striking feature of [2] is that the *OCEAN* traits are, successively, split off one line of refining non-*OCEAN* traits, the traits $T'_1 > T'_2 > T'_3 > T'_4$ (where $>$ indicates refinement as shown in the figure). In [3], by contrast, the *OCEAN* traits form two groups, C, N versus A, E, O , where A and E form a subgroup of the latter group A, E, O . These two major groups are in fact separated by *two* levels of more general traits: by the trait T_1 (which the two groups refine in different ways), but also by its two children, the traits T_2 and T_3 . Given the fact that T_2 and T_3 are themselves traits, this seems quite remarkable – and potentially interesting from an interpretation point of view. While we limit ourselves to structural comparisons here, we shall see in Section 5.3 that it is possible to use the information from the Appendix to write down how these traits orient their distinguishers. This provides experts in psychology with direct interpretable information.

In [3], the traits C and N have a direct common generalisation, the trait T_3 . In [2], they also have a common generalisation that does not also generalise any of the other *OCEAN* traits, the trait T'_2 . This, however, is not a direct generalisation just of the two traits C and N : there is an intermediate trait T'_3 , which is already a generalisation of N (though not of C), and the common generalisation T'_2 of C and N also generalises T'_3 . The study of [2] thus sees more structure than [3] in the relationship between the traits C and N .

Another interesting difference between [2] and [3] is that, in [2], the trait O does not further refine, while in [3] it splits into traits T_6 and T_4 , which in turn

splits into T_5 and T_7 . Thus, [3] reveals substantial structure inside the trait O , which [2] did not see.

5.3 Interpretation

Tangles, unlike clusters found by some more traditional clustering algorithms, are in principle open to interpretation: every tangle is a set of oriented partitions of Q , and these are known explicitly. An expert psychologist can draw interpretative conclusions from the way our tangle traits orient these partitions.

In our particular case, we can even say something interpretative about the tangle traits we found without any expert knowledge of psychology. This is because we were able to relate some of these traits, O , C , E , A and N , to five groups of questions in Q , as indicated in Section 5.1. This already yields interpretative information for these five tangle traits themselves, if one considers these five sets of questions as understood at the interpretation level. And this is a reasonable assumption: recall that these five sets of questions were *designed* to flesh out some *interpretation* of the Five Factors, see Section 1.1 and the Introduction. Our five tangle traits named O, C, E, A, N therefore inherit those interpretations in the way made precise in Section 5.1.

So what about interpreting the other ten of our fifteen tangle traits in [3], say? Our structural tangle analysis in Section 5.2 yields that some of those other ten traits are common generalisations of specific subsets of the *OCEAN* traits, and thereby inherit their interpretations. But even those traits that are not common generalisations of *OCEAN* traits can be interpreted succinctly, as follows.

Each of the 15 traits in [3] is determined uniquely by which of the efficient tangle distinguishers it orients, and how. While there are seven efficient distinguishers in total, one for every branching node of the tree of traits (Figure 4), the tangles in a given trait orient only those distinguishers whose order is smaller than their complexity. But different traits are distinguished by at least one of these distinguishers, so each trait is identified uniquely by how it orients the efficient distinguishers that it does orient.

And every oriented distinguisher is directly accessible to interpretation: it is one concrete subset of Q . For the efficient trait distinguishers, these subsets are listed explicitly in the Appendix, in Section 7.3. From these lists one can combine, for each of our 15 traits T other than the all-encompassing T_1 , how the tangles in T orient the distinguishers s which they orient at all, those of order less than the complexity of T . Conversely, given any two traits T and T' , Figure 4 tells us which $s \in S$ is their unique efficient distinguisher, and its list in Section 7.3 then tells us how T and T' orient s towards its opposite sides and which concrete subsets of Q those sides are.

A more direct, and simpler, route to interpreting our tangle traits is to identify, for each given trait, the questions in Q that represent it best – in a precise, formal sense. We discuss this next.

5.4 Which questions represent the traits best?

Recall that our questionnaire Q was designed for the purpose of testing the five *OCEAN* traits, originally in the pre-formal sense of ‘trait’ discussed and criticised in Section 1. In Section 5.1 we gave a formal definition of when we denote any of our tangle traits (which, in contrast, are a formal concept)

by those *OCEAN* labels. For these five traits, but also for all the other six (respectively, ten) traits we found in [2] and [3], our tangle framework allows us to determine explicitly which of the 50 questions in Q are best suited to test for that particular trait. This works as follows.

Every tangle in \vec{S} , by definition, orients some of the partitions of Q towards one of their two sides. Every element of Q does that too: we can think of it as orienting every partition of Q towards the side that contains it. We can therefore compare the questions q in Q directly with the tangles τ in \vec{S} , simply by counting the number $\sigma(q, \tau)$ of partitions they orient in the same way (Definition 8).

In order to assess how well a question q represents an entire trait T , now, rather than a given tangle, it would be natural to simply use $\sigma(q, \tau)$ for some $\tau \in T$. But which $\tau \in T$ should we choose for this comparison? Fortunately, it does not matter: as long as we compare q and τ only on partitions $s \in S$ that really matter, the result will be the same.

Indeed, given $q \in Q$ and a tangle trait T , write k^- for the complexity of T and k^+ for its cohesion. Then the tangles in T agree on all the partitions in

$$S_T := \{s \in S : k^- \leq |s| + 1 \leq k^+\}$$

which they orient, and they disagree, with any tangle that can be distinguished from them at all, on all partitions in S_T that such a tangle also orients. The partitions in S_T , therefore, are those on which the trait T comes into its own: those on which its tangles differ from tangles in all traits that are essentially different from T .

We now define $T(s) := \tau(s)$ for all $s \in S_T$, where τ is any tangle in T that orients s . (We do not have to specify which τ , since they all agree on s .) Then

$$\sigma(q, T) := |\{s \in S_T : q(s) = T(s)\}|$$

measures what we want: the similarity between q and the tangles in T on the partitions that reflect T best.

Section 7.2 lists, for each of the traits found in [2] and [3], the three questions q in Q that represent this trait T best as measured by $\sigma(q, T)$.¹⁴ These tables contain a wealth of information that is directly accessible to interpretation. Not only are the individual traits each represented by three concrete questions; it is also instructive to compare how the structural relationship between the traits, as expressed by the trees of traits shown in Figures 4 and 5, compares with the sets of questions that represent them.

For example, Figure 4 shows that the trait O in [3] splits into the subtraits T_4 and T_6 . These are represented in Section 7.2 by disjoint sets of questions: T_4 by O_4, O_5, O_{10} and T_6 by O_1, O_7, O_8 , which helps us distinguish them also at the interpretation level. The subtrait T_4 then splits again, into T_5 and T_7 . But these are both represented best by O_5, O_6, O_{10} , the same set of questions: their difference as subtraits of T_4 is more subtle than what their representing questions indicate. But this difference is captured succinctly by one explicitly known partition: their efficient distinguisher. This is listed in Section 7.3, associated with T_4 .

We must leave any further analysis to the experts in psychology, contenting ourselves with pointing out how the information on this can be based is made explicit in our tangle data.

¹⁴The questions themselves are listed in Section 7.1.

5.5 Robustness of our methods

Recall that by *robustness* of our methods we mean the invariance of our findings under small changes to the input, while we keep those methods unchanged. To test this, we repeated our tangle analysis of the large study [3] described earlier in this section on ten subsets of its participants.

We first divided its entire set of about a million participants into 5 disjoint subsets of 200.000 participants and analysed their tangles. We then chose five further random subsets of about 10.000 participants each, and again analysed their tangles.

The trees of traits we found in the five large subsets were very similar to that of the entire dataset (Figure 4). Two of them were in fact identical to it. Another two were identical to it except that trait A lacked the two children it had in Figure 4, denoted there as T_9 and T_{10} . The fifth tree of tangles is shown in Figure 6. It differs from the other four in that the two long branches hanging off T_2 and T_3 in Figure 4 are swapped. So there is now a joint trait generalising C, N, O but not A, E , whereas in Figure 4 there was one generalising A, E, O but not C, N . The *OCEAN* traits themselves are independent, as before.




Figure 6: The only different tree of traits on a large subset of [3]

The trees of traits we found on the five smaller subsets of [3] are depicted in Figure 7. They show a little more variation – not unexpected for smaller datasets. The first is again identical to that of Figure 4. The second is identical to it except that A has no children, as earlier. The other three are variants of the tree in Figure 6.¹⁵

Thus, the traits we found in the ten subsets of [3], as well as their tree structure, are all very similar to those for the entire dataset (Figure 4). In each of them there are five traits that merit being denoted as O, C, E, A, N by our rule in Section 5.1. Their relative structure is also the same as before: traits A and E combine to a more comprehensive trait not generalising any other *OCEAN* traits. And so do C and N , except in the last tree of Figure 7. In six of the ten trees, O combines with the supertrait of A and E to a trait not

¹⁵Note that for the tree structure it is irrelevant which two branches emanating from a node is drawn left and which right. The branch containing C , N and O in last tree in Figure 7 is in fact identical to that in the tree before, except that the parent trait of C and N has split into two traits.




Figure 7: The trees of traits for the five small subsets of [3]

generalising C or N , as in Figure 4, while in the other four it combines with C and N to a trait not generalising A or E .

Overall, our tangle analysis of the ten subsets of [3] yields essentially the same results as that of the entire dataset. This means that our tangle analysis is robust in the sense defined earlier: the essential aspects of its results are invariant under changes of data that are not due to other factors, such as a different time when the study was conducted, or a different sample of participants.

If we take this as confirmation of the robustness our methods also with respect to our analysis of the data from [2], the implication is that the difference our analysis in Section 5.2 found between the studies of [2] and [3] reflects genuine difference in the data. This, in turn, may have interesting implications for their interpretation, since the differences between [2] and [3], as well as what they have in common (such as being based on the same questionnaire Q), are known explicitly.

5.6 Stability

Recall that by *stability* of our findings we mean their invariance under small changes of our methods, performed on the same input.

In our methods we initially made some reasonable choices that we kept fixed throughout our tangle analysis. Those included the use of the ratio cut weight order function (based, however, on a similarity function which we did vary), to require our tangles to have a minimum agreement of 2, and the various normalisations described in Section 4.1.

However we tested variation of two parameters: of our choice of the similarity function used by the ratio cut weight order function (see Section 4.2), and of our choice of the matrix (L or J , see Section 4.3) whose eigenvectors provided the initial set of partitions of Q from which we then built our final set S of partitions.

We found that replacing L with J made no structural difference to the traits we found and their relationships. The only difference lay in the complexity and cohesion of the traits, but with very little resulting difference for their visibilities. So our findings were perfectly stable in this respect.

This is not unremarkable, and we interpret it as an indication of reliability for the tangle method in clustering in general. Its fundamental idea is to ignore just the right amount of detail in a fuzzy environment to identify the ‘true’ clusters and their mutual relationship. Both these turned out to be identical in our tangle analysis even when its input data, the starting set S of partitions of Q , was very different as we obtained it from the eigenvectors of the matrices L and J , respectively.

Let us now report briefly what we found when we based the ratio cut weight order, which we used throughout, on cosine similarity (Definition 18) rather than on entropy similarity (Definition 17). The trees of traits we found in [2] and [3], respectively, are shown in Figure 8.




Figure 8: The trees of traits computed with cosine similarity. The tree for the large study [3] is shown on the left, that for the small study [2] on the right.

Once more, we used the labels of O, C, E, A, N for traits that corresponded to the subsets of Q bearing these names, according to the rule from Section 5.1. This time, no trait found in the smaller study of [2] could be named as A . But it had traits corresponding to O, C, E and N , while in the larger study of [3] all five *OCEAN* traits occurred.

Unlike for the traits based on entropy similarity, the trees of traits were now strikingly similar for the smaller and the larger study, [2] and [3]. The only structural difference between the two is that the sibling of E is A in [3], while in [2] it is an unmarked trait. Both A and this unmarked trait have low visibility in their respective tree. More detailed analysis shows that the unmarked trait is in fact quite close to meeting our standards for being named as A , whereas trait A in [3] meets these only just.

An explicit look at the partitions of Q we collected into S using cosine rather than entropy similarity indicates that they are rougher: the entropy-based low-order partitions conform more readily to what, from an interpretation point of view, we would expect to be relevant partitions of Q . But this roughness is smoothed again by the tangles they form, and even more so by the traits to which these combine: we get a similar number of traits as before, the expected *OCEAN* traits are among them, and there are further traits that refine or generalise these *OCEAN* traits, thereby endowing them with additional structure.

Finally, there is a striking structural similarity between the two trees of traits based on cosine similarity (Figure 8) with the tree of traits based on entropy similarity (Figure 5) for the smaller study [2]. Indeed, in structural terms the tree of traits shown on the left in Figure 8 (which is based on the

larger study [3]) differs from that of Figure 5 (for the *smaller* study [2]) only in that O splits into two subtraits in the former but not in the latter, while A splits into two subtraits in the latter but not in the former. Otherwise, the two trees are structurally identical.¹⁶ The similarity between these two trees seems remarkable because they are based on different studies, [2] and [3], which in our earlier analysis we found to exhibit more distinct trait structures (which, moreover, we found to be robust, see Section 5.5).

We cannot explain this phenomenon. The fact that the trees of traits are more similar for [2] and [3] under cosine similarity (Figure 8) than they are under entropy similarity (Figures 4 and 5) can be interpreted as an indication that tangle analysis based on cosine similarity is a blunter tool than tangle analysis based on entropy similarity. But the fact that the more sensitive tool yields, for one of these studies [2], results very similar to those returned by the blunter tool for both studies, while it returns reliably different results for the other study [3], is harder to interpret. It might be seen as an indication of an anomaly in the design of [3], which analysis based on cosine similarity is (perhaps beneficially) too blunt to pick up, or it may be an artefact that indicates nothing significant.

6 Conclusion

The problem this paper aimed to address was to bridge the gap between the mathematical notion of a *factor*, in the factor analysis performed on personality data from Cattell [5] and Fiske [10] by Tupes and Christal [16], and interpretations of these (‘big five’) factors on which questionnaires for personality assessment are currently based. The gap stems from the fact that factors, by definition, are abstract linear combinations of answers obtained for the studies analysed, while questionnaires testing for such factors are formulated based not on those linear combinations of the original questions but on new questions designed to test for an *interpretation* of those factors, conceived of as a ‘trait’.

To do this, we developed quantitative criteria for groups of questions to merit being thought of as testing for any common trait at all. This required us to provide a formal definition of ‘trait’, against which such questionnaires could then be tested.

We first proposed, in general clustering terms, two criteria for a group of questions to achieve this: as ‘extensional cohesion’ and ‘extensional completeness’. Our use of the novel clustering method of *tangle analysis* then enabled us to propose a rigorous notion of ‘trait’, on which we based our further investigations: that of an equivalence class of tangles under some generic equivalence relation from tangle theory, whose origin lies in the tangle analysis of visual data. This formal notion of (tangle) trait satisfies our earlier requirements of extensional cohesion and completion; indeed these are built in to the very notion of a tangle and of tangle equivalence. But it goes further, in that it identifies the structural essence of clusters in a fuzzy environment, while filtering out unnecessarily detailed but less essential data as ‘noise’.

We analysed the data returned by the separate studies of [2] and [3], both for the same questionnaire Q designed to test for the ‘big five’ OCEAN traits. Our

¹⁶In mathematical terms: once we delete the two pairs of children named above, there exists a graph isomorphism between the two trees that fixes their OCEAN traits, i.e., which maps corresponding OCEAN traits to each other.

analysis confirms that the five groups of questions in Q designed to test for those five ‘traits’ do indeed each test for something that satisfies our formal notion of a tangle trait, with some reservations in the case of trait A (‘agreeableness’). But in addition to the five *OCEAN* traits we also found some others, both common generalisations and refinements of *OCEAN* traits: six in the smaller study of [2], and ten in the larger study of [3].

These traits, including the five *OCEAN* traits, are not equally cohesive and complete, and our analysis quantifies to what extent they are. *They are also interrelated in ways that were not, to the best of our knowledge, known before.* These relationships are exhibited by *trees of traits* (Section 5.2).

For each of the five *OCEAN* traits, but also for the other traits we found, our analysis can identify the questions in Q that represent this trait best, in a rigorous, quantitative sense. In Section 7.2 we list the top three questions for each trait. This is one aspect in which the tangle traits we found are directly interpretable – unlike the factors found in larger studies by factor analysis.

Further interpretability of the traits we found can be read off explicitly from their *efficient distinguishers*: five in [2], and seven in [3]. Each of these is one concrete partition of Q into two sets of questions. These are listed explicitly in Section 7.3, and are thus directly accessible to expert interpretation.

Finally, we tested the robustness of our methods (Section 5.5) for the larger study of [3], as well as the stability of our findings under changes to our tangle parameters (Section 5.6). We found that both the traits found and their relationships as expressed in the trees of traits were essentially invariant under changes of input data taken disjointly and independently from the same study. We found that our results were perfectly stable under changes to the parameters L or J that determined the partitions of Q for our tangles, and largely stable when we replaced the entropy-based similarity measure we used with the cruder measure of cosine similarity. Both these are state-of-the-art similarity measures in information theory and data science.

In summary, we found that the method of tangle-based clustering can be used reliably to find previously unknown traits in existing data obtained for FFM-based personality tests, and to relate known traits to each other, all in a rigorously defined quantitative sense.

References

- [1] Tangles: e-book and software. Home page at tangles-book.com.
- [2] Answers to the Big Five Personality Test, constructed with items from the International Personality Item Pool. Item 5/18/2014 on <http://openpsychometrics.org/tests/IPIP-BFFM/>, May 2014.
- [3] Answers to the IPIP Big Five Factor Markers. Item 11/8/2018 on https://openpsychometrics.org/_rawdata/, November 2018.
- [4] T. Buchanan, J. Johnson, and L. R. Goldberg. Implementing a five-factor personality inventory for use on the internet. *Europ. J. Psych. Assessment*, 21(2), 2005.
- [5] R. Cattell. The description of personality: Principles and findings in a factor analysis. *Amer. J. Psychol.*, 58:69–90, 1945.

- [6] P. T. Costa and R. R. McCrae. *Revised NEO Personality Inventory (NEO PI-R) and NEO Five-Factor Inventory (NEO-FFI)*. Psychological Assessment Resources, 1992.
- [7] R. Diestel. Tangles: a new paradigm for clusters and types. arXiv:2006.01830.
- [8] R. Diestel. *Graph Theory* (6th edition). Springer-Verlag, 2024. Electronic edition available at <http://diestel-graph-theory.com/>.
- [9] R. Diestel. *Tangles: A structural approach to artificial intelligence in the empirical sciences*. Cambridge University Press, 2024. e-book: tangles-book.com/book/.
- [10] D. Fiske. Consistency of the factorial structures of personality ratings from different sources. *J. Abnorm. Soc. Psychol.*, 44(3):329–344, 1949.
- [11] L. R. Goldberg. The development of markers for the Big-Five factor structure. *Psychological Assessment*, 4:26–42, 1992.
- [12] International Personality Item Pool. Administering IPIP measures, with a 50-item sample questionnaire. https://ipip.ori.org/new_ipip-50-item-scale.htm.
- [13] R. McCrae and P. Costa. A five-factor theory of personality. In L. Pervin and O. John, editors, *Handbook of personality: theory and research*, pages 139–153. Guilford Press, 1999.
- [14] R. McCrae and O. John. An introduction to the five-factor model and its applications. *J. Personality*, 60(2), 1992.
- [15] Open-source Psychometrics Project. Big five personality test. <http://openpsychometrics.org/tests/IPIP-BFFM/>.
- [16] E. Tupes and R. Christal. Recurrent personality factors based on trait ratings. *USAF ASD Tech. Rep. 61 97*, 1961.

7 Appendix

7.1 The questionnaire Q

Here is a list of the 50 questions in the questionnaire Q used in [2] and [3]:

Openness

- O_1 I have a rich vocabulary.
- O_2 I have difficulty understanding abstract ideas.
- O_3 I have a vivid imagination.
- O_4 I am not interested in abstract ideas.
- O_5 I have excellent ideas.
- O_6 I do not have a good imagination.

O_7 I am quick to understand things.

O_8 I use difficult words.

O_9 I spend time reflecting on things.

O_{10} I am full of ideas.

Conscientiousness

C_1 I am always prepared.

C_2 I leave my belongings around.

C_3 I pay attention to details.

C_4 I make a mess of things.

C_5 I get chores done right away.

C_6 I often forget to put things back in their proper place.

C_7 I like order.

C_8 I shirk my duties.

C_9 I follow a schedule.

C_{10} I am exacting in my work.

Extraversion

E_1 I am the life of the party.

E_2 I don't talk a lot.

E_3 I feel comfortable around people.

E_4 I keep in the background.

E_5 I start conversations.

E_6 I have little to say.

E_7 I talk to a lot of different people at parties.

E_8 I don't like to draw attention to myself.

E_9 I don't mind being the center of attention.

E_{10} I am quiet around strangers.

Agreeableness

A_1 I feel little concern for others.

A_2 I am interested in people.

A_3 I insult people.

- A_4 I sympathize with others' feelings.
 A_5 I am not interested in other people's problems.
 A_6 I have a soft heart.
 A_7 I am not really interested in others.
 A_8 I take time out for others.
 A_9 I feel others' emotions.
 A_{10} I make people feel at ease.

Neuroticism

- N_1 I get stressed out easily.
 N_2 I am relaxed most of the time.
 N_3 I worry about things.
 N_4 I seldom feel blue.
 N_5 I am easily disturbed.
 N_6 I get upset easily.
 N_7 I change my mood a lot.
 N_8 I have frequent mood swings.
 N_9 I get irritated easily.
 N_{10} I often feel blue.

7.2 All traits, represented by typical questions

In Section 5.4 we saw that tangles in \vec{S} , and by implication also tangle traits T , can be compared directly with individual questions q from Q (see Section 7.1), by counting how many of the partitions in S they orient in the same direction. We said that when this number is high the question q *represents* the trait T well.

The following two tables list the top three representatives from Q , in order (best first), for each of the traits we found in [2] and [3]:

O : O_{10}, O_6, O_1	C : C_5, C_6, C_1	E : E_7, E_5, E_2	A : A_9, A_1, A_8	N : N_6, N_8, N_9
T_2 : O_{10}, C_5, N_3	T_3 : C_5, N_9, N_8	T_4 : O_{10}, O_5, O_4	T_5 : O_{10}, O_6, O_5	T_6 : O_8, O_1, O_7
T_7 : O_{10}, O_5, O_6	T_8 : A_5, A_7, E_2	T_9 : A_5, A_7, A_4	T_{10} : A_2, A_7, A_5	

The questions best representing the traits in [3]

O : O_2, O_1, O_8	C : C_5, C_6, C_9	E : E_7, E_5, E_1	A : A_5, A_2, A_4	N : N_6, N_9, N_8
T'_2 : A_5, A_6, A_2	T'_3 : O_{10}, C_5, N_3	T'_4 : E_1, E_2, E_3	T'_5 : A_5, A_4, A_8	T'_6 : A_2, A_7, A_5

The questions best representing the traits in [2]

7.3 All traits, represented by the tangle-distinguishers

In Section 5.3 we noted that traits are described most succinctly by the small collection of partitions in S that distinguish them efficiently from other traits. In this appendix we list these distinguishers explicitly, and then list for every trait how it orients them.

The efficient distinguisher s of two traits T_1 and T_2 is naturally associated with their unique common ancestor T in the tree of traits that displays them. Indeed, by definition s is the partition of Q in S of lowest order that T_1 and T_2 orient differently. So they orient all partitions of lower order than $|s|$ in the same way. The set of these oriented partitions is the largest-order tangle in T , and T has cohesion $|s|$.

In the tables below we denote this partition s of Q as $s(T)$. Note that s is the efficient distinguisher not only for T_1 and T_2 , but for all pairs of distinct traits that both refine T but no tangle of higher cohesion. The two sides of a partition $s = s(T)$ are stated, as subsets of Q , in the left and right half of the box below s .

First, the efficient distinguishers from [3], one associated with each branching node of the tree of traits in Figure 4:

$s(T_1)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A3 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 A1 A2 A4 A5 A6 A7 A8 A9 A10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10
$s(T_2)$	
E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10	N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10
$s(T_3)$	
C1 C2 C3 C4 C5 C6 C7 C8 C9 C10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10
$s(O)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A3 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10 O1 O8	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 A1 A2 A4 A5 A6 A7 A8 A9 A10 O2 O3 O4 O5 O6 O7 O9 O10
$s(T_4)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10 O1 O2 O4 O7 O8 O9	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10 O3 O5 O6 O10

$s(T_8)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10

$s(A)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A1 A3 A4 A5 A6 A8 A9 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 A2 A7 A10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10

The 14 traits in [3] other than T_1 orient these distinguishers as follows. Left and right arrows correspond to the sides of $s(T)$ shown left and right above.

T_2 : $\vec{s}(T_1)$
 T_3 : $\overleftarrow{s}(T_1)$ $\vec{s}(T_2)$
 O : $\vec{s}(T_1)$ $\vec{s}(T_2)$ $\overleftarrow{s}(T_3)$
 T_4 : $\vec{s}(T_1)$ $\vec{s}(T_2)$ $\vec{s}(O)$ $\overleftarrow{s}(T_3)$
 T_5 : $\vec{s}(T_1)$ $\vec{s}(T_2)$ $\vec{s}(O)$ $\overleftarrow{s}(T_4)$ $\overleftarrow{s}(T_3)$
 T_6 : $\vec{s}(T_1)$ $\vec{s}(T_2)$ $\overleftarrow{s}(O)$ $\overleftarrow{s}(T_4)$ $\overleftarrow{s}(T_3)$
 T_7 : $\vec{s}(T_1)$ $\vec{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\overleftarrow{s}(T_3)$
 C : $\overleftarrow{s}(T_1)$ $\vec{s}(T_2)$ $\overleftarrow{s}(O)$ $\overleftarrow{s}(T_4)$ $\overleftarrow{s}(T_3)$
 T_8 : $\vec{s}(T_1)$ $\overleftarrow{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\vec{s}(T_3)$
 A : $\vec{s}(T_1)$ $\overleftarrow{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\overleftarrow{s}(T_8)$ $\vec{s}(T_3)$
 T_9 : $\vec{s}(T_1)$ $\overleftarrow{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\overleftarrow{s}(T_8)$ $\overleftarrow{s}(A)$ $\vec{s}(T_3)$
 T_{10} : $\vec{s}(T_1)$ $\overleftarrow{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\overleftarrow{s}(T_8)$ $\vec{s}(A)$ $\vec{s}(T_3)$
 N : $\overleftarrow{s}(T_1)$ $\vec{s}(T_2)$ $\overleftarrow{s}(O)$ $\overleftarrow{s}(T_4)$ $\overleftarrow{s}(T_8)$ $\overleftarrow{s}(A)$ $\vec{s}(T_3)$
 E : $\vec{s}(T_1)$ $\overleftarrow{s}(T_2)$ $\vec{s}(O)$ $\vec{s}(T_4)$ $\vec{s}(T_8)$ $\vec{s}(A)$ $\vec{s}(T_3)$

Next, the efficient distinguishers from [2], this time associated with the branching nodes of the tree of traits in Figure 5:

$s(T'_1)$	
O1 O2 O3 O4 O5 O6 O7 O8 O9 O10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10 C1 C2 C3 C4 C5 C6 C7 C8 C9 C10

$s(T'_2)$	
E1 E2 E3 E4 E5 E6 E7 E8 E9 E10 N1 N2 N3 N4 N5 N6 N7 N8 N9 N10 A1 A2 A3 A4 A5 A6 A7 A8 A9 A10	C1 C2 C3 C4 C5 C6 C7 C8 C9 C10 O1 O2 O3 O4 O5 O6 O7 O8 O9 O10

$s(T'_3)$	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10	E1 E2 E3 E4 E5 E6 E7 E8 E9 E10
C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	A1 A2 A3 A4 A5 A6 A7 A8 A9 A10
O1 O2 O3 O4 O5 O6 O7 O8 O9 O10	
$s(T'_4)$	
E1 E2 E3 E4 E5 E6 E7 E8 E9 E10	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10	A1 A2 A3 A4 A5 A6 A7 A8 A9
A10	
C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	O1 O2 O3 O4 O5 O6 O7 O8 O9 O10
$s(A)$	
E1 E2 E3 E4 E5 E6 E7 E8 E9 E10	
N1 N2 N3 N4 N5 N6 N7 N8 N9 N10	A1 A3 A4 A5 A6 A8 A9
A2 A7 A10	
C1 C2 C3 C4 C5 C6 C7 C8 C9 C10	O1 O2 O3 O4 O5 O6 O7 O8 O9 O10

The 10 traits in [2] other than T_1 orient these distinguishers as follows. Left and right arrows correspond to the sides of $s(T)$ shown left and right above.

$$\begin{aligned}
T'_2: & \vec{s}(T'_1) \\
T'_3: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \\
C: & \vec{s}(T'_1) \ \vec{s}(T'_2) \ \overleftarrow{s}(T'_3) \\
O: & \overleftarrow{s}(T'_1) \ \vec{s}(T'_2) \ \overleftarrow{s}(T'_3) \\
T'_4: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \vec{s}(T'_3) \\
A: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \vec{s}(T'_3) \ \vec{s}(T'_4) \\
T'_5: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \vec{s}(T'_3) \ \vec{s}(T'_4) \ \vec{s}(A) \\
T'_6: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \vec{s}(T'_3) \ \vec{s}(T'_4) \ \overleftarrow{s}(A) \\
N: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \overleftarrow{s}(T'_3) \ \overleftarrow{s}(T'_4) \ \overleftarrow{s}(A) \\
E: & \vec{s}(T'_1) \ \overleftarrow{s}(T'_2) \ \vec{s}(T'_3) \ \overleftarrow{s}(T'_4) \ \overleftarrow{s}(A)
\end{aligned}$$