Age-Related Nuances in Knowledge Assessment

Ulrich Schroeders¹, Luc Watrin², Oliver Wilhelm²

¹ Institute of Psychology, University of Kassel, Germany

² Institute for Psychology and Pedagogy, Ulm University, Germany

Author Note

Ulrich Schroeders <https://orcid.org/0000-0002-5225-1122>

Luc Watrin <https://orcid.org/0000-0003-4343-3781>

Oliver Wilhelm <https://orcid.org/0000-0001-7980-1166>

We have no conflicts of interests to disclose. We confirm that the work conforms to Standard 8 of the American Psychological Association's Ethical Principles of Psychologists and Code of Conduct. We thank Diana Steger for developing the large item pool that we partially used in this study and Johannes Zimmermann for his helpful comments on an earlier draft of the manuscript.

Correspondence concerning this article should be addressed to Ulrich Schroeders, Psychological Assessment, Institute of Psychology, University of Kassel, Hollaendische Str. 36-38, 34127 Kassel, Germany, Email: schroeders@psychologie.uni-kassel.de

Please note that this manuscript is currently under review.

Abstract

Although crystallized intelligence (gc) is a prominent factor in contemporary theories of individual differences in intelligence, its structure and optimal measurement are elusive. Analogously to the personality trait hierarchy, we propose the following hierarchy of declarative fact knowledge as a key component of gc: a general fact knowledge factor at the apex, followed by broad knowledge areas (e.g., natural sciences, social sciences, humanities), knowledge domains (e.g., chemistry, law, art), and nuances. In most scientific contexts we are predominantly concerned with aggregate levels, but we argue that the sampling of knowledge items strongly affects distinctions at higher levels of the hierarchy. We illustrate the magnitude of item-level heterogeneity by predicting chronological age differences through knowledge differences at different levels of the hierarchy. Analyses were based on an online sample of 1,629 participants between 18 and 70 years old who completed 120 broadly sampled declarative knowledge items across twelve domains. The results of linear and elastic net regressions, respectively, demonstrated that the majority of the age variance was located at the item level, and the strength of the prediction decreased with increasing aggregation. Knowledge nuances seem to tap important variance that is not covered with aggregate scores (e.g., sum or factor scores) and that is useful in the prediction of age. In turn, these effects extend our understanding how knowledge is acquired and imparted. On a more general stance, to gain new insights into the nature of knowledge, its optimal measurement and psychometric representation, item and person sampling issues should be considered.

Keywords: crystallized intelligence, declarative knowledge, nuances, age trajectories, psychometric modeling, elastic net

Age-Related Nuances in Knowledge Assessment

Crystallized intelligence (gc) is a prominent factor in contemporary theories of the structure of intelligence (Carroll, 1993; Cattell, 1987; Horn & Noll, 1997; McGrew, 2009), albeit its nature and content remain elusive (Keith & Reynolds, 2010). What researchers understand by the term gc is very different, and it seems to be „both a store of acquired knowledge (e.g., lexical knowledge, general information, information about culture) as well as a collection of processing abilities (e.g., communication ability, listening ability)“ (Flanagan & Dixon, 2014). This combination of knowledge and language abilities traces back to the different conceptualizations in the literature. In the Three-Stratum-Theory, Carroll (1993) subsumed verbal abilities such as language comprehension, listening comprehension, and lexical knowledge below a gc factor. He pointed to the close relation between knowledge and speculated that this relationship may be due to similarities in the operationalization and the acquisition processes (Schipolowski, Wilhelm, & Schroeders, 2014). In contrast, Cattell (1943, 1963) described gc as a blend of skills, knowledge, and language-related abilities, with a focus on the declarative knowledge part, whereas fluid intelligence (gf) is the ability to “arrive at understanding relations among stimuli, comprehend implications, and draw inferences”. In his later work, Cattell focused more clearly on the knowledge part, which is the result of learning, education, and acculturation. Accordingly, gc should be measured with “tasks indicating breadth and depth of the knowledge of the dominant culture” (Horn & Noll, 1997, p. 69). In addition to the compelling dichotomy of Cattell’s gf-gc framework, which draws upon the reasoning—knowledge distinction dating back to Aristotle, especially the developmental perspective has contributed to the dissemination of the theory. Cattell’s work has inspired many researchers and theoretical frameworks, including PPIK Theory (Intelligence-as-Process, Personality, Interests, and Intelligence-as-Knowledge, Ackerman,

1996), and the distinction between cognitive mechanics and cognitive pragmatics (Baltes, 1997; Baltes et al., 1999, p. 489). Cognitive mechanics (mainly *gf*) subsume the "speed, accuracy, and coordination of elementary processing operations", whereas cognitive pragmatics reflects "the acquisition and lifelong practice of culturally transmitted bodies of declarative and procedural knowledge that are made available to individuals in the course of socialization and lifetime experiences" (Baltes et al., 1999, p. 489 ff.). Some of these experiences are normative in nature and culture-specific (e.g., formal education), while others are more informal (e.g., occupational knowledge), and still others are highly person-specific and idiosyncratic (e.g., avocational knowledge and expertise). In the present paper, we seek to further our understanding of age-related differences in *gc*—in the conceptualization of Cattell as declarative knowledge—by demonstrating the relevance of the hierarchical level on which such differences are investigated. Specifically, we discuss the development of *gc* across the life span, introduce a hierarchical structural model of knowledge, and combine both aspects to predict age differences by means of knowledge differences.

A Developmental Perspective on *Gc*

The question of how knowledge accumulates over time and what facilitates its acquisition is extremely important in numerous contexts, including formal learning environments such as schools. In his seminal book, Cattell (1987) outlined what became known as the investment hypothesis, which states that fluid intelligence (more precisely, *historic gf*) is invested to acquire crystallized abilities in various domains, which are hence positively correlated. Kan and colleagues (2011) pointed to an alleged paradox in a verbatim interpretation of *gc*: According to this view, *gc* is a statistical abstraction rather than a psychological capacity, because from a measurement perspective, a causal relationship underlying a psychological capacity presupposes

that the cause (i.e., *gc* as a latent variable) can be distinguished from its effects (i.e., the indicators). The authors argue that such an understanding is at odds with the investment theory with *gf* being the ultimate source for the commonality between crystallized abilities. Following this line of reasoning, factor analyses should not reveal a separate factor apart from *gf*, which is, however, assumed in many theories of the structure in intelligence. In last consequence, *gc* would only be understood as a psychological capacity by pointing to the relevance of other factors that are exhaustively decisive for its existence as a statistical abstraction. Please note that Cattell (1987, p. 140) preempted the objection, when he discussed other factors besides *gf* that exert an influence on knowledge, such as the "school curriculum, and of the social, familial, and personal influences which create interest and time for learning". Surprisingly, the dispute regarding whether or not *gc* exists in the sense of being more than just a function of prominent predictors such as *gf*, learning interest, learning time, quality of learning materials, et cetera has not been settled. In Cattell's terminology, *gc* is almost detached from the use of *present gf*, while historic *gf* has an indirect effect through learning. However, the empirical results regarding the validity of this investment theory are inconclusive (e.g., Ferrer & McArdle, 2004; Ferrer et al., 2007) and a recent synthesis (Peng et al., 2019) came to the conclusion that other aspects such as socio-economic status, genetics, and their interaction have to be taken into account.

Previous research suggests a quite homogeneous picture of the typical age trajectories of *gc*: Horn's initial research (e.g., Horn & Cattell, 1967; Horn, 1972, 1976) and a plethora of more recent studies (e.g., McArdle, 2002; Tucker-Drob et al., 2014; Tucker-Drob, 2019) usually show a constant increase for many *gc* indicators through age 65 years with a beginning decline only at advanced ages. Some studies even report that *gc* seems to be exempted from the so-called terminal decline (Bäckman & MacDonald, 2006). Please note that the research design influences

the age trajectories. Cross-sectional estimates of age trends are known to be confounded by cohort differences (Baltes et al., 1977), while longitudinal estimates suffer from bias through practice effects and selective attrition (Salthouse & Tucker-Drob, 2009; Tucker-Drob, 2019). Analyses that corrected for the drawbacks of the implemented research design tend to agree on the more general pattern described above (Rönnlund & Nilsson 2006). However, such a general statement for gc does not entail a homogeneous trend across different indicators of gc. Even for different aspects of declarative knowledge such as occupational knowledge, avocational knowledge, and traditional knowledge (see Ackerman, 1996; Beier & Ackerman, 2005) the age trajectories might be somewhat different. Situational factors related to schooling or vocational training can be expected to have an impact on the composition of gc. Because “cultural change, shift of mixture of areas intellectually fashionable, or a change in the school curriculum can weaken its identity and unity as a discernible factor”, Cattell (1987, p. 149) attested gc a “precarious existence”. Cattell’s focus on the school curriculum comes from the observation that the content of a gc test is hard to define or justify prescriptively after formal education has ended and knowledge acquisition is no longer canonized. Such a substantial practical problem concerning the measurement of gc also raise the question of how it is organized.

The Structure and Organization of Gc

In prominent models on the structure of personality such as the Five-Factor Model (e.g., Costa & McCrae, 1992) or the HEXACO model (Ashton & Lee, 2007), personality is conceptualized as a higher-order construct. At the apex, there are broad domains (Neuroticism, Extraversion, Conscientiousness, Agreeableness, and Openness to Experience), which are composed of different facets (e.g. Neuroticism encompasses the facets Anxiety, Vulnerability, Depression). Recently, an even narrower level of personality below the facets has been

introduced—the personality “nuances” (McCrae, 2015)—that are often operationalized as single personality items (Möttus et al., 2017; Seeboth & Möttus, 2018). With reference to this hierarchical conceptualization of personality, we propose the following levels of a knowledge hierarchy: gc (as declarative knowledge) at the apex, followed by broad knowledge areas (e.g., humanities, natural sciences, social sciences), knowledge domains (e.g., chemistry, medicine, art), and nodes or nuances. These levels of knowledge represent different levels of granularity, with the proposed hierarchy intended as a heuristic representing larger, more distinct units on an otherwise dimensional continuum. Of course, meso-levels within this hierarchy are plausible, such as subdomains below the domain level (e.g., organic, inorganic, and physical chemistry) and even disciplines below the subdomain level (e.g., polymers, aromatic compounds, amines as components of organic chemistry). At the lowest level of the hierarchy, which was labeled the nuance level in personality research (McCrae, 2015; Möttus et al., 2019), are nodes or themes, which are not necessarily identical to specific items, but can be seen as variations on a topic (e.g., the contributions of Elsa von Freytag-Loringhoven to Dadaism or key differences between Mussorgsky’s and Tchaikovsky’s compositions). Thus, items capture unique knowledge-relevant variance which is not accounted for by domains, areas, or general knowledge. We like to think of knowledge areas and domains as agglomerations and clusters of knowledge items. In this sense, the similarity between any two knowledge items is a continuous quantity in a high-dimensional space, and the hierarchical levels we use have a mostly functional character to identify groups of items. This conceptualization is akin to representations of information in computational sciences, in which clusters like principal components are convenient tools to reduce the information space. Concerning the neurological foundation of gc, the degree to which the fibre network is

efficiently connected seems to be moderately associated with knowledge, even after controlling for age- and sex-related differences ($r = .19$, Genç et al., 2019).

A non-trivial question that follows from these considerations is how many areas and domains can be found in a specific knowledge test. The answer mainly depends on the composition and scope of the knowledge measure but also on the epistemological settings (not to say researchers' degrees of freedom), according to which some distinctions might be useful and relevant but others distorting and wrong. The question of dimensionality is further complicated by the fact that items may belong to more than one domain. For example, "Who was the first female Fields Medal winner?¹" might be considered a question from the mathematical domain, but it also touches on history and even politics. In a recent smartphone-based study (Steger et al., 2019), more than 1,000 participants worked on up to 4,000 knowledge items from 34 subject domains such as geography, medicine, and music. In a bass-ackwards analyses, the authors illustrated the unfolding factor structure of crystallized intelligence and concluded that the number and interpretation of the broad knowledge areas that are extracted depend strongly on how these knowledge areas are compiled. The same logic also applies to the level of knowledge domains and the selected set of items in each domain. Moreover, knowledge items often do not exhibit a simple structure, but rather considerable cross-loadings (or within-item multidimensionality) with other knowledge domains and correlations of item residuals.

Combining Questions of Development and Structure

In intelligence research, we are predominantly concerned with scores on an aggregate level, whereas research at the item level remains sparse. This is because, from a psychometric perspective, it is often assumed that the indicators of a specific test can be used interchangeably.

¹ Maryam Mirzakhani—an Iranian mathematician.

This statement is easy to accept when it comes to decontextualized content (e.g., matrices test items), which is for example also reflected in the rational item construction of parallel matrices items (e.g., Irvine & Kyllonen, 2002). But when it comes to a knowledge test or other highly contextualized tasks (such as reading comprehension tasks), indicators might not be used interchangeably. For example, Schroeders, Wilhelm, and Olaru (2016) demonstrated that sex differences in knowledge manifest mainly on the item level, which can also affect the domain scores in specific item samples. More precisely, the authors used *Ant Colony Optimization*, a metaheuristic strategy to compile uni-dimensional and reliable short scales that can either favor male or female depending on the item sampling strategy implemented. This result was surprising because most previous research addressed the question of gender differences at the domain level rather than the item level, even though it turned out that substantial information was located at the level of specific items. In contrast to the assumptions of reflective measurement models (i.e., items are interchangeable and exhibit a simple structure within their domain), gender-associated differences can be more or less pronounced depending on item sampling. Instead of treating such differences on an item level as a nuisance or aiming for random item sampling (which is doomed to fail, because item selection is “almost invariably expert selection rather than sampling”, Loevinger, 1965, p. 147), they might provide hints to the source of gender-related knowledge differences, such as differences in interest (Halpern et al., 2007). This line of reasoning can be applied not only to gender, but also to other person characteristics, including age.

The Present Study

In the present study, we try to locate where within the knowledge hierarchy age-related differences occur. To this end, we predicted chronological age in a series of linear regression models using (a) a global gc score, (b) four broad knowledge area scores, and (c) twelve

knowledge domain scores as predictors. These models with aggregate scores are compared with (d) individual responses at the level of dichotomous items to cover all levels of the proposed knowledge hierarchy. Choosing chronological age as the dependent variable seems odd, given causality would usually be conceived to run in the opposite direction. However, a prediction model is not necessarily a causal model. There are noteworthy, earlier publications predicting age based on personality nuances (Möttus & Rozgonjuk, 2019), Facebook likes (Kosinski et al., 2013), and WhatsApp messages (Koch et al., 2020). In the present context, age—technically, the time expired since birth—is only a proxy for events that are associated with knowledge acquisition. Age can, therefore, be seen as opportunities for knowledge acquisition, which has an intuitive appeal for cognitive pragmatics. We used item level regressions as a mere statistical abstraction to complement the typical life span trajectories of crystallized intelligence that have been previously reported in the literature (e.g., Baltes et al., 1999, 2006; Horn & Cattell, 1967; Li et al., 2004; Tucker-Drob, 2009; Tucker-Drob & Salthouse, 2008). Usually, we study the development of gc or knowledge areas, but what if many of the age-related differences are located at the item level?

In correspondence with similar investigations from personality psychology (e.g., Achaa-Amankwa et al., 2020; Möttus & Rozgonjuk, 2019; Seeboth & Möttus, 2018), it seems likely that specific knowledge items behave differently across the life span, which resembles the concept of differential item functioning (DIF). Instead of treating differential item functioning as nuisance, it could be seen as an expression of construct-relevant information. We proposed knowledge domains and areas as agglomerations and clusters of broadly sampled knowledge items. Adhering to a strict unidimensional modeling of latent variables is, however, restricting the nomothetic breadth of the construct. Instead a broad (best representative) sampling from the

item universe is necessary to abstain from trimming and artificially homogenizing item samples which would result in oversimplified models. These two views in dealing with item uniqueness are also addressed by Rocca and Yarkoni (2020), who distinguished between a model evaluation that follows an explanatory tradition and those that follows a prediction-oriented approach. That is, instead of developing models that try to explain the underlying processes (i.e., data modeling), researchers could also evaluate models based on their ability to predict a specific outcome (i.e., algorithmic modeling), regardless of how this aim is achieved.

In the current case, we deal with age-related differences, but several moderators are conceivable such as gender, socio-economic status, ethnicity or cultural background and the chance of an item functioning differently increases with every additional moderator considered. Such moderating effects might sometimes be unwanted from a normative point of view (e.g., gender-fair measures), but they can be informative. From a psychometric point of view, heterogeneous age-related trajectories at the item level would be at odds with prevalent measurement models (i.e., latent variable modeling) because they assume a common cause that underlies a set of item responses (Markus & Borsboom, 2013). Therefore, after establishing the level on which age differences predominantly appear, we take a closer look and ask what the results can tell us about how knowledge is organized and acquired over the life course.

Method

Design and Participants

Data were collected using a web-based survey administered between June 2020 and July 2020. Participants were recruited through social media, mailing lists, and the German online panel *Respondi*, which allowed for an age-stratified sample. Participants provided the following demographic information: age in years, gender, job status and sector and command of the

German language. Subsequently, they worked on a 120-item knowledge test, a 30-item RIASEC interest questionnaire (a German translation of the 30-item O*Net Mini Interest Profiler Mini-IP; Rounds et al., 2016), and a short version of the openness scale of the Trait Self-Description Inventory (German translation, 8 items; Olaru et al., 2015). A total of 1,721 participants completed the questionnaires, of which 1,629 (94.7 %) were included in the final analyses after data cleaning (for dropout at the different stages of data cleaning see Table OS 1 in the online supplement). The sample size was determined based on a power analysis conducted for another publication using the same dataset; the sample size is ample for the present purpose. About half of the analysis sample were female (46.8%, 51.9% male, and 1.3% diverse or did not report their gender); the average age was 45.3 years ($SD = 14.7$) with very good coverage across the lifespan (see Figure OS 1 for the age distribution). Concerning their current work-related situation, the sample was made up as follows ($n = 1,629$): 52.8% private-sector employees, 17.3% retired, 10.7% school/university students, 6.4% self-employed, 3.4% civil servants, and 9.4% others (e.g. unemployed, maternity leave, no response). All participants in the analysis sample stated that they understood German at least *well* on a four-point scale (with the response options *very poor*, *poor*, *well*, and *very well*).

Measure

The gc measure was compiled from a larger item pool (Steger et al., 2019) covering a broad range of knowledge across 34 knowledge domains (e.g., math, nutrition). Based on previous findings concerning the dimensionality of gc (Steger et al., 2019), we sampled items from four broad knowledge areas (humanities, life sciences, natural sciences, and social sciences). Each knowledge area contained three domains with ten items each, resulting in a total of 120 items. Items were selected to have a wide range of difficulty and to broadly and deeply

cover the content domain (all items are provided in German and English in the online supplement). The multiple-choice items with four response options were presented in random order. Correct answers were coded with 1, wrong answers with 0. The skipping of knowledge items was technically prevented. In addition to the gc measure, participants answered a short sociodemographic questionnaire as well as short measures of openness and vocational interest, which are not considered in the current manuscript.

Statistical Analysis

To answer the question to what extent knowledge and age were associated, we used knowledge scores at each level of the knowledge hierarchy to predict age-related differences. In more detail, we estimated three linear regressions, with a) a single gc factor score, b) four factor scores for the knowledge areas, and c) twelve factor scores for the knowledge domains as predictors to predict participants' age. For the item-level models with the 120 items as predictors, we applied d) elastic net regressions to avoid overfitting and to increase the robustness of the regression parameters. The R package *caret* (Kuhn, 2008, 2020) was used as the interface for the *glmnet* R package (Friedman et al., 2010). We provide the syntax used within the Open Science Framework (Soderberg, 2018) at https://osf.io/gcwku/?view_only=04565742dc574bea9d2c13d0fe9d17e5 to foster transparency and reproducibility.

Elastic net regression. Elastic net regression combines two common forms of regularization (Hastie et al., 2008)—ridge regression (Hoerl & Kennard, 1970) and the LASSO (= least absolute shrinkage and selection operator) regression (Tibshirani, 1996). Ridge regression shrinks the coefficients of correlated predictors towards each other (Friedman et al., 2010) and less important coefficients towards 0, without setting them exactly to 0, thus retaining

the total number of predictors (formulae are provided in the Appendix). In contrast, LASSO regression uses a penalization method which combines continuous shrinkage of the regression coefficients and automatic variable selection to calculate sparser prediction models (Tibshirani, 1996, 2011), which often makes it the preferred procedure in many applied settings (Zou & Hastie, 2005). However, there are also drawbacks to LASSO regression (e.g., when predictors are highly correlated, see also Friedman et al., 2010). Elastic net regression, which was used in the subsequent analyses at the item level, offers a compromise between ridge and LASSO regression by averaging the coefficients of strongly correlated predictors and keeping them in the model. At the same time, it provides a more parsimonious model by applying variable selection (Friedman et al., 2010; Zou & Hastie, 2005).

Model accuracy. To quantify the prediction accuracy of the models, we first split the total sample into a training-validation sample (80%) and an independent test sample (20%), also known as a holdout sample. Because item-level models with many predictors tend to overfit, that is, explain more variance than models on the area or domain level for purely statistical reasons (see also Seeboth & Möttus, 2018), we used ten-fold *cross-validation* (Yarkoni & Westfall, 2017). More precisely, we used a tune length of 21 to select the best parameters within the training-validation sample to minimize the *Mean Square Error* (MSE) term, which quantifies both the accuracy and the precision (Helwig, 2017). In a second step, we repeated this procedure of cross-validation and testing (inner sampling) 100 times (i.e., outer resampling). As central outcomes for evaluating performance of the regression models in both the training and the test samples, we used the following indices (see also Kvålseth, 1985): explained variance (R^2), the *Root Mean Squared Error* (RMSE), and the *Mean Absolute Error* (MAE, formulae are given in the Appendix).

Results

Descriptive Results

Table 1 shows the proportion correct and values describing the distribution (*SD*, skewness and kurtosis) at each level of the knowledge hierarchy (gc, area, domain). Item difficulty on the item level ranged between .24 and .94, with good coverage across all areas and domains (for more details, please see Table OS 2). The item-total correlations varied between .08 and .61, with an average of .38. To examine whether the items from a given knowledge area or domain can be meaningfully aggregated into scales (Little et al., 2002), we estimated uni-dimensional τ -congeneric measurement models with the *Weighted Least Squares Mean and Variance adjusted* (WLSMV) estimator, which is appropriate for binary scales (Beauducel & Herzberg, 2006). Overall, model fit indices were good for the knowledge areas and domains (Yu, 2002), evaluated according to the following criteria: *Comparative Fit Index* (CFI) $\geq .96$ and *Root Mean Square Error of Approximation* (RMSEA) $\leq .06$. Although the CFI was below the cut-off value for the knowledge domain *Nutrition* and *Law*, no item selection was conducted.

McDonald's ω (1999)—the factor saturation in a τ -congeneric uni-dimensional factor model (Bentler, 2009)—was calculated as a reliability estimate in addition to Cronbach's α , because the latter presupposes essential τ -equivalence, an assumption that is commonly not met in real data sets (McNeish, 2018). In the present case, we also preferred the τ -congeneric model with McDonald's ω , which varied between .60 and .81, and used the respective factor scores for the subsequent prediction models. Correlations between knowledge scores and chronological age were small to moderate (see Table 1), that is, slightly negative for the knowledge area of natural sciences and moderately positive for humanities. This pattern was also reflected in the corresponding knowledge domains, with a maximum of $R^2 = .14$ for music.

Table 1. Descriptive Statistics, Model Fit, Reliability Estimates, and Correlation with Age for Knowledge Scales

Scale	<i>M</i>	<i>SD</i>	skew.	kurt.	χ^2	<i>df</i>	CFI	RMSEA [CI _{90%}]	α	ω	$r_{fs,age}$	R^2
Gc	.57	.14	.15	-.60	12,650.1	7,020	.87	.02 [.02; .02]	.96	.96	.14	.02
Area												
Humanities	.56	.20	.16	-.79	969.0	405	.96	.03 [.03; .03]	.91	.92	.34	.12
Life Sciences	.60	.16	-.12	-.54	679.9	405	.95	.02 [.02; .02]	.86	.86	.09	.01
Natural Sciences	.51	.19	.45	-.54	784.1	405	.96	.02 [.02; .03]	.90	.90	-.16	.03
Social Sciences	.60	.17	.08	-.58	860.9	405	.93	.03 [.02; .03]	.86	.87	.10	.01
Domain												
Music	.58	.22	-.10	-.70	60.6	35	.99	.02 [.01; .03]	.76	.77	.37	.14
Literature	.57	.22	.10	-.74	117.1	35	.96	.04 [.03; .05]	.79	.79	.28	.08
Art	.52	.24	.18	-.75	72.5	35	.98	.03 [.02; .03]	.80	.81	.23	.05
Medicine	.62	.20	-.13	-.51	61.8	35	.98	.02 [.01; .03]	.73	.74	.05	.00
Health	.60	.19	-.28	-.32	36.0	35	1.00	.00 [.00; .02]	.59	.60	.03	.00
Nutrition	.59	.20	-.12	-.48	154.8	35	.88	.05 [.04; .05]	.71	.72	.16	.03
Physics	.57	.21	.08	-.68	75.5	35	.96	.03 [.02; .03]	.69	.70	-.08	.01
Math	.47	.23	.40	-.51	64.7	35	.98	.02 [.01; .03]	.78	.78	-.20	.04
Chemistry	.48	.24	.31	-.72	104.0	35	.97	.03 [.03; .04]	.78	.79	-.12	.01
Economics	.60	.21	-.07	-.73	96.4	35	.96	.03 [.03; .04]	.76	.77	.03	.00
Politics	.55	.20	.21	-.44	50.1	35	.98	.02 [.00; .03]	.65	.66	.14	.02
Law	.64	.20	-.28	-.45	84.1	35	.93	.03 [.02; .04]	.66	.66	.09	.01

Note. $N = 1,629$. skew. = skewness; kurt. = kurtosis; $r_{fs,age}$ = correlations are between age and factor scores from a uni-dimensional τ -congeneric model.

Comparison of Age Differences Across Levels of the Knowledge Hierarchy

Figure 1 shows the means (and standard deviations) of the prediction accuracy across 100 random splits of the total sample into a training-validation and a testing sample across the four levels of the knowledge hierarchy (global, area, domain, and item). The outcome R^2 is calculated as the correlation between predicted and observed age values. The difference between the training-validation and independent test sample indicates overfit. The degree of overfitting was low for the aggregate levels and it was still acceptable for the knowledge item level ($\Delta R^2 = .06$, see also Table OS 3). In addition to the elastic net regularization and cross-validation as statistical countermeasures, a large sample size provides a “natural guard against overfitting” (Yarkoni & Westfall, 2017, p. 9)². While the total gc score and age were almost unrelated in the holdout sample, 24% of the age variance could be explained with scores for the broad knowledge areas and 28% with scores for the knowledge domains. At the item level, 54% of the variation in age could be predicted, showing that a substantial amount of information is conveyed at the lowest level of the hierarchy. Put differently, age-associated differences are strongly related to differences in the individual item response vector, and these item effects are much stronger than what can be found with aggregate scores.

² In the online supplement, we demonstrate with a simulation study that the large sample size in particular is an important factor against overfitting (see Figure OS 2). We prefer elastic net over simple linear regressions, first because it is the appropriate and versatile method for models with a large number of predictors, second because the overfit is still smaller in comparison to linear regression, even in the total sample, and third the resulting models are more parsimonious and thus easier to interpret.

Figure 1. Accuracy Predicting Age Based on Knowledge Scores and Items. $N = 1,629$ Left side: The boxplot reflects the interquartile range, the solid line represents the median, and the whiskers represent minimum/maximum values within 1.5 times the interquartile range. Right side: Jittered point plot of 50 randomly drawn values.

The same finding expressed in terms of *MAE*—the mean absolute difference between the predicted and reported age—provides an intuitive measure of the prediction accuracy in the original metric of the criterion. For the item-level model, $MAE = 8.06$ years in the holdout sample (see Table OS 3), which is quite accurate given participants' large age range between 18 and 70 years old. In contrast, for the overall gc score, the mean absolute error was almost identical to the standard deviation of age ($MAE = 14.6$ vs. $SD = 14.7$ years). Moreover, prediction at the item level constitutes a significant improvement compared to the aggregate predictors, with an MAE of more than 10 years for the knowledge area and knowledge domain levels.

Variable Importance and the Development of Knowledge

Having demonstrated that a large share of the age differences in knowledge are located at the item level rather than higher levels of the hierarchy, subsequent questions include whether or not certain items are more suitable for differentiating between younger and older participants. Which items have the largest variation across age; do they belong to a specific domain, or what do they have in common? To tackle these questions, we plotted the average variable importance for all items in the test sample across all iterations (see Figure 2). Variable importance is a simple transformation of the absolute weights in the linear regression ($= |w_j|/\max(w_j) \times 100$), so that the values range between 0 and 100 (Grömping, 2009). The main findings are as follows: First, several items contribute incrementally to age differentiation. Second, the most informative items cannot be assigned to a specific knowledge domain. Again, it is evident that age variability is located at the level of knowledge items, not knowledge domains (i.e., no row-wise clustering) or areas. For example, the five most influential items are presented in Figure 2. Comparing these items, one is tempted to ask what these items have in common, and a preliminary answer might be that they deal with the “fine arts” and that theatre, opera, and classical music concerts tend to attract an older audience. But such explanations are prone to hindsight bias, because the same holds true for other items within these domains: While the item “Who composed Boléro?” has an average variable importance of 71.2, the item “Who composed The Magic Flute?” has a variable importance of only 1.1, and this with quite similar difficulty ($P_i(x) = .66$ vs. $.55$). Items with high variable importance are those that are arguably subject to historical and social change, but this can only be determined empirically. Scientometric analyses, for example with Google Books Ngram Viewer, can be used to depict cultural-societal change, but methodically such comparisons are not trivial (Pettit, 2016).

Figure 2. Average Variable Importance in Elastic Net Regression Across 100 Iterations.

Each set of three domains belongs to a common knowledge area (from top to bottom):

humanities, life sciences, natural sciences, social sciences.

Age Trajectories at the Item Level

Using (elastic net) regressions, we have so far shown that age-related variance within the knowledge hierarchy is mainly located at the item level. In order to exemplify the results and to present age trajectories at the item level, we plotted the percentage correct for two variables that are strong explanatory variables across age (Figure 3). Thus, the underlying data are scores aggregated across age groups (with a median sample size of 31 participants at each specific age). To illustrate the item trajectories we used LOWESS (*locally weighted scatterplot smoothing*), which is a generalization of a moving average and polynomial regression. The solid line represents the question concerning the author of the controversial erotic/pornographic novel

“Wetlands”, which was first published in 2009. The curve peaks shortly above 30 years of age, corresponding to readers who were in their early 20s when the book debuted. Younger and older participants were less familiar with the novel. In contrast, the dashed line represents the Boléro item, which peaked at an age above 60 years. The reasons for the shape of the item trajectories are reviewed in more detail in the discussion.

Figure 3. Trajectories of Percentage Correct Across Age for Two Example Items. Smoothing was conducted with LOWESS (locally sample-size weighted scatterplot smoothing). The size of the circles corresponds to the sample size at a specific age/year. Total $N = 1,629$.

Discussion

In this paper, we have investigated at which level of the knowledge hierarchy age differences are most pronounced. The results are unequivocal—the majority of age variance is located at the item level ($R^2 = .54$), and prediction accuracy decreases with increasing aggregation. Apparently, declarative knowledge measures have substantial item uniqueness, which conveys important age-related information. In personality research, the importance of

item-level variation is well-documented; personality nuances have been used to predict age differences (Möttus & Rozgunjuk, 2019), country of residence (Acha-Amankwaa et al., 2020), and a variety of socio-demographic and health-related variables (Seeböth & Möttus, 2018). In intelligence research, however, similar analyses are still pending. We have all reason to assume that especially for knowledge assessment, item sampling effects are not restricted to age or gender (Schroeders et al., 2016). For example, knowledge about unsaturated fatty acids or flavonoids could be related to the body mass index or myocardial infarction beyond the aggregate knowledge in nutritional science or medicine. The predictive power is assumed to be maximum if the granularity and width on the side of the criterion corresponds to the granularity and width of the measurement on the side of the predictor. This is also described in *Brunswick's Symmetry Principle* (Wittmann et al., 1988, 1999), which describes the relations between two hierarchically structured constructs at different levels of aggregation. Underestimation of the relations occur when the levels of aggregation do not match (see also Kretzschmar et al., 2018). This means, for instance, to predict myocardial infarction a comprehensive knowledge assessment should cover multiple risk or protective factors (healthy eating, exercises, etc.). Depending on the subject matter, this would require a criterion-targeted sampling of knowledge items (which we are tempted to label a coalition-of-the-willing approach), which does not tally with a “librarian-style” arrangement of knowledge as expressed in the knowledge domains and areas.

Please note that we do not imply to equate predictive accuracy with validity. The fact that unique item variance—variance that is not tapped with broader knowledge domains (or areas)—was predictive, does not invalidate or belittle existing findings on the age trajectories of aggregate scores (see also Möttus et al., 2019). They remain informative and meaningful.

Nonetheless, item-level variation might help us to understand which effects knowledge differences have on individuals in real life. Moreover, the finding that some items convey large age effects, while others show next to no differences does also tell us something about the way knowledge is acquired and imparted. In typical definitions, gc is conceptualized as “the breadth and depth of knowledge of the language, concepts, and information of the dominant culture” (Horn, 2008, p. 197) or “knowledge that is considered a cultural good and important enough to pass it on to the next generation” (Schroeders, Schipolowski, et al., 2016, p. 88). To understand the transmission of knowledge it is therefore essential to delineate what is considered the dominant culture and what is deemed important enough to pass on to the next generation.

“Communication as Culture”

A predominant view is that knowledge is transmitted from one person to another, that knowledge is imparted from one generation to the next. However, in a compelling essay, the American communication theorist James Carey (2009) described “culture” as a conversation that started long ago, a conversation to which we can contribute throughout our lifetime and that will continue after our deaths. If we replace the word “culture” with “knowledge”, we better understand what Cattell meant with his original definition of crystallized intelligence. The process of intergenerational knowledge transmission can also be seen as an iterated learning process with persons as Bayesian agents (Kalish et al., 2007). Thus, what knowledge a person acquires depends on prior knowledge and the density and intensity of new incoming information. If a sufficiently large part of a society contributes to the process of teaching, learning, and spreading knowledge, the respective information will survive, while other information will fade away and decay eventually. This also applies to the general knowledge items used in this study. These items were selected from a larger item pool to allow for reliable measurement across a

broad age range—not for maximum discrimination between different age groups. One can easily imagine a knowledge test addressing certain epoch-shaping events (e.g., the Cuban Missile Crisis and the moon landing in the 1960s) or issues that are especially important for a certain age range (e.g., life insurance or early retirement schemes). Such a test of current events knowledge (Beier & Ackerman, 2001; Hambrick et al., 2007, 2008) would likely outperform a broad knowledge test in the prediction of age. Bearing this in mind, it is even more astonishing to what extent it was possible to predict age using a set of items that was not selected for these reasons. These results tell us a lot about the process of acquiring knowledge and the nature of gc.

Some knowledge is of a general, indexical nature, while other knowledge is strongly bound to a specific period of life or time. As in all purely cross-sectional designs, these effects cannot be disentangled because age is confounded with cohort (see also the double x-axis in Figure 3). Moreover, the present study includes questions that are intrinsically more associated with ageing processes (“Which organ is supported with dialysis?”) and other questions that are more closely associated with cohort effects (“Who wrote the novel ‘Wetlands’?”). For still other items, both age and cohort effects may be at work: For example, the answer to “Who composed Boléro?” might depend on the respondents’ chronological age, but is also related to the general social trend in which traditional forms of entertainment and education (i.e., classical music concerts, opera, theatre) are becoming less common. In addition, confounding variables such as respondents’ financial situation (attending concerts is expensive) might be relevant, rendering post-hoc explanations difficult. Cattell (1987, p. 189) outlined a remedy for separating age and cohort effects, that is, to divide a given curve “into an abstracted ‘normative curve’—the typical age curve in typical biological and cultural conditions, sampled across several epochs—and an *epogenic* curve, that is, the special part due to the particular historical conditions in that epoch”.

Such a design is now labeled multi-cohort longitudinal study, which combines the advantages of cross-sectional and longitudinal studies. In summary, knowledge assessment, whether it be cross-sectional, longitudinal, or a combination of both, depends on an individual's characteristics such as reasoning or interests. But it is also bound to a specific period of time, culture, geographic region, etc.—which is also an essential part in common definitions of gc. In analogy to the idea of communication as culture, knowledge can also be conceived as constantly changing as part of an ongoing societal communication.

This dependence of knowledge on time, place, learning environments, et cetera relates to another discussion in the literature: In 2012, Schneider and McGrew introduced the new broad ability labeled gkn, somewhat unfortunately labeled “general (domain-specific) knowledge”. Gkn comprises four narrow abilities previously attributed to the gc domain (i.e., foreign language, geography achievement, general science information, and knowledge of culture). The idea was to account for acquired knowledge from specialized domains vs. general knowledge, but we argue that the distinction is arbitrary. Since there is no universal canon of knowledge (or even a consensus about the topics to be included), a classification of items cannot be made on the basis of its content, but would always follow a supposedly sample-dependent empirical investigation. Classifying or bundling knowledge items faces two major challenges: First, traditional factor-analytic tools might not satisfactorily handle the assignment of items to multiple domains (which violates Thurstone's principle of simple structure). Second, new dimensions can arise depending on the composition and amount of items in a knowledge test but also as a function of person sampling and sample composition. New methodological approaches such as exploratory graph analysis in combination with entropy fit indices might be useful

(Golino et al., 2020), but such approaches will still suffer the problems of accommodating the nuance perspective proposed here.

The Psychometrics of Knowledge Assessment

The present study's results are also informative with respect to the psychometric model of knowledge assessment, specifically reflective vs. formative modeling. The reflective model is considered the genuine model of the psychological world, in which constructs represent dispositions (e.g., intelligence, motivation) that cannot be measured directly. Items are indicators of a latent ability, which serves as a common cause influencing participants' responses to the indicators (Markus & Borsboom, 2013). In the context of knowledge assessment, this corresponds to the view that the latent ability fully accounts for the correlations among individual knowledge items, while their unique variance reflects sources of random errors (unless the model allows for domain-, area-, or item-specific variance). If someone is more knowledgeable, then the probability of solving any sampled knowledge items is increased relative to a less knowledgeable person. This mechanism implies a kind of transfer effect of knowledge (i.e., Matthew effect) or at least some mutual dependence (Savi et al., 2019). Indeed, prior knowledge can be essential (and even causal) for the acquisition of new knowledge (Hambrick et al., 2008, but see also Schroeders, Schipolowski, et al., 2016, for empirical evidence in the applied setting of formal education). Recently, Savi et al. (2019, p. 1052) proposed an interesting idiographic explanation for this mutualism, in which intelligence is conceptualized as an evolving network: Units of knowledge "that are more related have a higher probability of becoming connected" during development.

In formative models, indicators cannot be used interchangeably, and they together determine an emerging composite variable (van der Maas et al., 2014). According to this view,

knowledge is exhaustively represented by its indicators, to which specific weights are assigned. The formative approach has been criticized for substantial shortcomings, such as a measurement free of error or the fact that a formatively measured construct does not exist independently of its indicators (please see Bollen & Diamantopoulos, 2017 for a rebuttal of these points of criticism). Driving license tests, pilot license tests or the civics questions of naturalization tests applied by WEIRD countries arguably constitute examples in the field of knowledge assessment. In these tests, knowledge has been canonized or the elements of a test are bijectively tied to the specific measurement intention. School achievement tests that reflect a mandatory school curriculum also fall within this category. In this type of measurement model, aggregate scores represent an index that can be used to optimally capture the increase in knowledge, which can in turn be used to distinguish groups (e.g., test-takers who achieve vs. did not achieve the predefined knowledge requirements for driving a motor vehicle). Reckase (2017) pointed out that developing a test using the domain sampling mode (i.e., educational perspective of test construction) needs a detailed description of the domain from which representative samples can be taken. In most cases, however, knowledge in a given field is not defined or circumscribed with such clarity, that is, there is no exhaustive list of indicators that belong to a domain (for which reason Loevinger, 1965, objected to the term “item population”). Accordingly, *gc* cannot be understood as a fixed quantity assessed with a limited pool of knowledge items, although Markus and Borsboom (2013, p. 87) argued that “for infinite domains the formative composite takes on the properties of the latent variable in a reflective model, including properties like conditional independence”. However, this is no remedy for the original problems of formative assessment.

Knowledge Measurement—What For?

In our opinion, when it comes to prediction, both the formative and the reflective modeling approaches fall short because they do not strive for indicators (sampled or not) to go beyond the aggregate scores. A reflective model does not pursue it due to its implication that the covariance between any two indicators is (only randomly different from) zero once it is controlled for the latent factor(s). This also applies to external variables. The formative model presupposes that the sampled items are a specific and fixed set of indicators that represent the construct in question. It is not admissible to add items, specifically if they modify relations with external variables. In contrast to these considerations, what the present and other studies (Schroeders et al., 2016) demonstrate is that item sampling effects can be essential for knowledge tests in comparison to other cognitive ability tests—presumably due to the infinite and dynamic diversification of knowledge. Whereas the assumption of interchangeable items might hold for a fluid intelligence test like a matrix test (see also Spearman’s concept of the *indifference of indicator*), it does make a difference which items are compiled to form a knowledge scale. Unfortunately, such construct-relevant variance on the item level is not only understudied, it is also largely overlooked when considering aggregate knowledge scores. Consequently, as Möttus (2016) pointed out for self-report items, associations that are specific to a particular level in the hierarchy (e.g., items) should not be generalized to aggregate levels (e.g., domains), without taking appropriate account of item sampling issues.

The purpose of any scientific model is either to explain observed data (i.e., explanatory modeling) or to predict data accurately (predictive modeling), whereby the former seems to be much more popular in psychology (Shmueli, 2010). We conclude that although the result that item level models yield more accurate predictions seems obvious from a statistical point of view,

it challenges traditional views in psychology that tend to value explanation over prediction (Yarkoni & Westfall, 2017; regarding the reasons and consequences for the research see Jolly & Chang, 2019). Moreover, intelligence research predominantly relies on factor analytic methods (e.g., Carroll, 1993), a tool for capturing commonality between items rather than item uniqueness. In this respect, item level predictions are unusual in intelligence research. Therefore, a more nuanced view on knowledge assessment might improve our understanding of the development of knowledge, its optimal measurement and psychometric representation.

References

- Achaa-Amankwaa, P., Olaru, G., & Schroeders, U. (in press). Coffee or tea? Examining cross-cultural differences in personality nuances across former colonies of the British Empire. *European Journal of Personality*.
- Ackerman, P. L. (1996). A theory of adult intellectual development: Process, personality, interests, and knowledge. *Intelligence*, 22, 227–257. [https://doi.org/10.1016/S0160-2896\(96\)90016-1](https://doi.org/10.1016/S0160-2896(96)90016-1)
- Ashton, M. C., & Lee, K. (2007). Empirical, theoretical, and practical advantages of the HEXACO model of personality structure. *Personality and social psychology review*, 11(2), 150–166. <https://doi.org/10.1177/1088868306294907>
- Bäckman, L., & MacDonald, S. W. S. (2006). Death and cognition: Synthesis and outlook. *European Psychologist*, 11(3), 224–235. <https://doi.org/10.1027/1016-9040.11.3.224>
- Baltes, P. B., Cornelius, S. W., & Nesselroade, J. R. (1978). Cohort effects in behavioral development: Theoretical and methodological perspectives. In W. A. Collins (Ed.), *Minnesota Symposia on Child Psychology* (pp. 1–63). Lawrence Erlbaum Associates.
- Baltes, P. B. (1997). On the incomplete architecture of human ontogeny: Selection, optimization, and compensation as foundation of developmental theory. *American Psychologist*, 52(4), 366–380. <https://doi.org/10.1037//0003-066x.52.4.366>
- Baltes, P. B., Lindenberger, U., & Staudinger, U. M. (2006). Life span theory in developmental psychology. In R. M. Lerner & W. Damon (Eds.), *Handbook of child psychology: Theoretical models of human development* (pp. 569–664). Wiley.

Baltes, P. B., Staudinger, U. M., & Lindenberger, U. (1999). Lifespan psychology: Theory and application to intellectual functioning. *Annual Review of Psychology*, 50, 471–507.

<https://doi.org/10.1146/annurev.psych.50.1.471>

Beauducel, A., & Herzberg, P. Y. (2006). On the performance of maximum likelihood versus means and variance adjusted weighted least squares estimation in CFA. *Structural Equation Modeling: A Multidisciplinary Journal*, 13, 186–203.

https://doi.org/10.1207/s15328007sem1302_2

Beier, M. E., & Ackerman, P. L. (2001). Current-events knowledge in adults: An investigation of age, intelligence, and nonability determinants. *Psychology and Aging*, 16(4), 615–628.

<https://doi.org/10.1037/0882-7974.16.4.615>

Bentler, P. M. (2009). Alpha, dimension-free, and model-based internal consistency reliability.

Psychometrika, 74(1), 137–143. <https://doi.org/10.1007/s11336-008-9100-1>

Bollen, K. A., & Diamantopoulos, A. (2017). In defense of causal-formative indicators: A minority report. *Psychological Methods*, 22(3), 581–596.

<https://doi.org/10.1037/met0000056>

Carey, J. W. (2009). *Communication as culture: Essays on media and society*. Routledge.

Carroll, J. B. (1963). *A model of school learning*. Teachers College Record.

Carroll, J. B. (1993). *Human cognitive abilities: A Survey of factor-analytic studies*. Cambridge University Press.

Cattell, R. B. (1943). The measurement of adult intelligence. *Psychological Bulletin*, 40(3), 153–193. <https://doi.org/10.1037/h0059973>

Cattell, R. B. (1963). Theory of fluid and crystallized intelligence: A critical experiment. *Journal of Educational Psychology*, 54, 1–22. <https://doi.org/10.1037/h0046743>

- Cattell, R. B. (1987). *Intelligence: Its structure, growth and action*. North Holland.
- Costa, P. T., & McCrae, R. R. (1992). *The Revised NEO personality inventory (NEO-PI-R) and NEO Five-Factor inventory (NEO-FFI) professional manual*. Psychological Assessment Resources.
- Ferrer, E., & McArdle, J. J. (2004). An experimental analysis of dynamic hypotheses about cognitive abilities and achievement from childhood to early adulthood. *Developmental Psychology*, 40(6), 935–952. <https://doi.org/10.1037/0012-1649.40.6.935>
- Ferrer, E., McArdle, J. J., Shaywitz, B. A., Holahan, J. M., Marchione, K., & Shaywitz, S. E. (2007). Longitudinal models of developmental dynamics between reading and cognition from childhood to adolescence. *Developmental Psychology*, 43(6), 1460–1473. <https://doi.org/10.1037/0012-1649.43.6.1460>
- Flanagan, D. P., & Dixon, S. G. (2014). The Cattell-Horn-Carroll Theory of cognitive abilities. In C. R. Reynolds, K. J. Vannest, & E. Fletcher-Janzen (Eds.), *Encyclopedia of Special Education*. John Wiley & Sons. <https://doi.org/10.1002/9781118660584.cse0431>
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1). <https://doi.org/10.18637/jss.v033.i01>
- Genç, E., Fraenz, C., Schlüter, C., Friedrich, P., Voelkle, M. C., Hossiep, R., & Güntürkün, O. (2019). The neural architecture of general knowledge. *European Journal of Personality*, 33(5), 589–605. <https://doi.org/10.1002/per.2217>
- Golino, H., Moulder, R., Shi, D., Christensen, A. P., Garrido, L. E., Nieto, M. D., Nesselroade, J., Sadana, R., Thiagarajan, J. A., & Boker, S. M. (2020). Entropy fit indices: New fit measures for assessing the structure and dimensionality of multiple latent variables.

Multivariate Behavioral Research, 1–29.

<https://doi.org/10.1080/00273171.2020.1779642>

Grömping, U. (2009). Variable importance assessment in regression: Linear regression versus random forest. *The American Statistician*, 63(4), 308–319.

<https://doi.org/10.1198/tast.2009.08199>

Halpern, D. F., Benbow, C. P., Geary, D. C., Gur, R. C., Hyde, J. S., & Gernsbacher, M. A.

(2007). The science of sex differences in science and mathematics. *Psychological Science in the Public Interest*, 8(1), 1–51. <https://doi.org/10.1111/j.1529-1006.2007.00032.x>

Hambrick, D. Z., Meinz, E. J., & Oswald, F. L. (2007). Individual differences in current events knowledge: Contributions of ability, personality, and interests. *Memory & Cognition*, 35(2), 304–316. <https://doi.org/10.3758/BF03193451>

Hambrick, D. Z., Pink, J. E., Meinz, E. J., Pettibone, J. C., & Oswald, F. L. (2008). The roles of ability, personality, and interests in acquiring current events knowledge: A longitudinal study. *Intelligence*, 36(3), 261–278. <https://doi.org/10.1016/j.intell.2007.06.004>

Hastie, T., Tibshirani, H., & Friedman, J. H. (2008). *The Elements of Statistical Learning*. Springer.

Helwig, N. E. (2017). Adding bias to reduce variance in psychological results: A tutorial on penalized regression. *The Quantitative Methods for Psychology*, 13(1), 1–19. <https://doi.org/10.20982/tqmp.13.1.p001>

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67. <https://doi.org/10.1080/00401706.1970.10488634>

- Horn, J. L. (1972). State, trait and change dimensions of intelligence. *British Journal of Educational Psychology*, 42(2), 159–185. <https://doi.org/10.1111/j.2044-8279.1972.tb00709.x>
- Horn, J. L. (1976). Human abilities: A review of research and theory in the early 1970s. *Annual review of psychology*, 27(1), 437–485.
<https://doi.org/10.1146/annurev.ps.27.020176.002253>
- Horn, J. L. (2008). Spearman, g, expertise, and the nature of human cognitive capability. In P. C. Kyllonen, R. D. Roberts, & L. Stankov (Eds.), *Extending intelligence: Enhancement and new constructs* (pp. 185–230). Lawrence Erlbaum Associates.
- Horn, J. L., & Cattell, R. B. (1967). Age differences in fluid and crystallized intelligence. *Acta Psychologica*, 26(2), 107–129. [https://doi.org/10.1016/0001-6918\(67\)90011-X](https://doi.org/10.1016/0001-6918(67)90011-X)
- Horn, J. L., & Noll, J. (1997). Human cognitive capabilities: Gf-Gc theory. In D. P. Flanagan, J. L. Genshaft, & P. L. Harrison (Eds.), *Contemporary intellectual assessment: Theories, tests and issues* (pp. 53–91). Guilford Press.
- Irvine, S. H., & Kyllonen, P. C. (2002). *Item generation for test development*. Lawrence Erlbaum Associates.
- Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, 14(2), 288–294. <https://doi.org/10.3758/BF03194066>
- Kan, K.-J., Kievit, R. A., Dolan, C., & der Maas, H. van. (2011). On the interpretation of the CHC factor Gc. *Intelligence*, 39, 292–302. <https://doi.org/10.1016/j.intell.2011.05.003>

- Keith, T. Z., & Reynolds, M. R. (2010). Cattell–Horn–Carroll abilities and cognitive tests: What we’ve learned from 20 years of research. *Psychology in the Schools*, 47(7), 635–650.
<https://doi.org/10.1002/pits.20496>
- Koch, T., Romero, P., & Stachl, C. (2020). *Predicting age and gender from language, emoji, and emoticon use in WhatsApp instant messages*. PsyArXiv.
<https://doi.org/10.31234/osf.io/92ydh>
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 110(15), 5802–5805.
<https://doi.org/10.1073/pnas.1218772110>
- Kretzschmar, A., Spengler, M., Schubert, A.-L., Steinmayr, R., & Ziegler, M. (2018). The Relation of Personality and Intelligence—What Can the Brunswik Symmetry Principle Tell Us? *Journal of Intelligence*, 6(3), 30. <https://doi.org/10.3390/jintelligence6030030>
- Kuhn, M. (2008). Building predictive models in r using the caret package. *Journal of Statistical Software*, 28(5). <https://doi.org/10.18637/jss.v028.i05>
- Kuhn, M. (2020). *Caret* (6.0-85) [Computer software]. <https://cran.r-project.org/web/packages/caret>
- Kvålseth, T. O. (1985). Cautionary note about R^2 . *The American Statistician*, 39(4), 279–285.
<https://doi.org/10.1080/00031305.1985.10479448>
- Li, S. C., Lindenberger, U., Hommel, B., Aschersleben, G., Prinz, W., & Baltes, P. B. (2004). Transformations in the couplings among intellectual abilities and constituent cognitive processes across the life span. *Psychological Science*, 15(3), 155–163.
<https://doi.org/10.1111/j.0956-7976.2004.01503003.x>

- Little, T. D., Cunningham, W. A., Shahar, G., & Widaman, K. F. (2002). To parcel or not to parcel: Exploring the question, weighing the merits. *Structural Equation Modeling*, 9(2), 151–173. https://doi.org/10.1207/S15328007SEM0902_1
- Loevinger, J. (1965). Person and population as psychometric concepts. *Psychological Review*, 72(2), 143–155. <https://doi.org/10.1037/h0021704>
- Markus, K. A., & Borsboom, D. (2013). Reflective measurement models, behavior domains, and common causes. *New Ideas in Psychology*, 31(1), 54–64. <https://doi.org/10.1016/j.newideapsych.2011.02.008>
- McCrae, R. R. (2015). A more nuanced view of reliability: Specificity in the trait hierarchy. *Personality and Social Psychology Review*, 19(2), 97–112. <https://doi.org/10.1177/1088868314541857>
- McDonald, R. P. (1999). *Test theory: A unified treatment*. Lawrence Erlbaum Associate.
- McGrew, K. S. (2009). CHC theory and the human cognitive abilities project: Standing on the shoulders of the giants of psychometric intelligence research. *Intelligence*, 37(1), 1–10. <https://doi.org/10.1016/j.intell.2008.08.004>
- McNeish, D. (2018). Thanks coefficient alpha, we'll take it from here. *Psychological Methods*, 23(3), 412–433. <https://doi.org/10.1037/met0000144>
- Möttus, R. (2016). Towards more rigorous personality trait-outcome research: Personality traits and outcomes. *European Journal of Personality*, 30(4), 292–303. <https://doi.org/10.1002/per.2041>
- Möttus, R., Kandler, C., Bleidorn, W., Riemann, R., & McCrae, R. R. (2017). Personality traits below facets: The consensual validity, longitudinal stability, heritability, and utility of

- personality nuances. *Journal of Personality and Social Psychology*, 112(3), 474–490.
<https://doi.org/10.1037/pspp0000100>
- Mõttus, R., & Rozgonjuk, D. (2019). Development is in the details: Age differences in the Big Five domains, facets, and nuances. *Journal of Personality and Social Psychology*. Advance online publication. <https://doi.org/10.1037/pspp0000276>
- Mõttus, R., Sinick, J., Terracciano, A., Hřebíčková, M., Kandler, C., Ando, J., Mortensen, E. L., Colodro-Conde, L., & Jang, K. L. (2019). Personality characteristics below facets: A replication and meta-analysis of cross-rater agreement, rank-order stability, heritability, and utility of personality nuances. *Journal of Personality and Social Psychology*, 117(4), 1–16. <https://doi.org/10.1037/pspp0000202>
- Olaru, G., Witthöft, M., & Wilhelm, O. (2015). Methods matter: Testing competing models for designing short-scale Big-Five assessments. *Journal of Research in Personality*, 59, 56–68. <https://doi.org/10.1016/j.jrp.2015.09.001>
- Peng, P., Wang, T., Wang, C., & Lin, X. (2019). A meta-analysis on the relation between fluid intelligence and reading/mathematics: Effects of tasks, age, and social economics status. *Psychological Bulletin*, 145(2), 189–236. <https://doi.org/10.1037/bul0000182>
- Pettit, M. (2016). Historical time in the age of big data: Cultural psychology, historical change, and the Google Books Ngram Viewer. *History of Psychology*, 19(2), 141–153.
<https://doi.org/10.1037/hop0000023>
- Reckase, M. D. (2017). A tale of two models: Sources of confusion in achievement testing. *ETS Research Report Series*, 2017(1), 1–15. <https://doi.org/10.1002/ets2.12171>
- Rocca, R., & Yarkoni, T. (2020). *Putting psychology to the test: Rethinking model evaluation through benchmarking and prediction*. PsyArXiv. <https://doi.org/10.31234/osf.io/e437b>

Rounds, J., Ming, C. W. J., Cao, M., Song, C., & Lewis, P. (2016). *Development of an O*NET® Mini Interest Profiler (Mini-IP) for Mobile Devices: Psychometric Characteristics*.

Department of Labor O* NET Resource Center.

Savi, A. O., Marsman, M., van der Maas, H. L. J., & Maris, G. K. J. (2019). The wiring of intelligence. *Perspectives on Psychological Science*, 14, 1034–1061.

<https://doi.org/10.1177/1745691619866447>

Schneider, W. J., & McGrew, K. S. (2012). *The Cattell-Horn-Carroll model of intelligence*. In D. P. Flanagan & P. L. Harrison (Eds.), *Contemporary intellectual assessment: Theories, tests, and issues* (pp. 99–144). The Guilford Press.

Schroeders, U., Schipolowski, S., Zettler, I., Golle, J., & Wilhelm, O. (2016). Do the smart get smarter? Development of fluid and crystallized intelligence in 3rd grade. *Intelligence*, 59, 84–95. <https://doi.org/10.1016/j.intell.2016.08.003>

Schroeders, U., Wilhelm, O., & Olaru, G. (2016). The influence of item sampling on sex differences in knowledge tests. *Intelligence*, 58(3), 22–32.

<https://doi.org/10.1016/j.intell.2016.06.003>

Seeboth, A., & Möttus, R. (2018). Successful explanations start with accurate descriptions: Questionnaire items as personality markers for more accurate predictions: items as personality markers. *European Journal of Personality*, 32(3), 186–201.

<https://doi.org/10.1002/per.2147>

Shmueli, G. (2010). To explain or to predict? *Statistical Science*, 25(3), 289–310.

<https://doi.org/10.1214/10-STS330>

Slavin, R. E. (1994). Quality, appropriateness, incentive, and time: A model of instructional effectiveness. *International Journal of Educational Research*, 21(2), 141–157.

[https://doi.org/10.1016/0883-0355\(94\)90029-9](https://doi.org/10.1016/0883-0355(94)90029-9)

Soderberg, C. K. (2018). Using OSF to share data: A step-by-step guide. *Advances in Methods and Practices in Psychological Science*, 1(1), 115–120.

<https://doi.org/10.1177/2515245918757689>

Steger, D., Schroeders, U., & Wilhelm, O. (2019). On the dimensionality of crystallized intelligence: A smartphone-based assessment. *Intelligence*, 72, 76–85.

<https://doi.org/10.1016/j.intell.2018.12.002>

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.

<https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>

Tibshirani, R. (2011). Regression shrinkage and selection via the lasso: A retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(3), 273–282.

<https://doi.org/10.1111/j.1467-9868.2011.00771.x>

Tucker-Drob, E. M. (2009). Differentiation of cognitive abilities across the life span.

Developmental Psychology, 45, 1097–1118. <https://doi.org/10.1037/a0015864>

Tucker-Drob, E. M., & Salthouse, T. A. (2008). Adult age trends in the relations among cognitive abilities. *Psychology and Aging*, 23(2), 453–460. [https://doi.org/10.1037/0882-](https://doi.org/10.1037/0882-7974.23.2.453)

[7974.23.2.453](https://doi.org/10.1037/0882-7974.23.2.453)

van der Maas, H. L. J., Dolan, C. V., Grasman, R. P. P. P., Wicherts, J. M., Huizenga, H. M., & Raijmakers, M. E. J. (2006). A dynamical model of general intelligence: The positive

manifold of intelligence by mutualism. *Psychological Review*, 113(4), 842–861.

<https://doi.org/10.1037/0033-295X.113.4.842>

van der Maas, H. L. J., Kan, K.-J., & Borsboom, D. (2014). Intelligence is what the intelligence test measures. Seriously. *Journal of Intelligence*, 2(1), 12–15.

<https://doi.org/10.3390/jintelligence2010012>

Wilhelm, O., Hildebrandt, A., & Oberauer, K. (2013). What is working memory capacity, and how can we measure it? *Frontiers in Psychology*, 4.

<https://doi.org/10.3389/fpsyg.2013.00433>

Wittmann W.W. (1988) Multivariate Reliability Theory. In J.R. Nesselroade, & R.B. Cattell (Eds.), *Handbook of Multivariate Experimental Psychology* (pp. 505–560). Plenum Press.

https://doi.org/10.1007/978-1-4613-0893-5_16

Wittmann, W. W., & Süß, H.-M. (1999). *Investigating the paths between working memory, intelligence, knowledge, and complex problem-solving performances via Brunswik symmetry*. In P. L. Ackerman, P. C. Kyllonen, & R. D. Roberts (Eds.), *Learning and individual differences: Process, trait, and content determinants* (pp. 77–108). American Psychological Association. <https://doi.org/10.1037/10315-004>

Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100–1122.

<https://doi.org/10.1177/1745691617693393>

Yu, C. Y. (2002). *Evaluating cutoff criteria of model fit indices for latent variable models with binary and continuous outcomes* [Doctoral dissertation]. University of California, Los Angeles. <https://www.statmodel.com/download/Yudissertation.pdf>

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320.

<https://doi.org/10.1111/j.1467-9868.2005.00503.x>

Appendix

The loss function of the standard linear regression is denoted by:

$$\sum_{i=1}^N (y_i - \hat{y}_i)^2 = \sum_{i=1}^N \left(y_i - \sum_{j=1}^p w_j \times x_{ij} \right)^2 \quad (1)$$

with x_{ij} as the manifest score of a person i on an item j and w_j as weights or regression coefficients. Ridge regression restricts the coefficients by introducing the sum of the squared weights weighted by the shrinkage factor λ as a penalty term.

$$\sum_{i=1}^N (y_i - \hat{y}_i)^2 = \sum_{i=1}^N \left(y_i - \sum_{j=1}^p w_j \times x_{ij} \right)^2 + \lambda \sum_{j=1}^p w_j^2 \quad (2)$$

In contrast, LASSO regression uses the sum of the absolute weights instead of the squared weights the penalty term, which results in a more parsimonious (and often preferred) procedure (Zou & Hastie, 2005).

$$\sum_{i=1}^N (y_i - \hat{y}_i)^2 = \sum_{i=1}^N \left(y_i - \sum_{j=1}^p w_j \times x_{ij} \right)^2 + \lambda \sum_{j=1}^p |w_j| \quad (3)$$

Elastic net regression relies on two regularization parameters—the shrinkage parameter λ and the penalty parameter α , which balances between ridge and LASSO regression:

$$\sum_{i=1}^N (y_i - \hat{y}_i)^2 = \sum_{i=1}^N \left(y_i - \sum_{j=1}^p w_j \times x_{ij} \right)^2 + \lambda \sum_{j=1}^p (\alpha w_j^2 + (1 - \alpha) |w_j|) \quad (4)$$

Model performance was evaluated with the explained variance (R^2), the Root Mean Squared Error ($RMSE$), and the Mean Absolute error (MAE , see also Kvålseth, 1985):

$$R^2 = \text{cor}(y, \hat{y})^2 \quad (5)$$

$$RMSE = [\sum (y - \hat{y})^2 / n]^{1/2} \quad (6)$$

$$MAE = \sum |y - \hat{y}| / n \quad (7)$$

Content

Table OS 1. Data Cleaning Procedure.....	2
Table OS 2. Item difficulties.	3
Table OS 3. Accuracy Indices in the Prediction of Age with Knowledge Scores and Items.....	5
Figure OS 1. Distribution of Age in the Analysis Sample.	6
Figure OS 2. Overfit of Linear Regression vs. Elastic Net Depending on Sample Size.....	7
Figure OS 3. Distribution of Residuals for the Elastic Net Regression.	8

Table OS 1. Data Cleaning Procedure.

<i>n</i>	Reason for exclusion
1,721	Total sample. (The survey took place in June and July 2020)
– 35	Data recording problems
– 2	Attention check: „What is the second letter in the alphabet?“
– 20	Duration < 15 min.
– 20	Insufficient language skills
– 7	Older than 70 years ^a
– 8	Total score below guessing probability
1,629	Analysis sample

Note. ^a Participants older than 70 years and participants that failed the attention check were automatically removed by the panel provider.

Table OS 2. Item difficulties.

Humanities	Art		Literature		Music	
	Item	<i>P</i> (x)	Item	<i>P</i> (x)	Item	<i>P</i> (x)
	kun048	.31	lit099	.24	mus028	.33
	kun007	.32	lit021	.30	mus008	.33
	kun009	.46	lit121	.46	mus073	.38
	kun018	.42	lit100	.54	mus097	.55
	kun098	.55	lit039	.59	mus013	.62
	kun012	.49	lit091	.44	mus045	.58
	kun100	.50	lit005	.72	mus048	.69
	kun030	.71	lit030	.76	mus046	.79
	kun028	.58	lit092	.88	mus003	.66
	kun081	.84	lit042	.77	mus060	.88
	.52 (.16)		.57 (.20)		.58 (.18)	

Life Sciences	Nutrition		Health		Medicine	
	Item	<i>P</i> (x)	Item	<i>P</i> (x)	Item	<i>P</i> (x)
	ern097	.36	gsu028	.34	med043	.25
	ern102	.38	gsu012	.37	med032	.48
	ern089	.43	gsu008	.56	med040	.52
	ern015	.50	gsu057	.58	med093	.50
	ern070	.50	gsu015	.54	med075	.48
	ern100	.48	gsu083	.63	med100	.70
	ern034	.74	gsu013	.60	med009	.72
	ern044	.82	gsu055	.71	med053	.76
	ern014	.86	gsu011	.71	med045	.86
	ern091	.87	gsu009	.93	med106	.89
	.59 (.19)		.60 (.16)		.62 (.19)	

Natural Sciences	Chemistry		Math		Physics	
	Item	$P(x)$	Item	$P(x)$	Item	$P(x)$
	che009	.32	mat025	.29	phy105	.32
	che010	.30	mat070	.38	phy022	.42
	che030	.39	mat032	.40	phy081	.41
	che013	.47	mat085	.33	phy003	.52
	che079	.44	mat090	.32	phy108	.54
	che041	.52	mat019	.42	phy015	.54
	che078	.50	mat024	.45	phy034	.57
	che091	.61	mat027	.67	phy066	.74
	che017	.57	mat031	.65	phy032	.85
	che011	.69	mat047	.75	phy082	.79
	.48 (.12)		.47 (.16)		.57 (.16)	

Social Studies	Law		Politics		Economics	
	Item	$P(x)$	Item	$P(x)$	Item	$P(x)$
	jur105	.31	pol103	.25	wir106	.36
	jur074	.50	pol069	.43	wir090	.51
	jur042	.59	pol046	.37	wir061	.45
	jur020	.61	pol015	.43	wir107	.50
	jur008	.63	pol038	.56	wir033	.46
	jur104	.72	pol076	.50	wir088	.60
	jur107	.73	pol003	.68	wir082	.61
	jur012	.83	pol008	.70	wir062	.70
	jur092	.78	pol062	.72	wir010	.89
	jur056	.69	pol082	.91	wir018	.94
	.64 (.14)		.55 (.19)		.60 (.18)	

Note. $N = 1,629$.

Table OS 3. Accuracy Indices in the Prediction of Age with Knowledge Scores and Items.

Knowledge Level	Data Set	R^2 $M (SD)$	$RMSE$ $M (SD)$	MAE $M (SD)$
Gc	Train	.02 (.00)	14.54 (.06)	12.53 (.06)
	Test	.02 (.01)	14.57 (.22)	12.56 (.20)
Area	Train	.24 (.01)	12.78 (.09)	10.56 (.09)
	Test	.24 (.04)	12.82 (.35)	10.61 (.33)
Domain	Train	.30 (.01)	12.33 (.09)	10.11 (.09)
	Test	.28 (.04)	12.46 (.37)	10.23 (.34)
Item	Train	.60 (.01)	9.36 (.11)	7.52 (.10)
	Test	.54 (.03)	10.02 (.37)	8.06 (.31)

Note. $N = 1,629$. The table shows the means (and standard deviations) of the prediction accuracy across 100 random splits of the total sample into a training-validation and a testing sample across the four levels of the knowledge hierarchy (gc, area, domain, and item). $RMSE = [\sum(y - \hat{y})^2/n]^{1/2}$; $MAE = \sum|y - \hat{y}|/n$

Figure OS 1. Distribution of Age in the Analysis Sample.

Note. $N = 1,629$.

Figure OS 2. Overfit of Linear Regression vs. Elastic Net Depending on Sample Size.

Note. The figure shows the average overfit for linear regression vs. elastic net regression across 100 random splits of samples of varying sizes (80% training-validation and 20% testing).

Figure OS 3. Distribution of Residuals for the Elastic Net Regression.

Note. Figure OS 3 shows exemplarily the distribution of the fitted values with their respective residuals for the elastic net regression in the holdout sample. Since participants' age ranged between 20 and 70 years, the residual plot has the typical diamond shape of a double truncated distribution (see left panel). Please note that even in the current case where homoscedasticity is violated, the estimated parameters are still unbiased estimates for the population parameters (while standard errors are biased, see White, 1980). The assumption of normally distributed residuals is met (see right panel).