



Analysing postprandial amino acid responses in crossover studies with the aaresponse package for R

Ron Wehrens¹ · Jasper Engel¹ · Jurriaan Mes² · Aard de Jong² · Diederik Esser²

Received: 3 July 2023 / Accepted: 22 February 2024
© The Author(s) 2024

Abstract

Nowadays, a healthier and more sustainable lifestyle is the subject of much research. One example is the use of crossover trials to investigate the uptake of proteins, usually from alternatives to animal-based sources, by healthy volunteers. The data analysis is complex and requires many decisions on the part of the scientists involved. Such a process can be streamlined and made more objective and reproducible through bespoke software. This paper describes such a software package, **aaresponse**, for the R environment, available as open source. It features ample visualization functions, supports consistent curation strategies, and compares amino acid uptake of different protein meals (interventions) through the use of mixed models analysing parameters of interest like the area under the curve (AUC). The defining feature is the use of parametric curves to fit the amino acid levels over time, increasing the robustness of the approach and allowing for more strict quality control strategies.

Keywords Protein digestibility · Amino acid levels · Blood samples · Data analysis · Statistical software

Introduction

In the search for a more sustainable way of life, much attention is focused on plants as an alternative source of proteins. Not all plant proteins can be taken up easily by the human body, and much research has been devoted to assess digestibility characteristics of plant proteins into amino

acids (AAs). The nine essential or indispensable amino acids (EAAs: phenylalanine, valine, threonine, tryptophan, methionine, leucine, isoleucine, lysine, and histidine) are of special importance, since humans cannot synthesize these themselves.

Several methodologies have been developed to report protein digestibility. The Protein Digestibility-Corrected Amino Acid Score (PDCAAS) assesses food protein quality taking into account human requirements (by focusing on the limiting type of amino acid) and digestibility, determined over the total digestive tract. In general, PDCAAS estimates are too optimistic, and better estimates are obtained by the more recent Digestible Indispensable Amino Acid Score (DIAAS) (FAO Expert Consultation 2013), adopted by the FAO/WHO and by the US FDA as the preferred best method for determining protein quality. Both methods, however, require animal experiments, and form less-than-perfect models for human digestion. Whereas in vitro models may partly replace animal studies (Butts et al. 2012), the results are still not good enough for a reliable comparison of the food values of different protein sources. Finally, true amino acid digestibility in humans can be measured using stable-isotope techniques, but this is rather challenging from a methodological perspective and, in addition, expensive.

A much simpler and yet direct approach is to measure postprandial amino acid concentrations in blood to estimate

Communicated by F. Polticelli.

✉ Ron Wehrens
ron.wehrens@wur.nl
Jasper Engel
jasper.engel@wur.nl
Jurriaan Mes
jurriaan.mes@wur.nl
Aard de Jong
aard.dejong@wur.nl
Diederik Esser
diederik.esser@wur.nl

¹ Biometris, Wageningen University and Research, Droevendaalsesteeg 1, 6708 PB Wageningen, The Netherlands

² Wageningen Food and Biobased Research, Wageningen University and Research, Bornse Weiland 9, 6708 WG Wageningen, The Netherlands

the digestibility potential of a protein source, and this is the technique we are focusing on in the current paper. Note that the levels of amino acids in blood only partially reflect the digestibility of a protein source: they are influenced by many factors, including breakdown and endogenous synthesis. The latter factor is not present for the EAAs, which therefore constitute a valuable measure for protein quality comparisons.

To compare protein digestibilities with a reference protein, typically crossover experiments are performed, in which healthy volunteers are given different protein meals, separated by a significant time so that carryover effects are avoided (see, e.g. Mes et al. 2022; Wegrzyn et al. 2022; Liu et al. 2019). During several hours after the meal, blood samples are taken and analysed for amino acid levels. Statistical analysis of these time course data then allows to draw conclusions about amino acid kinetics and protein uptake, indicated by parameters of interest (PoIs), typically AUC, the area under the curve, peak height, and the speed of uptake (the time to the peak maximum). Whereas these values per se may not be easy to interpret, comparison with a reference protein, often easily digestible, allows one to put things in perspective and, e.g., set up a ranking of the digestibility of proteins from different sources. When the same reference is used in different studies, the corresponding results can be directly compared, even when the individual studies differ in their design or measurement setup.

Data analysis for such experiments can be done in many ways, but is certainly not trivial. This paper proposes a simple and principled approach, implemented in open-source software, and comprising several steps. The first, to be executed for all amino acids in all participants separately, is to obtain estimates for the PoIs. Alternatively, one can also focus on AA totals, either considering all AAs, only the EAAs, or a specific AA subset. The simplest possible approach would approximate peak height by the differences between the highest and lowest values in each time series and take the time of the largest value as the time to the peak maximum. A parameter like AUC can be assessed by the trapezoid rule. These estimates, however, are very sensitive to noise and artefacts in the data. A more sophisticated strategy, less affected by experimental variation, is to fit parametric curves through the amino acid time series, and to extract estimates of the PoIs from the estimated curve parameters.

In the second analysis step, these PoIs are compared between the different proteins in the experiment. This can be done in several different ways, the simplest being an analysis of variance. A more powerful approach is based on linear mixed models. Their benefits include being applicable in unbalanced situations, and being much more flexible in the description of the model and the contrasts of interest.

In such an analysis, quality checks are needed at several stages to prevent individual outlying observations to have

a large effect on the eventual outcome of the experiment. Sometimes, the raw data show individual measurements that deviate quite dramatically, or points showing a behaviour that is consistent across amino acids, but inconsistent in time, indicating that perhaps a sample is compromised. In other cases, strange values are observed at the level of the PoIs (e.g. implausible peak heights, maxima at the very first, or very last time point). The curve fits used to describe the amino acid time courses provide an intermediate possibility for curation: it is usually feasible to define ranges of suitable values for curve parameters, and to automatically flag cases that are outside these ranges for further inspection.

This paper describes a principled approach to analysing the data of protein digestibility experiments. To make this widely accessible, we have provided an implementation as an open-source package for the R language (R Core Team 2022), **aaresponse**. The package includes several visualization tools for quality assessment and curation of the data as well as intermediate results, curve fit procedures to describe the amino acid time courses, and mixed-model functionality to compare proteins. Although developed with protein digestibility in mind, the package can also be used to analyse data from other types of nutrient-uptake experiments (glucose, fat, micronutrients,...) for which time course data are available.

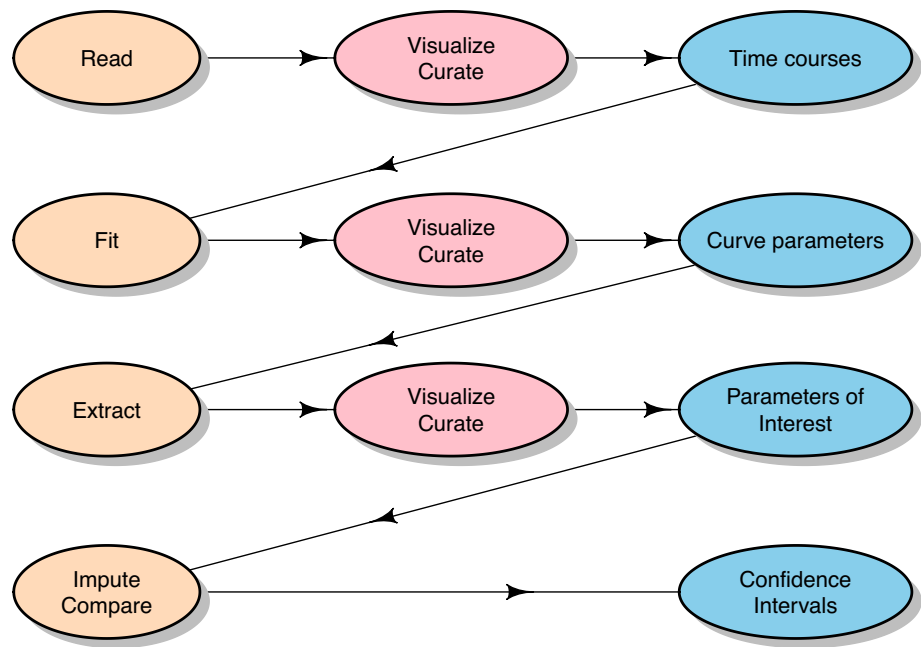
The **aaresponse** package for R

This section describes the **aaresponse** package, written in R, and developed for the analysis of protein uptake experiments using crossover designs. The methodology has been developed and used in several projects leading to peer-reviewed papers (Mes et al. 2022; Esser et al. 2023; van Dam et al. 2023; Ummels et al. 2023; Roelofs et al. 2024). The examples focus on a data set from one of these, which is distributed with the package and analysed more rigorously elsewhere (Esser et al. 2023). The R code leading to the plots and tables is discussed in the R package “vignette”, basically a walkthrough to obtain the results. An overview of the analysis strategy (and the package) is given in Fig. 1—note that at any point in time, earlier results may be revised by the user, and the analysis sequence can be continued from that point onwards (such feedback loops have not been included in the figure to avoid clutter). The following paragraphs describe the main functionalities of the package in more detail.

Data visualization

Data coming from crossover experiments are complex: typically 10–20 participants are given protein meals in two to three periods, and eight to ten blood samples are taken and analysed for up to 20 amino acids or amino acid aggregates.

Fig. 1 Flowchart of the analysis strategy in **aaresponse**. The left column contains the main steps, the middle column contains curation and associated visualization steps, and the right column contains the results of each step



This leads to thousands of data points. An appropriate visualization of the raw time courses is very powerful, in that it very quickly gives an idea of data quality, differences between proteins, differences between participants, and differences between amino acids. It also can point to anomalies, such as outlying measurements, and unexpected patterns: in one case, we were able to immediately spot a potential label mix-up which was later confirmed by going back to the original laboratory data. The **aaresponse** package contains several standard visualizations for amino acid time courses, allowing users to easily zoom in on areas of interest. In addition, several types of principal component analysis (PCA, Jackson 1991; Wehrens 2020), allowing a multivariate view of the data, are supported.

Fitting curves for amino acid time courses

Manual analysis of the amino acid time courses is labour-intensive, prone to errors, and, perhaps most importantly, not very robust against outlying observations. In our experience, these occur quite frequently, so our data analysis pipeline should be able to handle them. It is easy to imagine other potential problems—even the level of the baseline is hard to estimate in many cases. The crucial advantage of fitting a curve with a predefined shape through such data is that it includes an inherent smoothing, basically using information from all time points to determine appropriate curve parameters. The approach using curve fitting has a couple of additional advantages:

1. time points need not be spaced equidistantly, and may even vary within the experiment;

2. individual deviating values usually only have a minor effect on the parameters of a fitted curve;
3. missing values are allowed;
4. curves can be fit even when the values have not completely returned to baseline;
5. PoIs can be calculated analytically from the curve parameters.

Of course, the parametric form of the curve needs to be chosen carefully. A suitable model is given by the Wood curve, originally used to describe lactation behaviour in cattle (Wood 1967):

$$y(t) = d + at^{mc} e^{-ct},$$

where $y(t)$ is the level of a particular amino acid at time t . An example is shown in Fig. 2. Parameter d is the baseline, not present in the original paper by Wood, but added in other publications (see, e.g. Engel 2003). Parameter a is related to the height of the peak, and the other parameters determine the shape of the curve. Parameter m itself determines the time to the peak maximum, and mc (in Wood's publications indicated as b) is related to the increasing part of the curve; c itself describes the return to the starting position. Since the curve parameters in themselves have meaning, it can be useful to explicitly inspect them, e.g. to identify cases where curve fitting has led to unexpected results. Their values may even be a topic of study in particular situations.

Note that many other methods can be used for drawing smooth curves through a series of points (e.g., splines): the Wood curves employed here have the advantage of providing interpretable curve parameters, as well as having analytical

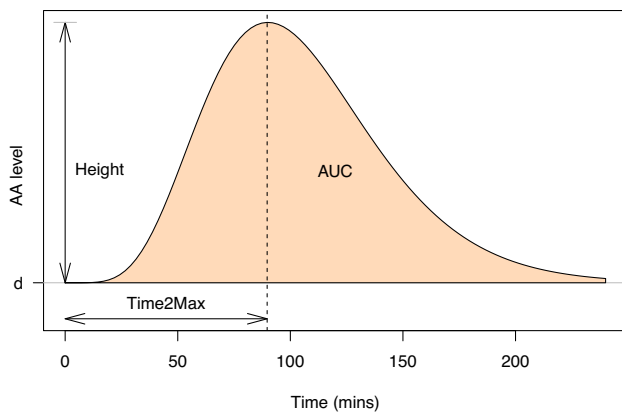


Fig. 2 A Wood curve and the parameters extracted from it. In this case, the AUC value is calculated up to 4 h

solutions for the PoIs. Moreover, they lead to curves that always return to the fitted baseline (although this may happen only after the last measurement), something that with more flexible curve fitting methods like splines is not necessarily the case.

If there would be no baseline ($d = 0$), the model could be fitted by taking logarithms and applying simple least-square regression. Here, we have to take account of the fact that all participants will have a non-zero (and different) baseline, corresponding to different AA levels in the sober state in which they start the experiment. This necessitates an optimization approach for finding the optimal values of the model parameters. In **aaresponse**, a nonlinear least-squares approach is used, performing 500 optimizations with different random starts and choosing the eventual best result.

Given that four parameters are estimated (a , m , c and d), we need to have at least five time points to obtain a fit in the first place, and preferably eight or more for a reliable result. However, even with a reasonable number of time points, it is important to assess the qualities of the fits. In some cases, the data may not conform to the Wood curve: in practice, one may see, e.g., amino acid levels that drop below the starting value, either due to natural variability or through metabolic processes not directly related to the experiment. In such a case, the fitted curves will be a compromise, usually putting the estimated baseline somewhere between the starting and ending values. Whether this is a useful approximation needs to be assessed on an ad hoc basis. Other situations potentially causing problems in curve fitting are very narrow peaks, e.g. containing just one point above the baseline, or completely flat profiles.

The optimization process will limit the ranges of possible values for the curve parameters, but it is unwise to choose them very narrow—that could lead to a failure to converge. It is better to start with liberal ranges, and to narrow down to plausible values in a subsequent curation step. Any values

outside these plausible ranges can be replaced by NA, the R coding for “not available”. If needed, one can always repeat the fit for specific cases with more narrow boundaries. If these plausible ranges are not known, or for some reason not appropriate, one can always revert to visual inspection: by plotting the sorted values for each of the four curve parameters one can easily identify any extreme observations, and these in turn can be related to specific time courses, which again can be inspected. Here, we only report results obtained with automatic curation using the defaults of **aaresponse**, eliminating nothing but the most glaring outliers—by necessity, defined by the user. However, it is important to note that it is actually preferable to err on the safe side: removing one time curve too many will likely not influence the result very much, whereas including a really strange observation may wreak havoc.

PoIs

Once we have the curve estimates, the PoIs follow analytically (Rook et al. 1993). The three PoIs we consider here are AUC, peak height, and time to peak maximum (see Fig. 2). We have already seen that the time to the peak maximum (Time2Max) is given by m ; the AUC value, relative to the baseline d ,¹ is given by

$$\text{AUC} = a/c^{(b+1)}\Gamma(ct_f, b + 1),$$

where $\Gamma(x, q) = \int_0^x z^{q-1} e^{-z} dz$ is the lower incomplete gamma function and t_f the time point up to which the AUC value should be computed. The peak height of curve $y(t)$ (relative to baseline d) is given by

$$\text{Height} = \max(y(t)) - d = a(b/c)^b e^{-b}.$$

Also here, the PoIs may be compared to plausible ranges, or potential outlying values may be identified by visual inspection. Note that plausible ranges may depend on the type of variable—for amino acid totals, e.g., one would expect much larger peak heights and AUC values than for individual amino acids. Again, in cases where the values are clearly not within plausible regions, they could be replaced by NA.

Imputation

In some cases, no suitable fit can be found, either because the curve fitting optimization does not converge, or because the curve parameters or PoIs are outside the optimization boundaries. In such cases, values will be stored as missing (NA) and will be ignored in the subsequent statistical analysis. It can be argued, however, that a major cause of these missing values

¹ This is often referred to as the incremental AUC, or iAUC.

is no uptake by the participant, or an uptake that is too low to measure, and that these missing values are actually highly informative. If that is the case, very low values for AUC and peak height would be more appropriate than missing values.

The default analysis considers only cases without missing values (responders), and this should be clearly stated to avoid incorrect interpretation of the results. An alternative is to impute the missing values. For the time to the peak maximum, this is hard to imagine (if there is no peak, where is the maximum?), but for AUC and peak height, estimates based on the distribution of the raw data are shown to be highly effective, especially for low-amplitude time courses. Basically, the idea is to use the 0.2 and 0.8 quantiles of the raw data (basically a robust estimate of the spread in the time course for one amino acid in one participant) to set up a regression model for either peak height or AUC value. After the imputation, the comparison of protein interventions can proceed in the normal way, now presenting results for the whole population (responders as well as non-responders). Note that with large numbers of non-response, the assumptions of the mixed model may be violated, so this needs to be checked. In any case, it is advisable to perform analyses both including and excluding non-responders, and to compare the results: where the differences are large or unexpected, further inspection is needed. For more details, we refer to the vignette included in the R package.

Further statistical analysis

In some cases, simply obtaining good estimates of the PoIs and the corresponding confidence intervals is the end goal of the experiment. However, often one is interested in comparisons, e.g. between different protein meals. The statistical analysis in the final step of the `aareponse` pipeline is geared towards digestibility of different proteins, both in terms of completeness and speed of uptake. Note that many more analyses are possible, and are usually easily implemented in the R environment. Here, we focus on the comparison of one or more proteins with a reference protein.

For every PoI, the same approach is followed. The first basic step is to average the PoIs over all participants, leading to tables with averages and standard deviations. An analysis of variance could be used to investigate whether there are differences between different protein sources, but as stated in the introduction section, linear mixed models are more suited here. These are fitted describing the PoI, in the simplest possible form, as a function of the protein meal and the participant:

$$\text{PoI} = \text{Intervention} + (1 \mid \text{Participant}).$$

In this case, we are not interested in separate coefficients for participants, but rather in the variability in the responses of the participants. This is formalized by taking `Participant` as a random effect—we use the Wilkinson–Rogers notation

employed by the `lme4` package (Bates et al. 2015), hence the $(1 \mid \text{Participant})$ term in the model definition. Parameters such as `Intervention` (describing the protein meal) are fixed effects: here we *are* interested in the coefficients.

The equation can be extended in several different ways. Often, it is wise to include a fixed factor for `Period`, to catch any time dependencies—this is also the default in `aareponse`. We do not expect such effects to be significant, since the different protein meals are consumed far apart in time, and carryover or learning effects are very unlikely, but it is good to check. In the formulation mentioned above, standard model assumptions hold: residuals are independently and identically distributed according to a normal distribution, and independent from the participant effects. More complicated mixed models allowing, e.g., different variances for different proteins, can be fit as well. The resulting model contains coefficients and confidence intervals for all protein PoIs. In addition, specific contrasts may be defined, by default comparing a test protein with the reference protein. This has been implemented using an *F*-test with the Kenward–Roger approximation (Kenward and Roger 1997).

The final decision concerns the type of contrast that is employed: for peak height, and the time to peak maximum, this is commonly the difference. For AUC, however, often the ratio is investigated. In `aareponse`, both differences and ratios are supported—the latter is the default for AUC, and the former for the other two PoIs. In the plots, the line of “no difference” will be set accordingly: i.e., at zero for differences, and at one for ratios.

Implementation and availability

The `aareponse` is written in R, using several packages for specific elements of the workflow. The main ones are the following: fitting Wood curves is done using package `nls.multstart` (Padfield and Matheson 2020), the mixed models for comparing different protein interventions are due to `lme4` (Bates et al. 2015), supported by the `emmeans` (Lenth 2023) and `multcomp` (Hothorn et al. 2008) packages, and the graphics are based on `lattice` (Sarkar 2008). The `aareponse` package is freely available from github at the address: <https://github.com/Biometris/aareponse>.

The installation is done using standard R practice—the only requirement is a working and not-too-outdated version of R. The package contains an extensive vignette showing the use of the functions in the package, using the same data as the current paper.

Case study: the SuPro data

Here, results are presented from the SuPro study (Esser et al. 2023), an investigation comparing two proteins, bovine plasma concentrate (BP) and corn protein concentrate (CP), to whey protein concentrate (WP), which is used as the reference protein. Four men and eight women, between 55 and 70 years old, consumed the three different protein shakes with 1 week between each pair of meals, and ten blood samples were taken within 3 h after each intervention.

The raw data are stored as a table with columns for Participant, Period, Time, Intervention, columns for different individual amino acids and, possibly, aggregated amino acid totals. For illustration purposes, the analysis here focuses on the essential amino acids only.

One example of the raw data visualizations provided by the **aaresponse** package is shown in Fig. 3, concentrating the essential amino acids in three participants. It is clear that the three proteins give rise to different responses, and also that there are differences between the participants. Curves for the non-essential amino acids show very similar patterns, indicating that at this timescale the variation in blood levels is dominated by the protein meal, something that is consistently the case in all data sets we have examined so far.

As an additional visualization tool, the **aaresponse** package offers several different ways of applying PCA to the data. The differences lie in the way the data are organized. One extreme is to create a matrix where each row contains all time courses for a participant, and the other extreme is to simply put all time courses (one AA from one participant having had one protein meal) in separate rows. An example of an intermediate approach is to define each row as the concatenation of the time courses for a particular participant–protein combination. The PCA score plot for the latter approach is shown in Fig. 4: each symbol represents the time course for a participant having had a particular protein meal. Clearly, one can see that the three proteins show consistent differences in uptake—the effect is the dominant factor and completely determines the first PC. At the same time, the variation between participants, in this case comprising the effects on all amino acids simultaneously is still appreciable. In this case, each time course was centered individually, so that differences in baseline are eliminated to a large extent; other options are available, too.

Figure 5 highlights the benefits as well as the limitations of using curve fits to estimate PoIs. The figure shows fitted curves for the WP and BP data of two of the participants from Fig. 3. In the left panels, one can immediately recognize that the curves around the third time point are less close to the data points. The effect of these deviations on the PoIs, however, is mostly limited: Height and Time2Max are almost unchanged. For AUC the effects are

larger, since the outlying point leads to peaks that seem to start later, and therefore have a lower surface area. With other approaches, this could even be much more extreme: when using the trapezoid rule for determining AUC values, for example, it is not uncommon to stop the integration as soon as a value drops below the first data point which in most of the examples in Fig. 5 would lead to extremely small values. Curve fit-based PoIs are more robust, but even so it could be appropriate to consider removing such points from the data, especially when the same effect is visible in several other amino acids as well, suggesting a problem at the sample level.


The right panels show a case in point: also there we can see outlying values at the third time point in several AAs, but because this is in the positive direction the curve fit algorithm tries to fit a very narrow curve through the outlying point alone in several cases. Here, one could definitely consider to remove the whole time point—if one decides to keep it in, the subsequent curation stages will most likely remove all narrow fits. This curation prevents such outlying points from causing damage, but it also decreases the number of data points in the final comparison, and hence the power of the analysis. In that sense, eliminating a single time point would be preferable. Obviously, such manual curation steps should be used sparingly, with caution, and be well documented. Since in this paper the onus is on presenting the methodology, here we will leave both time points unchanged, and consider the complete data set in the analysis.

Fitted peak parameters can, as already mentioned, be curated. Since the curve parameters do not have obvious regions in which one could expect “good” results to fall, it is wise to assess this on a case-by-case base and use visualization. Typically, the boundary values can be set by plotting the parameter distributions over all curves—all values outside the “normal” areas could be replaced with NA. This would also lead to NA estimates for the corresponding PoIs. One would expect only a very small fraction of all values to be curated.

The curation of PoIs is easier, since these are directly interpretable, and the effects of the curation are shown in Fig. 6. The left panel shows estimated PoIs for the essential amino acids from Fig. 3—the right panel shows the same values, after curation using package defaults (positive AUC values, peak heights between 0 and 1000, and peak times between 15 and 200 min). Note that now PoIs are shown for all participants (y axis) and all three proteins. In both Leu and Phe, it is easy to see a removed point in the Time2Max column, leading to a much narrower range for the curated data. Note that similar plots are available to assess the effect of the curation of the curve fit parameters.

The (curated) results for the PoI estimates in Fig. 6 already show clear patterns: CP usually shows a later peak

Fig. 3 Visualizing the raw data in **aaresponse**. The nine essential amino acids (rows in the figure) are shown for three participants: 9, 10, and 11. Abbreviations: see text



than BP, and WP is the protein that shows the fastest uptake; in addition, WP usually shows the highest uptake (AUC as well as Height), although Leu seems to be an exception. The uptake for Phe is low in absolute terms—note that this

is also influenced by the protein compositions of the three protein concentrates.

Finally, the confidence intervals for all PoI comparisons (with WP as a reference) can be presented in a simple graphic such as the one shown in Fig. 7. Significant effects,

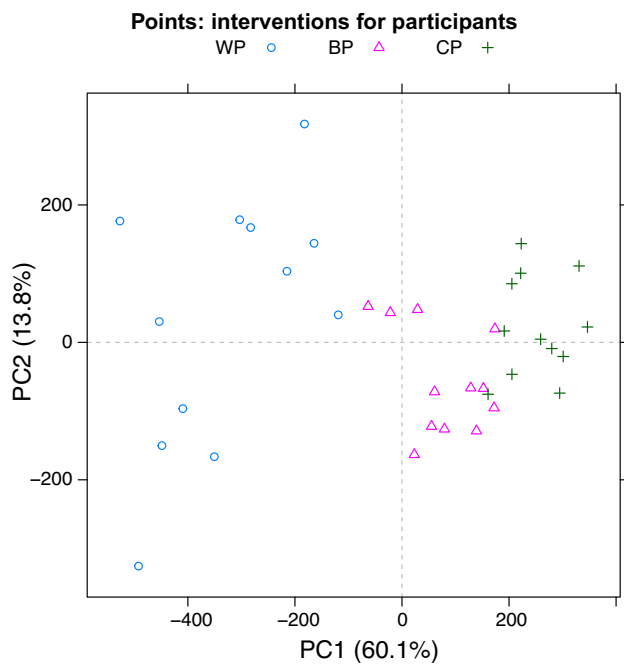


Fig. 4 PCA score plot: the 36 symbols correspond to the 12 participants having had one of the three meals. The differences between the three interventions dominate the first PC

not including the null hypothesis value indicated with a vertical dashed line in the confidence interval, are indicated in red; otherwise black is used. BP leads to lower AUC values for the essential amino acids, although the effect is not significant for Phe. CP shows two amino acids with increased AUCs, and a significantly lower AUC for most others. Clearly, all peak heights are significantly lower in the test proteins than in the reference protein. Nearly all estimates for the time to the peak maximum are larger in the test proteins than in the reference, although not all differences are significant. Plots like these give an immediate overview of all results, and although differences between amino acids will exist, sometimes leading to a confusing picture, general patterns should be quite obvious. It is worth noting that not all CIs are equally wide. This is the result of several effects. For one, the variation in the data is not constant—here, this is the dominating effect. In addition, if no imputation has been performed (as in the current example), the number of values taken into account may not be equal. In the current data set, very few missing values are present, and imputation would lead to virtually the same CIs as the ones shown here.

Discussion and conclusions

Performing crossover experiments in which postprandial amino acid levels are determined in the blood of the participants is a relatively simple and very straightforward

approach to assess digestibility characteristics of different proteins. We have described a principled strategy for analysing data from such experiments based on two main building blocks: curve fitting and linear mixed models.

The advantages of using fitted curves rather than the raw data are numerous. To start with, the influence of outlying points is greatly diminished because of the inherent smoothing characteristics associated with fitting parametric curves of a specific shape. Of course, this is no guarantee that all these problems will be avoided in the future: Fig. 5 is a case in point. Rather, the fitted curve parameters provide an additional checkpoint where the user can see whether these values are within plausible ranges, or not.

Furthermore, cases where no realistic curve can be fitted (often because there is no meaningful signal in the data; less often because of outliers or noise in the data) are easily identified, either in an automatic fashion or on the basis of visual inspection of the fits. Curves can also be fitted when the time points are not equidistant, or when the peak values have not quite returned to baseline level at the end of measurements—both cases that are problematic in approaches that are not based on parametric models. Parameters of interest, such as the location or height of the peak, can be calculated analytically.

Mixed models are used to compare the PoIs of the proteins in the experiment. In particular, methods have been implemented for comparing several test proteins with a reference protein. These comparisons have been set up in such a way as to provide maximal sensitivity, and in that sense are more flexible as well as more specific than alternative approaches like repeated-measures ANOVA: balanced data are no longer required, and specific contrasts such as the comparison to a reference intervention can be highlighted. The results of the comparisons can be presented as p values (suitably corrected for multiple testing) and as plots of confidence intervals. The latter are not only more easy to grasp but also provide more information: one particular advantage is that a confidence interval includes the information about the number of data points on which the comparison is based—this may be different for different amino acids and amino acid totals, depending on the number of successful curve fits and the amount of curation necessary. Although the functions in **aaresponse** are primarily focused on comparing different protein meals, many other statistical analyses can be performed.

All of this has been implemented in an open-source R package, **aaresponse**, which will continue to be developed as we apply it in more and more scientific projects. The package also provides several visualization methods—one thing to avoid is to see this as a push-button solution that automatically provides the required answers. In practice, it is a tool for the scientist that provides support for the many decisions to be taken on the route to the final outcome,




Fig. 5 Data and fitted curves for participants 9 and 11 in Figure 3, WP and BP, respectively

concerning, e.g. removal of outlying points, detection of strange patterns, and quality control checks throughout the analysis process. In no way do we mean to imply that this is the end-all solution for the analysis of crossover experiment data investigating protein uptake, but we do hope that the R package will prove useful for the scientific

community, and will form a focal point to which other researchers can contribute their views and suggestions.

Finally, it should be noted that this methodology is not just applicable to levels of amino acids in blood after a protein meal. Similar questions are ubiquitous in food science and can also relate to, e.g., glucose or fat levels. As long as the time courses that are measured after an intervention

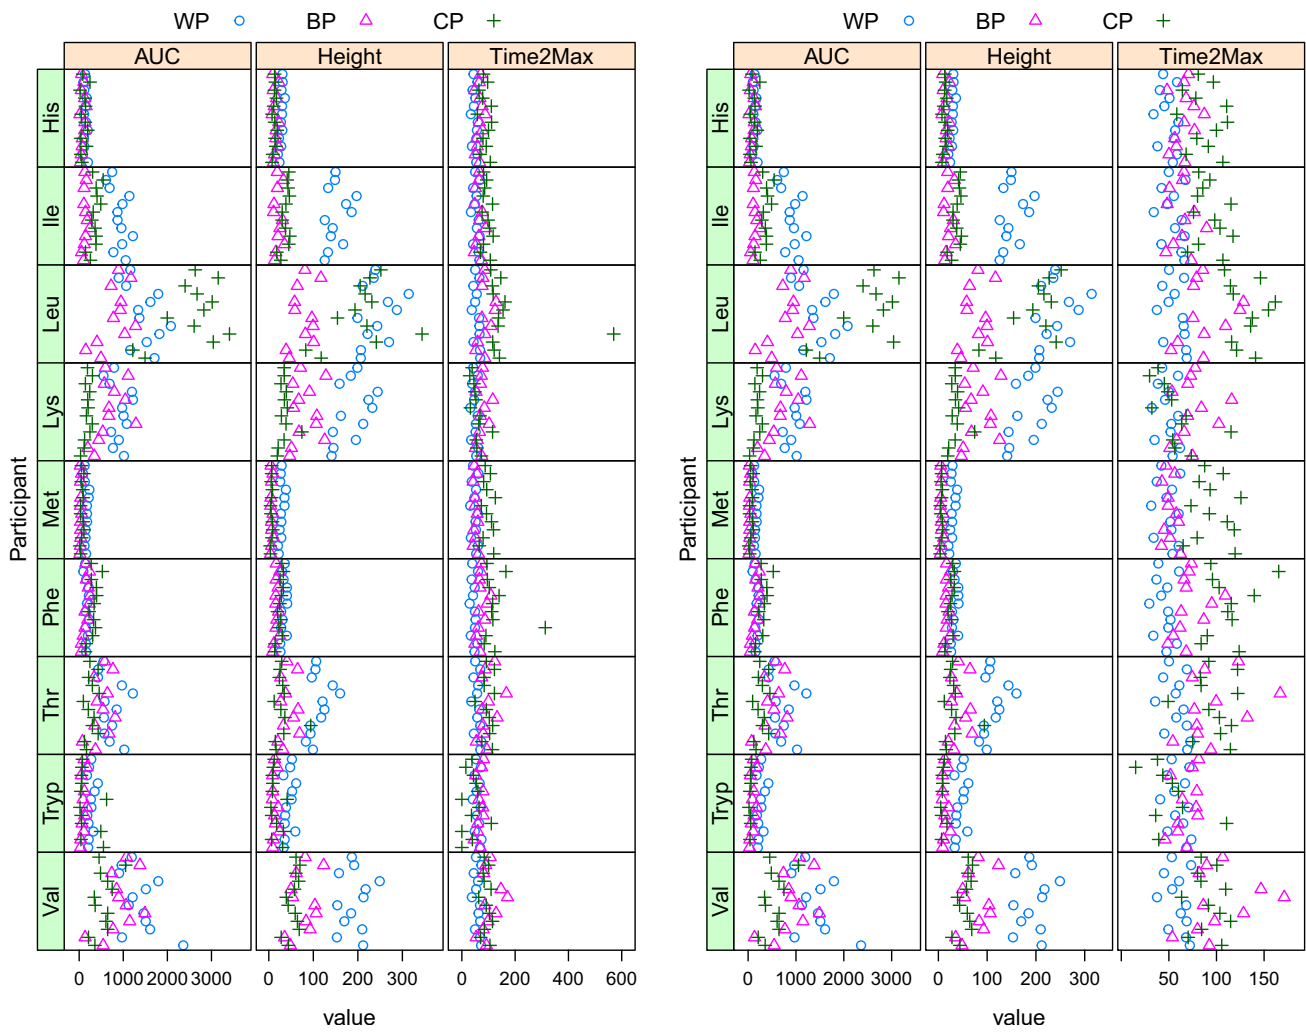


Fig. 6 Uncurated (left) and curated (right) estimates for PoIs (AUC, Height, Time2Max) for the essential amino acids also shown in Fig. 3. In each panel, results for all participants are shown




Fig. 7 Estimates and 95% confidence intervals for the comparison of BP and CP proteins with reference WP, all PoIs, and essential amino acids only. The vertical dashed lines correspond to the “no difference” situation (1 for the AUC ratio, 0 for the differences in Height or Time2Max). Significant differences (at the 95% level) are indicated in red, others in black

cover enough of the response curve to obtain meaningful fits, the methodology in principle can be applied. One requirement needs to be checked: if the changes in the levels of the compound of interest change significantly during the time course due to other causes than the intervention alone, the analysis will likely not lead to useful results. That, however, is true whatever method is used for the analysis of the data. We have tried to implement the package as generally as possible, and welcome any feedback from package users.

Data availability The aaresponse package contains all data necessary to reproduce the results in this paper.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bates D, Mächler M, Bolker B, Walker S (2015) Fitting linear mixed-effects models using lme4. *J Stat Softw* 67(1):1–48
- Butts CA, Monro JA, Moughan PJ (2012) In vitro determination of dietary protein and amino acid digestibility for humans. *Br J Nutr* 108(Suppl 2):282–287
- Dam LCHJ, Kardinaal A, Troupin J, Boulier A, Hiolle M, Wehrens R, Mensink M (2023) Postprandial amino acid response after the ingestion of pea protein, dairy proteins and their blend, in healthy older adults. *Int J Food Sci Nutr* 75:70–80
- Engel B (2003) Analysis of correlated series of repeated measurements: application to challenge data. *Biom J* 45:866–886
- Esser D, Mes J, Lenaerts K, Wehrens R, Engel J, Hermes G, Timmerman H, Dool R, Wichers H (2023) Comparing safety, nutritional quality and bio-functional activity of bovine-plasma, corn and whey protein concentrate consumption, outcomes of a randomised double blind controlled trial. *Curr Res Food Sci* 7:100588
- FAO Expert Consultation (2013) Dietary protein quality evaluation in human nutrition. Technical report food and nutrition paper no. 92, FAO, Rome, Italy (2013)
- Hothorn T, Bretz F, Westfall P (2008) Simultaneous inference in general parametric models. *Biom J* 50(3):346–363
- Jackson JE (1991) A user's guide to principal components. John Wiley & Sons Ltd, Chichester
- Kenward MG, Roger JH (1997) Small-sample inference for fixed effects from restricted maximum likelihood. *Biometrics* 53:983–997
- Lenth RV (2023) Emmeans: estimated marginal means, aka least-squares means. R package version 1.8.7. <https://CRAN.R-project.org/package=emmeans>
- Liu J, Klebach M, Visser M, Hofman Z (2019) Amino acid availability of a dairy and vegetable protein blend compared to single casein, whey, soy, and pea proteins: a double-blind, cross-over trial. *Nutrients* 11(11):2613
- Mes JJ, Esser D, Oosterink E, Dool RTM, Engel J, Jong GAH, Wehrens R, Meer IM (2022) A controlled human intervention trial to study protein quality by amino acid uptake kinetics with the novel lemna protein concentrate as case study. *Int J Food Sci Nutr* 73(2):251–262
- Padfield D, Matheson G (2020) Nls.multstart: robust non-linear regression using AIC scores. R package version 1.2.0. <https://CRAN.R-project.org/package=nls.multstart>
- R Core Team (2022) R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Roelofs JJM, Eijnatten EJM, Prathumars P, Jong J, Wehrens R, Esser D, Janssen AEM, Smeets PAM (2024) Gastric emptying and nutrient absorption of pea protein products differing in heat treatment and texture: a randomized in vivo crossover trial and in vitro digestion study. *Food Hydrocolloids* 149:109596
- Rook AJ, France J, Dhanoa MS (1993) On the mathematical description of lactation curves. *J Agric Sci* 121:97–102
- Sarkar D (2008) Lattice: multivariate data visualization with R. Springer, New York (ISBN 978-0-387-75968-5)
- Ummels M, Janssen Duijghuijsen L, Mes JJ, Van der Aa C, Wehrens R, Esser D (2023) Evaluating brewers spent grain protein isolate postprandial amino acid uptake kinetics: a randomized, cross-over, double-blind controlled study. *Nutrients* 15(14):3196
- Wegrzyn TF, Henare S, Ahlborn N, Ahmed Nasef N, Samuelsson LM, Loveday SM (2022) The plasma amino acid response to blended protein beverages: a randomised crossover trial. *Br J Nutr* 128:1555–1564
- Wehrens R (2020) Chemometrics with R—multivariate data analysis in the natural sciences and life sciences, 2nd edn. Springer, Berlin, Heidelberg (ISBN 978-3-662-62026-7)
- Wood PDP (1967) Algebraic model of the lactation curve in cattle. *Nature* 216:164–165

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.