An expert study on hierarchy comparison methods applied to biological taxonomies curation

Lilliana Sancho-Chavarria [Corresp., 1] , **Fabian Beck** [2] , **Erick Mata-Montero** [1]

[1] School of Computing, Costa Rica Institute of Technology, Cartago, Cartago, Costa Rica

[2] paluno, University of Duisburg-Essen, Duisburg, North Rhine-Westphalia, Germany

Corresponding Author: Lilliana Sancho-Chavarria
Email address: lsancho@tec.ac.cr

Comparison of hierarchies aims at identifying differences and similarities between two or more hierarchical structures. In the biological taxonomy domain, comparison is indispensable for the reconciliation of alternative versions of a taxonomic classification. Biological taxonomies are knowledge structures that may include large amounts of nodes (taxa), which are typically maintained manually. We present the results of a user study with taxonomy experts that evaluates four well-known methods for the comparison of two hierarchies, namely, edge drawing, matrix representation, animation, and agglomeration. Each of these methods is evaluated with respect to seven typical biological taxonomy curation tasks. To this end, we designed an interactive software environment through which expert taxonomists performed exercises representative of the considered tasks. We evaluated participants' effectiveness and level of satisfaction from both quantitative and qualitative perspectives. Overall quantitative results evidence that participants were less effective with agglomeration whereas they were more satisfied with edge drawing. Qualitative findings reveal a greater preference among participants for the edge drawing method. Also, from the qualitative analysis, we obtained insights that contribute to explain the differences between the methods and provide directions for future research.

1 **An expert study on hierarchy comparison methods applied to**
2 **biological taxonomies curation**
3

5 Lilliana Sancho-Chavarría[1], Fabian Beck[2], Erick Mata-Montero[1]
6
7 [1] School of Computing, Costa Rica Institute of Technology, Cartago, Costa Rica
8 [2] paluno, University of Duisburg-Essen, Essen, Germany
9
10 Corresponding Author:
11 Lilliana Sancho-Chavarría[1]
12 ITCR, Cartago, Costa Rica
13 Email address: lsancho@tec.ac.cr

Abstract

16 Comparison of hierarchies aims at identifying differences and similarities between two or more
17 hierarchical structures. In the biological taxonomy domain, comparison is indispensable for the
18 reconciliation of alternative versions of a taxonomic classification. Biological taxonomies are
19 knowledge structures that may include large amounts of nodes (taxa), which are typically
20 maintained manually. We present the results of a user study with taxonomy experts that
21 evaluates four well-known methods for the comparison of two hierarchies, namely, *edge*
22 *drawing*, *matrix representation*, *animation*, and *agglomeration*. Each of these methods is
23 evaluated with respect to seven typical biological taxonomy curation tasks. To this end, we
24 designed an interactive software environment through which expert taxonomists performed
25 exercises representative of the considered tasks. We evaluated participants' effectiveness and
26 level of satisfaction from both quantitative and qualitative perspectives. Overall quantitative
27 results evidence that participants were less effective with *agglomeration* whereas they were more
28 satisfied with *edge drawing*. Qualitative findings reveal a greater preference among participants
29 for the *edge drawing* method. In addition, from the qualitative analysis, we obtained insights that
30 contribute to explain the differences between the methods and provide directions for future
31 research.

Introduction

35 Visual comparison of hierarchies has been prevalent in information visualization research
36 because it is relevant for a wide range of domains such as tracking changes in software projects,
37 comparing budgets, and describing dynamics of organizational structures, among others. In this
38 work, we study the comparison of hierarchies in the domain of biological taxonomies.
39 Taxonomic information has been scattered in publications for centuries. In spite of integration
40 efforts of global initiatives in the last decades, there are numerous taxonomic databases around
41 the globe (Ball-Damerow et al., 2019). A visualization tool for the identification of differences
42 and similarities between two versions of a taxonomy would contribute to such integration efforts.

43
44 Taxonomies are hierarchies created by experts to classify living organisms. Through
45 classification, mutually resembling organisms are placed together in categories known as
46 *taxonomic ranks,* which, in turn, make up the levels of the hierarchy. The main taxonomic ranks
47 include *domain*, *kingdom*, *phylum*, *class*, *order*, *family*, *genus*, and *species*. Each node within the
48 hierarchy is referred to as a taxon, that is, a name given to a group of organisms; for
49 instanceinstance, *Vertebrates* and *Mammals* are two taxa, the former is placed at the *phylum*
50 taxonomic rank and latter at the *class* taxonomic rank. Species and sub-species are placed at the
51 lower level of the hierarchy and their scientific names are expressed with a *binomial system of*
52 *nomenclature* that uses a Latin grammatical form. The first part is the genus and the second part
53 is the *specific epithet*. For instance, the domestic cat's scientific name is *Feliz catus* (formerly
54 known, for many years, as *Felix domesticus*), where the epithet name is *catus* and the genus is
55 *Feliz*. Upper levels of the cat's taxonomy are: family *Felidae*, order *Carnivora*, class *Mammalia*,
56 phylum *Chordata*, and kingdom *Animalia*. For a taxon to be recorded in a taxonomy, it must
57 have been described in a publication, either as a new group or as a review of an existing group of
58 organisms. Taxonomy records may include various data; but for comparison, they are required to
59 include at least taxon name, author's name, and the year of publication. This allows users to
60 determine under which judgment the classification was devised. After almost three centuries
61 since modern taxonomy was first established by Carl Linnaeus (Linné & Gmelin, 1767), one
62 might think that most organisms on Earth have been identified and classified, and that
63 taxonomies are rather static. However, on one hand, it is estimated that only about 1.5 million
64 from approximately 11 million species of macro organisms haven been identified and described
65 (Larsen et al., 2017). On the other hand, the dynamics of taxonomic work has lead experts
66 worldwide to end up with different versions of the classifications. Taxa names represent concepts
67 whose definition depends on the authors' criteria, which eventually gives rise to conflicting
68 versions of a taxonomy. These multiple versions will require corrections and re-classifications in
69 order to come to an integrated version that can more accurately document biodiversity. That is
70 how taxonomists often face the problem of reconciling different versions of a taxonomy. For
71 such reconciliation efforts, biological taxonomists require to perform a series of curation tasks.
72
73 Sancho-Chavarría et al. (2016, 2018) characterized curation tasks that involve taxonomic
74 changes when comparing two versions T_1 and T_2 of a taxonomy. Such characterization involved
75 interviews to six experts from three different countries and followed a two-stage analysis. During
76 the first stage, the authors reviewed literature and interviewed experts in order to obtain a list of
77 preliminary tasks. In the second stage, the tasks were shared and discussed with the experts, in
78 order to obtain a final list of ten tasks. Table 1 provides a description of those derived ten tasks
79 organized into three categories, namely, *pattern identification*, *query*, and *edition*. Some of these
80 tasks are domain-independent, and are frequently mentioned in information visualization
81 research (e.g., filtering, focus, retrieving details). Other are domain-specific to biological
82 taxonomy curation work (e.g., tasks in the pattern identification category such as identify splits
83 or identify merges).
84
85 •   *Pattern identification* tasks include the identification of *congruent*, *merged*, *split*, *renamed*,
86   *moved*, and *added/excluded* taxa, as well as the *overview* of changes and the visualization of
87   a *summary* of the resulting comparison. Congruence refers to same taxonomic concepts
88   present in both versions of a taxonomy. A split occurs when taxonomists determine that a

89   group previously considered a unit actually consists of several groups of species that should
90   be described separately. Conversely, a merge happens when taxonomists decide that several
91   independent taxa should be combined into the same group. A change of name (i.e., rename)
92   is usually due to a typo that needs to be corrected. A taxon appears moved when it has been
93   re-classified and placed in a different location within the other version of the taxonomy. An
94   addition occurs when a new taxon is added to the taxonomy, either because it is a new
95   discovery or because it had not been previously recorded in the database. Exclusions refer to
96   taxa that are contained in version T_1 but that are missing in the alternative version T_2. It is
97   important to note that, from a taxonomic point of view, once a species is discovered, it is
98   kept in the taxonomy even if the species becomes extinct; however, in this work we consider
99   exclusions because it is important that taxonomists know when records are missing in the
100   database. The *overview changes* task refers to the possibility of globally overviewing all
101   differences between two versions of a taxonomy. The *summarize* task consists of obtaining
102   statistical information on changes.
103   •   Curation tasks in the *Query* category enable users to obtain detailed information on taxa.
104   The *retrieve details* task lets users obtain attributes of a taxon, for instance, the year of
105   publication and the authors' names. The *focus* task refers to the action that users perform
106   when focusing on a group of organisms. Through *filter*, users may find taxa that satisfy
107   some given conditions and through *find inconsistencies*, users may recognize differences due
108   to errors or missing information (e.g., typos or undefined names).
109   •   The *Edition* category comprises just one task that is rather ample, namely, the process
110   through which experts make changes to T_1 and/or T_2 after analyzing the results of a
111   comparison.
112
113   For this research, we considered tasks that are most relevant for the identification of changes
114   between two versions of a taxonomy, that is: 1. *Identify congruence*, 2. *Identify corrections*
115   (splits, merges, moves and renames), 3. *Identify additions/exclusions*, and 4. *Overview changes*.
116   Changes are more likely to occur at lower level taxonomic ranks, for instance, species level
117   (Vaidya, Lepage & Guralnick, 2018), therefore for this study the identification of changes will
118   be visualized only at species level.
119
120   Previous research has contributed with visual models and tools to support the comparison of
121   alternative versions of a taxonomy (Contian et al., 2016; Graham, Craig & Kennedy, 2008; Dang
122   et al., 2015). However, in practice, most taxonomists still rely on simple indented lists to carry
123   out the curation process with little computational assistance. Information visualization provides
124   visual representations that help people perform their tasks more efficiently (Munzner, 2015);
125   therefore, we believe that the identification of differences can be eased with the support of a
126   hierarchy comparison visualization system.
127
128   Graham and Kennedy (Graham & Kennedy, 2010) surveyed the comparison of hierarchies and
129   organized the visualization methods into five categories, namely, *edge drawing*, *coloring*,
130   *animation*, *matrix representation*, and *agglomeration*. The *edge drawing* method presents the
131   two hierarchies as separate structures where differences and similarities are represented by edges
132   from nodes in T_1 to the associated nodes in T_2. The *coloring* method represents similar nodes
133   with the same color. *Animation* shows changes as smooth transitions from one hierarchy to the
134   other. In a *matrix* representation, one hierarchy is placed along the vertical axis and the other one

135   along the horizontal axis; matrix cells indicate relationships between nodes of the compared
136   hierarchies. The *agglomeration* method visually merges both hierarchies into an integrated list.
137

138   The open question that we address in this work is how well these methods support the above
139   curation tasks between two versions T_1 and T_2 of a taxonomy. From the five methods, we leave
140   out coloring as an independent condition since color can be used across all methods. We
141   designed and conducted a user study where twelve expert taxonomists evaluate those four
142   methods. We wanted participants to interact with each of the four methods in a close-to-reality
143   scenario. We developed an interactive software environment that integrates the four methods and
144   allows users to easily navigate from one method to another while doing the assessment exercises
145   (see figures 1 through 4). We wanted to capture the essence of each method and avoided the
146   introduction of features that could potentially favor any particular method. We developed the
147   software taking into consideration the importance of reaching a balanced implementation; thus,
148   the software included same functionality and a common user interface design for all methods.
149   We also carefully selected the data. Datasets contain sufficient types of changes to carry out
150   exercises for all tasks and were also selected to avoid the introduction of any bias due to the
151   potential prior knowledge of the data by the experts. Participants performed the same exercises
152   with each method; however, the target taxa were not the same, also to avoid bias from a learning
153   effect. Immediately after performing the exercises related to a task, participants were asked to
154   answer a user satisfaction questionnaire in order to evaluate each method in relation to the
155   completion of that task. We registered the participants' answers along with their interactions and
156   thinking out-loud comments. We performed a quantitative analysis on the participants' responses
157   to the exercises (i.e., whether they answered correctly or not) and also on their user satisfaction
158   assessment. Additionally, we obtained qualitative findings based on the participants' feedback
159   throughout the session. The software, the data, the questionnaire, a written guide to the tool, a
160   link to a demo video, and the analysis materials are publicly posted in Github at
161   https://github.com/lsanchoc/MethodsTasksUserStudy.
162

163   Our contribution with this work is twofold. On one hand, we assessed the effectiveness and level
164   of participants' satisfaction with each visualization method. Participants were likely effective
165   with three methods: *edge drawing*, *matrix*, and *animation*, and also preferred *edge drawing* over
166   the other methods. On the other hand, we obtained a set of themes that contribute to explain the
167   differences between the methods and provide valuable insights for future work on the design of
168   software tools for the comparison of biological taxonomies.
169

170   The rest of this paper is organized as follows. Section 2 presents related work. Section 3 details
171   the study design. It includes a description of the interactive software environment, the
172   participants' profile, the characteristics of the datasets, the user study protocol, and the
173   questionnaire. Section 4 describes the results of the study. In Section 5, we discuss the results
174   and present limitations and implications of the study. Finally, in Section 6, we present
175   conclusions and future work.
176

177

178   **Related Work**

179 Comparison—understood as the examination of two or more items to determine similarities and
180 differences—is a means that facilitates the process of interpreting information. Gleicher et al.
181 (2011) provide a comprehensive survey of visual comparison approaches focusing on comparing
182 complex objects. They analyze a number of publications, systems and designs, looking for
183 common themes for comparison, and propose a *general categorization of visual designs for*
184 *comparison* that consists of three general categories, namely, juxtaposition, superposition, and
185 explicit encoding. The juxtaposed designs place the objects to be compared separately, in time or
186 space. The superposed designs place the objects to be compared one over the other in an overlay
187 fashion. The explicit encoded designs show the relationships between objects explicitly and is
188 generally used when the relationships between objects are the subject of comparison. Hybrid
189 designs are also possible and usually combine two categories. More recently, a set of
190 considerations to understand comparison tasks and their challenges has been discussed (Gleicher,
191 2018), as well as comparison in the context of exploratory analysis (von Landesberger, 2018).
192 Guerra-Gómez et al. (2012) provide a classification of trees for comparison. The trees used in
193 this work are similar to Type 0 trees, which display only a label; however, nodes also contain
194 information about the author and year of publication of the species. Each node also contains the
195 list of synonyms. All these data are necessary in order to identify the types of changes.
196
197 As mentioned above, for the comparison of biological taxonomies we are considering the
198 methods surveyed by Graham and Kennedy (2010) for the comparison of two hierarchies,
199 namely, *edge drawing*, *matrix*, *animation* and *agglomeration*. Each of these methods can be
200 mapped to the mentioned *general categorization of visual designs for comparison*. So, the *edge*
201 *drawing* method comprises characteristics from both juxtaposition (hierarchies are placed
202 separately side by side) and explicit encoding (edges encode the relations between nodes); the
203 *matrix* layout corresponds to an explicit encoding design since the matrix cells can explicitly
204 indicate the relations among taxa; and *agglomeration* corresponds to a superposed design.
205
206 Previous works on hierarchy comparison match these categories. For instance, TreeJuxtaposer
207 (Munzner et al., 2003) compares phylogenetic trees by using a juxtaposed layout. It presents a
208 novel focus+context technique for guaranteed visibility and comparison is approached by
209 coloring. Holten and van Wijk (Holten & van Wijk, 2008) present a visualization method where
210 hierarchies are structured as icicle plots placed in juxtaposition. Relations are explicitly
211 represented by edges arranged through hierarchical edge bundles to reduce cluttering. The
212 Taxonomic Tree Tool (Contian et al., 2016) uses a juxtaposed layout to compare biological
213 taxonomies. It combines glyphs to explain the relations between taxa. ProvenanceMatrix (Dang
214 et al., 2015) compares two taxonomies using a matrix representation. Relations are explicitly
215 displayed through two mechanisms: glyphs and edges. Beck and Diehl (2010) compare two
216 software architectures that use a matrix. Hierarchies here are represented as icicle plots.
217
218 Examples that use animation for comparison are scarcer. Ghoniem and Fekete (2001) use
219 animation to visualize the transition between two alternative representations of the same tree laid
220 out as treemaps. Considering agglomeration-based designs, Beck et al. (2014) present a nested
221 icicle plot approach for comparing two hierarchies and Guerra-Gomez et al. (2012) contrast two
222 trees for the visualization of both node value changes as well as topological differences for the

223 comparison of budgets. Lutz et al. (2014) compared directory structures and conducted a
224 qualitative user study to identify usage strategies. Unlike the above-mentioned work, which did
225 not focus on biological taxonomies, Graham and Kennedy (2007) propose an agglomerated
226 visualization based on directed acyclic graphs for the comparison of multiple biological
227 taxonomies. They also analyzed set-based hierarchies and agglomerated graph-based
228 visualizations for the comparison of botanical taxonomies (Graham et al., 2000).
229
230 Our work differs from the above-mentioned previous studies in that, for the first time, four
231 visualization methods described in (Graham & Kennedy, 2010), are assessed for the comparison
232 of pairs of biological hierarchies with respect to typical curation tasks.
233
234

235 ## Study Design

236 The research question addressed is: "*How well does each method support carrying out biological
237 taxonomy curation tasks?*" This is assessed both quantitatively and qualitatively. On one hand,
238 we aim at obtaining a quantitative understanding of the participants' effectiveness and level of
239 satisfaction. On the other hand, we also aim at obtaining qualitative insights on the capacity of
240 each method to carry out tasks for the comparison of biological taxonomies. We therefore
241 explore how users interact with the visualizations and what their judgment of each method is. We
242 opted for a within-subject design study that involves four experimental conditions, one for each
243 method (*edge drawing*, *matrix*, *animation*, and *agglomeration*). In this way, each participant
244 could test and contrast all methods. The study fits into the category "Evaluating Visual Data
245 Analysis and Reasoning (VDAR)" (Lam et al., 2012). In this approach, the goal is to assess how
246 a visualization tool supports the analytic process for a particular domain. Accordingly, in our
247 study we want to evaluate how each implemented method supports the identification of
248 similarities and differences for the curation of biological taxonomies. In order to assess each
249 method for each task individually, participants were asked to solve a task with each method, then
250 giving feedback, before moving to the next task. We favored this design over one where
251 participants would first complete all tasks with one method (before moving to the next method)
252 because we wanted to gain task-specific insights on the differences of the tested methods.
253

254 **The Software Environment**
255 We developed an interactive web-based software environment that integrates in the same
256 environment the four methods for visual comparison of biological taxonomies described in the
257 Introduction section. The software environment was designed to investigate how these four
258 methods support the taxonomy curation tasks described in Table 1. We tried to conceive a
259 balanced design within the same environment, in order not to favor any method. We also decided
260 to develop a functional system in which participants get a realistic impression. We provided
261 remote web-based access to the software because many of the participants were located in
262 different parts of the world. Figure 1 illustrates the user interface for the agglomeration method,
263 Figure 2 presents the animation method, Figure 3 presents the edge drawing method, and Figure
264 4 presents the matrix method implementation. The two taxonomies to be compared, T_1 and T_2,
265 are displayed as indented lists. Each method implementation is accessible by easily clicking on a
266 tab. Users can inspect the data through the provided zooming features and by vertically scrolling
267 for all methods. Additionally, horizontal scrolling is provided for *matrix*. The visualization

268 layout is of course method-dependent. For *edge drawing* and *animation*, taxonomies are placed
269 juxtaposed. T_1 is placed on the left side of the screen and T_2 on the right side. For *matrix*,
270 taxonomy T_1 is also placed on the left side but T_2 is at the top of the matrix. Finally, for
271 *agglomeration*, T_1 and T_2 are interleaved and centered horizontally.
272
273 The main menu is common to all methods, which adds uniformity to the user interface; however,
274 the implementations of the different visualizations were optimized to best reflect the intrinsic
275 characteristics of each method. The main menu is located at the top of the window and contains
276 eight toggle buttons that display the changes induced by each type of curation task that we are
277 considering, namely, *congruence*, *splits*, *merges*, *moves*, *renames*, *new*, *exclusions*, and an
278 additional *all* button. For example, when the *splits* switch is on, the visualization shows how
279 each taxon with a split in T_1 is divided into taxa contained in T_2. The system is flexible enough to
280 allow users to turn several buttons on at the same time, in case they want to have several types of
281 changes displayed simultaneously. For *animation* (see Fig. 2) additional controls to play and stop
282 animations were added.
283
284 The color-coding scheme of the toggle buttons (i.e., tasks) is also the same across all methods
285 and defines the types of changes to be visualized. We tried to use hues that were easily
286 distinguishable to the eye in order to avoid confusion: blue for congruent taxa, pink for splits,
287 orange for merges, light green for moves, light brownish purple for renames, red for exclusions,
288 and green for added taxa. Familiar codes were also used; that is, red for exclusions and green for
289 new taxa added. The representation of relations depends on the comparison method. For
290 agglomeration, relations have to be inferred since data is interleaved and no explicit additional
291 marks or lines can be included easily. Hence, for this method we decided to use an augmented
292 color code in order to have a cue that would make it easier for participants to recognize to which
293 taxonomy a node belongs and to highlight the types of changes between T_1 and T_2 (i.e., the
294 relations between nodes). For this, we use the same hues but with different intensity, so the light
295 nuanced nodes in the agglomerated structure indicate that they belong to the taxonomy of origin
296 T_1 while the darker nuanced nodes indicate that they belong to the taxonomy of destination T_2.
297 For example, a split of a taxon x in T_1 into taxa p, q, and r in T_2 would show x in light pink color
298 and p, q, and r in a darker hue of pink. A legend was added to explain this color-coding. In the
299 *agglomeration* method relations are permanent but not explicit.
300
301 In the *edge drawing* and *matrix* methods, relations are explicit and permanent. For instance, with
302 the *edge drawing* method, a split of a taxon x in T_1 into taxa p, q, and r in T_2 is shown as three
303 pink edges going horizontally from taxon x to p, q, and r in taxonomy T_2. In the *matrix* method,
304 the same split case is shown as marked colored cells $(x; p)$, $(x;q)$, and $(x; r)$ respectively.
305
306 For *animation*, we considered two design choices: "animation by movement" and "animation by
307 emergence". In the former, an animation consists of moving the target taxon from T_1 to its new
308 position in T_2. In the latter, the target taxon would fade out from T_1 and would gradually appear
309 in T_2. In either case, relations are explicit although temporary because they disappear when the
310 animation is finished. We chose the first option because the paths followed by each moving
311 taxon provides better traceability cues than the second one. Considering the split case described
312 above, the animation would show x moving towards taxonomy T_2. On its way, x splits and
313 disappears to let p, q and r appear and keep moving until each of them reaches its definitive

314   position in T_2. The *animation* method per se does not necessarily involve leaving an explicit trace
315   (as *edge drawing* does).
316
317   Two curation tasks do not involve relations between nodes in the alternative taxonomies,
318   namely, *identify new taxa added* and *identify excluded taxa*. Given that inclusions and exclusions
319   take place only in one of the taxonomies, the system visualizes these situations only in the
320   taxonomy in which they occurred. Thus, excluded taxa are visualized in red color in taxonomy
321   T_1 and included taxa in green color in taxonomy T_2. Without the use of color, asking users to
322   visually infer which taxa were excluded from T_1 and which ones were included into T_2 would
323   require too much mental effort, especially when taxonomies are large.
324
325   The software included interactions that we considered crucial, such as navigating, selecting, and
326   zooming. For zooming, we took advantage of the browser's functionality and it is offered via
327   three different alternatives: a) by using the zoom feature on the browser's menu, b) using the
328   browser controls, that is, by pressing the <ctrl> key and moving the mouse scroll wheel up or
329   down, or c) pressing the key *I* on the keyboard for *zoom in* and the key *O* for *zoom out*. Functions
330   such as undo/redo, edit, and statistics were left out of the prototype since they were not crucial
331   for the study. We developed the software incrementally through several iterations until we
332   reached balanced implementations of the four methods. At the end of each iteration, computer
333   science students tested the software. Tests were also conducted involving an experienced
334   taxonomist and a PhD student in computer science.
335

Participants

337   Twelve experts were recruited from our professional network. They received no compensation.
338   Table 2 summarizes the participants' profiles. Each participant was given an identification
339   number, ranging from E1 to E12. Eight of them are botanists (three of which are also Forestry
340   engineers), two biologists (one entomologist and one ichthyologist), one ecologist, and one
341   computing engineer (with 21 years of experience in biodiversity informatics). In addition, three
342   of them reported Biodiversity Informatics as a second area of expertise. One participant holds an
343   Engineering degree, five have a Master's degree, and six have a PhD degree. Their average
344   professional experience was 28 years and their average experience in the taxonomy field was 23
345   years; this includes taxonomic classification, taxonomy nomenclature, and curation of biological
346   taxonomies. Ten participants are male and two female. Three participants worked as full time
347   university professors and the rest worked full time at herbaria, museums, or biodiversity
348   conservation initiatives. Participants came from three different countries and their expertise was
349   with different taxonomic groups of organisms.
350

Datasets

352   We carefully selected and designed the datasets, taking into account the level of familiarity
353   participants might have with the data. Although taxonomists, in general, have extensive
354   knowledge on certain groups of species, in practice, a taxonomist is only expert on a limited
355   group of organisms; therefore, not all taxonomists need an overview of large subtrees. In
356   addition, because of the large number and complexity of groups of species, their expertise is also
357   geographically focused. Thus, despite of having ten botanists in our group of experts, all of them
358   specialize in different groups of plants. In order not to favor any participant and avoid the
359   eventual bias, we did not choose groups of species that were known by any of the experts.

360   Therefore, we chose an unfamiliar taxonomy. It should not be very large since we did not want
361   to burden participants by spending too much time performing the user study. However, at the
362   same time, the dataset should be large enough to contain representative cases of all types of
363   changes. We therefore used a small-size real taxonomy and derived artificial variants from it.
364

365   We downloaded a set of 66 species of amphibians from the Catalogue of Life website
366   (http://www.catalogueoflife.org/) with a total of 96 nodes. We called this the *seed taxonomy*
367   from which we derived variations on the datasets (derived taxonomies) to be used for each
368   method. It was important that the data sets were different, but at the same time similar enough to
369   be comparable. Therefore, we programmed an *artificial taxonomy generator* to which we input
370   as parameters the percentage of splits, merges, movements, renames, additions, and exclusions
371   that we wanted to add to the seed taxonomy T_1. The generator randomly selected the taxa to
372   introduce the changes and verified that only one change was introduced to each taxa to be
373   modified. In this way, we prevented data conflicts, since more than one change to a taxon could
374   generate inconsistent data. We also verified that, although questions were identical, the datasets
375   would produce different answers for each method. The number of nodes in the derived datasets
376   varied between 78 to 116 nodes. Table 3 describes the main characteristics of the four derived
377   datasets, that is, the number of nodes, of species, of splits cases, merges, moves, renames, new,
378   and excluded taxa. The goal of this setup was to ask experts to visualize changes in four pairs of
379   datasets: (T_1, T_2), (T_1, T_3), (T_1, T_4), and (T_1, T_5), with respect to *edge drawing*, *matrix*, *animation*,
380   and *agglomeration*, respectively. The derived datasets T_2, T_3, T_4 and T_5 are similar because they
381   are all obtained from the seed taxonomy and have roughly the same number of changes. We
382   avoided the use of the same pair of datasets across all visualization methods in order to neutralize
383   a potential bias introduced by a learning effect.
384

385   **User Study Protocol**
386   We planned the user study for a 2-hour session with each participant. During the session,
387   participants would work with the interactive software environment to perform some exercises
388   and to answer questions from a questionnaire. Seven out of the twelve experts lived overseas;
389   therefore, the session was conducted remotely via a video call for them. For the rest of
390   participants, sessions were face-to-face. For participants in remote sessions, at the beginning of
391   the session, we shared a link were the software and data were hosted. In case of the face-to-face
392   interviews, we supplied a laptop computer. In both settings, access to the software environment
393   was via web browser. We followed the same interview protocol for all participants.
394

395   A written guide and a 15-minute descriptive video of the software environment were available to
396   the participants at least two days before the session, so that they could get familiar with it.
397   Access to the software environment, datasets, and questionnaire was not provided before the
398   interview session.
399

400   A moderator was leading the session and assisted participants, while an observer was taking
401   notes. Participants did not have to write down the answers; both the moderator and the observer
402   would write the participants' answers on an answer sheet that they had previously prepared.
403   Audios of the interviews were recorded for later confirmation of answers and analysis. At the
404   beginning of the session, we checked to see if the participants had studied the guide and video
405   beforehand and if they had any questions. In case they had not done so or if they needed to

406 clarify any aspect, the moderator offered a demonstration of the software and resolved the
407 doubts. Exercises were not started until both the participant and moderator felt they were ready;
408 only then did the moderator provide the link for participants to access the interactive
409 environment. Working speed was not to be measured and participants were made aware of the
410 fact that they had no time limit to answer the questions and were able to express any inquiry,
411 doubt or suggestion at any time. Participants were also asked and reminded to think aloud while
412 solving the questions. Our goal was to get insights on how they carry out the data exploration
413 and the tasks.
414
415 We designed an instrument that consists of twelve task performing exercises, nine method
416 assessment questions and one open-ended comments section. The task performing exercises have
417 clearly correct answers and were intended to measure the participants' effectiveness. The method
418 assessment questions were intended to obtain participants' perception. The purpose of the open-
419 ended question was to obtain additional feedback on user satisfaction and suggestions for a
420 future design of an interactive visualization system. The study started with an exercise where
421 participants had to identify the most common type of change (overview task). Next, exercises
422 were targeted to identify splits, merges, renames, moves, added or excluded taxa, and ended with
423 an overview question again. Each task-performing exercise had to be answered with each
424 method. For instance, instructions such as "Use the Matrix method: Explore the visualization and
425 find into what taxa *Babina caldwelli* was split?" were followed by the same question for all
426 methods. However, the taxon to be used in each exercise (*Babina caldwelli*) was different for
427 each method.
428
429 We randomized the order in which participants used each method on each question. Participants
430 performed the exercises related to one task (for instance, identification of splits) and then were
431 asked to assess each method to perform such task. The nine method-assessment questions
432 consisted of five-level Likert scale items that assessed how good each method was to carry out
433 the task. In the course of the session, participants had access to a copy of the questions and
434 instructions, especially because taxa names were in Latin, and we wanted to avoid any
435 confusion.
436

Analysis

438 For the analysis of the results, we organized the participants' responses into a spreadsheet. We
439 gathered three types of data: a) the effectiveness data, i.e., whether the participants answered
440 each question correctly or incorrectly, b) the user satisfaction data, i.e., the Likert-scale ratings
441 that participants gave to each method after accomplishing each task, and c) the qualitative data,
442 i.e., the thinking-aloud comments and the suggestions that participants provided during the
443 session. Quantitative analysis was performed on data of types a) and b) by using a statistical
444 package. We used non-parametric statistics with alpha=0.05 and compared medians to determine
445 that differences are not due to chance. For the analysis of effectiveness, we used the Cochran's Q
446 test, which can be used when you have a group of people performing a series of tasks where the
447 outcome is dichotomic (e.g. success or failure). For the analysis of participants' satisfaction, we
448 used the Friedman test, which is appropriate for within-subjects designs that have three or more
449 conditions, and particularly it can be used for the analysis of ordinal data, such as the Likert-
450 scale responses (MacKenzie, 2013). When necessary, both tests were followed by pairwise
451 comparison using Dunn's test with Bonferroni correction.

452

453 For the qualitative analysis, we applied the following procedure. The responses were first placed
454 in the same order in which the questions were presented to the participants, and, then, they were
455 sorted by method. Two columns were designated for each participant, one to record their
456 comments and suggestions (e.g., E1) and another one to afterwards register the codes generated
457 during the qualitative analysis (e.g., E1-codes). Secondly, we listened to the audio recordings
458 checking for additional feedback from the participants, which we added to the spreadsheet.
459 Thirdly, we conducted a qualitative analysis: the first author made several coding passes using
460 *open coding* (Charmaz, 2006) to obtain a first coding version that was then shared with the other
461 authors. We coded participants' interactions and feedback. Repeated or related topics were
462 grouped together, revised and re-grouped through several refinement cycles until we reached an
463 agreement with twelve categories to finally conclude with four meaningful themes. During the
464 process, we also organized the positive and negative comments, as well as the participants'
465 suggestions for improving the methods.

466

Results

468

469 The study took 2:15 hours on average per participant. We first present quantitative results on
470 participants' effectiveness and satisfaction, and then findings from the qualitative analysis.

471

Effectiveness

473 The results of the participants' effectiveness on the task-performing exercises are summarized in
474 Table 4. Overall results indicate that participants obtained more correct answers with *matrix*
475 (94%), then with *edge drawing* (88%), followed by *animation* (87%) and then with
476 *agglomeration* (73%). We tested for statistical significance by using Cochran's Q test for N=12
477 and DF=3. We did not find significant differences on participants' responses between pairs of
478 methods (*matrix*, *edge*), (*matrix*, *animation*), and (*edge*, *animation*). However, we did find
479 differences (χ^2=40.480, p-value = 0.05) between agglomeration (73%) and the other methods,
480 meaning that participants were less effective with the agglomeration method.

481

482 We also did a quantitative analysis on responses to each exercise. We did not find significant
483 differences among participants' responses when identifying: a) into which taxa a taxon was split
484 (exercise 2), b) whether species were merged and how (exercises 4 and 5), c) whether species
485 were renamed (exercises 6 and 7), d) whether any species were added to a version of the
486 taxonomy (exercise 9), and e) whether any species were excluded (exercise 10). We found
487 significant differences in participants' responses in identifying: a) an overview of changes
488 (exercises 1 and 12), b) which species were most divided (exercise 3), c) moved taxa (exercise
489 8), and d) all changes on a taxon (exercise 11). For these cases, a post hoc pairwise comparison
490 was performed in order to determine where the differences occurred. Resulting less-effective
491 methods are highlighted in orange color in Table 4:
492 • Exercise 1. Overview of changes. We found differences (χ^2=25.500, p-value < 0.05) between
493   the following pairs of methods: (*animation*, *edge drawing*), (*matrix*, *edge drawing*) and
494   (*agglomeration*, *edge drawing*). This indicates that the effectiveness with *edge drawing* (8%)
495   was lower with respect to *agglomeration* (75%), *animation* (92%) and *matrix* (92%).
496 • Exercise 3. Identification of splits. We found differences (χ^2=17.571, p-value <0.05) between
497   pairs of methods (*edge drawing*, *agglomeration*) and (*matrix*, *agglomeration*). These results

498   indicate that the effectiveness with *agglomeration* was different with respect to the other
499   methods. Participants were less effective with *agglomeration* (33%) and more effective with
500   *edge drawing* (100%) and *matrix* (92%).
501 • Exercise 8. Identification of moved taxa. We found differences (χ^2=21.000, p-value < 0.05)
502   between pairs of methods (*animation*, *agglomeration*), (*edge drawing*, *agglomeration*), and
503   (*matrix*, *agglomeration*). This indicates that participants were less effective with
504   *agglomeration* (42%) whereas they were more effective with the other methods (100%).
505 • Exercise 11. Focus on a taxon. We found differences (χ^2=9.692, p-value < 0.05) between the
506   pair of methods (*agglomeration*, *matrix*). This indicates that participants were less effective
507   with *agglomeration* (50%) and more effective with *matrix* (100%).
508 • Exercise 12. Overview of changes. We found differences (χ^2=10.714, p-value < 0.05)
509   between pairs of methods (*animation*, *agglomeration*), (*edge drawing*, *agglomeration*), and
510   (*matrix*, *agglomeration*). This indicates that participants were less effective with
511   *agglomeration* (17%) than with *animation* (58%), *edge drawing* (58%), and *matrix* (58%).
512
513 **Satisfaction Level**
514 Right after carrying out the task-performing exercises for each task, participants answered a
515 Likert-scale questionnaire to assess the methods. The questions had the following structure:
516 "How good do you think each method is in order to perform task t? For each method provide a
517 rating between 1 and 5, where 1 stands for 'poor', 2 for 'fair', 3 for 'good', 4 for 'very good',
518 and 5 for 'excellent'". We performed a statistical analysis on the participants' ratings using the
519 Friedman test. Table 5 summarizes the results for N=12 and DF=3. We did not find any
520 difference in participants' responses to accomplish the task for the identification of excluded
521 species (question 10e). Neither we found differences regarding the identification of added
522 species (question 9e) after running the post pairwise comparison. On the contrary, we found
523 differences in participants' responses to carry out tasks for the identification of the most common
524 type of change (1e), splits (3e), merges (5e), renaming (7e), moves (8e), changes to a taxon (12e)
525 and the general methods assessment question (13). The post hoc pairwise comparison gave the
526 following results (in Table 5, higher satisfaction level is shown in green and less satisfaction
527 level is highlighted in orange):
528
529 • Question 1e-overview (χ^2=12.588, p-value =0.05). We found differences between pairs of
530   methods (*animation*, *edge drawing*). Participants gave a better rating to the *edge drawing*
531   method (median = 3.29 ~ 'good/very good') than *animation* (median = 1.79 ~ 'fair').
532 • Question 3e-splits (χ^2=33.055, p-value = 0.05). We found differences between pairs of
533   methods (*agglomeration*, *matrix*) (*agglomeration*, *edge drawing*) (*animation*, *matrix*), and
534   (*animation*, *edge drawing*). There was no difference between *agglomeration* and *animation*,
535   and neither *matrix* and *edge drawing*. Participants ratings for *agglomeration* (median = 1.50
536   ~ 'poor/fair') and *animation* (median = 1.50 ~ 'poor/fair') were the lowest while for *matrix*
537   (median = 3.21 ~ 'good') and e*dge drawing* (median = 3.79 ~ 'good/very good') were the
538   highest ones.
539 • Question 5e-merges (χ^2=20.050, p-value = 0.05). We found differences between pairs of
540   methods (*animation*, *edge drawing*) and (*agglomeration*, *edge drawing*). Participants
541   assessed the *edge drawing* method with the highest rating (median = 3.71 ~ 'good/very
542   good') compared to *agglomeration* (median = 2.0 ~ 'fair') and *animation* (median = 1.71 ~
543   'poor/fair').

- Question 7e-renames (χ^2=21.559, p-value < 0.05). We found differences between pairs of methods (*animation*, *edge drawing*) and (*agglomeration*, *edge drawing*). Participants assessed the *edge drawing* method with the highest rating (median = 3.67 ~ 'good/very good') than *agglomeration* (median = 2.13 ~ 'fair') and *animation* (median = 1.50 ~ 'poor/fair').
- Question 8e-moves (χ^2=24.295, p-value = 0.05). We found differences between pairs of methods (*agglomeration*, *edge drawing*) and (*animation*, *edge drawing*). Participants assessed the *edge drawing* method with the highest rating (median = 3.83 ~ 'very good') than *agglomeration* (median = 1.54 ~ 'poor/fair') and *animation* (median = 1.88 ~ 'poor/fair').
- Question 12e-focus (χ^2=16.057, p-value = 0.05). We found differences between pairs of methods (*animation*, *edge drawing*) and (*agglomeration*, *edge drawing*). Participants assessed the *edge drawing* method with the highest rating (median = 3.46 ~ 'good/very good') than *agglomeration* (median = 1.88 ~ 'fair') and *animation* (median = 1.79 ~ 'fair').
- Question 13-general assessment of all methods. (χ^2=16.057, p-value = 0.05). We found differences between pairs of methods (*animation*, *edge drawing*) and (*agglomeration*, *edge drawing*). Participants assessed the *edge drawing* method with the highest rating (median = 4.97 ~ 'excellent') than *agglomeration* (median = 2.97 ~ 'good') and *animation* (median = 2.84 ~ 'good').

It is important to notice that response for question 13 summarizes the participants' level of satisfaction.

Findings from Qualitative Analysis

Regarding *edge drawing*, all participants referred to this method throughout all exercises with expressions such as: "easy", "very direct", "I can easily relate taxa", "it is very fast", "I do not have to think too much", and "you can see ... at a glance". One participant (E6) said that it was the best because "you can clearly see the origin and the destination". Another participant (E7) considered that *edge drawing* "is familiar, it is similar to an 'associate' type of exercise".

Participants' feedback on *matrix* highlighted this method as good for the visualization of general overviews and the identification of patterns, eight participants mentioned it (E2, E3, E7, E8, E9, E10, E11, and E12). Five participants (E6, E7, E10, E11 and E12) mentioned that *matrix* was the fastest one. Another participant (E2) recognized that "it is easy to see in a row the changes to a taxon". Three participants (E10, E11 E12) considered that the required vertical and horizontal scrolling add complexity and two other participants (E2 and E10) mentioned that scaling could be a problem. One participant (E1) complained that he had to use his fingers on the screen to follow the relations in the two dimensions. Eleven participants complained about the vertical name implementation in the top hierarchy (all except E2), and two participants (E1 and E10) found navigation difficult because parts of the hierarchies were off the screen.

Regarding *animation*, two participants (E1 and E8) rated it positively indicating that it was "dynamic", and therefore "fun". However, eight out of the twelve participants described this method in negative terms such as "difficult", "ineffective", "hard to follow", "complicated", "not intuitive", and "waiting until the end of the animation is a waste of time". Five participants considered that the *animation* was not necessary. Participants emphasized that changes between the two taxonomies were very difficult to follow because they could very soon forget what

590 happened, especially if taxonomies were large. Five participants indicated that, while taxa were
591 moving, it was easy to lose track of the relation between origin and destination because the taxa
592 were moving. Most participants speeded up the animation, giving the impression that they
593 wanted it to get to an end quickly, but some had to execute it several times before being able to
594 solve the exercise.
595
596 Participants indicated that it was very difficult to carry out the tasks with the *agglomeration*
597 method, except for the identification of excluded taxa. Eight participants (E2, E5, E7, E8, E9;
598 E10, E11, E12) referred to *agglomeration* as very good when looking for specific taxa or to
599 focusing on a small part of the taxonomy. However, all participants also described it in negative
600 terms, such as "difficult", "very complicated", "requires too much effort", "not evident",
601 "confusing", and "very difficult to know origin and destination". Participants complained that
602 this method involved many variables that were difficult to remember (that is, many color hues)
603 and that it required considerable effort to recognize differences. However, two participants
604 thought that the *agglomeration* view could be complementary to *edge drawing*, and that it could
605 work well for small taxonomic groups.
606
607 We coded and organized the participants' feedback until we reached themes that we considered
608 meaningful. Our observations show four specific issues that are relevant when performing tasks
609 for the curation of biological taxonomies:
610
611 •   **Explicit representation of changes**. Changes are visualized through relations among taxa. Nine
612   out of twelve participants clearly indicated that being able to identify the origin and
613   destination of relations was very important to recognize changes when comparing biological
614   taxonomies. Participants' suggestions such as; "add edges to animation", "add edges to
615   matrix", or "add numbers to each change in the agglomeration method" are indications that
616   they would prefer to see relations explicitly and, therefore, prefer methods that explicitly
617   represent the changes.
618 •   **Efficiency**. Participants often commented about speed and time needed to solve the exercises.
619   Across exercises, they referred to the importance of understanding what is going on at a
620   glance. They expressed feeling frustrated when having to wait for the animation to end. They
621   speeded up the animation when they felt that solving the exercise was taking too much time.
622   Participants considered that having to scroll horizontally and relate rows and columns of the
623   matrix or having to interpret different colors as in agglomeration were steps that consumed
624   time.
625 •   **Multiple views**. Several participants commented that the methods could be complementary;
626   for instance, that the *edge drawing* and the *matrix* methods could be used to visualize all
627   cases at once whereas the *animation* and the *agglomeration* methods could be useful when
628   analyzing smaller groups of species. They explained that, by combining several methods, the
629   advantages of one method could overcome the disadvantages of another one. On the other
630   hand, the experts also emphasized the convenience of having both overview and detailed
631   views; the first one to obtain a general understanding of changes and the second one to obtain
632   detailed information on a focused part of the taxonomies.
633 •   **Visual and numerical summaries**. When asked for number of taxa that match a certain
634   condition, participants expressed their frustration because they had to count manually and
635   suggested to add statistics to the software environment. Although obtaining statistical

636 information is one of the tasks for the curation of biological taxonomies, we decided to leave
637 it out of this study on purpose in order to force participants focus on the visualizations. The
638 intention of quantity-related questions was to see if participants were able to visually identify
639 magnitude of changes (for instance, matrix resulted good). We obtained confirmation on the
640 importance of providing numerical understanding of changes.
641
642 **Suggested Improvements**
643 The methods that received most suggestions for improvement were the ones that had the lowest
644 participants' effectiveness and preference. Suggestions for agglomeration focused on
645 mechanisms that would make the relations explicit somehow and allow them to be recognized
646 quickly. For instance: a) add numbers instead of different color hues to indicate taxa of the origin
647 and taxa of the destination, b) use a different color hue for each change (not only for each type of
648 change), c) instead of using several color hues, consider glyphs or some other visual cues, d)
649 color the background of the text instead of the text, e) visually separate the taxonomies on user's
650 demand, f) separate the legend so that hues associated to origin are placed on the left side and the
651 hues associated to the destination are placed on the right side of the screen. For animation
652 participants provided suggestions such as: a) add a time slider and a rewind button, b) identify
653 each change with a number, c) add traces as in edge drawing, d) identify each specific change
654 with a different color, e) maybe consider the use of animation for comparing a small part of the
655 taxonomies. For matrix, participants also made suggestions to improve the visualization of
656 relations; for instance, a) add a feature to freeze rows or columns in place and ease the
657 visualization of relations when vertical or horizontal scrolling is needed, b) add horizontal and
658 vertical guiding lines to ease following the relations, c) add a colored rectangle around the
659 excluded or added taxa in order to highlight these changes, d) add edges in order to make the
660 relations more explicit, e) consider the matrix method as a way to feed a database with the
661 relations between the two taxonomies. Suggestions for edge drawing included: a) use more
662 intense color hues and b) provide features that ease the comparison between taxa of higher
663 taxonomic ranks (for instance, at the family level). Three participants (E3, E5, E6) mentioned on
664 several occasions that the identification of relations was easier when the involved taxa were
665 closer together, within the same view. They expressed this thought as they were solving matrix
666 and agglomeration exercises. On the other hand, four participants (E3, E6, E10, E12) mentioned
667 that vertical scrolling was fine in edge drawing, since "it is very familiar". Feedback obtained
668 from the open-ended comments section of the questionnaire also included suggestions for
669 enhancements to the implementation of the methods. Regarding the representation of hierarchies,
670 participants' suggestions included the elimination of lines that indicate hierarchical structure and
671 to use only indentation as a visual cue to recognize hierarchy (some participants believed that the
672 visualization could look cleaner). They also mentioned that the visual clutter caused by long
673 names might be overcome by using abbreviations when possible (for instance, for the genus part
674 of the species names).
675

676 # Discussion

677

678 The research question we investigated was how well each of the four methods of hierarchy
679 comparison supports the tasks of contrasting two versions of a taxonomy. The quantitative and
680 qualitative results revealed differences among the methods. The difference in effectiveness
681 occurred only with respect to agglomeration, as the participants were the least effective with this

682   method. One likely explanation is that in all methods, the changes between the two taxonomies
683   were consistently represented throughout predetermined colors for each type of change, but in
684   agglomeration each change was represented by two tonalities of the predetermined color, one to
685   indicate the taxon of origin and the other to indicate the taxon of destination. This might have
686   added complexity that affected the participants' effectiveness with this method. Overall,
687   participants were very effective with the other three methods, which might be because the
688   participants could take as much time as they considered necessary to solve the exercises. The
689   number of correct answers with animation and with edge drawing was quite similar (125 and 126
690   respectively). In spite of this similarity in the participants' accuracy, and that both methods used
691   juxtaposed layouts, the user satisfaction results indicate greater participants' preference for edge
692   drawing. Comparing the amount of correct answers, participants showed similar performance in
693   many exercises with edge drawing and matrix. The difference between these two methods comes
694   mainly from the responses to the overview exercise where they had to identify "1. Which is the
695   most common type of change?" where only one participant answered correctly with edge
696   drawing and 11 participants with matrix. This might indicate that matrix works well to get a
697   general overview of changes. This is reflected also in the participants' feedback when they
698   highlighted that this method was good for pattern recognition.
699
700   The effectiveness on recognizing new taxa or excluded taxa was similar with all methods. This is
701   explained by the fact that both new and excluded species are visualized only in one of the
702   taxonomies, and require no relations between the involved taxonomies.
703
704   In another respect, we noticed that participants were more effective at identifying changes at the
705   lowest level of the taxonomy (i.e., species level) than when trying to recognize changes at upper
706   levels (such as at the genus level, exercise 12). This might suggest the convenience of having
707   summary overviews on changes at higher-level taxa.
708
709   Both the quantitative and the qualitative analysis coincide that the agglomeration method ranked
710   last. In spite that the results on effectiveness did not show clear differences among animation,
711   edge drawing and matrix, both the participants' feedback on satisfaction and the qualitative
712   findings suggest edge drawing in the first place.
713
714   **Threats to Validity**
715   Various factors can limit the validity of the results of visualization user studies, including the
716   number of participants, the choice of the datasets, the design of the study, and the data analysis.
717
718   The study included twelve expert users. A small number of participants can affect the quality of
719   results, especially with respect to quantitative results. However, with respect to qualitative
720   findings, the number was sufficient: After completing a first set of nine interviews, we recorded
721   the data on a spreadsheet and did a preliminary processing. Afterwards, when we finished all
722   twelve interviews, we noticed that the qualitative results repeated (that is, with 33% more
723   participants than in the preliminary processing). Such consistency is an indication of
724   dependability regarding the qualitative results.
725
726   In spite of the tasks randomization, given the small number of participants, bias might have been
727   introduced by the learning effects.

728

729   A restriction of the study is the small size of the datasets, which contained between 78 and 116
730   nodes. As the data were unknown to all participants, the datasets had to be restricted to a small
731   size in order to carry out the study in reasonable time. Still, the datasets contained all types of
732   required changes to perform the tasks. In the case of the dataset for matrix, it turned out a little
733   smaller than the other ones, but large enough to allow users to experience both horizontal and
734   vertical scrolling, so we considered it was fine for the study; however, the difference in size
735   compared to the other datasets might have added some bias to the study. For visualizing larger
736   excerpts or whole taxonomies, other visualizations might be necessary, and our results may not
737   generalize to this. Also, the binding of datasets to methods could have introduced a slight
738   confounding factor but seemed an acceptable risk, which we took to simplify the evaluation
739   logistics.

740

741   Although the transferability of results is limited by a) the specific domain application, b) the
742   tasks studied, and c) the data sets, the comparison of hierarchical structures is independent of the
743   application domain; thus, some features of the study might contribute to other contexts where
744   users need to identify divided, joined, moved or renamed nodes between two hierarchies.
745   The within-subject design is applied in studies with a small number of participants. All
746   participants interacted with every method. In this way, we expected to reduce biases associated
747   with individual differences. However, the within-subject design may bias participants because of
748   the carry over effect; that is, once participants perform a task with one method, they may expect
749   certain conditions to happen in the next method to evaluate. We tried to counterbalance bias from
750   the learning and tiring effects by asking participants to interact with the methods in a randomly
751   established order. In addition, although the questions were identical, the data sets would produce
752   different answers for each method in order to prevent the learning effect.

753

754   A bias may be introduced by the design of the study or the design of the interactive software
755   environment. We tried to implement the essence of the methods as well as to keep a standard
756   user interface for all visualizations (same main menu and color codes throughout all methods),
757   but limitations may come from the design choices and implementation decisions. For instance,
758   zooming and scrolling features were limited and not designed for very large amounts of data;
759   also, participants were unsatisfied regarding the vertical text orientation of the matrix
760   implementation. Applying coloring was more difficult in the agglomeration approach because, in
761   this method, color is also required to indicate the different versions of the taxonomy. We decided
762   for a mixed brightness and color encoding as a tradeoff, which seemed introducing a smaller
763   bias, instead of avoiding the use of color for encoding types of changes or introducing an
764   inconsistent use of color across the methods.

765

766   We aimed to objectively examine the collected data. For the quantitative part of the study, we
767   used statistical tests to analyze if the differences between the medians were significant. For the
768   qualitative part of the study we carefully organized the data and coded the participants' feedback
769   and interactions through several refinement cycles. By counting and grouping similar feedback
770   from participants, we were able to define the codes and themes. Nonetheless, the interpretation
771   of the data may be subject to the perspective of the researchers.

772

773   **Implications**

774 The participants performed well with edge drawing and consider it, in general, the best method;
775 despite they did not have the best performance with it. For overview tasks, participants showed
776 similar effectiveness with matrix and animation, however they preferred matrix.
777
778 Results indicate that identifying explicitly the origin and destination of taxa is very relevant for a
779 more efficient identification of changes; edge drawing and matrix methods seem to have
780 facilitated it. The participants' need to determine origin and destination may explain that edge
781 drawing outperformed animation in both effectiveness and user satisfaction. Both methods
782 present the taxonomies in a juxtaposed layout, however, relations are not explicit in animation.
783 Some participants considered that animation could be useful to focus on changes in small
784 taxonomic groups, which reaffirms the scope of animations, as indicated by Graham and
785 Kennedy [11].
786
787 During the sessions, we noticed that sometimes taxonomists wanted to see the big picture and
788 then focus on a smaller group of organisms of their interest. Also, sometimes they wanted to go
789 directly to the group they want to inspect. Thus, future research should consider easily toggling
790 between overview and detailed views as well as search and filtering functions. Text (that is, taxa
791 names) is crucial when comparing taxonomies. Visual cues such as color, size, shapes or glyphs
792 are not enough to recognize the differences and similarities. Text must be legible. The users
793 would have to read names, which would need to be accommodated efficiently avoiding
794 cluttering. Unlike other studies, we were not assessing how participants use the tool [18], neither
795 we were measuring the prototype efficiency for comparing hierarchies [12]. Instead, our
796 contribution lays in the assessment of the four visualization methods for the comparison of pairs
797 of biological hierarchies with respect to curation tasks.
798
799 A final remark about methods and tools for the comparison of biological taxonomies is that the
800 methods we evaluate in this work are those that have been mainly discussed in the scientific
801 literature; however, their implementation in tools is limited. The Taxonomic Tree Tool (Lin et
802 al.) is, as far as we know, one of the few tools available online that allows us to explore
803 relationships (such as ancestor-descendant relationships) when comparing two biological
804 taxonomies. In a post-study contact with some participants, we asked about the software they use
805 for taxonomy comparison. We obtained five responses and only one indicated to occasionally
806 use the Taxonomic Name Resolution Service (http://tnrs.iplantcollaborative.org) to compare a
807 list of plant names with an authoritative database of published names. We do not know the
808 precise reasons why the methods are not in widespread use, but some possibilities could be that
809 the developments only reached prototype versions and are not available online, that the
810 visualizations could be too complex, or the difficulty that comes with the lack of standardization
811 of data in different communities.
812
813
814 **Conclusions and Future Work**
815
816 This study contributes insights on the capacity of four visualization methods for hierarchy
817 comparison in typical biological taxonomy curation tasks. Twelve expert taxonomists took part

818 in a study and provided feedback. We performed quantitative as well as qualitative analysis. The
819 results clearly show differences among the methods, on both users' effectiveness and
820 satisfaction: the edge drawing method was preferred over other methods.
821
822 In this study, the data sets were selected to avoid bias, all participants used the same datasets, and
823 participants were able answer the questions in reasonable time. However, it will be interesting to
824 design a similar study with larger datasets. Another approach would be to design a study in
825 which the data would be specific and familiar to each participant.
826
827 Enhancements such as providing multiple views, adding visual cues at inner taxonomic rank
828 levels, and avoiding overloading caused by long names and hierarchical structure lines, are
829 insights for future research. Functions for searching, statistics and queries to retrieve the
830 information of a taxon will be considered in a future design of a visualization environment. We
831 also plan to further research on visual summary views to facilitate the comparison at different
832 taxonomic rank levels. Finally, it will be important to make biological taxonomy comparison
833 visualization tools available for use. This implies overcoming challenges such as promoting the
834 standardization of data to facilitate data sharing and comparison.
835
836

Acknowledgements

838 The authors would like to thank all expert taxonomists who collaborated with this research,
839 dedicated the time to participate and provided very valuable feedback. We would also like to
840 thank Eros Hernández-Romero and Ronald Andrés Bolaños-Rodríguez for their collaboration
841 during the software development and the interview sessions, and Shivam Agarwal for his
842 collaboration as test participant.

843
844

References

846
847 Ball-Damerow, J.E, Brenskelle, L, Barve, N, et al. Research applications of primary biodiversity
848   databases in the digital age. PLOS ONE 2019;14. Ed. by Silva, DdP:e0215794.
849 Beck F, Diehl S. 2010. Visual comparison of software architectures. In: *Proceedings of the 5th*
850   *international symposium on Software visualization - SOFTVIS '10*. New York, New York,
851   USA: ACM Press, 183. DOI: 10.1145/1879211.1879238.
852 Beck F, Wiszniewsky F-J, Burch M, Diehl S, Weiskopf D. 2014. Asymmetric visual hierarchy
853   comparison with nested icicle plots. In: *Joint Proceedings of the Fourth International*
854   *Workshop on Euler Diagrams and the First International Workshop on Graph Visualization*
855   *in Practice*. GraphVIP. 53–62.
856 Charmaz K. 2006. *Constructing grounded theory : a practical guide through qualitative*
857   *analysis*. London; Thousand Oaks, Calif.: Sage Publications.
858 Contian L, Huijie Q, Jiangning W, Liqiang J.Taxonomic Tree Tool -. *Available at*
859   *http://ttt.biodinfo.org/TF/* (accessed January 1, 2016).
860 Dang T, Franz N, Ludäscher B, Forbes A. 2015. ProvenanceMatrix. *International Workshop on*

861   *Visualizations and User Interfaces for Ontologies and Linked Data, VOILA 2015 - co-*
862   *located with 14th International Semantic Web Conference, ISWC 2015.*
863   Ghoniem M, Fekete J-D. 2001. Animating treemaps. In: *Proceedings of 18th HCIL Symposium-*
864   *Workshop on Treemap Implementations and Applications.*
865   Gleicher M. 2018. Considerations for Visualizing Comparison. *IEEE Transactions on*
866   *Visualization and Computer Graphics* 24:413–423. DOI: 10.1109/TVCG.2017.2744199.
867   Gleicher M, Albers D, Walker R, Jusufi I, Hansen CD, Roberts JC. 2011. Visual comparison for
868   information visualization. *Information Visualization* 10:289–309. DOI:
869   10.1177/1473871611416549.
870   Graham M, Craig P, Kennedy J. 2008. Visualisation to Aid Biodiversity Studies through
871   Accurate Taxonomic Reconciliation. In: *Sharing Data, Information and Knowledge*. Berlin,
872   Heidelberg: Springer Berlin Heidelberg, 280–291. DOI: 10.1007/978-3-540-70504-8_29.
873   Graham M, Kennedy J. 2007. Exploring Multiple Trees through DAG Representations. *IEEE*
874   *Transactions on Visualization and Computer Graphics* 13:1294–1301. DOI:
875   10.1109/TVCG.2007.70556.
876   Graham M, Kennedy J. 2010. A survey of multiple tree visualisation. *Information Visualization*
877   9:235–252.
878   Graham M, Kennedy JB, Hand C. 2000. A comparison of set-based and graph-based
879   visualisations of overlapping classification hierarchies. In: *Proceedings of the working*
880   *conference on Advanced visual interfaces - AVI '00*. New York, New York, USA: ACM
881   Press, 41–50. DOI: 10.1145/345513.345243.
882   Guerra-Gómez JA, Buck-Coleman A, Plaisant C, Shneiderman B. 2012. TreeVersity: Interactive
883   Visualizations for Comparing Two Trees with Structure and Node Value Changes. In: *Proc.*
884   *Conference of the Design Research Society-DRS2012*. 10.
885   Holten D, van Wijk JJ. 2008. Visual Comparison of Hierarchically Organized Data. *Computer*
886   *Graphics Forum* 27:759–766. DOI: 10.1111/j.1467-8659.2008.01205.x.
887   Lam H, Bertini E, Isenberg P, Plaisant C, Carpendale S. 2012. Empirical Studies in Information
888   Visualization: Seven Scenarios. *IEEE Transactions on Visualization and Computer*
889   *Graphics* 18:1520–1536. DOI: 10.1109/TVCG.2011.279.
890   von Landesberger T. 2018. Insights by Visual Comparison: The State and Challenges. *IEEE*
891   *Computer Graphics and Applications* 38:140–148. DOI: 10.1109/MCG.2018.032421661.
892   Larsen BB, Miller E, Rhodes MK, Wiens JJ 2017. Inordinate Fondness Multiplied and
893   Redistributed: the Number of Species on Earth and the New Pie of Life. *Quarterly Review*
894   *of Biology* 92.
895   Linné Carl von, Gmelin Johann Friedrich. 1767. *Systema naturae per regna tria naturae :*
896   *secundum classes, ordines, genera, species, cum characteribus, differentiis, synonymis,*
897   *locis*. Lipsiae :impensis Georg. Emanuel. Beer,.
898   Lutz R, Rausch D, Beck F, Diehl S. 2014. Get Your Directories Right: From Hierarchy
899   Visualization to Hierarchy Manipulation. In: *Proceedings of the 2014 IEEE Symposium on*
900   *Visual Languages and Human-Centric Computing*. VL/HCC. IEEE, 25–32. DOI:
901   10.1109/VLHCC.2014.6883017.
902   MacKenzie IS. 2013. *Human-Computer Interaction: An Empirical Research Perspective*. San
903   Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
904   Munzner T. 2015. *Visualization Analysis and Design*. CRC Press.
905   Munzner T, Guimbretière F, Tasiran S, Zhang L, Zhou Y. 2003. TreeJuxtaposer: Scalable Tree
906   Comparison Using Focus+Context with Guaranteed Visibility. *ACM Trans. Graph.* 22:453–

907   462. DOI: 10.1145/882262.882291.

908 Sancho-Chavarria L, Beck F, Mata-Montero E, Weiskopf D. 2016. Visual Comparison of
909   Biological Taxonomies: A Task Characterization. In: *Poster aceptado para ser presentado*
910   *en EuroVis*.

911 Sancho-Chavarria L, Beck F, Weiskopf D, Mata-Montero E. 2018. Task-based assessment of
912   visualization tools for the comparison of biological taxonomies. *Research Ideas and*
913   *Outcomes* 4:e25742. DOI: 10.3897/rio.4.e25742.

914 Vaidya G, Lepage D, Guralnick R. 2018. The tempo and mode of the taxonomic correction
915   process: How taxonomists have corrected and recorrected North American bird species over
916   the last 127 years. *PLoS ONE* 13.

917

Biological taxonomy curation tasks.

Category	Task
Pattern Identification	*1. Identify congruence*: Identify same taxonomic concepts.
	2. Identify corrections: Identify splits, merges, moves, renames.
	3. Identify additions/exclusions: Identify new or missing taxa.
	4. Overview changes: Obtain a global view of changes.
	5. Summarize: Obtain numerical understanding of change.
Query	*6. Find inconsistencies*: Recognize violation of rules (e.g. repeated names)
	7. Filter: Find cases that satisfy certain conditions.
	8. Retrieve details: Retrieve the attributes of a particular concept.
	9. Focus: Navigate and see the information in detail.
Edit	*10. Edit*: Make changes to the taxonomies.

Participants' profile.

Id	Area	Degree	Professional Experience (years)
E1	Botany, Bioinformatics	PhD	32
E2	Botany, Forestry	MSc	10
E3	Informatics, Bioinformatics	Engineer	21
E4	Botany, Forestry	PhD	28
E5	Botany	MSc	15
E6	Botany, Bioinformatics	PhD	31
E7	Botany, Forestry	Master	21
E8	Botany	MSc	21
E9	Botany	PhD	23
E10	Biology	MSc	30
E11	Ecology	PhD	12
E12	Ichthyology	PhD	32

Derived datasets.

	T_2	T_3	T_4	T_5
Nodes	116	78	105	114
Species	86	55	75	84
Splits	6	6	7	9
Merges	7	3	5	3
Moves	6	4	4	6
Renames	7	4	6	6
New	6	4	3	6
Excluded	4	3	4	8

Figure 1

Agglomeration method implementation.

Taxonomies are agglomerated into one hierarchical structure. Changes are indicated by relationships between the two taxonomies. In order to distinguish to which taxonomy the taxon belongs to, a dual use of color was introduced. For each taxon's specific change, light nuanced nodes indicate that the taxa belongs to the taxonomy of origin and the darker nuanced nodes indicate that they belong to the taxonomy of destination. Through the colored toggle buttons on the menu, users can visualize either one or several types of changes at once. Users can select a specific taxon in order to visualize its changes.
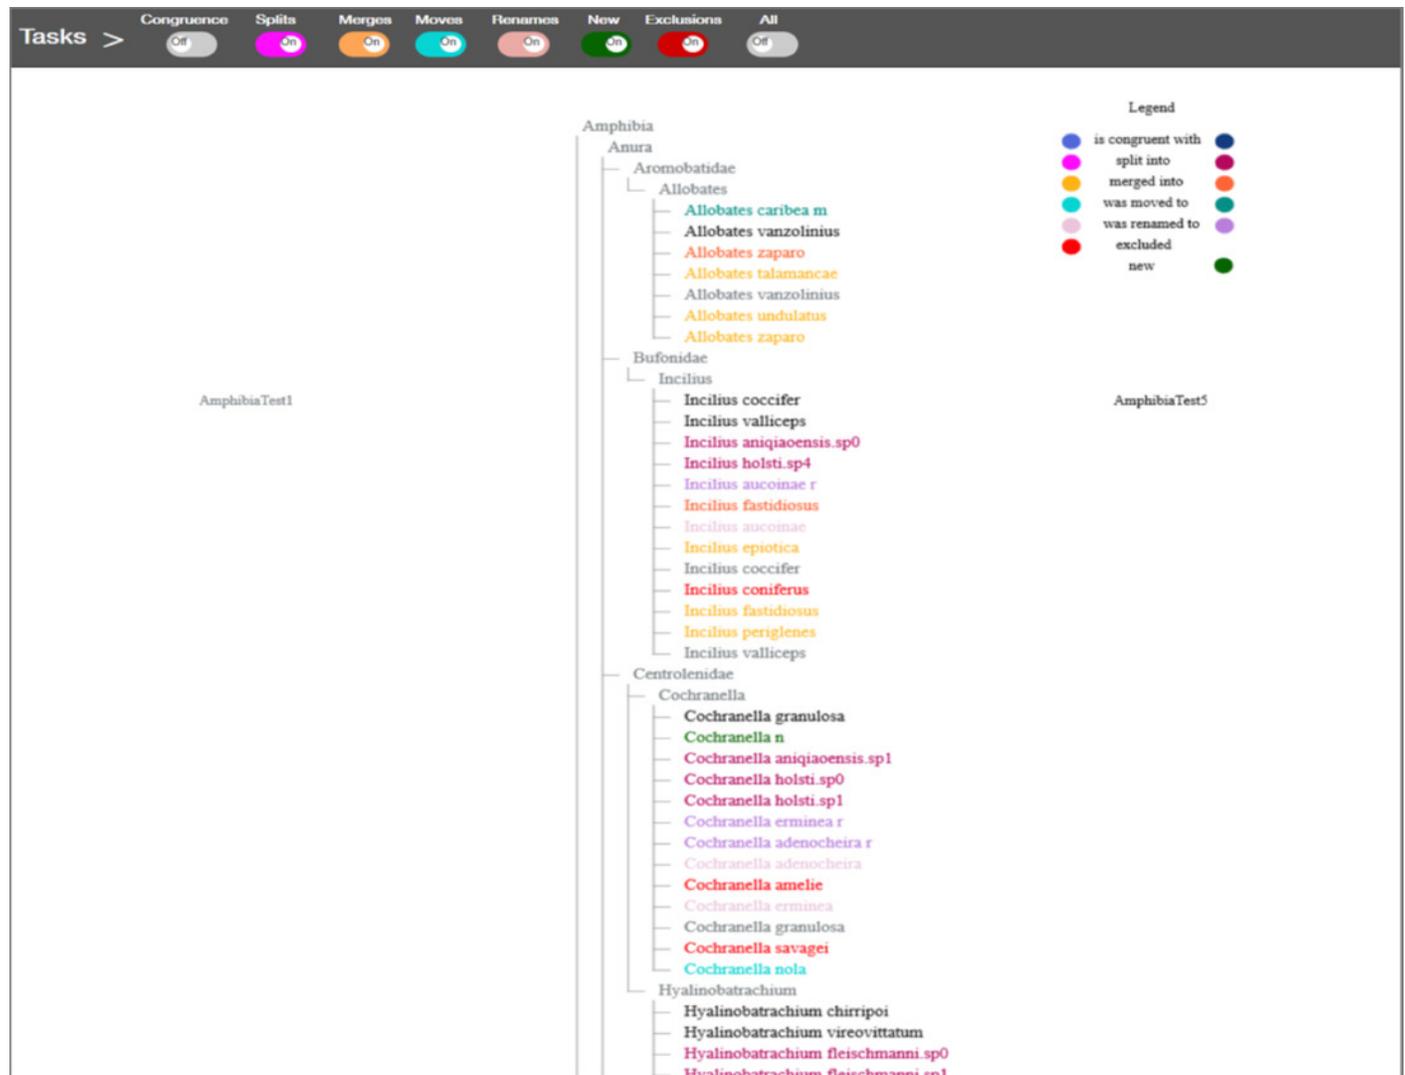
Figure 2

Animation method implementation.

Changes are presented through animation. Besides the common toggle buttons on the main menu for all methods, the animation method includes a play/stop button, a speed slider, and a button that controls whether the animation will present the changes either one by oner or simultaneously. During animation, taxa moves from the origin taxonomy on the left side to the destination taxonomy on the right, displaying the transformation. Through the colored toggle buttons on the menu, users can animate either one or several types of changes at once. Users can select a specific taxon in order to visualize its changes.

Figure 3

Edge drawing method implementation.

Changes are indicated by relationships between the two taxonomies. In the edge drawing implementation, relationships are visualized by color-coded edges that go from one taxonomy to the other. Through the colored toggle buttons on the menu, users can visualize either one or several types of changes at once. Users can select a specific taxon in order to visualize its changes.
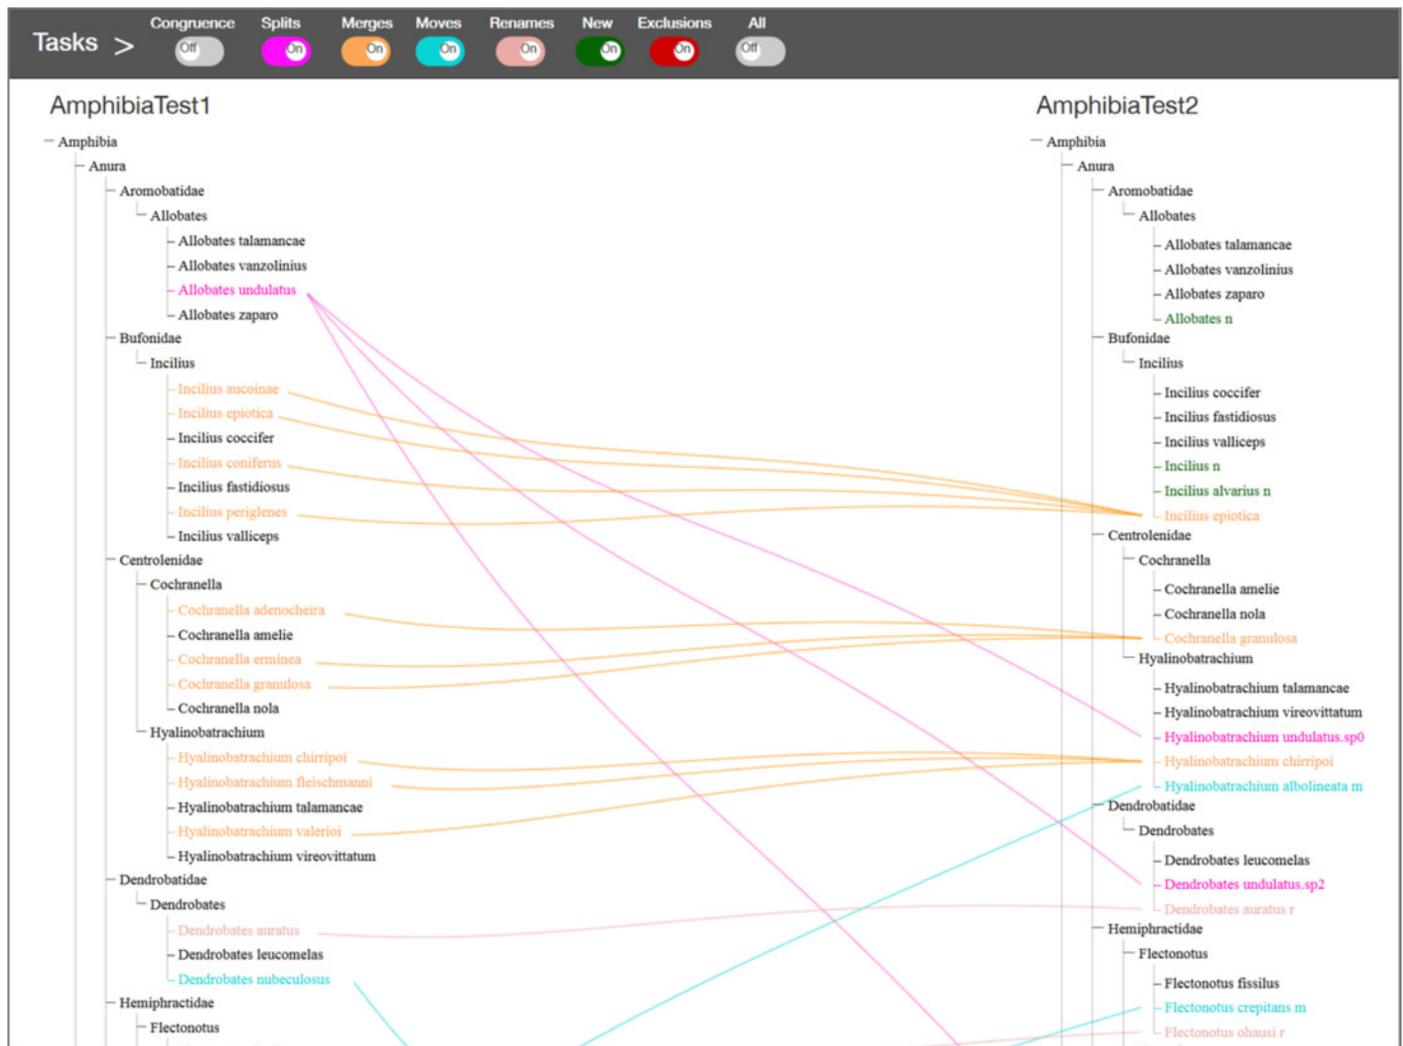
Comparison of participants' effectiveness using Cochran's Q test.

Each row corresponds to a task-performing exercise. Each cell of the row contains the results of the participants' effectiveness with each method for the evaluated exercise, as well as the chi-square and p-value resulting from the Cochran's Q test. Resulting less-effective methods for each exercise are shaded in orange.

Exercise	Frequency (%)				χ^2	p-value
	Agg	Ani	Edg	Mat		
1. Which is the most common type of change?	9 (75)	11 (92)	1 (8)	11 (92)	25.500	0.000
2. Into what taxa was taxon "t" split?	11 (92)	10 (83)	11 (92)	11 (92)	3.000	0.392
3. Which species was split most?	4 (33)	7 (58)	12 (100)	11 (92)	17.571	0.001
4. Was species "s" merged with any other species?	11 (92)	11 (92)	12 (100)	12 (100)	3.000	0.392
5. With which other species was taxon "t" merged?	12 (100)	11 (92)	12 (100)	12 (100)	3.000	0.392
6. Which is the new name of taxon "t"?	12 (100)	12 (100)	12 (100)	12 (100)	–	–
7. Which was the previous name of "t"?	11 (92)	12 (100)	12 (100)	12 (100)	3.000	0.392
8. Which species were moved to genus "g"?	5 (42)	12 (100)	12 (100)	12 (100)	21.000	0.000
9. Which family has the most species added?	11 (92)	11 (92)	12 (100)	12 (100)	2.000	0.572
10. Genus to which more than one species were excluded?	11 (92)	12 (100)	12 (100)	12 (100)	3.000	0.392
11. What types of changes occurred to taxon "t"?	6 (50)	9 (75)	11 (92)	12 (100)	9.692	0.021
12. Identify which genus shows most changes.	2 (17)	7 (58)	7 (58)	7 (58)	10.714	0.013
Effectiveness (overall)	**105 (73)**	**125 (87)**	**126 (88)**	**136 (94)**	**40.480**	**0.000**

Figure 4

Matrix method implementation.

The two taxonomies are arranged along the horizontal and vertical axes of the matrix. Changes between the two comparable taxonomies are highlighted thorugh the color-coded cells. Through the colored toggle buttons on the menu, users can visualize either one or several types of changes at once. Users can select a specific taxon in order to visualize its changes.
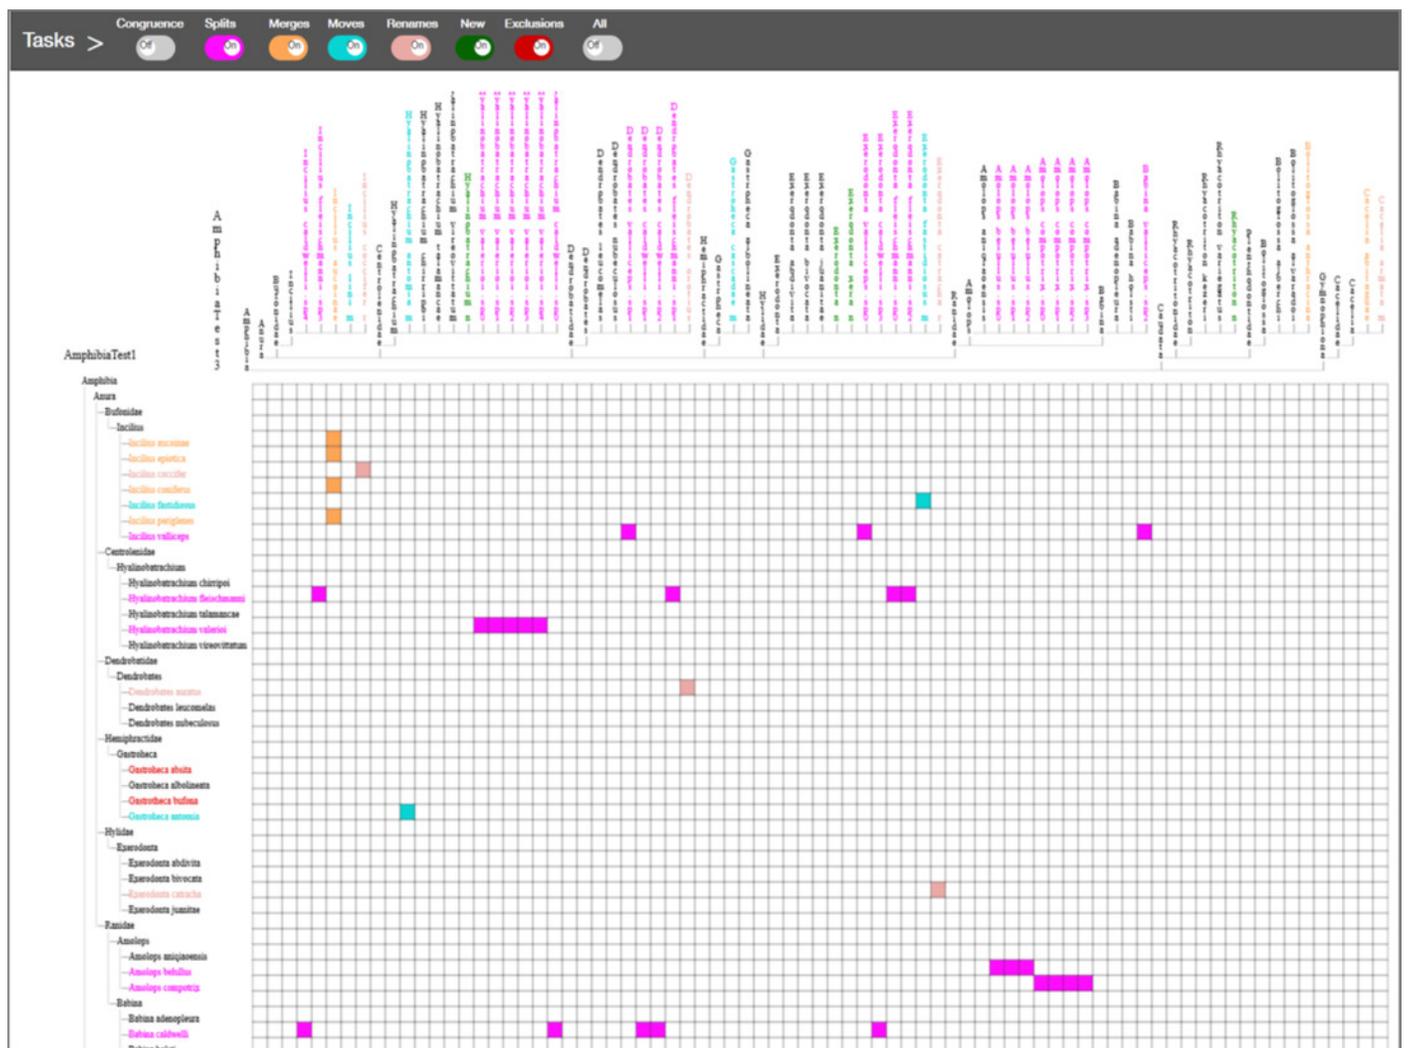
Comparison of participants' satisfaction using Friedman test.

Each row corresponds to a task-performing exercise. Each cell of the row contains the results of the participants' satisfaction level with each method for the evaluated exercise, as well as the chi-square and p-value resulting from the Friedman test. Higher satisfaction level is shaded in green and less satisfaction level in orange.

Question How good do you think is each method in order to …	Median				χ2	p-value
	Agg	**Ani**	**Edg**	**Mat**		
1e- … identify the most common type of change?	2.08	1.79	3.29	2.83	12.588	0.006
3e- … identify splits?	1.50	1.50	3.79	3.21	33.055	0.000
5e- … identify merges?	2.00	1.71	3.71	2.58	20.050	0.000
7e- … identify renaming of taxa?	2.13	1.50	3.67	2.71	21.559	0.000
8e- … identify moves?	1.54	1.88	3.83	2.75	24.295	0.000
9e- … identify new species added?	2.75	2.63	2.96	1.67	9.539	0.023
10e- … identify excluded taxa?	2.67	1.88	2.79	2.67	7.062	0.070
12-e- … identify changes to a taxon?	1.88	1.79	3.46	2.88	16.057	0.001
13- … visualize differences and similarities between two taxonomies? (Overall)	2.97	2.84	4.97	4.09	25.064	0.000