RESEARCH ARTICLE

Open Access

Helium: visualization of large scale plant pedigrees

Paul D Shaw^{1,2*}, Martin Graham², Jessie Kennedy², Iain Milne¹ and David F Marshall¹

Abstract

Background: Plant breeders use an increasingly diverse range of data types to identify lines with desirable characteristics suitable to be taken forward in plant breeding programmes. There are a number of key morphological and physiological traits, such as disease resistance and yield that need to be maintained and improved upon if a commercial variety is to be successful. Computational tools that provide the ability to integrate and visualize this data with pedigree structure, will enable breeders to make better decisions on the lines that are used in crossings to meet both the demands for increased yield/production and adaptation to climate change.

Results: We have used a large and unique set of experimental barley (*H. vulgare*) data to develop a prototype pedigree visualization system. We then used this prototype to perform a subjective user evaluation with domain experts to guide and direct the development of an interactive pedigree visualization tool called Helium.

Conclusions: We show that Helium allows users to easily integrate a number of data types along with large plant pedigrees to offer an integrated environment in which they can explore pedigree data. We have also verified that users were happy with the abstract representation of pedigrees that we have used in our visualization tool.

Background

The effects of climate change and ensuring food security in a world with an increasing population is becoming ever more pertinent [1-3]. The exploitation of pedigrees in plant breeding allows breeders to target specific plant crosses to maximise the potential of achieving desirable agriculturally important characteristics such as yield, drought/water tolerance and disease resistance which will be required if new varieties are to be bred to cope with increased demand in a changing environment.

The ability to predict and visualize the inheritance of alleles that facilitate resistance to pathogens or any other commercially important characteristic is crucially important to experimental plant genetics and commercial plant breeding programmes. Derivation of the inheritance of such traits by traditional molecular techniques is expensive and time consuming, even with recent developments in high-throughput technologies. This is especially true in industrial settings where, due to time

constraints relating to growing seasons, many thousands of plant lines may need to be screened quickly, efficiently and economically every year.

Due to their complexity, there is a cognitive limitation in conceptualising large pedigree structures.

While it may not be achievable or indeed necessary to understand every mating relationship between related individuals, an overall picture can lead to insight into the data and any patterns it may contain. This can also aid in the identification of problems (both biological and data handling issues) within datasets when coupled with expert domain knowledge.

This is particularly important when looking at pedigree data as the context in which each line sits may hold additional and important information (such as the inheritance of particular genome regions from ancestral varieties). It is because of this that a combination of visual and statistical analytics would allow geneticists and commercial breeders to gain a deeper understanding of the transmission of genetic elements within a pedigree based framework but there is currently a lack of suitable tools to analyse these data types.

Software tools that offered improvements in the speed at which this analysis can be carried out, and increase

* Correspondence: paul.shaw@hutton.ac.uk

¹Information and Computational Sciences, The James Hutton Institute, Invergowrie, Dundee DD2 5DA, UK

²School of Computing, Edinburgh Napier University, 10 Colinton Road, Edinburgh EH10 5DT, UK

users' ability to conceptualise large pedigrees would bring both time and cost gains to breeding companies.

Using a unique and extensive barley dataset covering pedigree, genotypic and phenotypic data for UK elite germplasm which has been through the UK National List Testing procedures [4], we discuss the challenges of visualizing the transmission of alleles encoding traits and characteristics of agricultural importance in a pedigree-based framework. We then describe the subsequent development of a pedigree visualization tool that was implemented in close collaboration with domain experts.

Plant pedigrees

A pedigree (Figure 1) is a representation of how genetically discrete individuals are related (usually but not exclusively) in time to one another. It is therefore a representation of the genetic relationship between individual plant lines, their parents and progeny (predecessors and successors). Pedigrees are often used in human contexts to show the transmission of alleles responsible for genetic conditions of medical importance. In plants they are used as a framework along with environmental data, on which statistical analysis can be used to determine factors such as mode of inheritance (Identity by Descent, IBD and Identity by Association, IBA). Additionally, they are often used to check for potential genotyping errors, since these errors, by the very nature of Mendelian inheritance, are constrained by the pedigree structure in which they exist [5]. The accurate representation of pedigrees is therefore becoming increasingly important in plant breeding and genetics.

While there are defined standard nomenclatures for human pedigrees [7] there is no single formal system for plant pedigrees, however, there are moves towards

defining standards. There are valid biological reasons for this including: the hermaphrodite nature of most plant species, the complexity of mating designs possible in plant genetics and, finally, the absence of any overseeing coordinating organisation.

While plant and animal breeding share routine breeding techniques such as standard crossing and back-crossing, pedigrees used in plant breeding display some subtle but important differences, often involving key shorthand conventions that are unique to plant mating designs leading to complex textual based records which can be difficult to read (see 'Pedigree formats' subsection). Firstly, the named entities in plant pedigrees may, but not always, represent a population of genetically identical individuals, not a single plant. While it is relatively simple to grow many plants from seed, potentially many decades after production, in humans and animals this is understandably not the norm. The generation of these genetically identical (homozygous) varieties is possible through doubled haploidy, inbreeding, or crossing of pairs of inbred lines to achieve what is termed an F1 hybrid. Successive inbreeding by self-pollination of these F1 generation plants leads to individual plants that are close to homozygous across all alleles. The exploitation of homozygous lines in crop species such as barley is a powerful tool in genetic analysis, removing some of the genetic complexities associated with species (such as humans) where there is a high level of heterozygosity.

Pedigree formats

- a. A/B//C//D
- b. ((A * B) *C) *D
- c. [A × [(B × C) * D] × E] * [F × A] × C

Genetic transmission

Pedigree formats can be complex with no standard nomenclature; a. Purdy Notation System [8] was put forward by Purdy as a common format for representing small grain cereal pedigrees. Forward slashes '/' are used to delimit lines. In this case A is crossed with B which is then crossed with C whose progeny is crossed with D. b. Lamacraft and Finlay notation [9] which was put forward as a format which could be more easily parsed by computers. The example here is the same as in the Purdy notation above. c. A typical pedigree that can be found in old records where a mixture of notations are used. These mixed notation systems are common and most breeders will use shorthand that is unique to them. These records are sometimes difficult to read and would benefit from being represented in a more user friendly way.

Data sets

There are a number of different data types used in this work. Our primary data set is composed of a large barley pedigree data set for 803 UK Elite cultivars as well as Single Nucleotide Polymorphism (SNP) genotypic data for 750 of these lines across 4,769 genetic markers. In addition, phenotypic data for these lines for 33 Distinctiveness, Uniformity and Stability (DUS) characters [10] across multiple years and sites was used (1980 - present which equates to 601,148 data points). Datasets covering UK wheat (*Triticum* spp.) and Asian rice (*Oryza sativa*) were also used in this work although these are more limited in size. Data are stored in the Germinate 2 database system. [11]. The ability to connect to Germinate was an important design decision as allowing users to access all background information on plant lines that we had available was important.

Pedigree definitions

The nucleus of pedigree data are a series of parent/child relationships defined as encoded strings (see 'Pedigree formats' subsection) [8,9]. Data was atomised into simple parent/child definitions which were used to dynamically reconstruct the pedigree. In addition there may also be information identifying whether the parent was male or female and the type of genetic cross performed. Something unique in plant breeding is where a plant can be both male and female parents in the same cross.

Complications may arise from either older pedigree data which is error prone and may be difficult to verify without expert guidance and from the re-use of names to describe varieties creating false relationship joins. It is not uncommon for a breeder's favourite name to be used multiple times until a line is adequately different, and has sufficient performance to be accepted for wider distribution into the UK recommended list programme.

Genotypic data

The genotypic data set for our study is based on a set of SNP markers which are mapped to known chromosome positions in the barley genome. Each plant line within the test set has been genotyped for a set of 7,000 of these markers.

A given plant variety will have an allele call for each of a series of loci represented as a pair of nucleotide bases e.g. AA, GG (which are homozygous) or AG (which are heterozygous), for a locus. Due to the inbred nature of our barley germplasm there are low levels (less than 0.5%) of residual heterozygosity present.

Phenotypic data

The phenotypic data in our study has been either collected in field experiments or by molecular testing. Though many of the agriculturally important traits are controlled by many genes of small effect (quantitative traits) for simplicity we concentrated on traits under simple genetic control. Examples of such traits include DUS characteristics which are used in the varietal registration and seed certification process and allele data on disease resistance genes such as *Mlo* and *Mla*.

Previous work

The ability to visualise data is imperative in modern experimental plant genetics, with volumes of data being routinely produced far exceeding the ability for humans to digest and identify underlying phenomena. Until now, pedigree visualization, with few exceptions [12,13] has primarily been focussed on work carried out in the human genetics domain. Because plant breeding programmes involve phenomena not normally seen in human populations, such as routine inbreeding, there are additional visualization challenges that need to be overcome. There are often large numbers of plant lines involved in any pedigree, many more so than in an average human pedigree due to factors such as generation time/time to sexual maturity which is far lower in most plant species than that of their mammalian counterparts. This section will look at the various visualization techniques used to represent pedigree based data and highlight the problems and strengths that these techniques exhibit.

Table-based approaches

Table-based visualization tools such as Flapjack [14] address some of the problems associated with visualizing large datasets and are optimized for efficient sorting and querying of genotypic and phenotypic data, but currently lack the ability to display data on a pedigree-based scaffold.

While other tools such as PedStats [15] offer statistical validation of users' pedigree data without visualization of the actual pedigree structure, it is difficult if not

impossible to conceptualize pedigree structure for complex data sets without some visual representation.

Matrix-based visualizations to represent pedigrees use the intersection of the x and y edge to define relationships. Matrix-based visualizations have advantages over node-link or graph-centred layout approaches including the ability to create compact graph representations and the ability to remove edge overlapping. However, tests generating matrix visualizations using our pedigree data have shown that the data density is so low the resulting representations are not particularly insightful. The ability to easily track flow and identify paths is also removed.

Tools such as GeneaQuilts [16], offer a new visualization technique suitable for use with thousands of individuals but offer limited scope for addition of complex genotypic and phenotypic data and discussions with our users showed that they found it difficult to easily interpret such representations.

Finally, tools such as VIPER [17] offer novel pedigree visualization and genotypic error checking capabilities. VIPER is essentially a stack of nested table representations of generations where rows represent sires, dams or children and columns represent individuals which can span multiple columns where they are parents. VIPER's primary use is in identification of genotyping problems in farmed animals and would be unsuitable for visualizing the complex crossing relationships that exist between crops where selfing is not uncommon. VIPER requires both separate male and female parents which is the norm in any applications handling animal or human data, but not always the case in plant breeding.

Graph-based

Unlike trees, graphs allow for the precise modelling of the complexity of a plant breeding programme. Techniques such as node link diagrams have long been used as a way of representing graph-based data and recent work has examined how effective the node-link model performs representing graph data when compared to matrix-based visualizations [18]. Work carried out by Purchase [19,20] and Bennett [21] also indicated that while graph layout played an important part in a user's understanding, it was not the major focus; this focus perhaps being the use of other aesthetics relating to node colour and shape.

Most of the current tools have been developed for human pedigrees where consanguineous mating events are negligible. This is not the case in plant and animal breeding which cannot be properly modelled using tools that use node-link or tree hierarchies such as Pedfiddler and Madeline [22].

Cranefoot [23] reports the use of mathematical graph structures to deal with between-relative mating but the approach is limited in its current form in the amount of information that can be attached to a node. Finally,

HaploPainter [24] allows the drawing of genetic haplotypes, but suffers from being restricted in the number of individuals it is able to display.

A commonly used two-dimensional pedigree visualization tool is Peditree [13] which offers a tree-based view of data in a pedigree but this is not suited to our requirements as plant pedigrees are not trees (inbreeding and the use of older lines in more modern crosses prevents us from treating them as such). Other tools such as the Pedigree Visualizer by Wong [25] offer new layout algorithms. Wong suggests introducing duplicate "alias" lines in representations with multiple matings from the same individuals, phenomena that are commonplace in plant data. PyPedal [26] not only offers rudimentary graph drawing tools, restricted to changing node shape to represent male and females, but also error checking algorithms to try and identify potential pedigree errors where appropriate genotypic data exists.

Visualization techniques such as sunbursts [27] which are space filling versions of a node-link diagram have the advantage that a node's position in a hierarchy is maintained. Additionally, Fan Charts [28] and H-trees [29] have also been described as a means for recounting human genealogy; these techniques however assume no inbreeding (they are trees and not graphs) and thus rule themselves out for use with plant pedigrees.

While the main problems with these additional techniques are that they are not appropriate for observing a pedigree in its entirety (indeed the complexity of the data may rule many of them out), they may be useful when trying to visualize a sub-section of data such as a sub-pedigree for specific lines.

Layout algorithms

Plant pedigrees often form what we describe as a *pedigree net*, whereby there is structure to the graph but it's not as simple as traditional top-down pedigree representation that is seen in humans and to a lesser extent in farmed animals (Figure 2).

This abstract representation does include a time component in the form of generations, but due to the viability of seed, and the existence of varieties and landraces that may be many hundreds of years old, there is the potential to use these older varieties in modern crosses. This situation leads to nodes at the top of the graph having edges connecting to nodes at the bottom - this is not common in animals and would be extremely unlikely in humans. The existence of a time component means that the use of a layout algorithm that preserves topology (top-down generations) is nonetheless important as most (but not all) crossing will be between newer varieties. Because of this, layout methodologies such as force-directed algorithms (Figure 3B) would not offer the ability for us to arrange our pedigree based on time. Force

directed layouts are not well suited to our requirements. The lack of a visually identifiable pedigree structure is strikingly apparent.

The problem of very large pedigrees in humans has been identified and solutions proposed in tools such as PVin [30] which looks at windows on large datasets but only offers pedigree drawing with no scope for addition of other information onto the visualization. In addition, its traditional human family tree output is not the most efficient use of space for plant pedigrees which form a more dense net due to the nature of reproduction which is not seen in humans or animals (Figure 2A)

Although there are problems associated with 2D node-link layouts such as a lack of horizontal space and problems with crossing of edges [31] they are still well suited to displaying data of this type. 3D tools also have

their problems, including visual occlusion and that they tend to visualise high-level features and not specifics, so while some trends are easy to spot, the actual detail is hidden from the user. From this point of view they are limited in use for our purposes and offer no advantages over their 2D counterparts. Notable examples of such tools are Walrus [32] and Celestial3D [31] but their success lies in alternate problem domains.

Discussion

It is clear that these techniques and tools contain many features that are useful, but none meet the exact requirements (including data abstraction) of our problem to be able to overlay genotypic and phenotypic data onto a complex pedigree structure.

Figure 3 Layout types. Barley pedigree for Quench which is a commonly used high yielding spring malting variety from Syngenta. It is a cross between Sebastian and Drum. **A.** shows the standard Sugiyama-style layered layout using dot and **B.** the same data using the fdp force-directed layout. Both layout tools are from GraphViz. The traditional pedigree top to bottom topology is lost using force-directed layout algorithms.

There is a need for the development of tools that are tailored for the unique needs of plant breeding with the ability to explore pedigree structure, and paint additional genotypic and phenotypic data on top, to allow breeders to make informed decisions and visualize the way in which alleles for agriculturally important traits are transmitted through previous and subsequent generations. Such tools do not currently exist.

Through the examination of methodologies to display pedigree data we suggest that the best method to visualize plant pedigree data is a layered layout (Sugiyama-style) based approach (Figures 2A and 3A). Not only does this allow us to accurately map the exact specifics of how breeding programmes run (including inbreeding) but also provides a well-established framework onto which a visualization can be built. The use of graphs as our data structure means that features such as standard graph-traversal algorithms can be used to bring greater functionality to our pedigree structure in locating ancestors and descendants and as a logical framework which can be used to look for problems with underlying datasets. The layered layout representation also brings a coherent structure to sparse relationships and generations and topological layout are clearer compared to matrix style layouts. This is not the case with animal (Figure 2B) and human pedigrees whose top-down fan type shape is not well suited to a layered layout as they quickly become very large, consuming large volumes of horizontal space [17].

Tools that allow exploration of data to try and bring a greater understanding of complex relationships between individuals should bring greater insight into how plant breeding programmes operate at the genetic level and how to bring maximum potential benefit from them. The ability to detect patterns and associations (or even anomalies) within these datasets such as; the identification of

problems with inheritance of alleles, the identification of lines from which additional information would allow inference of data on large parts of the pedigree, simple typos and errors, or looking for lines which are similar to unknown lines, will lead to increased depth of domain knowledge for plant breeders and geneticists.

Methods

Paper-based visualization

We wanted to test if our use of a DAG based data structure and layered layout approach would work with our barley pedigree data and would be accepted by our users. In order to do this a paper-based layout was implemented, overlaying basic character data on to the graph nodes represented by colour and sizing nodes based on the number of times they had been used in crosses in our data. In this prototype (which was implemented in Perl and the Graphviz dot library) our pedigree was modelled as graph nodes to represent plant lines and edges to show mating/parentage. While GraphViz has been used before in pedigree drawing [33], examples focus on a small number of individuals.

While initially this prototype was run by users as a command-line computer program which generated images based on input `_les` and generated an image which could be viewed on their computer monitors it was decided that printing this static representation (2.5 m × 1 m see Figure 4) would allow domain experts to better interact with the visualization. We overlaid, by means of colouring nodes, the winter/spring ecotype category on this dataset as (along with the 2-row/6-row ecotype) it is the most commonly used physiological means of differentiating barley varieties, and one that all of our test users were familiar with. This tool was also implemented as a web-service which allowed us to include static (but

Figure 4 Helium static prototype. Users interacting with the large static prototype implementation of our pedigree visualization. In this example 2 and 6 row ecotypes are coloured green and blue respectively and varieties (nodes) are sized based on their contribution to the overall pedigree in terms of the number of lines that are derived from them. This shows which lines are most commonly used in barley breeding in the UK. The 2/6 row division is one of the important characteristics used to differentiate barley worldwide. 2 row barley is primarily of spring type and used in brewing and distilling while 6 row varieties are used in feed due to decreased quality characteristics for the brewing and distilling industry. Consent was obtained from both participants for publication of their images.

dynamically generated) pedigree representations within our internal barley information portal.

Feedback on paper-based prototype

Through observation and talking to twelve geneticists and plant breeders while they interacted with our wall-mounted visualization it was clear that there were a number of issues associated with this implementation. Firstly, it was almost impossible to trace edges between nodes when the data was dense (even at a large output size) so we found ourselves falling back on examining text based records to confirm lineage. Secondly, it is incredibly challenging to quickly locate specific plant lines with this density of data. Commonly used lines are immediately identifiable due to the use of size to represent the number of uses in breeding crosses but these are not always what users are most interested in. Users used these larger nodes as reference points, almost as if they were notable points on a map [34,35] and attempts at using slightly different layouts or orientations were not well received.

It was also clear that users were beginning to quickly spot pedigree problems. These problems related to the parentage of lines and in some cases the assignment of ecotype. These types of errors would be extremely difficult for a user without extensive experience to pick up on and this has not only shown that it is an effective technique for visualization but also an effective way of identifying errors with underlying datasets.

Users liked this representation of large pedigrees. Not only is it visually attractive, but geneticists were using it

to identify problems with the underlying pedigree and phenotypic data in a way that is more interactive, social, and tactile compared to the examination of records.

When presented with our results, plant breeders told us that it gave them an overview of their data that was not currently available to them; indeed these representations uncovered interesting information relating to the relative frequency of use of particular 'key' lines in the UK Elite Barley germplasm that would have been difficult to see from textual records in the format seen in 'Pedigree formats' subsection, such records have not been collated like this before. Missing data was also easily spotted thus allowing us to update our underlying datasets.

Problems do however exist, especially in the inability to search for particular plant varieties and tracing of edges to establish lineage. In order to try and address these, it was quickly realised that we would need to move towards the development of a more interactive software tool - Helium - named after the balloon type appearance of our static prototype.

The Helium prototype

Taking the feedback obtained from our initial informal user testing, an interactive detail and overview [36] prototype pedigree visualization system using Java and the yFiles library from yWorks [37] (Figure 5) was implemented. This prototype maintained the same visual metaphors (nodes and edges) to describe pedigree structure but now could add features to allow users to search and explore the data and link in plant passport, phenotype and background data from our Germinate database. One of the design decisions to use Germinate was that we can ensure that researchers working on our barley data will all be using the same data from the same source.

While our paper prototype included a single static image it was clear that when users were viewing our visualization on computer monitors there would be a limitation on the number of nodes that could be displayed while still retaining legibility of line names. To address this our main visualization panel (Figure 5A) can be zoomed and panned to allow users to explore data. An overview panel was added (Figure 5B) which would allow users to track where they were in the main visualization window and give a high-level overview of the pedigree structure. The overview would act as a common reference point for our users that would not change as the main visualization window was manipulated. Feedback from our paper implementation also showed that users would want to get as much background information as possible on lines and so a detail panel was added (Figure 5C) which displays passport and general background information. Data from Germinate is displayed in the detail panel and is pulled on demand

based on a user's selection in the main visualization window.

Germinate also includes phenotypic data of both nominal and ordinal types which were colour-coded in Helium using ColorBrewer2 palettes [38,39]. Hue was used to differentiate nominal data and saturation to distinguish ordinal data classes for phenotypes and genetic similarity metrics within our visualization [40] (Figure 6).

User testing of the Helium prototype

User testing is an important aspect of the development lifecycle of visualization [41-43]. Both Munzner and Lam lay out the requirements for testing, specifically relating to visualization studies in both contemplation and reflection of user studies. A subjective evaluation was performed to establish user perception/acceptance and understanding of the visualization methods within Helium. This was to establish empirically if users were happy with representing data as graphs, moving away from the traditional family-tree type methods, and whether the use of graphs fits in with a user's perception of pedigree structure and function. Could our users perform basic pedigree operations such as accurately tracking back through generations and find information they require using our visualization? We

also wanted to ensure that users were able to interact well with our methods which allow much greater data density and increased plant line density.

The testing data was obtained through a questionnaire and comment-based feedback based on how intuitive our users found the main features of the prototype to be. We also asked how our tool could be improved relating to general usage or new features. This is important as while initial user-requirements were gathered, when our users actually started using our software we had expected them to come up with new ideas on features or utility that would benefit their research.

This feedback allowed us to improve our interface and visualization to help increase our users understanding of the system and underlying biological concepts.

User testing methodology

A pre-screening questionnaire, user tasks, and a follow up questionnaire centred on predefined tasks that users would be asked to perform was developed. The initial questions were to gain an overall impression of the length of experience the user has had in this field, and to classify their job title. There are two distinct groups of potential users: bioinformaticians/computational biologists and

Figure 6 The Helium interface. Our visualization tool has been split into a number of distinct areas which are shown here. Our choice in using these coordinated multiple views is a common design choice in visualization interfaces [36]. The use of a high-level overview window (Figure 6A) and main view (Figure 6B) helps users maintain orientation and provides a filtering mechanism for a detailed local view (Figure 6C). This allows easier tracing of lineage whilst maintaining greater context with the main visualization window. **A)** Data selection and overview panel. This contains a high level overview of the entire pedigree and does not change (apart from node colouring and resizing) during a user's interaction with the application. This is intended to be used as a reference point for large pedigrees. This panel also contains a set of tabbed panes that let a user select phenotypic or genotypic data to overlay as well as further information from our Germinate database. **B)** The main pedigree visualization panel is where most manipulation happens. Users can pan and zoom across this window. The nodes are hotspots both for selection using the mouse or by hovering additional information is displayed about both nodes and edges. **C)** The local view panel contains our local view of a selected line and offers tools to allow a user to define how many generations (forwards and backwards) they want to view. The local view removes much of the visual clutter associated with the main pedigree visualization and allows edges to be more easily tracked. **D)** The detail panel contains search functionality and overview statistics. This example shows colouring for the ordinal data type "Anthocyanin Colour" which is a DUS character.

plant geneticists (experimental)/breeders (applied). User tasks were developed using our initial application requirements and were designed to force the users conducting the test to explore our experimental test datasets. The follow up questionnaire was clearly split into two sections; the first taking the form of attitude-scale questions on the user's opinion on the software and visualization in terms of both their use of it (assuming comparison to their current method of viewing these data types), and follow up subjective open-ended questions to get additional information that could be used to drive development of this software tool.

The questions assume that a comparison is being made to other methods that test subjects are, or have been using to obtain the same information, and we can use these to signify if our visualization and user interface brings significant improvements in visual representation and understanding of pedigree structure. Throughout

the study, notes were taken and screen and audio capture was used to further examine a user's interaction with the interface and to aid in recount of the tests.

Each test was scheduled to take around 45 minutes;

- 5 minutes - pre-questionnaire
- 5 minutes - familiarisation
- 25 minutes - test
- 10 minutes - post-test questionnaire

After completion of the main interaction study our users completed an attitude scale where they indicated their preference on a 5 point scale between "Very Difficult" (1) to "Very Easy" (5) relating to a number of statements about their use of this software.

The questionnaire asked users to detail features or concepts that they found to be confusing, those they found to be clear, and features that they feel would add

value to their research. Finally users were asked to provide general comments about their use of our software; this would be used to allow us to tweak and fine-tune the Helium interface to aid our users with their research.

Results

General background profiling

The 16 expert users that undertook this study break down as follows; 5 bioinformaticians, 10 plant geneticists and breeders and 1 statistician. Out of the users 94% were educated to PhD/MSc level and the average length of time working in their areas was 17 years. The minimum experience was 1 year, maximum 36 years giving a median length of experience of 13.5 years.

While all users were familiar with pedigree data, 69% used it on a day-to-day basis as part of their research and 38% regularly used alternative tools.

It should be noted that through verbal feedback it was established that the researchers who were using pedigree data were using paper records and spreadsheets to curate and maintain pedigree data used in their work and not a specific pedigree tool.

Main user interaction study

There were eight questions that users were asked to answer in using our pedigree interface. The questions were assigned an overall category and can be seen in Table 1 where we show the question classification along with the number of correct and incorrect responses.

Our user testing uncovered some interesting problems with our visualization. For example, the category “Identifying Children” from Table 1 asked our participants to identify the progeny of a specific barley variety. In 44% of completed questionnaires this answer was incorrectly given. However, when examining “Tracing Lineage” from Table 2 which related to this question, users thought that it was easy to trace lineage by following graph edges. Our test users were continually missing the same progeny (one of three) of the line; the one whose complete edge was not

Table 2 Post study questions (Scaled/Likert 5 very easy, 1 very difficult)

Question classification	Median (M)	Mode
Colour coding	3	3
Phenotype classes	4	4
Maintaining position	4	4
Clarity of relationships	4	5
Tracing lineage	4	4
Understanding data	4	4
Background information	4.5	5
Ease of use	4	4
Finding parents	4.5	5
Navigation	5	5
Children	5	5
Finding lines	5	5

immediately visible, and disappeared off the right-hand side of their computer display. When talking to a selection of users after the test had been carried out and asking them to perform the same question they did so without error (obviously suspicious to the reasons behind the request).

Post-study questionnaires (attitudinal and open ended)

After carrying out our main interaction study the users were asked to fill in a series of questions that asked them to compare Helium to pedigree tools, or methods of handling pedigree data that they are familiar with using, and to get feedback on what they found easy and difficult to understand or perform with Helium. These results are presented in Table 2.

Test results discussion

The most common responses have been detailed by dividing them into features users liked and disliked. These were obtained from feedback gained in our post-study questionnaire.

Features users liked

1. Layout was easy to understand and made scientific sense to users. 2. It was easy to follow edges. 3. Searching for plant lines was simple. 4. Bringing together additional data sources was extremely helpful.

Features users disliked or found confusing

1. Sometimes difficult to differentiate colour coding. 2. Long edges are disorientating. 3. No auto-selection of lines when performing a search. 4. Clearer explanations of ordinal data categories.

Table 1 Interaction study correct answers

Question classification	Correct (%)	Incorrect (%)
Unexplained concepts	50	50
Simple grandparent tracking	93.75	6.25
Identifying children	56.25	43.75
Complex grandparent tracking	50	50
Phenotype classes	100	0
Great-grandparent tracking	37.50	62.50
Finding additional information	93.75	6.25
Colour coding perception	56.25	43.75

Our test users liked the speed at which they could find data, the ease of tracing lineage through complex graphs (although our testing has shown that there were issues with this) and the intuitive layout of our visualization and supporting application. Our testing did highlight some issues, mainly around the use of colour gradients used in ordinal lists which are ineffective and difficult for our users to distinguish when there are more than eight phenotype classes.

Development of Helium

Feedback from the user evaluation allowed us to address issues that our users had with our prototype in order to develop a more refined and useful visualization application. We needed to work to increase understanding of concepts, representations and visual metaphors that our users found difficult to understand during testing.

The main feedback gained from our initial prototype was that it was difficult to track lineage with overlapping edges and that the ability to interactively overlay, query and retrieve various data types from our internal barley database would be important. Our users also had problems with identifying phenotype classes. Other issues were with the complexity of the graphs and problems identifying children.

Any subsequent development would need to address these points if it was going to offer a usable and effective tool for users.

The interface was re-designed to show 4 main areas: a) the overview panel and data selection panel, b) the main pedigree visualization panel, c) the local view panel and finally d) the details panel. These are described below.

Overview and data selection panel

This panel (Figure 6A) also includes selection mechanisms for choosing ordinal and nominal categorical phenotypic classes as well as tools for visualizing genetic similarity data (Figure 7). Users can use the overview to navigate to a particular region within the main visualization window if required.

Interactive sliders allow users, in the case of similarity data, to set a percentage similarity value and in real time highlight lines which match the search criteria (Figure 7). In this way it is possible to see lines which should not be closely related appearing on the peripheries of our visualization as the slider is moved, which may indicate problems with pedigree definition or genotyping. Histograms have also been included, where appropriate, to show data distribution which can be an aid in the identification of problem markers. While the number of markers that have this problem is limited, it is nonetheless important to address.

Other features included in this panel are the ability to select more than one phenotype then recolour nodes

based on the merged phenotype classes. While originally it had been intended to show each phenotype as a different section on a node it was decided, through speaking to users, that they would be interested in finding exact combinations and so it was decided to go with the single node colour to reduce clutter and keep the visualization clearer. There are however problems as the number of colours that may have to be used can be around 20. Such a high number has been shown to be ineffectual at differentiating between classes [40,44,45].

Main visualization panel

The main visualization window (Figure 6B) was modified in a number of ways from our prototype. Firstly, we have moved away from bundled orthogonal edge routing (Figure 5) which will make the tracing of lineage easier. Slightly modified colour palettes were used to account for the situation where there are more than eight categorical classes. The new colour palette will help with the problem our testing showed where adjacent classes were too similar in colour for users to accurately distinguish. In Table 1 the incorrect responses to "Identifying Children" were high at 43.75%. In order to address this visual prompts when hovering over a node were added which display the number of ingoing and outgoing edges from a node and the names of the line's progeny (Figure 6B). This makes the number of progeny immediately obvious, which will help prevent some of the problems seen in testing. When a user selects a node the edges connecting nodes of interest are made more prominent by both removing edges which are not associated with the selected node, its ancestors, or successor, and by darkening the edges which are left. Hovering over a graph edge will show the names of the two nodes that it connects, in this way with long edges, while using the main visualization window, it is easier to track their origin and destination.

Local view panel

Our testing also showed that while users reported they found it easy to identify lineage there were some issues. These problems could be addressed by including a "local" implementation of our graph showing only the line of interest and its lineage (Figure 6C). This would be shown when a user selects a node in our visualization. This view was implemented below the main visualization window. The local view can be panned and zoomed in the same way as the main visualization window. Within the local view the user has control of how many generations, forwards and backwards, they want to go. This addresses the problems highlighted in Table 1 where there were 50% and 62.5% of users incorrectly answering the "Complex Grandparent Tracking" and "Great-Grandparent Tracking" questions respectively. With appropriate selection of generation level, grandparents, or indeed any other generation,

Figure 7 Helium showing genetic similarity data. Genetic similarity data is stored in our Germinate database (all-by-all pairwise comparisons) which is displayed to the user by selecting a base node then showing similarity of lines in relation to this node. When the user selects another node the new data is retrieved and displayed. **A)** A slider allows users to select a cut-off or the similarity values (45-100%) and the results are shown in a sortable table below. The histogram shows the data distribution for the selected line which is one of many indications on the quality of the data. Selecting a row from the table jumps to that line, then updates the visualization accordingly. **B)** Coloured nodes to show similarity to the base line, and node sizing to show the number of times a line has been used in subsequent crosses (the larger the node the greater the number times it has been used as a parent). **C)** When a line is selected from the main display only edges joining the selected edge with predecessors or ancestors are shown which reduces display clutter.

are now immediately obvious in the simplified pedigree. Additionally, the ability to layout the graph using a number of edge routing algorithms was added. Any changes made to the main pedigree visualization are propagated to the local view. While the local view includes another copy of a portion of the main visualization, it will increase the accuracy of tracing lineage when unnecessary lines are removed and edges between nodes shortened, thus addressing the problems highlighted in testing and reducing the need to “chase edges”.

Detail panel

The details panel (Figure 6D) shows information about either the current selected phenotype(s) or information from Germinate about specific selected plant lines. This example shows the distribution of the DUS character “Anthocyanin Colour”. The histogram has been coloured in the same way as the phenotype classes in the main visualization window.

The details panel also houses a search functionality which allows searching for lines with usual search features such as wild-card matching and an option which we have called the “follow me” mode which jumps to a search hit, selects it and subsequently updates the detail panel and main visualization window.

During discussions with users it was also apparent that the ability to export line names would be a useful feature to allow scientists to make up lists for sending samples off for genotyping based on phenotypic or genotypic characteristics so the ability to allow users to export lists has been implemented. Users can select nodes then add them to an export list which can be saved to a text file.

Finally, a user history panel has been included which records the lines and phenotypes that have been selected over a session so that if required, users can go back and see what they had been doing previously. This is important as with large quantities of data it is easy for users to forget what they have been doing over time.

Examples of the layout and features offered by Helium can be seen in Additional file 1.

Discussion

An interesting outcome of the development of Helium is trying to quantify if this tool actually make a user's decision making better and does the software influence users into making more informed decisions about their data. One of the outcomes from our testing was to assure ourselves that the decisions that had been made around the design of the tool were actually good foundations that our target users can build knowledge on and to that end we seem to have made an impact. While we have used standard approaches to the visualization tool we have developed we have applied it directly to a specific domain, and tailored our application appropriately.

While users requested as much information as possible in the interface we need to be careful that we only include necessary information and do not turn Helium into a tool that presents so much unnecessary information to users it in itself becomes unusable or difficult to comprehend; we need to avoid a situation where we overload users with information. While this may seem like a problem that scientists would love to have it could have detrimental effects; do we need to actually present raw data or are overviews enough? Would a user's understanding be affected by what we present them with?

Users have told us that the overlaying of data onto the pedigree structure has in some ways more impact than showing the division of data in a bar chart or as a table. Having areas of colour in your face brings insight both into the location of clusters of similar data and visual impact of nodes changing from one colour to another, it brings the representation of data to life and in logical an understandable ways.

Examples of the sorts of things that users wanted to be able to do with our tool include a) given genotype data for a line identify possible matches and b) basic error checking based on genotypic or phenotypic data. These are detailed below.

Given genotype data for a line identify possible matches

Helium will take a string of genotypic data and identify possible matches from data held in our Germinate database then display the possible hits on the pedigree display. This is useful as it is not uncommon for errors to be introduced through mislabelling or handling errors in the lab when genetic material is sent for genotyping. Using the pedigree framework may give users other ways of trying to identify what unknown or problem lines are, or they may point geneticists and breeders in the right direction as to their source, if for example two similar lines are mislabelled we may be able to deduce the correct naming through examination of pedigree records. Further investigation would be required to correctly identify the correct

Figure 8 Pedigree visualization static prototype. This was one of our first attempts at visualizing our entire barley pedigree. The colours of nodes were used to distinguish between the winter/spring ecotype (red shows spring barley, blue shows winter barley and the cream coloured nodes are lines that are in both winter and spring pedigrees - qualitative data type) and node size to show the number of times the line has been used in crosses that have given rise to progeny that have been successful in National List trialling in the UK - quantitative data type. To the best of our knowledge this is the first time that a pedigree involving this number of commercially released lines has been brought together in one place and sparked interest with commercial plant breeders when they were presented with it.

source of this germplasm as there is a possibility either it, or the genotyping is wrong. These types of error are not uncommon.

Basic error checking based on genotypic or phenotypic data

We can use the interface to look for potential errors with a given line. We know that the alleles of a line must be from either parent, so we can use this in basic error checking. For example, if two lines have been genotyped for allele A at given locus but the progeny has allele B then we know there is a problem. Additionally, we can expand this type of search to look at multiple loci within a dataset. Taking this a step further we can use genotypic data to highlight potential parents of a line and if one parent is known, make a guess at possible candidates for the second parent.

Conclusions

We have shown through the development of Helium that visualization of our example pedigrees along with genotypic and phenotypic data provides users with new insights into crop breeding.

The representation of our unique barley test dataset shows that the pedigree structure takes the form of what we have coined a pedigree net. Our visualization has shown that there are three main classes of plant lines seen when viewed in Helium which we have named; a) principal lines which are commonly used to generate new cultivars due to their possession of desirable characteristics b) flanking cultivars brought in to increase the genetic diversity of subsequent lines and less commonly used in crosses and finally c) terminal varieties that are released, but have had little subsequent use.

One of the more hard-hitting measures of success of our first paper-based prototype came from the presentation of data to a meeting of UK plant breeders. While the pedigree data that we demonstrated was available to all in the room as written records, (like those in 'Pedigree formats' subsection), the representation that we showed (Figure 8) had a major impact through the provision of new insights as to how germplasm was very closely related. When written as a text string it is difficult to construct the bigger picture, but when displayed in our tool, the relationships between competing breeders lines was much more striking. While this was privately known to the individual breeders, having it presented to them when they were all in the same room was very enlightening. This not only highlights the value of visualization but that we have implemented a visualization tool with real-world impact.

While Helium has been tailored to specific data types (genotypic/similarity, nominal and ordinal phenotypic data and pedigree definitions) it is intended to be a

framework on to which, over time, additional data types can be added and we are working with worldwide plant scientists and breeders to develop the Helium platform further.

For more information on Helium please visit our website <http://ics.hutton.ac.uk/helium>.

Additional file

Additional file 1: Helium features movie. This movie shows the main layout and features of the Helium system along with basic interface interaction.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

PDS wrote this paper, developed the pedigree visualizations and database infrastructure, and designed and conducted the user study. IM gave advice on programming. MG, JK and DFM contributed towards the writing and editing of the manuscript and provided constant advice over the course of this work. All authors read and approved the final manuscript.

Acknowledgements

The authors gratefully acknowledge funding from the Scottish Government's Rural and Environment Science and Analytical Services (RESAS) division and Edinburgh Napier University. We would also like to thank colleagues at The James Hutton Institute, in particular Bill Thomas and Luke Ramsay for help and advice with pedigree data. We would also like to thank colleagues from NIAB (National Institute of Agricultural Botany) and the AGOUEB (Association Genetics of UK Elite Barley) consortium for the use of experimental data. Additionally, we would like to thank those who were generous enough with their time and enthusiasm to participate in the user evaluation of this software tool.

Received: 21 April 2014 Accepted: 10 July 2014

Published: 1 August 2014

References

- Gregory PJ, Ingram JSI, Brklacich M: **Climate change and food security.** *Philos Trans R Soc Lond B Biol Sci* 2005, **360**(1463):2139–2148.
- Knox JW, Morris J, Hess TM: **Identifying future risks to UK agricultural crop production: putting climate change in context.** *Outlook Agric* 2010, **39**(4):249–256.
- Mahmuti M, West JS, Watts J, Gladders P, Fitt BDL: **Controlling crop disease contributes to both food security and climate change mitigation.** *Int J Agric Sustain* 2009, **7**(3):189–202.
- DEFRA Page on NL Testing Procedures. <http://www.fera.defra.gov.uk/plants/plantVarieties/nationalListing/>.
- Paterson T, Graham M, Kennedy J, Law A: **Evaluating the VIPER pedigree visualisation: Detecting inheritance inconsistencies in genotyped pedigrees.** In *2011 IEEE Symposium on Biological Data Visualization BioVis*. 2011:119–126.
- Fishbeck G: **Diversification through breeding.** In *Diversity in Barley - Hordeum vulgare*. 2003:29–52. Chapter 3.
- Bennett RL, Steinhaus KA, Uhrich SB, O'Sullivan CK, Resta RG, Lochner-Doyle D, Markel DS, Vincent V, Hamanishi J: **Recommendations for standardized human pedigree nomenclature.** Pedigree Standardization Task Force of the National Society of Genetic Counsellors. *Am J Hum Genet* 1995, **56**(3):745–752.
- Purdy LH, Loegering WQ, Konzak CE, Peterson CJ, Allan RE: **A proposed method for illustrating pedigrees of small grain varieties.** *Crop Sci* 1968, **8**(4):916.
- Lamacraft RR, Finlay KW: **A method for illustrating pedigrees of small grain varieties for computer processing.** *Euphytica* 1973, **22**(1):56–60.
- Plant Variety Testing SASA (Science & Advice for Scottish Agriculture)*. <http://www.sasa.gov.uk/plant-variety-testing>.
- The Germinate Platform*. <http://ics.hutton.ac.uk/germinate/>.

12. Voorrips RE, Bink MCAM, van de Weg WE: **Pedimap: software for the visualization of genetic and phenotypic data in pedigrees.** *J Hered* 2012, **103**(6):903–907.
13. van Berloo R, Hutten RCB: **Peditree: pedigree database analysis and visualization for breeding and science.** *J Hered* 2005, **96**(4):465–468.
14. Milne I, Shaw P, Stephen G, Bayer M, Cardle L, Thomas WTB, Flavell AJ, Marshall D: **Flapjack-graphical genotype visualization.** *Bioinformatics (Oxford, England)* 2010, **26**(24):3133–3134.
15. Wigginton JE: **PEDSTATS: descriptive statistics, graphics and quality assessment for gene mapping data.** *Bioinformatics* 2005, **21**(16):3445–3447.
16. Bezerianos P, Dragicevic JD, Fekete J, Bae B, Watson B: **GeneaQuilts: a system for exploring large genealogies.** *IEEE Trans Vis Comput Graph* 2010, **16**(6):1073–1081.
17. Paterson T, Graham M, Kennedy J, Law A: **VIPER: a visualisation tool for exploring inheritance inconsistencies in genotyped pedigrees.** *BMC Bioinformatics* 2012, **13**(Suppl 8):5.
18. Ghoniem M, Fekete JD, Castagliola P: **A Comparison of the Readability of Graphs Using Node-Link and Matrix-Based Representations.** In *IEEE Symposium on Information Visualization*. 2004:17–24.
19. Purchase HC, Cohen RF, James M: **Validating Graph Drawing Aesthetics.** In *Graph Drawing, Lecture Notes in Computer Science, Volume 1027*. Edited by Brandenburg Franz J. Berlin Heidelberg: Springer; 1996:435–446. doi:10.1007/BFb0021827.
20. Purchase H, Carrington D, Alder J: **Empirical Evaluation of Aesthetics-based Graph Layout.** *Empirical Software Engineering*. 2002:233–255.
21. Bennett C, Ryall J, Spalteholz L, Gooch A: **The Aesthetics of Graph Visualization.** *Computational Aesthetics in Graphics, Visualization and Imaging*. 2007:1–8.
22. Trager EH, Khanna R, Marrs A, Siden L, Branham KEH, Swaroop A, Richards JE: **Madeline 2.0 PDE: a new program for local and web-based pedigree drawing.** *Bioinformatics* 2007, **23**(14):1854–1856.
23. Makinen V-P, Parkkonen M, Wessman M, Groop P-H, Kanninen T, Kaski K: **High-throughput pedigree drawing.** *Eur J Hum Genet: EJHG* 2005, **13**(8):987–989.
24. Thiele H, Nurnberg P: **HaploPainter: a tool for drawing pedigrees with complex haplotypes.** *Bioinformatics (Oxford, England)* 2005, **21**(8):1730–1732.
25. Wong L: **Visualization and manipulation of pedigree diagrams.** *Genome Inform Workshop Genome Inform* 2000, **11**:63–72.
26. Cole JB: **PyPedal: a computer program for pedigree analysis.** *Comput Electron Agric* 2007, **57**(1):107–113.
27. Bostock M, Heer J: **Protovis: a graphical toolkit for visualization.** *IEEE Trans Vis Comput Graph* 2009, **15**(6):1121–1128.
28. Draper GM: **Interactive Fan Charts: A Space-saving Technique for Genealogical Graph Exploration.** In *Proceedings of the 8th Annual Workshop on Technology for Family History and Genealogical Research (FHTW 2008)*. 2008.
29. Tuttle C, Nonato LG, Silva CT: **PedVis: a structured, space-efficient technique for pedigree visualization.** *IEEE Trans Vis Comput Graph* 2010, **16**(6):1063–1072.
30. Wernert EA, Lakshminpathy J: **PViN - A Scalable and Flexible System for Visualizing Pedigree Databases.** In *Proceedings of the 2005 ACM symposium on Applied computing*. ; 2005:115–122.
31. Loh AM, Wiltshire S, Emery J, Carter KW, Palmer LJ: **Celestial3D: a novel method for 3D visualization of familial data.** *Bioinformatics* 2008, **24**(9):1210–1211.
32. **Walrus - Graph Visualization Tool.** http://www.caidda.org/tools/visualization/walrus/.
33. Zhao JH: **Pedigree-drawing with R and graphviz.** *Bioinformatics (Oxford, England)* 2006, **22**(8):1013–1014.
34. Dieberger A, Frank AU: **A city metaphor to support navigation in complex information spaces.** *J Vis Lang Comput* 1998, **9**(5):597–622.
35. Muller Michael J, Kuchinskaya O, Minassian Suzanne O, Tang John C, Danis C, Zhao C, Harrison B, Moran Thomas P: **Shared Landmarks in Complex Coordination Environments.** In *CHI '05 Extended Abstracts on Human Factors in Computing Systems, CHI EA '05*. New York, NY, USA: ACM; 2005:1681–1684. doi:10.1145/1056808.1056996.
36. Cockburn A, Karlson A, Bederson BB: **A review of overview + detail, zooming, and focus + context interfaces.** *ACM Comput Surv* 2009, **41**(1):2–1231.
37. **yWorks - The Diagramming Company.** http://www.yworks.com/en/index.html.
38. Brewer CA, Hatchard GW, Harrower MA: **ColorBrewer in print: a catalog of color schemes for maps.** *Cartography Geogr Inform Soc* 2003, **30**(1):5–32.
39. Harrower M, Brewer CA: **ColorBrewer.org: an online tool for selecting colour schemes for maps.** *Cartogr J* 2003, **40**(1):27–37.
40. Ardi C, Tan S: **Color Palette Generation for Nominal Encodings.** http://vis.berkeley.edu/courses/cs294-10-fa08/wiki/images/6/66/FP-STCAKYpaper.pdf.
41. Sedlmair M, Meyer M, Munzner T: **Design study methodology: reflections from the trenches and the stacks.** *IEEE Trans Vis Comput Graph* 2012, **18**(12):2431–2440.
42. Munzner T: **A nested model for visualization design and validation.** *IEEE Trans Vis Comput Graph* 2009, **15**(6):921–928.
43. Lam H, Bertini E, Isenberg P, Plaisant C, Carpendale S: **Empirical studies in information visualization: seven scenarios.** *IEEE Trans Vis Comput Graph* 2011, **18**(9):1520–1536.
44. Ware K: *Information Visualization: Perception for Design*. ; 2004:486.
45. Kelly KL: **Twenty-two colors of maximum contrast.** *Color Engineering* 1976, **3**:26–27.

doi:10.1186/1471-2105-15-259

Cite this article as: Shaw et al.: Helium: visualization of large scale plant pedigrees. *BMC Bioinformatics* 2014 **15**:259.

Submit your next manuscript to BioMed Central and take full advantage of:

- **Convenient online submission**
- **Thorough peer review**
- **No space constraints or color figure charges**
- **Immediate publication on acceptance**
- **Inclusion in PubMed, CAS, Scopus and Google Scholar**
- **Research which is freely available for redistribution**

Submit your manuscript at
www.biomedcentral.com/submit