Recommendations Based on the Florida Citrus Rootstock Selection Guide Using a Web-server Application of Artificial Intelligence

Steven Rogers*[1], William S. Castle[2], Stephen H. Futch[2], and Andrew S. Persaud[2]

[1]*Ecostat, Inc., P.O. Box 237, Highland City, FL, 33846*

[2]*University of Florida/IFAS, Citrus Research and Education Center, 700 Experiment Station Road, Lake Alfred, FL 33850*

Additional index words. artificial intelligence, software as a service, SAAS, DataTables, Metronic, Apache Tomcat, Apple, Mac OS X, Twitter Bootstrap, expert system, decision support, Exsys Corvid Core

The Third Edition of the Florida Citrus Rootstock Selection Guide was recently released online. This website is a valuable resource containing updated information on rootstock options. The site is unique in that several technologies there complement the rootstock table, including: 1) an interactive online version of the table; 2) an extensive bibliography containing over 100 rootstock references; and 3) an expert system to help focus on the best rootstock candidates given certain user-selected criteria. The expert system is a backward-chaining platform that interviews users about their planting and site requirements. The system uses artificial intelligence technology to infer the best candidate rootstocks based on those criteria. The results are presented in an ordered list from top to bottom showing rootstocks that might be considered. The expert system is built on the Exsys Corvid® Core for Mac OS® X platform, which has the advantages of providing robust development features at a reasonable cost. This paper presents the expert system, provides details on the development process, and discusses the results of a focus group presentation for real-world user feedback.

One key decision in new citrus planting and replanting situations is deciding which rootstock(s) are good choices for a grower's situation. New technologies can help manage some of the complexity and uncertainty in those decisions. Growers should arrive at their rootstock decisions by focusing on their specific site, and then take into consideration the performance of their particular scion-rootstock combination at that site (Castle and Futch, 2015).

Several rootstock selection guides have been published by the University of Florida, beginning in 1989 with "Rootstocks for Florida Citrus" (Castle et al., 1989), which was revised in 1993 (Castle et al., 1993). In 1998, the "rootstock wheel" was introduced to the industry and consisted of 19 rootstocks of which 11 were considered "commercial" and 8 were "new" to the industry (Castle and Tucker, 1998). This new wheel format contained information from the previous two publications, but was condensed into a single-page, wheel format highlighting 17 rootstock characteristics or traits. Five new rootstocks from the U.S. Department of Agriculture program were added in 2006 (Castle et al., 2006). At this point, the wheel format contained 24 rootstocks, and 19 characteristics or traits ranging from tolerances, horticultural traits, and pest and disease considerations.

The most recent Third Edition of the Florida Citrus Rootstock Selection Guide, published in 2016, is a comprehensive resource enabling growers to research and select rootstocks specific to their site based on factors the grower considers most appropriate (Castle et al., 2016). There is also is an interactive online version of this same guide. Growers can find this website on the University of Florida, IFAS Citrus Research and Education Center website <flrootstockselectionguide.org> (to date, the website has had more than 290,000 visits). The site hosts the citrus rootstock selection guide, custom query options, and an extensive collection of articles focusing on rootstocks in Florida. The system is web-based and free to use anywhere in the world with an internet connection. This interactive online version of the guide has three new features to easily find relevant rootstock information. The first feature is a printable online edition of the guide itself. The second feature, found by clicking on the "Query the Rootstock Selection Guide" link at <flrootstockselectionguide.org>, is a rootstock expert system. Developed by Drs. Steven Rogers, Bill Castle, Steve Futch, and Mr. Andrew Persaud, this system uses artificial intelligence to help growers identify the best rootstocks. (In this paper, we use the term, "artificial intelligence," to refer to knowledge automation and a computer's ability to learn about and solve problems similarly to a human expert.) The third feature

[3]Corresponding author; email: steve@rsteve.com; phone: 863- 646-5187

of the online rootstock guide is a collection of over 100, virtually all Florida-based, downloadable references related to each rootstock and trait.

The expert system complements the rootstock table. Its essence is a series of questions asked in the form of a short interview to identify a user's specific set of circumstances. The system can reason through their answers to hone in on the best rootstock options. The expert system helps growers in five main ways:

- provides a measure of confidence that certain rootstocks are good matches for their site;

- easily test different rootstock planting scenarios;

- provides a "second opinion" regarding initial rootstock choices;

- provides rootstock suggestions growers may not have considered; and

- helps growers consider all rootstock options.

Our purpose in this paper is to describe certain technical aspects of how the online technology for the Florida Citrus Rootstock Guide and its complementary expert system were developed.

Materials and Methods

An expert system is a knowledge-based computer application that uses an inference engine to reason like a human expert. [An inference engine is a computer algorithm that processes a set of rules to solve a problem (Nikolopoulos, 1997)]. Expert systems assist non-experts with decisions when variables become numerous or potentially complex. The system we designed is a rule-based expert system (Exsys, Inc., 2016), in which an inference engine processes the set of rules to diagnose the grower's rootstock needs. The most recent 2016 edition of the rootstock guide contains 45 rootstocks and 20 horticultural traits, so it can be difficult for growers to compare all these factors across all rootstocks. Our system contains heuristic (common-sense) knowledge and thus makes these comparisons easier by using a conversational question-and-answer interface. The system was designed to complement and interact with the existing rootstock table and bibliographic collection. Using all three modules together provides growers a powerful suite of tools to research, understand and interpret their rootstock situations.

Website Architecture

The online version of the Third Edition of the Florida Citrus Rootstock Selection Guide is comprised of five core technologies: 1) Metronic responsive admin theme; 2) DataTables plugin; 3) Zoom Search Engine®; 4) Apache Tomcat®; and 5) Exsys Corvid® Core for Mac OS® X. DataTables and Apache Tomcat are open-source, while Metronic, Zoom, and Exsys Corvid Core are reasonably-priced commercial software applications. Integrating these applications together provided robust programming capabilities, accelerated development time, and helped contain costs.

METRONIC HTML 5 THEME. Metronic is a responsive, multi-purpose admin-powered website theme using the Twitter Bootstrap 3.3.5 & AngularJS 1.3 frameworks (Metronic, 2016). Metronic offers an interface for bridging HTML and DataTables to make it more user friendly. Due to its included plugins and components, Metronic reduced our web framework development time and costs.

DATATABLES. DataTables, a plug-in for the jQuery Javascript library (DataTables, 2016), provided our framework for representing the rootstock table on the web. DataTables is a flexible formatting tool that adds enhanced, interactive controls to HTML tables. We use a subset of the DataTable control complement to provide navigation for the online version of the rootstock table. DataTables supports a number of data sources, including DOM (Document Object Model), Javascript, and server-side processing. In our application, however, rootstock data are hard-coded into the Javascript.

ZOOM SEARCH ENGINE FOR MAC. Zoom is software that creates a search engine for websites, intranets, CDs, or DVDs (Zoom Website Search Engine, 2016). Zoom provides full-text searching by indexing websites with a desktop application. The resulting index is uploaded to the website to provide robust search features. In our application, Zoom was used instead of Google Custom Search primarily to circumvent Google's 100-reference limit on search results.

APACHE TOMCAT. Apache Tomcat is an "open source implementation of the Java Servlet, JavaServer Pages, Java Expression Language, and Java WebSocket technologies" (Apache Tomcat, 2016). The advantages of using Tomcat are that it is: 1) lightweight, 2) open-source, 3) flexible, 4) stable, and 5) secure (FutureHosting, 2016). Tomcat provides the basic functionality needed to run the expert system without a lot of overhead. This code-light framework means the system should run fast in most user environments, including slow mobile connections. Web pages served from Tomcat can be customized using standard HTML5, making updates to the site interface easy to perform. The Tomcat server is hosted on a low-cost shared platform managed by HostGeneral (hostgeneral.com).

EXSYS CORVID CORE FOR MAC OS X. Exsys Corvid Core for Mac OS X is an expert system shell for automating expert knowledge solutions (Exsys, Inc., 2016). Exsys Corvid allows for: 1) capturing decision-making logic; 2) processing a domain expert's reasoning pathway; 3) wrapping the system in a user interface with a customized look-and-feel; and 4) integrating the expert system with other resources (Exsys Corvid Core for Mac OS X, 2016). Corvid provides several ways to describe an expert's logic using variables. It uses heuristic IF...THEN rules based on those variables to arrive at its conclusions. Variables have several properties, allowing them to be used in many ways. Corvid contains seven types of variables, including numeric and string variables to make it easy to build probabilistic and confidence-based systems (Exsys, Inc., 2007).

Expert System Architecture

The architecture of the expert system is grounded on grower and researcher experience (Fig. 1, adapted and redrawn from Luger and Chakrabarti, 2008). Generally, the expert system is comprised of three main parts: 1) a user interface; 2) a knowledge base; and 3) an inference engine (Exsys, Inc., 2016). The user interface is that part of the application that directly interacts with the user. It provides the interface for questions and answers, and displays the final report of rootstock rankings. The interface is built in HTML5, and served to the user through the web using Apache Tomcat. Second, the knowledge base is the collection of knowledge, facts, and rules. The rootstock table serves as the primary source of expert information programmed into the application's rules. Third, the inference engine takes the user's input and matches that to its knowledge base. The inference engine is the bridge between the rules that the computer processes and that users understand (Exsys, Inc., 2016).
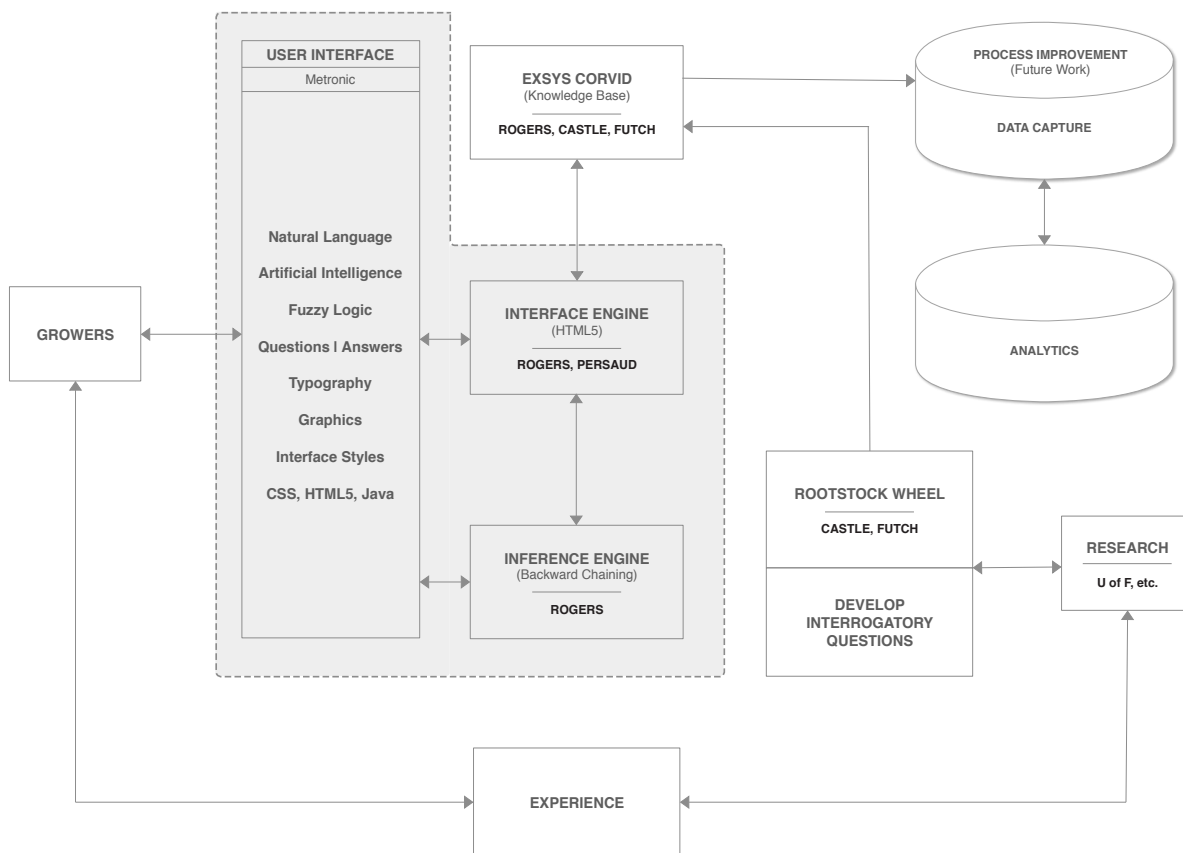
Fig. 1. Expert system architecture (redrawn and further adapted from Luger and Chakrabarti, 2008). The architecture of the expert system is grounded on grower and researcher experience. The expert system is comprised of three main parts: 1) a user interface; 2) a knowledge base; and 3) an inference engine.

Additional Web Development

Website HTML and CSS programming were accomplished on both the Mac OS X 10.10.5 (Apple, Inc., Cupertino, CA) and Microsoft® Windows® Version 7 (Microsoft, Inc., Redmond, WA) operating systems using Panic Coda Version 2.5.16 (Panic, Inc., Portland, OR) and Adobe® Dreamweaver® CS5 ver 11.0 (Adobe, Inc., San Jose, CA). Exsys Corvid Core development was carried out in Windows Version 7 in a Parallels Desktop Version 11.2.0 virtual instance running in Mac OS X Yosemite 10.10.5. Graphic and image manipulation was with Adobe Bridge CC (6.2.0.179), Adobe Photoshop® CS5 and CC (2015), Adobe Illustrator® CC (2015), and Adobe InDesign® CC (2015) (Adobe, Inc., San Jose, CA). Cross browser testing was performed using Browsershots (browsershots.org) and the iOS Version 9.3.2 operating systems on the iPhone® 6 Plus and iPad® 2 mobile devices. Adobe Acrobat® PDF renderings were created using Adobe InDesign (2015) and Adobe Acrobat Pro DC Version 2015 (Adobe, Inc., San Jose, CA). Web-page load-speed metrics were tested using PageSpeed Insights (Google, Inc., Mountainview, CA) and Pingdom Website Speed Test <tools.pingdom.com>. The HTML5 payload size was optimized using HTML Compressor <htmlcompressor.com>. Rootstock rank calculations were developed and tested in CalcSpreadsheet Version 5.1.3.2 or earlier (LibreOffice Document Foundation, <libreoffice.org> and Panorama (ProVUE Development, Inc., Huntington Beach, CA). Statistical calculations were performed in Stata® Version 14.1 (StataCorp, College Station, TX).

Results

Users who navigate to the expert system's landing page at flrootstockselectionguide.org are presented with an introductory screen informing them about the system and its capabilities. The system then begins by asking growers if they want to query the system based on horticultural traits, site condition, or pests and diseases. Users may select one, two, or all three options. Depending on their selections, the system asks questions to gather information about the grower's specific situation. The line of questioning is determined on a case-by-case basis depending on how the interview develops. For example, a grower concerned with clay soils could be presented with a different set of questions than a grower more concerned with *Phytophthora-Diaprepes* problems. Currently, the questions are basic, but we anticipate improving them after receiving feedback on system performance from users.

While users are entering their information through interactive questioning, the system uses its rules to match and score their input with the 900 possible factors in the guide. The result is that the user is presented a report with a ranked list of rootstocks that ON AVERAGE are top candidates for their situation. The rankings are based on a "Selection Score." This score is not a statistical value. Instead, it refers to a relative ranking of the expert system's confidence that its suggested rootstocks are suitable for the site situation described to it during the interview process. The closer in value the selection score for a rootstock is to 100, the closer the match of that rootstock is for the grower's described situation.

The current version of the expert system does not permanently store or archive information provided by users, but they can print a copy of the report to keep handy. Any information retained is maintained during the session only. Future versions of the system could provide more advanced capabilities for storing and sharing information.

Determination of Selection Score

The "bottom line" of each session using the expert system, i.e., the desired outcome, is the list of rootstocks and their associated selection scores. To achieve that outcome, each rootstock is first assigned a confidence value for each trait ranging from –1.0 (definitely false) to +1.0 (definitely true). This is similar to a system used in the Mycin bacterial infection expert system developed at Stanford University in the 1970s (Shortliffe and Buchanan, 1975; Buchanan and Shortliffe, 1984). Such assignments become a "rule". The numeric values assigned (–1 to +1) generally indicate whether that choice is "bad (–1)" or "good (+1)" based on the logic in the rules. The backend of the system is comprised of hundreds of rules relating confidence to site and horticultural traits. These rules can be combined and layered on one another to arrive at conclusions relating to multiple traits.

After entering the expert system, and at the beginning of the interview, the user selects which trait or traits are most important in their situation. As the interview progresses, the expert system considers each of the 45 possible rootstock choices and assigns confidence values associated with each user-selected trait to each rootstock. For example, Volkamer lemon would be assigned a low confidence value if a user needed a rootstock producing small trees.

These confidence values are combined using an additive model based on all the rules that match during the interview. The result of this process is the selection score for the rootstock. The selection scores generally range between –100 and +100. The rules leading to the selection scores can easily be updated based on new information and expert opinion. This makes the system adaptable to new situations.

Before the interview is concluded, there may be qualifier questions that help refine results for specific users. As a result, the final selection score is adjusted up or down as needed based on answers to the qualifier questions. The confidence value of each trait can be thought of as "points" added to or taken from the selection score of each rootstock as the interview progresses (Exsys, Inc., 2007). This helps determine if a rootstock is a good match and where it fits in the hierarchy. Before the resulting selection scores are presented, they are normalized so the maximum value is 100. This makes it easy to compare results within and between separate runs of the system.

Growers should always rely on their own judgement to make final rootstock decisions. It is very important to note that this expert system is a decision aid and is thus just one of several references growers should consult in their planting decisions. Other tools include publications, other growers, and experience. Nevertheless, an expert system such as this one can provide greater insight into making planting decisions.

Focus Group Tests

We have conducted several formal and informal focus group tests to demonstrate the system. The main goals of these focus group tests were to: 1) get feedback regarding ease of use of the user interface, and 2) determine if attendees generally agreed or disagreed with the ranking of rootstocks presented in their reports. So far, the feedback has been positive. We have also received several constructive comments how the system might be improved. Overall, however, most users and focus group attendees have generally agreed with the rankings provided to them in the reports generated by the system. We have not yet encountered any strong disagreement with the report results.

Discussion

Why did we take this particular web application approach to the matter of rootstock selection? Here are the primary reasons. This expert system vastly improves the rootstock selection process primarily by reducing the challenges of sorting through a large amount of information, and by offering customized options. Yet, it still only scratches the surface of what is possible with this kind of technology. As new knowledge and feedback are obtained, we expect to improve the accuracy, precision, and reliability in the rootstock rankings. This can be accomplished in several ways. The results of field knowledge and experience can go toward improving the confidence values assigned in the rules used by the system. This will depend on actively engaging users in providing feedback. Secondly, the platform itself can be improved by including features users find useful, such as saving and sharing reports, automatically determining soil types from GPS location, etc. Nevertheless, the current version of the system has a number of advantages outlined below (see also Durkin, 1990).

WEB APPLICATION. This system was designed to be as accessible to as many growers as possible. A web application was the best solution. A web application as used in this expert system has benefits over a locally-hosted (desktop) application, including that it does not require growers to download, install, or maintain software. Our system is platform-independent, so we do not have to enforce version checks in client machines. Thus, growers will always be using the latest version. The main disadvantage of a web application is that it requires an Internet connection, but ours is a code-light application that is not affected much by variations in connection speed.

RAPID ADAPTATION AND DEPLOYMENT. The expert system can be rapidly adapted to include new knowledge and innovation. For example, if new research or grower consensus is obtained that a certain rootstock shows better tolerance to HLB than all others, it is relatively straightforward to modify relevant rules in the system to include this new information. There is no executable compiling or other packaging of the application needed to generate web-ready working software. So, updating the application is as simple as modifying rules and uploading updated files to the server. The total time to modify a small set of rules to having the new application online can be less than 30 minutes.

NON-ALGORITHMIC PROGRAMMING. Many computer software applications are algorithmic in that they are executed in step-by-step fashion. Frequently, changing one section of code in such software means a cascading series of additional checks or coding need to be performed to ensure proper software operation. In contrast, Exsys Corvid Core as used here allows for inserting, changing, or deleting rules in non-linear fashion. Performing such changes does not necessarily touch or affect other rules within the system. Plus, the inference engine will continue to process any new or changed rules as required by its internal logic or based on information provided by the user to arrive at appropriate conclusions.

Adaptive. The system learns dynamically from information gained from its interview with the user and adapts its questions accordingly. For example, different users will interact with the system differently depending on their particular situation. Further, different users are presented with varying qualifier questions the system uses to better understand a given situation. For example, a user asking about a rootstock for a high-clay-containing soil will be asked a different set of questions than will a user with a *Phytophthora-Diaprepes* problem.

Unbiased. The expert system is unbiased in that it does not have preconceptions about the suitability of any of the rootstocks for particular situations. It evaluates each situation objectively based on information provided by the user. The system derives its recommendations based on the Florida Citrus Rootstock Selection Guide, and additional expert and grower information as required for the situation.

Weighted coding. Certain rootstocks may have overriding criteria that should be weighted into any final recommendation. For example, 'sour orange' rootstock is one of the more tolerant rootstocks to HLB, but it may not be suitable for many situations because of its susceptibility to Citrus Tristeza Virus (Stover and Castle, 2002). Users have the ability to activate or deactivate such weighted considerations for the list of recommendations based on their custom needs.

Similar to human. The system interviews the user in a way that is similar to that of a human consultant. It asks questions to get information about the specific needs of the user. It temporarily stores that information while matching it to rules defined in its programming. This essentially follows the same problem-solving logic as a human expert to arrive at its ranked list of rootstock recommendations. Given the same information, it is more likely than not that a human expert would arrive at a ranking generally similar to the expert system.

Backward chaining. The knowledge in the system is written in the form of numerous IF...THEN rules. If the expert system needs information for what is currently doing, it will check the rules to see if there are any that could inform it as to that value (Exsys, Inc., 2007). If it finds such a rule, the system suspends what it is doing and will evaluate that rule; then it will return to what it was originally doing. This happens recursively so that the system continuously collects information it needs to accomplish its task. If a new rule requires a value that can be obtained from other rules, those would be tested too. All possible conclusions are iteratively reviewed until a goal that can be supported by the premises is encountered.

Future plans. Our plans for the immediate future of this system include providing demonstrations to industry groups. This will allow us to conduct de facto focus group tests that can provide guidance on improving the user interface. As user feedback is obtained, we expect to improve the design and content of the printed reports. Several industry articles have also been produced, with more planned for the future. Further, we anticipate being able to obtain higher confidence in the rootstock selection scores and rankings based on general agreement or disagreement with the ranked reports. Improvements in application features are also anticipated. One example would be adding the ability to save or share links of ranked reports with other users. Many other features may involve improvements in the backend, so they will not necessarily be apparent in the user interface per se. Finally, publicizing the expert system along with the Florida Rootstock Selection Guide will be a component of keeping the system in front of its potential users. Another goal will be ultimately to train other people in how to code in the Exsys Corvid Core platform as a way to ensure continuity of operation. We hope the above activities, along with their strategic coordination, will continue to improve the value of the expert system to the industry.

Literature Cited

Apache Tomcat®—Welcome! 2016. Apache Software Foundation Accessed 4 June 2016. <http://tomcat.apache.org/>.

Buchanan, B. G. and E. H. Shortliffe. 1984. Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project. Addison-Wesley, Reading, MA.

Castle, W.S. and D.P.H. Tucker. 1998. Florida citrus rootstock selection guide. Extension Publ. SP-248. Florida Coop. Ext. Serv., University of Florida, Gainesville.

Castle, W.S., D.P.H. Tucker, A.H. Krezdorn, and C.O. Youtsey. 1989. Rootstocks for Florida citrus. Extension Publ. SP-42. Florida Coop. Ext. Ser., University of Florida, Gainesville.

Castle, W.S., D.P.H. Tucker, A.H. Krezdorn, and C.O. Youtsey. 1993. Rootstocks for Florida. 2nd ed. Extension Publ. SP-42. Florida Cooperative Extension Service, University of Florida, Gainesville.

Castle, W.S. and S.H. Futch. 2015. Choose the right rootstock. Extension Publ. HS-932. Florida Coop. Ext. Ser., University of Florida, Gainesville. <http://edis.ifas.ufl.edu/>.

Castle, W.S., K.D. Bowman, J.H. Graham, and D.P.H. Tucker. 2006. Florida citrus rootstock selection guide, revised. Extension Publ. SP-248. Florida Cooperative Extension Service, University of Florida, Gainesville.

Castle, W.S., K.D. Bowman, J.W. Grosser, S.H. Futch, and J.H. Graham. 2016. Florida citrus rootstock selection guide, 3rd ed. Extension Publ. HS1260. Florida Coop. Ext. Serv., University of Florida, Gainesville. < http://edis.ifas.ufl.edu/>.

DataTables—Table Plugin for jQuery. 2016. DataTables.net. Accessed 4 June 2016. <http://datatables.net/>.

Durkin, J. 1990. Research review: Application of expert systems in the sciences. Ohio J. of Sci., 90(5):171–179.

Exsys Corvid® Core for Mac OS® X. 2016. Exsys, Inc., Albuquerque, NM. Accessed 4 June 2016. <http://www.exsys.com/CorvidCore/overview.html>.

Exsys, Inc. 2007. Exsys Corvid Advanced Tutorial. Exsys, Inc., Albuquerque, NM.

Exsys, Inc. 2016. Exsys Corvid: The Industry Proven Software and Services for Knowledge Automation Expert Systems. Exsys, Inc. Albuquerque, NM.

FutureHosting. 2016. Five reasons you should use Tomcat. Accessed 4 June 2016. <https://www.futurehosting.com/blog/five-reasons-you-should-use-tomcat/>

Luger, G. and C. Chakrabarti. 2008. Knowledge-based probabilistic reasoning from expert systems to graphical models. Accessed 1 Mar. 2017. <https://www.researchgate.net/publication/228589548_Knowledge-Based_Probabilistic_Reasoning_from_Expert_Systems_to_Graphical_Models>.

Metronic–Premium Bootstrap Admin Themes. 2016. SpryMedia. Accessed 4 June 2016. <https://themeforest.net/item/metronic-responsive-admin-dashboard-template/4021469>.

Nikolopoulos, C. 1997. Expert systems: Introduction to first and second generation and hybrid knowledge based systems. Marcel Dekker, Inc. New York.

Shortliffe, E.H. and B.G. Buchanan. 1975. A model of inexact reasoning in medicine. Math. Biosci. 23:351–379.

Stover, E. and W. Castle. 2002. Citrus rootstock usage, characteristics, and selection in the Florida Indian River Region. HortTechnology 12(1):143–147.

Zoom Website Search Engine. 2016. Surry Hills, NSW, Australia: WrenSoft, Division of PassMark Software Pty Ltd. Accessed 4 June 2016. <https://www.wrensoft.com/zoom/>.