UCLA

UCLA Previously Published Works

Title

IDOCS: intelligent distributed ontology consensus system--the use of machine learning in retinal drusen phenotyping.

Permalink

<https://escholarship.org/uc/item/78f766tb>

Journal

Investigative ophthalmology & visual science, 48(5)

ISSN

0146-0404

Authors

Thomas, George
Grassi, Michael A
Lee, John R
et al.

Publication Date

2007-05-01

DOI

10.1167/iovs.06-1022

Peer reviewed

IDOCs: Intelligent Distributed Ontology Consensus System—The Use of Machine Learning in Retinal Drusen Phenotyping

George Thomas,^{1,2} Michael A. Grassi,^{2,3} John R. Lee,⁴ Albert O. Edwards,⁵ Michael B. Gorin,⁶ Ronald Klein,⁷ Thomas L. Casavant,^{8,9,10} Todd E. Scheetz,^{8,9,10} Edwin M. Stone,^{9,11} and Andrew B. Williams¹²

PURPOSE. To use the power of knowledge acquisition and machine learning in the development of a collaborative computer classification system based on the features of age-related macular degeneration (AMD).

METHODS. A vocabulary was acquired from four AMD experts who examined 100 ophthalmoscopic images. The vocabulary was analyzed, hierarchically structured, and incorporated into a collaborative computer classification system called IDOCs. Using this system, three of the experts examined images from a second set of digital images compiled from more than 1000 patients with AMD. Images were annotated, and features were identified and defined. Decision trees, a machine learning method, were trained on the data collected and used to extract patterns. Interrelationships between the data from the different clinicians were investigated.

RESULTS. Six drusen classes in the structured vocabulary were largely sufficient to describe all the identified features. The decision trees classified the data with 76.86% to 88.5% accuracy and distilled patterns in the form of hierarchical trees composed of 5 to 15 nodes. Experts were largely consistent in their characterization of soft, and to a lesser extent, hard drusen, but diverge in definition of other drusen. Size and

crystalline morphology were the main determinants of drusen type across all experts.

CONCLUSIONS. Machine learning is a powerful tool for the characterization of disease phenotypes. The creation of a defined feature set for AMD will facilitate the development of an IDOCs-based classification system. (*Invest Ophthalmol Vis Sci.* 2007;48:2278–2284) DOI:10.1167/iovs.06-1022

Age-related macular degeneration (AMD) is the most common cause of severe vision loss in the developed world, affecting more than 10 million people in the United States alone.¹ More than 7 million people in the United States have macular drusen of sufficient size and number that they are at substantial risk of severe visual loss.¹

Genetic predisposition plays a significant role in the pathogenesis of AMD.^{2–4} Recently, several genetic variants have been implicated as predisposing factors for this condition.^{5–12} Discovery, characterization, and eventual therapeutic control of the influence of these and other genes associated with the pathogenesis of AMD represent important goals of the vision research community. Many forms of non-AMD have already been molecularly characterized.^{13–16} All these entities have recognizable ophthalmoscopic appearances and distinct phenotypic characteristics. Similarly, we hypothesize that discrete genetic causes will be associated with differing prevalences of identifiable AMD subtypes.

Computational methods offer several advantages in collaborative ontology generation and complexity reduction through their ability to distill patterns in multidimensional data. They can be used to define a systematic process for capturing, organizing, and analyzing knowledge from distributed experts. This knowledge-acquisition-based approach for formalizing a clinical nomenclature may have greater consistency and reproducibility than do existing methodologies used to develop classification schemes. Moreover, machine learning techniques have already been successfully used to identify patterns in other types of complex data sets.^{17,18} The application of machine learning to analyze the ophthalmoscopic features of AMD may similarly yield insights into the critical defining attributes and features of this disease that would greatly facilitate the clinical training of ophthalmologists. This approach accelerates the creation of a classification system with improved discriminative power resulting in greater availability of more homogenous AMD patient populations.

The objective of this study was to use the power of knowledge acquisition and machine learning to facilitate the collaborative development of a novel AMD classification system. Herein, we discuss the initial steps used in retinal drusen analysis and present the prototype for IDOCs: the Intelligent Distributed Ontology Consensus System.¹⁹

From the ¹Department of Computer Science, the ⁸Center for Bioinformatics and Computational Biology, and the Departments of ⁹Ophthalmology and Visual Sciences and ¹⁰Biomedical Engineering, University of Iowa, Iowa City, Iowa; the ³Department of Ophthalmology, University of Chicago, Chicago, Illinois; ⁴Assistive Intelligence, Inc., Iowa City, Iowa; the ⁵Department of Ophthalmology, Mayo Clinic, Rochester, Minnesota; the ⁶Department of Ophthalmology, David Geffen School of Medicine—UCLA, Jules Stein Eye Institute, Los Angeles, California; the ⁷Department of Ophthalmology and Visual Sciences, University of Wisconsin, Madison, Wisconsin; the ¹¹Howard Hughes Medical Institute, Chevy Chase, Maryland; and the ¹²Department of Computer and Information Sciences, Spelman College, Atlanta, Georgia.

²Contributed equally to the work and therefore should be considered equivalent authors.

Supported by National Eye Institute Grant EY013688-03, EY014467, NIH/RIMI (Research Infrastructure in Minority Institutions) Grant 5 P20 MD000215 and the Heed Foundation. MAG was a Heed Fellow from 2003–2005. EMS is a Howard Hughes Medical Investigator.

Submitted for publication August 29, 2006; revised November 17, 2006; accepted March 13, 2007.

Disclosure: G. Thomas, None; M.A. Grassi, None; J.R. Lee, None; A.O. Edwards, None; M.B. Gorin, None; R. Klein, None; T.L. Casavant, None; T.E. Scheetz, None; E.M. Stone, None; A.B. Williams, None

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be marked “advertisement” in accordance with 18 U.S.C. §1734 solely to indicate this fact.

Corresponding author: Andrew B. Williams, Department of Computer and Information Sciences, Spelman College, 350 Spelman Lane SW, Campus Box 1257, Atlanta, GA 30314; williams@spelman.edu.

FIGURE 1. A 30° color fundus image centered on the macula. A box is used to annotate the desired feature that is described by the investigator as calcified drusen. *Right:* attribute values for the drusen feature.

METHODS

Patients

The recruitment and research protocols for human subjects were reviewed and approved by the University of Iowa's institutional review board and informed consent was obtained from all study participants, in accordance with the guidelines set forth in the Declaration of Helsinki. More than 1000 patients were ascertained from the University of Iowa's Department of Ophthalmology, examined by an ophthalmologist, and found to have signs consistent with the clinical diagnosis of AMD. All patients were over the age of 50 years (average age, 75.5 years). Only white individuals were enrolled in the study. All participants were ascertained during the same period by the same clinic.

Four retina specialists with clinical and research expertise in AMD who see patients at institutions that are physically distant from the University of Iowa were invited to participate in the study (RK, MBG, AOE, ACB). Collaborating retina specialists were involved in the study to generate a vocabulary that was broad, diverse, and encompassed the entire range of AMD terminology, so as not to be limited by either institutional or geographic mores.

Vocabulary Generation

A retina specialist at the University of Iowa (MAG) selected 100 representative images from the cohort. The images were displayed on a user-friendly interface for computer-based assessment. Using a digital dictating machine, each of the four collaborators described in detail the ophthalmoscopic appearance of each image.

The data from the dictation were transcribed and analyzed by MAG. Transcriptions were reviewed, and all key words and phrases were extracted, grouping the vocabulary into a hierarchical structure with attributes and associated values that were extensive and descriptive enough to cover the aggregate vocabulary of the clinicians. This structured vocabulary was distributed to the clinicians for their review and approval as representative of the spectrum of AMD vocabulary. The semantic definitions of attribute values were not specified.

Data Collection via User Interface

A second image set was created for analysis that consisted of representative digital images (one to seven per patient) from 100 different patients. In all, there were 232 images. An attempt was made to incorporate the broad spectrum of features found in AMD within this sample set. The database included 30° and 60° digital photographs of

the posterior pole, stereo macula photos, red-free images, and selected frames from the fluorescein angiogram (typically from an early frame from the arterial phase and a late frame from the venous recirculation phase).

An online web interface that incorporated the images and our structured vocabulary was designed. Features were identified by dragging and resizing a box over a specific image area, a process termed *annotating*. A feature name and associated attributes were provided by the expert. Drusen were classified according to a predefined list of terms. Users could also enter their own specific feature names. Features could be designated as exemplars, or classic stereotypic examples, of that drusen feature. No definition of attributes or other keywords was presented and the interpretation was left to the individual user. The drusen interface for a given patient is shown in Figure 1. All images could be enlarged to facilitate review. There was also an "index" allowing access to any patient (nonsequentially) as well as a help section that explained various features of the tool. When the clinicians finished entering a feature they would be taken to a "summary" page (Fig. 2), where all the various features and annotations were listed. From this page, they could continue to the next patient. Figure 3 shows the annotations from all three experts for a particular fundus image.

Using this interface, three of the collaborating retinal specialists (RK, MBG, AOE) described the ophthalmoscopic features of the patients over multiple sessions. Henceforth, the three retinal specialists are individually referred to anonymously as E1, E2, and E3 and collectively referred to simply as experts.

Ontology Analysis

Supervised machine learning algorithms work on a set of preclassified examples, called training sets, and create a model that can best capture the classification presented in these examples. New instances, whose classifications are unknown, are presented to the model which then predicts the class based on what was learned during training. In practice, a process known as 10-fold cross validation is used. The training set is divided into 10 folds or chunks; 1 fold is kept aside for testing (its classifications are ignored) and the model is trained on the other 9 folds and then tested on the 10th. After testing, the classes predicted by the model are compared with the actual classes of those instances as determined by the experts, and accuracy rates are generated for the model. This process is repeated 10 times so that every fold is used once as a test set and 9 times as part of a training set. Total accuracy rates are then averaged over the 10 test sets.

http://www.idocs.org - Intelligent Distributed Ontology Consensus System - Microsoft Internet Explorer

Click on the image to enlarge

7 of 100 INDEX HELP LOGOUT

DRUSEN RPE OTHER SUMMARY

Below is a summary of everything you have described for this patient.

You have described the following drusen features for this patient.

Name	Annotation	Morphology	Color	Size	Number	Pattern	Exemplar
Dystrophic		fine crystalline	pale	large	many	clustered	YES
Basal Lamellar/Cuticular		indistinct coalescent	pale	medium/intermediate	many	scattered homogenous	NO

Retinal Pigmented Epithelium - Atrophy

Annotation

Type geographic

Extent Involving Loss of Choriocapillaris Yes

Miscellaneous - Optic Disc

Peripapillary Atrophy Yes

Miscellaneous - Other

Photo Quality good

Comments

FIGURE 2. Summary page with three funduscopy images displayed.

Four different machine learning algorithms from the WEKA suite²⁰ (<http://www.cs.waikato.ac.nz/~ml/weka/> provided in the public domain by the University of Waikato, Hamilton, NZ): (1) J48 decision trees (C4.5, revision 8), (2) Bayes Net, (3) Lazy Weighted Learning, and (4) Ensemble Bagging were trained with 10-fold cross validation on the data of each expert. Based on overall accuracy of performance across all experts, the J48 decision tree, a specific type of machine learning in which patterns are distilled into rules, was selected as the algorithm for training the classifiers. From the decision tree results, rules capturing the ontologies were thus produced.

Ontology Reconciliation

The data for all experts were aggregated into one set, and classifiers were trained on this set by using J48 decision trees with 10-fold cross validation. To examine the relationships between the different ontologies, the J48 decision tree classifier trained on the individual data from each expert was used to classify the data from each of the other two experts. The predictions of the classifier were then compared with the assigned drusen names of the other two experts. Accuracy rates were generated from this step. Accuracy rates were also computed for testing each expert's data against the paired, aggregate classifier of the other two experts.

Data Sets

Experiments were performed on three data sets in all. In the first data set, DS1, the drusen data from the experts were used without modification. In the second data set, DS2, nominal attributes of size and number were mapped to a numeric format, and multivalued attributes of morphology and pattern were converted to multiple single-valued attributes each. Apparent synonyms among the attribute values such as "many" and "extensive" were merged and any omitted values were converted to missing data in WEKA format. In the third data set, DS3, all instances of reticular drusen and any other drusen that had only one instance in the data set were excluded. Consequently, DS3 was composed of drusen instances of only soft, hard, basal laminar/cuticular, dystrophic, and calcified.

RESULTS

Vocabulary

Figure 4 shows the structured drusen vocabulary deduced from the vocabulary-generation step. In addition to the drusen feature name, which can also be user defined, there were five attribute classes consisting of *number*, *color*, *morphology*, *size*, and *pattern*. Of these, morphology and pattern are multivalued attributes, whereas the remaining are single valued. Multivalued attributes are those that can have multiple values simultaneously.

Filename: 8R_1.jpg | Patient: 96

FIGURE 3. Annotations by three experts identifying crystalline morphology. Each expert's annotations are represented by a unique color.

FIGURE 4. AMD drusen vocabulary with attributes and their possible values. Morphology and pattern are attributes that can take multiple values simultaneously.

Figure 5 shows the statistical distribution of the various drusen types across the data of the four experts. The distribution indicates that a representative sample of the AMD drusen spectrum was captured with six main drusen types.

Soft were the most numerous type observed, with hard having a large number of examples for two of the experts. The remaining four drusen types had example instances ranging from 1 to 15. There were also five user-defined

FIGURE 5. Statistical distribution of drusen types across experts.

TABLE 1. Accuracy Rates of Classifiers Trained on Different Data Sets of Experts Using 10-fold Cross Validation

	DS1*				DS2†				DS3‡			
	<i>n</i>	%	κ	%RAE	<i>n</i>	%	κ	%RAE	<i>n</i>	%	κ	%RAE
E1	100	81.97	0.5817	48.50	103	84.43	0.6550	42.14	100	88.50	0.7221	36.06
E2	93	76.86	0.5807	51.57	94	77.67	0.6146	51.25	94	78.33	0.6263	50.16
E3	147	77.37	0.5784	58.34	146	76.84	0.5931	54.67	150	85.23	0.7158	44.42
All	325	75.06	0.5214	64.70	339	78.29	0.6135	54.47	339	82.89	0.6791	46.08

n, and % are the number and percentage of accurate classifications, respectively.

* Drusen data with no modification.

† Quantification of attributes; conversion of multivalued attributes to single-valued; missing data representation.

‡ Exclusion of reticular, flecked, granular, peau-d and reticulated.

drusen types by expert E3 with a total membership of seven instances.

The scope of our generated, structured IDOCS vocabulary was validated by the fact that only one expert needed more than the six predefined drusen names, creating five additional drusen names that were used in only 7 of 190 instances. The overall capture rate of the vocabulary hierarchy was 98.4% for all three experts.

Machine Learning

Table 1 shows the results of 10-fold cross-validation for each expert over the different data sets. For each expert and data set, the number of correct instances, percentage accuracy rate, κ statistic and relative absolute error (RAE) rate are given. The RAE is a normalization of the absolute error. The accuracy rates for the three experts increased, in general, across data sets DS1 to DS3 (76.86%–88.5%). The κ statistic increased (0.5784–0.7221) and error rates consistently decreased across the data sets for each expert (58.3%–36.1%). The accuracy range among the experts across data sets is 76.84% to 88.5%. The accuracy rate for cross-validation on the aggregate data of all three experts ranges from 75% to 82%. The most commonly misclassified drusen type by IDOCS was reticular drusen. In 6 of 9 instances, E1 classified reticular drusen as soft; E2 classified a single instance as soft; and E3 classified all 10 instances as soft. For the aggregate classifier, we see that hard (90/99) and soft (232/248) classifications had a very high accuracy; dystrophic (13/19) had a high accuracy; and calcified (5/22), basal laminar (0/18), and reticular (1/18) had a low accuracy.

The first three data rows of Table 2 show the results of training a classifier on one expert's data and using the other two experts' instances as test data for this classifier. The number of correct instances and the percentage accuracy rates are given. Expert E1 performed the worst with 30% and 41% accuracy, respectively. E2 performs roughly the same against

both the other experts (49% and 50%, respectively). E3, however, performs significantly better with an accuracy of 70% and 75%. The next three data rows of the table show the results of the classifier trained on the combined data of two experts and tested against the data of the remaining expert. E3 combined with the other two experts performed better against E1 (72.95%) and E2 (76.03%) than did the combined E1-E2 classifier against E3 (57.37%). To eliminate the possibility that the differing sizes of the test sets affected the accuracy rates for each expert, we repeated these tests on 10 sets of 100 instances randomly chosen from each expert's data set. The averaged results are consistent with our prior findings.

AMD Drusen Ontology

With the use of J48 decision trees, our tool achieved 77% to 82% accuracy (on data set DS1) in correctly predicting the drusen classes of the instances defined by each expert. A graphic depiction of these rules as a tree composed of multiple levels and nodes is shown in Figure 6. The elliptical nodes represent attributes, branches represent attribute values, and rectangular leaf nodes represent the predicted drusen types. The numerical values in the rectangular nodes indicate how many instances were classified by that node and how many were incorrectly classified by our tool. At each level of the tree, J48 uses the concept of entropy as a differentiating measure, to select the attribute that best partitions the data (or reduces the chaos). The trees of experts E2 and E3 were very similar. For each, the attributes *size* and *crystalline morphology* served as the deciding characteristics. The tree of expert E1 was larger but it also had these same attributes as key nodes. Overall, the trees indicate that *size*, *crystalline morphology* and *color* were the deciding attributes that achieved approximately 80% accuracy.

DISCUSSION

Computer engineering may have an important role in the process of AMD classification. Previous studies of knowledge acquisition have used system designers to define a set of criteria, which experts then assessed over a set of given examples. In IDOCS, the experts themselves implicitly defined the criteria through vocabulary generation, which was further structured by another expert. We are not aware of any previous performance-evaluation studies of knowledge elicitation that are comparable to our work.

Machine learning techniques have been successfully used in a variety of domains to extract patterns hidden within complex data that would be impossible to discern by human observation alone. In the biological sciences, these techniques have been applied to the fields of genomics, proteomics, and systems biology.²¹ Even in the field of AMD genetics, machine learning has found early application in creating models that can clarify

TABLE 2. Accuracy Rates from Classifying Data Instances of an Expert Using Classifier Trained on Another Expert's Data

Classifiers	E1 Test Set	E2 Test Set	E3 Test Set
E1		50 (41.32)	57 (30.00)
E2	61 (50.00)		93 (48.95)
E3	86 (70.49)	91 (75.21)	
E1-E2			109 (57.37)
E1-E3		92 (76.03)	
E2-E3	89 (72.95)		

The first three rows show the number and percentage (in parentheses) accuracy rates of a classifier trained on one expert's data to classify data instances of the other two experts. The next three rows show the same data for a classifier trained on the combined data of two experts to classify data instances of the third expert.

FIGURE 6. (A) Decision tree rules for the machine learning classifier for expert E1. Elliptical nodes are attributes, with the branches giving the attribute values on that path. Rectangular nodes at the end of a path are the drusen type deduced, with the first number in the parentheses being the number of instances that came down that path, and the second number being the number of instances that ended up misclassified by the decision tree on that path. (B) Decision tree rules for the machine learning classifier for expert E2. (C) Decision tree rules for the machine learning classifier for expert E3.

the significance of interactions between multiple genetic risk factors.²² If one were to envision the complexity of this experiment represented in graphic form, the data would appear the following way: more than 120 feature instances per expert plotted on five different attribute axes comprising more than 600 total data points for each expert. Our tool was able to distill all the information with higher than 80% accuracy into a flow chart that, in the case of E3, had as few as two decision nodes. Such remarkable findings reflect the power of this approach in identifying patterns in complex datasets, which has potential for significant applications in novice training and as an integrated part of clinical medical training to develop more standardized phenotyping for a variety of conditions.

Characterizing the relationships between the ontologies of the various experts is the first step in developing a classification scheme for more homogenous AMD subtypes. In our experiments, the results of testing each expert's classifier on the data instances of the other two experts and the results of testing on pair-wise combined classifiers seem to indicate that E3's classifier was the most general (it also had the lowest individual accuracy) and was able to classify the other experts' data the best, due to this generality. The other experts had higher individual accuracy rates but lower generalization power. E3's generalization power was further supported by

the fact that its decision tree was the smallest of the three. E1's low accuracy rates (30%–40%) can be explained by the fact that E1 had very few hard drusen instances ($n = 3$) and hence would heavily misclassify those instances from the other experts. The high number of matches for soft drusen among semantically equivalent features and the high accuracy rates for hard and soft drusen for the classifier trained on the aggregate data seem to indicate that the experts have very similar ontologies with respect to soft and, to a lesser extent, hard drusen. The lower accuracy rates for the other drusen types may indicate that the ontologies differ here or that there were just insufficient instances of each of them. In general, the more example instances that are available for a class, the better the decision tree can learn all the different patterns for the membership of that class.

There are several limitations to our study. Only three retina specialists participated in this initial experiment. A greater number of experts and a larger sampling of images would give us more confidence in our findings and a better sense of the generalizability of the study results. Future experiments will easily accommodate the involvement of more retina specialists. We envision the evolution of this classification system as being an iterative process incorporating both independent image analysis and group consensus meetings. This phase of IDOCS

entailed only the first half of this model. Future work will include the latter half as well. Although IDOCS was able to learn and classify soft and hard drusen satisfactorily, classification of less frequently used AMD feature names such as reticular had a paucity of examples as well as from the lack of consensus among the three graders regarding the definition of these rarer drusen types. Future iterations will benefit from a more uniform distribution of drusen types. IDOCS also exhibited a certain degree of bluntness in its descriptive power, in terms of its inability to capture the difference between reticular and soft drusen through the defined vocabulary. Finally, only one aspect of AMD was studied in this phase, retinal drusen characteristics, due to the limited data generated for the different retinal pigment epithelium characteristics like geographic atrophy and pigmentary abnormalities noted in our cohort. Future studies will be improved by the inclusion of the entire range of manifestations encompassed by AMD.

In summary, a structured vocabulary was generated, and three recognized AMD experts used it in a Web-based tool to formalize their knowledge by describing individual ontologies for retinal drusen. Decision tree machine learning algorithms were trained on these data with approximately 80% classification accuracy, and the ontologies were distilled into a set of simple rules. From these rules, it is observed that *size* and *crystalline morphology* (Fig. 6) of the drusen are critical attributes in classification. The experts were largely consistent among themselves in defining soft drusen; and expert E3, whose ontology could be described by the most general set of rules, performed the best in classifying instances described by the other experts, with a 70% to 75% accuracy. Future work will examine the IDOCS ontology in the context of novice training and improving clinical consistency of reporting diagnostic features of AMD.

The IDOCS prototype serves as a useful starting point toward the creation of more homogenous AMD subtypes. Having a defined feature set for AMD is a critical step toward an IDOCS-based classification system.

Acknowledgments

The authors thank the subjects and their families for their participation in the study; Alan C. Bird for participation in the vocabulary-generation phase; John Fingert, Robert Mullins, Michael Abramoff, and Stephen Russell for many helpful discussions; Stephen Russell, James Folk, and Culver Boldt for assistance in recruiting subjects; Victor Chong and Andrew Lotery for assistance in preliminary stages and prototype design; Andrea Lawrence, Ebony Smith, and Zijian Ren for technical assistance; and Ed Heffron for lending photographic expertise.

References

1. The Eye Diseases Prevalence Research G. Prevalence of Age-Related Macular Degeneration in the United States. *Arch Ophthalmol*. 2004;122:564-572.
2. Seddon JM, Ajani UA, Mitchell BD. Familial aggregation of age-related maculopathy. *Am J Ophthalmol*. 1997;123:199-206.
3. Klaver CCW, Wolfs RCW, Assink JJM, van Duijn CM, Hofman A, de Jong PTVM. Genetic risk of age-related maculopathy: population-based familial aggregation study. *Arch Ophthalmol*. 1998;116:1646-1651.
4. De Jong PT, Klaver CC, Wolfs RC, Assink JJ, Hofman A. Familial aggregation of age-related maculopathy. *Am J Ophthalmol*. 1997;124:862-863.
5. Edwards AO, Ritter R III, Abel KJ, Manning A, Panhuysen C, Farrer LA. Complement factor H polymorphism and age-related macular degeneration. *Science*. 2005;308:421-424.
6. Hageman GS, Anderson DH, Johnson LV, et al. A common haplotype in the complement regulatory gene factor H (HF1/CFH) predisposes individuals to age-related macular degeneration. *Proc Natl Acad Sci USA*. 2005;102:7227-7232.
7. Haines JL, Hauser MA, Schmidt S, et al. Complement factor H variant increases the risk of age-related macular degeneration. *Science*. 2005;308:419-421.
8. Klein RJ, Zeiss C, Chew EY, et al. Complement factor H polymorphism in age-related macular degeneration. *Science*. 2005;308:385-389.
9. Conley YP, Thalamuthu A, Jakobsdottir J, et al. Candidate gene analysis suggests a role for fatty acid biosynthesis and regulation of the complement system in the etiology of age-related maculopathy. *Hum Mol Genet*. 2005;14:1991-2002.
10. Jakobsdottir J, Conley YP, Weeks DE, Mah TS, Ferrell RE, Gorin MB. Susceptibility genes for age-related maculopathy on chromosome 10q26. *Am J Hum Genet*. 2005;77:389-407.
11. Rivera A, Fisher SA, Fritsche LG, et al. Hypothetical LOC387715 is a second major susceptibility gene for age-related macular degeneration, contributing independently of complement factor H to disease risk. *Hum Mol Genet*. 2005;14:3227-3236.
12. Zarepari S, Buraczynska M, Branham KEH, et al. Toll-like receptor 4 variant D299G is associated with susceptibility to age-related macular degeneration. *Hum Mol Genet*. 2005;14:1449-1455.
13. Petrukhin K, Koisti MJ, Bakall B, et al. Identification of the gene responsible for Best macular dystrophy. *Nat Genet*. 1998;19:241-247.
14. Weber BH, Vogt G, Pruett RC, Stohr H, Felbor U. Mutations in the tissue inhibitor of metalloproteinases-3 (TIMP3) in patients with Sorsby's fundus dystrophy. *Nat Genet*. 1994;8:352-356.
15. Nichols BE, Sheffield VC, Vandenburgh K, Drack AV, Kimura AE, Stone EM. Butterfly-shaped pigment dystrophy of the fovea caused by a point mutation in codon 167 of the RDS gene. *Nat Genet*. 1993;3:202-207.
16. Allikmets R, Singh N, Sun H, et al. A photoreceptor cell-specific ATP-binding transporter gene (ABCR) is mutated in recessive Stargardt macular dystrophy (published correction appears in *Nat Genet*. 1997;17:122). *Nat Genet*. 1997;15:236-246.
17. Golub TR, Slonim DK, Tamayo P, et al. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*. 1999;286:531-537.
18. Sotiriou C, Neo SY, McShane LM, et al. Breast cancer classification and prognosis based on gene expression profiles from a population-based study. *Proc Natl Acad Sci USA*. 2003;100:10393-10398.
19. Williams AB, Krygowski T, Casavant T. I-DOCS: distributed agent-assisted knowledge fusion for disease gene discovery. *Proceedings of the Eighth International Conference on Parallel and Distributed Systems (ICPADS 2001)*, Kyongju City, Korea, June 26-29, 2001. Washington, DC: IEEE Computer Society Press; 2001:698-705.
20. Witten I, Frank E. *Data Mining: Practical Machine Learning Tools and Techniques*. 2nd ed. New York: Elsevier; 2005.
21. Larranaga P, Calvo B, Santana R, et al. Machine learning in bioinformatics. *Brief Bioinform*. 2006;7:86-112.
22. Gold B, Merriam JE, Zernant J, et al. Variation in factor B (BF) and complement component 2 (C2) genes is associated with age-related macular degeneration. *Nat Genet*. 2006;38:458-462.