

Automated processing and identification of benthic invertebrate samples

David A. Lytle^{1,7}, Gonzalo Martínez-Muñoz^{2,8}, Wei Zhang^{2,9},
Natalia Larios^{3,10}, Linda Shapiro^{3,4,11}, Robert Paasch^{5,12},
Andrew Moldenke^{6,13}, Eric N. Mortensen^{2,14}, Sinisa Todorovic^{2,15}, AND
Thomas G. Dieterich^{2,16}

¹ Department of Zoology, Oregon State University, Corvallis, Oregon 97331 USA

² School of Electrical Engineering and Computer Science, Oregon State University,
Corvallis, Oregon 97331 USA

³ Department of Electrical Engineering, University of Washington, Seattle, Washington 98195 USA

⁴ Department of Computer Science and Engineering, University of Washington,
Seattle, Washington 98195 USA

⁵ Department of Mechanical Engineering, Oregon State University, Corvallis, Oregon 97331 USA

⁶ Department of Botany and Plant Pathology, Oregon State University, Corvallis, Oregon 97331 USA

Abstract. We present a visually based method for the taxonomic identification of benthic invertebrates that automates image capture, image processing, and specimen classification. The BugID system automatically positions and images specimens with minimal user input. Images are then processed with interest operators (machine-learning algorithms for locating informative visual regions) to identify informative pattern features, and this information is used to train a classifier algorithm. Naïve Bayes modeling of stacked decision trees is used to determine whether a specimen is an unknown distractor (taxon not in the training data set) or one of the species in the training set. When tested on images from 9 larval stonefly taxa, BugID correctly identified 94.5% of images, even though small or damaged specimens were included in testing. When distractor taxa (10 common invertebrates not present in the training set) were included to make classification more challenging, overall accuracy decreased but generally was close to 90%. At the equal error rate (EER), 89.5% of stonefly images were correctly classified and the accuracy of nonrejected stoneflies increased to 96.4%, a result suggesting that many difficult-to-identify or poorly imaged stonefly specimens had been rejected prior to classification. BugID is the first system of its kind that allows users to select thresholds for rejection depending on the required use. Rejected images of distractor taxa or difficult specimens can be identified later by a taxonomic expert, and new taxa ultimately can be incorporated into the training set of known taxa. BugID has several advantages over other automated insect classification systems, including automated handling of specimens, the ability to isolate nontarget and novel species, and the ability to identify specimens across different stages of larval development.

Key words: bioassessment, water quality, automated insect identification, bioindicators, Plecoptera.

⁷ lytle@oregonstate.edu

⁸ martinez@eecs.oregonstate.edu

⁹ zhangwe@eecs.oregonstate.edu

¹⁰ nlarios@u.washington.edu

¹¹ shapiro@cs.washington.edu

¹² paasch@engr.orst.edu

¹³ moldenka@science.oregonstate.edu

¹⁴ enm@eecs.oregonstate.edu

¹⁵ sinisa@eecs.oregonstate.edu

¹⁶ tgd@oregonstate.edu

Automated taxonomic identification is an area of rapid innovation, and computer vision methods are being developed for visual classification of spiders (Do et al. 1999), butterflies (Watson et al. 2004, Bhanu et al. 2008), plants (Clark 2007), plankton (Rodenacker et al. 2006), and other groups (reviewed in MacLeod 2007). Many of these approaches achieve high levels of accuracy with specific data sets, but considerable challenges remain before automated classification can be deployed on a large scale. First, automated methods for specimen handling and image processing

are needed in addition to automated image classification. Most approaches currently rely on significant input from human users, such as landmarking of morphological features on each image by hand, which considerably slows the classification process. This step requires too much time to be practical for insect studies that might have hundreds of specimens per sample. Second, new methods must be able to recognize and reject novel species that lie outside their training set of known species. With current approaches, samples must be prescreened to ensure that nontarget specimens are removed, a requirement that necessitates some amount of a priori specimen classification and that undercuts the utility of automated classification systems in the first place. Last, methods must be robust to variability arising from ontogenetic changes within species, damaged specimens, and variable imaging conditions.

These challenges are especially apparent in benthic invertebrate samples collected as part of bioassessment projects. A benthic invertebrate sample can be collected in a matter of minutes, but sample processing might require hours or even days of laboratory work. Thus, an automated method for benthic invertebrate sample processing and classification could greatly increase the number of samples used for monitoring and conservation efforts. Sample processing usually involves separating target taxa from detritus (fragments of leaves, sand, and other debris) and nontarget species and then classification of specimens to an appropriate taxonomic level (usually genus or species), so the ability to isolate unknown specimens is important. Last, invertebrate specimens span a large range of ontogenetic diversity because most taxa are collected during the larval stage, and specimens are often bent, broken, and discolored with sediment.

Our BugID approach automates specimen handling, image capture, and image classification into a single process (Sarpola et al. 2008). The method is not fully integrated from sample jar to identified specimens, but we have automated several key steps. We begin with an apparatus that automatically positions specimens under a microscope and captures images. Images are then rescaled and processed to remove background noise. The classification method uses feature-based classification techniques developed by our group (Larios et al. 2007, Mortensen et al. 2007, Martínez-Muñoz 2009). Rather than focus on specific diagnostic features typically used by human taxonomists, this feature-based approach applies machine learning techniques to find multiple regions of interest within images that are informative for species classification. These regions are analyzed simulta-

neously to provide a classification; the method is similar to an experienced taxonomist integrating many subtle visual cues to sight-identify a specimen by *gestalt*, or overall appearance. This approach minimizes reliance on specific diagnostic characters that might require manipulation, dissection, or other labor-intensive handling.

Here we describe the general BugID approach. We trained it using images from 9 larval stonefly taxa commonly found in Pacific Northwest streams and rivers, and then tested its ability to classify images from novel specimens of these 9 taxa. We then tested BugID's accuracy in the context of distractors, images from 10 common invertebrates that were not in the training set.

Methods

Sample collection and processing

Benthic macroinvertebrates were collected from Oregon rivers and streams with standard Oregon Department of Environmental Quality protocols (Hafele and Mulvey 1998). Collection sites were distributed across several biotic provinces, including temperate coastal rainforest, alpine forest, high desert, and valley floodplain. In the laboratory, specimens were separated and identified to genus or species. All larval stages and damaged specimens (missing legs, antennae) were included, and each specimen was identified independently by 2 taxonomists. All specimens were kept in individual vials, assigned unique numbers, and accessioned into the Oregon State Arthropod Collection after imaging (OSAC lot #0278).

Our main focus was on the larval stage of 9 stonefly taxa spanning 7 families. Each taxon was common in streams in the Pacific Northwest: *Calineuria californica*, *Doroneuria baumanni*, *Hesperoperla pacifica* (Perlidae), *Isoperla* sp. (Perlodidae), *Moselia infuscata* (Leuctridae), *Pteronarcys* sp. (Pteronarcyidae), *Sweltsa* sp. (Chloroperlidae), *Yoraperla* sp. (Peltoperlidae), and *Zapada* sp. (Nemouridae). These names were abbreviated in our analyses as CAL, DOR, HES, ISO, MOS, PTE, SWE, YOR, and ZAP, respectively. We obtained ~100 specimens for most species, except *Moselia* (24 specimens) and *Pteronarcys* (45 specimens).

To test the ability of the algorithms to discern known stonefly species from unknown or novel specimens, we used larval specimens of 10 distractor taxa that are commonly found in benthic invertebrate samples (10 specimens per taxon): the ephemeropterans *Baetis* (Baetidae), *Ameletus* (Ameletidae), *Ephemerella* (Ephemerellidae), *Caudatella* (Ephemerellidae), *Ironodes* (Heptageniidae), and *Rhithrogena* (Heptageniidae); the trichopterans *Brachycentrus* (Brachycentri-

dae) and *Neophylax* (Uenoidae); the plecopteran *Taenionema* (Taeniopterygidae); and the amphipod *Hyaella* (Hyaellidae). These specimens were considered distractors in the sense that we did not include them in our initial training set. However, they could be incorporated into the BugID system as known taxa at a future date.

Image capture

We processed and imaged specimens with an updated version of the BugID system described in Sarpola et al. (2008). Specimens are dropped singly into a Plexiglas® tube filled with 70% ethanol, and a recirculating pump moves individuals to the stage of a standard dissecting microscope (Leica MZ9.5) fitted with a 5-megapixel camera. An infrared sensor automatically detects a specimen in the field of view and causes the pump to cease. A pulse jet rotates specimens until a dorsal view is available, and the operator chooses when to take an image. We captured 4 to 5 dorsal images per specimen for our 9 stonefly taxa, rotating specimens between each image so that no 2 images were exactly alike. Partial images (antennae, legs, or cerci out of the field of view) were retained in our image set. We collected a total of 349 images of the distractor taxa (3–4/specimen).

We divided specimens at random into 3 sets (called folds) with approximately the same number of specimens per set. The folds were generated by selecting the specimens at random, but with the constraint that all images of a single specimen were always in the same fold. The experiments were carried out with 3-fold cross-validation. This step involved 3 iterations of experiments where, in each round, 2 of the folds were used for classifier training and estimation of rejection parameters, and the 3rd fold was used for testing. Hence, every image was used twice for training and once for testing, but never for both training and testing in the same iteration.

Classification process

The first step before training any classification algorithm is to apply interest operators to the images. In our case we computed Hessian affine regions (Mikolajczyk and Schmid 2002), saliency regions (Kadir and Brady 2001), principal curvature-based regions (PCBRs; Deng et al. 2007), and regularly sampled points (Martínez-Muñoz et al. 2009). These operators find points of interest in the images, such as corners, ridges, lobes, patterning, and other morphological features. The detected regions were then described using the Scale Invariant Feature Transform (SIFT) descriptor (Lowe 2004). Each SIFT describes a

region as a vector of 128 numerical values that are approximately viewpoint-invariant and illumination-invariant. The classification system was built by analyzing these extracted SIFT descriptors.

The first step of the learning process involved learning random forest classifiers directly from the descriptors. A random forest classifier (Breiman 2001) is an ensemble of decision trees, where each tree is a computer-generated structure analogous to a dichotomous key. Each step in the tree compares a chosen element of the SIFT vector against a chosen threshold, and branches on the result. In a random forest, each tree is constructed by a process that incorporates a random component to introduce variation amongst the trees. Each of our random forest classifiers contained 100 decision trees. A separate 100-tree random forest was trained for each of the 3 interest operators. The random forest classifier took as input a single SIFT vector. Each tree analyzed this SIFT vector and predicted the taxon of the specimen based on it. The classifications from all of the SIFTs and all of the trees were accumulated as votes. Hence, there were $(D_1 + D_2 + D_3) \times 100$ votes, where D_1 denotes the number of detected Hessian Affine regions, D_2 denotes the number of detected Kadir regions, and D_3 denotes the number of detected PCBR regions. These votes were combined by a 2nd-level stacked classifier to make the final prediction of the taxon (see Martínez-Muñoz et al. 2009 for more details on the algorithm).

Distinguishing between unknown distractor taxa and known taxa that were included in the training set presented a number of challenges. Because this classification system was trained only on 9 stonefly taxa, it could categorize images only as belonging to one of these taxa. Hence, before classifying an image, the system had to determine whether the image belonged to this group of taxa. This task was challenging because only the 9 taxa of interest were available during training, so the system had no way of modeling the distractor taxa. Our approach was to estimate the probability of a specimen belonging to any of the stonefly taxa of interest, rather than modeling all possible distractor species (probably an impossible task). Unlike classification, which seeks to find parts of the stoneflies that discriminate among the different taxa, rejection of distractors required us to locate parts of the training images that were shared across taxa. We accomplished this task with the following steps. First, the SIFT descriptors were clustered (via K-means clustering) to form a visual dictionary (Csurka et al. 2004). Next, each SIFT descriptor in an image was mapped to the nearest dictionary entry, and those indices were accumulated

TABLE 1. Confusion matrix for the 9 stonefly taxa with no distractors present. Rows show true taxonomic identities and columns show the identities output by the classifier. Each entry records the number of images whose true taxonomic group corresponds to the row and whose predicted group corresponds to the column. Overall, 94.5% of images were correctly identified. (CAL = *Calineuria californica*, DOR = *Doroneuria baumanni*, HES = *Hesperoperla pacifica*, ISO = *Isogenoides* sp., MOS = *Moselia infuscata*, PTE = *Pteronarcys* sp., SWE = *Sweltsa* sp., YOR = *Yoraperla* sp., ZAP = *Zapata* sp.).

Taxon	CAL	DOR	HES	ISO	MOS	PTE	SWE	YOR	ZAP	Correct
CAL	459	11	3	1	1	2	12	0	3	93.3%
DOR	10	505	2	3	0	1	10	0	1	94.9%
HES	5	2	464	2	0	1	17	0	0	94.5%
ISO	2	7	1	460	0	0	26	0	4	92.0%
MOS	0	0	0	0	107	1	3	0	8	89.9%
PTE	0	0	0	0	0	218	5	0	0	97.8%
SWE	0	12	0	13	2	3	442	0	7	92.3%
YOR	1	1	0	0	0	0	4	483	3	98.2%
ZAP	1	0	1	4	2	1	8	3	478	96.0%

into a histogram that represented the image in terms of the number of interest regions belonging to each dictionary entry. Last, a Naïve Bayes probabilistic model was fitted to these histograms. This model could be applied to a new image to estimate the probability that the image belongs to *any* of the 9 stonefly taxa. New specimens were rejected as being nonstoneflies based on a threshold or operating point (OP) for this probability. For the results reported in this paper, we selected the OP at the equal error rate (EER; the EER OP is where the proportion of distracters rejected equals the proportion of stoneflies accepted). We estimated the OP threshold using only the training set (without distractors) by averaging over the 9 EER thresholds obtained by treating 1 stonefly taxon at a time as a distractor and using the Naïve Bayes model trained on the remaining 8 taxa.

Given a new image, the system operated as follows. First, the Naïve Bayes model was evaluated, and the output probability was compared to the OP threshold. If it was less than the threshold, the image was rejected as being a nonstonefly. Otherwise, the SIFT descriptors were passed to the classifier, which assigned the image to one of the 9 stonefly taxa.

We explored 2 scenarios commonly encountered with sample processing. First, we tested BugID's performance when all specimens belonged to known species (i.e., species present in the training set). This situation occurs, for instance, in samples from experiments where the species pool is controlled or in samples that have been presorted by technicians to remove groups that are not in the training set. Second, we tested BugID's ability to remove unknowns automatically and then to identify the remaining specimens correctly. This scenario is encountered with most field-collected samples, where novel species (not in training set) often are found in samples.

To put our automated results into perspective, we compared them with results from humans asked to identify the same images (Larios et al. 2007). Twenty-six graduate students and faculty (entomologists, ecologists, computer scientists, and mechanical engineers) were trained by showing them 50 labeled images of *Calineuria* and *Doroneuria* drawn at random from the image set. Their ability to identify 50 unlabeled images was then tested.

Results

BugID correctly identified 94.5% of stonefly images when distractors were excluded (Table 1). The similar perlids *Calineuria* and *Doroneuria* were most often confused, and *Sweltsa* was confused with many of the other taxa. *Yoraperla*, which has a distinctive body shape that is different from most of the other taxa, had the highest correct classification rate at 98.2%.

All of the receiver-operator characteristic (ROC) curves in our experiments approached the upper left corner of the graph, suggesting a favorable tradeoff between correct rejection of distractors and incorrect rejection of stoneflies (Fig. 1). The total area under the curve for each of the 3 folds was 87.5%, 87.5%, and 85.9% (perfect discrimination is 100%). Results from the 1st fold are shown in Fig. 1. Operating point P₂ represents the EER, where the percentage of stonefly images misidentified as distractors equaled the percentage of distractors misidentified as stoneflies. The EER was similar for each of the 3 folds (20.6%, 20.1%, and 23.2%), and at this point, the overall accuracy with distractors included was 89.4%. However, 96.4% of the nonrejected stonefly images were correctly identified, suggesting that some of the more problematic stonefly images were rejected and not classified (Table 2). Each point in the ROC curve is computed by setting a different threshold on the




FIG. 1. Receiver-operator characteristic (ROC) curves showing the tradeoff between correct rejection of distractors and incorrect rejection of selected stonefly taxa. Operating point P_2 is the equal error rate (ERR); P_1 represents greater inclusion of images for classification; P_3 represents more conservative rejection of images prior to classification. *Pteronarcys* and *Doroneuria* had the best and worst ROC curves, respectively, with all other taxa occurring in between.

probability of being a distractor. The points P_1 and P_3 were obtained by multiplying the EER threshold by 0.95 and 1.05 respectively. Operating point P_1 represents a scenario where inclusiveness is more important (it is more important to identify many stonefly specimens than to reject distractors). In this case, the accuracy decreased to 86.0%. Operating point P_3 represents the opposite scenario where removing distractors is paramount, even at the expense of rejecting many stoneflies as distractors. Here, 38.9% of stonefly images were rejected incorrectly, but 89.8% of distractors were also removed prior to classification.

In the human study, the 26 graduate students and faculty successfully identified an average of 78.6%

(SD = 8.4) of *Calineuria* and *Doroneuria* images. This performance was substantially less accurate than the BugID system, which correctly identified >90% of both taxa, even when faced with rejecting distractors and classifying 9 rather than 2 taxa.

Discussion

Our experiments emphasize the importance of detecting and removing distractors prior to classification. We found a clear tradeoff between rejection of unknown taxa and incorrect rejection of stonefly specimens, although the ROC curves suggest a reasonable compromise is possible. In our experi-

TABLE 2. Confusion matrix for the 9 stonefly taxa plus distractors (Dis). Rows show true taxonomic identities. Results are from experiments at the equal operating rate. Of the nonrejected stonefly images, 96.4% were identified correctly. (CAL = *Calineuria californica*, DOR = *Doroneuria baumanni*, HES = *Hesperoperla pacifica*, ISO = *Isogenoides* sp., MOS = *Moselia infuscata*, PTE = *Pteronarcys* sp., SWE = *Sweltsa* sp., YOR = *Yoraperla* sp., ZAP = *Zapata* sp.).

Taxon	CAL	DOR	HES	ISO	MOS	PTE	SWE	YOR	ZAP	Dis	Rejected	Correct
CAL	373	10	3	1	0	2	5	0	2	96	19.5%	94.2%
DOR	7	398	2	2	0	0	8	0	1	114	21.4%	95.2%
HES	5	1	357	0	0	1	2	0	0	125	25.5%	97.5%
ISO	1	1	0	361	0	0	11	0	0	126	25.2%	96.5%
MOS	0	0	0	0	99	0	3	0	8	9	7.6%	90.0%
PTE	0	0	0	0	0	188	4	0	0	31	13.9%	97.9%
SWE	0	6	0	5	0	1	348	0	2	117	24.4%	96.1%
YOR	1	1	0	0	0	0	1	402	1	86	17.5%	99.0%
ZAP	1	0	1	0	2	1	6	2	432	53	10.6%	97.1%
Dis	5	25	27	61	1	9	41	3	68	807		

ments, the ratio of distractor to stonefly images was quite high, ~1 distractor image for every 3.5 stonefly images. The operating thresholds probably would improve (ROC curves would move toward 100% area under curve) if distractors were less frequent in samples. One advantage of our method is that the user can set the threshold for rejection to be high or low. As more taxa are included in the training set, this threshold could be relaxed to ensure that a high proportion of specimens is classified. When accuracy is desired, a stringent rejection threshold could be used, and the rejected specimens could be identified later by a specialist. This practice would be most important when processing samples that contain a high proportion of unknown taxa or poor-quality specimens. In our experiments, we saw that accuracy of nonrejected stoneflies jumped up to 96.4% at the EER, probably because many poor-quality images and difficult-to-identify specimens had been rejected prior to classification.

The BugID approach represents several steps forward in automated specimen classification. We relied primarily on automated methods to handle and image the specimens (Sarpola et al. 2008). Our specimen data set also was realistic because we did not exclude specimens based on size or condition. Differences in coloration, size, posture (curled abdomens, legs extended vs contracted), and lighting were commonplace. Thus, our specimens were representative of those found in real samples collected for bioassessment purposes. Last, we confronted the system with distractor specimens representing unknown taxa. Despite these challenges, the BugID system correctly identified >90% of images, and had >96% accuracy when only nonrejected specimens were considered.

On average, BugID was able to discriminate between *Calineuria* and *Doroneuria* much more accu-

rately than humans. Part of the explanation for this result is that the dorsal view images collected by our apparatus do not capture many of the characters normally used by taxonomic experts to identify specimens, and for this reason, separating *Calineuria* from *Doroneuria* using a dorsal view is a difficult task for either humans or the BugID system. In the absence of these characters, the humans had to discover new characters that were present in the dorsal views (see Fig. 2). We hypothesize that humans were able to find only a few such pattern characters, whereas BugID was able to perform a systematic search for a large number of such pattern characters and combine them all simultaneously to make the classifications. For some of the specimens, especially smaller specimens, the diagnostic dorsal features used by humans were not always visible, which resulted in errors.

Our system can automatically identify macroinvertebrate specimens, but we predict that automated approaches will increase, rather than diminish, the need for skilled taxonomists. Sample verification, incorporation of new taxa and populations, and identification of specimens rejected as distractors always will require taxonomic expertise. Systems such as BugID will serve more to increase the number of samples, and the number of specimens per sample, that can be incorporated feasibly into research or biomonitoring projects. The visually based BugID method is distinct from DNA-based methods that identify unique sequences in aggregate samples of specimens (e.g., DNA barcoding; Savolainen et al. 2005). Unlike current DNA-based methods, visually based methods count the number of individuals in each taxon in a sample (DNA methods identify unique haplotypes or provide only relative estimates of haplotype frequencies), and specimens are not destroyed (aggregate DNA analysis generally requires homogenization of samples).



FIG. 2. Examples of specimen images used in experiments. Top row to bottom: 3 images each of *Calineuria californica*, *Doroneuria baumanni*, *Hesperoperla pacifica*, *Isogenoides* sp., *Moselia infusata*, *Pteronarcys* sp., *Sweltsa* sp., *Yoraperla* sp., *Zapata* sp.

Prospects for future improvements include: 1) automated sorting of identified specimens into separate containers, which would facilitate sample vouchering and verification by experts, 2) better methods for detecting and rejecting distractors, and 3) algorithms for identifying specimens in photographs that contain background clutter, which would enable the system to identify images collected remotely and then submitted electronically via email or web-upload.

Acknowledgements

We thank Justin Miles, Frank Drake, Kate Boersma, Mike Bogan, Laura McMullen, and Asako Yamamuro for help with specimen collection and identification. Funding was provided by grants from the National Science Foundation (IIS-0326052, IIS-0705765, and DEB-0445366). GM-M acknowledges support from the Fulbright Program and the Spanish Ministerio de Ciencia e Innovación, project TIN2007-66862-C02-02.

Literature Cited

- BHANU, B., R. LI, J. HERATY, AND E. MURRAY. 2008. Automated classification of skippers based on parts representation. *American Entomologist* 54:228–231.
- BREIMAN, L. 2001. Random forests. *Machine Learning* 45:5–32.
- CLARK, J. Y. 2007. Plant identification from characteristics and measurements using artificial neural networks. Pages 207–224 in N. MacLeod (editor). *Automated taxon identification in systematics: theory, approaches and applications*. CRC Press, Boca Raton, Florida.
- CSURKA, G., C. R. DANCE, L. FAN, J. WILLAMOWSKI, AND C. BRAY. 2004. Visual categorization with bags of keypoints. *Workshop on Statistical Learning in Computer Vision, European Conference on Computer Vision 2004*:1–22.
- DENG, H., W. ZHANG, E. MORTENSEN, T. DIETTERICH, AND L. SHAPIRO. 2007. Principal curvature-based region detector for object recognition. *IEEE Conference on Computer Vision and Image Processing 2007*:1–8.
- DO, M. T., J. M. HARP, AND K. C. NORRIS. 1999. A test of a pattern recognition system for identification of spiders. *Bulletin of Entomological Research* 89:217–224.
- HAFELE, R., AND M. MULVEY. 1998. Stream macroinvertebrate protocol. Oregon plan for salmon and watersheds. March 1998 draft. Oregon Department of Environmental Quality, Portland, Oregon, (Available from: Oregon Department of Environmental Quality, 811 SW 6th Ave., Portland, Oregon 97204 USA.)
- KADIR, T., AND M. BRADY. 2001. Scale, saliency and image description. *International Journal of Computer Vision* 45(2):83–105.
- LARIOS, E., H. DENG, W. ZHANG, M. SARPOLA, J. YUEN, R. PAASCH, A. MOLDENKE, D. A. LYTLE, S. R. CORREA, E. MORTENSEN, L. SHAPIRO, AND T. DIETTERICH. 2007. Automated insect identification through concatenated histograms of local appearance features: feature vector generation and region detection for deformable objects. *Machine Vision and Applications*. doi: 10.1007/s00138-007-0086-y.
- LOWE, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2):91–110.
- MACLEOD, N. 2007. *Automated taxon identification in systematics: theory, approaches and applications*. CRC Press, Boca Raton, Florida.
- MARTINEZ-MUÑOZ, G., W. ZHANG, N. PAYET, S. TODOROVIC, N. LARIOS, A. YAMAMURO, D. LYTLE, A. MOLDENKE, E. MORTENSEN, R. PAASCH, L. SHAPIRO, AND T. DIETTERICH. 2009. Dictionary-free categorization of very similar objects via stacked evidence trees. *Computer Vision and Pattern Recognition 2009*:549–556.
- MIKOLAJCZYK, K., AND C. SCHMID. 2002. An affine invariant interest point detector. *European Conference on Computer Vision 2002*:128–142.
- MORTENSEN, E. M., E. L. DELGADO, H. DENG, D. A. LYTLE, A. MOLDENKE, R. PAASCH, L. SHAPIRO, P. WU, W. ZHANG, AND T. DIETTERICH. 2007. Pattern recognition for ecological science and environmental monitoring: an initial report. Pages 189–206 in N. MacLeod (editor). *Automated taxon identification in systematics: theory, approaches and applications*. CRC Press, Boca Raton, Florida.
- RODENACKER, K., B. HENSE, U. JUTTING, AND P. GAIS. 2006. Automatic analysis of aqueous specimens for phytoplankton structure recognition and population estimation. *Microscopy Research and Technique* 69:708–720.
- SARPOLA, M., R. PAASCH, E. MORTENSEN, T. DIETTERICH, D. A. LYTLE, A. MOLDENKE, AND L. SHAPIRO. 2008. An aquatic insect imaging system to automate insect classification. *Transactions of the American Society of Agricultural and Biological Engineers* 51:2217–2225.
- SAVOLAINEN, V., R. S. COWAN, A. P. VOGLER, G. K. RODERICK, AND R. LANE. 2005. Towards writing the encyclopaedia of life: an introduction to DNA barcoding. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences* 360:1805–1811.
- WATSON, A. T., M. A. O'NEILL, AND I. J. KITCHING. Automated identification of live moths (Macrolepidoptera) using Digital Automated Identification SYstem (DAISY). *Systematics and Biodiversity* 1:287–300.

Received: 23 June 2009

Accepted: 22 April 2010