Correction

PLANT BIOLOGY

Correction for "Robust predictions of specialized metabolism genes through machine learning," by Bethany M. Moore, Peipei Wang, Pengxiang Fan, Bryan Leong, Craig A. Schenck, John P. Lloyd, Melissa D. Lehti-Shiu, Robert L. Last, Eran Pichersky, and Shin-Han Shiu, which was first published February 5, 2019; 10.1073/pnas.1817074116 (*Proc Natl Acad Sci USA* 116:2344–2353).

The authors note that the following statement should be added to the Acknowledgments: "This publication was made possible in part by a predoctoral training award to B.L. from Grant T32-GM110523 from the National Institute of General Medical Sciences of the National Institutes of Health."

Published online March 4, 2019.

www.pnas.org/cgi/doi/10.1073/pnas.1902386116

Robust predictions of specialized metabolism genes through machine learning

Bethany M. Moore[a,b], Peipei Wang[a], Pengxiang Fan[c], Bryan Leong[a], Craig A. Schenck[c], John P. Lloyd[d,e], Melissa D. Lehti-Shiu[a], Robert L. Last[a,c], Eran Pichersky[f], and Shin-Han Shiu[a,b,g,1]

[a]Department of Plant Biology, Michigan State University, East Lansing, MI 48824; [b]Ecology, Evolutionary Biology, and Behavior Program, Michigan State University, East Lansing, MI 48824; [c]Department of Biochemistry and Molecular Biology, Michigan State University, East Lansing, MI 48824; [d]Department of Human Genetics, University of Michigan, Ann Arbor, MI 48109; [e]Department of Internal Medicine, University of Michigan, Ann Arbor, MI 48109; [f]Department of Molecular, Cellular and Developmental Biology, University of Michigan, Ann Arbor, MI 48109; and [g]Department of Computational Mathematics, Science, and Engineering, Michigan State University, East Lansing, MI 48824

Edited by Julia Bailey-Serres, University of California, Riverside, CA, and approved December 17, 2018 (received for review October 10, 2018)

Plant specialized metabolism (SM) enzymes produce lineage-specific metabolites with important ecological, evolutionary, and biotechnological implications. Using *Arabidopsis thaliana* as a model, we identified distinguishing characteristics of SM and GM (general metabolism, traditionally referred to as primary metabolism) genes through a detailed study of features including duplication pattern, sequence conservation, transcription, protein domain content, and gene network properties. Analysis of multiple sets of benchmark genes revealed that SM genes tend to be tandemly duplicated, coexpressed with their paralogs, narrowly expressed at lower levels, less conserved, and less well connected in gene networks relative to GM genes. Although the values of each of these features significantly differed between SM and GM genes, any single feature was ineffective at predicting SM from GM genes. Using machine learning methods to integrate all features, a prediction model was established with a true positive rate of 87% and a true negative rate of 71%. In addition, 86% of known SM genes not used to create the machine learning model were predicted. We also demonstrated that the model could be further improved when we distinguished between SM, GM, and junction genes responsible for reactions shared by SM and GM pathways, indicating that topological considerations may further improve the SM prediction model. Application of the prediction model led to the identification of 1,220 *A. thaliana* genes with previously unknown functions, each assigned a confidence measure called an SM score, providing a global estimate of SM gene content in a plant genome.

specialized metabolism | machine learning | predictive biology | data integration

Gene duplication and subsequent divergence/loss events led to highly variable gene content between plant species (1, 2). These differential gain and loss events have given rise to diverse metabolic enzymes ranging from those involved in generally conserved, primary metabolic processes found in most species [referred to as general metabolism (GM) genes], to those that function in lineage-specific, specialized metabolism (SM) (3–6). The proliferation of SM genes in plants has resulted in an overall far larger number of specialized than general metabolites. These specialized metabolites are important for niche-specific interactions between plants and environmental agents that can be harmful (e.g., herbivores) or beneficial (e.g., pollinators) (3, 7, 8). They are also the basis for thousands of plant-derived chemicals, many of which are used for medicinal and/or nutritional purposes, such as carotenoid derivatives with antioxidant properties in tomato (9–11). Thus, identification of the genes encoding enzymes that produce specialized metabolites (referred to as SM genes) is key to understanding the causes underlying the diversity of plant specialized metabolites as well as for engineering plant-derived chemicals and pharmaceuticals.

Despite their importance, most plant metabolites and the enzymes and genes involved in their biosynthesis are yet to be identified (12). Although many SM genes arise by duplication of GM genes (13, 14) or other SM genes (15), duplication itself is not sufficient for pinpointing SM genes for four reasons. First, genes encoding GM or SM enzymes can belong to the same family. Second, duplicated GM genes may not necessarily become specialized (1), and minor sequence changes can lead to substantially altered enzyme functions (16, 17). Third, SM genes may arise through lineage-specific loss of the GM function without duplication. Finally, convergent evolution may explain the presence of unrelated enzymes in different lineages that use the same substrate to make similar products (5). Consequently, it remains unresolved whether most plant enzyme genes are involved in GM or SM pathways, even in the best annotated plant species, *Arabidopsis thaliana* (3, 5, 18, 19). Therefore, in recent years there has been an enhanced focus on identifying SM genes (20, 21). Multiple properties have been shown to differ between SM and GM genes (4, 20–22). For example, whole genome duplications (WGDs) and tandem duplications both contribute to metabolic innovations in glucosinolate biosynthesis genes (23). In addition, compared with GM genes, SM genes tend to have a more restricted phylogenetic distribution, a higher family expansion rate, tandem clustering of paralogs, a propensity for genomic clustering (close physical proximity of genes encoding enzymes in the same pathways), higher degrees of expression variation, and

Significance

Specialized metabolites are critical for plant–environment interactions, e.g., attracting pollinators or defending against herbivores, and are important sources of plant-based pharmaceuticals. However, it is unclear what proportion of enzyme-encoding genes play a role in specialized metabolism (SM) as opposed to general metabolism (GM) in any plant species. This is because of the diversity of specialized metabolites and the considerable number of incompletely characterized pathways responsible for their production. In addition, SM gene ancestors frequently played roles in GM. We evaluate features distinguishing SM and GM genes and build a computational model that accurately predicts SM genes. Our predictions provide candidates for experimental studies, and our modeling approach can be applied to other species that produce medicinally or industrially useful compounds.

Author contributions: B.M.M., M.D.L.-S., R.L.L., E.P., and S.-H.S. designed research; B.M.M., P.W., and S.-H.S. performed research; B.M.M., P.W., J.P.L., and S.-H.S. contributed new reagents/analytic tools; B.M.M., P.W., P.F., B.L., C.A.S., and S.-H.S. analyzed data; and B.M.M., P.W., P.F., B.L., C.A.S., M.D.L.-S., R.L.L., E.P., and S.-H.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: shius@msu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1817074116/-/DCSupplemental.

Published online January 23, 2019.

higher degrees of coexpression. Coexpression with known SM genes (20, 24) or genomic neighborhood and gene–metabolite correlation (25) were also used to predict SM pathway genes.

With the influx of more biochemical and omic data, there is an increasing number of gene properties that can be evaluated for their utility in distinguishing SM/GM genes. Furthermore, the studies published to date have mainly focused on specific SM or GM pathways, raising the question of how SM/GM genes differ globally. This prompted us to examine 10,243 gene properties (referred to as features), including new features and those evaluated in early studies, falling into five categories (gene function, expression/coexpression, gene networks, evolution/conservation, and gene duplication), and evaluate the ability of each feature to distinguish SM genes from GM genes. Earlier studies revealed that the association between features and SM genes is far from absolute (26), and—in most cases—the effect sizes (i.e., the extent to which these specific features can distinguish SM and GM genes) were not reported. To build on these studies, a machine learning approach (21), which jointly considers all five categories of heterogeneous features, was used to distinguish SM and GM genes. This approach led to machine learning models that were used to predict if an *A. thaliana* enzyme gene is likely an SM gene. Furthermore, we examined the properties of enzyme genes in cases where the annotations and predictions differed. Our findings provide a global estimate of SM gene content in the *Arabidopsis thaliana* genome, and the identified features may pave the way for further improvement of the modeling approach.

Results and Discussion

Benchmark SM and GM Genes. Currently, there are two major resources for plant SM and GM gene annotations: Gene Ontology (GO; ref. 27) and AraCyc (28). For SM genes, we started with the 357 genes with the GO term "secondary metabolic process" and 649 enzyme-encoding genes in 129 AraCyc "secondary metabolism" pathways (Dataset S1). Initial GM genes included 2,009 annotated with the GO term "primary metabolic process" and 1,557 enzyme-encoding genes in 490 AraCyc nonsecondary metabolism pathways (Dataset S1). Although 32.4% of GO- and 41.8% of AraCyc-annotated GM genes overlapped, only 35 SM genes (15% of GO- and 8.3% of AraCyc-annotated SM genes) overlapped (Fig. 1*A*). Although this is a significantly higher degree of overlap than expected by chance (*SI Appendix*, Fig. S1 *A* and *B*), it indicates a greater inconsistency in SM annotation criteria than in GM annotation criteria between the GO and AraCyc datasets. Furthermore, 152 and 261 genes were annotated as both SM and GM in GO and AraCyc, respectively. This indicates that although SM and GM genes may have distinct properties, several genes can be both, and their properties may not be distinct. Here we focus on cases that are not ambiguous, but later we delve into this gene set to see if genes involved in both SM and GM pathways can be uniquely classified.

To further assess the differences in AraCyc and GO annotations, we asked whether SM and GM genes annotated based on these two sources have different functional and pathway annotations and Pfam protein domains. We found that GO- and AraCyc-annotated SM genes have substantially different enriched GO categories (Fig. 1*B* and Dataset S1), AraCyc pathways (*SI Appendix*, Fig. S1*C*, and Dataset S1), and protein domains (*SI Appendix*, Fig. S1*D*, and Dataset S2). In contrast to SM genes, GO- and AraCyc-annotated GM genes tend be overrepresented in the same functional categories and pathways (Fig. 1*B*). Considering the above findings, we defined three benchmark sets (Dataset S1). The first (benchmark 1) was defined to include as many annotated SM genes as possible. Here 393 benchmark 1 SM genes were defined as the union of GO and AraCyc SM annotations that have Enzyme Commission (EC) numbers. Similarly, 2,226 benchmark 1 GM genes are from the union of GO and AraCyc primary metabolism gene annotations associated with EC numbers. In the second set (benchmark 2), we used only AraCyc annotations, which were likely better annotated

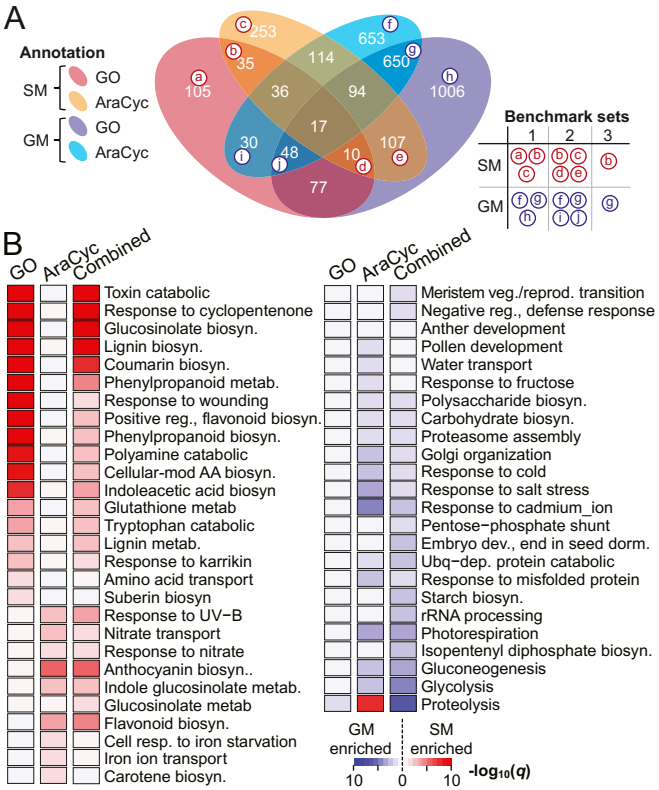

Fig. 1. Gene Ontology and AraCyc annotation of specialized and primary metabolism genes. (*A*) Overlap between Gene Ontology (GO)/AraCyc primary metabolism (PM) and secondary metabolism (SM) gene annotations. The number of genes in an intersection or in a complement set are shown. Three benchmark SM/GM gene sets were defined: benchmark 1 (Union), benchmark 2 (AraCyc), and benchmark 3 (Intersection) (*Methods*). The table to the right shows the genes (labeled with lowercase letters in the Venn diagram) included in each benchmark set. (*B*) GO term enrichment in SM genes (*Left*) and in GM genes (*Right*). The three columns show statistics for GM/SM genes that are GO-annotated, are AraCyc-annotated, or belong to a combined set (union between GO and AraCyc). Rows indicate GO terms. Color represents the *q* value (multiple testing corrected *p* value) of the Fisher's exact test for a GO term enriched in either GM (blue) or SM (red) genes (Dataset S1). White indicates no significant enrichment.

because the focus of AraCyc is on metabolic pathways (SM = 411, GM = 1,306; Fig. 1*A*). In the third set (benchmark 3), we used the intersection between GO and AraCyc annotations (SM = 35, GM = 650; Fig. 1*A*). When we examined which gene feature could distinguish benchmark SM and GM genes (described in the following four sections; Dataset S2), the *p* values from testing >10,000 features were highly correlated among the three benchmark definitions [$R^2 \geq 0.55$ (*SI Appendix*, Fig. S1 *E*–*G*) and all Pearson correlation coefficients (PCCs) ≥0.74 (Dataset S2)]. Therefore, we focus on comparing benchmark 1 (union-based) and benchmark 2 (AraCyc-only) genes, particularly when the conclusions (whether a feature can distinguish between SM and GM genes) were inconsistent.

Differences in Gene Expression and Epigenetic Marks Between SM and GM Genes. A previous study showed that the expression of genes in some SM pathways tends to be more variable than the expression of genes in "essential pathways" (22). To further assess differences in SM and GM gene expression, we examined transcriptome datasets encompassing 25 tissue types (development dataset) and 16 abiotic/biotic stress conditions [stress dataset (*Methods*); for all test *p* values, see Dataset S2]. In addition to confirming that benchmark 2 SM genes tend to have higher expression variability ($P = 0.003$; Fig. 2*A*), we examined 23 additional

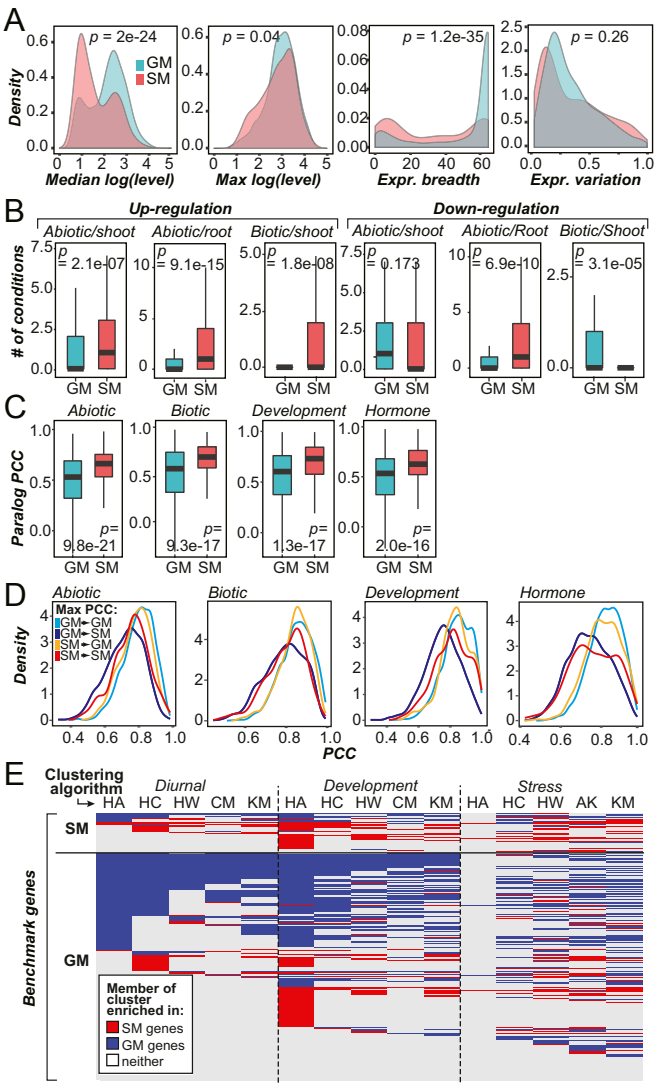

Fig. 2. Differences in expression and coexpression characteristics of benchmark 1 SM and GM genes. (*A*) Distributions of SM (red) and GM (blue) gene expression-related values calculated from the development dataset. Level is microarray intensity. Expression breadth is the number of tissues/developmental stages in which a gene is expressed. Expression variation is median absolute deviation/median. (*B*) Distributions of the number of conditions in which a gene is up- or down-regulated in the abiotic stress (root and shoot) and biotic stress (shoot) datasets. (*C*) Distributions of maximum PCC values between SM or GM genes and their paralogs in four expression datasets. All test statistics from *A* to *C* were generated using Mann–Whitney *U* tests. (*D*) Distributions of maximum PCC between GM–GM (light blue), GM–SM (dark blue), SM–GM (orange), and SM–SM (red) gene pairs using the same expression datasets as in *C*. (*E*) Clustering of SM and GM genes based on their expression patterns in the diurnal development and stress datasets using six algorithms: HA (hierarchical, average linkage), HC (hierarchical, complete linkage), HW (hierarchical, Ward's method), CM (*c*-means), KM (*k*-means), and AK (approximate *k*-means). Row indicates a benchmark SM/GM gene. Blue and red shading indicates that the gene belongs to a cluster with an over-represented number of GM genes and SM genes, respectively, compared with the background ($P < 0.05$, Fisher's exact test).

expression features. We found that SM genes had significantly narrower breadths of expression (Mann–Whitney *U* tests, for all benchmark sets, $P <$ 1e-35; Fig. 2*A*), lower median expression levels ($P =$ e-24; Fig. 2*A*), and lower maximum expression levels ($P = 0.04$; Fig. 2*A*). These findings are consistent with the fact that SM genes have more specialized roles, whereas GM genes are involved in basic cellular functions (3, 6). As expected with the established roles of some specialized metabolites in environmental interactions (e.g., refs. 8 and 29), we found that benchmark 1 SM genes tend to be up-regulated under a higher number of abiotic and biotic stress conditions compared with GM genes (all $P <$ 2e-7; Fig. 2*B*), largely similar to the results based on benchmark 2 ($P = 0.24 \sim$ 1e-8). Relatively fewer SM genes were down-regulated in the shoot under stress compared with GM genes ($P = 0.18 \sim$ 3.1e-5; Fig. 2*B*), likely reflecting a growth–defense tradeoff (30) where GM genes involved in housekeeping functions are down-regulated under stress and SM genes with roles in abiotic and biotic interactions are not. We do not, however, see the same trend in roots. Because CG methylation and histone modification can influence gene expression (31, 32), we compared the numbers of these sites between SM and GM genes. We found that SM genes tend to have a lower degree of gene body CG-methylation than GM genes (Fisher's exact tests, $P <$ 3e-4; Dataset S2). On the other hand, the extent of histone modification did not significantly differ between SM and GM genes for seven of the eight histone marks (*Methods* and *SI Appendix*, Fig. S2*A*).

Previous studies used expression correlation to evaluate how well genes in distinct SM pathways are correlated (20, 21). To see if similar correlation measures could be used to distinguish SM and GM genes, we used maximum PCCs to evaluate expression correlation between each SM/GM gene and its paralogs (Fig. 2*C*) as well as to other SM and GM genes (Fig. 2*D*) in each of four expression datasets (abiotic stress, biotic stress, development, and hormone treatment). We found SM paralogs to have a significantly higher expression correlation than GM paralogs in all four data sets (Mann–Whitney *U* test, all $P < 0.05$; Fig. 2*C*). Because SM genes have undergone more recent expansion than GM genes (2, 4) and the degrees of sequence and expression divergence are positively correlated (33, 34), the higher expression similarities between SM paralogs than between GM paralogs may be partly explained by the more recent timing of SM duplication. We next looked at the maximum expression correlation between each SM gene and other SM genes (SM–SM) or GM genes (SM–GM), as well as between each GM gene and other GM genes (GM–GM) or SM genes (GM–SM). The expression correlations ranked as follows: GM–GM > SM–GM > SM–SM > GM–SM (all benchmark 1 $P < 0.05$, but all benchmark 2 $P > 0.05$ for correlation in the development and biotic stress datasets; Fig. 2*D*). The higher expression correlation for GM–GM compared with SM–SM is likely because GM genes tend to be more broadly expressed and at higher levels than SM genes (Dataset S2). The ratio between expression level variance and mean is higher for genes with lower expression levels, such as SM genes, which contributes to the comparatively lower correlation between these genes. Taken together, our findings indicate that expression correlation features can distinguish SM and GM genes.

Because pathway genes tend to be coexpressed and belong to the same coexpression cluster (20, 21), we next assessed if benchmark 1 SM and GM genes that belong to distinct pathways were members of distinct coexpression modules (Fig. 2*E* and Dataset S2). Among these modules, 99 and 125 contained significantly more SM genes than randomly expected ($\alpha = 0.05$) and are referred to as SM modules. Similarly, 125 GM modules were significantly enriched in GM genes ($P < 0.05$). Therefore, a subset of benchmark GM and SM genes tend to be coexpressed with other GM and SM genes, respectively. However, >50% of SM and GM genes did not belong to SM/GM modules (gray, Fig. 2*E*). In addition, 0.3–14.0% of GM genes were found in SM modules, and 0–32% of SM genes were found in GM modules, depending on the dataset and algorithm (Fig. 2*E*). This pattern reflects the fact that GM genes which function immediately upstream of an SM pathway may be coexpressed with genes in the SM pathway in question. Examples include 208 "junction" genes interfacing GM and SM pathways based on AraCyc annotations (Dataset S2). These findings further highlight the challenges in differentiating SM and GM genes globally using coexpression patterns alone.

Network Properties of SM and GM Genes. SM genes tend to have specialized functions and are involved in one or a few pathways, leading us to hypothesize that SM genes would have fewer connections in biological networks than GM genes. To test this prediction, we first assessed the connectivity among SM genes and among GM genes in a protein–protein interaction network (35) and found that SM genes have a significantly smaller number of physical interactions (mean = 1.25) than GM genes (1.84; benchmark 1, $P = 0.03$; benchmark 2, $P = 3.85e\text{-}8$; *SI Appendix*, Fig. S2*B*). The smaller number of SM gene interactions is not because SM genes have shorter coding regions (SM > GM, all $P = 0.004$; *SI Appendix*, Fig. S2*C*) but is possibly due to the presence of fewer protein domains (SM < GM; benchmark 1, $P = 0.35$; benchmark 2, $P = 4.3e\text{-}6$; *SI Appendix*, Fig. S2*D*). Our finding that significantly fewer protein–protein interactions are known for SM proteins is consistent with SM genes having more specific functions than GM genes (6). It is also possible that there have been more interaction experiments for GM genes or that GM genes tend to function in larger pathways compared with SM genes. Although GM genes tend to have more interactions than SM genes, SM genes with certain domains, such as cytochrome P450, have a higher median number of gene interactions (99.5) compared with their P450 GM counterparts (15.0). Thus, proteins in some domain families may deviate from the general trend we uncovered.

Next, we examined the same relationships using the AraNet functional network (36), which connects genes with likely similar functions through the integration of multiple datasets, including expression and protein–protein interaction datasets. Although the number of protein–protein interactions was significantly higher for GM genes relative to SM genes (*SI Appendix*, Fig. S2*B*; all $P < 0.05$), the differences in network connectivity between GM and SM genes in benchmark 1 were not significant ($P = 0.139$; *SI Appendix*, Fig. S2*E*) but were significant for benchmark 2 genes ($P = 0.027$). AraNet considers multiple gene features including protein interactions, coexpression, shared domains, and homologous genes to construct gene networks, so it is not surprising that this result differs from that for analysis of only protein–protein interactions. These findings suggest that the amount of network connectivity is dependent on the type of network, and this may be useful for distinguishing between SM and GM genes. We should also note that the results from the benchmark 1 and 2 sets are inconsistent, highlighting the effect of the benchmark definition on our analyses. In particular, benchmark 1 p values were higher than those of benchmark 2, despite the fact that benchmark 1 was substantially larger and would have lower P values compared with a smaller dataset with the same effect sizes. This suggests that the AraCyc-only–based benchmark 2 is likely of higher quality.

Evolutionary Rates of SM and GM Genes Based on Within- and Cross-Species Comparisons. SM genes are frequently involved in plant adaptation to variable environments (8, 29, 37). In contrast, GM genes, which are involved in ancient and stable metabolic functions such as photosynthesis, are expected to be more highly conserved (38) and experience stronger negative selection (39, 40). An earlier study found a high degree of genetic variation in glucosinolate genes across *A. thaliana* accessions (21). Here, by comparing SM to GM genes globally, we found that SM genes tend to have higher nucleotide diversities than GM genes ($P = 3.9e\text{-}19$; *SI Appendix*, Fig. S3*B*). In addition, we analyzed 15 evolutionary features based on within-species and across-species comparisons of SM and GM genes. First, we searched for *A. thaliana* SM and GM paralogs as well as homologs across six plant species spanning more than 300 million years of evolution (*Methods*). A significantly higher proportion of SM genes have paralogs than GM genes ($P = 1.2e\text{-}10$; *SI Appendix*, Fig. S3*A*). However, consistently fewer SM genes (14.8–54%) have homologs across species than GM genes (27–76%) (all $P < 2e\text{-}4$; *SI Appendix*, Fig. S3*A*). In addition, as expected for lineage-specific functions, only 0.94% of SM genes have homologs in core eukaryotic genomes (41) compared with 14.7% of GM genes (*SI Appendix*, Fig. S3*A*).

Finally, we determined the timing of GM and SM duplications over the course of land plant evolution using sequence similarity to determine the most recent duplication point (*Methods*). We found that 75% of SM genes were products of duplication events after the divergence between the *A. thaliana* and *Brassica rapa* lineages compared with only 40% of GM genes (Fig. 3*A*), indicating that SM genes tend to be more recently duplicated relative to GM genes. Additionally, 25% of SM genes were duplicated after the *A. thaliana–Arabidopsis lyrata* split, compared with only 7% of GM genes (Fig. 3*A*). Thus, SM genes have higher duplication rates but do not persist in the long run, leading to the observation of fewer homologs across species.

We also found that SM genes and their homologs had significantly higher nonsynonymous (*dN*) to synonymous (*dS*) substitution

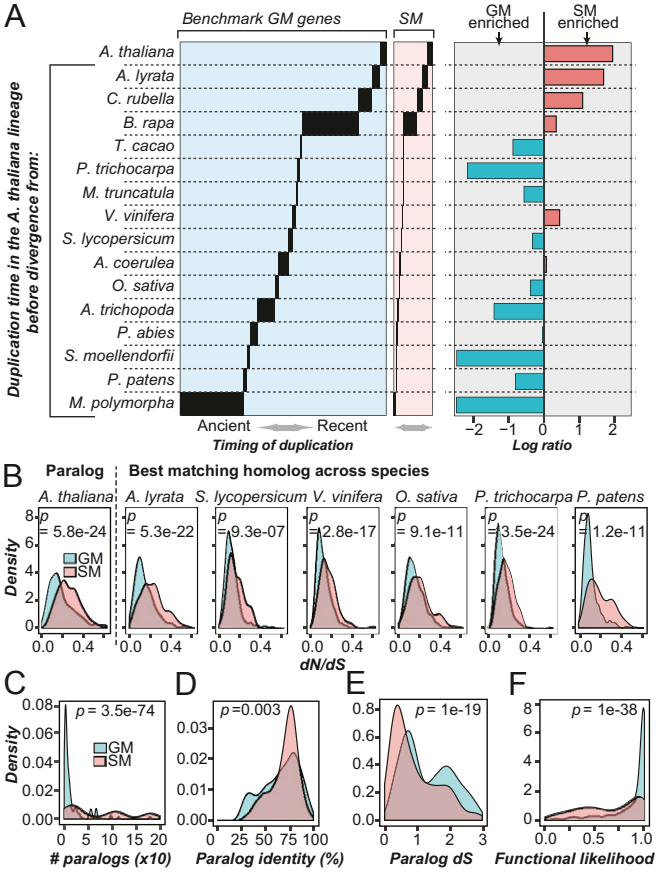

Fig. 3. Differences in the duplication timing, degree of selective pressure, paralog-related features, and functional likelihood between benchmark 1 SM and GM genes. (*A*) The distribution of duplication time points (*y* axis) for each GM/SM gene (*x* axis). (*Left* and *Middle*) A black line indicates that the GM (*Left*) or SM (*Middle*) gene in question likely duplicated before the divergence between the *A. thaliana* lineage and the species lineage to the left of the black line. Species order is based on the time of divergence from *A. thaliana*. (*Right*) Each bar represents the log2 ratio (*x* axis) between the proportions of SM and GM genes duplicated at each duplication time point (*y* axis). For full species names, see *Methods*. (*B–F*) Density plots showing SM (pink) and GM (blue) gene feature distributions. Test statistics were generated using Mann–Whitney *U* tests. (*B*) Median nonsynonymous substitution rate/synonymous substitution rate (*dN*/*dS*) values between *A. thaliana* SM/GM genes and their *A. thaliana* paralogs or best matching homologs in six other species, arranged based on the time of divergence from *A. thaliana*. (*C*) The number of *A. thaliana* paralogs of SM or GM genes. (*D*) The maximum percent identity of an SM or GM gene to its paralogs. (*E*) The *dS* distribution between each SM or GM gene and its paralog. (*F*) The functional likelihood ranging from 0 to 1, which indicates the likelihood that a gene is under selection.

rate ratios (all $P < 1e-06$; Fig. 3*B*) compared with GM genes. Together with other measures of selection (*SI Appendix*, Fig. S3 *C* and *D*), both within- and cross-species comparisons suggest that SM genes are under weaker negative selection relative to GM genes. One reason for this pattern may be that these SM genes initially experienced positive selection (higher rate than GM) followed by negative selection (similar to GM). This would result in SM genes having a higher rate of evolution than GM genes, with the appearance of weaker negative selection. Another possible reason for this pattern is that some of these SM genes may have experienced strong negative selection (similar to GM) but are now neutrally evolving. This may be because the selective agent (e.g., a particular environmental factor) previously contributing to the selection no longer exists. This is consistent with the roles of SM genes mostly in the production of metabolites important for tolerance to rapidly changing abiotic stress conditions and defense against biotic agents (6).

Duplication Mechanisms and Genomic Clustering of SM and GM Genes. Gene duplication mechanism, such as whole genome duplication (WGD), tandem duplication, and dispersed duplication, may affect subsequent functional divergence and ultimately influence whether a duplicate is under selection and retained (1). For example, genes in a few SM pathways, such as aliphatic glucosinolate biosynthesis, tend to be tandemly duplicated and have a higher degree of expression variation (22). To assess if SM and GM genes differ in their post-WGD retention rate, we compared the number of GM and SM WGD duplicates in the *A. thaliana* lineage. Although two different glucosinolate pathways arose in the α WGD event ∼50 Mya (15), these two pathways do not lead to a significantly higher number of SM WGD duplicates compared with the number of GM WGD duplicates. This indicates that SM genes from multiple SM pathways (not just those involved in glucosinolate metabolism) are not more likely to be derived from WGDs than GM genes (benchmark 1 $P = 0.1$, and benchmark 2 $P = 0.85$; *SI Appendix*, Fig. S4*A*). This suggests that the likelihood of long-term retention of SM and GM WGD duplicates does not appear to differ significantly. In contrast, significantly more SM genes tend to be tandem duplicates than GM genes ($P < 2e-43$; *SI Appendix*, Fig. S4*A*). Genes involved in response to the environment are more likely to be tandem duplicates (2, 42), and tandem duplication potentially allows for rapid evolution of SM gene families that are subject to selection in variable environments.

The numbers of paralogs and pseudogenes were used as measures of the degree of SM and GM gene gains and losses, respectively. Our analysis revealed that SM genes tend to have more paralogs ($P < 3e-72$; Fig. 3*C*), higher sequence similarities to their paralogs (benchmark 1, $P = 3e-3$, and benchmark 2, $P = 0.3$; Fig. 3*D*), and lower synonymous substitution rates (dS) ($P < 2e-19$; Fig. 3*E*) compared with GM genes. Furthermore, a higher percentage of SM genes duplicated since *A. thaliana* diverged from *A. lyrata* ($P < 4e-8$; *SI Appendix*, Fig. S4*B*), and SM genes tended not to be found in single copies ($P < 1e-3$; *SI Appendix*, Fig. S4*C*). These findings all point to more recent expansion of SM gene families. We also compared the functional likelihood, which is a measure of how likely it is that a gene is functional and, thus, under selection (39), between SM genes, GM genes, and pseudogenes. The functional likelihoods of SM genes are significantly lower than those of GM genes but higher than those of pseudogenes (ANOVA, Tukey's test, $P < 2e-16$; Fig. 3*F* and *SI Appendix*, Fig. S4*E*). Genes under strong negative selection have high functional likelihoods that are close to 1, whereas pseudogenes tend to have values close to 0 (39). In addition, most pseudogenes are eventually removed from the genome (43) and tend not to be under selection (44). Our finding that SM genes tend to have lower functional likelihood is consistent with the hypothesis that some SM genes are under weaker selection and may be in the process of becoming pseudogenes. The proportion of pseudogene paralogs for SM genes (between benchmarks, 9.8–11.1%) compared with GM genes (6.1–6.5%) is not significant overall ($P = 0.04 \sim 0.2$; *SI Appendix*, Fig. S4*D*). Considering that SM genes tend not to have cross-species homologs (*SI Appendix*, Fig. S3*A*), this finding suggests that pseudogenes are too short lived to be adequate indicators of gene loss.

SM and GM genes that function in the same pathway are sometimes found in genomic clusters (21, 45–47), and we used two approaches to compare the occurrence of SM and GM genes in close physical proximity. In the first approach, we asked whether SM and GM genes tend to be located near other SM and GM genes, respectively, regardless of whether the neighboring genes are paralogous or not. We found that SM genes cluster near other SM genes (benchmark 1, $P = 9.5e-121$, and benchmark 2, $P = 0.02$; *SI Appendix*, Fig. S4*F*) and GM genes tend to be close to GM genes ($P < 2e-5$; *SI Appendix*, Fig. S4*G*). It is surprising that the P values for SM clustering differ so greatly between benchmark sets. This may indicate that AraCyc annotation (benchmark 2) is of higher quality. In the second approach, we defined metabolic clusters identified using Plant Cluster Finder (21), but the identified clusters were not enriched in either SM or GM genes (*SI Appendix*, Fig. S4*H*). Taken together, SM genes are more likely to be tandemly duplicated and tend to belong to large gene families. Our findings provide genome-wide confirmation of earlier studies (e.g., refs. 2, 15, and 22) that focused on a relatively small number of SM genes or pathways. These characteristics may be useful features in distinguishing SM and GM genes.

Machine Learning Model for Predicting SM and GM Genes. In total, we examined 10,243 features (summarized in Dataset S2) that differ widely in their ability to distinguish benchmark SM and GM genes. For example, the best performing single feature—gene family size—led to a model with an area under receiver operating characteristic curve (AuROC) of 0.8. An AuROC of 0.5 indicates the performance of random guesses, and a value of 1 indicates perfect predictions. However, using this high-performing feature alone as the predictor resulted in a 43% false positive rate (FPR) and a 58% false negative rate (FNR). In addition, the majority of the features are not particularly informative (Dataset S3), because the average AuROC for single feature-based models was extremely low (0.5) with an average FPR of 89%. These findings indicate that SM and GM genes are highly heterogeneous and cannot be distinguished with high accuracy using single features. To remedy this, we next integrated all 10,243 features, regardless of whether they were significantly different between SM and GM genes or not, to build machine-learning models for predicting SM and GM genes. We used machine learning because it allowed us to build an integrated model where multiple features were considered simultaneously. Integrated models offer better predictive power than individual features by lowering FNR and FPR.

Two machine learning algorithms, Support Vector Machine and Random Forest, were used to build predictive models using all three benchmark datasets (*Methods*; Fig. 4*A*; *SI Appendix*, Fig. S6; and Dataset S3). The best performing SM gene prediction model was based on benchmark 2 (AraCyc only) and Random Forest (AuROC = 0.87, FPR = 29.4%, FNR = 14.8%; Fig. 4*A*). Randomizing SM/GM labels but maintaining the same feature values associated with the benchmark genes as the initial model resulted in AuROCs = 0.51 ∼ 0.57, as expected for random guesses (Dataset S3). Note that the performance measures reported above were based on models built with a 10-fold cross-validation scheme where 90% of the data were used for training the models and 10% for testing them. Based on the prediction outcomes, each gene was given an SM score ranging from 0 to 1 indicating the likelihood that the gene is an SM gene. Based on a threshold SM score defined by minimizing false predictions (*Methods*), 85.6% of the training SM genes (Fig. 4*B*) and 73.1% of the training GM genes were correctly predicted (Fig. 4*B*), which reflects an improvement over the individual feature-based, naïve models.

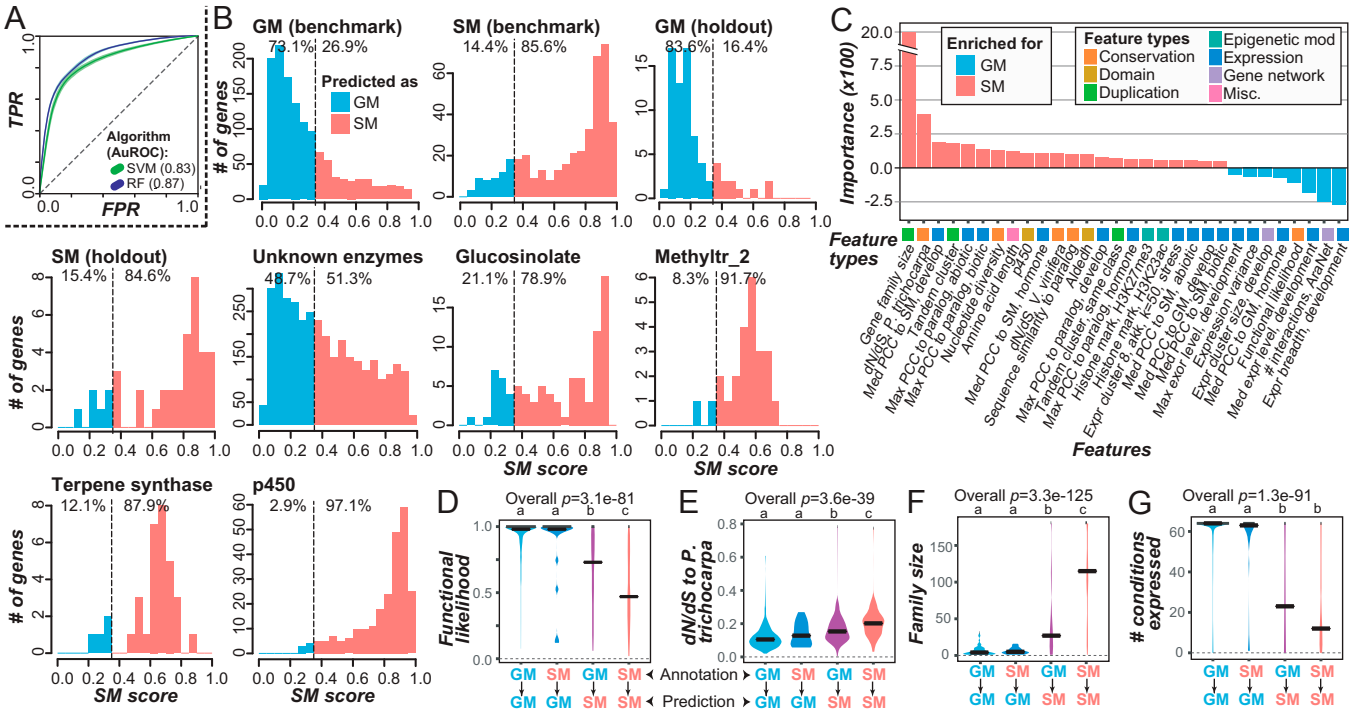

Fig. 4. SM gene prediction model performance based on benchmark 2. (*A*) AuROC curves of binary SM/GM prediction models built with Support Vector Machine (SVM) and Random Forest (RF) algorithms. FPR, false positive rate; TPR, true positive rate. (*B*) SM score distributions for benchmark GM, benchmark SM, holdout SM (not included in models), unannotated enzymes from AraCyc, glucosinolate pathway, p450, terpene synthase, and methyltransferase 2 (methyltr_2) domain-containing genes. Dotted line indicates SM score threshold (*Methods*). Red and blue shading indicate genes predicted to be SM and GM genes, respectively. (*C*) The most important features for SM (red) and GM (blue) gene predictions. (*D*–*G*) Distributions of the values of representative, predictive features for correctly and incorrectly predicted SM and GM genes. Black horizontal bar indicates median. Overall *P* values are from Kruskal–Wallis tests used to evaluate differences among classes. The Dunn post hoc test was used to test differences between classes (Dataset S3). (*D*) Functional likelihood. (*E*) *dN/dS* between *A. thaliana* and *P. trichocarpa* homologs. (*F*) Sizes of the gene families the four categories of genes belong to. (*G*) Expression breadth in the development dataset.

Features Important for SM Gene Prediction and Model Application to Unannotated Enzyme Genes. In addition to the SM score, the machine learning result included a list of feature importance values, where features with more positive values are more informative for predicting SM genes. In contrast, more negative feature weights are more informative for predicting GM genes (Fig. 4*C* and Dataset S3). Based on the AraCyc-only (benchmark 2) model, the most informative features for predicting SM genes included specific protein domains as well as multiple gene duplication-related features, such as duplication mechanism (higher degree of tandem duplication), gene family expansion (larger family size), and higher degrees of correlation in expression between an SM gene and other SM genes or its paralogs (Fig. 4*C*). In addition, higher evolutionary rates were among the most informative for predicting *A. thaliana* SM genes based on comparison of an SM gene to its *Populus trichocarpa* and *Vitis vinifera* homologs but not to homologs from more closely related species. This pattern may reflect the fact that at these time points (postdivergence between *A. thaliana* and the *P. trichocarpa* or *V. vinifera* lineages) a number of SM genes experienced accelerated, potentially positive, selection that contributed to the diversification of major SM pathways. In contrast, wider expression breadth, measured using the development expression dataset, and higher connectivity in gene networks were among the most important features for predicting GM genes, indicating the more generalizable functions of GM genes and the tendency to interact with a greater number of genes/gene products relative to SM genes. Finally, specific histone marks as well as hierarchical, *k*-means, and approximate *k*-means coexpression clusters based on the stress, diurnal, and development datasets were informative for predicting both SM and GM genes (Dataset S2). While earlier studies established that genes belonging to a particular SM pathway tend to be coexpressed (20, 25), our findings demonstrate that there are global differences in expression patterns and properties between SM and GM genes.

With the accuracy of the SM gene prediction models assessed through cross-validation and prominent features identified, we next applied these machine learning models to make predictions for 2,085 known enzymatic genes (with an EC number) not annotated to be SM or GM genes (Dataset S1). Of these genes, 58% (1,220 genes) were predicted to be SM genes. We took three approaches to assess the accuracy of these SM and GM gene predictions. First, we intentionally held out 10% of both known SM and GM genes (Fig. 4*B* and Dataset S1) from any model training. Upon application of the machine learning model, 84 and 85% of withheld GM and SM genes were correctly predicted, respectively, indicating that the model has an 84% true positive rate (or 16% FNR). Second, we tested how well genes in well-known SM pathways involved in glucosinolate biosynthesis (38, 39) could be predicted. To do this we built a model using the benchmark SM and GM genes but excluding genes from glucosinolate biosynthetic pathways (*Methods*) (Fig. 4*B* and Dataset S1). When applying this model to glucosinolate genes, 79% of known glucosinolate pathway genes were correctly predicted as SM genes. The FNR was 16% overall, which is much better than the 58% FNR when using the single best feature, gene family size.

Finally, methyltransferase, terpene synthase, and cytochrome P450 families were identified based on their respective protein domains (*Methods*) and analyzed to test model performance within a specific family (Fig. 4*B* and Dataset S1). These families were chosen because they tend to be associated with SM. To this end, we built three models using our benchmark sets, excluding holdout

genes from the families we planned to predict. Upon applying this model to each enzyme family, 97% of P450, 88% of terpene synthase, and 92% of methyltransferase genes were predicted as SM genes (Fig. 4*B*). Thus, these models predicted the majority of holdout genes with known SM functions, glucosinolate pathway genes, and genes in enzyme families whose members predominantly play roles in SM pathways. In summary, our models allowed assessment of the relative importance of features in distinguishing SM and GM genes, as well as provided predictions for 1,220 SM genes among enzyme genes with no known SM/GM designation. In addition, our findings indicate that our models and this general approach are valuable for predicting unknown enzymes.

Characteristics of Mispredicted Genes. Although our SM prediction model performed well, 354 (27.1%) AraCyc annotated GM genes were mispredicted as SM genes. In addition, 60 (15.3%) AraCyc annotated SM genes were mispredicted as GM genes. To assess the properties of mispredicted SM/GM genes, we determined how the values of a subset of the most informative features (Fig. 4*C* and Dataset S3) differed between four gene classes defined based on the consistency between the gene annotation and the benchmark 2 (AraCyc only)-based model prediction. These four classes included (*i*) annotated GM predicted as GM (GM [annotation] → GM [prediction]), (*ii*) annotated SM predicted as SM (SM → SM), (*iii*) annotated GM predicted as SM (GM → SM), and (*iv*) annotated SM predicted as GM (SM → GM). Genes in the mispredicted classes (3 and 4) tend to have feature values between those of genes in correctly predicted classes (1 and 2). For example, the median values of the functional likelihood among these four gene classes follow the order GM → GM > SM → GM > GM → SM > SM → SM (Fig. 4*D*). The opposite pattern (SM → SM has the highest value) was observed for *dN*/*dS* values (Fig. 4*E*), gene family size (Fig. 4*F*), the number of conditions expressed (Fig. 4*G*), and values for other gene features we examined (*SI Appendix*, Fig. S5 *A*–*J*). Thus, in the SM → GM mispredicted class, the annotated SM genes in fact possess multiple properties that are more similar to those of GM genes and vice versa, but no single feature can fully explain why these genes were mispredicted.

These observations suggest that some of the mispredicted benchmark genes (Fig. 4*B*) may in fact be misannotated, or alternatively, they may point to a deficiency in our model (addressed in the next section). To assess how many of the mispredictions are due to misannotation, we collated information from 25 genes with predictions (from the benchmark 2-based model) matching the AraCyc annotations (GM → GM = 4, SM → SM = 21), and for 32 genes with predictions that were not consistent with their AraCyc annotations (SM → GM = 20, GM → SM = 12) (*SI Appendix*, *SI Text*, and Dataset S1). We focused on genes in the P450/terpene synthase families because there is substantial biochemical and functional information available for these genes (3, 48, 49). For mispredicted genes, which were manually examined, five (42%) genes in the GM → SM class had supporting SM evidence (Dataset S1). In addition, 16 (80%) genes in the SM → GM class have supporting GM evidence (Dataset S1). These findings indicate that a subset of these genes (66%) are mispredicted due to misannotation, not due to prediction errors.

For the benchmark 1 set, which is based on the union between AraCyc and GO annotations, a similar percentage of the mispredicted genes [5 of 11 GM → SM (45%) genes examined] were likely misannotated (Dataset S1). This is consistent with our finding that some SM genes enriched in AraCyc pathways and GO terms—such as carotene, leucine, suberin, and wax biosynthesis—are found across all major land plant lineages and should be considered GM genes (Fig. 1*B* and *SI Appendix*, Fig. S1*C*). It is also possible that some of the erroneous annotations are based on in vitro biochemical activity and/or sequence similarities alone, criteria that may not accurately represent their in vivo functions. We should note that genes that were manually examined were mostly from the P450 and terpene synthase families. More enzyme families should be evaluated to obtain a more complete picture of the reasons behind inconsistent annotation and prediction. Together with the finding that nearly all (24/25) benchmark 2 genes with consistent annotations and predictions had biochemical evidence supporting their SM or GM classification (*SI Appendix*, *SI Text*), these results further demonstrate the feasibility of using the model prediction outcome to prioritize future experiments to determine the in planta role of SM or GM genes, including those that may be misannotated or have functions in addition to their annotated activities.

Effect of Dual-Annotated Genes on Model Performance. The number of genes mispredicted with our model that were not misannotated according to our manual examination (*SI Appendix*, *SI Text*, and Dataset S1) indicate that our model can be further improved. Our original model focused on distinguishing SM and GM genes as binary classes, but genes with both SM and GM functions were excluded. However, there are 261 genes (Fig. 1*A*) annotated as belonging to both SM and GM pathways in AraCyc (dual-annotated or DA genes; Fig. 5*A*). We thus explored the possibility that DA genes have properties distinct from SM or GM genes and should be considered a distinct class. We first compared the SM scores between SM, GM, and DA genes based on our AraCyc-only binary model. If DA genes belong to a distinct class that is neither SM nor GM, the SM scores of DA genes should have a unimodal distribution with a median close to 0.5. Contrary to this expectation, the SM score distribution of DA genes is bimodal, where some DA genes resemble SM genes and others resemble GM genes (Fig. 5*B*). Thus, based on a GM vs. SM binary model, DA genes do not appear to belong to a distinct class. These findings raise the question whether the dual annotation is valid.

To assess whether our inability to distinguish DA genes from SM/GM genes is because the binary model is inadequate, we built a multiclass model assuming SM, GM, and DA genes as three distinct classes and plotted the SM scores for each class in a ternary plot (Fig. 5 *C*–*F*). If the three classes of genes can be perfectly separated, then the highest gene density areas will be toward different corners of the ternary plots. Although the GM/SM/DA model has an F1 score of 0.51 (higher than the F1 of 0.33 for a random model) and an accuracy of 0.53, the inclusion of DA genes as a third class significantly diminished the ability of the model to separate SM (Fig. 5*C*) and GM (Fig. 5*D*) genes. Note that SM and GM genes are not well separated in the ternary plots (Fig. 5 *C* and *D*), but in the binary model, their SM score distributions are highly distinct (Fig. 5*B*). In addition, the DA gene distribution in the ternary plot overlapped with the distributions of both SM and GM genes (Fig. 5*E*), consistent with the bimodal SM score distribution observed among DA genes. Thus, the DA genes belong to two subclasses, with each subclass resembling SM or GM genes, again raising the question whether the dual annotations in AraCyc are valid. Curiously, GM genes separate into two populations in the GM/SM/DA model where one population is located toward the GM corner of the ternary plot (arrow g1) and the second population (arrow g2) overlaps with areas of high SM (arrow s) and DA (arrow d) gene density (Fig. 5*C*). Therefore, although this three-class model does not separate SM and GM genes well, it raises the question of how the two GM gene populations (g1/g2 peaks) differ and should be further examined.

Consideration of Junction Genes in Predictive Model Building. Another potential way to improve our model is to consider metabolic network topologies. We hypothesized that SM and GM genes closer to pathway junctions (*Methods* and Fig. 5*A*) are more likely to be mispredicted. We identified junction reactions connecting 15 GM (upstream) and 20 SM (downstream) pathways. The 212 genes encoding enzymes responsible for junction reactions were referred to as junction (JC) genes. By further classifying JC genes based on the connectivity of their associated reactions, four topological subclasses of junction genes were defined: 1 → J → 1, junction reactions, each connected with one reaction upstream and one reaction downstream; n → J → 1, multiple upstream reactions but only one downstream reaction; 1 → J → n, one

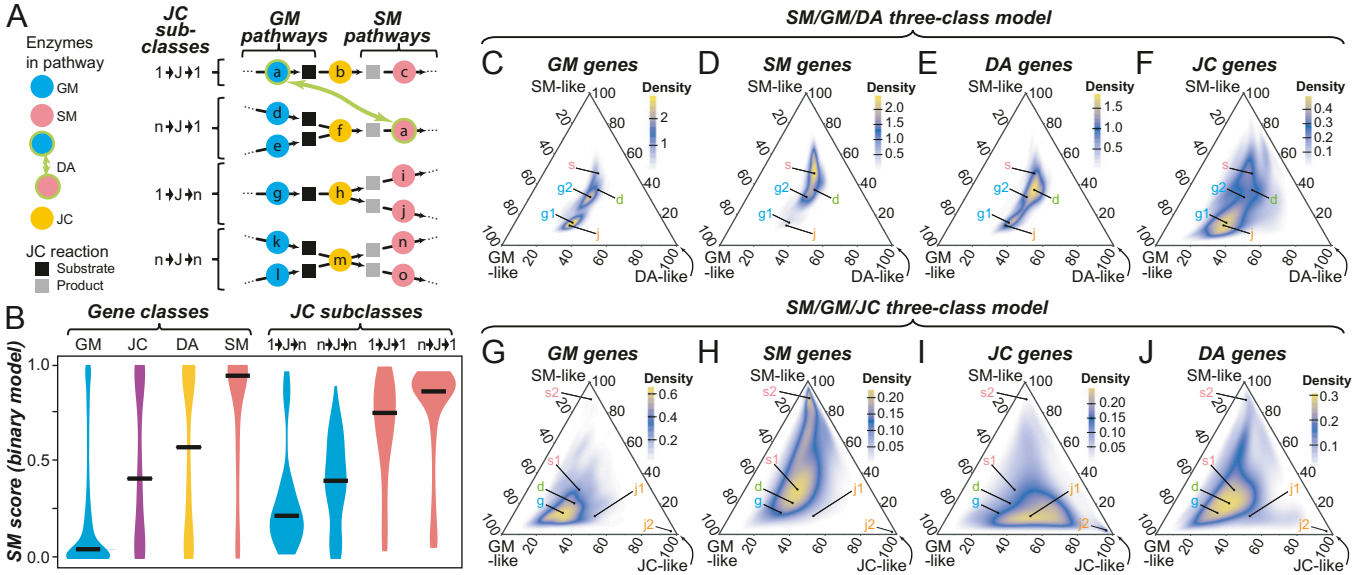

Fig. 5. Three-class models for classifying SM/GM/DA and SM/GM/JC genes. (*A*) Definition of DA (dual annotation) and JC (junction) genes. For JC genes, four subclasses were defined based on the degree of connectivity, defined as the number of connecting reactions in the metabolic network based on AraCyc annotations. The letters a–o indicate SM/GM enzymes that are annotated as GM (blue), SM (red), or DA (green outline) or are defined as JC (orange). JC reaction substrates and products are in black and gray, respectively. (*B*) Distributions of SM scores based on the binary model built using benchmark 2 data for GM, SM, DA, JC (all), and JC subclass genes. (*C–F*) Ternary plots showing the SM/GM/DA model-based score distributions for GM (*C*), SM (*D*), DA (*E*), and JC (*F*) genes. The g (blue), s (red), d (green), and j (orange) labels indicate the peak gene density areas (brighter yellow) occupied by GM, SM, DA, and JC genes, respectively. (*G–J*) Ternary plots showing the SM/GM/JC model-based score distributions for GM (*G*), SM (*H*), JC (*I*), and DA (*J*) genes. The color scheme follows that in *C–F*.

upstream and multiple downstream reactions; and n → J → n, multiple upstream and downstream reactions (Fig. 5*A*). Although junction genes as a whole also have a bimodal SM score distribution similar to that of DA genes (JC all; Fig. 5*B*), the score distributions were distinct among the four topological subclasses, indicating that network topology is a distinguishing characteristic between SM and GM genes. Considering that products of GM pathways serve as substrates for many other pathways, it is expected that GM genes functioning in junction reactions would be connected to multiple downstream pathways. Consistent with this, JC genes in the n → J → n and 1 → J → n subclasses where $n > 1$ tend to be more similar to GM genes (Fig. 5*B*). In contrast, SM enzymes are more likely involved in incorporating substrates from multiple reactions and serve as the committed step for producing specialized metabolites with an expected n → J → 1 topology. In addition, a typical SM pathway mostly contains a series of nonbranching reactions that lead to specialized metabolite products and is also expected to have a 1 → j → 1 topology. Consistent with these expectations, JC genes in the n → J → 1 and 1 → J → 1 subclasses are the most similar to SM genes (Fig. 5*B*).

The GM/SM/JC three-class model separated SM and GM genes significantly better (F1 score = 0.65, accuracy = 0.65; Fig. 5 *G–J*) than the GM/SM/DA model (Fig. 5 *C–F*), indicating that junction genes have unique characteristics and that some genes intersecting annotated SM and GM pathways can be considered a separate class. In addition, the four topological subclasses of JC genes are located in different areas in the ternary plots for the SM/GM/DA (*SI Appendix*, Fig. S7*A*) and GM/SM/JC (*SI Appendix*, Fig. S7*B*) models. We should emphasize that JC genes were defined based on a network constructed using AraCyc pathway annotations where the criteria for defining pathway boundaries may differ between research groups and/or annotators. Despite this, the GM/SM/JC model predictions demonstrate that JC genes are by and large distinct from GM/SM genes. Although we cannot be certain which JC genes were key enzymes in the committed steps entering SM pathways, the JC genes in the n → J → 1 subclass are clearly a class of their own with most genes at the JC-like corner (*SI Appendix*, Fig. S7*B*). Taken together, these findings demonstrate that further categorization of SM and GM genes based on biologically motivated criteria, such as network topology, could help further distinguish different types of SM or GM genes leading to modest improvement of our models. In addition, the binary classification of SM and GM genes, while meaningful, can be an oversimplification. Finally, the consideration of additional topological characteristics (e.g., pathway depth and terminal reaction) and additional biochemical features (e.g., substrate and product identities) may lead to further improvements in SM and GM predictions.

Conclusions

Machine learning models built using genomic features show considerable promise in predicting the functions of unclassified or unannotated genes (21, 39). Before establishing such models for predicting SM and GM genes, we first explored how SM and GM genes in *A. thaliana* differ in >10,000 conservation, protein domain, duplication, epigenetic, expression, and gene network-based features. Most of these features have not been examined by other studies contrasting SM and GM genes. We demonstrated that machine learning models in which these features are integrated to predict SM and GM genes perform well based on cross-validation performed using three benchmark datasets, three predominantly SM gene families, glucosinolate biosynthesis pathway genes, and 39 AraCyc-annotated SM genes that were deliberately withheld from the model building process. Focusing on the AraCyc-only benchmark (benchmark 2), although 380 individual features significantly differed between SM and GM genes, the effect sizes are small, and any individual feature does a poor job of distinguishing SM and GM genes compared with the machine learning models. In addition, machine learning models allow the global prediction of SM and GM genes in a plant genome. Based on the SM scores derived from these models, candidate SM genes can be prioritized for further experimental studies.

Although the binary SM/GM gene prediction model performed well, the FPR and FNR were substantial at 28% and 19%, respectively. Through closer examination of experimental evidence for 10 genes annotated as GM genes but predicted as SM genes, we found ~50% had evidence supporting classification as SM genes, indicating that a subset of the mispredictions is

likely due to misannotation. Thus, in addition to predicting likely GM/SM functions of unannotated enzymes, our models can be used to pinpoint potentially misannotated GM/SM genes. Mispredictions can be avoided by further improving the model in two areas: the classes defined and the features used. Classifying enzyme genes as GM and SM may be an oversimplification. By building two three-class models (GM/SM/JC and GM/SM/DA), we found that SM and GM genes could be further categorized based on the metabolic network topology and, to a lesser extent, based on their dual-annotated roles in both SM and GM pathways. Future studies distinguishing genes at the pathway level can be carried out using similar multiclass modeling methods. Additional features that can distinguish SM and GM genes may also be needed to further improve model performance. One possibility is to incorporate topological information as features. Another possibility is to examine feature combinations (e.g., combining an expression and a duplication feature linearly or nonlinearly) using approaches such as deep learning.

In summary, we have conducted a global analysis of gene features that are useful to distinguish SM and GM genes. We also established well-performing machine learning models that provide a global estimate of the SM gene content within a plant genome. The great majority of the predicted SM genes have not been assigned to pathways, highlighting the important next step of combining the GM/SM prediction scheme described here with approaches for pathway discovery and assignment. Considering that the most important features are related to gene duplication, evolutionary rate, and gene expression and that these types of data are readily available for an ever-expanding number of plant species, the machine learning workflow we have developed can be readily applied to any other species for predicting SM genes or, more generally, gene functions. Nonetheless, there is room for further improvement. Our prediction model serves as a baseline model for future studies incorporating additional features and algorithms that are anticipated to further improve the accuracy of predictions.

Methods

Specialized and General Metabolism Gene Annotation and Enrichment Analysis. Gene sets were identified based on GO (ref. 27; www.geneontology.org/ontology/go.obo), and/or AraCyc (ref. 28; https://www.plantcyc.org/) annotations but not MapMan (50). We did not analyze MapMan annotations because all GO and AraCyc SM genes, which include a large number of well-known SM examples, were annotated as GM in MapMan, raising questions about the utility of MapMan SM/GM designations. GO annotations for *A. thaliana* were downloaded from The *Arabidopsis* Information Resource (TAIR) (51), and genes annotated with the secondary metabolism term (GO:0019748) and primary metabolism term (GO:0044238) were selected as potential SM genes and GM genes, respectively. Genes that were associated with more specialized primary and secondary metabolism child GO terms were also classified as GM and SM genes, respectively. Only genes annotated with either SM or PM terms, but not both, were included in the analysis, and only those with experimental evidence codes IDA, IEP, IGI, IPI, and/or IMP were included. For AraCyc genes, the v.15 pathway annotations were retrieved from the Plant Metabolic Network database (https://www.plantcyc.org/) (28). Potential SM genes were those associated with secondary metabolites biosynthesis pathways. Potential GM genes were those found in nonsecondary metabolite biosynthesis pathways. In addition, genes without experimental evidence in AraCyc (EV-EXP) were not included in the benchmark. Some genes were annotated in both SM and non-SM pathways and were defined as dual-annotated (DA) genes, not as SM or GM. Potential SM and GM genes from GO or AraCyc were required to have an EC number annotation from AraCyc or from Pfam v.30 (pfam.xfam.org/) (52). Five benchmark gene sets were defined. In addition, glucosinolate pathway genes were also defined to test model performance. The criteria for defining benchmarks and glucosinolate pathway genes are detailed in *SI Appendix, SI Methods*. Terpene synthase, P450, and methyltransferase genes were identified from *A. thaliana* annotated protein sequences using the following domain matches from Pfam: terpene_synth, p450 and methyltr_2. Details of gene set enrichment analysis are available in *SI Appendix, SI Methods*.

Expression Dataset Processing and Coexpression and Gene Network Analysis. Expression datasets were downloaded from TAIR. Target datasets included plant development (53), biotic stress (54), abiotic stress (54, 55), hormone treatment (56), and diurnal expression (57). Genes that were considered significantly expressed relative to the background in the development expression dataset were those with a $\log_2$ microarray hybridization intensity value of ≥ 4 (the cutoff value is based on our earlier study, ref. 39). The median and maximum expression levels and expression variation and breadth across the developmental expression dataset were calculated as previously described (39). Differentially expressed genes under biotic stress, abiotic stress, and hormone treatments were defined as those that had an absolute $\log_2$ fold change ≥ 1 and adjusted $P < 0.05$ following analysis using the affy and limma packages in R (58, 59). For each gene, the number of conditions in which the gene in question was significantly differentially regulated was also calculated. This resulted in 16 expression values that were used as model features (Dataset S2).

For each expression dataset (development, abiotic, biotic, and hormone), PCCs were calculated between each gene and genes in the same paralogous cluster as defined by ORTHOMCL v1.4 (60). For the gene in question, the maximum PCC < 1 for genes in the paralog cluster was used as the PCC value. In addition to examining expression correlation, coexpressed genes in the biotic stress, abiotic stress, diurnal, and developmental datasets were classified into coexpression clusters using *K*-means, approximate kernel *K*-means, c-means, and hierarchical clustering algorithms as described in our earlier study (26) resulting in 5,303 binary features. For *K*-means–related analyses, the within-cluster sum of squares was plotted against the number of clusters, and *K* was chosen based on the number of clusters at the elbow or bend of the plot. Gene clusters that were significantly enriched in SM or GM genes were identified using Fisher's exact tests (adjusted $P < 0.05$). The number of AraNet gene network interactions (ref. 36; www.functionalnet.org/aranet/), number of protein interactions (35), domain number, and amino acid length were calculated in our earlier study (39). There were 23 model features related to PCC values, significant cluster membership, and gene network data (Dataset S2).

Conservation, Duplication, Methylation, Histone Modification, and Genome Location Related Features. Nonsynonymous (*dN*)/synonymous (*dS*) substitution rates between plant homologs, core eukaryotic gene status, nucleotide diversity data, Fay and Wu's H, and MacDonald–Kreitman test statistics were the same as used in our earlier studies (39, 61, 62). Details on determining the timing of duplication of an *A. thaliana* gene are available in *SI Appendix, SI Methods*. Pseudogenes were defined using a published pipeline (53). The lethal gene scores, which represent the relative likelihood that a mutation in a gene is lethal, and additional gene duplication-related features, including gene family size, rates of synonymous substitutions, α and β/γ whole genome duplication status, and tandem duplication status (Dataset S2), were obtained from (39). CG methylation and $\log_2$ fold change of histone marks relative to background were taken from ref. 39 (detailed in *SI Appendix, SI Methods*). Three approaches were used to evaluate the degree of metabolic gene clustering (*SI Appendix, SI Methods*).

Machine Learning Classification of SM and GM Genes. The prediction models were built based on 10,243 features using the Random Forest (RF) and Support Vector Machine (SVM) algorithms implemented using the Python package sci-kit learn (63). To build binary machine learning models, we used three benchmark sets (benchmarks 1, 2, and 3). For each benchmark set, SM and GM genes were first divided into a modeling set (90%) and a holdout set for independent validation (10%). Because there were significantly more GM genes than SM genes, 100 balanced datasets were constructed by randomly selecting GM genes equal to the number of SM genes in each balanced set. Additionally, 10-fold cross validation was performed for 100 random draws of a balanced dataset for each machine learning run, and grid searches were performed to obtain the best performing parameters for each model. Details for the performance measure are available in *SI Appendix, SI Methods*. A confidence score between 0 and 1 was produced by the model and was used as the SM prediction score. For the procedure to define threshold SM score classifying a gene as SM or not, the performance measures used, and the random background model, see *SI Appendix, SI Methods*. Dual-annotation (DA) genes are genes annotated as both GM and SM pathway genes in AraCyc. Junction (JC) genes were defined based on the pathway annotation data (pathway.dat) from the PlantCyc *A. thaliana* v.12 dataset. Two three-class models were built. The first SM/GM/DA model used SM, GM, and DA genes (benchmark 4) as the three classes. The second SM/GM/JC model used SM, GM, and JC genes (benchmark 5). Additional information for defining the JC gene type is available in *SI Appendix, SI Methods*.

ACKNOWLEDGMENTS. We thank Christina B. Azodi, Joshua J. Moore, Nicholas L. Panchy, and Sahra Uygun for helpful discussion and support. We also thank the editor and anonymous reviewers for critical comments

that led to new findings, particularly those on dual function and junction genes. This work was partly supported by grants from the National Science Foundation (NSF) IOS-1546617 (to R.L.L., E.P., and S.-H.S.), NSF DEB-1655386 (to S.-H.S.), and the US Department of Energy (DOE) Great Lakes Bioenergy Research Center (DOE Office of Science BER DE-SC0018409) (to R.L.L. and S.-H.S.).

1. Panchy N, Lehti-Shiu M, Shiu S-H (2016) Evolution of gene duplication in plants. *Plant Physiol* 171:2294–2316.
2. Hanada K, Zou C, Lehti-Shiu MD, Shinozaki K, Shiu S-H (2008) Importance of lineage-specific expansion of plant tandem duplicates in the adaptive response to environmental stimuli. *Plant Physiol* 148:993–1003.
3. Chen F, Tholl D, Bohlmann J, Pichersky E (2011) The family of terpene synthases in plants: A mid-size family of genes for specialized metabolism that is highly diversified throughout the kingdom. *Plant J* 66:212–229.
4. Chae L, Kim T, Nilo-Poyanco R, Rhee SY (2014) Genomic signatures of specialized metabolism in plants. *Science* 344:510–513.
5. Pichersky E, Lewinsohn E (2011) Convergent evolution in plant specialized metabolism. *Annu Rev Plant Biol* 62:549–566.
6. Hartmann T (2007) From waste products to ecochemicals: Fifty years research of plant secondary metabolism. *Phytochemistry* 68:2831–2846.
7. Ehrlich PR, Raven PH (1964) Butterflies and plants: A study in coevolution. *Evolution* 18:586–608.
8. Ali JG, Agrawal AA (2014) Asymmetry of plant-mediated interactions between specialist aphids and caterpillars on two milkweeds. *Funct Ecol* 28:1404–1412.
9. Howat S, et al. (2014) Paclitaxel: Biosynthesis, production and future prospects. *N Biotechnol* 31:242–245.
10. Zhong J-J (2002) Plant cell culture for production of paclitaxel and other taxanes. *J Biosci Bioeng* 94:591–599.
11. Giuliano G, Tavazza R, Diretto G, Beyer P, Taylor MA (2008) Metabolic engineering of carotenoid biosynthesis in plants. *Trends Biotechnol* 26:139–145.
12. Milo R, Last RL (2012) Achieving diversity in the face of constraints: Lessons from metabolism. *Science* 336:1663–1667.
13. Shoji T, Hashimoto T (2011) Recruitment of a duplicated primary metabolism gene into the nicotine biosynthesis regulon in tobacco. *Plant J* 67:949–959.
14. Ning J, et al. (2015) A feedback-insensitive isopropylmalate synthase affects acylsugar composition in cultivated and wild tomato. *Plant Physiol* 169:1821–1835.
15. Hofberger JA, Lyons E, Edger PP, Chris Pires J, Eric Schranz M (2013) Whole genome and tandem duplicate retention facilitated glucosinolate pathway diversification in the mustard family. *Genome Biol Evol* 5:2155–2173.
16. Schenck CA, et al. (2017) Molecular basis of the evolution of alternative tyrosine biosynthetic routes in plants. *Nat Chem Biol* 13:1029–1035.
17. Moghe GD, Last RL (2015) Something old, something new: Conserved enzymes and the evolution of novelty in plant specialized metabolism. *Plant Physiol* 169:1512–1523.
18. Initiative AG; Arabidopsis Genome Initiative (2000) Analysis of the genome sequence of the flowering plant Arabidopsis thaliana. *Nature* 408:796–815.
19. D'Auria JC, Gershenzon J (2005) The secondary metabolism of Arabidopsis thaliana: Growing like a weed. *Curr Opin Plant Biol* 8:308–316.
20. Wisecaver JH, et al. (2017) A global coexpression network approach for connecting genes to specialized metabolic pathways in plants. *Plant Cell* 29:944–959.
21. Schläpfer P, et al. (2017) Genome-wide prediction of metabolic enzymes, pathways, and gene clusters in plants. *Plant Physiol* 173:2041–2059.
22. Kliebenstein DJ (2008) A role for gene duplication and natural variation of gene expression in the evolution of metabolism. *PLoS One* 3:e1838.
23. Edger PP, et al. (2015) The butterfly plant arms-race escalated by gene and genome duplications. *Proc Natl Acad Sci USA* 112:8362–8366.
24. Wei H, et al. (2006) Transcriptional coordination of the metabolic network in Arabidopsis. *Plant Physiol* 142:762–774.
25. Higashi Y, Saito K (2013) Network analysis for gene discovery in plant-specialized metabolism. *Plant Cell Environ* 36:1597–1606.
26. Uygun S, Peng C, Lehti-Shiu MD, Last RL, Shiu S-H (2016) Utility and limitations of using gene expression data to identify functional associations. *PLoS Comput Biol* 12: e1005244.
27. Ashburner M, et al.; The Gene Ontology Consortium (2000) Gene ontology: Tool for the unification of biology. *Nat Genet* 25:25–29.
28. Rhee SY, Zhang P, Foerster H, Tissier C (2006) AraCyc: Overview of an Arabidopsis metabolism database and its applications for plant research. *Biotechnology in Agriculture and Forestry* (Springer, Berlin).
29. Steppuhn A, Baldwin IT (2007) Resistance management in a native plant: Nicotine prevents herbivores from compensating for plant protease inhibitors. *Ecol Lett* 10: 499–511.
30. Huot B, Yao J, Montgomery BL, He SY (2014) Growth-defense tradeoffs in plants: A balancing act to optimize fitness. *Mol Plant* 7:1267–1287.
31. Cedar H, Bergman Y (2009) Linking DNA methylation and histone modification: Patterns and paradigms. *Nat Rev Genet* 10:295–304.
32. Chan SW-L, Henderson IR, Jacobsen SE (2005) Gardening the genome: DNA methylation in Arabidopsis thaliana. *Nat Rev Genet* 6:351–360, and erratum (2005) 6:590.
33. Das M, et al. (2016) Expression pattern similarities support the prediction of orthologs retaining common functions after gene duplication events. *Plant Physiol* 171: 2343–2357.
34. Liu S-L, Baute GJ, Adams KL (2011) Organ and cell type-specific complementary expression patterns and regulatory neofunctionalization between duplicated genes in Arabidopsis thaliana. *Genome Biol Evol* 3:1419–1436.
35. Arabidopsis Interactome Mapping Consortium (2011) Evidence for network evolution in an Arabidopsis interactome map. *Science* 333:601–607.
36. Lee T, Lee I (2017) A network biology server for Arabidopsis thaliana and other non-model plant species. *Plant Gene Regulatory Networks. Methods in Molecular Biology* (Humana Press, New York).
37. Brachi B, et al. (2015) Coselected genes determine adaptive variation in herbivore resistance throughout the native range of *Arabidopsis thaliana*. *Proc Natl Acad Sci USA* 112:4032–4037.
38. Puthiyaveetil S, et al. (2010) Transcriptional control of photosynthesis genes: The evolutionarily conserved regulatory mechanism in plastid genome function. *Genome Biol Evol* 2:888–896.
39. Lloyd JP, Seddon AE, Moghe GD, Simenc MC, Shiu S-H (2015) Characteristics of plant essential genes allow for within- and between-species prediction of lethal mutant phenotypes. *Plant Cell* 27:2133–2147.
40. De Smet R, et al. (2013) Convergent gene loss following gene and genome duplications creates single-copy families in flowering plants. *Proc Natl Acad Sci USA* 110: 2898–2903.
41. Tatusov RL, et al. (2003) The COG database: An updated version includes eukaryotes. *BMC Bioinformatics* 4:41.
42. Rizzon C, Ponger L, Gaut BS (2006) Striking similarities in the genomic distribution of tandemly arrayed genes in Arabidopsis and rice. *PLoS Comput Biol* 2:e115.
43. Balakirev ES, Ayala FJ (2003) Pseudogenes: Are they "junk" or functional DNA? *Annu Rev Genet* 37:123–151.
44. Moghe GD, et al. (2014) Consequences of whole-genome triplication as revealed by comparative genomic analyses of the wild radish Raphanus raphanistrum and three other brassicaceae species. *Plant Cell* 26:1925–1937.
45. Osbourn A (2010) Secondary metabolic gene clusters: Evolutionary toolkits for chemical innovation. *Trends Genet* 26:449–457.
46. Qi X, et al. (2004) A gene cluster for secondary metabolism in oat: Implications for the evolution of metabolic diversity in plants. *Proc Natl Acad Sci USA* 101:8233–8238.
47. Sakamoto T, et al. (2004) An overview of gibberellin metabolism enzyme genes and their related mutants in rice. *Plant Physiol* 134:1642–1653.
48. Matsuno M, et al. (2009) Evolution of a novel phenolic pathway for pollen development. *Science* 325:1688–1692.
49. Renault H, Bassard J-E, Hamberger B, Werck-Reichhart D (2014) Cytochrome P450-mediated metabolic engineering: Current progress and future challenges. *Curr Opin Plant Biol* 19:27–34.
50. Thimm O, et al. (2004) MAPMAN: A user-driven tool to display genomics data sets onto diagrams of metabolic pathways and other biological processes. *Plant J* 37: 914–939.
51. Berardini TZ, et al. (2015) The Arabidopsis information resource: Making and mining the "gold standard" annotated reference plant genome. *Genesis* 53:474–485.
52. Finn RD, et al. (2016) The Pfam protein families database: Towards a more sustainable future. *Nucleic Acids Res* 44:D279–D285.
53. Schmid M, et al. (2005) A gene expression map of Arabidopsis thaliana development. *Nat Genet* 37:501–506.
54. Wilson TJ, Lai L, Ban Y, Ge SX (2012) Identification of metagenes and their interactions through large-scale analysis of Arabidopsis gene expression data. *BMC Genomics* 13:237.
55. Kilian J, et al. (2007) The AtGenExpress global stress expression data set: Protocols, evaluation and model data analysis of UV-B light, drought and cold stress responses. *Plant J* 50:347–363.
56. Goda H, et al. (2008) The AtGenExpress hormone and chemical treatment data set: Experimental design, data evaluation, model data analysis and data access. *Plant J* 55: 526–542.
57. Mockler TC, et al. (2007) The DIURNAL project: DIURNAL and circadian expression profiling, model-based pattern matching, and promoter analysis. *Cold Spring Harb Symp Quant Biol* 72:353–363.
58. Gautier L, Cope L, Bolstad BM, Irizarry RA (2004) Affy—Analysis of Affymetrix GeneChip data at the probe level. *Bioinformatics* 20:307–315.
59. Ritchie ME, et al. (2015) Limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Res* 43:e47.
60. Chen F, Mackey AJ, Stoeckert CJ, Jr, Roos DS (2006) OrthoMCL-DB: Querying a comprehensive multi-species collection of ortholog groups. *Nucleic Acids Res* 34: D363–D368.
61. Lehti-Shiu MD, et al. (2015) Molecular evidence for functional divergence and decay of a transcription factor derived from whole-genome duplication in *Arabidopsis thaliana*. *Plant Physiol* 168:1717–1734.
62. Moghe GD, et al. (2013) Characteristics and significance of intergenic polyadenylated RNA transcription in Arabidopsis. *Plant Physiol* 161:210–224.
63. Pedregosa F, et al. (2011) Scikit-learn: Machine learning in Python. *J Mach Learn Res* 12:2825–2830.