

New meta-analysis tools reveal common transcriptional regulatory basis for multiple determinants of behavior

Seth A. Ament^{a,1,2}, Charles A. Blatt^{b,1}, Cedric Alaux^{c,d,1}, Marsha M. Wheeler^c, Amy L. Toth^c, Yves Le Conte^d, Greg J. Hunt^e, Ernesto Guzmán-Novoa^f, Gloria DeGrandi-Hoffman^g, Jose Luis Uribe-Rubio^h, Gro V. Amdam^{i,j}, Robert E. Page, Jr.ⁱ, Sandra L. Rodriguez-Zas^{k,l}, Gene E. Robinson^{a,c,l,3}, and Saurabh Sinha^{b,l,3}

^aNeuroscience Program, Departments of ^bComputer Science, ^cEntomology, and ^kAnimal Sciences, and ^lInstitute for Genomic Biology, University of Illinois at Urbana-Champaign, Urbana, IL 61801; ^dUnité de Recherche (UR) 406 Abeilles et Environnement, Institut National de la Recherche Agronomique (INRA), Avignon cedex 9, France; ^eDepartment of Entomology, Purdue University, West Lafayette, IN 47907; ^fDepartment of Environmental Biology, University of Guelph, Guelph, ON, Canada N1G 2W1; ^gCarl Hayden Bee Research Center, US Department of Agriculture-Agricultural Research Service, Tucson, AZ 85719; ^hCentro Nacional de Investigación en Fisiología Animal, Instituto Nacional de Investigaciones Forestales, Agrícolas y Pecuarias, Ajuchitlan, Qro, 76280, Mexico; ⁱSchool of Life Sciences, Arizona State University, Tempe, AZ 85287; and ^jDepartment of Chemistry, Biotechnology and Food Science, Norwegian University of Life Sciences, 1432 Aas, Norway

Contributed by Gene E. Robinson, April 15, 2012 (sent for review February 12, 2012)

A fundamental problem in meta-analysis is how to systematically combine information from multiple statistical tests to rigorously evaluate a single overarching hypothesis. This problem occurs in systems biology when attempting to map genomic attributes to complex phenotypes such as behavior. Behavior and other complex phenotypes are influenced by intrinsic and environmental determinants that act on the transcriptome, but little is known about how these determinants interact at the molecular level. We developed an informatic technique that identifies statistically significant meta-associations between gene expression patterns and transcription factor combinations. Deploying this technique for brain transcriptome profiles from ca. 400 individual bees, we show that diverse determinants of behavior rely on shared combinations of transcription factors. These relationships were revealed only when we considered complex and variable regulatory rules, suggesting that these shared transcription factors are used in distinct ways by different determinants. This regulatory code would have been missed by traditional gene coexpression or *cis*-regulatory analytic methods. We expect that our meta-analysis tools will be useful for a broad array of problems in systems biology and other fields.

honey bee | transcriptional regulation

A hallmark of complex phenotypes is that they have many different intrinsic and extrinsic determinants, and phenotypic variation between individuals is shaped by the interplay between heritable genotypic factors and environmental conditions (1). Understanding a complex phenotype, therefore, requires knowing whether its different determinants are sub-served by common or distinct molecular mechanisms. For a given phenotype, do the multiple determinants act in fundamentally similar or different ways?

Both intrinsic and extrinsic determinants influence phenotypes through their effects on gene expression. For instance, intraspecific behavioral differences between individuals are associated with changes in the expression of hundreds to thousands of genes in the brain (2). Also, manipulations of genetic, environmental, and hormonal determinants of some of these behavioral differences can induce similar changes in gene expression (3). This manipulation leads to the hypothesis that multiple intrinsic and extrinsic determinants that influence the same behavior exert shared effects on shared transcriptional regulatory mechanisms, for instance, relying on the activities of the same sequence-specific transcription factors (TFs).

Testing this hypothesis poses serious challenges: transcriptional mechanisms linking genes to behavior are not well-understood, and although it is increasingly common to obtain

transcriptomic profiles for a variety of determinants of a complex phenotype, there is no existing tool that can directly test if specific *cis*-elements are activated in multiple transcriptomic states in a statistically significant manner. Some existing tools (4–7) have been widely used to search for associations in specific transcriptomic states, but they are not designed for detecting meta-associations across multiple states.

Here, we describe computational tools for identifying a statistically significant relationship between single or combinations of predefined *cis*-regulatory elements and multiple transcriptomic states. These tools combine evidence of association with each transcriptomic state into a single *P* value that reflects a recurring role for the *cis*-element in those states. No existing tools provide this functionality, although there are available methods for integrative analysis of multiple expression datasets. For example, “biclustering” tools attempt ab initio discovery of “gene modules” that display coordinated expression across experiments (8, 9). Other tools additionally require these modules to have common promoter motifs (10–12). These algorithms are faced with a large space of potential modules that (i) makes finding the best modules prohibitively slow and (ii) likely leads to loss of statistical power because of the large number of hypotheses tested. By contrast, our technique focuses on gene modules defined by predetermined *cis*-elements and statistical evaluation of meta-associations involving those modules. Its search space is relatively small, enabling an exhaustive search for the best meta-associations and keeping the extent of multiple hypothesis testing at a relatively modest level. We deploy our tools to analyze brain transcriptome profiles and promoter scans from ca. 400 individual honey bees, and we show that diverse determinants of

Author contributions: S.A.A., C.A.B., C.A., G.E.R., and S.S. designed research; C.A.B., C.A., M.M.W., A.L.T., Y.L.C., G.J.H., E.G.-N., G.D.-H., J.L.U.-R., G.V.A., R.E.P., and S.S. performed research; S.A.A., C.A.B., S.L.R.-Z., G.E.R., and S.S. analyzed data; and S.A.A., C.A.B., G.E.R., and S.S. wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

Data deposition: The microarray datasets reported in this paper have been deposited in the ArrayExpress Archive, <http://www.ebi.ac.uk/arrayexpress> (accession nos. E-TABM-953, E-TABM-952, E-MEXP-3079, E-MTAB-507, and E-MTAB-490).

¹S.A.A., C.A.B., and C.A. contributed equally to this work.

²Present address: Institute for Systems Biology, Seattle, WA 98109.

³To whom correspondence may be addressed. E-mail: generobi@illinois.edu or sinhas@illinois.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1205283109/-DCSupplemental.

behavioral maturation rely on specific, shared combinations of transcription factors.

Behavioral maturation is a common feature in vertebrates [including humans (13)] and social insects (14). Worker honey bees do not grow up and reproduce, but they do grow up. They perform brood care (nursing) and other tasks inside the hive for the first 2–3 wk of adult life, then switch to foraging outside for nectar and pollen (15). The age at onset of foraging is a major maturational milestone in bee life, and it is determined by a variety of genotypic and environmental factors (henceforth called maturation determinants) (Fig. 1A) that act, in part, by inducing differential expression of thousands of genes in the brain. For example, some subspecies of honey bees mature faster than others (16), whereas environmental determinants such as pheromones and nutrition also affect maturation (17) through downstream endocrine and molecular processes (17, 18). Brain transcriptome profiling of some maturation determinants has been conducted previously (16, 19–22), revealing strong similarities in their underlying brain gene expression profiles. However, existing statistical methodologies are insufficient to infer whether these maturation determinants act through shared transcriptional regulatory mechanisms.

Results

Datasets and Initial Analysis. We studied brain transcriptome profiles for *ca.* 400 individual honey bees. The profiles included a comparison of nurses and foragers plus responses to 10 known determinants of maturation. For simplicity, we refer to a total of 11 maturation determinants. Some profiles are presented here, and others are from previously published microarray studies (16, 20, 21). These determinants relate to genotypic differences, environmental signals and cues, and manipulations of hormonal

and molecular pathways (Fig. 1A). Each maturation determinant either speeds up or slows down behavioral maturation and was associated with hundreds to thousands of differentially expressed genes (DEGs) in the brain (Table S1); each determinant was studied in sufficient numbers of bees to obtain robust *P* values for the DEGs. We refer to the DEGs that are more highly expressed in the faster or slower maturing bees in each experiment as the “fast” or “slow” maturation genes, respectively. Because all bees were sampled before foraging, the expression differences could not be because of the effects of foraging *per se*, which also have been shown (16, 23).

We first assessed the extent to which the 11 maturation determinants cause shared changes in gene expression in the brain. Most (41 of 55) pairs of determinants influenced significantly more genes in common than expected by chance (Fig. S1), providing initial support for a common basis to their effects on behavior. We next asked if pairs of determinants also induced correlated changes to the transcriptome. We measured the correlation in expression fold changes between pairs of determinants for each of 1,000 randomly sampled “bootstrap” gene sets, each with 500 genes. We assumed pairwise correlations limited to relatively few bootstrap sets reflect shared effects that are localized to specific gene modules (24), whereas correlations that are observed in most of the bootstrap sets more likely correspond to shared effects encompassing much of the transcriptome.

Based on the above assumptions, a few pairs of maturation determinants induced strongly correlated effects on gene expression that involved much of the transcriptome (Fig. 1B, solid colors and Figs. S2 and S3). For example, bees from a genetic strain artificially selected to store large amounts of pollen in the hive (PH experiment) had brain gene expression profiles that were very similar to the profiles of wildtype bees experimentally deprived of pollen (Diet experiment). Both pollen-hoarding and nutritionally deprived bees show accelerated behavioral maturation (17, 25). This result suggests that selection for high pollen hoarding resulted in genetic assimilation (26) of transcriptional changes associated with the need for pollen and that the maturational effects for this pair of inherited and environmental determinants involve a shared transcriptional program.

By contrast, many pairs of maturation determinants induced correlated effects detectable in relatively few bootstrap gene sets, hinting at the existence of gene modules that are used by multiple determinants (Fig. 1B, stripes and Fig. S3). Surprisingly, several of these pairs induced correlations in a direction opposite to that predicted by their effect on behavior. For instance, Africanized honey bees initiate foraging at a younger age than bees from European subspecies (27) (AvE experiment); a molecular manipulation, Vg RNAi, also cause bees to forage at a younger age (like Africanized bees) but causes brain transcriptome responses more similar to European subspecies (18). This observation suggests that shared gene modules may be manipulated in different ways by distinct maturation determinants.

Metalysis and *cis*-Metalysis. The above analysis did not reveal any insights into the specific gene modules that may underlie the common phenotypic effect of distinct maturation determinants. To address this issue, we sought to identify biological features that are shared by DEGs across many or all maturation determinants. Surprisingly, we found that existing tools were not suitable for this task (Discussion). We, therefore, developed a tool, called Metalysis, to test for statistically significant relationships between a biological feature and multiple transcriptomic states. In a specialized version of this tool, called *cis*-Metalysis, the biological features considered are defined by the presence of single or combinations of *cis*-elements in gene promoters.

Metalysis can be used broadly to conduct meta-analyses of gene expression data from multiple transcriptomic states, identify gene properties that are correlated with expression in several




Fig. 1. Determinants of adult worker honey bee behavioral maturation and their correlated effects on brain gene expression. (A) Bees switch from hive work (primarily nursing brood) to foraging outside for nectar and pollen. We studied the regulation of new and previously reported brain transcriptomic differences between nurses and foragers and responses to 10 determinants of the timing of the age at onset of foraging (+ stimulating or – inhibiting foraging onset; e.g., rich diet delays onset of foraging). Previous studies also have characterized stimulatory (+) and inhibitory (–) relationships among some of the maturation determinants themselves. For genetic comparisons (AvE, LvM, and PH), the first genotype shows faster maturation than the second one. AvE, Africanized vs. European subspecies; BP, brood pheromone; cGMP: cyclic-guanosine monophosphate treatment; diet, rich vs. poor diet; JHA, juvenile hormone analog treatment; LvM, Northern (*A. mellifera mellifera*) vs. Southern European (*A. mellifera ligustica*) subspecies; Mat, maturation (nurses vs. foragers); Mn, manganese treatment; PH, high vs. low pollen-hoarding genetic strains; QMP, Queen Mandibular Pheromone; Vg, vitellogenin RNAi (details in Table S1). (B) Pairs of maturation determinants with effects on gene expression that were correlated in a direction consistent with their shared effects on behavior (green) or opposite to their effects on behavior (red) across most (solid colors) or relatively few (stripes) bootstrap gene sets.

transcriptomic states, and quantify the statistical significance of such meta-associations. The inputs to Metalysis are (i) a matrix with information on DEGs, indicating which genes are up- or down-regulated in each of several experiments, and (ii) a matrix of gene annotations, listing which of several different annotations (biological properties) apply to each gene. Each annotation type defines a gene set (i.e., genes that have that annotation). Given the lists of DEGs from multiple studies and an annotation-based gene set, Metalysis tests the association between that gene set and the DEGs from each study, and it combines the *P* values from these tests into a “meta-*P* value for the aggregate of all tests (we compare this statistic with related tests and procedures in *Discussion* and empirically show its advantages in *Methods*). The meta-*P* value is a newly designed statistic that can reveal meta-associations involving many moderately significant *P* values as well as those meta-associations involving a few highly significant *P* values. Metalysis uses empirical extreme value distributions (EVDs) to eliminate weak meta-associations, which controls for the effects of multiple hypothesis testing. Thus, Metalysis combines the common practice of testing for associations between gene sets (4–7) with an intuitive and powerful statistic that aggregates statistical evidence from multiple tests.

cis-Metalysis is specially designed for meta-analysis at the level of *cis*-regulatory elements in gene promoters. It first scans promoter regions to define “motif modules” [i.e., sets of genes with promoters that all contain the same *cis*-regulatory sequence pattern (motif) recognized by a specific TF (*Methods*)]. Here, the term promoter is used loosely to refer to the 5-kbp region upstream of the gene, not only the “core promoter” (28) (*Discussion*). *cis*-Metalysis then searches for meta-associations between each motif module and lists of DEGs from multiple experiments (Fig. 2*A*) (*Methods*) using the same meta-analytic approach as Metalysis. In fact, this functionality of *cis*-Metalysis is equivalent to using Metalysis with the input gene annotations being based on motif presence in promoters. However, *cis*-Metalysis is capable of greater flexibility in its search for meta-associations. For instance, it also tests motif modules defined by logical combinations of multiple motifs. Moreover, it can be configured so that different logical combinations of the same motifs can be associated with DEGs in different experiments, therefore offering a flexible model of regulatory mechanisms shared by multiple transcriptomic states. This flexibility (explained in detail below and *Methods*) and the powerful statistical framework of Metalysis lead to the broad applicability of *cis*-Metalysis. Figs. S4 and S5 and Table S2 show an illustrative application of *cis*-Metalysis to an entirely different topic, regulatory changes underlying cancer.

Metalysis and *cis*-Metalysis Reveal Common Transcriptional Regulatory Basis for Multiple Determinants of Behavioral Maturation. We used Metalysis to explore the gene modules underlying behavioral maturation. We first used Metalysis in conjunction with Gene Ontology (GO) annotations, revealing biological processes that were enriched in brain transcriptomic responses to as many as 8 of 11 maturation determinants (Fig. 3*A*). We only examined meta-associations with EVD *P* value ≤ 0.05 , which means that, in a carefully constructed negative control (*Methods*), there is less than a 5% chance of finding any meta-association as strong as the examined ones. These meta-associations included processes occurring in the brain related to macronutrient and energy metabolism (translation, mitochondrial electron transport, and glycolysis), neuronal plasticity (synaptic transmission and nervous system development), and stress responses (protein folding and response to heat). All of these biological processes had known associations with particular individual maturation determinants (29–31), but our results suggest that they actually are driven in common by many determinants. This finding suggested to us the hypothesis that multiple maturation determinants operate through the actions of a common set of TFs.

We used *cis*-Metalysis to test this hypothesis using 602 previously characterized motifs from vertebrate and *Drosophila* databases. Because transcriptional regulation in many biological contexts involves combinatorial interactions between TFs, we used multiple modes of *cis*-Metalysis to test both simple and complex models of regulation, which is described next (Fig. 2*B*). In our analysis, we tested meta-associations involving individual motifs or motif pairs, although *cis*-Metalysis can be used to examine all *k*-tuples of motifs for a user-defined *k*.

Single motif and identical logic modes. We first used *cis*-Metalysis to test the hypothesis that maturation determinants use regulatory rules that are identical across many or all conditions. In the single motif mode, *cis*-Metalysis defines a motif module as the set of genes with promoters that have matches to a motif and then tests for a meta-association involving this motif module, just as the GO gene sets were tested above. In the identical logic mode, the motif module is defined by a logical combination of the presence and/or absence of two motifs (e.g., motif A and not motif B) in promoters. These simple regulatory rules were shared by the responses to some, but not all, of the maturation determinants. The most significant meta-associations involving single motifs (EVD *P* value < 0.01) (Fig. 3*B*) spanned four to six determinants and were often bidirectional. This finding means that most of the 22 individual motifs reported here do not consistently predict fast vs. slow maturation. In fact, when we configured *cis*-Metalysis to only seek unidirectional meta-associations, a much smaller number of motifs was reported at the same significance threshold (Fig. 3*B*, asterisks). Similarly, meta-associations involving combinations of two motifs interacting with identical logic spanned no more than 4 of 11 determinants (Fig. S6). These results suggest that more complex regulatory logic may be required to explain the shared effects of maturation determinants on behavior; this suggestion proved to be true, which is seen in the next analysis.

Role-consistent logic mode. Allowing for greater complexity in the combinatorial action of TF pairs revealed shared regulatory logic across most or all maturation determinants. We characterized higher-complexity regulatory rules by using *cis*-Metalysis in its role-consistent logic and flexible logic modes. In the role-consistent logic mode, *cis*-Metalysis tests if a motif pair is associated with multiple determinants, but unlike the identical logic mode, the motif module associated with differential expression can vary from one determinant to another, and is not defined by a predetermined logical combination. Thus, a module defined by the presence of motif A and absence of motif B may be associated with fast maturation genes of one determinant, whereas a module defined by the presence of motif B and absence of motif A is enriched in slow maturation genes of a second determinant. However, the role consistency requirement mandates that, if a motif's presence associates with fast maturation genes in one experiment, its presence cannot be associated with slow maturation genes in a second experiment. In the above example, an association of a module defined by the presence of both motifs A and B with fast maturation genes of a third determinant will not be considered for combining with the two mentioned associations.

This role-consistent logic mode revealed pairs of motifs that were associated with differential gene expression for as many as 10 determinants (Fig. 4*A*). In contrast to the simpler rules above, these rules necessarily predict unidirectional responses—each of the implicated TFs is predicted to function consistently to either speed up or slow down maturation. We note that the increased complexity of possible meta-associations is accounted for when EVD *P* values are computed, and therefore, it is not obvious a priori that more complex combinations will show more widely shared associations.

The motif combinations identified by the role-consistent logic mode featured *cis*-regulatory sequences recognized by TFs that




Fig. 2. Informatics techniques for the discovery of meta-associations across transcriptomic experiments. (A) Flowchart for *cis*-Metalysis. Steps 1 and 2: Motif modules are defined based on the presence (green cells) of single motifs (e.g., M_1) or their Boolean combinations (e.g., M_1 and M_2) in the 5-kb promoter sequences of each gene. Steps 3 and 4: Sets of up- (orange) and down- (blue) regulated genes are identified from experimentally profiling gene expression for each determinant. Step 5: Statistical enrichments are tested between motif modules and expression sets producing all of the association P values (shaded cells) that are combined by *cis*-Metalysis. In its most flexible mode (shown), *cis*-Metalysis combines the best P value (bordered in yellow) from each determinant when considering all combinations of motifs. Step 6: Meta- P value calculation. *Top* shows the best association P value from each condition ($D_1 \dots D_4$). For each $K = 1 \dots 4$, the statistic Φ_K combining the best K P values is computed and translated to a P value $P(\Phi_K)$ in *Middle*. The minimum $P(\Phi_K)$ over all K is the meta- P value (highlighted in yellow within the red border) and considers the number and strength of the combined P values. The meta-association is represented in *Bottom* with selected significant conditions colored and the remaining conditions in gray. (B) Multiple modes of *cis*-Metalysis enable discovery of simple or combinatorial identical or plastic forms of regulatory logic shared across transcriptomic states (conditions) by selecting specific subsets of association P values to combine. For each mode, an example subset of associations (shaded cells) and the selected best P values (bordered in yellow) are depicted. In single motif *cis*-Metalysis, P values are selected from the association tests with the gene modules defined by a single motif (M_1). In the identical logic configuration, a rigidly defined combination of two motifs ($!M_1$ and M_2 in this example) defines the gene module that is tested for association with a fixed direction of regulation (down). In role-consistent and flexible *cis*-Metalysis, a meta-association of a motif pair (M_1 and M_2) is allowed to use different Boolean combinations of the two motifs and different directions of regulation in different determinants. In role-consistent mode, each motif-expression association must use a particular motif in the same role (activator or repressor); here, M_1 is an effective activator, and M_2 is a repressor. This constraint is relaxed in flexible *cis*-Metalysis (example in A, Step 5).

are well-known as regulators of neuronal plasticity (*Creb*) (32), stress responses [*Xbp1* (33) and *dl*, the *Drosophila* homolog of NF- κ B (34)], and juvenile hormone (JH) signaling (*broad*) (35). JH is known to mediate the effects of both genotypic and environmental maturation determinants (16, 17) (Fig. 1A). Intriguingly, the discovered *cis*-regulatory rules often paired these regulators and other well-known regulators of behavior and physiology with additional TFs (e.g., *Ets*) that have not previously been associated with behavior (Fig. 4B). This analysis, therefore, suggests that combinations of known and previously

unknown regulators of behavior are shared across many or all maturation determinants.

Flexible logic mode. This mode, like the role-consistent logic mode above, tests for meta-associations for a motif pair while allowing their specific logical combination to vary from one determinant to another; however, the role consistency requirement is eliminated now. With this greater flexibility, *cis*-Metalysis reported meta-associations spanning all 11 determinants (Fig. S7). These meta-associations predict that the same pair of motifs is involved in the responses to all of the determinants, but the roles of




Fig. 3. Meta-associations between GO terms, *cis*-regulatory motifs, and genotypic, environmental, and physiological determinants of honey bee behavioral maturation. (A) Meta-associations reported by Metalysis between maturation determinants and GO biological processes, with EVD P value ≤ 0.05 . The EVD P value for each meta-association is based on comparing its meta- P value to the smallest meta- P values returned from each of many randomization tests (negative controls). Orange and blue cells indicate the individual maturation determinants contributing to a meta-association; orange and blue denote associations with genes up-regulated in the fast maturation and slow maturation conditions, respectively. (B) Meta-associations (at EVD P value ≤ 0.01) revealed by single motif *cis*-Metalysis (Upper) and subjection of SAMBA biclusters to single motif enrichment tests (Lower). No significant meta-associations were discovered using BiMax biclusters. Asterisks denote motifs reported by *cis*-Metalysis (at EVD P value ≤ 0.01) when configured to search only unidirectional meta-associations. (C, Top) The best meta-associations (red ticks) between *cis*-regulatory motifs and determinants were much stronger than the best found in many randomized datasets (blue ticks; best-fit γ -distribution shown as the blue curve). Each blue tick represents the strongest meta-association in a distinct (randomized) dataset. Same plots for the results from BiMax (C, Middle) and SAMBA (C, Bottom) biclusters subjected to motif enrichment tests.

specific TFs with respect to maturation seem to vary among determinants. One such meta-association featured a bHLH motif (*USF*) paired with a motif bound by *mtTFA* (Fig. 4C), which regulates mitochondrial function (36). This finding reinforces our finding above (Fig. 3A) that mitochondrial functions such as energy metabolism are influenced by most or all maturation determinants and proposes *mtTFA* regulation, combined with a bHLH TF, as part of the underlying mechanism. This combination and 16 other motif combinations (involving 20 TFs overall) (Fig. S7) were found to be associated with all 11 determinants, with meta-associations having an EVD P value $< 2E-16$. These results are the most statistically significant among all results produced by *cis*-Metalysis, but they are intriguing in their use of the same TFs with switching functional roles.

Comparisons Between *cis*-Metalysis and Biclustering Tools. A popular approach to analyze expression data from multiple experiments is to solve the so-called biclustering problem, which is to

find subsets of genes and experiments such that the chosen genes are coordinately expressed in the selected experiments. After such a coexpressed gene set has been discovered, it is common to test for *cis*-elements overrepresented in their promoters, therefore inferring meta-associations between *cis*-elements and the selected experiments. To test this strategy, we applied the BiMax tool (37) to find biclusters across the 11 maturation determinants. About 500 biclusters were discovered, each including at least five genes and spanning at least three determinants. The genes in each of these biclusters were then tested for enrichment of motifs in our collection of 602 motifs as well as every logical combination of motif pairs. The top five motif-bicluster associations are shown in Fig. S8A with their P values. We then repeated the entire analysis 50 times on randomized datasets exactly as described for Metalysis to obtain empirical EVD P values that correct for multiple hypothesis testing. Fig. 3C shows the empirical EVD [i.e., the top P value (blue ticks) from each of the 50 negative controls]. We see that even the smallest P value




Fig. 4. Combinatorial regulation of maturation-related gene expression. (A) Top three most significant meta-associations between maturation determinants and pairs of motifs that interact with role-consistent logic. Motif module–gene expression associations with fast maturation and slow maturation genes in each experiment are denoted by orange and blue, respectively. Role-consistent meta-associations of motifs with fast and slow maturation are indicated by green and red, respectively. For example, the third row indicates that the presence of the DI motif and absence of the ETS motif are together associated with fast maturation genes in the QMP, Mat, BP, AvE, LvM, and JHA experiments, whereas the presence of the ETS motif and absence of the DI motif are associated with slow maturation genes in the Mn, diet, Vg, and cGMP experiments. (B) Combinations of motifs recognized by the TFs *DI*, *Xbp1*, and *ETS* are associated with fast and slow maturation, respectively, and together, they predict gene expression patterns across all 11 maturation determinants. A wiring diagram shows the most significant combination of motifs for each determinant. “And” and “not” logic gates are depicted by rounded rectangles and triangles, respectively. (C) Combinations of motifs recognized by *mtTFA* and *USF* are associated with the responses to all 11 maturation determinants, with each motif bidirectionally associated with up- and down-regulated genes in different experiments. Same notation as in B.

from the original dataset (red tick) is not statistically significant after multiple testing correction, with an EVD *P* value of 0.053. A similar examination using *cis*-Metalysis (Fig. 3C) reveals 22 meta-associations that are to the right of the entire empirical EVD and have EVD *P* values ≤ 0.01 . We also observed a relative lack of statistically significant meta-associations in the BiMax analysis for combinations of motifs (Fig. S8B).

We repeated the above analysis with another biclustering tool called SAMBA (9). Here, we used the fold-change values of genes, rather than their discretization into one of three categories (up, down, or neither). Twelve biclusters were discovered and then subjected to motif enrichment tests and multiple hypothesis corrections. Only two associations were found to have empirical EVD *P* value ≤ 0.01 (Fig. 3B and C and Fig. S8C and D); again, this is a much smaller number of associations than discovered by *cis*-Metalysis at the same significance levels.

Discussion

Biological Significance. With a set of 10 known intrinsic and extrinsic determinants, behavioral maturation in honey bees provides a compelling example of the challenge of elucidating the underlying transcriptional architecture of a complex trait. Our results present a striking view of the combinatorial *cis*-regulatory code underlying behavioral maturation. We found 16 meta-associations, involving 20 *cis*-regulatory motifs, that spanned all determinants but only when we considered pairs of motifs that varied in the ways in which they are predicted to interact with each

other. These results suggest that many different maturation determinants use the same TFs to exert common effects on behavior but that different determinants use some of these TFs in distinct ways. Those TFs that showed role-consistent patterns of association across the determinants might be involved in speeding up or slowing down behavioral maturation, whereas those TFs that showed bidirectional patterns of association might play a more general role in neural and behavioral plasticity.

Our results go beyond prior demonstrations of the strong relationship between behavioral plasticity and brain gene expression (2, 3, 16, 29, 38), and they hint at considerable complexity in the ways in which the same TFs interact with each other to influence gene expression. Our findings also show that diverse determinants can exert common influences, at least in part, by shared transcriptional mechanisms that involve a relatively small number of TFs. We focused on finding only the most robust associations involving only known *cis*-regulatory motifs, and therefore, the complete transcriptional regulatory network underlying behavioral maturation in honey bees is likely considerably more complex.

Our results suggest hypotheses for specific TFs as regulators of behavioral maturation, including those TFs with established roles in neuronal plasticity, stress responses, hormonal signaling, and metabolism. However, assignments of specific TFs to the *cis*-regulatory motifs identified in our analyses are limited by the fact that a single motif is often recognized by multiple TFs, and therefore, roles for TFs implicated by our analyses must be

confirmed by experimental manipulations. However, several of the TFs identified by *cis*-Metalysis also were implicated as high-level regulators of maturation-related gene expression in a bee brain transcriptional regulatory network reconstructed from some of the same gene expression datasets used here. These TFs include *Creb*, *br*, *dl*, *Xbp1*, and others (22), and they, thus, are particularly promising candidates for functional characterization in future experiments. We recently showed a functional role in behavioral maturation for *ultraspiracle* (39), a JH-related TF with a *cis*-regulatory motif that is enriched in the promoters of many maturation-regulated genes expressed in the brain (19), showing the use of *cis*-regulatory analysis in predicting regulators of bee behavior.

Our findings may reflect a significant theme in the regulation of complex phenotypes. The fact that multiple determinants of the same phenotype use common regulatory components is not surprising, and it is reminiscent of the diverse mechanisms for segmentation from flies to vertebrates, all of which rely on a common toolkit of regulatory genes (40). Less appreciated is the possibility that these common components may be wired differently in the different regulatory networks, reflecting the different adaptive forces that shaped the evolution of each of those networks and posing significant challenges to characterizing an underlying regulatory code for the phenotype.

Methodological Significance and Limitations. It seems that, despite the substantial literature on meta-analysis of transcriptomic experiments, there has been no rigorous treatment of the following intuitive task: determine if a given gene module shows significant statistical association with a differentially expressed gene set in all or most of several datasets. The Metalysis tool enables us to perform this task with rigorous statistics given any number of predefined gene modules, whereas *cis*-Metalysis extends this functionality to accommodate flexible combinations of modules.

An important contribution of our work is that we formulated a basic challenge of meta-analysis—combining a set of *P* values from multiple tests—in an intuitive way and solved the formal problem analytically. This statistic has important technical advantages for the meta-analysis problem addressed here compared with alternatives such as the Fisher combined probability test used in previous studies (41–43). For instance, the meta-*P* value statistic that we propose is better suited to situations where only a subset of the multiple *P* values being aggregated provides evidence against the null hypothesis. We believe that our meta-analysis tools will be useful for a broad array of statistical problems in systems biology and other fields.

We solved a fundamental statistical problem in meta-analysis, built on it to develop informatics tools that analyze sequence and expression data, and used these tools to reveal a flexible *cis*-regulatory code underlying a complex phenotype. We used a meta-analytic strategy where multiple assays of the same experimental condition are analyzed first and gathered robust information on differentially regulated genes of that condition, and this information is then integrated across multiple conditions. Prior applications of this strategy (12, 44, 45) were mostly devoted to finding a core set of genes with coordinated expression in multiple experiments (41, 46–50). In contrast, our approach bears on a more fundamental notion of what is shared across multiple transcriptomic responses—not a list of common genes but a common biological process or a common regulatory logic. This idea was mentioned in the work by Subramanian et al. (7), which used Gene Set Enrichment Analysis (GSEA) to discover that the same gene sets were associated with two separate transcriptomic profiles of lung cancer. However, this observation was made anecdotally, and GSEA does not provide a way to quantify such meta-analytic observations.

Our work has some similarities to biclustering, which is an existing technique to analyze transcriptomic data from multiple

experiments and studies. It typically involves ab initio discovery of a subset of genes and a subset of conditions that show strong association with each other. Such a gene set can then be subjected to functional characterizations (e.g., by examining its enrichment for GO terms, protein-protein interaction network colocalization, or even promoter motifs) (8–10). Biclustering-based approaches first find a core gene set activated in multiple transcriptomic states and then test for associations with predefined modules. Our approach switches the order of the two operations: modules are tested for association with each experiment first, and this information is then aggregated in the second meta-analytic step. This approach has several advantages. (i) It is a more direct way to search for meta-associations, because it skips the challenging step of finding a core gene set common to many transcriptomic states and instead sets up the meta-analysis as a well-defined and efficiently computable task of combining *P* values from multiple experiments. (ii) Evaluations of the statistical significance of biclusters (e.g., using extreme value distributions) are often difficult because of a large search space. Our strategy of testing meta-associations only for specific gene modules reduces the search space drastically, thus promising to provide greater statistical power. (iii) Our approach allows us to define meta-associations with much greater flexibility than the biclustering-based strategy. In particular, *cis*-Metalysis can test different logical combinations of the same motif pair for their association with different transcriptomic states when inferring a meta-association for that motif pair. This flexibility, which is crucial if the shared molecular mechanisms are themselves plastic, is not allowed in the biclustering-based approach. We believe that these are some of the reasons why *cis*-Metalysis was able to find more meta-associations than BiMax and SAMBA-based approaches.

A variant on the biclustering theme in meta-analysis strategies is the variant of integrative biclustering, where genes in a bicluster are required to be not only coordinately expressed but also show enrichment of other annotations such as motif presence. Tools in this genre such as cMonkey (10), Coalesce (12), and Allegro (11) integrate the two steps mentioned above: assessing the coexpression of gene sets and testing motif module enrichment in those gene sets. These tools, like other biclustering tools, are designed for ab initio discovery of biologically meaningful gene modules rather than testing for meta-associations involving predetermined modules, which is our goal. As such, they are likely to have less statistical power if subjected to rigorous multiple hypothesis correction. Because of the exponentially large space of potential biclusters, these tools rely on heuristics that may miss the best biclusters. Moreover, these three tools rely on ab initio motif finding to improve the quality of reported biclusters, and they do not provide the means to test hypotheses about known motifs and combinations thereof. This finding is important for metazoan genomes, where ab initio motif discovery, especially without comparative genomics, is problematic and to a large extent, can be replaced with the use of databases of experimentally characterized motifs (Table S3).

Our approach has some parallels to the approach in the work by Segal et al. (24), where reported gene modules exhibit associations with several transcriptomic states (cancer types) that have a common annotation (e.g., metastatic cancer). However, their approach requires a priori annotation of the transcriptomic states into two categories and thus, cannot be used to reveal meta-associations in a framework such as the one described here. Moreover, it uses meta-analytic statistics that count the number of significant associations at an arbitrary threshold as opposed to Metalysis, which integrates the strengths and number of associations without imposing thresholds. Also, their approach does not report gene modules representing combinatorial *cis*-regulatory codes.

A potential limitation of our work is that *cis*-regulatory analysis was based on analysis of promoter regions, whereas regula-

tory information in metazoa is frequently located more distally (51). We note, however, that our definition of a promoter is the region 5 kbp upstream of a gene, which is expected to include a substantial fraction of regulatory elements. In a genome where additional clues about regulatory locations, such as DNA accessibility and cofactor binding, are available, one may be able to scan more comprehensively for motif matches. Here, we made a practical tradeoff between sensitivity and specificity by focusing on 5-kbp promoters. We also note that *cis*-Metalysis may be run with any definition of the sequence space, not just the particular choice that we used here. Another point of future improvement is the completeness of the motif collection used. We have relied on multiple sources of experimentally and computationally predicted motifs (from *Drosophila* and vertebrates) to build a comprehensive motif collection. Nevertheless, a large number of relevant motifs are likely not known, and future high-throughput characterization of TF motifs will improve the sensitivity of our approach. We have used *cis*-Metalysis to be able to efficiently explore a large motif collection and detect and appropriately handle redundant pairs of motifs, which are expected to arise more frequently in a larger collection. Finally, we anticipate that future versions of our approach will incorporate information about TF expression levels when predicting regulatory roles for motifs. The work by Chandrasekaran et al. (22) has reported on such an analysis, where TF-gene interactions were predicted based on expression data alone; an integration of that approach (22) and our sequence-based approach is the logical next step for honey bee behavioral maturation, which has been done for *Drosophila* development (52).

Methods

Maturation Determinants. We profiled brain gene expression for naturally occurring behavioral differences between young and old *Apis mellifera* honey bees [nurses vs. foragers; maturation (20)], genotypic comparisons between subspecies and artificially selected strains [Africanized vs. European subspecies (AvE), Northern (*A. mellifera mellifera*) vs. Southern European (*A. mellifera ligustica*) subspecies (LvM), and high vs. low pollen-hoarding genetic strains (PH)], social and nutritional environmental determinants [brood pheromone (BP) (20), Queen Mandibular Pheromone (QMP) (21), and diet], and hormonal/molecular manipulations [vitellogenin RNAi (Vg), juvenile hormone analog treatment (JHA) (16), manganese treatment (Mn) (16), and cyclic-guanosine monophosphate treatment (cGMP) (16)]. All of these factors have been shown previously to influence the timing of maturation. Datasets listed with a reference were published previously. Details are provided in Table S1.

Sample Processing and Microarrays. Brain gene expression profiling was performed using a 70-mer oligonucleotide spotted microarray (20), with 13,440 unique experimental probes (Mat, AvE, LvM, PH, BP, diet, and Vg) or a cDNA spotted microarray (16), which assayed a total of 5,736 genes (QMP, JHA, Mn, and cGMP) with 4,662 genes matching a feature on the oligo array. Features from the two platforms were considered to match if they both mapped to the same Official Gene Set 2 (derived from the honey bee genome sequencing project) (53) gene model or if they mapped to overlapping regions outside an Official Gene Set 2 model. We profiled expression from individual brains (10–20/group) on the oligo array and material pooled from several brains on the cDNA array. Studies were implemented as loop designs, each with 20–200 arrays. Experimental procedures for cDNA microarrays and oligo arrays were as described (16, 20). Briefly, partially lyophilized, frozen brains were dissected out of head capsules. RNA extracted from individual or pooled brains (RNeasy; Qiagen) was subjected to one round of linear amplification and labeled with fluorescent dye (Cy3 or Cy5) using the Amino-Allyl Message AmpII kit (Ambion) or the Message Amp II kit combined with a Universal Labeling System (Kreatech). Slides were scanned with an Axon 4000B scanner, and images were analyzed with GENEPIX software (Agilent Technologies).

DEGs. Statistical analyses of differential gene expressions were performed as in ref. 20. Genes abundantly expressed in hypopharyngeal glands (a potential source of tissue contamination in brain samples) were filtered before analysis. A Loess transformation was performed using the software program Beehive (<http://stagbeetle.animal.uiuc.edu/Beehive>) to normalize expression

intensities. A linear mixed effects model implemented by using restricted maximum likelihood was used to describe the normalized log₂-transformed gene intensity values, including the effects of experimental variables, dye, bee, and microarray. Effects were evaluated with an *F*-test statistic, and the *P* values were adjusted for multiple hypothesis testing by using a false discovery rate (FDR) criterion. We studied DEGs that were significant at FDR < 0.05.

Meta-*P* value. Let a given set of *P* values from *n* independent tests be denoted, in ascending order, by $\{p(i)\}$. For each $k \in [1 \dots n]$, we compute (Eq. 1)

$$\phi_k = 1 - \prod_{i=1}^k (1 - p(i)) \quad [1]$$

and its *P* value $P(\phi_k)$ conditional on the fact that the *k* smallest *P* values were chosen from a set of *n*. Note that $\phi_k \in [0, 1]$ is a combination of the *k* $p(i)$ values that assumes a low value only if every $p(i)$ is small. A significant *P* value of this statistic, thus, corresponds to a situation where all of the individual $p(i)$ values are low to an extent that is unlikely to happen by random chance. We call $\text{mink}(P(\phi_k))$ the meta-*P* value corresponding to $\{p(i)\}$. We consider the minimum among all $P(\phi_k)$ values, because we want to allow for a subset of the tests to carry evidence against the null hypothesis but do not know a priori the number of tests. Computing $P(\phi_k)$ involves a newly derived analytical calculation given by (Eq. 2)

$$\Pr(\phi_k \leq \tau) = 1 - \int_0^1 \binom{n}{k+1} (k+1)(1-t)^k t^{n-k+1} A(t, k, 1-\tau) dt, \quad [2]$$

where $A(t, k, x)$ is a function with a calculation that was provided in the work by Dettmann and Georgiou (54) (SI Methods). Calculation of the meta-*P* value and previous work on the problem (46) assume that the combined *P* values come from independent tests. In practice, this assumption may be violated; however, the EVD *P* value calculation outlined below addresses this issue.

Metalysis. Inputs to this program are (i) a gene set *M* (for example, defined by a shared biological feature such as a GO annotation or motif presence) and (ii) for each experimental condition *C_i*, annotation of subsets of genes that are up- or down-regulated (*G_i+ and G_i-*, respectively) in that condition. The steps are (i) calculate *P* values of a Hypergeometric test of association between *M* and *G_i+ and G_i-*, respectively, (ii) select the lower *P* value *p_i* for each condition *C_i*, and (iii) compute the meta-*P* value from the resulting $\{p_i\}$ as the significance of meta-association between *M* and the set of conditions.

In future implementations of Metalysis, we will allow for testing individual associations between *M* and *G_i** using alternative procedures [e.g., GSEA (7) or other available tests (4, 55)]. There are two main reasons why the current implementation uses Hypergeometric tests after discretizing the expression information on each gene into one of three categories (up- or down-regulated or neither). First, in our application, this expression information is obtained from multiple biological replicates (bees), providing us with a *P* value for differential expression. This finding naturally leads us to work with genes with *P* values that are significant at some predetermined level, but it does not offer an equally rigorous way to choose an expression fold-change value that is representative of all replicates. Second, we noted that different transcriptomic studies often choose different methodologies to estimate expression levels, but almost all of these studies finally report lists of differentially regulated genes for additional analysis. Therefore, a meta-analysis framework with a starting point of sets of DEGs is likely to be more widely useable by the community.

Multiple Hypothesis Correction. Because the procedure is repeated for each given gene module *M* and because step ii amounts to performing two tests for each *C_i*, a multiple hypothesis correction is required. We, therefore, repeated the entire analysis for a randomized dataset where the matching between genes and promoters has been permuted randomly (SI Methods) and noted the smallest *P* value reported by Metalysis. Repeating this exercise many times, we constructed an empirical EVD of meta-*P* values, which provides an EVD *P* value corresponding to each meta-*P* value in the original dataset. The last step uses a γ -distribution fit to the empirical EVD, which has been reported previously in the context of ab initio motif discovery tools (56). We use the EVD *P* value for multiple hypothesis correction in Metalysis, because our goal in the present study was to obtain only the most reliable

meta-associations. Alternative deployments of the software may choose to use the less-conservative FDR approach (57).

Comparison of Meta-*P* Value with Previous Work. Prior work has examined the problem of combining *P* values from multiple tests of the same hypothesis (58), the most popular approach (41, 52) being Fisher's combined probability test, which uses the statistic (Eq. 3)

$$\chi_F^2 = -2 \sum_{i=1}^n \ln p_i. \quad [3]$$

This method is meant to test the significance for the aggregate of a number of independent tests, which is also the motivation for the meta-*P* value that we propose. However, unlike the Fisher method and other approaches to the same problem (59, 60), our method provides a composite test in the scenario where only a subset of the combined tests carries evidence against the null hypothesis. Moreover, the statistic ϕ_k and thus, the meta-*P* value is designed to be most sensitive to the largest of the *k* best *P* values, similar to the *P*_{max} statistic of ref. 61 and unlike the Fisher method, which is most sensitive to the smallest *P* value. In Figs. S9 and S10, we perform a comparison between Fisher's method and our meta-*P* value in terms of their ability to score meta-associations on synthetic as well as real sets of *P* values. We noted that the meta-associations retrieved from our datasets based on the meta-*P* value were more diverse than those meta-associations based on the Fisher method (Fig. S9) (e.g., they included more cases where only 2–4 of the 11 *P* values combined were individually significant and more cases where all 11 *P* values were individually significant). In positive control experiments where we artificially simulated 11 tuples of *P* values (Fig. S10), we found the meta-*P* value to correlate better with the strength of the meta-association that was used in the simulation procedure. Finally, we note that the Fisher combined test (or our meta-*P* value) can also be used to detect if at least one of the tests can reject its null hypothesis [i.e., as a family-wise error control procedure akin to the family-wise error rate (FWER)] (62); however, this detection is not the motivation for our work. The meta-*P* value is not intended to be a multiple hypothesis correction procedure—its goal is to aggregate evidence from multiple tests of the same hypotheses rather than select a subset of tests that meets predetermined criteria of significance. Nevertheless, it is worth noting that, in our experiments (Fig. S11), we found a meta-*P* value based approach to better select the subset of tests where the null hypothesis was false than was possible with a procedure based on FDRs and *q* values (63).

Meta-analysis of Associations Between GO Terms and Maturation Determinants. GO terms (Biological Process) with ≥ 10 and $\leq 1,000$ genes in *D. melanogaster* were considered, which resulted in 613 GO term gene sets used that were mapped to their *A. mellifera* orthologs. Metalysis was used to find associations between each GO gene set and the differentially regulated gene sets from subsets of the 11 experimental conditions.

cis-Metalysis. cis-Metalysis is an extension of the Metalysis program designed to find meta-associations involving motifs (or motif combinations) present in gene promoters. It scans gene promoters with each motif to predict which genes may be regulated by that motif and then performs meta-analysis using this information. The software is modular, with the motif scanning step being separate from the meta-analysis. Therefore, the motif scanning step may be skipped as long as motif presence information is provided to the software in the appropriate format.

i) Motif module prediction. A motif module is defined as the set of genes with significant presence of the motif in their promoters, which is determined as described in ref. 29. The method searches up to 5 kbp upstream of a gene for a motif's presence using the SWAN program (29), which captures the presence of one or more strong or weak matches to the motif in the segment and accounts for the local G/C composition as well as the global frequency of motif occurrence. Motif modules were predicted for 602 motifs obtained from FlyREG (*D. melanogaster*), Transfac (*D. melanogaster* and *Homo sapiens*), Jaspar (*H. sapiens*), and ref. 64.

(Table S3 presents our list of experimentally characterized motifs as well as our rationale for using them exclusively and not performing ab initio motif discovery.) The result of this step is a genes $\times$ motifs matrix of Boolean values set to true if a motif is deemed to be present in a gene's upstream region.

ii) Meta-analysis of motif-condition associations. The genes $\times$ motifs matrix computed by the previous step or through alternative means is used in this meta-analysis step. The user may configure the program to operate in one of four different modes corresponding to different ways of defining the cis-regulatory logic shared by multiple conditions. The modes are described briefly here, and a more formal description is in SI Methods. (i) Single motifs: the meta-*P* value of each motif module is computed in a manner identical to the manner for a GO gene set (above). (ii) Identical logic: For any motif pair (*m*₁ and *m*₂), the combinations *m*₁ $\wedge$ *m*₂, *m*₁ $\wedge$ $\neg$ *m*₂, and *m*₂ $\wedge$ $\neg$ *m*₁ are analyzed separately by constructing the respective motif modules from the motif modules of *m*₁ and *m*₂ and calculating the meta-*P* value of each of these three derived modules; a meta-association must involve either the up-regulated (*G*_i,+) in every experimental condition *C*_i or the down-regulated gene set (*G*_i,−) in every experimental condition. The best meta-*P* value among the three motif combinations is reported for the motif pair. (iii) Role-consistent logic: for any motif pair (*m*₁ and *m*₂), associations with different conditions may involve any of the following motif modules: *m*₁, *m*₂, *m*₁ $\wedge$ *m*₂, *m*₁ $\wedge$ $\neg$ *m*₂, and *m*₂ $\wedge$ $\neg$ *m*₁; however, all associations must be mutually role consistent in the sense that, if *m*_i is associated with up-regulated genes in one condition, then *m*_i may not be associated with down-regulated genes in another condition and $\neg$ *m*_i may not be associated with up-regulated genes in any condition. (iv) Flexible logic: for any motif pair (*m*₁ and *m*₂), associations with different conditions may involve any of the following motif modules: *m*₁, *m*₂, *m*₁ $\wedge$ *m*₂, *m*₁ $\wedge$ $\neg$ *m*₂, and *m*₂ $\wedge$ $\neg$ *m*₁. No additional constraints are imposed here in defining a valid meta-association.

Biclustering and Motif Associations. We created a binary membership matrix with rows for each of 9,272 genes and columns for each of 22 DEG sets (the fast and slow genes of each of the 11 maturation determinants). With the R package biclust, we ran the Bimax Biclustering method (37) on the membership matrix, searching for submatrices of logical ones; 492 biclusters were discovered with this method while requiring the resulting biclusters to contain at least five genes and at least three DEG gene sets. For the genes contained in each bicluster, we tested for enrichment in each of the 602 motifs in our collection as well as every logical combination of motif pairs. EVD *P* values were calculated reusing the 492 biclusters and shuffling the motif target gene sets. We also created 12 SAMBA (9) biclusters using the built-in tool in the Expander software (version 5 for Windows). We performed the SAMBA biclustering on log₂ fold change values of the genes from the 11 maturation experiments. Missing data were assigned a value of zero. Although we tried several settings of the parameters (overlap prior factor, hash kernel size, option file type, and number of genes corresponding to hash), we were most satisfied with the biclusters produced from the default SAMBA parameters (0.1, 4, valsp_3ap, and 100.). Each of our 12 SAMBA biclusters contained between 19 and 90 genes and covered four or five maturation determinants. Association tests and EVD analysis were performed using the same approach as with the Bimax biclusters.

ACKNOWLEDGMENTS. We thank K. S. Pruiett for beekeeping; T. C. Newman for assistance with molecular techniques; C. J. Fields for mapping cDNA array features to the oligo array; J. J. Collins, V. Jongeneel, J. Ma, P. Ma, N. D. Price, L. J. Stubbs, and S. Zhong for reviewing the manuscript; E. Carefoot for producing the bee cartoons; and O. Georgiou for sharing a C routine used in our implementation of meta-*P* values. Funding was from National Science Foundation Frontiers in Biological Research Award EF25852 (B. Schatz, principal investigator); National Institutes of Health Director's Pioneer Awards 1DP1OD006416 (to G.E.R.), 1R01DK082605-01A1 (to G.E.R.), and 1R01GM085233-01 (to S.S.); and National Science Foundation Grant DBI-0746303 (to S.S.). Table S4 shows accession numbers for each experiment (<http://www.ebi.ac.uk/microarray-as/ae/>).

- West-Eberhard MJ (2003) *Developmental Plasticity and Evolution* (Oxford Univ Press, New York).
- Robinson GE, Fernald RD, Clayton DF (2008) Genes and social behavior. *Science* 322: 896–900.
- Whitfield CW, Cziko AM, Robinson GE (2003) Gene expression profiles in the brain predict behavior in individual honey bees. *Science* 302:296–299.
- McLeay RC, Bailey TL (2010) Motif Enrichment Analysis: A unified framework and an evaluation on ChIP data. *BMC Bioinformatics* 11:165.

- Warner JB, et al. (2008) Systematic identification of mammalian regulatory motifs' target genes and functions. *Nat Methods* 5:347–353.
- Sinha S, Adler AS, Field Y, Chang HY, Segal E (2008) Systematic functional characterization of cis-regulatory motifs in human core promoters. *Genome Res* 18: 477–488.
- Subramanian A, et al. (2005) Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci USA* 102:15545–15550.

8. Smith DD, et al. (2008) Meta-analysis of breast cancer microarray studies in conjunction with conserved cis-elements suggest patterns for coordinate regulation. *BMC Bioinformatics* 9:63.
9. Tanay A, Sharan R, Kupiec M, Shamir R (2004) Revealing modularity and organization in the yeast molecular network by integrated analysis of highly heterogeneous genomewide data. *Proc Natl Acad Sci USA* 101:2981–2986.
10. Reiss DJ, Baliga NS, Bonneau R (2006) Integrated biclustering of heterogeneous genome-wide datasets for the inference of global regulatory networks. *BMC Bioinformatics* 7:280.
11. Halperin Y, Linhart C, Ulitsky I, Shamir R (2009) Allegro: Analyzing expression and sequence in concert to discover regulatory programs. *Nucleic Acids Res* 37:1566–1579.
12. Huttenhower C, et al. (2009) Detailing regulatory networks through large scale data integration. *Bioinformatics* 25:3267–3274.
13. Rice D, Barone S, Jr. (2000) Critical periods of vulnerability for the developing nervous system: Evidence from humans and animal models. *Environ Health Perspect* 108(Suppl 3):511–533.
14. Wilson EO (1971) *The Insect Societies* (Belknap Press, Cambridge, MA).
15. Robinson GE (2002) Genomics and integrative analyses of division of labor in honeybee colonies. *Am Nat* 160(Suppl 6):S160–S172.
16. Whitfield CW, et al. (2006) Genomic dissection of behavioral maturation in the honey bee. *Proc Natl Acad Sci USA* 103:16068–16075.
17. Ament SA, Wang Y, Robinson GE (2010) Nutritional regulation of division of labor in honey bees: Toward a systems biology perspective. *Wiley Interdiscip Rev Syst Biol Med* 2:566–576.
18. Nelson CM, Ihle KE, Fondrk MK, Page RE, Amdam GV (2007) The gene vitellogenin has multiple coordinating effects on social organization. *PLoS Biol* 5:e62.
19. Sinha S, Ling X, Whitfield CW, Zhai C, Robinson GE (2006) Genome scan for cis-regulatory DNA motifs associated with social behavior in honey bees. *Proc Natl Acad Sci USA* 103:16352–16357.
20. Alaux C, et al. (2009) Regulation of brain gene expression in honey bees by brood pheromone. *Genes Brain Behav* 8:309–319.
21. Grozinger CM, Sharabash NM, Whitfield CW, Robinson GE (2003) Pheromone-mediated gene expression in the honey bee brain. *Proc Natl Acad Sci USA* 100(Suppl 2):14519–14525.
22. Chandrasekaran S, et al. (2011) Behavior-specific changes in transcriptional modules lead to distinct and predictable neurogenomic states. *Proc Natl Acad Sci USA* 108:18020–18025.
23. Lutz CC, Rodriguez-Zas SL, Fahrback SE, Robinson GE (2012) Transcriptional response to foraging experience in the honey bee mushroom bodies. *Dev Neurobiol* 72:153–166.
24. Segal E, Friedman N, Koller D, Regev A (2004) A module map showing conditional activity of expression modules in cancer. *Nat Genet* 36:1090–1098.
25. Pankiw T, Page RE (2001) Genotype and colony environment affect honeybee (*Apis mellifera* L.) development and foraging behavior. *Behav Ecol Sociobiol* 51:87–94.
26. Bell AM, Robinson GE (2011) Genomics. Behavior and the dynamic genome. *Science* 332:1161–1162.
27. Giray T, et al. (2000) Genetic variation in worker temporal polyethism and colony defensiveness in the honey bee, *Apis mellifera*. *Behav Ecol* 11:44–55.
28. FitzGerald PC, Sturgill D, Shyakhthenko A, Oliver B, Vinson C (2006) Comparative genomics of *Drosophila* and human core promoters. *Genome Biol* 7:R53.
29. Alaux C, et al. (2009) Honey bee aggression supports a link between gene regulation and behavioral evolution. *Proc Natl Acad Sci USA* 106:15400–15405.
30. Withers GS, Fahrback SE, Robinson GE (1993) Selective neuroanatomical plasticity and division of labour in the honeybee. *Nature* 364:238–240.
31. Huang ZY, Robinson GE (1992) Honeybee colony integration: Worker-worker interactions mediate hormonally regulated plasticity in division of labor. *Proc Natl Acad Sci USA* 89:11726–11729.
32. McClung CA, Nestler EJ (2008) Neuroplasticity mediated by altered gene expression. *Neuropsychopharmacology* 33:3–17.
33. Yoshida H, Matsui T, Yamamoto A, Okada T, Mori K (2001) XBP1 mRNA is induced by ATF6 and spliced by IRE1 in response to ER stress to produce a highly active transcription factor. *Cell* 107:881–891.
34. Koo JW, Russo SJ, Ferguson D, Nestler EJ, Duman RS (2010) Nuclear factor-kappaB is a critical mediator of stress-impaired neurogenesis and depressive behavior. *Proc Natl Acad Sci USA* 107:2669–2674.
35. Riddiford LM (2008) Juvenile hormone action: A 2007 perspective. *J Insect Physiol* 54:895–901.
36. Larsson NG, et al. (1998) Mitochondrial transcription factor A is necessary for mtDNA maintenance and embryogenesis in mice. *Nat Genet* 18:231–236.
37. Prelic A, et al. (2006) A systematic comparison and evaluation of biclustering methods for gene expression data. *Bioinformatics* 22:1122–1129.
38. Renn SCP, Aubin-Horth N, Hofmann HA (2008) Fish and chips: Functional genomics of social plasticity in an African cichlid fish. *J Exp Biol* 211:3041–3056.
39. Ament SA, et al. (2012) The transcription factor *ultraspiracle* influences honey bee social behavior and behavior-related gene expression. *PLoS Genet* 8:e1002596.
40. Pick L (1998) Segmentation: Painting stripes from flies to vertebrates. *Dev Genet* 23:1–10.
41. Rhodes DR, Barrette TR, Rubin MA, Ghosh D, Chinnaiyan AM (2002) Meta-analysis of microarrays: Interstudy validation of gene expression profiles reveals pathway dysregulation in prostate cancer. *Cancer Res* 62:4427–4433.
42. Tchagang AB, Pan Y, Famili F, Tewfik AH, Benos PV (2010) *Handbook of Research on Computational and Systems Biology: Interdisciplinary Applications* (IGI Global, Hershey, PA), pp 148–186.
43. Bailey TL, Gribskov M (1998) Combining evidence using p-values: Application to sequence homology searches. *Bioinformatics* 14:48–54.
44. Ulitsky I, Shamir R (2007) Identification of functional modules using network topology and high-throughput data. *BMC Syst Biol* 1:8.
45. Müller FJ, et al. (2008) Regulatory networks define phenotypic classes of human stem cell lines. *Nature* 455:401–405.
46. Zhou XJ, et al. (2005) Functional annotation and network reconstruction through cross-platform integration of microarray data. *Nat Biotechnol* 23:238–243.
47. Rhodes DR, et al. (2004) Large-scale meta-analysis of cancer microarray data identifies common transcriptional profiles of neoplastic transformation and progression. *Proc Natl Acad Sci USA* 101:9309–9314.
48. McCarroll SA, et al. (2004) Comparing genomic expression patterns across species identifies shared transcriptional profile in aging. *Nat Genet* 36:197–204.
49. Hong FX, Breitling R (2008) A comparison of meta-analysis methods for detecting differentially expressed genes in microarray experiments. *Bioinformatics* 24:374–382.
50. Ramasamy A, Mondry A, Holmes CC, Altman DG (2008) Key issues in conducting a meta-analysis of gene expression microarray datasets. *PLoS Med* 5:e184.
51. Perry MW, Boettiger AN, Levine M (2011) Multiple enhancers ensure precision of gap gene-expression patterns in the *Drosophila* embryo. *Proc Natl Acad Sci USA* 108:13570–13575.
52. Kazemian M, et al. (2010) Quantitative analysis of the *Drosophila* segmentation regulatory network using pattern generating potentials. *PLoS Biol* 8:e1000456.
53. Honeybee Genome Sequencing Consortium (2006) Insights into social insects from the genome of the honeybee *Apis mellifera*. *Nature* 443:931–949.
54. Dettmann CP, Georgiou O (2009) Product of n independent uniform random variables. *Stat Probab Lett* 79:2501–2503.
55. McLean CY, et al. (2010) GREAT improves functional interpretation of cis-regulatory regions. *Nat Biotechnol* 28:495–501.
56. Ng P, Keich U (2008) GIMSAN: A Gibbs motif finder with significance analysis. *Bioinformatics* 24:2256–2257.
57. Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J R Stat Soc B Stat Methodol* 57:289–300.
58. Whitlock MC (2005) Combining probability from independent tests: The weighted Z-method is superior to Fisher's approach. *J Evol Biol* 18:1368–1373.
59. Stouffer SA, Suchman EA, DeVinney LC, Star SA, Williams RM, Jr. (1949) *The American Soldier, Vol. 1: Adjustment During Army Life* (Princeton Univ Press, Princeton).
60. Rice WR (1990) A consensus combined P-value test and the family-wide significance of component tests. *Biometrics* 46:303–308.
61. Tippet LH (1931) *The Methods of Statistics* (Williams and Norgate, London).
62. Storey JD, Taylor JE, Siegmund D (2004) Strong control, conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *J R Stat Soc B Stat Methodol* 66:187–205.
63. Storey JD, Tibshirani R (2003) Statistical significance for genomewide studies. *Proc Natl Acad Sci USA* 100:9440–9445.
64. Stark A, et al. (2007) Discovery of functional elements in 12 *Drosophila* genomes using evolutionary signatures. *Nature* 450:219–232.