Abstract Eigenstuff

Square matrices are coordinate representations of automorphisms. Automorphisms are simply linear transformations from a vector space V into itself. Given a basis $\mathscr{B} = \{v_1, v_2, \ldots, v_n\}$ for V, any $v \in V$ can be written uniquely as a linear combination of the v's:

$$v = \sum_{i=1}^{n} c_i v_i.$$

The coordinate vector of v wrt $\mathscr{B}$ is $[v]_{\mathscr{B}} = (c_1, c_2, \ldots, c_n) \in \mathbb{R}^n$ (or $\mathbb{C}^n$ or such).

Now if $L : V \to V$ is a linear transformation, we can represent L as a matrix by looking at the coordinate vectors of its action: $[Lv]_{\mathscr{B}} \equiv [L]_{\mathscr{B}}[v]_{\mathscr{B}}$. We take this to be the definition of $[L]_{\mathscr{B}}$ and since $[v_k]_{\mathscr{B}} = e_k$, we can find $A = [L]_{\mathscr{B}} = [a_1, a_2, \ldots, a_n]$ as

$$a_k = Ae_k = [Lv_k]_{\mathscr{B}}.$$

Why talk about coordinates here? If we had chosen a different basis, say $\mathscr{D} = \{w_1, w_2, \ldots, w_n\}$, then $B = [L]_{\mathscr{D}}$ would have columns $b_k = Be_k = [Lw_k]_{\mathscr{D}}$. Now A and B are two different matrices representing the same transformation. Coordinates (matrices) are coordinates wrt a basis, but L is L, is L. How do we know L except through A or B? A and B are clearly related to each other, but how? Is every square matrix a representation of L wrt some basis? Is there a basis for V that displays L in a particularly suggestive way?

These questions suggest we partition the $n \times n$ matrices into equivalence classes given by automorphisms. Two matrices A and B are *similar* if they are coordinate representations of the same automorphism. However you first met similarity, this is her true face. The standard definition of similarity (matrices A and B are similar if there exists S with $A = S^{-1}BS$) follows from our definition and the basis discussion above. So what do all representations of L have in common? Properties of L that are shared by all matrix representation of L are called *similarity invariants.*

The fundamental similarity invariants for an automorphism are the subspaces that it doesn't move. Let $U \le V$. If $Lu \in U$ for every $u \in U$, then U is an invariant subspace for L. Some invariant subspaces have large dimension (e.g. V itself is an invariant subspace for every L), and some are small (e.g. 0 is an invariant subspace for every L). The smallest interesting (nontrivial) invariant subspaces have dimension 1, and for most operators, all invariant subspaces (including V itself) can be built from these.

If U is an invariant subspace of L of dimension 1, then every $u \in U$ is of the form αu_0, where α is a scalar and u_0 is *any* nonzero vector in U. Then saying $Lu \in U$ is the same as saying $L(\alpha u_0) = \beta u_0$. But L is linear, so we can pull out the scalars and let $\lambda = \beta/\alpha$, giving

$$Lu = \lambda u.$$