Expertise Tracking

Raymond D'Amore
The MITRE Corporation
McLean, VA
rdamore@mitre.org

ABSTRACT
Information sharing and collaboration are critical to effective analysis; especially where the focus crosses mission areas and organizational boundaries. While, new methods are emerging to handle large-scale, multi-source data; awareness tools that track work and support coordination lag. This paper sets the context for using expertise detection and organizational analysis as enablers for the next generation of awareness tools operating at the group and community levels. Initial results from the *Expert Locator* prototype are described as well as exploratory efforts to exploit *Expert Locator* functionality into team building, contact network generation, and work tracking tools; all useful in coordinating analysis.

1. INTRODUCTION
The problem of effectively exploiting massive information flows within an organization is compounded by the fundamental question of "who knows what", Hinds and Pfeffer (2002). It is becoming increasingly difficult for organizations *to know what they know*. In diverse organizations, expertise may be widely distributed and not restricted to only a few individuals Huber, (1999). Further, finding a single expert may be suboptimal as it may not address diversity across experts and preclude organization-wide consensus building. As such, experts can play critical roles in assessing recommendations from other experts; this is a critical need in many areas such as technology assessments, product testing, and intelligence analysis.

There are a number of systems that make attempt to make expertise explicit and promote sharing using standard skill registries or electronic marketplaces. While there is value in encoding "expertise" the supporting process is not scalable; especially in large heterogeneous, dynamic environments. In part as a response to the cost to build and maintain expert repositories, there is increased interest in automating the expert finding task and this has given rise to a new class of search engines known as *expert finders*, Yimam, (1999). Many expert finder systems are architecturally similar to standard retrieval systems but use relatively simple counting schemes to relate the number of relevant artifacts to the likelihood that a person is an expert. These systems use automatic indexing techniques to exploit publications, email, and other artifacts as sources of expertise indicators. While this approach requires some knowledge of the tie between a person and an artifact (e.g., author), they tend to incorporate shallow expertise models, and do not easily extend to handle multiple sources of evidence. In addition, while there has been some work to characterize the emergence of expertise groups or communities D'Amore in Maybury, M., D'Amore, R. and House, D. (2002), a more structured framework for evidence combination is needed to detect isolates, groups, and communities-of-practice.

2. Expertise
Merriam-Webster[1] defines an expert as "one with the special skill or knowledge representing mastery of a particular subject". While this definition is elegant in its simplicity, it belies the true complexity as reflected in the extensive literature on the nature of expertise. Chi, M. T. H., Hutchinson, J. E., and Robin, A. F., (1989) focused on the definition of knowledge structure

[1] Merriam Webster reference goes here.

within a specific domain, and the relationship between structure and use. Glaser, R., (1986) was able to show that high levels of competence result from the interaction between knowledge structure and processing capabilities. Bedard, J., and Chi, M. T. H., (1992) assessed the influence of domain knowledge on perceptual processes and strategies in problem solving. However, while there is some cross-study convergence regarding the relationship between domain knowledge, knowledge structuring, and processing methods, the performance of experts on certain tasks is more variable.

Performance-related research is divided as to how well experts perform on a range of tasks. Research in the decision sciences suggests that experts perform poorly across a number of decision analysis tasks. Experts make flawed decisions and employ heuristics that introduce significant biases in the analysis task. Foss, Wright, and Coles (1975) discussed the low validity of expert assessments in judging livestock, even when compared to novices. In Dawes and Corrigan (1974) experts were shown to under perform simple linear models across a range of forecasting problems. While experts were effective in determining the key variables or factors in the prediction problem they often relied on heuristics and that resulted in a number of biases such as anchoring and availability, Kahneman, Slovic, and Tversky (1982). On the other hand, cognitive science research suggests that experts are competent and have both knowledge and functional skills that are distinct from novices. In this area, much of the literature on expertise modeling is domain specific, Chase and Simon (1973), and emphasizes experts' use of domain knowledge and specific processing methods as key discriminators between experts and non-experts; Chi, Glaser, and Farr (1988). The difference in findings between decision science and cognitive research suggests that other factors may be involved.

Shanteau, (1992) suggests that the different view held by decision scientists and cognitive scientists is explained by differences in task characteristics. Essentially, Shanteau presents a "theory of expert competence" that suggest that both analyses are correct but incomplete. He lays out five components of competence, (sufficient domain knowledge, psychological traits, cognitive skills needed to make decisions, use of appropriate decision strategies, and tasks characteristics) and concludes that the difference between the decision science and cognitive science literatures is related to differences in task domains studied. While expert performance may vary with task characteristics, there are some general behaviors that are transferable across domains.

While, much of the literature related to expert's behavior focuses on domain-dependency in problem knowledge and methodology used, there is a behavioral constant: *experts signal their expertise*. Experts exhibit behaviors consistent with making explicit there skill areas. A key hypothesis here is that experts signal expertise much like firms do. Experts may advertise their expertise through artifacts produced, honorifics, roles, and by embedding themselves within expertise networks, they establish reputation and build trust.

3. Expertise Modeling

Detecting experts and tracking their behavior requires a network view of the underlying workspace. Evidence of experts in an enterprise may be scattered across a number of dynamically changing online forums to include organizational websites, personal home pages, project workspaces, news groups, chat rooms, email, and others. The premise here is that corporate workspaces can be used to glean expertise. The enterprise model developed is general in that observable work (expert signaling) is embedded within an organizational mesh. Notionally, portions of the mesh are selected for instantiation in the enterprise model based on whether the workspace supports: collection, feature extraction, and corporate privacy policies. Here we use activity spaces D'Amore (2004), D'Amore et al (2003) to establish work space context. Activity spaces are characterized by actors, events, interactions, and artifacts.

For example, ListServ discussion threads provide postings and links between posters that may be used as social context. Essentially, artifact evidence and social context are signaling mechanisms that provide a basis for measuring expertise. This follows, in part, from the notion that expertise (or trust) is communicated through structural or relational embeddedness, Granovetter (1973): experts tend to work within groups or communities-of-practice consistent with their area of specialization and in doing so signal their expertise.

The expertise model developed, D'Amore (2004) combines query-specific artifact evidence and social context to infer expertise. In the basic model, Equation 1, social context and artifact evidence are weighted as a function of the distribution of artifacts with an activity space and social structure or ties between embedded experts. Evidence from each activity space is "fused" into an overall ranking using an information-theoretic weighting. Here,

$$I(p \mid q) = \sum \alpha_i \cdot E_{i,\bullet,\bullet,p} \qquad (1)$$

where $I(p \mid q)$ is the importance of person, p, for query, q, α_i is the weight assigned to activity space, i. $E_{i,\bullet,\bullet,p}$ is the aggregate (artifact and social) evidence for all subspaces within all activity spaces (for example, a subspace might be a particular project within the project activity space) and is computed as

$$E_{i,\bullet,\bullet,p} = \sum_{\forall k} \omega_k \bullet \sum_{\forall j,k} \beta j \bullet e_{i,j,k,p} \qquad (2)$$

where $e_{i,j,k,p}$ is the evidence type, k, associated with a particular person, p, within a particular subspace of an activity space; ω_k is the weight assigned to evidence type, k, and βj is the weight for subspace, j. Composite expert score ranks from each space are merged using a linear fusion model:

$$I(p \mid q) \Rightarrow I_R(p \mid q) \qquad (3)$$

where,

$$I_R(p \mid q) = \mathrm{N}^{\gamma} \bullet \sum \alpha_i \bullet B(Ei,\bullet,\bullet,_p) \qquad (4)$$

Note that equation (4) follows from equation (1), where raw scores are transformed into Borda counts, $B(Ei,\bullet,\bullet,_p)$, Montague, and Aslam, (2002), and N^{γ} is the number of populated activity spaces to the power (-1,0, or 1), as in CombMNZ , Fox et al, (1993), α_i, equation (5), is an information-theoretic weighting generally applied in ecological modeling McCune and Grace, (2002), that here assigns weights of importance to each activity space based on the distribution of experts across spaces in the current query.

4. Evaluation

The evaluation uses snowball sampling to establish relevance sets for each query tested. The HITS algorithm, Kleinberg (1999), was used to assign a relevance weight to each graph member (based on various combinations of authority or hub scores). *Expert Locator* rankings were then compared to the snowball generated relevance rankings. The overall results (N=32 queries) are presented in the table, below. The mean *R-precision* was 37% across all test queries. Hub detection was on average higher than Authority detection. However, since user acceptance is biased towards high precision, the mean Pr (Top 5) is more revealing as it measures the likelihood that any person ranked in the top five is a known expert; here, precision is between 65% and 78%.

	mean r-prec	mean Pr(Top 5)
Auth	0.319	0.650
Hub	0.429	0.544
A|H	0.365	0.781

A sensitivity test was run to test system robustness to variation in the number of activity spaces used. Using *R-precision*, *Expert Locator* performance was always lowest when a single space was used; for example ListServs. The best results were obtained when using all three activity spaces, although roughly 1/3 of the time use of two spaces outperformed three spaces.

5. New Applications

Expert Locator is being assessed for use in several new applications to include team building, expertise resource allocation, and domain detection and tracking. There is special interest in detecting the emergence of new technology areas within the enterprise. For example, interest in Biocomputing may emerge in niche areas, say, with a strong focus on the core science areas of "biology"; however, over time the *core* work may extend to include supporting technologies on the *periphery*. Figure 1 depicts the intersection between two domains: *biology* and *language processing*. This is essentially a collaboration graph reflecting the coevolution of the two expertise domains. The blue and yellow colored nodes represent scientists from one domain or the other and edges between represent collaboration (joint project work). Possibly more interesting are the green nodes, centrally positioned in the network, which represent scientists with specialties in both domains. This is a form of "organizational learning" where several scientists have developed a new specialty; e.g., a language processing expert gaining some expertise in the biological sciences. Ongoing work is focused on identifying key players or groups within single or multi-domain work environments as a basis for risk assessment (i.e., critical skill areas), resource allocation, and information dissemination.

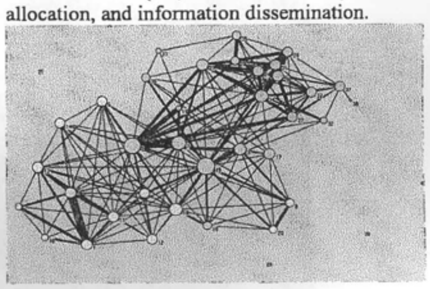

Figure 1: The Intersection of Biology and Language Processing expertise areas. Individual classified as having expertise in both domains are colored green. Node size reflects betweeness centrality—how central they are to information flow. Edge thickness is proportional to levels of joint work between pairs of experts.

6. Conclusions

An expert finder has been developed and evaluated within an enterprise environment. The results are promising in terms of overall retrieval performance, the potential applicability of the expertise model to new environments, and in the potential for using expertise detection as an enabler for organizational analysis and coordinate work or analysis. Planned future work, is focused on extending expert finding and organizational analysis to address analysis coordination: tracking analysis tasks, analysts roles and skills, and dynamic resource allocation.

REFERENCES

[1] References forthcoming pending public release acceptance. Final submission will include all references.