Robustness and Adaptation in "Information Ecosystems"

Stephen Racunas, Stanford University

Abstract

We have been developing techniques for monitoring and "proofreading" knowledge resources, and these techniques proved to scale remarkably well. We are easily able to dynamically proofread the knowledge base containing all known and curated human metabolic pathways at speeds that enable real-time human interaction with both the data and with continuous evaluations of that data's internal consistency and reliability.

Recently, both manual and automated techniques for the attachment of machine-understandable meta-data to incoming knowledge streams have proved to be effective ways of verifying and cross-checking information validity at multiple levels of resolution during the evaluation of medical and biological data. We illustrate the power of such techniques by detailing our construction of model verification and hypothesis evaluation software for S. cerevisiae.

We outline how one might combine these two technologies to achieve a context-aware system capable of recognizing and mitigating both accidental and malicious information loss and conflict. By attaching meta-data to broad classes of information records to indicate what an adversary would be able to do with such information if it were compromised and what an adversary might do to such information if it were in fact misrepresented, we can use our techniques to automatically formulate and evaluate hypotheses about historical and potential threats.

As attacks become more commonplace and information infrastructure becomes more complex, we believe it will make less and less sense to try to eradicate, forstall or backtrace each individual attack. Organizations produce information which is of value to others. Various allied entities must consume this information in a timely manner for both participants to function optimally, and various antagonistic entities must be prevented from interfering. Rather than considering only individual information transactions, we believe it makes sense to apply biologically-inspired techniques for evaluating and maximizing the fitness of information strategies on the level of entities competing for resources – the level of the "information ecosystem."