# Research Statement

Will Bridewell

The holy grail of medical informatics rests in the holistic patient model, currently incarnated as a variety of electronic health records. The predominant ideal would let physicians and clinical researchers examine patients from every perspective at every moment in time. Additional technologies would extract the contextually relevant details, building a variety of meaningful views of each patient. The reality is a fractured collection of records, some well structured, others in free form, often composed of missing, erroneous, and contradictory information. My goal is to fashion from this bundle of noise and signal computational artifacts that better approximate the manipulable patient model.

Thankfully the growing prevalence of electronic health records has led to an abundance of computationally available data, raising the obvious question, "What do we do with all this information?" In a perfect world, we could form the holistic patient model from the contents of the structured fields, reconstructing an accurate view for use in later analyses. For numerous reasons, the structured data fail to capture many of the important clinical details, leaving us with little choice but to combine them with unstructured physician notes. Typical solutions to this challenge emphasize concept extraction and range from rudimentary string matching to sophisticated natural language processing. In any case, the goal is to find the relevant details in the haystack of the clinical narrative, blending the results with the well formed data.

While in graduate school, I noticed that early work on tapping the unstructured content of electronic health records suffered from a severe oversight. Although these systems were proficient in extracting clinical concepts from free text, they almost uniformly failed to distinguish between truly present concepts and those that appear in a negative context (e.g., "rules out pneumonia"). The mechanisms for systems that could reason about negation were opaque. Seeing this state of affairs, my colleagues and I developed a novel approach for identifying negated concepts. This system, called *NegEx*, relies on the predictability of clinical language and searches for phrasal patterns that indicate negated concepts (e.g., "rules out X"). To date, this work has been highly cited and the clarity of the solution has inspired others both to extend NegEx and to develop new solutions. Furthermore, a group within the University of Pittsburgh Medical Center adopted the system to service the hospital community. This research improved the accuracy of concept extraction and as a result, the patient model.[1]

Of course, creating coherent views of patients requires more than a collection of concepts and numbers. In particular we need to conjoin the individual atoms of information to construct a model of the person. To this end, I have worked to transition disconnected data into cohesive explanations. My research on *inductive process modeling*, undertaken in the Computational Learning Laboratory at Stanford University, addresses this challenge through systems that automate model construction and revision. Given multivariate time-series and domain-specific knowledge, these systems search the dual space of candidate mathematical structures and their numerical parameters to discover the most plausible explanations of the data. The underlying process formalism, which combines high-level conceptual knowledge with lower-level mathematical expressions, lets researchers readily draw from libraries of predefined components or easily create their own. Each application of the system produces a list of ranked models that account for the data in terms of established theoretical mechanisms.[2]

The ability to search large spaces of potential explanations brings a common obstacle to the forefront: models of man and molecule are underdetermined by measurable data. As a result, for any one system described by a set of observations, we may find multiple equally accurate and plausible models. Here I saw opportunity. Adapting techniques from machine learning, my colleagues and I developed approaches to uncover the commonalities in the best and worst explanations that let us discover constraints on the form of a solution. While the exact set of processes controlling a dynamic system may elude detection, we can identify those that must appear in a reasonable model and those that should not. Most recently, we have developed techniques to detect mutually exclusive processes and those that must appear together. As a result, even when engulfed in uncertainty, we can identify deep knowledge about complex systems. This approach mitigates the hazards of missing and noisy data in the patient record.[3]

Moving forward, models offer a principled approach to data integration. By connecting models that cut across systems or levels of biological organization, we can effectively integrate data from multiple sources. Consider clinical and molecular data that are quite different and disconnected now. The future promises to bridge that gap, and in this case we can see explanations of mechanisms that connect metabolic pathways to physiologic processes. Moreover, models provide a principled approach for connecting knowledge (e.g., quantitative processes, semantic web ontologies) to data. These rigid structures instantiate general knowledge to explain scientific measurements and other observations, restricting potential interpretations to ones fitting a plausible narrative. In the qualitative case of natural language, models in the forms of plans (general behaviors, guidelines, etc.) help disambiguate the intended meaning of a phrase, providing strong semantic and pragmatic constraints.

Following this direction, at the Stanford Center for Biomedical Informatics Research my interests turned back toward unstructured data exemplified by natural language. Here the driving challenge is to build a system that can navigate and engage in human-level dialog, providing expert advice to combat medics in the field. My piece of the puzzle is a system that can reason about the beliefs, goals, and intentions of itself and other agents and use that information to fill in the gaps implied during conversations. To this end, I am developing a general approach to *socially aware inference*. Potential uses of this research extend beyond the admittedly lofty goals of this current project. Specifically, I intend to apply the developed system to reconstruct treatment plans from clinical narratives, a procedure that requires the joint use of billing codes and unstructured text. Fortunately an extension of NegEx should help reduce errors when extracting concepts related to the physician's plans. Furthermore, much like the knowledge in inductive process modeling, the information stored in well developed clinical guidelines can help define the space of potential solutions and protect against missing and conflicting information. Success on this front will enable the reconstruction of treatment plans, a crucial component of the patient record otherwise buried in text.[4]

Extending the notion of socially aware inference to the macro scale, I continue to work closely with Professor Amar Das to investigate the implications of a physician's referral network on clinical practice variation. This project is in its early stages, but my ultimate goal is to build a system that guides patients toward physician networks whose care philosophies match their own. Along the way, I plan to construct predictive models of patient care, tracing their movement through the social network coupled with the treatments and services received. Moreover, this research should uncover the influence of social interactions on patterns of technology and procedure adoption.[5]

Returning to the holistic patient model, multiple obstacles—computational, political, social, and financial—impede success, but current electronic health records contain a wealth of information that remains untapped by existing informatics systems. The approaches described above all contribute to richer views of patients capitalizing on available data, potentially leading to insights about individual cases useful for clinical research and practice. This broad vision depends on systems that can consider patients from an assortment of perspectives and assemble plausible explanations, plans, or narratives that tell each patient's story. Ultimately, I look forward to a time when medical records move beyond an accounting of facts to interactive and even collaborative tools for exploration and discovery.

## Notes

[1]See Chapman, W.W., Bridewell, W., Hanbury, P., Cooper, G.F., Buchanan, B.G. (2001). A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of Biomedical Informatics*, 34, 301–310.

[2]See Bridewell, W., Langley, P., Todorovski, L., Dzeroski, S. (2008). Inductive process modeling. *Machine Learning*, 71, 1–32.

[3]Bridewell, W., Borrett, S.R., Todorovski, L. (2007). Extracting constraints for process modeling. *Proceedings of the Fourth International Conference on Knowledge Capture*, 87–94.

[4]See the preliminary report on socially aware inference in Bridewell, W., Langley, P. (2011). A computational account of everyday abductive inference. *Proceedings of the Thirty-Third Annual Conference of the Cognitive Science Society*, 2289–2294.

[5]See early findings in Bridewell, W., Das, A.K. (2011). Social network analysis of physician interactions: the effect of institutional boundaries on breast cancer care. *Proceedings of the 2011 AMIA Annual Symposium*, 152–160.