

Experimental Assessment of a Process-Workflow Redesign Agent

Mark E. Nissen

Naval Postgraduate School

MNissen@nps.navy.mil

Abstract

Today, a second generation of computer-based reengineering tools employ knowledge systems technology to automate and support key, intellectual activities required for effective process-workflow redesign. But a central question remains as to the effectiveness of redesign through such knowledge systems. The research described in this paper is focused on testing the effectiveness of knowledge-based, process-workflow redesign systems. We employ one such system called "KOPeR-lite" as a platform for experimentation to assess the relative efficacy of redesigns generated by computer versus those developed by people. In this sense, we conduct a modified Turing Test to compare redesign performance of reengineering analysts with that of the knowledge system. A number of important conclusions have implications for process-workflow redesign and practice, including issues associated with leading adoption of knowledge system technology and extension of redesign automation systems such as KOPeR-lite.

1. Process-workflow redesign support

As we approach the ten-year mark for business process reengineering (BPR)—based on seminal articles by [8] and [14]—the phenomenon has lost much of the hyperbole and attention in the management press that surrounded its early years. Yet few would argue the opportunities for performance improvement have been exhausted or the need for process redesign has diminished. This is particularly the case in the area of workflow automation, as many firms are simply inserting enterprise workflow technology into their operations without first redesigning the underlying processes themselves. This practice is colorfully referred to as "paving the cowpaths" [14].

Today, a second-generation of computer-based reengineering tools employ knowledge systems technology to automate and support key, intellectual activities required for effective process-workflow redesign. By "process workflow," we mean the ordering, sequencing, organization and technology associated with the flow of work through an enterprise process. This is

consistent with the process view articulated in [10] and corresponds with the customary reengineering focus on the process as a central unit of analysis (e.g., see [6, 15]).

Unlike first-generation BPR tools, knowledge systems directly address the kinds of deep, distributed knowledge and specialized expertise associated with knowledge work such as process redesign [34], and they can augment first-generation tools to reduce reengineering time and expense [16]. Important questions associated with feasibility of such tools and their effect on reengineering efficiency are now being addressed through research, in which the use and utility of process-workflow redesign systems is demonstrated in "industrial strength" reengineering engagements (e.g., see [31]). A number of such reengineering engagements now have knowledge-based tools employed to analyze and redesign major processes, some consisting of a thousand activities or more and many hundreds or thousands of employees. It is in these, major reengineering projects that the potential of second-generation BPR tools appears greatest.

But a central question remains as to the effectiveness of process-workflow redesign through such knowledge systems. It does little good to quickly complete a poor redesign, for example. Another question pertains to the experience level to be supported through process-workflow redesign systems. For example, do such systems offer greater potential as "assistants" to reengineering experts or as "experts" themselves to support relative novices in the reengineering domain? And in which BPR contexts and environments should a process-workflow redesign system be preferred to a human, and vice versa?

The research described in this paper is focused on testing the effectiveness of knowledge-based, process-workflow redesign systems. We employ one such system called "KOPeR-lite" as a platform for experimentation to assess the relative efficacy of process-workflow redesigns generated by computer versus those developed by people. In this sense, we conduct a modified Turing Test (see [35]) to compare redesign performance of reengineering analysts with that of the knowledge system. In a set of experiments toward this end, the first of which is described in this paper, knowledge system performance is compared directly with that of reengineering novices who employ conventional redesign methods and tools. We

focus initially on this group because, *ex ante*, novices appear to have the most to gain through redesign automation tools such as KOPeR-lite. For instance, absent some kind of knowledge-based system, the novice is almost totally dependent upon external BPR consulting services. In follow-on experiments—beyond the scope of this paper—we are also interested in examining human versus machine redesign performance at the expert level, for clearly, most BPR consultants perform at this expert level.

Following this introduction, we provide background information pertaining to redesign automation and support through the KOPeR-lite application and outline the experimental design conceived to test its redesign efficacy. We then report experimental findings and generalize from laboratory results to professional redesign practice in the field. A number of important conclusions have implications for process-workflow redesign and practice, and we use these conclusions to refine an agenda for further experimental research and knowledge system development in the workflow area.

2. The KOPeR application

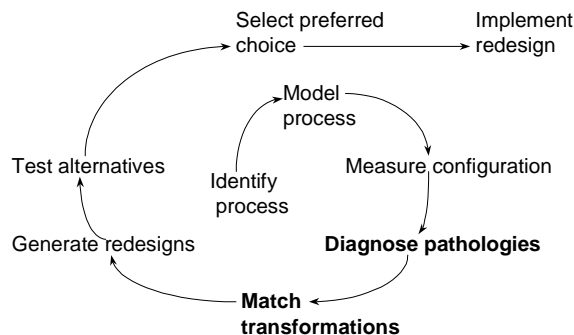


Figure 1. General Redesign Process

Previous research has described the concept *knowledge-based organizational process redesign* [27] and articulated how the proof-of-concept KOPeR agent [32] provides automated redesign support through measurement-driven inference. KOPeR automation and support focus specifically on the redesign phase of the BPR Life Cycle (cf. [13, 20]), as illustrated through the general redesign process in Figure 1. The sequence of process-redesign activities delineated in the figure represents a blend of expert reengineering methodologies—particularly those of [1, 6, 15, 17, 19]—synthesized together to compose an analytical method supporting measurement. The path through these steps is delineated as a spiral in the figure,

which represents a common notation for evolutionary processes.

KOPeR supports process redesign through automation of two key, intellectual steps highlighted in bold in the figure: diagnose pathologies and match transformations. Pathology diagnosis is focused on identifying problems and shortcomings associated with a process, and transformation matching concentrates on selecting the most appropriate technologies or other enablers of dramatic performance improvement. Automation of such intellectual activities represents a sharp contrast with the kinds of process modeling and simulation support generally available through first-generation tools. In the remainder of this section, we provide a high-level overview of the KOPeR architecture and then briefly describe the mechanics of this knowledge-based, process-workflow redesign system.

2.1. KOPeR architecture

The original version of KOPeR was implemented to contain and reason with considerable domain-specific knowledge, in addition to more general, domain-independent redesign rules. As a UNIX system predicated on the KnowledgeCraft tool, however, this original version became difficult to maintain and impossible to port to a PC environment. KOPeR has since been partially re-implemented and extended to provide automated redesign support through a Web interface [29].

To provide some contrast between the systems, we refer to this re-implementation as "KOPeR-lite" to depict the domain-independent nature of its knowledgebase; that is, KOPeR-lite lacks the considerable, domain-specific knowledge formalized through the original KOPeR, and its redesign capabilities are notably less strong. Nonetheless, even KOPeR-lite serves to demonstrate the concept of intelligent workflow-redesign problem solving. And this re-implementation delivers its redesign functionality to the Web through a commercial bridge, which provides direct access to a local client and includes an interface for Internet access by remote clients. With this, KOPeR-lite has the capability to access process data from anywhere on the Internet and remotely deliver intelligent reengineering support through ordinary HTML-based forms.

The high-level KOPeR-lite architecture is presented in Figure 2. Two principal KOPeR-lite functions depicted in the figure include pathology diagnosis and transformation matching. These functions correspond directly to the twin taxonomies labeled in the figure as part of the knowledgebase (KB)—one taxonomy for process pathologies, another for redesign transformations. The KB also incorporates a repository of process models used to represent various processes (e.g., from the domains of software engineering, finance, accounting and procurement). These are predicated on the process ontology first developed in [22] and later extended in [25]

through work on the Articulator, a knowledge system for general process diagnosis and repair (see [23, 24]). Process models supply instance-level information or “facts” pertaining to a process, which are used to support automated reasoning.

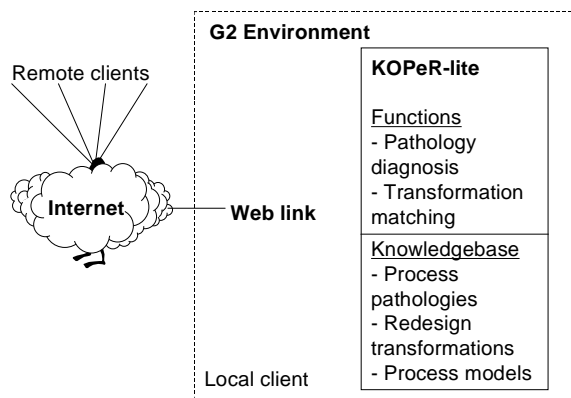


Figure 2. KOPeR-lite Architecture

The functions and knowledgebase comprise the core of KOPeR-lite and operate within the large box labeled "G2 environment" in the figure. KOPeR-lite is implemented using the G2 development tool [9], which also provides an inference engine, user interface functions, file and KB housekeeping and other services expected from a shell tool. The fourth KOPeR-lite element is labeled "Weblink." This bridge facility links KOPeR-lite reasoning to remote clients through the Internet, functioning effectively as an application programming interface (API) for TCP/IP networking and servicing Web browser HTTP requests. It enables KOPeR-lite to dynamically generate Web pages providing intelligent redesign recommendations that are tailored to each specific process submitted for KOPeR-lite analysis.

2.2. KOPeR-lite mechanics

Table 1 summarizes a sample of process measures used to drive the KOPeR-lite inference, along with graph-based definitions to operationalize the metrics. These KOPeR-lite measurement constructs conform to the graphical nature of most tools used for process modeling today (see [5]) and are naturally suited for process redesigns associated with workflow automation. Notice all measures defined in the table are operationalized in terms of simple graphical elements (e.g., nodes, edges, attributes, paths). KOPeR-lite's inferential power derives from heuristic linkages between such graph-based measures and corresponding process pathologies that can be diagnosed (see [33]).

Table 1. Example Process Measures

Measure	Graph-Based Definition
Process Length	Nodes in longest path
Process Breadth	Distinct paths
Process Depth	Process levels
Process Size	Nodes in process model
Process Feedback	Cycles in graph
Parallelism	Process Size divided by Length
IT Support	IT-support attributes
IT Communication	IT-communication attributes
IT Automation	IT-automation attributes
Organizational Roles	Unique agent role attributes
Process Handoffs	Inter-role edges
Organizations	Unique organization attributes
Value Chains	Unique Value Chain attributes

Pathology diagnosis is predicated upon a taxonomy of process pathologies used for classification of problems and shortcomings. This taxonomy formalizes some of the deep reengineering knowledge required for process redesign. The idea is to use process configuration measurements to detect and classify a variety of common process pathologies. The taxonomy is constructed from the BPR literature, as classes and instances of pathologies are synthesized from the various process problems and shortcomings noted in the expert reengineering methodologies from above (e.g., [1, 6]). Problematic conditions described in the many published redesign cases (e.g., [11, 21, 37, 38]) are similarly used to organize and populate the taxonomy. The class-level taxonomy of process pathologies is presented in Table 2, along with a sample instance from each of the ten classes.

KOPeR-lite employs a set of IF-THEN rules to classify pathologies on the basis of process measurements. For instance, the *parallelism* measure characterizes the extent to which process activities are performed in a sequential (low parallelism) versus concurrent (high parallelism) manner. If the measured value of parallelism for a process falls below some empirically determined threshold, then KOPeR-lite fires the rule corresponding to the "sequential process flows" pathology. This is listed as an instance of the pathology class "problematic process structure" in Table 2. A large number and variety of process pathologies can be diagnosed in this manner.

Table 2. Taxonomy of Process Pathologies

Pathology Class	Sample Instance
Problematic process structure	Sequential process flows
Bureaucratic organization	Job specialization
Fragmented process flows	Process friction
IT infrastructure	Manual process
“Checking” approach to	Review-intensive process

quality	
Centralized authority	Long decision chains
Under-utilized human potential	Training emphasis
Inhibitive leadership	Directive supervision
Centralized information	Central database architecture
Deficient core competency	Low IT expertise

Transformation matching is also automated through KOPeR-lite. A second taxonomy of redesign transformations is used for matching with pathologies. This taxonomy formalizes additional, deep reengineering knowledge required for process redesign. The idea is to use measurement-driven, diagnostic information from the steps above to match appropriate transformations. This taxonomy is also constructed by drawing from the BPR literature, as classes and instances of redesign transformations are synthesized from the various enabling technologies, organizational changes, workflow modifications and like interventions noted in the expert reengineering methodologies. Redesign transformations described in the many published BPR cases are similarly used to organize and populate the taxonomy. The class-level taxonomy of redesign transformations is presented in Table 3, along with a sample instance from each of the seven classes.

Table 3. Taxonomy of Redesign Transformations

Transformation Class	Sample Instance
Workflow reconfiguration	Process de-linearization
Information technology	Shared database system
Organizational design	Case manager
Human resource	Team-based compensation
Information availability	Informate agents
Inter-organizational alliance	Supplier-managed inventory
Management & culture	Employee stock ownership

KOPeR-lite also employs a set of IF-THEN rules to match pathologies with the appropriate redesign transformations. For instance, the "sequential process flows" pathology from above will cause KOPeR-lite to fire the rule corresponding to "process de-linearization," an instance of the transformation class "workflow reconfiguration" in Table 3. A large number and variety of redesign transformations can be matched in this manner, individually as well as in combination; that is, KOPeR-lite "recommends" redesign transformations independently in a first pass through its reasoning and then suggests effective combinations and partial orderings of such transformations. For example, KOPeR-lite's knowledgebase includes concepts such as the prerequisite

need for IT infrastructure (e.g., network technology) to support workflow automation. When redesigning a paper-based process, for instance, KOPeR-lite may recommend network technologies and electronic communications as necessary prerequisites to a workflow application for automatic routing and support of process activities.

To reiterate, KOPeR-lite automates two key, intellectual redesign activities: pathology diagnosis and transformation matching. The graph-based process measures enable KOPeR-lite to naturally accommodate most process redesigns associated with workflow automation. And KOPeR has been demonstrated—in both the laboratory and field—to provide a feasible and efficient approach to redesigning process workflows. As noted above, however, the question of redesign effectiveness remains to be answered. The research design that follows addresses this question directly.

3. Research design

The research design centers around laboratory experimentation. Experimentation offers a number of research benefits, including control over variables, mitigation of confounding, enhanced replicability and others [18]. And this technique can overcome many acknowledged threats to validity (see [3]). Experimentation is employed extensively in the physical sciences for these reasons, but perceived problems with external validity and generalizability of laboratory experiments complicate use of this research method in the business domain [26].

In this investigation, we help promote external validity and generalizability by having representative subjects perform widely practiced process-workflow redesign tasks using common BPR methods and tools. Aside from some artificiality imposed through the laboratory environment, this closely matches the manner in which redesign is accomplished in the enterprise (see [6, 15]). Thus, the experimental task associated with this study should externalize well to practice. The subjects are also representative of the target professional group addressed through this research. And the experimental task is representative of redesign activities commonly performed in the field. Thus, results should also generalize well beyond the laboratory. These aspects of the research design are described in greater detail below.

The specific task assignment studied through this experiment involves generation of redesign alternatives for a process workflow. As the name implies, this task requires subjects to identify and describe one or more possible alternatives to be considered for redesigning a process workflow. Subjects are assigned background reading about reengineering (e.g., [28]) and required to read the "Credit Financing Case" [15, pp. 36-39], [30]) several days before the experiment. The case description includes textual discussion and a "rich pictures" representation (see [4]),

along with supplemental and complementary process information (e.g., the kind of graph-based information used by KOPeR-lite). Subjects also receive a one-hour reengineering lecture prior to being tested in the laboratory. Together, these activities constitute our training example. This approach is similar to the pre-experimentation techniques employed to assess the performance of a maritime expert system [12].

Further, subjects are encouraged to think about how to redesign processes for several days preceding the experiment, as they are directly exposed to the credit financing case and have ample opportunity to ask questions pertaining to reengineering in general—or credit financing in particular—before the experiment begins. Moreover, the experiment is conducted within the context of a graduate information systems course, which ensures the background readings and other preparations are completed. Additionally, subjects are motivated to perform well on the experimental task through competition for peer recognition and prizes.

The experiment directly compares the performance of KOPeR-lite with that of a group of reengineering novices tasked to generate redesign alternatives from a common process workflow description. As noted above, the reengineering novice group is composed of graduate students in the information systems curriculum. In this particular group of subjects, every student is professionally employed, has between five and fifteen years' experience and holds a full-time, middle-management position involving knowledge and information work. This subject group is quite typical of the work experience and responsibility level associated with process redesign in practice.

Subjects are screened before the experiment on the basis of prior BPR experience, and to ensure "novice" status, the experiment excludes anyone who has worked previously in a reengineering capacity. All eligible students in the class are required to participate in the experiment, so we avoid many problems associated with self-selection bias. A total of thirty-two people participated in the experiment.

Under laboratory conditions, in which availability of information, access to tools and assistance, problem-solving time and other key variables are strictly controlled, the redesign performance of each subject is evaluated. Subjects are provided a description of the process to be redesigned during the experiment—this is the same "Credit Financing Case" reviewed as homework—and instructed to write a short description for every distinct redesign alternative they can generate. Aside from the homework assignments, the entire experiment takes place within one two-hour class session. Each subject is allotted roughly half the session for the redesign task itself. The limited time allotted for the experiment therefore implicitly tests the speed at which redesign alternatives can be generated, in addition to the

effectiveness criteria outlined below. Speed is often touted as an advantage of hiring external consultants.

KOPeR-lite performs the same process-redesign task as the subjects, subject to the same time constraints and information in the form of process diagrams and measurements. KOPeR-lite requires about one minute to initialize, accept the process information and generate its redesign alternatives. Evaluation of redesign performance is conducted by the course instructor (i.e., the researcher in this study). This understandably introduces some bias, as only one instructor's views are reflected in the redesign performance data. A follow-on study using a second evaluator can be conducted to improve the experiment's reliability and allow for statistically testing inter-rater agreement.

Two operationalized criteria are consistently applied for assessment of both human subjects and KOPeR-lite performance: 1) quantity of redesign alternatives generated during the experiment, and 2) potential of each redesign alternative to improve process performance. The basis for operationalization of these criteria is straightforward. Specifically, the *quantity* criterion is operationalized as the number of distinct redesign alternatives generated by a subject in the time period allowed by the experiment. Each redesign alternative must be unique and clearly distinguishable from the process baseline to be classified as "acceptable." Also, each corresponding process description must clearly describe how the redesigned activities would be performed. And each redesign alternative must be feasible, in that the enabling redesign transformations must be based on demonstrated technologies and approaches that are available or in use today. Thus, redesign transformations must be realistic and attainable by an enterprise today (although perhaps involving novel technologies and/or high implementation risk). This feasibility requirement is characterized as the "no Star Wars" restriction.

The *potential* criterion is operationalized as the potential of each, distinct redesign alternative to improve process performance. Performance-improvement potential is rated and scored according to four categories: 1) unacceptable redesign (score = 0)—does not meet acceptability criteria from above; 2) negligible improvement (score = 1)—reflecting expected single-digit or negative performance improvement over the process baseline; 3) moderate improvement (score = 2)—reflecting expected cost and/or cycle time gains on the order of tens of percentage points; and 4) dramatic improvement (score = 3)—reflecting expected order-of-magnitude gains in terms of cost and/or cycle time. Ratings are based on three factors: 1) the redesign transformations coincide with one or more of the same redesigns cited in the credit financing case; 2) the researcher's BPR experience; and 3) current reengineering guidance suggesting redesign transformations applied in combination lead to greater performance gains.

3.1. Experimental hypotheses

We develop two primary hypotheses to be tested through the experiment. First, KOPeR-lite is expected to exceed the redesign performance of the average reengineering novice. The knowledge system is designed and developed to integrate deep reengineering knowledge and redesign expertise captured from experts in the field. Given that process redesign represents a knowledge-intensive activity, we would not expect a novice who lacks such knowledge and expertise to perform at the same level. The first hypothesis follows. Because we are using two dimensions—quantity and potential—to define redesign performance, this hypothesis has two parts.

Hypothesis 1a: KOPeR-lite can generate a greater quantity of redesign alternatives than the average novice.

Hypothesis 1b: KOPeR-lite can generate redesign alternatives with greater potential for performance improvement than the average novice.

Second, a question then arises whether people pooling redesign alternatives can outperform KOPeR-lite, even if the average person individually cannot match the performance of this knowledge system. This leads to a related, second hypothesis, which similarly has two constituent parts.

Hypothesis 2a: KOPeR-lite can generate a greater quantity of redesign alternatives than the combined novice group.

Hypothesis 2b: KOPeR-lite can generate redesign alternatives with greater potential for performance improvement than the combined novice group.

3.2. Experimental data

Subjects provide hand-written descriptions of redesign alternatives generated during the experiment. They are asked to write short paragraphs describing each redesign alternative (no more than one page per redesign). Space precludes inclusion of the subjects' responses in full text, but here, we summarize selected responses to provide some characterization of the raw data captured and analyzed through this experiment.

For instance, Subject 3 proposes to "automate all [the process] steps." This redesign alternative (Alternative 3-2) is assessed as "infeasible," however, because it fails to specify how all such steps could be performed without human intervention. Notwithstanding the power and potential of expert systems, workflow applications, intelligent agents and other modern automation capabilities, we are not convinced the entire process can be automated. Without more information about how this redesign could be

implemented, this input is classified as infeasible (i.e., "unacceptable") and is assigned a score of zero.

Another subject (number 4) describes a portion of one redesign alternative (Alternative 4-1) as follows.

The credit approval process can be integrated with the application for financing. Upon approval, desired terms can also be taken from the application and matched with pricing information to give a quotation. This is essentially 2 steps. Approval and development of terms.

This redesign alternative proposes to combine steps that are performed separately in the process baseline. This constitutes a redesign alternative that is distinct from the baseline, and the proposed integration appears to be feasible, so this constitutes an acceptable redesign alternative. Despite the proposed integration of steps and reduction of handoffs, however, absent automation or some other change enabler, the process steps appear to be performed in essentially the same manner as specified by the process baseline. Because little performance improvement would be expected from this change, we assess the performance potential of the corresponding redesign alternative as "negligible" and assign a score of one.

Subject 6 also develops a redesign alternative (Alternative 6-1) assessed to offer "negligible" potential for performance improvement (score = 1). This redesign calls for transmitting all process information electronically, as opposed to using paper and telephonic media. Although we view this as a positive step that can significantly decrease the transport time for process work, simply installing a network and e-mail application, for instance, to speed communications is unlikely to offer more than a negligible increase in process performance. As noted above in the context of workflow automation, such redesign transformations are colorfully referred to as "automating the mess" in the reengineering literature (cf. [14]).

For comparison, KOPeR-lite generates a similar redesign alternative (Alternative KOPeR-lite-3) calling for network-enabled e-mail and shared databases. For the same reason, the potential of this KOPeR-lite-generated redesign is also assessed as "negligible" (score = 1). As a note, one may disagree with any particular assessment such as this—and another researcher may classify and score this and other redesigns higher or lower in terms of performance potential—but such absolute measurements are not as important in this study as the *relative* performance of KOPeR-lite with respect to human subjects. The key is for human- and computer-generated redesigns to be assessed consistently. We ensure such consistency through assessment based on operationalized criteria.

Subject 6 develops a second redesign alternative (Alternative 6-2), which is assessed to offer "moderate" potential for performance improvement (score = 2). This

latter redesign is based on the "case manager" concept, which is widely noted in the reengineering literature as appropriate for a process such as this (e.g., linear sequence of relatively straightforward, administrative steps; see [15]) and is actually recommended in the credit financing case. We note KOPeR-lite also recommends use of a case manager (Alternative KOPeR-lite-2) and is scored accordingly (i.e., offering moderate potential, score = 2). Again, we ensure consistency of assessments to support comparative analysis of KOPeR-lite and the novice group.

As an example of a redesign offering potential for "dramatic" performance improvement, Subject 5 develops an alternative (Alternative 5-2) combining a case manager arrangement with network-enabled electronic communications, a Web-based application to remotely support field sales agents and workflow automation for the process. Current reengineering experts emphasize the potential for dramatic performance improvement through applying such "change levers" or enablers in combination (e.g., see [7, 36]). This reinforces our basis for classifying the redesign as having dramatic potential for process improvement (score = 3). Likewise, KOPeR-lite generates a redesign (Alternative KOPeR-lite-6b) based on a case manager, workflow application and intelligent supply chain agents (see [2]). Consistently, although the details differ, this KOPeR-lite-generated redesign is also assessed to offer dramatic potential for process improvement (score = 3). Assessments of the other novice- and KOPeR-lite-generated redesigns are made in a similar fashion.

4. Knowledge system assessment

Tables 4 and 5 present summary statistics for the novice group, with relatively low scores (e.g., "unacceptable" and "negligible" redesigns) in the former and comparatively high scores (e.g., "moderate" and "dramatic" redesigns) in the latter table. The first column of Table 4 (labeled "unacceptable") summarizes novice-generated redesign alternatives deemed to be infeasible or indistinct (score = 0). The mean value (0.47) reflects a total of fifteen unacceptable redesigns (sum) generated by the thirty-two subjects; in other words, on average, nearly every other subject in this experiment generates one infeasible or indistinct redesign alternative.

Table 4. Statistics - Novice Group (Low)

Unacceptable		Negligible	
Score	0	Score	1
Mean	0.47	Mean	1.00
Standard Dev.	0.67	Standard Dev.	1.37
Minimum	0	Minimum	0
Maximum	3	Maximum	7
Sum	15	Sum	32

Observe that while many subjects have no such unacceptable redesigns (minimum = 0), at least one generates three infeasible or indistinct alternatives (maximum = 3). The same summary statistics for the "negligible" column (score = 1) indicate the average subject generates one redesign assessed to have negligible impact in terms of potential to improve process performance. One subject generates seven redesigns at this level of potential for performance improvement (maximum = 7). Statistics corresponding to the eighteen moderate-potential redesign alternatives (score = 2) are shown in the first two columns of Table 5, with their dramatic-potential counterparts in the right-hand columns. Notice seven redesign alternatives are generated with potential to effect dramatic performance improvement (score = 3), and one subject even generated two dramatic redesigns (maximum = 2).

Table 5. Statistics - Novice Group (High)

Moderate		Dramatic	
Score	2	Score	3
Mean	0.56	Mean	0.22
Standard Dev.	0.72	Standard Dev.	0.49
Minimum	0	Minimum	0
Maximum	3	Maximum	2
Sum	18	Sum	7

Table 6 summarizes comparative statistics for the novice group as a whole and KOPeR-lite. Notice KOPeR-lite does not generate even one unacceptable redesign alternative, whereas the novice group generates fifteen. Comparing KOPeR-lite output (6 redesign alternatives) with the novice group *average* quantity (2.25), we see KOPeR-lite is more prolific in its generation of redesign alternatives than the average novice. This result offers evidence to support Hypothesis 1a; that is, KOPeR-lite outperforms the average novice in terms of quantity of redesign alternatives generated. A 99% confidence interval (1.47-3.03) indicates redesign performance of the average novice is statistically quite different than KOPeR-lite.

Table 6. Comparative Statistics

Category	Novice	KOPeR-lite
Unacceptable (0)	15	0
Negligible (1)	32	1
Moderate (2)	18	3
Dramatic (3)	7	2
Total Quantity	72	6
Ave. Quantity	2.25	6
Ave. Score	1.23	2.17

Alternatively, comparing KOPeR-lite's output (6 redesign alternatives) with the novice group's *total* quantity

(72 redesigns), as a pooled group, the thirty-two subjects clearly generate a greater number of redesigns across all categories. On the surface, this result offers evidence against Hypothesis 2a; that is, the novice group collectively outperforms KOPeR-lite in terms of quantity of redesign alternatives generated. However, we should note the 72 redesigns contain many redundancies, which effectively bias the results in favor of the novice group. By adjusting the data to remove redundant redesigns, this number drops to 9, relatively close to the number generated by KOPeR-lite (6). And if we neglect the (2) non-redundant redesigns classified as “unacceptable” from above, the number generated by KOPeR (6) and the novices (7) converges quite closely.

Further, notice the distribution of novice-generated redesigns is skewed toward the negligible end of the (potential-performance) range, whereas KOPeR-lite-generated redesigns are distributed more toward those assessed with moderate and dramatic potential. The average score shown at the bottom of Table 6 reflects the weighted average of all redesign alternatives across the four categories. The KOPeR-lite average score (2.17) is nearly twice that of the novice group (1.23). This suggests KOPeR-lite-generated redesigns offer greater potential for performance improvement than alternatives generated by the average novice. This result offers evidence to support Hypothesis 1b; that is, KOPeR-lite outperforms the average novice in terms of performance-improvement potential of redesign alternatives generated. A 99% confidence interval (0.86-1.60) indicates redesign performance of the average novice is statistically quite different than KOPeR-lite.

Alternatively, the novice group generates eighteen moderate and seven dramatic redesigns, compared with only three and two for KOPeR-lite, respectively. On the surface, this result offers evidence against Hypothesis 2b; that is, the novice group collectively outperforms KOPeR-lite in terms of potential of redesign alternatives generated. However, removing redundancies as above, KOPeR-lites’ five moderate/dramatic redesigns actually exceed the total number generated by the novice group (3).




Figure 3. Comparative Redesign Performance

Figure 3 further illustrates comparative redesign performance for the average novice and KOPeR-lite. Two histograms delineate the average frequency of redesign alternatives falling into the four categories for both the

novice group and KOPeR-lite. Notice KOPeR-lite generates more redesign alternatives overall (6) than the average novice (2.25). This graphically depicts the numerical result discussed above and helps illustrate the evidence with respect to Hypothesis 1a. The two histograms also delineate the relative distribution of redesign alternatives across the four categories. As summarized in the table above, observe that KOPeR-lite generates no unacceptable redesign alternatives (i.e., with score of 0), and the majority of KOPeR-lite-generated redesigns fall in the higher-potential categories (i.e., “moderate” and “dramatic” with respective scores of 2 and 3). Visual inspection of these histograms confirms the statistical difference noted above and helps illustrate the evidence with respect to Hypothesis 1b.

5. Conclusions and future research

This investigation focuses on testing the effectiveness of a knowledge-based, process-workflow redesign agent and employs KOPeR-lite as a platform for experimentation to assess the relative efficacy of redesigns generated by computer versus those developed by reengineering novices. In addition to implicitly testing redesign speed, reengineering performance is explicitly assessed along two dimensions: 1) quantity of (feasible and distinct) redesign alternatives generated, and 2) potential for process improvement through the generated redesign alternatives. We hypothesize KOPeR-lite can outperform the average novice as well as the collective novice group. Experimental results provide statistically-significant support for the first hypothesis (i.e., KOPeR-lite outperforms the average novice). And once redundant novice-generated redesigns are removed from the dataset, KOPeR-lite compares favorably with the novice group in terms of the second hypothesis as well (e.g., KOPeR-lite roughly matches the collective performance of the novice group). Thus, at the novice level, KOPeR-lite’s redesign performance is much better than the average novice and comparable to the entire novice group ($n = 32$). In the sense of our modified Turing Test, this suggests KOPeR-lite performs redesign activities at an overall level of effectiveness exceeding that of the reengineering novice.

These results accentuate the question, at what level—now arguably beyond that of the reengineering novice—does KOPeR-lite perform its process redesign? Future research to address this question may involve a similar experiment with more experienced subjects (e.g., reengineering experts from professional BPR consulting firms). But given our comments above pertaining to KOPeR-lite’s limited (e.g., domain-independent) KB, we would hypothesize its redesign performance to fall somewhere in between that of the novice and the expert.

Regarding practical application, the enterprise manager today may not want to wait for the follow-on studies noted above. The leading adopter of this

reengineering technology may have some opportunity to employ the kinds of knowledge systems discussed in this paper for competitive advantage. Indeed, the manager may want to experiment on his or her own—for example, using KOPeR-lite to support its own employees in the operational enterprise.

Also noting opportunity for competitive advantage through knowledge system technology, the professional BPR consultant may similarly want to examine this technology—for example to directly support its reengineering consultants. The use of such redesign-automation technology would certainly represent a differentiating factor between BPR consulting firms.

Finally, the efficacy of such knowledge systems may not be limited to the reengineering domain. Although the KOPeR-lite application examined in this study contains knowledge focused on the reengineering domain, conceivably, a knowledge system such as this can be developed to focus on a variety of knowledge and information work. For instance, a knowledge system such as KOPeR-lite may be adaptable to address problems in domains such as knowledge management and like areas of current, intense interest and attention. Indeed, these second-generation, knowledge-based tools may effectively perform at or beyond the expert level in a number of important domains. This represents an exciting possibility, one that will call for many future studies along the lines of this investigation.

6. References

- [1] Andrews, D.C. and Stalick, S.K. *Business Reengineering: the Survival Guide* New York, NY: Yourdon Press Computing Series (1994).
- [2] Barbuceanu, M. and Fox, M.S., "The Design of COOL: A Language for Representing Cooperation Knowledge in Multi-Agent Systems," WWW: <http://www.ie.utoronto.ca/EIL/ABS-page/ABS-intro.html> (1993).
- [3] Campbell, D.T. and Stanley, J.C. *Experimental and Quasi-Experimental Designs for Research* Chicago, IL: Rand McNally (1973).
- [4] Checkland, P. *Systems Thinking, Systems Practice* New York, NY: Wiley (1981).
- [5] Curtis, B., Kellner, M.I. and Over, J. "Process Modeling," *Communications of the ACM* 35:9 (September 1992), pp. 75-90.
- [6] Davenport, T.H. *Process Innovation: Reengineering Work through Information Technology* Boston, MA: Harvard Press (1993).
- [7] Davenport, T.H. "Business Process Reengineering: Where It's Been, Where It's Going," in V. Grover and W. Kettinger (eds.), *Business Process Change: Reengineering Concepts, Methods and Technologies* Middletown, PA: Idea Publishing (1995), pp. 1-13.
- [8] Davenport, T.H. and Short, J.E. "The New Industrial Engineering: Information Technology and Business Process Redesign," *Sloan Management Review* 31:4 (Summer 1990), pp. 11-27.
- [9] Gensym. G2 product documentation; World Wide Web address: <http://www.gensym.com> (1999).
- [10] Georgakopoulos, D., Hornick, M. and Sheth, A. "An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure," *Distributed and Parallel Databases* 3:2 (1995), pp. 119-153.
- [11] Goldstein, D. K. "Hallmark cards," Harvard Business School Case No. 9-186-044 (July 1986).
- [12] Grabowski, M. and W.A. Wallace, "An Expert System for Maritime Pilots: Its Design and Assessment Using Gaming," *Management Science* 39:12 (1993), pp. 1506-1520.
- [13] Guha, S., Kettinger, W.J. and Teng, J.T.C. "Business Process Reengineering: Building a Comprehensive Methodology," *Information Systems Management* 10:3 (Summer 1993), pp. 13-22.
- [14] Hammer, M. "Reengineering Work: Don't Automate, Obliterate" *Harvard Business Review* 68:4 (July-August 1990), pp. 104-112.
- [15] Hammer, M. and Champy J. *Reengineering the Corporation: A Manifesto for Business Revolution* New York, NY: Harper Business Press (1993).
- [16] Hamscher, W. "AI in Business Process Reengineering: A Report on the 1994 Workshop," *AI Magazine* 15:4 (Winter 1994), pp. 71-72.
- [17] Harrington, H.J. *Business Process Improvement: the Breakthrough Strategy for Total Quality, Productivity, and Competitiveness* New York, NY: McGraw-Hill (1991).
- [18] Jenkins, A.M. "Research Methodologies and MIS Research," in: E. Mumford, R. Hirscheim, G. Fitzgerald and A. Wood-Harper (Eds.), *Research Methods in Information Systems* Amsterdam: North-Holland (1984), pp. 103-117.
- [19] Johansson, H.J., McHugh, P., Pendlebury, A.J. and Wheeler, W.A. III, *Business Process Reengineering: Breakpoint Strategies for Market Dominance* Chichester, UK: Wiley (1993).
- [20] Kettinger, W.J., Guha, S. and Teng, J.T.C. "The Process Reengineering Life Cycle Methodology: A Case Study," in V. Grover and W. Kettinger (eds.), *Business Process Change: Reengineering Concepts, Methods and Technologies* Middletown, PA: Idea Publishing (1995), pp. 211-244.

- [21] King, J. L., and Konsynski, B. "Singapore Tradenet: A Tale of One City," Harvard Business School Case No. 9-191-009 (1990).
- [22] Kling, R. and Scacchi, W. "The Web of Computing: Computer Technology as Social Organization," in: M. Yovits (Ed.), *Advances in Computers* 21, New York, NY: Academic Press (1982), pp. 3-85.
- [23] Mi, P. *Modeling and Analyzing the Software Process and Breakdowns* doctoral dissertation, University of Southern California (1992).
- [24] Mi, P. and Scacchi, W. "Modeling, Integrating, and Enacting Complex Organizational Processes," *Proceedings of the Fifth International Symposium on Intelligent Systems in Finance, Accounting and Management* Volume 1, Stanford University (December 1993).
- [25] Mi, P. and Scacchi, W. "A Meta-Model for Formulating Knowledge-Based Models of Software Development," *Decision Support Systems* 17:4 (1995), pp. 313-330.
- [26] Mumford, E., Hirscheim, R., Fitzgerald G., and A. Wood-Harper *Research Methods in Information Systems* Amsterdam: North-Holland (1984).
- [27] Nissen, M.E. *Knowledge-Based Organizational Process Redesign: Using Process Flow Measures to Transform Procurement* doctoral dissertation, University of Southern California (1996).
- [28] Nissen, M.E. "A Focused Review of the Reengineering Literature: Expert Frequently Asked Questions (FAQs)," *Quality Management Journal*, 3:3 (1996).
- [29] Nissen, M.E. "Toward Intelligent Web-based Redesign Support," *Proceedings of the American Association for Artificial Intelligence (AAAI) 1997 Conference, Workshop on Using AI in Electronic Commerce, Virtual Organizations and Enterprise Knowledge Management to Reengineer the Corporation*, Providence, RI (1997).
- [30] Nissen, M.E. "Dr. Mark's Credit Financing Case," Naval Postgraduate School Instructional Case, World Wide Web address: <http://web.nps.navy.mil/~menissen/haas/ba248d/kopercredit.html> (1998).
- [31] Nissen, M.E. Redesigning Reengineering through Measurement-Driven Inference," *MIS Quarterly* 22:4 (1998), pp. 509-534.
- [32] Nissen, M.E. "An Intelligent Agent for Web-based Process Redesign," in: R. Sprague (Ed.), *Proceedings Hawaii International Conference on System Sciences*, Maui, HI (January 1999).
- [33] Nissen, M.E. "A Configuration-contingent Enterprise Redesign Model," in: T.R. Gullledge, D.J. Elzinga and C.Y. Lee (Eds.), *Process Engineering: Advancing the State of the Art* Boston: Kluwer (1999).
- [34] Rich, E. and Knight, K. *Artificial Intelligence (Second Edition)* New York, NY: McGraw-Hill (1991).
- [35] Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach* Englewood Cliffs, NJ: Prentice-Hall (1995).
- [36] Stoddard, D.B. and Jarvenpaa, S.L. "Business Process Redesign: Tactics for Managing Radical Change," *Journal of Information Management Systems* 12:1 (1995), pp. 81-107.
- [37] Stoddard, D.B., and Meadows, C.J. "Capital Holding Corporation - reengineering the Direct Response Group," Harvard Business School Case No. 9-192-001 (1992).
- [38] Talebzaheh, H.; Mandutianu, S; and Winner, C.F. "Countrywide Loan-Underwriting Expert System," *AI Magazine* 16:1 (1995), pp. 51-64.