Via Teldok

ISSN 0283-5266 • Nr 12 • Oktober 1988

# Den automatiserade experten

## En uppsats om expertsystem och några intryck från en AI-konferens

Olov Östberg, Randall Whitaker & Benjamin Amick, III

Via TELDOK

ISSN 0283-5266 • Nr 12 • Oktober 1988

# Den automatiserade experten

## En uppsats om expertsystem och några intryck från en AI-konferens

Olov Östberg, Randall Whitaker & Benjamin Amick, III

Via TELDOK och TELDOK Referensdokument måste normalt styckbeställas (gratis, dygnet runt, från DirektSvar – f d TeleSvar, 08-23 00 00), medan TELDOK Rapport och TELDOK-Info skickas automatiskt till den som så vill.

Det här numret av Via TELDOK skickas till alla som angett att de önskar få TELDOK Rapport m m fortsättningsvis. Observera att innehållet i Via TELDOK nr 12 kan förefalla mer svårtillgängligt än i en "vanlig" TELDOK Rapport – därav valet av publikationsserie!

En förteckning över några senare utgivna TELDOK-publikationer återfinns på omslagets baksida. Använd kupongen längst bak i skriften om Du vill beställa rapporter från TELDOK eller lämna andra meddelanden!

ISSN 0283-5266



Publikationerna kan beställas gratis,
dygnet runt, från DirektSvar (f d TeleSvar), 08-23 00 00

Tryckeri: Nyströms, Bollnäs 1988

# Introduktion & Sammanfattning

## Om denna Via TELDOK

Föreliggande Via TELDOK består av två uppsatser som författats av samma grupp experter och gäller samma ämnesområde – **expertsystem**.

Automatiserade expertsystem ingår i vad som kallas kunskapsstödssystem, eller kanske mer korrekt på engelska: **knowledge based systems**. Expertsystem räknas till det alltmer aktuella och delvis kontroversiella teknikområdet artificiell intelligens, **AI**.

Olov Östberg – numera tillbaka vid Televerket (Pla, telefon 08-713 3897) – samt Randy Whitaker och Ben Amick skrev först manus till den omfattande och engelskspråkiga rapporten THE AUTOMATED EXPERT (som här återges på sidorna 27–122). När författarna sedan tillsammans skulle besöka den internationella konferensen "Culture, Language and Artificial Intelligence" i Stockholm, erbjöd TELDOK dem gärna att skriva ett ireferat från denna. Referatet, som presenterar och diskuterar några viktiga utvecklingsdrag ifråga om expertsystem, återfinns på sidorna 1–23.

Eftersom det huvudsakliga bidraget, THE AUTOMATED EXPERT, är skrivet på engelska och innehåller en del snåriga passager, har vi valt att närmast återge några kärnpunkter ur uppsatsen på svenska. Med detta svenskspråkiga kort-referat som bakgrund kan det visa sig lättare att *dels* förstå bakgrunden till några av de diskussioner som återges i referatet KULTUR, SPRÅK OCH ARTIFICIELL INTELLIGENS, *dels* skumma, läsa och förstå THE AUTOMATED EXPERT.

## The Automated Expert

Rapporten THE AUTOMATED EXPERT är indelad i två delar. Den första delen är en översikt av utrustning (hårdvara), program (mjukvara), principer, utbildningsinsatser m m som individer behöver och använder i anknytning till expertsystem. Den andra delen av rapporten granskar och diskuterar expertsystem med hänsyn till förhållandena inom den arbetsplats där de utnyttjas.

Författarna har haft svårt att hitta faktiskt fungerande expertsystem, som skulle kunna beskrivas i enlighet med TELDOKs syfte att bidra till "dokumentation vid tidigast möjliga tidpunkt av praktiska tillämpningar av teleanknutna informationssystem i arbetslivet". Rapporten ger exempel på några informationssystem – där bara ett fåtal är teleanknutna! – av det här slaget som flyttat ut från laboratorierna, men ännu är expertsystemen i praktiskt bruk inte så många ens i USA.

Dels är det svårt att definiera vad expertsystem egentligen är. Definitioner som baseras på den teknik som används är vanskliga eftersom de ofta visat sig vara motstridiga och tekniken kan förändras. Författarna väljer att se expertsystem ungefär som ett sätt att överföra kunskap ("the process of knowledge transfer"), oavsett hur maskinen som hjälper till med detta är konstruerad.

Dels består enligt författarna mycken av den information som man kan hitta om expertsystem i böcker och tidskrifter av överdrifter, ogrundad optimism och försäljningsprat. De expertsystem som beskrivs i press eller litteratur – ibland med stor entusiasm, som t ex att de kan öka de anställdas produktivitet *minst* tio gånger – är oftast bara laboratorieprodukter eller leksaksprototyper som aldrig använts i verkliga livet eller åtminstone inte nått någon framgång där. Till exempel: General Electric har utvecklat systemet DELTA/CATS, som ofta framställs som framgångsrikt – utan att någonsin ha tagits i verkligt bruk efter de första laboratorieförsöken.

Det betyder inte att man idag kan avfärda expertsystemen med en axelryckning. För det **första**: forskning om expertsystem är ett tillväxtområde. Expertsystem ägnas nu investeringar på ca 60 miljoner USD, och siffran förväntas stiga till 350 miljoner USD 1990. (Observera dock att en stor del av investeringarna – all finansiering av någon betydelse, utom i Japan – har militär anknytning och militära ändamål.)

För det **andra**: det finns redan expertsystem som används rutinmässigt och som betraktas som mycket framgångsrika. Honeywell har utvecklat systemet COOKER, som beräknas ha sparat miljoner dollar – trots att användarna i vissa fall struntar i hälften av COOKERs rekommendationer. Digital Equipment Corporation, DEC, har utvecklat XCON, vilket sägs spara företaget 18 miljoner USD – per år.

Idag finns inte längre några hårdvarumässiga begränsningar som kan hindra utvecklingen av expertsystem. Persondatorer har spritts på kontoren och har sådan kapacitet att de kan användas för expertsystem.

När det gäller programvara finns dels ett antal högnivåspråk som är deklarativa (dvs användaren talar om för programmet vad som skall göras, inte hur det ska göras), dels ett antal "expertsystem-skal" som anpassats speciellt för utveckling av kunskapsdatabaser och som hämtat inslag från olika programspråk. Författarna konstaterar att expertsystemen får alltmer modesta uppgifter, och att de är exempel på komplexa (högnivå-) produkter som levererar triviala (lågnivå-) "kunskaper" och rekommendationer.

Kunskapsingenjör är den som först skall locka kunskaper ur levande experter och sedan använder ett expertsystemskal eller programmeringsspråk för att utifrån dessa kunskaper definiera ett (ibland mycket) stort antal "regler" som gäller inom det kunskapsområde där expertsystemet skall användas. (Ett åttiotal regler har fungerat bra för ett system, medan i ett annat fall tolvhundra regler har visat sig vara för få!) Expertsystemet används sedan av ett antal användare i en serie konsultationer, där systemet ger rekommendationer eller gör "bedömningar" som skall bygga på områdesexperternas samlade kunskaper.

De bästa expertsystemen har utvecklats av kunskapsingenjörer som själva är eller har blivit experter inom det aktuella kunskapsområdet. I de flesta fall är tyvärr kunskapsingenjörerna experter enbart på att uttrycka redan definierade kunskaper i sökbara regler. Kunskapsingenjörerna är däremot sällan experter på att (1) leta fram kunskap, (2) anpassa expertsystemens gränssnitt till ergonomiska och användarmässiga aspekter eller (3) testa de expertsystem som de varit med om att utveckla.

Författarna diskuterar dessa särskilt problematiska arbetssteg och drar bl a slutsatsen att det ännu saknas bra testmetoder och pålitliga testresultat. — Dessutom: Arbetet med ett expertsystem kanske **aldrig kan avslutas** eftersom kunskapsmassan växer och t o m reglerna som binder samman kunskapsområdet kan komma att ändras!

I motsats till vad kunskapsingenjörer vanligen gör, ser författarna inte kunskap eller **expertis** som en given storhet och gripbar massa utan som ett slags yrkesskicklighet ("skill") som är beroende av det företag och den omgivning där experten är verksam. Expertis i arbetslivet har blivit mindre rutinbaserad, mindre knuten till konkreta och avgränsbara arbetssteg. Expertis i det "post-industriella samhället" har blivit abstraktare och inriktad på förståelse av allt komplexare arbetsuppgifter. Expertis (och experten/den anställde) måste också vara dynamisk och föränderlig, eftersom förutsättningarna ständigt förändras för företagens verksamhet och de anställdas arbetsuppgifter.

Expertsystem i produktionen kan uppfattas som ytterligare en metod att rationalisera arbetet genom att dela in det i ett antal avgränsbara steg som kan upprepas t ex av en robot. I så måtto är expertsystemen en del av den "kontrollrevolution" (författarnas term) som t ex tayloristisk Scientific Management står för – arbetsprocessen kontrolleras med väldefinierade metoder och uppdelning av arbetsuppgifterna.

Men under senare tid – i dagens postindustriella eller "informationssamhälle" – uppfattar författarna att istället en "kvalitetsrevolution" har inträffat där tonvikten läggs på produktkvalitet och därför helt andra aspekter måste betonas, såsom kreativitet och innovation, livslång vidareutbildning och ökad inriktning på grupparbete.

I det sammanhanget måste expertsystem uppfattas som redskap vilka bör användas för uppgifter där de är mest lämpliga men som långtifrån kan utföra alla typer av arbetsmoment. Som komplement till automatisering av arbetet nämner författarna att "informatisera" – dvs att ge människan datorstöd och lämpliga verktyg – och att "humanisera" arbetet – dvs att låta en människa sköta arbetsprocessen helt på egen hand. I de senare fallen bör expertsystemens rekommendationer vara – just (artiga) rekommendationer, inte tvingande order. – Författarna menar att arbetstillfredsställelsen är störst där den anställdes uppgifter definieras så lite och så vagt som möjligt.

När expertsystem introduceras i arbetsmiljön, utgör de inte bara en serie nya verktyg för beslutsfattare och andra, utan innebär enligt författarna helt nya sätt att genomföra arbetsuppgifterna. Användarna måste vara motiverade och involverade redan under utvecklingen av systemen – det finns så många prototyper till medicinska expertsystem, men så mycket färre sådana i praktiskt bruk, eftersom kunskapsingenjörerna felaktigt antagit att läkare "automatiskt" (utan omfattande samarbete och diskussioner) kommer att vilja använda vilken ny teknik som helst för att rädda liv och bota plågor.

Företag som utvecklar och inför expertsystem måste göra en avvägning mellan i vilken utsträckning systemen skall användas för att **kontrollera** arbetsresultaten respektive för att **öka kvaliteten** på arbetsresultaten. Expertsystem kan användas för att öka beslutsfattarnas och användarnas kreativitet, genom att de ger fler individer möjlighet att få mer kunskaper om det egna företagets verksamhet och förutsättningar.

*P G Holmlöv* *Bertil Thorngren*
Sekreterare Ordförande
TELDOK Redaktionskommitté

Författarna har själva skrivit in sina manus och gjort egna illustrationer med sådana av dagens persondatorer som kan sägas öka användarnas kreativitet och höja kvaliteten på arbetsresultaten. Vi tackar dem för detta.

# Innehåll

# KULTUR, SPRÅK OCH ARTIFICIELL INTELLIGENS

**Redovisning av intryck från den internationella konferensen "Culture, Language and Artificial Intelligence", Stockholm, 30 maj – 3 juni, 1988.**

av
Randy Whitaker, Olov Östberg och Ben Amick

## Inledning: Konferensens upprinnelse

EG-Kommissionen (Commission of the European Communities) bedriver ett antal tväreuropeiska forskningsprogram inom området vetenskap och teknologi (COST). Under 1986 startades ett nytt programområde kallat "AI and Pattern Recognition" (COST 13), i vilket gavs utrymme för projektansökningar avseende icke-teknisk forskning. Detta ledde till att ett forskarteam från Sverige, Norge, England och Österrike beviljades medel för projektet *AI-Based Systems and the Future of Language, Knowledge and Responsibilitiy in Professions.*

Svenska representanter i detta projekt var i första hand Bo Göranzon och Ingela Josefson från Arbetslivscentrum (ALC). COST-medel är avsedda för internationalisering av nationella forskningsinsatser. Detta medförde i detta fall att COST-medel primärt användes för stöd till en längre tids englandsvistelse för Göranzon och Josefson, samtidigt som Göranzon och Josefson i sin COST-forskning anknöt till projektet *Perspektiv på Analysverktyg och Arbetsformer vid Systemutveckling* (PAAS), vilket har bedrivits vid ALC sedan 1977.

Genom COST-projektet fick således PAAS-projektets medlemmar tillfälle till viktiga internationell anknytningar. Detta visade sig vara en mycket fruktbar stimulans resulterande i en rad vetenskapliga artiklar och böcker. I Sverige avknoppade COST/PAAS samtidigt på Kungliga Dramatiska Teatern en seminarieserie om "dialoger", vilket med stöd från Forskningsrådsnämnden (FRN) löpande dokumenterats i den nystartade tidskriften *Dialoger.*

# Grundtonen i konferensen

Konferensen *Culture, Language and Artificial Intelligence* var finalen på ett internationellt forskningsprojekt. Det finansierades primärt av COST, ALC och FRN, men visst stöd utgick också från bl.a. Tage Danielssons Minnesfond, Televerket och Infologics.

Ekonomisk stöd var behövligt. Konferensen var på en gång både en kultur- och forskningspolitisk engångsmanifestation och ett led i en långsiktig och sökande dialog. För att åstadkomma detta hade engagerats ett mycket stort antal nationella och internationella forskare, debattörer, musiker och scenaktörer. Med internationellt mått och med tanke på konferensomfånget och de rika inslaget av okonventionella konferensaktiviteter var konferensavgiften på drygt 6000:- ett relativt modest belopp. Ett stort antal forskarstudenter erhöll dessutom stipendier för täckande av kostnader för resa, uppehälle och konferensavgift. Som framgår av konferensreferatet i bilaga fanns dock individer som inte uppskattade konferensformen och som därför fann konferensavgiften väl hög.

Om konferensens grundton kan i övrigt sägas att den var klart influerad av de kontakter PAAS-gruppen (med omnejd) under COST-projektets gång etablerat i England. Därtill var konferensen i mångt och mycket en fortsättning på temat *Tyst Kunskap och Ny Teknik*, som med Göranzon som redaktör publicerats i boken *Datautvecklingens Filosofi* (Carlsson & Jönsson Bokförlag AB, 1983). Speciellt kan följande ses som en av utgångspunkterna vid planeringen av konferensens:

> *"[Utmaningen för 'semi-datoriserade' arbetare och tjänstemän blir att] upprätthålla och utveckla de 'tysta kunskaperna' samtidigt som nya datoriserade hjälpmedel kommer in i arbetslivet som ett inslag i yrkeskompetensen. Kampen gäller att förebygga attacken från försöken inom Artificiell Intelligens att trivialisera de 'tysta kunskaperna' som möjliga att systematisera och överföra till en automat."*
>
> (Bo Göranzon, 1983)[1]

# Grundtonen i vårt referat av konferensen

Ett referat kan naturligtvis inte ge rättvisa åt form och innehåll av en okonventionell konferensens. I själva verket handlade konferensen till stor del om det omöjliga i att åstadkomma sådana objektiva avbildningar.[2] I en speciell workshop diskuterades för övrigt översättarens

---

1 Citerad i P. Docherty, C. Werngren och A. Widman, *Informationsteknologi och Verksamhetsutveckling*. EFI, 1984.

2 Band- och videoinspelningar finns att köpa för den som vill penetrera enskilda konferensaktiviteter.

problem och omöjligheten att vara objektiv och sann visavi källan. Ett konferensreferat är inte bara en översättning utan dessutom en komprimering.

Vi har därför valt att mer fritt lyfta fram områden som vi själva funnit betydelsefulla eller som vi uppfattat att konferensledningen fäst stor avsikt vid. Till del är detta fria förhållningssätt en nödvändighet, ty under hela konferensen pågick parallellt ett stort antal olika workshops (och denna redogörelses författare var knutna till workshop nr 8 behandlande "Frågor rörande användning av ny teknik för överföring av kunskap och yrkesskicklighet").

## Kunskaps sammanhang och bakgrund

Expertsystem är en del av teknikområdet Artificiell Intelligens (AI). Ett mer adekvat namn på typiska expertsystem är dock kunskapsdatabassystem. Diskussionen om "tyst kunskap" handlar om möjligheterna/omöjligheterna att fånga upp människors kunskap och lagra i form av datorprogram.

Ett återkommande tema under konferensen var den avgörande rollen för sambandsfaktorer vid organisation och kommunikation av kunskap. En 'talhandling' (talakt, speech act) uttrycker inte hela innebörden i en situation. Den betydelse vi tillskriver ett objekt, ett uttalande eller en handling, är beroende av det sammanhang det förekommer i. I en samtalssituation är man beroende av situationens omständigheter för att rätt kunna avgöra vilken av flera möjliga tolkningar av en situation som är den för tillfället riktiga. Ett kunskapsbaserat datasystem gestaltar och omfattar en delmängd av människans kunskap. Det har inriktats på att verka inom ett avgränsat område. De regler som konstituerar en sådan kunskapsbas är definitionsmässigt oberoende av andra sammanhangsfaktorer än de som gäller inom av själva problemområdet. Fakta som är uppenbara för ett 2-årigt barn är inte uppenbara för ett datorprogram.

Anta att en kunskapsingenjör till ett datorprogram har överfört en modell av hur en människa uppträder och löser problem av en viss natur. Detta betyder inte att programmet också har fångat och kan återge människans dynamik i dessa problemlösningssituationer. En menings betydelse (avseende t.ex. tvetydighet eller skärpa) kommer med största sannolikhet aldrig att kunna överföras till ett expertsystem tillsammans med de formella kunskapsstrukturer som kunskapsingenjören är ute efter. Kunskapsingenjörens dilemma är att avgöra vilken sorts kunskap som kan föras över till datorn, dvs vad som kan 'representeras' eller 'förekomma' i datorn, samt avgöra hur denna kunskap skall kunna återanvändas.

Stor del av konferensen ägnades den historiska debatten gällande den vetenskapliga behandling av begreppet sammanhang (context). Forskningsområdet AI (och tillämpningsområdet expertsystem) beskrevs som den logiska fortsättningen av försöken inom vetenskapen att behandla information som något som objektivt kan uttryckas oberoende av de sammanhang informationen är relaterad till. Stephen Toulmin ansåg att en sådan syn på sammanhangsoberoende bottnar i följande tre vetenskapstraditioner:

> **Det perfekta språket.** Det perfekta språket kan förmedla människans alla begrepp i sinnesbilder. Försök att utforma ett sådant språk har som mål att utveckla ett allmängiltigt semantiskt system – ett symbolsystem så allmängiltigt att varje tänkbar idé kunde formuleras med hjälp av systemet. Sökandet efter det perfekta språket kan spåras tillbaka till Leibniz. Misstaget med föreställningen om det perfekta språket är antagandet om att det finns semantiska strukturer vilka gäller över alla kulturgränser. Till dags dato har inget perfekt språk kommit i dagen, ty vi förstår inte vad kunskap är eller av vilka beståndsdelar det består.
>
> **Den rationella metoden.** Den rationella metoden som består av de grundsatser och metoder som vår vetenskapliga tradition vilar på spårades tillbaka till Descartes. Denna metod har karakteriserats som logisk och vetenskaplig, i motsats till den retoriska och konstnärliga. Även om denna metod har bidragit med mycket inom de naturvetenskapliga områdena, så kvarstår emellertid frågor om dess förmåga att kunna avslöja mysterierna rörande människans uppfattnings- och tankeförmåga (perception och kognition).
>
> **Enhetlig (holistisk) vetenskap.** Förutsatt att det finns ett perfekt språk och en rationell metod skulle målet vara att inbegripa alla mänskliga frågeställningar inom en övergripande ram. Den centrala tanken bakom en universell, enhetlig vetenskap är att det skulle kunna finnas ett begränsat antal begrepp vilka ligger till grund för alla ämnesområden. Avsaknaden av en enhetlig vetenskap har till stor det berott på att man inte haft det perfekta språket samt att den rationella metoden misslyckats med att ta reda på vad kunskap är.

Flera slutledningar kan fås från Toulmins historiska analys. Den allmänna åsikten var att de filosofiska frågorna vilka ligger till grund för AI-debatten inte i sig själv var ny. Mer specifik var den iakttagelsen att språk förutsätter sammanhang, eftersom det är en spegel av den tid det används i. Vartefter språket utvecklas måste det också ändras i lexikonet. En idé som formulerats vid ett tillfälle kommer vid en senare tidpunkt att verka föråldrat.

Precis som uttrycksformen förändras kommer innehållet eller perspektivet av en idé att kunna skifta. Sue Bassnett-McGuire gjorde oss i den avslutande paneldiskussionen uppmärksamma på, att översättningar blir föråldrade och oanvändbara efterhand som de kulturella och språkliga sammanhangen förändras. Yuri Lotman påpekade att varje historikers analys av sitt ämne är sammanhangsberoende. På vilket sätt en analys av historiska dokument utförs varierar från tid till tid allt efter det nya fakta grävs fram och nya tolkningar tillskrivs dem. *Historikerns arbete går ut på att göra profetior om det förflutna,* enligt Lotman.

Precis som historikern gör kunskapsingenjören en socialt tids- och rumsbunden analys. Han tolkar den kunskapsmassa som finns tillgänglig vid tillfället för kunskapsförvärvandet. Det finns därför skäl att anta att en kunskapsbas inte kommer vara stabil över tid. – Kanske filosofernas tolkning av AI på liknande sätt kommer att ändras inom de kommade åren.

Detta leder till frågan om sammahangsberoende kunskap någonsin med fördel kan insamlas med hjälp av rationalistiska metoder och sedan förvaras i en stabil kunskapsbas. Språk och textöversättningar har påvisat de problem som uppstår när mekanistiska modeller och metoder används till sådana sammanhangsberoende företeelser som litteratur. Lars Kleberg betonade att vårt dagliga tal är beroende av sammanhanget. Sådant beroende gör maskinöversättningar i det närmaste omöjliga. Finns det därför anledning anta att kunskapsingenjörer verkligen skall kunna klara av logiska begreppsöversättningar?

För övrigt fanns kunskpasproblem förutom förändring över tid. Även kunniga människor arbetande inom samma område har svårighet att nå överenskommelse när de använder personliga metaforer för att beskriva samma sak eller fenomen. Gustav Östberg framställde sina erfarenheter inom materialteknologiområdet som exempel på sådana problem. Metallurger med molekylär bakgrund upptäckte att det i det närmaste var omöjligt att kommunicera med ingenjörer med mekanik- eller metallbearbetningsbakgrund. Beroende på vitt skilda synsätten om vad saken gäller var ingen av grupperna i stånd att utveckla ett gemensamt språk för dialog.

Horace Engdahl påpekade att det användbara i en databas i sig själv är sammanhangsberoende. Samma information kan vara mer användbar för en person än för en annan. Alternativt kan den vara lika användbar för båda men på helt olika sätt. Som exempel nämnde han att en bok har olika kunskapsvärde för en teoretisk expert och en praktisk ingenjör. Modeller för att strukturera kunskap inom datavärlden är beroende av starkt begränsade och abstrakta formler. Sådana modeller fångar inte de sammanhangsberoenden som genomsyrar de naturliga språken. Resultatet blir en omöjlighet för de nuvarande datasystemen att troget återge vidden av den kunskapsmassa som finns bland människor. Exempel på ett sådant tillkortakommande vad gäller maskinöversättningar är maskinens oförmåga att hantera ironi, det vill säga att använda ett uttalande för att uttrycka två olika meningar.

## Frågor om nomenklatur

Filosoferna diskuterade på ett abstrakt plan idéerna om sammanhangsberoende och "det perfekta språket". Samtidigt blev dessa problem tydliga och levandegjorda under själva konferensen. Under veckans gång fanns det tillfällen när dialogprocessen försämrades av att antingen olika meningar tillskrevs samma term eller av felaktiga antagandet om att man var överens om uppfattning av en term.

Ett sådant exempel åskådliggjordes i en workshop. Där kommenterade en av bidragsgivarna riskerna för auktoritativ påverkan i kunskapsbaserade system, vilka han benämde "beslutsstödssystem". Han kom från en cybernetisk bakgrund och uppfattade "beslutsstödssystem" som system kapabla till ett visst mått av bestutsfattande (om så ock bara på signalbehandlingsnivå). En annan deltagare ansatte honom genom hänvisning till sina egna erfarenheter av benämningen "beslutsstödssystem" såsom gällande system vilka stödde en beslutsfattare snarare än att gälla funktionen som beslutsfattande. De båda deltagarna var precisa vid användningen av ordet, men deras enskilda betydelser stod inte i överensstämmelse med varandras. Följande är ytterligare exempel på terminologisk förvirring:

**Intelligens.** Termen "intelligent" användes för att beskriva datorsystem som uppvisade flexibla beteenden samt system vilka verkligen försökte efterlikna mänskliga förståndsprocesser. Under konferensen diskuterades "intelligenta inlärningssystem", vilka för vissa av deltagarna var liktydigt med system som bara kunde klara av klasser av användarbeteenden medan andra deltagare hänvisade till dessa system som vore de intelligenta i en mer allmän AI-betydelse.

**Interaktion.** Denna term användes ofta felaktigt[3] för att beskriva växelverkan inom själva kunskapsbasen till skillnad från växelverkan mellan själva systemet och omgivningen. Några deltagare hade föreställninen att inom ett interaktivt system, kunde användaren verkligen åstadkomma faktiska förändringar i själva kunskapsmassan. De tycktes förvånade när de fick höra att denna interaktion bara sträckte sig till att låta datamaskinen välja vilket av flera fastställda datablock som skulle komma till användning.

**Regler.** Dataforskare använder detta ord för att beteckna en grupp av kunskapsstrukturer. Andra forskare använde det i en mer ursprunglig betydelse; en oflexibel föreskrift för beslutsfattande (såsom regler för hur man skall handla). Dessa två olika användningsområden ledde till förvräng-

3 "Felaktigt" enligt referenternas mening.

ning vad gäller mening och slutsatser under konferesensen. Den senare användningen innebär att ett regelbaserat system nödvändigtvis är stelt och auktoritativt, medan det förra alls inte innebär något liknande.

**Generell AI visavi specifika AI-tillämpningar.** Mycket av den kritik som drabbade AI hänförde bara till den grupp av användningsområden som går under beteckningen expertsystem. Liten-uppmärksamhet gavs andra aspekter på AI-forskning såsom planeringssystem, maskinell perception och maskininlärning. Något av kritiken gällde även mer specifika projekt. Som exempel kan nämnas att det kunskapsbaserade system som krävs för ordergivning och kontroll av "Stjärnornas Krig" (SDI) anses omöjligt att förverkliga inom överskådlig framtid, men att detta inom vissa kretsar lett till antagandet att hela AI-området är en omöjlighet. John Searl påpekade att det föreligger en generell missuppfattning om nivån av AI-teknologin.

**Naturliga språk visavi programeringsspråk.** En av de tre viktiga punkterna på konferensen gällde språk. Emellertid användes ordet så löst att det hänförde både till naturliga och formella språk, när sådan allmängiltighet varken var avsedd eller riktig. En tänkbar förklaring till skillnaderna kom fram i slutet av veckan när Kristen Nygaard skilde mellan beskrivande och föreskrivande språk.

# Frågor om tyst kunskap

Tyst kunskap var konferensens huvudtema.[4] De som betonade ämnets vikt föreslog att tyst kunskap är en form av mänsklig kunskap som inte kan programmeras in i en dator; denna kunskap ansågs bara vara tillgänglig för människan. De hävdade att ett av AI-forskningens huvudmål är att kunna uttyda sådan kunskap. En sådan betoning av kodifiering av det okodifierbara skulle kunna få konsekvenser för den praktiska användningen av AI. På arbetsplatser där systemen inte klarar av att på ett fördelaktigt sätt göra en modell av arbetsuppgifterna, skulle resultatet kunna bli ohumana arbetsuppgifter.

Detta var motivet för att diskutera frågan om tyst kunskap visavi programmering; dvs att det inte kan anses vare sig möjligt eller ens önskvärt att i datorer försöka programmera in sådan distinkt mänsklighet som kunskap. Konferenstiteln angav att med "datorprogram" skulle menas AI-program. Emellertid var flertalet konferensbidrag

---

[4] För undvikande av eventuella reklamationer från konferensdeltagarna borde detta stått som en konsumentupplysande underrubrik till konferenstiteln.

antingen så generella att de endast idé-mässigt berörde AI, eller också handlade de om traditionella datorprogram utan AI-inslag.

Filosoferna betonade att tyst kunskap inte går att överföra till symboler, men de tog aldrig upp vilken kunskap som är icke-tyst och som således skulle vara möjlig att återge med symboler. Anna Hart ifrågasatte betoningen av tyst kunskap, när det inte hade förekommit något verkligt försök att definiera begreppet mer exakt. Hon ansåg att det finns två typer av kunskap som är beroende av tysta förhållningssätt – det som är osägbart och det som man inte hört talas om, eller det som är osagt (som utelämnats eller gått förlorat). Dianne Berry hänvisade till den dubbelbetydelse av ordet "tyst" som förekom genom hela konferensen. Att något inte kan uttydas vid ett visst tillfälle behöver inte innebära att det inte kan bli det vid ett senare. Detta sade hon inte för att förringa betydelsen av sådan outtalad kunskap. Faktiskt visade experiment hennes psykologiska laboratorium snarast ett negativt samband mellan personers förmåga att utföra en uppgift och den grad som denna kunde uttryckas i ord. Hon drog slutsatsen att människor fick svårare att formulera sina kunskaper i ord ju mer skickade de blev att hantera kunskapen. Kanske skulle man lite vitsigt kunna tolka Berrys forskningsresultat som att:

*Om man inte är expert* **vet** *man inte vad man* **talar** *om.*
*Om man är expert* **talar** *man inte vad man* **vet** *om.*

Avsaknaden från försök att definiera "tyst kunskap" skapade en klyfta mellan filosofiska och praktiska synsätt, vilket kom till uttryck under konferensen. Några försökte överbrygga denna klyfta. Hart föreslog att man till en början skulle uppskatta vad kunskap symboliserar och därefter arbeta utåt för att finna dessas begränsningar. En annan fallbeskrivningsmetod förespråkades av Maja-Lisa Perby. Hon föreslog att man genom dialog skulle försöka förstå vilken slags kunskap som kommuniceras och används vid arbetsplatser för att därigenom förstå tyst kunskaps betydelse inom arbetslivet.

Herbert Dreyfus diskuterade frågan om tyst kunskap i samband med förvärvandet av yrkesskicklighet. Han definierade datorns uppgift till att vara ett försök att efterlikna hjärnans eget sätt att hantera symboler. Han påpekade att Minsky för 20 år sedan förutsåg att inom en generation skulle alla AI-problem vara lösta. Detta misslyckades till stor del beroende på att man inte kunde formalisera betydelsen av vad sunt förnuft är.

Vad menar vi då med "sunt förnuft"? Denna term användes på ett lika odefinierat sätt som "tyst kunskap". Aristoteles definierade sunt förnuft som ett sjätte sinne – ett sinne som bl.a uppfattar och känner igen de varseblivningstecken som missas av de andra fem sinnenas begränsade möjligheter. Det blev inte klart hur somliga deltagare uppfattade sunt förnuft, så detta blev ytterligare exempel på den termförvirring som framkom under konferensveckans gång.

Dreyfus konstaterade att problemet angående frågan om det sunda förnuftet förhindrat alla framsteg inom teoretisk AI. Han föreslog att AI-teknik enbart skulle användas inom de tillämpningsområden där det sunda förnuftet spelar ringa roll. Exempel på sådana "mikrovärldar" är klart avgränsade tekniska områden, samt den typ av avgränsade laboratorievärldar som behandlats av Terry Winograds SHRDLU-program för stapling av byggklossar. Andra tillämpningsområden befriade från sunt-förnuft-förutsättningar är schack och diagnoser av sjukdomar. Enligt Dreyfus är hör dessa till dags dato faktiskt just till de områden där AI-teknik visat sig framgångsrik.

Dreyfus hävdade att vi skulle kunna börja förstå begränsningarna hos AI-tekniken genom att undersöka **hur kunskap blir tyst.**[5] I en rad artiklar har han detaljerat redogjort för de fem stegen ledande till utveckling av *människans* expertis och samtidigt ledande till att kunskapen inom människan "automatiseras" och blir tyst. Med nedanstående 5-stegsmodell som grund för hur yrkesskicklighet förvärvas konkluderar han att expertsystem i själva verket inte är experter inom sitt område. Som bäst är expertsystemen kompetenta (vilket väl får anses vara ett gott betyg från kritikern Dreyfus).

(1) **Nybörjarstadiet.** Sönderdelar uppgiften i delar vilka är oberoende av sammanhanget och använder dem till att bygga upp en uppsättning av stränga regler. En sådan strängt föreskrivande struktur tillåter nybörjaren att handla riktigt i en mängd situationer, allt medan han/hon vinner ny erfarenhet.

(2) **Nybörjaren som gjort framsteg.** Strukturen med de stränga reglerna utökas och berikas nu med ny information, som erhållits vid meningsfull erfarenhet. Resultatet blir en sammanhangsberoende struktur som ger stöd för beslut. Denna struktur specificerar de villkor som ger upphov till en given följd – dessa villkor kallas maximer. En maxim är en beslutsledande stuktur som bara kan användas av den som besitter viss erfarenhet.

(3) **Kompetens.** Det finns nu så mycket information och erfarenhet att personen enbart koncentrerar sig på problemet eller målet. Vid denna punkt i utvecklingen börjar personen få en intuitiv känsla för problemet. Under de tidigare stadierna använde han/hon sig av abstrakta regler och maximer. Nu börjar personen internalisera processen. Vad som internaliseras är inte beslutsledande strukturer utan 3-dimensionella bilder av situationen och händelser som de kan förekomma i.

---

[5] Ett dilemma är att *om* en process till fullo har kunnat automatiseras, *så* skulle automaten därmed kunna anses besitta (maskinlig) tyst kunskap.

**(4) Skicklighet.** På detta stadium har personen erhållit en uppsjö av erfarenhet så att han kan reagera på ett sätt som leds av de internaliserade minnena snarare än från abstrakta regler och maximer. Betydelsen av analys minskar medan betydelsen av intuition och förmågan att plocka fram ur minnet ökas.

**(5) Expertis.** På denna nivå erhålls lösningen utan användning av analys. Ingen medveten tankeprocess är nödvändig för att välja lösning. Beslutet fattas utan hänsyn till de abstrakta reglerna och/eller maximerna. Nu kommer personen inte längre ihåg reglerna. Detta innebär att den som påstår sig kunna förklara reglerna inte har full expertförståelse.

Experter kan som bäst formulera maximer, inte de mer abstrakta och stela reglerna. Detta är orsaken till att expertsystem på sin höjd kan vara kompetenta. Dreyfus hävdar därför att tyst kunskap och expertis är samma sak.

Denna filosofiska diskussion om tyst kunskap kompletterades med synpunkter från arbetsmiljöhåll. Det diskuterades om relationerna mellan användning av yrkesskicklighet, tyst kunskap och automation. Olov Östberg tog upp exemplet med styrmannen som använder tyst kunskap för att styra sitt skepp (se **Figur 1**). För att styra ett skepp krävs dynamisk skicklighet, ty styrmannen måste hela tiden reagera på vinden och havets rörelser. Styrmannen kan motverka de olika påverkningarna på skeppet genom att känna till ännu fler manövermöjligheter mot havet. Ordet "cybernetik" kommer från det grekiska ordet för en styrman. Inom cybernetiken har exemplet med styrmannen generaliserats till **lagen om nödvändig variabilitet**. Denna lag säger att olika förutsättningar bara kan klaras av om lika många eller ännu fler förutsättningar finns tillgängliga för att motverka de tidigare. Med avseende på AI kan en maskin aldrig ha samma variationsmöjlighet som en människa och kan än mindre ha ännu större förmåga till variation. Därför kan människoproblem (stor variabilitet) inte fullt ut hanteras av en dator (liten variabilitet).

Ett annat arbetsmiljöperspektiv presenterades av Howard Rosenbrock, som studerat vilken roll tyst kunskap har i ingenjörens arbete. Rosenbrock exemplifierade sin syn med en "kunskapssol" som metafor (se **Figur 2**). Den explicita (icke-tysta) kunskapen finns i solens kärna och den tysta kunskapen symboliseras av auran runt solen. Enligt gängse förhållningssätt strävar ingenjören ständigt efter möjligheter att utöka kärnan och göra allt större kunskapsmängd explicit och tekniskt tillgänglig (vänstra delen av Figur 2). Ingenjören söker ständigt efter metoder att *göra det tekniskt omöjliga möjligt* och att *göra det tekniskt möjliga enkelt*. I praktiken betyder detta att ingenjören skapar tekniska lösningar för att ersätta människan.

Försöken att till alla delar teknifiera det mänskliga är dömda att misslyckas, inte minst för att de arbetsuppgifter som trots allt blir kvar därmed blir omänskliga: "Det här kan vi göra med en maskin, men det där är för svårt eller för dyrt, så där måste vi nog ha kvar en operatör". Den tysta kunskapen försvinner emellertid inte, den flyttas till de produktionstekniska experterna och till företagsledningen.

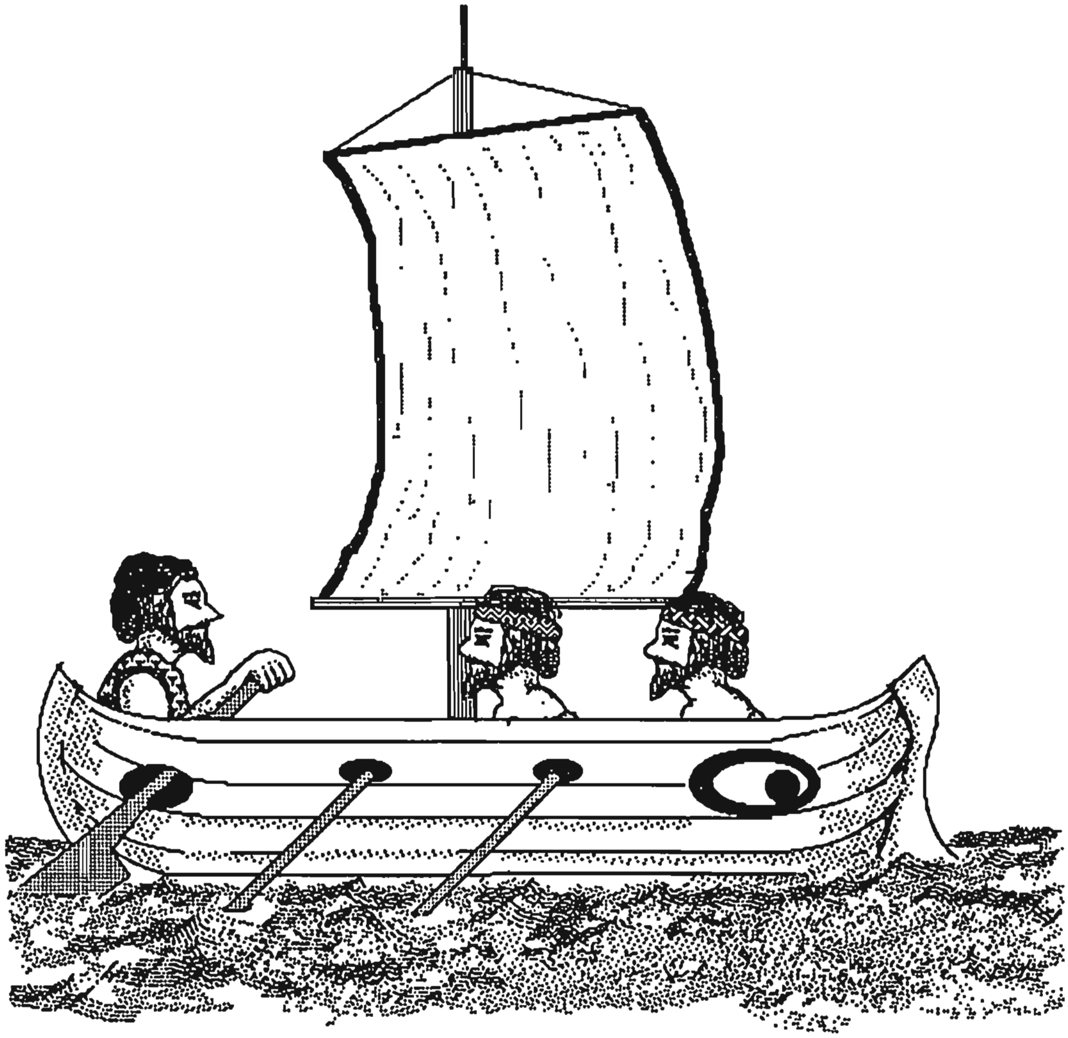

**Figur 1**

*Vind, vågor och ström skapar variabilitet i båtens gång. Enligt "lagen om den nödvändiga variabiliteten" måste cybernetikern vid styråran vara kapabel till minst lika stor motvariabilitet.*

Som alternativ föreslog Rosenbrock att en ingenjör skulle inrikta sitt kreativa tänkande i mer human-centrerade banor (högra sidan av Figur 2). Detta innebär att släppa ingenjörskravet på kontroll av en verksamhet till förmån för **kreativitet/flexibilitet** i verksamheten, samt att stödja människors arbete och acceptera och utveckla begreppet yrkesskicklighet. Detta är inte bara en fråga om hur industriarbete ska organiseras, utan är i lika hög grad en fråga om att ur ett ingenjörsperspektiv välja en konstruktions- och systemutvecklingsmetod. En tolkning av Rosenbrocks plädering är att *blomstringen av kreativitet och flexibilitet är beroende av en grogrund tyst kunskap.*

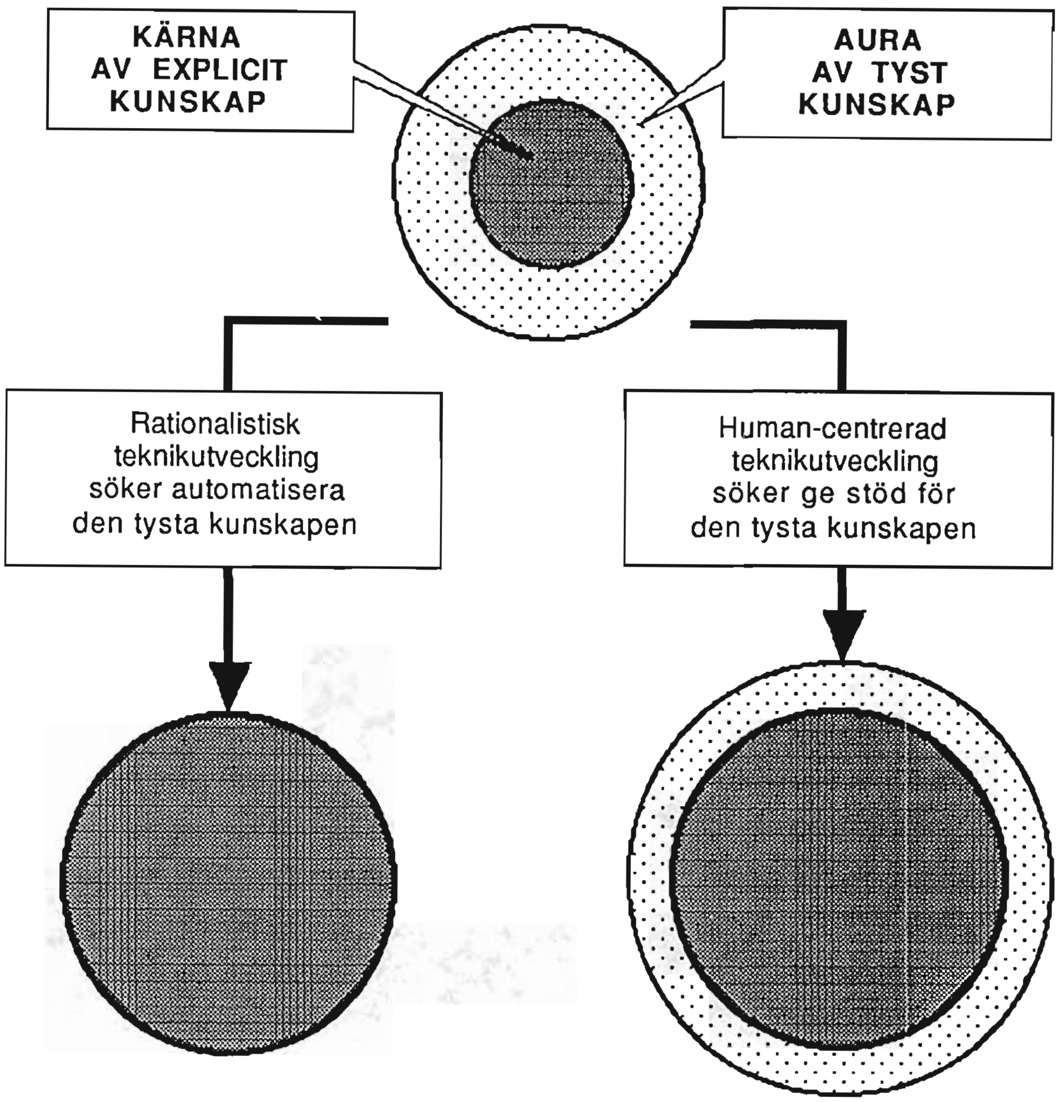


**Figur 2**

*Behandlingen av "tyst kunskap" i den rationalistiska, teknisk-vetenskapliga teknikutvecklingen, kontrasterad mot Howard Rosenbrocks human-centrerade teknikutveckling. Den rationalistiska paradigmen innebär att söka göra den tysta kunskapen explicit. Den human-centrerade paradigmen innebär att ge stöd för och se den tysta kunskapen som en tillgång.*

Betänk följande av Rosenbrock angivna exempel. En maskinarbetare använder sig av ett datasystem för en numeriskt styrd bearbetningsmaskin. Datorn/maskinen stöder arbetaren på två sätt. Maskinen tar fram de kvantitativa parametrar som behövs för att arbetet skall kunna utföras. Dessa parametrar kan sedan ändras efter arbetarens egen "tysta" syn på produktionsupplägget. Datorn/maskinen kontrollerar då denna nya uppsättning parametrar mot andra kända förutsättningar och indikerar vid behov frågor som arbetaren bör vara uppmärksam på. Detta illustrerar ett tekniskt system som är konstruerat för att stödja arbetaren i sin yrkesutövning, snarare än att få honom/henne att avhålla sig från att använda sig av tyst kunskap. Slutresultatet var bättre maskinanvändning, men, som Olov Östberg påpekade, inte genom att människan blivit "maskinvänlig" utan snarare genom att maskinen blivit mer "människovänlig". Detta kan bara uppnås när den övergripande konstruktionen är humancentrerad, dvs när målet för ingenjören inte är att programmera förbi den tysta kunskapen utan att ge arbetaren stöd till att kunna använda sin egen tysta kunskap. Detta diskuterades ytterligare i debatten om hur yrkesskicklighet överförs.

## Frågan om överföring/ förmedling av yrkesskicklighet

Det förekom en klar debatt om möjligheten för användning av AI vid överföring av yrkesskicklighet. Den kunde ta två riktningar, överföring från arbetaren till datorn eller användning av datorn för att förmå arbetaren att skaffa sig ny, kanske mer komplex skicklighet. Henrik Sinding-Larsen försökte på denna punkt konkretisera konferensdebatten på följande sätt:

Överföring av kunskap är:

| | |
|---|---|
| Möjlig och bra | (Optimistisk tekniker) |
| Möjlig/svår och bra | |
| Möjlig och dålig | |
| Omöjlig och dålig | (Pessimistisk humanist) |

Att konferensdeltagarna vidhöll dessa vitt åtskilda perspektiv försvårade de tidigare nämnda problemen angående nomenklatur. Inte desto mindre visade diskussionerna kring sammanhang och tyst kunskap på en vilja att försöka nå en gemensam grund att stå på som underlag för dialog. Den bästa möjligheten för en sådan gemensam grund tycktes vara genom uppslutning kring en forskningsinriktning angående arbetsmiljöfrågor.

Ett sådant program skulle kunna ta upp nuvarande tendenser vid datorisering av arbetsplatser, pågående arbetsmiljöforskning och framtida vägar för avgörande forskning gällande yrkesskicklighet och datorisering. Vad gäller industriproduktion identifierade Mike Cooley följande tre tendenser att beakta inom en sådan forskningsinriktning:

**(1). Förändring i "kapitalets organiska sammansättning".** Detta innebär en förändring från arbetskraft till information som bytesvara. Detta innebär att information i form av kunskap kan vara framtidens kapital.
**(2).** Det har skett en **ändring från det analoga till det digitala,** vilket fått två konsekvenser. Det första är den ökade kunnighet som krävs av de som måste använda maskiner. Detta kan leda till att det krävs större skicklighet för att klara av ett arbeta. Det andra är förmågan att kodifiera flera typer av arbete. Detta kan kan leda till en minskning eller ökning av den skicklighet som krävs för att klara av arbetet.
**(3). Maskinerna har tilldelats en aktivare roll i arbetsprocessen,** medan människan tilldelats en passivar roller. Som följd av att teknologin har tagit överhand ligger betoningen nu på att överföra kunskap från människan till maskinen. Den traditionellt logiska följden är att använda expertsystem för att kodifiera in experternas beslutsprocess i datamaskiner.

Datorer används som arbetsbesparande teknologi. Detta synsättet på elektronik/datorer såg Steven Deutsch som ett uttryck för strävanden att kopiera människan; robotar kopierar fysiskt arbete och expertsystem kopierar (snarare försöker kopiera) mänsklig varseblivning och förståelse. Slutresultatet av sådan kopiering blir att arbetaren avskaffas. Detta ger sken av att höja effektiviteten och att hålla kostnaderna nere. Enligt Cooley är det ett bedrägligt sken. Sett från ett bredare samhällsperspektiv skulle det inte bli någon besparing. Den synbara besparingen skulle gå åt till att betala ut arbetslöshetsunderstöd. Detta problem löses inte nödvändigtvis genom att omskola arbetaren till annat arbete. Deutsch tog nedläggningen av en bilfabrik i Kenosha i Wisconsin, USA, som exempel. Det upptäcktes vid nedläggningen att 40% av arbetarna i praktiken inte kunde läsa. Om denna händelse är en indikation på framtiden, kommer det att vara väsentligt att utveckla utbildnings- och träningssystem för att stödja människans egen vidareutveckling.

Avseende vägar för framtida forskning var det en allmän åsikt att en av tyngdpunkterna skulle vara på utvecklingen av human-centrerade system. Inom dessa system ses datorn som ett redskap till stöd för arbetaren. Cooley betonade att människan måste vara det centrala vid arbetsutformning. Detta förhållningssätt undviker att urholka arbetsinnehållet för de yrkesskickliga arbetarna och göra dem till maskinuppassare.

Som en parallell till studiecirklar och kvalitetscirklar introducerade Gerald Heidegger begreppet "designcirklar", genom vilka arbetaren blir delaktiga i själva designprocessen. Dessa cirklar skulle således koncentrera sig på att använda sig av teknologin till att förbättra arbetslivskvaliteten. Även om det inte tydligt sades, blev den logiska följden av dessa diskussioner om förhållningssättet till teknikanvändning, att kunskapsingenjörernas arbete med AI-teknik skulle vara human-centrerad.

Som konkret exempel på fördelarna med ett human-centrerat tillvägagångssätt påpekade Cooley kontrasten mellan maskinberoende inom produktionen i England och Tyskland. De tyska produktionssystemen var konstruerade med tanke på en större roll för människan. Engelsmännen lade huvudansvaret på maskinen, så att arbetaren ofta tvingades att avbryta arbetet vid tekniska processproblem. I Tyskland däremot kunde arbetaren ofta lösa problemet på ett annat sätt. Detta berodde på den tyske arbetarens flexibilitet, beroende på den mer aktiva rollen i produktionsprocessen. Som sammanfattning föreslog Heidegger att den viktiga frågan inte gällde teknologins användning vid arbetsträning utan själva arbetsträningen. Vi skall uppfatta expertsystem som arbetarens stödsystem hellre än skickliga experter i en låda.

En annan typ av värdefull undersökning skulle vara att titta på utvecklingen av yrkesskicklighet inom traditionella hantverksyrken. Detta för att förstå hur relationerna mellan yrkesskicklighet, expertis och tyst kunskap verkligen fungerar. Heidegger ansåg att den skicklige arbetaren följer rutiner och regler. Detta skulle innebära att den av Dreyfus lanserade 5-stegsprocessen skulle vara för begränsad. En arbetare kan vara en expert men inte en mästare (personen med expertis i Dreyfus schema). Den genomsnittlige arbetaren är inte medveten om 10-tusentals erfarenhtsbilder vid varje beslut. Han/hon är människa och som sådan begränsad av sitt minnes förmåga.

Peter Gullers föreslog att experten skulle kunna vara en "konstnär" i relation till tyst kunskap, överföring/förmedling av yrkeskunskap med hjälp mikroelektronisk teknologi. Konstnären inte bara avbildar eller omkombinerar element för att skapa ett konstverk, utan han/hon behärskar tradition och redskap och besitter yrkesskicklighet. Inom fotograferingskonsten är det stora variationer avseende både input och output av skapelseprocessen. Det samspel mellan olika beslutsfaktorer som gäller i det ögonblick bilden tas och det som gäller när bilden skall framkallas, avgör hur produkten blir. Detta återknyter till "lagen om nödvändig variabilitet", som illustrerats i Figur 1. Precis som den styrande rorsgängaren måste fortsätta att sitt båtfareri för att upprätthålla sin skicklighet, så måste fotografen ständigt fotografera för att upprätthålla sin skicklighet som konstnär. Nyckeln till att vara en skicklig fotograf är att föreställa sig den eftersträvade bilden, inte själva utförandet. Även om en kamera kan vara automatisk så kan inte ens den bästa automatiska kameran skapa en bild som är bättre än vad som är givet av komposition, ljus osv. För konstnären är tyst kunskap i form av yrkesskicklighet oberoende av teknologin, men kan alltid underlättas eller förbättras med hjälp av tekniska verktyg.

Tomas Tempte påpekade att det är hantverkaren som driver fram utvecklingen av verktyg inom de traditionella hantverksyrkena. Hantverkaren måste ha en holistisk föreställning om produktens avsedda form. När en gång en prototyp skapats, kan många kopior göras genom att (med hjälp av automation) imitera det ursprungliga sättet på vilket den första modellen kom till.

Ta det speciella exemplet om Temptes föresats att replikera en egyptisk stol från Tutankamons tid. Hans ursprungliga mål var att imitera. För detta krävdes att mot bakgrund av idag tillgänglig teknologi analysera originalstolen och den kreativa handling som gett stolen dess speciella form och utseende. Han erkände att för att kunna tillverka en sådan stol kunde alla deluppgifter åstadkommas av en maskin – allt utom vävningen av stolsitsen. Tempte förelog att hantverk var grunden för allt arbete, samt att det var ett led i en tradition gemensamt vidareförd av alla hantverkare. För honom var teknologin ovidkommande för själva arbetet.

Trots dessa diskussioner gjordes under konferensen aldrig några egentliga försök att ta upp frågan om AI, arbetsliv och överföring/förmedling av yrkesskicklighet. Istället koncentrerades uppmärksamheten på möjligheten att bestämma kriterier för att hantera och/eller utveckla begreppet tyst kunskap.

Detta lämnar en mängd frågor obegrundade och obesvarade. Vad är motsvarigheten till en traditionell, yrkesskicklig hantverkare i dagens teknologiska samhälle? Är programmeraren den nya hantverkaren? Vad händer med arbetsorganisationsfrågorna? Måste expertsystem alltid ersätta mänsklig insats? Vad är det unika med den kunskapsbaserade datorns roll när det gäller överföring/förmedling av yrkesskicklighet? Är dess huvudsakliga roll att utbilda/träna arbetare? Är den av Dreyfus lanserade 5-stegsskalan relevant vid diskussion av överföring/förmedling av yrkesskicklighet? (Flera av konferensdebattörerna ansåg att hans modell för "linjär" och att den inte förde frågeställningarna vidare.)

Man kan fråga sig om försöket att hitta vägen att använda AI vid överföring/förmedling av yrkesskicklighet har övergivits. Betänk att Cooley och Rosenbrock synes givit upp försöken att överföra/förmedla yrkesskicklighet. Istället koncentrerar de sig på att stödja den yrkesskicklige arbetaren genom att via bättre gränssnitt mellan människa och maskin betona relationen mellan människan och maskinen. Som många av konferensdebattörerna emellertid påpekade: endast genom att använda AI får vi möjlighet att verkligen utreda hur AI fungerar i arbetslivet.

Nygaard definierade AI som 1001 intressanta tekniska hjälpmedel på jakt efter något att hjälpa till med. Rosenbrock observerade att nuvarande rationalistiska teknikutveckling har resulterat i svåra konsekvenser för de arbetare som fått känna på resultaten av att AI-forskare och AI-tillskyndare jagat tillämpningar. Med facit i handen är det befogat att betona vikten av Rosenbrocks filosofi om human-centrerade konstruktions- och systemutvecklingsmetoder, som rätt tillämpade förhoppningsvis kan bidra till att förhindra ett upprepande av gårdagens (och dagens) misstag.

# Frågan om kreativitet

Begreppet kreativitet diskuterades livligt under hela konferensen. Så länge vi saknar kunskap om den kreativa processen kan den inte kodifieras och representeras i form av ett datorprogram. Delvis var diskussionen figurativ och behandlade kreativitet som den användbara mänskliga förmåga en maskin aldrig skulle kunna besitta eller överträffa. En illustration kan vara Lars Gyllenstens hänvisning till "Columbi ägg". Den historiska anekdoten är, att "de lärde männen" vid spanska hovet gick bet på det problem Columbus förelade dem, nämligen att få ett ägg att stå upp. Columbus egen lösning var att fatta om ägget och knacka det mot bordet – varefter det tillknackade ägget naturligtvis utan svårighet kunde stå på upp. Att lösningen fungerar demonstreras i **Figur 3**.

Denna anekdot anfördes som exempel på "insikt" som en del av en kreativ process. En datamaskin konfronterad med samma problem skulle måhända ha räknat ut förutsättningar för balans vid alla möjliga konfigurationer av kontakt mellan ägg och bord, inklusive hänsynstagande till bordets ojämnheter, äggulans position innanför skalet, osv. Datamaskinen skulle troligen ha redovisat att någon möjlighet för ett stabilt upprättstående ägg inte fanns, och skulle således aldrig kunnat uppvisa det insiktsfulla beteende som demonstrerats av Columbus.

Intressant nog togs denna anekdot med Columbi ägg åter upp under en senare paneldebatt. Debattören undrade om Columbus egentligen uppvisat kreativitet. Kanske var det frågan om att ha monopol på problemställningsdefinitionen, om att på ett opportunistiskt sätt demonstrera nyttan av att kunna ta i med hårdhandskarna ("kreativt våld"), samt om att känna till att ägget var hårdkokt. För övrigt var ju Columbus samtida lika bristande i kreativitet som de datorprogram som nu utdöms såsom varande icke-kreativa.

En annan diskuterad aspekt på begreppet kreativitet gällde dess roll i arbetslivet. I detta sammanhang frammanades bilden av att i den mån kreativitet var en komponent i en produktionsprocess, kunde förloppet inte reproduceras av enbart tekniska komponenter. Detta argument kommer från övertygelsen att grunden för kreativt tänkande och kreativt handlande inte kan uttryckas med symbolspråk. De kreativa processerna har studerats sedan årtusenden utan att någon klarhet har kunnat bringas om när kreativiteten träder fram och vad kreativiteten resulterar i. Denna gäckande skapelseprocess illustrerade Magnus Florin med ett citat från Om Poesi i Sak, som Carl Jonas Love Almquists skrev år 1839:

> "Vad skall man däremot säga om dessa besynnerliga varelser – de verkliga artisterna – som gråta när ordentligt folk ler, och ler när människor gråta? Äro de i sig själva melankoliska? eller löjliga sällar? Äro de nyktra av sig, eller månne de icke för mycket gärna taga till livs? Månne de äro rätt kloka? Hos dem finnes ingen plan, men allt vad de gör blir plan utav. Vem hjälper dem då och ställer allting så tillrätta för

dem? Den gör den därom, i vars osynliga sällskap de gå. Man kan aldrig säga att de äro borta, och aldrig hemma: ty på torget stå de som inne på kammaren, och sitta de någon gång i så kallat eget hem, så är det ändå för dem som vore de på gatan, i skogen eller på ängen. De dragas med en absolut relativitet, och tvärtom; men vilken begriper detta ord, som tycks innebära en contradictio in adjecto? Artisten är en sådan kontradiktion. Han gör på lediga stunder upp teorier, som han aldrig följer då han kommer till själva arbetet; när han arbetar följer han någonting, som han aldrig gjort upp. Ingenting är således opålitligare att läsa än hans teorier: de kunna säga mycket, som är rätt bra men röra dock aldrig det innersta. De kunna ej röra det. Ty vad vet konstnären egentligen. Ingenting. Och vad följer han då? I grunden blott den (agotho) dämon, som bemäktigat sig honom. Mot denne endast är han fullt trogen. Med honom går han, ty han är trogen. Konstnären är således den ärligaste skälm Gud har uppfunnit åt oss: likväl råda vi ingen att lita på honom så, som man förstå med att lita på. Han håller ofta mycket mer än han lovat; men jämt upp det, som han lovat, har han en inrotad avsmak för att hålla."

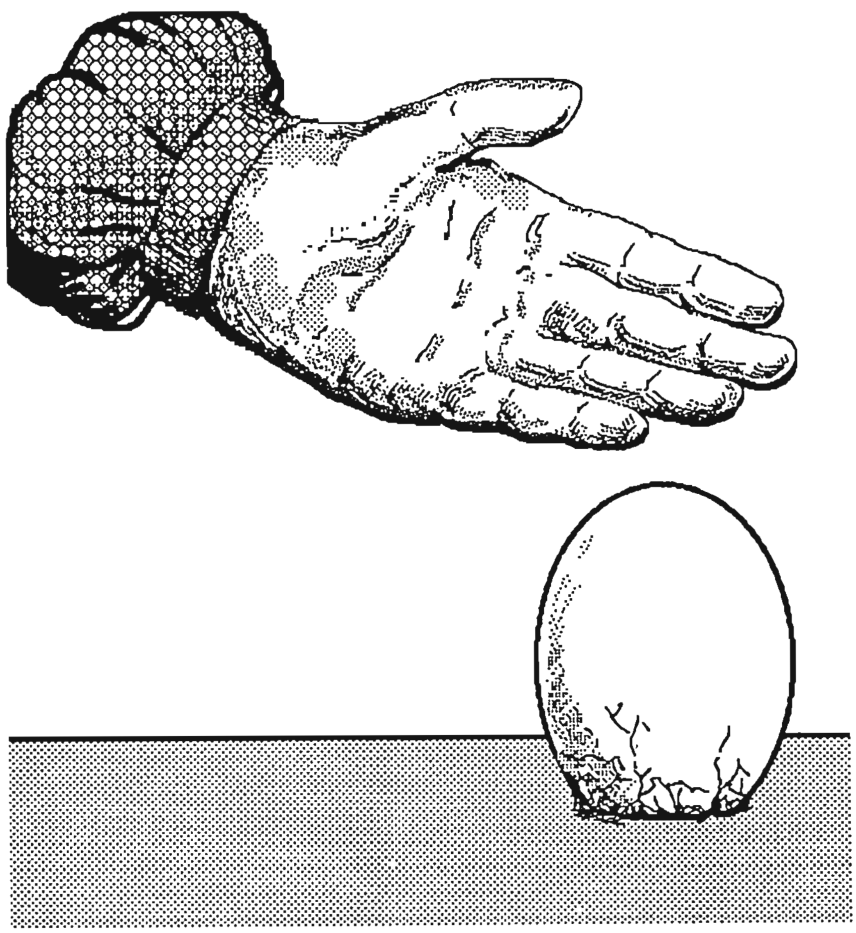

Figur 3

*Columbus och det stående ägget. Var lösningen en akt av kreativ insikt, burdust maktspråk eller kreativt våld, eller var det en fråga om att ha initiativet vid definitionen av problemet samt att välja ett hårdkokt ägg?*

Ann Buttimer rapporterade om sin studie från en konferens hållen i Sigtuna år 1977. Hennes arbete är ett försök att avmystifiera (inte nödvändigtvis att kodifiera eller kvantifiera) den kreativa processen. Undersökningen grundade sig på det sätt på vilket individer gått till väga för att utföra sitt arbete. Hennes undersökning utgjordes till stor del av anekdotiska observationer och självrapporter från hypotetiskt kreativa personer (mycket subjektivt utvalda), men är av intresse genom att erbjuda generalisering från företeelse till teori, istället för att utgå från en övergripande teori. Hennes studier ledde till följande 5-stegsprocess för kreativitet:

1. **Den initiala fasen:** Kreatören har en initial idé.
2. **Den intersubjektiva fasen:** Kreatören delar med dig av idéerna till andra, vilka värderar och förfinar idéerna till ett eller flera begrepp.
3. **Den analytiska fasen:** Kreatören upplöser och omorganiserar genom att använda sig av "arkiv" och/eller "experimentella undersökningar".
4. **Den syntetiska fasen:** Kreatören arrangerar om de analyserade komponenterna, så att de bildar en enhetlig enhet.
5. **Den kommunikativa fasen:** Kreatören rapporterar erfarenheter och resultat från den kreativa processen.

Buttimer fick i denna sin undersökning ingen antydan om att det skulle finnas något självskrivet standardsätt för hur en idé uppstår. I en uppföljande undersökning studerade hon närmare vilka omständigheter som ökar möjligheten till kreativ insikt. En grupp akademiker och yrkesverksamma undersöktes för att utröna var och under vilka omständigheter de tycktes vara mest kreativa. På basis av dessa personers berättelser föreslog Buttimer att det som befrämjar att en ny idé uppstår är (1) känslomässig stabilitet, (2) spontanitet, och (3) förekomsten av något som är välbekant men ändå ovanligt.

Eftersom man inte kan precisera vad kreativitet är, kan man inte heller ha någon grund för att bedöma kunskapsbaserade systems förmåga att stödja sådan aktivitet. Konferensdeltagarna tycktes ändå vara överens om att varje användbar AI-tillämpning inom detta område skulle ha formen av human-centrerat system (se Figur 3). I den utsträckning en datamaskin kunde användas som ett verktyg snarare än en övervakare, så kan den kreativa personen ha användning av dess kapacitet. Åtminstone, sades det, skulle den kombination av kunskap som erhölls på detta sättet kunde ge uppslag och alternativ.

Diskussionen om kreativitet och hantverkare fördes med hänvisning till yrkesverksamma personer inom redan etablerade konstarter och hantverksområden. Detta kan ha varit något korttänkt eftersom det vartefter tiden går, kan komma att utvecklas nya områden som ger utlopp för kreativitet. Varje förbättring inom teknologin skapar nya medier och nya härtill anpassade skickligheter. "Utövare" vilka är beroende av dessa nya färdigheter undersöker naturligt nog de nya medierna avseende nya uttrycksmöjligheter. Till skillnad från tidigare hantverkare (representerade av t.ex. Tempte), som koncentrerar sig på att upprätthålla traditionen och att arbeta inom vedertagna gränser, har dessa nya hantverkare ingen given tradition. De måste därför lära sig sina mediers gränser och möjligheter från början.

Under en paneldebatt påpekade Olov Östberg och Thomas Bernold att "hackers" ("dataknuttar") liknar gamla tiders hantverkare. Dessa har börjat behärska sitt medium – datamaskinen – till den grad att de kreativt och flexibelt kan producera mjukvara som tillåter och ibland utnyttjar (nedärvda) fel i maskinen.

## Frågeställningar i fokus

Fanns det en röd tråd eller en brännpunkt i konferensen? Om så, vem anslöt sig till den och vem kom den till godo? Verklig AI-användning diskuterades aldrig och det var ett envist fokuserande på det metafysiska. Den filosofiska AI-debatten koncentrerades på teoretiska överväganden om kunskapsingenjörernas möjligheter att i form av datorprogram fånga in mänsklig skicklighet. Överväganden om arbetslivskonsekvenser sträckte sig så långt som att i största allmänhet diskutera praktiska begränsningar av datateknologins förmåga att stödja arbetslivet, samt diskussioner om hur detta i så fall skulle kunna åstadkommas. Detta betyder inte att det med en annan konferens-inriktning skulle ha kunnat komma fram mer praktiskt nyttig kunskap angående AIs inflytande på arbetsplatserna. John Searl betonade att vi ännu inte vet vad teknologin kan åstadkomma. Många samhällsdebattörer har inte adekvat förmåga att rätt uppfatta AI-teknologin. Med denna oförmåga att rätt uppfatta det nuvarande läget, blir varje förutsägelse eller förväntan av diskutabelt värde.Varje beslut eller större förutsägelse grundad på osäkra uppfattningar snarare ökar osäkerhet.

De metafysiska diskussionerna resulterade i betraktelser av föga relevans för den process som ska leda fram till AI-system. Inte ens basbegreppet intelligens är tillräckligt väl definierat för att tjäna som utvärderingskriterium för "intelligenta" system. Att diskutera om människan är intelligentare än en dator, sade Searle, är som att diskutera om ett jetplan är snabbare än en fotosyntes. Att tillskriva datamaskinen mänskliga kvaliteter har bara förvirrat debatten kring AI. Här är några av dessa källor till förvirring:

> **Teknikmetaforer begränsar vår syn på tänkandet.** Eftersom vi inte vet hur vår hjärna fungerar, använder vi ständigt den nuvarande teknologin som metafor för att förklara och klargöra vår uppfattningsförmåga. Som med alla metaforer finns det gränser för dess användbarhet. Detta ger upphov till frågan, vilken emellertid inte togs upp under konfe-

rensen, om intresset för det centrala nervsystemet egentligen är ett exempel på det motsatta, det vill säga intresse för datamaskinens funktion.

**"Det mentala" ses som ett formellt, abstrakt system.** Allt vi egentligen vet om hjärnan är att den är biologisk. Det är inte säkert att uppkomsten av medvetandestrukturer, vilka anses utgöra människans intelligens, kan förklara vad varseblivning är. Som bäst kan sådan varseblivning tillhandahålla en beskrivning av intelligent beteende. Kanske kan forskning om hjärnans strukturer medföra att en biologisk modell av hjärnan skulle kunna bli vägledande inom AI-världen.

**Subjektiva kriterier används ej för att bedöma intelligent beteende.** Detta härstammar från den rationalistiska, vetenskapliga traditionen. Forskningen om AI har koncentrerat sig på att styrka en redan existerande formell modell, snarare än att simulera den verkliga världen. Man har glömt att det ofta diskuterade "Turingtestet"[6] för AI-intelligens inte är en formell bevisprocedur utan en subjektiv jämförelse mellan datamaskinen och människan.

**Maskinerna tillskrivs ett själsliv.** Vi säger ofta att en maskin handlar som en människa. Det är inte så tokigt, så länge vi kommer ihåg att vi menar att den handlar "som om" den vore mänsklig. Alltför ofta använder vi det emellertid ospecifikt, som om det betydde att en maskin faktiskt innehar mänskliga kvaliteter och/ eller kapaciteter.

Kanske är det bäst att acceptera omöjligheten med att jämföra människor och maskiner. Engdahl betonade att människor och maskiner är fundamentalt olika. Människor har en referenspunkt i världen där hans/hennes varseblivningar uppstår. Datamaskinen, som vanligtvis inte interagerar på ett sådant direkt sätt med sin omgivning, har inget sådant fokus. För att ge ett konkret exempel på denna speciella förgrening, påpekade Engdahl att varje begrepp om sanning inbegriper möjligheten för falskhet. Både sanning och falskhet bedöms med referens till omgivningen eller till en abstrakt modell. Eftersom datamaskinen saknar båda kan den inte utföra operationer som är beroende av sanning och falskhet, såsom att ljuga. Datorn kan bara vara riktig.[7] Å andra sidan kanske alla dessa jämförelser mellan människor och maskiner vare sig gör till eller från.

---

[6] Turingtestet säger ungefär, att om det inte går att avgöra om det är en människa eller en dator i den svarta lådan, ja då är datorn intelligent.

[7] En intressant parallell kommer från en domstol i England. En person blev anklagad för att med en falsk polett ha lurat en parkeringsautomat. Domstolen friade honom med motiveringen att det inte går att lura maskiner. Att lura är en term som bara gäller i människovärlden.

Nygaard lanserade ett Elfte Budord: *Du skall inte skapa maskiner till människans avbild* och ifrågasatte därmed värdet av "Turingtestet" för maskinintelligens. Han höll med om att det är ett intressant mål att *försöka* efterlikna mänsklig intelligens. Att ens komma i närheten av en verklig efterlikning kräver emellertid en sådana våldsam resursansamling att diskussionen om arbetslivstillämpningar är helt orealistiska. Fritt tolkat ansåg Nygaard att det är mycket viktigare att arbeta med "Turingtestets omvändning", dvs att utforma tekniken så att människa kan skilja mellan de gånger när de tilltalar en annan människa från de gånger när de vänder sig till en maskin.

Kognition diskuterades av dataloger för mer än 20 år sedan. Man väntade emellertid inte på frågornas lösning för att kunna fortsätta med den datalogiska forskningen. Nära nog samma frågeställningar ventilerades under konferens, och faktiskt ledde inte heller de fem konferensdagarna på Kgl. Dramatiska Teatern i Stockholm till att några frågor fann sin lösning (möjligen fann några lösningar sina frågor). Det är emellertid betänkligt att ett flertal av konferensdeltagarna levde i uppfattningen att datalogerna fortfarande diskuterar i samma banor som för 20 år sedan. Det fanns således tendenser till att slå in redan öppna dörrar. Mot detta kan ställas en under konferensdagarna växande samstämmighet om att de metafysika frågorna skulle nedtonas och att större uppmärksamhet skulle ägnas de problem som gäller införandet av expertsystem och andra AI-influerade komponenter och system till arbetsplatserna.

## Avslutande synpunkter

Under konferensens öppning berättade Erland Josephson om när en grupp forskare kom till Dramaten, just där konferensen avhölls. Forskarna var intresserade av teknologins betydelse inom teatern och utvecklingen av ett skådespel. De kunde emellertid aldrig förstå hur det var möjligt att under den korta tiden av 2 veckor, en så kreativ process kunde äga rum att en produktion faktiskt växte fram. Denna konferens involverade människor från vitt skilda bakgrunder, utan manuskript och med bara 1 veckas samvaro. Detta räckte tydligen inte, eftersom någon "produkt" i egentlig mening aldrig kom fram.[8] Kreativiteten kunde likväl föda fram uppslag till framtida insatser. Här några axplock:

> Maja-Lisa Perby föreslog en fallstudie för att ta reda på hur kunskaper fungerar på en arbetsplats. Detta för att koncentrera sig på praktiska dialoger i levande situationer.

---

[8] I enlighet med Yuri Loftmans syn på historiska data kanske framtidens AI-historiker hittar mängder med sådana konkreta "produkter".

Sue Bassnett-McGuire föreslog att ett sätt att närma sig översättningsfrågor var att arbeta på problem förankrade i nuet, såsom översättningar av serietidningar.

Karamjit Gill uppmanade deltagarna att vidareutveckla de framsteg som nåtts under veckan och att upprätthålla dialogen. Han föreslog speciellt att denna konferens skulle kunna vidareutvecklas genom att planera fortsatta möten i London vart tredje år. (Konferensledningen har redan börjat planera för en kommande konferens i London, men varför permanenta till just London?).

Yurji Masuda tillkännagav en gemensam konferens om tvärkulturella studier gällande yrkesskicklighet. En av huvudpunkterna skulle vara att undersöka den tysta kunskapens roll vid kulturella olikheter; tyst kunskap måste ha motsvarigheter som tyst språk och tyst kultur, något som betonats i flera av konferensbidragen. Han föreslog dessutom att det skulle startas ett internationellt projekt, för att därigenom fortsätta den dialog som påbörjats här i Stockholm. De första mötena skulle äga rum i Brighton (England) och Urbino (Italien).

Kristen Nygaard föreslog/krävde att om det skulle arrangeras ännu en konferens, skulle forskare inom data och andra involverade i AI utveckling vara bättre representerade.

Det fanns en genomgående liknelse till teatern under hela veckan. Denna konferens var faktiskt första akten i en process som skulle ge humanistiska synsätt större inflytande på AI-området. Avsikten med konferensen var att skriva ett manuskript med framtida dialoger om forskning. När nästa akt blir av, det vill säga akt nr. 2, skall ensemblen bestå av andra aktörer. Till fromma för dialogen bör då AI-praktikerna ges större utrymme.

DAGENS NYHETER Fredagen den 10 juni 1988

**Kultur**

**Jazzband, stråkkvartetter, paneldebatter, workshopar och happenings erbjöds på konferensen "Culture, Language and Artificial Intelligence" i förra veckan. Jag har nog aldrig varit på en konferens som i denna grad lyckats integrera vetenskap och underhållning, skriver Thomas Söderqvist.**

*Dramatenkonferensen om artificiell intelligens*

# Vetenskap är konst är show

UNDER efterkrigstiden har datoriseringen drabbat oss som en kulturell naturkraft. Datorn har blivit vår tids totempåle. Informationssamhället har blivit den stående programpunkten i Kurs- och Konferens-Sverige.

Först sent omsider har vi humanvetare, likt en flock minervska ugglor, börjat intressera oss för fenomenet. Den spirande debatten har fokuserats kring AI- (artificiell intelligens)forskningen och dess starka tes: Datorn ses som en tänkande maskin och t o m som modell för mänskligt tänkande.

Här delas humanvetarna i två läger. Det ena accepterar i princip den starka AI-tesen. Daniel Dennett, Douglas Hofstadter och Marvin Minsky tror att mänsklig kunskap kan representeras i datorernas symbolvärld.

En mer kritisk humanistisk strömning menar att detta är principiellt omöjligt. Man försöker nedvärdera AI:s betydelse, man talar om det AI *inte* kan, hur oformulerbart, outtalat och litet regelstyrt vårt handlande är.

## "Tyst" kunskap

Svenska forskare var tidigt framme. En grupp vid Arbetslivscentrum, med Bo Göranzon som entusiastisk samordnande kraft, började vid mitten av 1970-talet att studera datoriseringens konsekvenser för yrkesskickligheten. De frågade sig om det finns "tysta", eller underförstådda, yrkeskunskaper som inte kan formaliseras till de schematiska arbetsbeskrivningar som datasystemen kräver. De blev rädda för att människornas kunskaper och färdigheter ställs på undantag när kunskapsingenjörerna försöker utveckla sina expertsystem.

Detta utrednings- och forskningsarbete, som bedrivits i nära samarbete med arbetsplatser och fackliga organisationer (t ex i en välkänd studie av hur datoriseringen av sjukkasserutinerna påverkade personalens informella kompetens), sammanfattades preliminärt i boken "Datautvecklingens filosofi" (Göranzon, red.), 1983. Med "filosofi" menar man här i allt väsentligt den sene Wittgensteins arbeten. I hans anda ifrågasätter man vår kulturs tendens att sätta likhetstecken mellan sådant vetande som kan formuleras som påståenden (och således i princip representeras i en dator) och kunskaper över huvud taget.

## Kulturhändelse

Efter tio års förarbete har det nu blivit dags för internationell lansering. Under den vittfamnande titeln "Culture, Language, and Artificial Intelligence" samlades man i förra veckan till konferens på Dramaten i Stockholm. Det var i sanning storslaget. I den vetenskapliga världen markerar man ju revir genom att invitera till konferenser och seminarier. Ju större och internationellare sammankomsten är och ju större kanoner man lyckas engagera, desto bättre tror man sig kunna markera sin tolkning av världen. I denna mening var Dramatenkonferensen en lyckad reklamframstöt för Göranzon-gruppens idéer. Och de satte sannerligen inte sitt ljus under en skäppa.

Under fem förmiddagar kunde vi lyssna till tjugotalet föredrag varvade med jazzband, stråkkvartetter, paneldebatter och happenings. M A Numminen tolkade Wittgensteins "Tractatus" och Jan Malmsjö hade instuderat en scen ur Hugh Whitmores "Breaking the Code" (om Alan Turing, en av AI-forskningens pionjärer).

På eftermiddagarna kunde den intresserade deltagaren välja mellan fjorton parallella workshopar, allt från Karl Kraus och den österrikiska kritiska modernismen till medicinska expertsystem. Och om aftnarna serverades vi Stadshusbuffet och båtutflykt med teater ombord. Även mätt med den stockholmska kulturhorisontens allra strängaste kriterier var denna konferens således en av vårens stora kulturhändelser.

Min personliga favorit var den amerikanske språkfilosofen John Searle, som berättade hur han i slutet av 1970-talet sopade golvet med AI-forskarna vid Yale med hjälp av sina geniala tankeexperiment. AI-forskarna kan inte skilja mellan "inneboende" intentionalitet (som bara högre organismer har) och "som-om"-intentionalitet (datorn fungerar "som om" den tänker).

Stephen Toulmin gav den nödvändiga idéhistoriska förankringen. Han menade att AI-forskningens dröm om ett perfekt språk och en rationell vetenskaplig metod kan föras tillbaka till Leibniz. Dennes kamp för ett universellt matematiskt språk som kunde ersätta det oprecisa vardagsspråket borde förstås, menade Toulmin, mot bakgrund av dåtidens teologiska och kulturella kaos. Efter 30-åriga kriget var Europas folk vilda efter säkerhet, sade Toulmin, och antydde att vi också är det.

Berkeley-filosofen Hubert Dreyfus hävdade att AI-forskningen, som hade varit så framgångsrik under 1960- och 70-talen, nu hade utvecklats till ett "degenererat" forskningsprogram. AI-forskningen har inte bara misslyckats med att representera naturliga språk, man har t o m missförstått vad som rör sig inne i huvudet på experten. Han/hon tänker nämligen inte efter regler, utan i högsta grad intuitivt.

Och även om konferensens buttre gamle, den sovjetiske litteraturteoretikern Jurij Lotman, till synes inte hade upptäckt att veckans tema var språk och artificiell intelligens, var det ändå välgörande att höra honom tala. På en didaktisk och flärdfri ryska, långsamt tolkad till engelska, påminde han oss om att det är svårt att skriva historia. Historien är en serie överlagrade texttolkningar men samtidigt en irreversibel process av tillfälligheter.

## Jippoartat

Jag har nog aldrig varit på en vetenskaplig konferens (allra minst internationell) som till den grad lyckats med att integrera vetenskap och underhållning. När de vetenskapliga föredragen var som bäst utkonkurrerade de de officiella artisterna. Malmsjö reducerades till provinsiell dussinvara i jämförelse med Searles tjuvpojksaktiga professionalitet. Ur dramaturgen Magnus Florins perspektiv måste detta ha varit en mycket lyckad uppsättning.

Uppbådet av internationell vetenskaplig "show-biz" gjorde mig dock i ökande utsträckning illa till mods. Redan programmet, med dess blandade lista över "performers/lecturers", gav mig onda aningar. Bara inträdet gick på sex tusen, och med resor och hotell blev det omkring tolv tusen. Därmed har man uteslutit de flesta forskare och fria intellektuella från att delta. Visserligen fanns det ett antal "stipendier" till fattiga forskarstuderande, men därutöver var det bara folk med organisationsuppbackning som kunde vara med.

Men det var inte bara det med pengarna och tjusigheten som störde mig. Storleken är ett problem i sig. Det är tyvärr så att stora konferenser nästan aldrig lyckas förmedla den glädjens närkänsla och intellektets berusning som ett litet, intimt seminarium kan åstadkomma. Med över

Illustration: LARS-ERIK HÅKANSSON

30 personer blir det per automatik jippoartat. Väl att märka om man inte lägger ned ett mycket hårt arbete på att skapa stämningen av vi-som-kommit-hit-för-att-snacka-fritt-och-otvunget.

Det förarbetet hade man inte gjort. Och som dagarna skred fram insåg jag att man inte ens hade ambitionen. Tvärtemot – det var meningen att det skulle vara stort, dyrt och tjusigt. Det var Searle som gav mig lösningen fredag förmiddag. Han är själv en filosofisk entertainer av format (vilket han visade på sitt seminarium på filosofiska institutionen dagen före), och han gav järnet i sitt plenarföredrag. Men han visste också precis vad han gjorde: "Här är det mediet som är budskapet", sade han.

Det blev sanningens minut för mig. Form och innehåll smälter samman. Vetenskap är konst, konst är underhållning. I Göranzon-gruppens självförståelse förenas vetenskap och konst i dialogbegreppet. De har tidigare kört en serie seminarier på Dramaten om dialog. Men det var pinsamt att se hur konferensen misslyckades med att praktisera dialogen som kunskapsform. I stället blev vi utsatta för en i dessa sammanhang ovanligt mördande serie monologer.

Denna hybrid mellan vetenskaplig konferens och teaterföreställning svarar till uppsvinget för retoriken i humanvetenskaperna. Med retorik menar man då inte bara en mer eller mindre god förpackningskonst. Retoriken ses som en kunskapsform. Det finns mycket spännande i retorikens renässans. Men intresset för retoriken kan också medföra en alltmer cynisk inställning till kunskapen och makten.

## Massmöte

Inom vissa delar av svensk sektorsforskning har man i dag förlorat förmågan att skilja mellan vetenskap som kritisk diskussion och vetenskap som medel för makt och inflytande. Man tror att man kan lansera idéer på samma sätt som man säljer tandkräm eller rockband. Medan publiken får lov att sitta och klappa. När delar av publiken applåderade Horace Engdahl och buade åt Telecoms vetenskaplige chef Kurt Katzeff så visade man därmed att man hade förstått vad syftet egentligen var – vi hade blivit inkallade till ett sextusenkronors massmöte.

Att kalla sådant för konferens (som enligt min gamla uggleupplaga betyder "rådplägning öfver någon gemensam angelägenhet") är inte bara missbruk av språket. Det är ett hån mot idén om vetenskap som ett kritiskt, offentligt samtal. Om Göranzon-gruppen till en annan gång kunde tänka sig att utvidga sin filosofiska läsning utöver salig Wittgenstein, vill jag rekommendera ett studium av Karl Poppers "Det öppna samhället och dess fiender" (1980-81).

Workshoparna då? Var det inte där vi kunde få ta ställning till innehållet i Göranzon-gruppens teser? Men bortsett från att workshoparna också var schemalagda på förhand, så ändrar förekomsten av dem inte konferensens karaktär. En analogi kan belysa detta. Det är inte den fria diskussionen i utskotten som gör riksdagen till en demokratisk institution – det är förekomsten av kritisk diskussion i plenum som skiljer vårt demokratiska styresskick från totalitära regimers s k parlament.

Jag kan bara dra den slutsatsen att man inte hade ambitionen att ge kritiska röster en chans. Men inte ens den mest välregisserade föreställning kan kväsa vår lust att uppföra oss som intellektuella. Kritiken lyckades vid några få tillfällen bryta igenom. Både från panelen (Dianne Berry och Stephen Toulmin) och från salen (t ex Searle) luftades oppositionen mot Göranzon-gruppens teoretiska premisser. Begreppet "tyst" kunskap är vagt, mest ett slagord, sades det.

## Öppen dialog

Datoriseringen av Sverige skrider obönhörligt fram. AI-forskningens människo- och samhällssyn tränger in i organisationerna. Utan att behöva gripas av hysteri är det hög tid att vi humanister och samhällsvetare tar ställning till händelseförloppet.

Det kräver naturligtvis konferenser och seminarier. Men om dessa ska bli konstruktiva bidrag till det offentliga samtalet, måste de vara både kritiska och självkritiska. Inriktas på att öppet lokalisera problem och svagheter i våra nuvarande uppfattningar, i stället för att propagera den egna förträffligheten.

Tänk om Dramaten-konferensen hade satsat på att konfrontera olika forskningstraditioner med varandra i en öppen dialog! Gärna som nu med stöd av teaterns formvärld. Se det hade kunnat bli en riktigt lärorik föreställning.

THOMAS SÖDERQVIST
*Vetenskapsteoretiker, Göteborgs universitet*

# THE AUTOMATED EXPERT

## Technical, Human, and Organizational Considerations in Expert Systems Applications

Olov Östberg[1], Randall Whitaker,
and Benjamin Amick, III

Madison, Wisconsin, U.S.A.

FINAL REPORT SUBMITTED TO TELDOK, SWEDEN

1988

[1] Currently at Testing and Research Labs, Pla/C217, Swedish Telecommunications Administration, S–123 86 FARSTA, Sweden.

# TABLE OF CONTENTS

## PART II

# PREFACE

In **TELDOK Information Report No. 6,** *The Automated Office,* Helander and Östberg (1983) reviewed some human factors design issues pertaining to office automation. This report outlined US development trends from the perspective of Swedish work environments. A condensed and updated version of this report was later published by Helander (1985). The only artificial intelligence (AI) technology cited in the 1983 review was speech recognition. In the intervening years, another AI technology has come to the workplace – expert systems. Such systems hold the promise of capturing human expertise for preservation or distribution within an enterprise. The potential applications of this technology has been addressed by Hägglund (1986) in **TELDOK IReport No. 26,** *Computer Supported Knowledge Systems in the Office of the Future.*

The present report is a logical continuation of the 1983 work. The authors feel that the automated expert will be a routine component of the automated office. By "automated expert" we mean a knowledge-based expert system embodying the acumen of one or more recognized experts. Expansion of expert systems' penetration into the marketplace has been rapid.This would seem a prudent time to examine the human factors impacts of these systems. The agenda for this study was first described in a conference paper (Östberg, 1986). The study has since resulted in three scientific papers (Amick & Östberg, 1987; Östberg, 1988; Whitaker & Östberg, 1988), which should be seen as supplements to the present report.

This report was produced in Madison, Wisconsin, USA, where Olov Östberg spent an 18 months' leave of absence from The Swedish Telecommunications Administration (Televerket) as a visiting professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. Randall Whitaker is doing doctoral work on knowledge acquisition in the university's Department of Computer Sciences. His work on this project was supported by a doctoral fellowship from the Office of Naval Research, U.S. Navy. Benjamin Amick, formerly an analyst with the US Congress' Office of Technology Assessment, is an independent researcher and consultant in Madison.

The report reviews existing expert system implementations and analyzes the potential impact of such systems in the workplace. The focus is not restricted to office systems. Unfortunately, very few expert systems have been successfully installed in real world enterprises. As a result, the information base for this report derives from a combination of literature research, conference attendance, and visits to key researchers and expert system development sites. The authors would like to express their gratitude to the people who consented to discuss the state of the expert systems art. Unless otherwise stated, the views expressed in the report are those of the authors.

Olov Östberg
Stockholm, Sweden

Randy Whitaker
Madison, Wisconsin

Ben Amick
Madison, Wisconsin

# INTRODUCTION

## BACKGROUND FOR THE REPORT

To effect efficiency, human enterprises seek methods for both working harder and working smarter. Efficiencies can derive from a variety of strategies. The capacity for an individual's work can be expanded (for example, a farmer obtains a plow and the horse to pull it). The focus of production can be narrowed (the farmer specializes on one crop). Labor can be divided among the workforce (the farm wife and children adopt appropriate tasks). Labor can be multiplied via wage work (farm hands are hired for the harvest). Production resources can be pooled for the common good (farmers cooperatively build a steam mill). Similarly, pertinent knowledge can be centrally pooled for the common good (the state establishes an agricultural college). Accumulated expertise can be made widely available (agricultural extension agents are trained and stationed in farming areas).

The strategies listed above can been implemented by organizational changes, automation of physical labor, introduction of control and information facilities, or some combination thereof. The last two strategies – pooling and/or dissemination of knowledge – generally circumscribe the focus of the present survey. Our subject is the codification and/or transmission of knowledge via computers. Products attempting these functions are termed *knowledge-based systems*. We focus on the subset of these systems embodying human acumen within some tightly specified task domain – expert systems

Unfortunately, much of the available information on expert systems is of questionable value. It is rife with naive optimism, sales talk, and general hyperbole. There have been few attempts to systematically examine the range of issues entailed in developing, implementing, and using the technology. This report is a consideration of the technical, human, and organizational dimensions of expert systems. We intend to identify both current and prospective technological trends, with specific reference to impacts on the individual worker and his/her enterprise.

Over the past decade, these expert systems have begun migrating from the artificial intelligence research environment into the workplace. It has been estimated that the expert system share of the artificial intelligence market, currently $60 million, is expected to climb to $350 million by 1990 (DM Data, 1985). A specific motivation is that delivery of domain expertise to workers is expected to greatly enhance productivity. Feigenbaum has recently stated that his surveys indicate a tenfold minimum work speedup when a person is assisted by an expert system (Feigenbaum, 1988).

This technology, touted as the harbinger of a twenty-first century information society, may have unforeseen effects on the quality of work life. We feel there is some urgency in evaluating the technology due to its rapid proliferation and our belief that its application to date is largely antithetical to the trend toward sociotechnical perspectives on work design. This belief is based in part on the role of the military, especially in the United States, in managing the preponderance of expert systems research and development.

New technologies are commonly applied to military operations prior to domestic use; many innovations were specifically developed for warfare. In effect, war has served as a testbed for machines and methodologies. Whitby (1986) has pointed out that the computer was invented during World War II for military information tasks such as breaking codes and calculating ballistic trajectories. It is not surprising that the technical foundations for knowledge-based systems have been laid by artificial intelligence researchers underwritten by the Pentagon. In fact, all major research funding (outside of Japan) comes directly or indirectly from the military (Whitby, 1986). For an illustration of the military role in artificial intelligence research, refer to Figure 1.

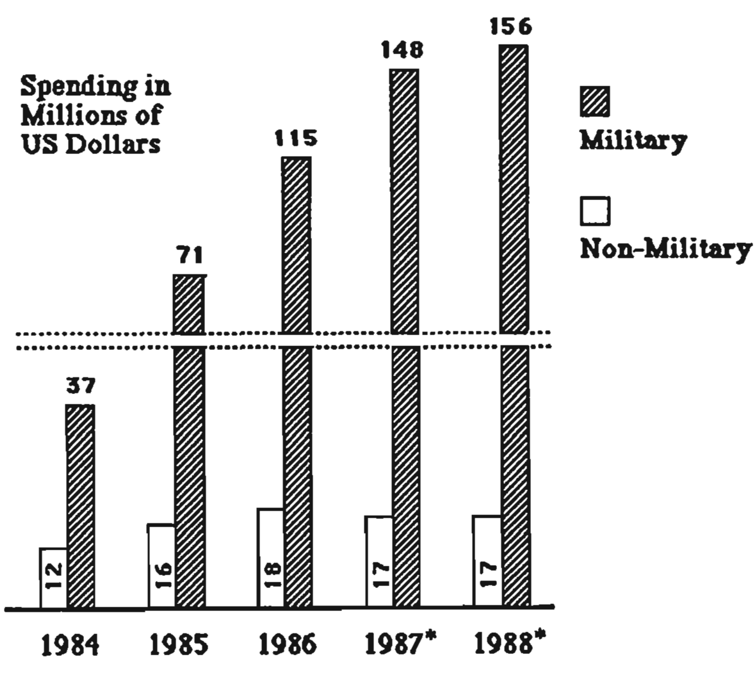


**Figure 1**

*Growth in U.S. federal funding for artificial intelligence, comparing military and non-military spending. These numbers exclude National Library of Medicine support, which rose from $1.5 million in 1979 to $4.0 million in 1985. (SOURCE: Office of Technology Assessment, 1987).*

* Estimated figure for spending

This heavy military involvement gives cause for concern on two counts. First, and most apparent, are issues pertaining to the efficacy of the systems themselves. As will be discussed later, in Part I, there exist no prescriptive methods for assuring expert system reliability. Nonetheless, knowledge-based systems are being developed for variety of real time tactical support purposes. Reliance on these systems means that a software bug could have disastrous effects for the battlefield. The most hotly debated such application is the command software envisioned for the Strategic Defense Initiative. A bug in this system could have disastrous effects for the entire planet, and the expected complexity of the program code would seemingly make bugs a certainty (Parnas, 1985).

The second cause for concern has to do with organizational issues. The strict hierarchical command structure of the military has served as the model for domestic enterprises. During the early days of the industrial revolution, strategists such as Clausewitz (1832) were consulted on ways to manage workers, because until that time only armies had cause to oversee large working populations. The hierarchical command structure worked in battle, so it was adopted for production. A reflection of this influence can be found in the drive to maximize organizational control beginning in the late nineteenth century. As will be discussed later, in Part II, this control revolution has had negative implications for the quality of work life.

So what has this to do with expert systems? Whitby (1986) has identified some typical features of military operations. Among them are emphases on reliability, rigid control, and total obedience. These features form a background for the manner in which expert systems have commonly been employed to date – to enforce consistency and control. By and large, industry has followed the military's lead, both in terms of the technology's form and in terms of its application. An expert system, applied as an instrument for organizational control, fits the needs of an army. However, there are questions concerning its appropriateness for the workplace. As we discuss in more detail in Part II, recent work suggests alternatives to a control orientation.

Bell (1973) has described a prospective post-industrial society – one organized around knowledge for the purposes of social control and directing innovation. He contrasts this vision with industrial society, which coordinated men and machines for the production of goods and services. Where labor was a key component of the industrial society, knowledge will be the key component of the post-industrial society. To reflect this increasing significance of knowledge, Bell identifies the growth of a new intellectual technology. This technology supplants intuitive human judgements with rational algorithms. His forecast presupposes that technical limitations would soon be overcome, and that the ability to make theory and research practical was imminent. This was, in effect, a promise of expert systems for society.

By attempting to understand the development of an information society, Bell and others hope to effect informed policies on research, development, and social issues. Like Bell and other forecasters, we seek to identify significant trends in expert system proliferation through examination of the current states of both the technology and its application to working life. Identification of relevant trends necessitates a broad, macro-scientific perspective. Unfortunately, expert systems to date vary widely among enterprises, there is much secrecy and/or exaggeration in the available data, and few systems have actually arrived in the workplace. This impels us to work from the micro-scientific issues of the technology's capabilities and issues of organizational dynamics.

## ORGANIZATION OF THE REPORT

This report is organized into two parts. The first is an examination of the status quo as of this date. We present a survey of the hardware, software, and methodological tools supporting expert system construction; issues of human/computer interaction; and the trend toward embedding expert systems within other technologies. The goal of this section is to familiarize the reader with the technology as it is, not as it has been optimistically portrayed. The second part of the report is an examination of the human and organizational issues surrounding expert system implementation. The goal of this section is to identify significant trends in the use of knowledge-based applications and to anticipate their consequences.

**Part I** of the report emphasizes the technological issues involved in expert systems. In 1973, when Bell wrote *The Coming of Post-Industrial Society*, expert systems were still objects of pure research. They were attempts to apply the techniques of artificial intelligence to concrete tasks. The literature was pretty much restricted to academic projects such as MYCIN and DENDRAL. In the ensuing fifteen years, these systems have left the laboratory for the workplace. This is a propitious time to address some questions about them. Is the technology really mature enough for dissemination? What are the capacities of the technology, both now and for the future? What is required to undertake development of such systems? How should such development be managed?

Doyle (1988) foresees major transformations in the field of artificial intelligence. This will be manifested primarily in the reorganization of disciplinary boundaries. Many of the fields which utilize artificial intelligence techniques will absorb these facilities into their own infrastructure. Some of the current concerns of artificial intelligence research (e.g., codification of knowledge and theoretical formalisms) will drift off as new specialties. Will expert systems, like the parent field of artificial intelligence, transform into something new? Will such systems dissipate as being unworkable? Or will the technology diffuse, becoming nearly indiscriminable among the applications it enhances?

We introduce a paradigm for expert systems based on the process of communicating knowledge. Structural definitions of expert systems fail to account for the systems' roles in organizations or the basis for their impacts on users. To date, developing expert systems has been described in terms of the implementation details. Such a perspective is ill-suited to an analysis of organizational issues. Organizationally, these systems represent the communication, augmentation, or outright transfer of skills. Using the metaphor of a communications channel, the paradigm directly addresses these organizational effects.

Part II of the report considers the human and organizational issues of expert system proliferation. Drucker (1980) has identified the embodiment of information and control in the machine as a fundamental characteristic of the second half of the twentieth century. This is considered as radical a change as the integration of the prime energy producer and machine tool at the close of the nineteenth century. With expert systems, we see an attempt to embody expertise itself within a machine. What is the relationship between skills and expertise? What are the consequences of transferring skills from the human to the machine?

Expert system design currently proceeds within a philosophical framework of rigorous analysis, seeking mechanistic interpretations of human actions. Casting the product as a representation of the human mind may be pretentious, but it is an allure of the technology (Barrett, 1987). Need this be the only choice, or are there alternative design perspectives? There is a great potential for alienating the worker from the work process and ultimately affecting the enterprise's effectiveness. Will the ability to controlthis technology be subverted by implicit design choices of the knowledge engineer (Gardell, 1988)?

We consider two fundamental design alternatives – either to automate knowledge work or to use knowledge to either better inform the worker or augment his/her skills. Examination of these alternatives are embedded in a discussion of the historical division of labor and the need for organizational control associated with differentiation and specialization. We contemplate the potential for extending the division of labor into a division of knowledge. Alternatively, a less mechanistic conceptualization of the labor process suggests other possibilities for organizational uses of expert systems, such as bridging islands of knowledge among workers.

# PART I

# EXPERT SYSTEMS: THE STATUS QUO, 1988

We initially planned this to be a report on installed products, such as those ostensibly working examples listed in Buchanan (1986). What we found in the field bore little resemblance to the success stories reported in high-priced insider newsletters. As of this date, we find that there are very few operational expert systems in use worldwide. What, then, of all the expert systems which the literature would have us believe are revolutionizing the workplace? Many of those constructed so far were academic research projects or "toy" prototypes. Of those intended for actual deployment, some were abandoned and many more have never in fact reached fruition. The literature, including the expensive and supposedly insightful expert systems newsletters, has consistently overstated the degree of expert system penetration into the workplace. However, even the outright failures have not dissuaded organizations from pursuing the technology; they have simply categorized previous efforts as learning experiences.

There has been no reluctance to bring knowledge-based techniques to bear on previously uncodified task domains. The imperfect development process has often been rationalized as an experiment which will hopefully produce useful results. Developers have been lacking in experience and success; this has not, however, prevented them from enthusiastically launching projects and optimistically reporting the outcomes. Apparently they believe that victory can be snatched from the jaws of defeat (Gillett, 1987). Leith (1988) noted that expert systems researchers are always keen to tell of systems which they have never actually used or even seen in full operation. He drew a parallel with the *Loch Ness Syndrome* – many folks claiming to have sighted the beast, while investigators are unable to confirm the phenomenon. The present authors can testify to the elusive nature of the living, operational expert system.

A good example of this Loch Ness scenario is General Electric's DELTA/ CATS – an expert system to support troubleshooting diesel electric locomotives (Bonissone & Johnson, 1984). DELTA/CATS is widely described as a success story similar to Digital Equipment's

R1/XCON. It has been a frequent subject in textbooks and journals. It was included in Buchanan's (1986) listing of installed systems. In short, DELTA/CATS seemed to be an ideal candidate for our field survey. We found that it had not in fact survived its field trial. When the umbilical cord from the development lab in New York to the locomotive shop in Nebraska was cut, the system expired.

This is not to say that DELTA/CATS was a poorly designed system. During its development, the builders realized that success would depend on the quality of the system's interactions with the locomotive technicians. To this end, much effort was directed toward effective graphics, friendly interface persona, and a help facility layered to allow for differential sophistication among the users (Dietz, 1986). The expert system interfaced with a video disk player, permitting the display of both detailed diagrams and instructive sequences taken from training films as necessary. In terms of anticipating the communicational needs of the end user, DELTA/CATS was an example of something done right.

All the good intentions and good ideas did not guarantee the eventual utility of the product. One explanation for DELTA/CATS' demise would be the marketplace. General Electric had not signed any new contracts on diesel locomotives for years. The real shortcomings, though, lay with the system itself, and these problems are typical of the immature technology. First, the developers viewed the project as a feasibility study. Second, critical restrictions on functionality were imposed by the mandated use of 16-bit microcomputers as delivery platforms. Third, the knowledge base (initially some 1200 rules) was not configured so as to allow worksite personnel to maintain and update the system. This last factor was symptomatic of the product's utter dependence on the development laboratory. When weaned from its "parents", DELTA/CATS was a nicely packaged static knowledge base covering a dormant market sector – another overspecialized, non-adaptive casualty of history.

In other cases, a system's shortcomings have not been fatal. Honeywell's COOKER process control advisor is considered a success, even though as much as 50% of the advice proffered is ignored by human operators. Attempts to isolate the reasons behind human resistance to the advice have been inconclusive. The system is still in use because even half the advice is sufficient to save the corporation "millions" (Cochran & Christopherson, 1987).

The Helena Laboratories' electrophoresis interpreter, implemented on a microchip, has been sold to over 1000 customers as a densitometer component. Contacts at Helena Labs were of the opinion that (1) the knowledge base, though adequate, no longer reflected the state of the art; (2) users were not relying on the advice given by the system; and (3) the expert system had not proven to be a significant factor in the marketability of the densitometers. Nonetheless, the contacts were pleased with the project as a learning experience, and they were optimistic about the prospects for subsequent generations of expert systems embedded within their instruments (Richmond & Landers, 1987).

The Swedish Agency for Administrative Development has established a unit responsible for expert systems development. It is known as the Greenhouse. To judge from our experience, this title is quite appropriate. The state of expert systems technology and knowledge engineering is one of infancy. Such systems (and perhaps their developers) still need a protective environment. Too frail to fulfill their promise yet, they will have to be fostered as nurslings for some time to come.

# TOOLS SUPPORTING EXPERT SYSTEM CONSTRUCTION

## WHY LOOK AT TOOLS ?

The current state of the expert system market is sometimes termed a plateau, where the early players are counting up their costs and profits while potential players are becoming more reluctant to jump into the game. Justification for this view is largely based on decelerating expenditures for knowledge-based programming environments (e.g., expert system shells) and high-performance hardware such as LISP machines. This could simply reflect the fact that enterprises have completed their initial resource acquisitions and are now exploring the technology. The proliferation of more sophisticated, cheaper development aids on lower-cost machines may also help explain the deceleration. This proliferation engenders optimistic outlooks such as the following:

> *"Using development shells designed to operate with personal computers, there can be relatively broad user participation in mini-systems creation, and this can lead to the collection and codification of a wide range of company expertise."*
>
> (Ernst & Ojha, 1986)

Our contacts, even those whose experiences had not been successful, retain confidence in the utility and the promise of expert systems. If there is indeed a broad decline in enthusiasm, we have not found strong evidence for it. In particular, we found no evidence that hardware, software, or procedural tools are generally constraining knowledge-based applications. While specific projects may have encountered problems involving tool limitations, there is no widespread claim of insufficient implements.

The subject of this report – human factors – is not completely independent of issues regarding tools. Whatever the capacity of the technology to adversely affect workers, its impact is not somewhere out in the future – it is occurring now. Our experience shows that the technology is spreading, and spreading quickly. There is an increase in the number of players and a broadening of the venues for their play. A tangible cause for this growth is the availability of tools sufficient to permit knowledge-based projects at ever decreasing cost. Let us examine the tools supporting expert system construction to both dispel fears of general insufficiency and illuminate specific points for prudent consideration.

# HARDWARE SUPPORTING EXPERT SYSTEMS DEVELOPMENT

## Memory – The Capacity for Storing Knowledge

Artificial intelligence programs have long been characterized as resource gluttons, requiring several megabytes of memory and vast quantities of secondary storage. There has been considerable debate regarding how much machine memory is sufficient for general AI programs. Norem (1985) argued that a one-megabyte chip would be sufficient to simulate the knowledge of a human with an IQ of 100. This argument was based on the assumption that human operational knowledge is a monotonic function of the number of active memories stored in the brain. Though relatively small by comparison, the Norem estimate is of the same order of magnitude as a sufficiency criterion set by the Japan Electronic Dictionary Research Institute (1987) – 200,000 words multiplied by one frame for each of a few languages. Also of the same order of magnitude are the encyclopedic knowledge criterion (Lenat, *et al.*, 1986) of 30,000 articles multiplied by 30 frames per article and Minsky's (1985) calculation of a rate of 4 entries per hour into human long term memory from birth to adulthood.

Lenat and Feigenbaum (1987) estimate that within a decade we will see silicon chips containing Very Large Scale Knowledge-bases – VLSK chips. However, the development of large expert systems need not wait for such chips. Presumably, the amount of memory necessary to support an expert system is less than the amount necessary to support a general AI program. Now that microcomputers boast over ten megabytes of RAM and hard disk capacities exceeding 100 megabytes, the availability of memory is no longer a constraint. All the memory capacity criteria listed above are achievable with current workstations and many microcomputers.

## Performance – The Capacity to Emulate an Expert

High-performance LISP machines were the workhorses of the AI practitioner twenty years ago. Over the past two decades, the operational power of state-of-the-art machines has filtered down into modestly priced workstations, even into personal computers. Specialized symbolic processors have progressed from exotic research testbeds to packaged commercial products. There are now drop-in boards for microcomputers which provide parallel processing capabilities. Figure 2 shows the trend toward microcomputer delivery platforms in the insurance industry.

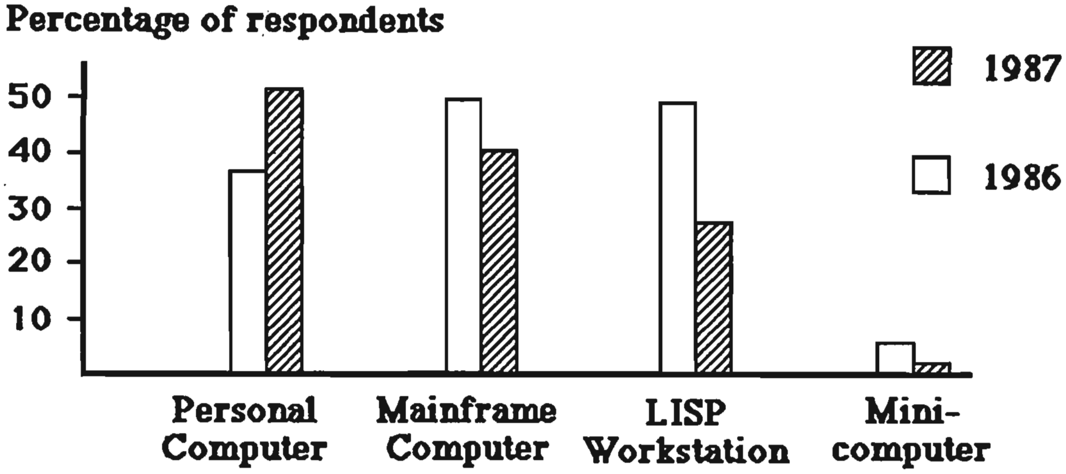


**Figure 2**

*Expert systems delivery hardware anticipated in the insurance industry. Note the trend toward the use of personal computers to support expert systems. (SOURCE: Coopers & Lybrand, 1987b).*

It would appear that at this point in time hardware performance is not a limiting factor to expert systems' proliferation. Of course, current performance constraints may preclude expert systems for certain high performance real time applications. This caveat is, however, dependent upon the performance demands of the intended task, the size of the requisite knowledge base, and any restrictions on the size or type of hardware to be utilized.

## Displays – The Capacity to Portray Knowledge

Hardware support for interface displays is good and getting better. High resolution color monitors provide powerful graphics capabilities. Rudimentary support now exists for speech recognition, speech generation, and even the input of handwritten text. An interesting application of display technology is the use of video disk players for the storage and display of images. This format allows for the use of both static diagrams and moving images. Video disk capabilities have been used in General Electric's DELTA/CATS and the National Library of Medicine's AI/RHEUM.

## Alternative Architectures

Connectionist (or parallel distributed processing – PDP) architectures are the latest development in a series of attempts to emulate the mechanics of neural structures. Ancestral research includes Hebb's (1949) network learning theories, Hull's (1943) unified mathematical models of learning, Rosenblatt's (1962) perceptron, and work undertaken in the area of adaptive control systems. The PDP concept of data representation is radically different from that of conventional von Neumann (sequential) computers. In the sequential computer, there is one central processor, and its data is contained in sets of bits at fixed locations in memory. The connectionist system consists of many processors. Information is contained in the pattern(s) of the states of these units. These patterns are a function of the degree of connectivity among the processor nodes and the strength of the connections (Rumelhart *et al.*, 1986). Since the thrust of this research has ostensibly been to model the human brain, PDP technology has been strongly identified with AI.

Parallel distributed processing networks display a host of intriguing properties. They are capable of learning by adjusting connection strengths in response to incoming stimuli. Ill-defined or noisy training inputs can be processed by PDP systems, whereas such inputs are problematical for sequential machines. These systems have therefore shown remarkable capacities for pattern recognition.

Unfortunately, it is unclear how symbols can be handled by PDP networks. This means that perspicuous knowledge representation may not be tractable in connectionist architectures. The PDP system's logic (if that be the right term) is contained in its connections, not in a language that can be examined and revised separate from the network. Consider this in light of knowledge base maintenance. In a rule-based expert system, behavior can be changed by modifying the symbolic content of the knowledge base. In a parallel distributed processing system, there is no declarative representation of the patterns (the knowledge) within the system. Modifications to the PDP system would be most easily done by retraining the network. For some task domains presentation of the necessary input and training stimuli could prove prohibitive.

Does connectionist technology offers anything to the knowledge engineer at this time? To the extent that this technology offers flexible high-performance pattern recognition in peripheral devices, the answer would certainly be "yes". Beyond that, answers are elusive. There has been work directed at building expert systems using connectionist architectures. This work has been limited, however, to problem domains where the task solution can be portrayed as a classification system. Tasks that are ill-structured or incorporate much procedurality are admittedly unsuited for these PDP approaches (Gallant, 1988).

### Summary

It appears that at this point in time hardware is not a constraint on the proliferation of expert systems. Not only do adequate machines exist, they are available with increasing power at decreasing cost. While PDP architectures are touted as providing representational and learning advantages over traditional sequential machines, it is unclear how they could support knowledge engineering as we now know it.

## SOFTWARE SUPPORTING EXPERT SYSTEMS DEVELOPMENT

The term "knowledge engineering" is used to describe the whole process of codifying domain expertise and encapsulating it within an information processing system. The codification of knowledge via expert systems is well underway, even though the process of doing so is still widely considered a black art. There has emerged no single set of resources, principles, or methodologies essential to the development of expert systems. This suggests that implementation details are indeed secondary to the quality of the knowledge transfer process.

The efficacy of this process depends on the representational power of the knowledge transmission medium. Knowledge domains modelled to date typically exhibit high degrees of complexity – complexity with respect to discriminable knowledge structures, their organization, and the operations performed over them. Complexity would seem a consequence of the knowledge being worthy of an expert; cynically, it is simply a consequence of the real world. Managing complexity is easier where large masses of information can be addressed and manipulated in forms closely resembling natural language representations. Programs supporting such manipulations exemplify *symbolic* processing (of complex conceptual structures) rather than *data* processing (of elementary data structures).

Development tools associated with expert systems emphasize ease of symbolic manipulation. These tools address information at a very high *level,* where the term "level" denotes a vertical dimension of conceptual abstraction. A lower level programming tool deals in bits, bytes, and elementary data structures (e.g., arrays, trees, and lists). A higher level tool can handle more sophisticated constructs (e.g., rules, frames, etc.). Higher level tools permit system programmers and end users to address the computer using syntax more akin to that of their mental models of the task domain.

## An Example

We turn to an example from Hardy (1987) to illustrate the advantages afforded by high level symbolic manipulation. Let us suppose we wish to model the expertise of an expert light bulb fixer. Let us also suppose that we are restricted to a non-computerized representation of the knowledge. We could devise a mapping from the possible fault situations to the correct procedure(s) for repair. In this case, we will model the fault situations using only two states – the state of the switch and the state of the bulb. Organizing this information into tabular form, we produce a decision table:

| SWITCH STATE | BULB STATE | REPAIR PROCEDURE |
|---|---|---|
| Off | (any) | Turn switch on |
| On | Rattles when shook | Replace bulb |
| On | Doesn't rattle | Check wiring |

This table could be distributed to other fixers in the form of a reference card or (for bigger tables) a diagnostic manual. The responsibility for finding the correct procedure lies with the human fixer. He or she must be able to read the table and associate the conditions and consequents via scanning the rows of the table. A literate fixer can add to or modify the table by writing in new lines.

Now let us relax the restriction against computer-based solutions. Given a microcomputer and a BASIC interpreter, we could translate the decision table into the following code:

```
 10   input "Is the switch ON or OFF?", switch$
 20   if switch$ = "off" then advice$ = "Turn switch on"; goto 100

 30   input "Does the bulb rattle when shook?", rattles$
 40   if rattles$ = "yes" then advice$ = "Replace bulb"; goto 100
 50   advice$ = "Check wiring"
100   print advice$
110   stop
```

The decision table loses much of its textual clarity during translation. The code itself could hardly be printed out and used as a guide for diagnosis. One reason is that the BASIC code uses a restricted syntax unlike the decision table's natural language description. Another reason is that the information contained in the code is not entirely related to the diagnostic task. We find line numbers, variable names, input/output commands, branching commands, and execution commands. None of these items is of any use to the human fixer. Finally, the end user cannot readily record changes in the knowledge; such modifications must be done by a programmer. Having enlisted the services of a silicon assistant, we must now rely on it to follow an increasingly cryptic specification.

The inclusion of control information intermixed with the task specification is typical of *procedural* programming. In the procedural paradigm, the programmer is responsible for outlining the detailed actions necessary to do a job. These actions are explicitly described in the implementation code. The alternative view is that of *declarative* programming. In this paradigm, the knowledge representation syntax resembles the sentential forms of natural language. A programmer specifies the nature of the task domain in a relatively perspicuous fashion, e.g., by stating appropriate rules in an IF-THEN format. The procedural aspects of the computer are hidden from view, because the knowledge representation is maintained separately from the control coding. To illustrate, consider the following declarative representation of our decision table:

```
IF switch = off
   THEN advice = "Turn switch on"

IF switch = on AND rattles = yes
   THEN advice = "Replace bulb"

IF switch = on AND rattles = no
   THEN advice = "Check wiring"
```

Here we have a depiction of the knowledge which can be understood by a human user *and* executed by the computer. The knowledge base can be perused without having to navigate through control expressions, and humans can modify the knowledge base by addressing and manipulating the rules without consideration of the control mechanisms underneath.

## Development Aids for Expert System Construction

Now that we have described the dimension of high/low level of information abstraction, let us consider existing expert systems in the light of these categories. Some expert systems built to date have been constructed using procedural languages such as FORTRAN or C. These languages are lower level tools – the implements of mainstream software engineering. An expert system necessitates symbolic manipulation, inferential ability, and (possibly) probabilistic representations. Programmers using these procedural languages must develop such capacities from scratch. Constructing expert systems in these environments is labor intensive and presumes a sophisticated (or masochistic) programming staff. The alleged payoff is in execution efficiency.

On the other hand, a majority of expert systems have been developed using higher-level programming languages (i.e., "AI languages" such as LISP and PROLOG) or construction tools called expert system shells. AI languages usually include the capacity for symbolic manipulations and/or inference. This obviates the need for the expert system builder to generate such facilities before undertaking the actual knowledge engineering task. Implementations of such languages are available for machines as small as personal computers. Such languages have been criticized for their relative inefficiencies in terms of execution speed and resource management. These disadvantages are disappearing due to both faster compiler implementations and more powerful hardware.

An expert system shell is a programming environment providing the developer with an array of knowledge representation and inference tools not easily obtained within any one language. One of the earliest such shells, EMYCIN (van Melle, 1980), was simply the expert system MYCIN with its domain-specific knowledge base removed. By adding a knowledge base, a MYCIN-like expert system could be devised for any appropriate domain. Similarly, the geological expert system PROSPECTOR begat the tool KAS (Reboh, 1981). Other shells, such as ART and KEE, are original products, not derivatives of some earlier development. These tools support the declarative paradigm, providing programmers straightforward access to the symbolic and inferential capacities required in knowledge engineering.

It is important to point out that an expert system shell may not be as representationally or logically complete (in the formal sense) as a programming language such as C, LISP, or PROLOG. These products may provide only a small subset of the available knowledge representations or inferential techniques. Developers using such tools will therefore be constrained in the manner in which they can model expertise. This limitation is most severe in the case of the empty expert system shells, such as MYCIN and KAS. Applying such tools to a project means that the developer must accept any innate biases (e.g., available knowledge structures, probabilistic mechanisms, user interfaces, and inferential strategies).

## How High Level Tools Facilitate Developers

With high-level programming tools there is little effort or time lost translating the original task specification (e.g., an expert's description or a decision table) into a new form (e.g., a rule base). Expert system builders can portray the domain knowledge in a form more congruent with the expert's own expressed model. This lessens the requisite sophistication of the developer to the point that some vendors suggest that their shells are friendly enough to permit the expert to model his own expertise!

One consequence is the conservation of effort. High-level tools ostensibly enable people to attack previously intractable problems. Certainly, this aids proficient programmers in expanding the range of tasks to which knowledge-based technology is applied. Perhaps just as important is that this affords less proficient folks the initial means to address jobs previously out of their reach. The result is a bandwagon effect – more and more practitioners trying out knowledge-based techniques on more and more problems. The optimistic view is that such tools stimulate creativity for all programmers (e.g., Tichy, 1987).

Such tools, although advertised for expert system construction, are more properly considered as extremely high-level programming environments whose primary application to date has been expert system construction. The power of these programming environments could be applied to construct software for algorithmic (rather than knowledge-based) tasks. The application of such knowledge-based tools to algorithmic (or at least highly deterministic) tasks has resulted in products which may not seem to embody much knowledge, but which are nevertheless labelled expert systems.

Like all major telecommunications groups, the Japanese telecommunications provider NTT is developing a whole series of expert systems. What makes NTT interesting is their work on systems which are designed for consultation over the telephone. ANGEL (Ohyama, 1987) will function like a telephone operator in answering callers' directory questions. DOCTORS (Tsumura *et al.*, 1986) will function like a medical doctor in diagnosing callers' headache symptoms. Both these systems, though of advanced design, will provide quite mundane assistance to the consultees. They are *high level* systems providing *low level* knowledge (Michie, 1983; Bobrow & Stefik, 1985; Brown, 1984).

Broadly speaking, there is a tendency for expert systems to become increasingly modest in scope. The task domains to which knowledge-based techniques have been applied are getting simpler. These latter day tasks require less reliance on expert knowledge, and they are less likely to be as ill-defined as the classic topics of early research (e.g., mineral prospecting or disease diagnosis). In the view of Steels (1987) this trend follows from the drive by expert system shell vendors to make as much money as quickly as possible. These programming aids are being vigorously marketed with claims that they are good for every type of development endeavor. Steels urged the AI community to fight the trivialization by focusing on sophisticated problems involving deep knowledge.

Steels' pessimistic view is not universally held. Others like Hewett & Sasson (1986) conclude that this proliferation of low level applications is a natural effect deriving from the transfer of ideas and technology from academia to industry. Programming personnel and data processing managers in industry are not commonly familiar with AI concepts or the purported range of uses to which AI techniques could be applied. They are likely to see the new technology as a new means for addressing old fashioned problems. As a result, algorithmic tasks are being attacked with high level programming environments initially designed and marketed for knowledge-based applications.

# PROJECT MANAGEMENT METHODOLOGIES

Another facilitating effect of high level programming aids is conservation of time in building any system. High level programming environments make system development easier – easy enough, in fact, that initial prototypes can be constructed without a massive investment of resources or a detailed specification. By testing and refining such prototypes, complex systems can be built in a progressive, evolutionary manner.

This sort of exploratory programming is unknown in traditional data processing circles. Over the last twenty years there have been attempts to delineate a proper sequence of design and construction steps leading to a software product. This work has resulted in a large body of literature on structured programming techniques. Such structured methods grew out of the milieu of the 1960's, when computer power was diminutive and computer time was expensive. These conventional methodologies emphasize a linear process, working from a definitive specification through modular construction to formal demonstrations of the finished system's correctness. Adherence to such practices affords developers a measure of discipline in managing complex projects.

The clearest comparison of exploratory programming with traditional structured approaches can be found in the work of Partridge (1986) and Partridge & Wilks (1986). The range of development regimens is divided into:

> SPIV – Specify / Prove / Implement / Verify paradigm. This corresponds to the linear stepwise process described in the literature on structured conventional programming. The key requirement for this SPIV paradigm is a definitive specification for the task. Such a specification controls the design agenda and provides the benchmark for final validation.
>
> RUDE – Run / Understand / Debug / Edit paradigm. This approach, common in AI, relies little on task specifications, because tasks addressed by AI programs are typically ill-defined. According to Partridge (1986, p. 34), it can be roughly defined as "incrementally developing an adequate approximation to some incompletely specified function". Due to the incremental nature of this paradigm, the course of work is portrayed as cyclical or spiral, as opposed to the linear path of the SPIV methods.

Conventional wisdom holds that the RUDE paradigm is simply symptomatic of a young field, and that over time these practices will give way to more structured procedures. At this point, there is no way to decide the issue once and for all. In fact, the alleged SPIV/RUDE distinction may be illusory, with one being a subset of the other. For the time being, expert systems development will necessarily imply some degree of exploratory work. While we await a canonical paradigm, it would be prudent to follow Heng (1987) in applying discipline and structure wherever possible within this evolutionary construction process.

## TRAINING CONSIDERATIONS

To date private sector expert system development has typically been accomplished by small groups of existing data processing personnel with no formal AI training (Sacerdoti, 1987). Descriptions of the knowledge engineering process typically delineate a one-directional flow of knowledge (expertise) from the expert to the end user's computer terminal. However, the constraints and requirements imposed by the enterprise and the individual users can be seen as knowledge (task parameters) flowing in the opposite direction. Knowledge engineers must recognize that a successful expert system can only be created via a two-directional interaction of expertise and task parameters. As Hart (1986) points out, the effective knowledge engineer will resemble the systems analyst – a jack-of-all-trades, conversant with both technicians and managers, capable of seeing data processing projects from the broadest perspective. Such a person would have to be a composite of an organizational analyst, a systems designer, an acute observer, a computer programmer, and an ergonomist. Clearly, this wide range of roles is not likely found in the typical development group.

How, then, can we ensure that knowledge engineers are sufficiently eclectic? It is unclear what sort of training would necessarily prepare someone as a general facilitator of projects. It is not even clear where to find examples of good knowledge engineers. Informally, such people can be identified by previous successes. This is of little help in assembling a knowledge engineering staff, except to the extent that such folks can be found in-house or recruited. In the meantime, the safe approach would be to actively involve representatives from all enterprise components affected by the transfer of knowledge.

There is a need to examine the training curricula for computer scientists and MIS professionals to see if graduates of existing programs are in fact prepared to shoulder the tasks required in constructing useful knowledge-based systems. Assuming there is indeed an insufficiency, we must ascertain whether it can be rectified by supplementing existing CS/MIS curricula. If not, it is possible that interdisciplinary programs will have to be established to produce knowledge engineers possessing the necessary breadth of vision and skills. This training issue may be the problem farthest from correction in terms of time required to effect change.

Let us see if we can pinpoint the deficiencies. Any knowledge engineering project entails certain procedural steps. Some of these are necessary for any software project; for such subtasks there already exist preferred practices. For other facets of the development task, there are few if any consensus guidelines. The following list illustrates some general stages in building an expert system. The X symbols following selected items represent the degree to which today's knowledge engineers are usually trained for carrying out the given development step.

| | | |
|---|---|---|
| (1) | Model the organizational use of the expert system. | |
| (2) | Model the end users of the expert system. | |
| (3) | Identify sources of relevant domain knowledge. | X |
| (4) | Elicit and capture the knowledge. | X |
| (5) | Formalize the knowledge into facts and rules | XX |
| (6) | Represent the formalized knowledge. | XXX |
| (7) | Devise control structures for search and inference. | XXXX |
| (8) | Devise an interface for addressing the end user. | X |
| (9) | Devise a facility for explaining results to users. | X |
| (10) | Design job tasks for end users of the expert system. | |
| (11) | Evaluate finished system performance. | |

Clearly, there is a need to make the knowledge engineer's background more uniformly adequate for the range of tasks included in transmitting knowledge. Note that knowledge engineers are most likely to be adequately trained in software engineering areas. This reflects the fact that knowledge engineers are usually drawn from the ranks of computer scientists and data processing professionals. The likely deficiencies are in the areas of knowledge acquisition, interface design, and system testing. These activities require developers to address humans as well as computers. This immediately suggests that people trained only in the nuances of computing are ill-equipped for these tasks. Let us examine these problematic areas in more detail.

## KNOWLEDGE ACQUISITION

First comes the job of discerning and organizing the expert's knowledge. This phase, known as *knowledge acquisition,* is largely left to the creativity of the knowledge engineer The need to elicit knowledge from which to construct a knowledge base has been recognized since the early 1970's. However, it wasn't until recently that the first textbook on knowledge acquisition appeared (Hart, 1986). In the meantime, knowledge elicitation has been commonly done via *ad hoc* methods. Those few instances of structured methodology usually entail adoption of information gathering techniques from the fields of psychology, business management, or statistics. Figure 3 illustrates the method of protocol analysis as applied to capturing the expertise of a woodcarver examining a wooden plate.

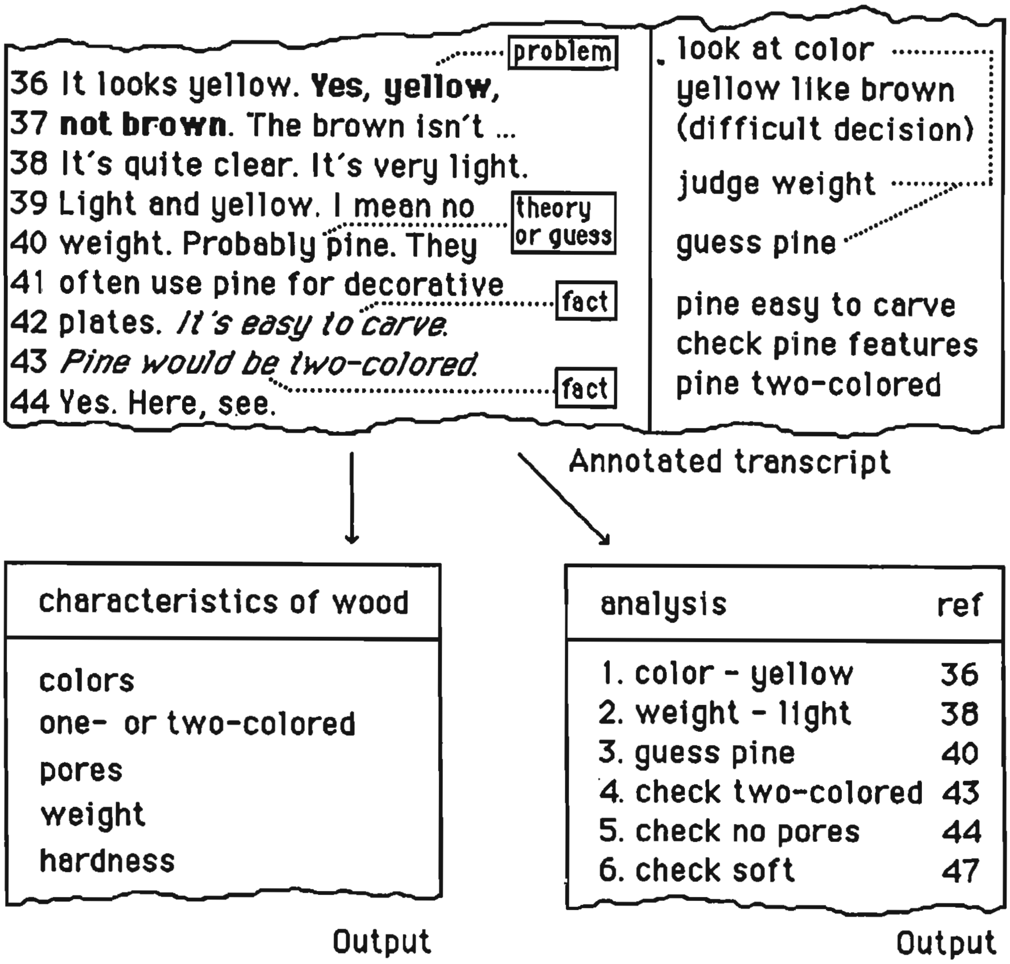


**Figure 3**

*Transcript and protocol analysis of a knowledge capturing session with a woodcarver examining a wooden plate. (SOURCE: Hart, 1986).*

Typical illustrations of knowledge acquisition portray a knowledge engineer studying a single expert. Unfortunately, it is not often the case that only one expert need be consulted to comprehensively outline the task domain knowledge. In the usual case of multiple sources for expertise, there arises the problem of reconciling discrepancies among those sources. Although there exist techniques for shoe-horning multiple views into one consensus model (Thesen *et al.*, 1987), their applicability must be determined on an *ad hoc* basis. The knowledge engineer is left to devise his/her own version of the composite expertise, thus coloring the resultant knowledge base with his/her views of the domain. Unless the knowledge engineer is personally familiar with the task domain, this introduces the potential for significant error.

In practice, the roles of the knowledge elicitor and the expert blur during the process of building a knowledge base. To model the acumen of his/her subject, the knowledge engineer must be conversant with the concepts and terminology of the task domain. He/she will to some extent become a "pseudo" domain expert. An illustration of this effect would be ACE, a well-known expert system for telecommunications fault diagnosis. ACE was designed by a knowledge engineer who was trained to proficiency in the application domain. Similarly, the domain expert may familiarize him/herself with the format of the representation or even participate actively in the construction of the knowledge base. In such a situation, the expert would become a "pseudo" knowledge engineer. Extensions to MYCIN were intended to develop tools to, in effect, aid a physician in developing knowledge engineering proficiency ( Davis, 1979).

The knowledge acquisition process requires skills in interviewing, data gathering,and data analysis. None of these skills are commonly taught to computer science students. It should come as no surprise that current practitioners are often hampered by lack of preparation and a corresponding ignorance of pertinent techniques from other fields. Although some would argue that automated knowledge acquisition is just around the corner, the prudent solution would be to train the prospective expert system builder instead of waiting for new technology.

# INTERFACE ISSUES

Human/computer interactions are typified by constrained syntax and minimal reference to tacit models. Since the computer cannot act as a human, the human must act like a computer for any productive communication to occur. This applies to programs in general and expert systems in particular:

> *"Knowledge based systems built so far share with their knowledge-free predecessors an intolerant rigidity of stylistic expression, vocabulary, and concepts. They rarely accept synonyms and pronouns, never metaphors, and only acknowledge users willing to wear a rigid grammatical strait-jacket."*
>
> (Lenat & Feigenbaum, 1987)

There has been widespread criticism of the user interfaces found on current expert systems, even though up to 50% of programming effort expended to date has been directed to user modelling and interface design (Bobrow *et al.*, 1986). Fisher and Stevens (1987) concluded that lack of communication capacity on the part of the expert system is the main reason that these systems have not moved beyond the research stage.

## The Role of Models in Interface Design

Natural language is not merely a vehicle for transmitting static information. Humans use language for effect in their social environment, via rhetoric, propaganda, poetry, and the like. The field of *pragmatics* within linguistics is the study of language usage for operational effect. In AI, a task domain may be so complex that a system must reason out a course of action which will accomplish a goal before enacting any of the steps leading to that end. The topic of *planning* covers such predictive operations. Whether we speak of a human considering the effects of his/her words on a social setting or a machine evaluating the incremental effects of actions within a task environment, the evaluating agent must have a model to manipulate.

For a system to effectively communicate expertise, it must incorporate knowledge about the end user and how to communicate with him/her as well as knowledge of the application domain (Young, 1984). A relatively simple solution (if feasible) would be to incorporate such user/communication knowledge into the system during programming. This approach (building a *passive* interface) necessitates an adequate conceptualization of the user's needs at the design stage and a willingness to commit to a built-in interface style which will not be easily modifiable. Techniques for passive interface construction are available now.

Unfortunately, no two users are exactly alike. They may vary in perspective, needs, and level of sophistication. To assume that one manner of discourse is appropriate for all consultees is dangerous. Furthermore, initially naive users will develop increased acumen over time. Human experts tailor the content, terminology, and style of their conversations to match the requirements and limitations of the consultee. It would be desirable to provide the expert system with the capacity to dynamically profile individual users and tailor its behavior to each of them. This approach (the *active* interface) is sometimes characterized as providing an expert subsystem whose task domain is managing the user dialogue.

A compelling reason for pursuing interface subsystems that model users is that an individual's performance may vary significantly with characteristics of the information exchange. An illustration would be the experiments reported in Pask and Gregory (1987). Students were categorized as "serialist" or "holist" processors of instructional materials, based on the manner in which they structured new information. Subjects given new materials organized in their respective structuring style retained 80% of the data six weeks after presentation. When serialists used holistically structured materials (and *vice versa*), retention performance dropped to 10% over the same time span. An interface management subsystem capable of classifying individual users as (for example) holists or serialists could tailor information displays thereafter for optimal effect.

Furthermore, a transparency of the expert system in providing the human a model of itself would enhance the human's ability to work in harmony with the machine. Concrete indication of this comes from the experiments of Lehner & Zirk (1987). They studied naive subjects working with a knowledge-based system in solving simple problems. In the optimal performance setting, the subject had been trained with an accurate model of the machine's problem solving procedures. In comparison with this best-case scenario, subjects lacking a good mental model of the system's operation performed between 30 and 60% worse.

An expert system containing such an active interface capacity would be an expert in the task domain and an "amateur" in user psychology and communication. Researchers at IBM concluded that this capacity is needed for usable expert systems (Thomas, 1984). Being expert systems themselves, these subsystems are subject to the limitations of the current technology. Such "intelligent" interfaces are the subject of ongoing research, and practical implementation of such interfaces is probably still far in the future.

Unfortunately, there is a practical limit on the degree to which efforts in this direction can enhance human/computer dialogues. The best that could be attained would be the development of what Branscomb & Thomas (1984) call *artificial personality*. A complete understanding between conversants as dissimilar as a human and a computer will probably never be achieved. After all, it is often difficult for two people to understand each other! This derives in part from the critical role of context in mediating the interchange of information.

## The Importance of Context

Much of the early work in computerized natural language processing concentrated on coding programs embodying a finitely specifiable grammar for utterances. The resulting work was based in large part on the presumption that the full expressiveness of a natural language could be modelled via its grammar, assuming the rules of the grammar were regular and discernible (e.g., Chomsky, 1965). Simply put, syntax was of primary importance in communications among humans.

This emphasis on the form of the communication in conveying a message is apparently borne out by the stringent restrictions on word usage and sentence construction. We learn early on that the meaning of a sentence can be analyzed by isolating its canonical components (i.e., the subject and predicate) and recursively analyzing these into even smaller parts. This structural primacy is not badly strained in accounting for additional effects of non-lexical cues such as tone shifts or gestures used in speech (e.g., in warning a child: **"NO! Don't touch!! It's shaaaarp."**).

Syntax, at least in terms of formal grammar, seems less fundamental in confronting linguistic acts which convey messages differing from their own words. Local traffic authorities may put up a sign saying "REDUCE SPEED TO 30" when their intent is that you slow down to 50. The phrase "This room is dirty", based on its syntax alone, is an idle statement of a condition. It conveys much more when spoken by a guest to a homemaker – maybe a taunting barb or an implied request to do some cleaning.

Communication is therefore not limited to the surface form of a message. The senders and receivers in an exchange rely on more than grammar rules to transfer meaning. Through formal training, socialization, etc., we are all brought up to be semiotic beings (from *semeiotikos*, a Greek expression for "sign/symbol"), assigning arbitrary connotations to symbol tokens. We consistently utilize nonverbal cues, intonation, and a host of implications based on shared understanding of innate and acquired structures which comprise our world models. In short, we actively apply our knowledge in processing information. Pylyshyn (1984, p. 47) summarizes this formally:

> *"...the way cognitive, or representational, processes unfold has a high degree of independence from the organism's causal interactions with the world. They separate the content of the representation from the stimulus conditions by an act of interpretation or encoding, which itself may involve an act of inference ... What regularities follow from the presence of a certain stimulus depend not only on what the stimulus was but on what it was taken to be (on what it was 'seen as'). The latter, in turn, depend on the system's other semantically interpreted structures (beliefs, goals, fears, imagination)."*

All of us have experienced the confusion of processing a message with imperfect reference to the speaker's operant model. Much of comedy relies on this for effect. Communication of even simple concepts may become difficult when we try to communicate with someone from a different culture. Expressions common for us may be nonsensical for their world models, and *vice versa*. This brings us back to the computer – a conversant devoid of the common sense, tacit knowledge, and cultural structures taken for granted in interactions among humans. If you were to key the message "I'M HUNGRY" into the most sophisticated supercomputer in the world, the best reply you could expect would be a hollow "I UNDERSTAND".

Restricted syntax is a constraint in communications between a human and a computer. However, given a sufficiently powerful syntax most information can be transferred intact, even if the form is not elegant. It is troubling that the best interfaces afford something less than the expressive power of natural language. More troubling is the lack of attention paid to the internal models and/or contextual settings involved in human/human and human/computer interactions. These background resources can be used to overcome deficiencies such as noise or gaps in the information flow. They allow some relief from the shortcomings of the syntax. Given a complete absence of these resources, conversants are limited to the expressiveness of the available grammar.

For an illustration of context and the transfer of knowledge via an expert system, refer to Figure 4. In the figure, there is a direct proportional correspondence between the number of symbols in each representation ("thought ballon") and the relative richness of the model. As used in the figure, the domain model contains both explicit knowledge and contextual information. Note that there is a progressive degradation in model richness from the expert to the expert system. This is largely due to the loss of contextual information.

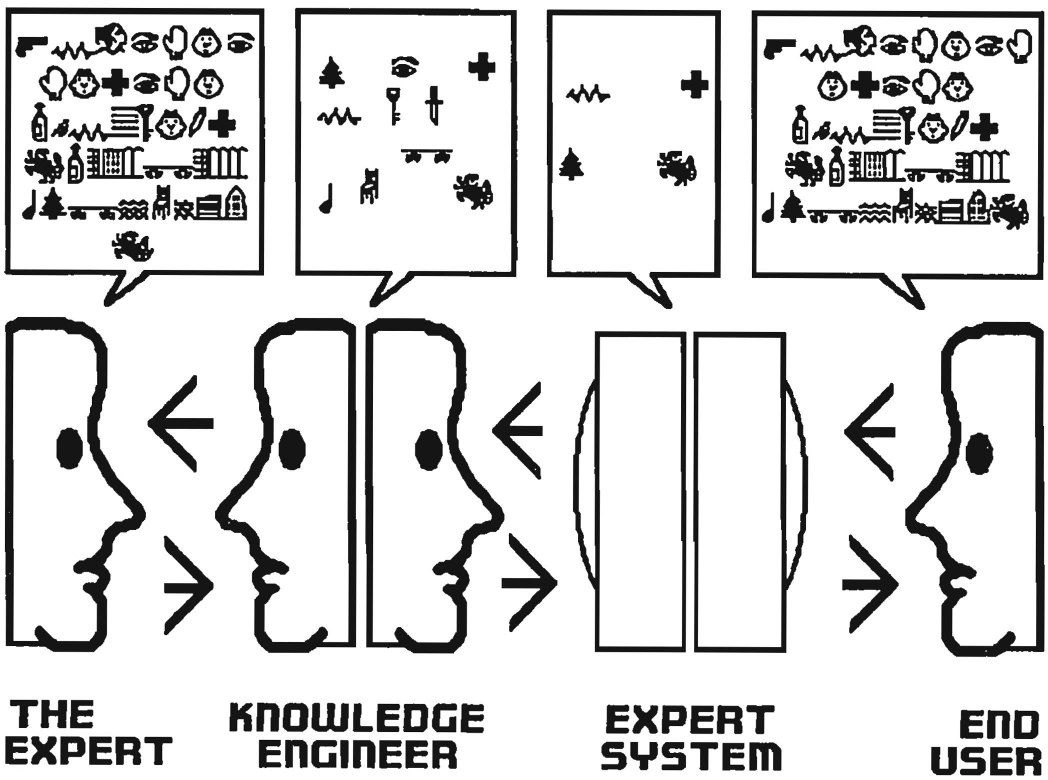


Figure 4

*The contextual flow in the transfer of knowledge. The "thought balloons" above the players are representations of the domain models held by each of the conversants.*

In Figure 4, also note that the user's model is richer than that of the system. This is due to his/her access to situational context. The implications are that (1) some of the potential communication is lost during the knowledge engineering process due to the decrement in contextual representation; (2) the user has a more detailed model of the task environment than the machine by virtue of his/her contextual information; and (3) the user must bridge the gap from his/her full contextual representation to the limited contextual range of the expert system.

In discussing expert systems as a medium for knowledge transfer, Collins *et al.* (1985) devised an analogy to chicken soup with dumplings. The soup itself is the knowledge and the dumplings represent the formalizable facts and rules that the expert can articulate. An expert system acts as a colander – the dumplings get transferred whereas the soup gets lost. The only reason that the end product resembles knowledge is that the users themselves provide so much broth that the dumplings are sufficient to complete the meal. Enhancing the representational power of an expert system may only be marginally helpful, because by analogy it is only reducing the size of the holes in the colander. To guarantee a full meal for all potential "diners", the "cook" must see that soup is provided with the dumplings. To ensure that a knowledge transfer system is useful to naive users, the designer must try to replicate implicit as well as explicit knowledge from the supplanted human/human consultation channel(s).

A more formal description can be derived from the work of Langefors (1966, 1987). He assayed the situation in terms of observational and conclusional pre-khowledge brought to the interaction by the participants. Transforming his "infological equation", we obtain the following facts concerning the communication of knowledge:

(1) Knowledge data is the representation of knowledge.
(2) Knowledge data must be decoded and interpreted to convey knowledge.
(3) Decoding and interpretation is based on pre-knowledge.
(4) Hence, knowledge is an increment to existing pre-knowledge.
(5) Observational pre-knowledge is needed for the interpretation process.
(6) Conclusional pre-knowledge is needed to draw conclusions.
(7) To design knowledge-based systems, knowledge engineers must consider the end users and their pre-knowledge.

This last point emphasizes that knowledge engineers must pay more attention to the individual and organizational contextual parameters of the task setting. To date, the focus has been on explicit knowledge representation; a continuation of this focus will hinder expert systems from reaching their potential. In Langefors' opinion, methodologies for correcting this overly narrow perspective will have to come from social science as well as from computing fields.

## A Modelling Dilemma

A theoretical limitation on the prospects for true dialogue derives from a paradox. The ability to model some agent or process implies a higher degree of complexity relative to the entity being modelled. This follows from the modeller's need to view the modelled from a level of descriptive abstraction higher than the level(s) necessary for simple recognition. The system would ideally be able to model the user. The user would ideally be able to model the system. Each must be of a higher degree of representational complexity to model the other. In effect, each must outmodel the other.

This impasse introduces a dilemma for the knowledge engineer, both in terms of knowledge acquisition and interface design. We term the situation a dilemma because it confronts us with a pair of choices, both of which are awkward. If the machine is capable of more representational complexity than the human (e.g., if the knowledge base surpasses the end user's own understanding), the person is subservient to the system. On the other hand, a situation where the human's complexity outstrips that of the system implies that the knowledge transferred must be constrained.

This modelling dilemma is discussed in relation to new technology, but its theoretical basis is relatively old. A related insight dates back to early research in cybernetics. The term "cybernetics" comes from the Greek word for a helmsman, i.e., the person steering a ship. The "law of requisite variety" (Ashby, 1958), applied to steering, states that it is not possible to steer a system possessing some degree of inherent variety unless the steering agent has access to an even greater degree of variety. The modelling dilemma can be recast as a struggle between two agents (the expert system and the human user), each trying to steer the other.

A knowledge engineer can not be expected to model the expert's knowledge except to the extent that his/her conceptual abilities can handle the complexity therein. An intelligent interface can not effectively model an end user save to the extent that it can portray the complexities of human behavior. In these cases, the modelling dilemma constrains a representation supporting the communication of knowledge. The dilemma also affects the act of communication itself. The end user can not employ the expert system as a tool except to the degree that he/she can steer it – taking initiatives in a consultation, controlling the dialogue, and deciding the course of the problem solving process.

The modelling dilemma induces an absolute limit on the prospects of any participant in the knowledge transfer process ever steering the person or machine with which interaction occurs. Luckily, while the dilemma precludes complete modelling on the part of both developer, system, and user, it does not preclude some modelling ability sufficient to the task. Research in this area is therefore not wasted, so long as recognition of this limitation tempers expectations. Given this caution, let us proceed to consider the modalities of interaction between consultees and expert systems.

## Modes of Interaction and Their Effects

Whether it be fixed (passive) or dynamic (active), the user model implicit in an expert system determines the range of possible consultative dialogue modes. Users can be assigned a variety of roles in the consultation process, with a given role being dependent on the modality for that consultation. There may be one or multiple modes relevant to a given task, with multiple modes being assigned to different classes of users (e.g., novices and adepts). One end result is to influence the ability of the user to effectively use the system. This effect is discussed here. Another effect is to prescribe the relative authority of the user with respect to the system. This will be addressed later, in Part II of the report.

Choice of interactional modality affects the usability of the system by imposing a priori constraints on the end users. Consider the following example of multiple modalities and a conflict thereof. It comes from the MYCIN project. A common desideratum for any expert system is an explanation facility. By invoking explanation, the user is entitled to learn why the system came to the conclusion that it did. It has been commonly suggested that an explanation facility could provide new users with detailed step-by-step training instances. The reasoning is that explanations of conclusions might, if presented to trainees, serve as adequate education in the problem domain.

William Clancey set out to utilize a combination of MYCIN's extensive rule base and explanation facility as a vehicle for tutoring medical students. The result was the system GUIDON (Clancey, 1979). GUIDON was not an effective teaching tool, because the contextual premises underlying explanations were not the same as those required for instruction. The explanation facilities were not general enough to support two different interactional modalities. In this particular case, the root cause of the conflict extended deeper than the interface. The MYCIN knowledge base, constructed for the use of clinicians, did not reflect the underlying principles of the diagnostic domain (e.g., basic assumptions or strategies) that needed to be taught to the students (Clancey, 1984).

At this point in time, expert system developers must exercise prudence and creativity in fashioning interfaces. Interface design is a new field, still in the research stage. It incorporates elements of computer science, ergonomics, and cognitive psychology. Due to its interdisciplinary nature and recency of demand, there is little structure to the design process. While this may explain the resultant problems noted above, it is no excuse for continuing to provide users with difficult, conceptually opaque interfaces. Knowledge engineers are deficient in skills for modelling tasks and users. Recognizing this deficiency, Gaines & Shaw (1986) have suggested that a cadre of dialogue engineers should be trained to augment existing hardware and software builders. Whether such a training initiative would help or hurt the situation is a matter of debate.

# SYSTEMS TESTING: ASSURING SUFFICIENCY PRIOR TO INSTALLATION

## The Most Stringent Approaches: Verification and Validation

Once an expert system has been constructed, it hopefully fulfills its intended function. At some point prior to installing the product in the workplace, it would be nice to know whether it was a useful tool or merely "...100,000 lines in Pascal code for your million dollar piece of diagnostic equipment which finds 30% of the faults..." (DeKleer, 1984). This can be verified only by testing the system. Testing methodologies have been the subject of much enquiry among software engineers (an excellent introductory text would be Myers, 1979). Unfortunately, all this study has only elevated testing to the status of an art. The goals of testing can be differentiated into *validation* and *verification*. Although these two terms are often used interchangeably, they carry unique connotations:

> *"Simply stated, validation refers to building the right system (that is, substantiating that a system performs with an acceptable level of accuracy), whereas verification refers to building the system 'right' (that is, substantiating that a system correctly implements its specifications)."*
>
> (O'Keefe *et al.*, 1987, p. 82).

For algorithmic problems, the idealized evaluation process is verification – the demonstration by formal proof that the program is both correct (i.e., it produces the right output in response to a given input state) and complete (i.e., it always produces an output in response to an appropriate input state). This sort of rigorous prescription is advanced by proponents of strict, disciplined development methodologies (e.g., Dijkstra, 1976). Unfortunately, the complexity of algorithmic programs containing even a few tens of lines of code may render formal proof procedures impossible due to intractable combinatorics. To date, such formal proofs have been successfully employed in verification of few programs of relatively small size (Parnas, 1985).

Expert systems are commonly applied to problem domains for which no algorithmic solutions exist, compounding the testing difficulties. Furthermore, the use of exploratory programming practices often means that there is no rigorous specification to which the finished product can be compared. Rather than the program being derived from a specification, the program is itself the specification of a (hopefully) adequate solution. This unification of problem specification and program code is both the promise and the curse of the declarative programming paradigm. In the absence of a rigorous specification, formal verification is a futile effort.

Validation is therefore the feasible evaluation practice for such systems. Unfortunately, this feasibility is obtained at the cost of rigor. The goal of system sufficiency, rather than system correctness, implies a relaxation of stringency to the point of accepting *ad hoc* evaluation procedures applied to loosely defined parameters. Sufficiency is not likely to mean the same thing to both developers, managers, and users. Without formal task specifications and a consensus perspective, adequacy becomes more a matter of negotiation than a matter of quantification.

Presumably, the best method of evaluating sufficiency would be to gauge the degree to which the codified knowledge accounted for the circumstances under which the system would be consulted. However, there are problems in determining such a measure of comprehensivity. For example, the size of the knowledge base is not a useful statistic. Twelve hundred rules were not enough to ensure the viability of DELTA/CATS, while the Helena Laboratories electrophoresis analyzer works with a set of 82 rules (Weiss & Kulikowski, 1984). Similarly, the complexity of the knowledge base (e.g., the degree of interaction among rules) is no indicator of a faithful domain model. Clearly, the comprehensivity of the knowledge base can not be measured without reference to the task domain context. This suggests that validation testing must incorporate trials.

Care must be exercised in trying to demonstrate sufficiency via test trials. Selection of appropriate test cases by random sampling induces problems which may require the services of a statistician. Cases used for validation should not include cases utilized in designing the system, and this reduces the number of cases available for trial use. In highly specialized domains where the modelled expertise is typically employed for idiosyncratic situations, there may not be enough cases for validation (O'Keefe *et al.*, 1987). For some applications, such as the Strategic Defense Initiative (Parnas, 1985), there may be no cases to evaluate, forcing the developers of the expert system to develop a simulation testbed for their product. Finally, applying the system to an adequate population of test cases may prove prohibitively expensive.

Even if the number of trials is deemed adequate, the proportion of situations correctly handled by the system may fail to reflect the actual comprehensivity of the knowledge base. Consider the light bulb fixer example discussed earlier. The rule "GIVEN no light, THEN replace the bulb and put the switch in the ON position" is a one-rule advisor capable of solving 99.99% of all such problems. Clearly this does not indicate that the system is only 0.01% away from being as intelligent as a human light bulb fixer. The conditions for the rule are difficult to evaluate in the case of a blind consultee. The consequents of the rule are difficult to implement where electrical equipment is not available.

## The Most General Approach: The Turing Test

Modern computing owes much to the British mathematician Alan Turing. He outlined the general purpose framework for symbolic manipulation – the Turing machine. He foretold a day when computers would be capable of displaying intelligent behavior indistinguishable from that of humans. To ascertain when that day had come, he proposed a test for machine intelligence.

In this test, an interrogator would conduct a question and answer dialogue directed to (1) a human and/or (2) a computer. The interrogator would not know whether his/her interactions were with the person or the machine, as the interactions would be periodically rerouted during the course of the trial. If the interrogator judged there to be no difference between the human and machine dialogues, the computer would have demonstrated intelligence. This trial format has come to be known as the Turing test (Turing, 1950).

The Turing test has become the textbook benchmark for any general purpose AI program. So far, no such program has passed the test. Let us examine the Turing test within the narrower perspective of knowledge-based systems. Given an expert system, the Turing trial would consist of dialogues arbitrarily alternating between a consultee and either the system or a human expert. Evaluation would consist of judging the degree to which conversation with the contrivance functionally duplicated conversation with the human expert. In consulting a human expert, one would expect to be accessing a reference source competent in the subject area and capable of effective dialogue. To mimic a human consultant, an expert system would therefore have to display:

(1) **Authoritativeness**, i.e., competence and reliability in addressing caseswithin the domain of discourse; and
(2) **Affability** with respect to offering easy and useful communications, presumably in a conversational fashion. The basis of the Turing trial is the degree of correspondence between the behavior of the machine and that of a human. To pass the trial, a system would need to be "socialized", i.e., taught good manners, conversational arts, and social skills.

The relevance of the Turing test criteria to design decisions is straightforward where the developers cast the expert system as a communications medium rather than a structured object. In this perspective, there already exist one or more (possibly null) channels of consultative communications between the population of users and the population of knowledge references. The goal of an expert system is to partially or wholly supplant these with a new computer-based channel (Whitaker & Östberg, 1988). Viewing the expert system as a rerouted dialogue affords us both a specific model for prototyping the system (communications along the prior channels) and a useful framework for directly evaluating the resulting product on the criteria of authoritativeness and affability via specific comparisons to the characteristics of the supplanted channel.

To date, the evaluation procedure most akin to the Turing test would be the third evaluation study of MYCIN (Yu *et al.*, 1984). This study had human specialists judge the correctness of therapies prescribed by both MYCIN and other clinicians. The study was blinded, i.e., the evaluators did not know which prescriptions were from the machine and which were from humans. In effect, this Turing-like test only addressed authoritativeness, since the evaluators never worked directly with the expert system. Given the limitations of current interfaces, there is little reason to believe that the affability criterion will be satisfied in the foreseeable future.

## The State of Testing Today

The state of the art at this time is little more than platitudes and a plea for prudence. It is widely recognized that much work remains to be done in this area. With regard to evaluation by analysis of the system itself, well-designed field trials are probably the most practical approach to testing usability of expert systems at this time. Such trials have been successful in giving valuable feedback to developers in finalizing their system (e.g., Cochran & Hutchins, 1987). Where the system replicates or supplants access to a human expert, the general framework of the Turing test may help in designing an evaluation protocol. In both cases, the quality of the results relies upon the vision and creativity of the testers. As for the future:

> *"...Expert systems developers need a* ***prescriptive*** *methodology; that is, one explaining how to validate expert systems under certain conditions...and under certain constraints. At present, expert system validation experience is limited. A methodology, or methodologies, will evolve only in the light of future collective experience and critical appraisal of that experience."*
>
> (O'Keefe *et al.*,1987, p. 88)

In any event, it is questionable whether any number of satisfactory trials translates into meaningful validation. The MYCIN developers performed three evaluation studies in an attempt to assess the accuracy of their system versus human physicians (Shortliffe, 1974; Yu *et al.*, 1979; Yu *et al.*, 1984). The National Library of Medicine's AI/RHEUM has been tested on over a thousand clinical cases, and testing is still incomplete (Kingsland, 1987). The result of all this testing is inconclusive – there remain questions of MYCIN's adequacy, and AI/RHEUM has yet to be fielded.

There is reason to believe that no amount of testing will be enough. If the system is being created via an ongoing, evolutionary process (the RUDE paradigm described previously), there will be no invariant benchmark against which to gauge its adequacy. Furthermore, each iteration of the run/ understand/debug/edit cycle requires additional validation, so as to ensure that the incremental changes are correct. In other words, the system will require validation testing for at least as long as it requires exploratory development.

## KNOWLEDGE BASE MAINTENANCE: ASSURING SUFFICIENCY AFTER INSTALLATION

Once installation has occurred, a general lack of complaints does not imply that the system is sufficient to the task. Users may rely on their own experience rather than the system's counsel. Recall the example of Honeywell's COOKER system, considered beneficial even though 50% of its advice was ignored (Cochran & Christopherson, 1987). Even more striking is the example of the Helena Laboratories electrophoresis analyzer. Users apparently do not rely on the system's suggestions at all (Richmond & Landers, 1987). Infrequent use of an explanation facility is no sure indication of a well determined human/computer symbiosis. It could just as well indicate a mismatching of knowledge or models between the machine and the user (Miyake & Norman, 1979; Whitaker & Östberg, 1988).

Expert systems are not products that can be constructed, debugged, delivered, and forgotten. Domain knowledge is not static, and neither are user or organizational needs. Even if the rules in a knowledge base are sufficient to handle the task, there may be need to modify the priority criteria by which the rules are considered. Consider the situation in more conventional applications of Management Information Systems (MIS). A manager's knowledge base is estimated to change by some 20% annually (Trappl, 1986). Given a one-to-one mapping onto a set of rules, this would imply an annual overhaul of up to 20% of the rules. Maintenance is a perpetual process of redesign and refinement; it is one of the main reasons that MIS is often translated as meaning (throwing) "Millions In the Sea".

Knowledge base maintenance entails all the problems of the original knowledge acquisition, plus some new ones. Much of the information which indicates where changes are needed is hard to collect because it is often propagated through informal channels such as social contacts in a variety of settings (Rasmussen, 1987). Integration of new (as opposed to original) knowledge is made difficult by the very existence of the previous knowledge. Any new items must be checked for interactions and/or conflicts with existing ones. Conflicts must be resolved to preserve the accuracy and integrity of the knowledge base. As the knowledge base expands, there is an increasing cost for assuring this integrity. Incorporation of new knowledge will cost ever more per increment as time goes on.

In conventional data processing applications, it is a rare program that is used for an extended length of time without modification. For complex systems, the design and operations are no longer discernible as separate activities decoupled by a commissioning test period (Rasmussen, 1986). Changes in the system require adherence to the structures and biases introduced by the designers. This will likely be more constraining in expert systems, because the knowledge base is in fact the knowledge engineer's interpretation of the expert's acumen. Preserving the efficacy of the system will require maintaining the original model biases. Even if the users themselves are able to debug, update, and improve the system (an extremely unlikely prospect at this time), the knowledge engineers' invisible intentions and conceptualizations will persist. Apparently DELTA/CATS was not an isolated example of dependency. The umbilical cord between the development lab and the application site (as illustrated in Figure 5) may never be cut.

# THE DIFFUSION OF KNOWLEDGE-BASED SYSTEMS

Earlier, we discussed that expert system shells were being utilized for applications which were largely algorithmic, and that expert systems were being applied to increasingly modest domains. Those points were made with reference to whole systems. In this section, we address another dimension to expert systems' proliferation – one which we will call *diffusion*.

By "diffusion" we mean the spread of knowledge-based technology into subsystems within other, larger application programs. By incorporating expert subsystems, such application programs can provide enhanced functionality to their users. By "application programs" we include such low-level software as the microcode in instrumentation, e.g., the Helena Laboratories' electrophoresis analyzer. Although such subsystems are knowledge-based programs in their own right, they are not addressable as independent units. They provide an aura of intelligence for their applications, and that aura may be the only visible trace of their existence.

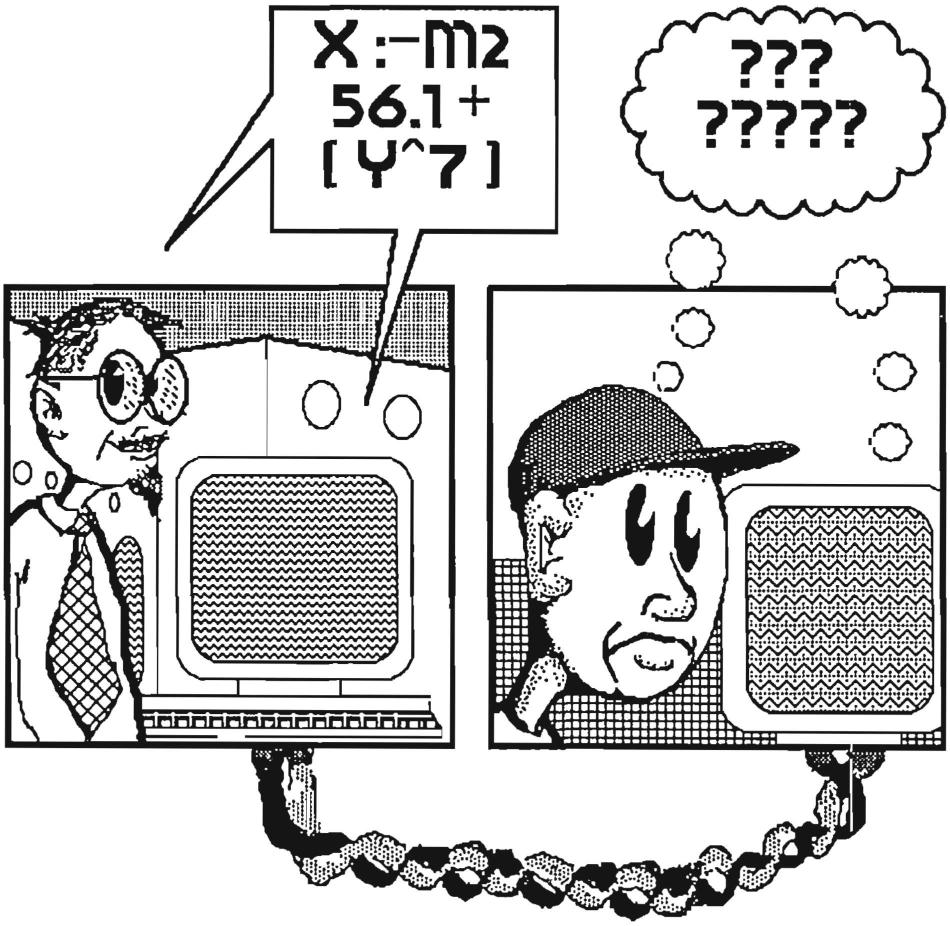


**Figure 5**

*The "umbilical cord" between the expert system development laboratory and the application site can never be cut completely for non-trivial applications.*

To date, the major effort to incorporate knowledge-based technology with an older computing specialty has been in the area of databases. The goal is to provide database management applications with the capacity for inference. An entire new class of systems – expert database systems – has been suggested (Missikoff & Wiederhold, 1986; Smith, 1986). Such a merger allows a system to give users both inferentially-guided retrieval of information and maximum efficiency. The former is arguably the essence of expert systems, and the latter is a benefit of a well-organized database.

There is as yet no good concise definition of an expert database system. This derives from the lack of a good definition for expert systems, and from a lack of consensus regarding why and how to incorporate inferential capacities into existing database models. The preponderance of the literature has been written by database specialists who see knowledge-based technology as another feature to ornament their products. Like current database management systems, expert database systems would serve as development frameworks for constructing workplace applications. This implies knowledge engineering (of the expert subsystems) being absorbed by the existing processes of systems analysis and database design.

If the diffusion of knowledge-based capabilities occurs throughout a wide range of computing applications, there is every reason that it will diffuse into itself. Small knowledge-based components will augment expert systems themselves. Such augmentation might appear as the incorporation of expert subsystems whose domains are knowledge structuring/acquisition, communications, and maintenance. This view is illustrated in Figure 6.

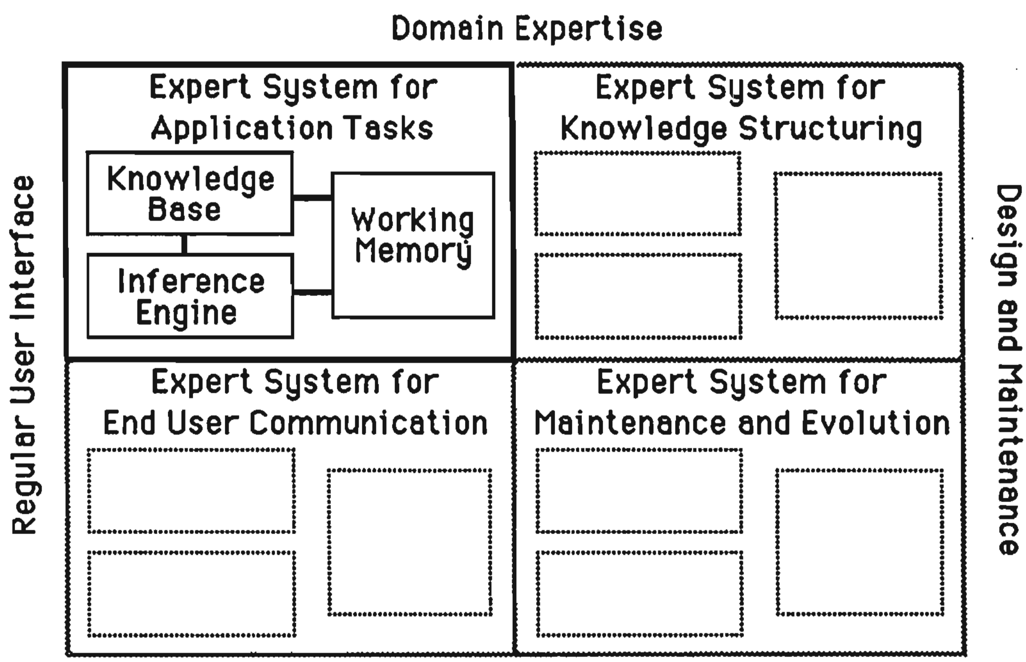


**Figure 6**

*An expert system augmented by means of (modular) expert subsystems.*

The view of Figure 6 is that of extreme modularity. If such modularity were successful, we would hope for the day when the subsystems would be interchangeable across applications. Only the knowledge base for the specific task domain would need to be built from scratch. While this ideal is still in the future, work toward its realization is already underway. Some limited tools for knowledge structuring (i.e., knowledge acquisition) are available now. Mechanisms for adapting the user interface dynamically (i.e., the intelligent interface paradigm discussed earlier) are the subject of ongoing research. No work is yet reported on facilities available for updating knowledge bases *per se*, but much of the research in machine learning has addressed learning in rule bases analogous to those employed in expert systems.

The idea of diffusion suggests that both the power and the problems of knowledge-based programs will be distributed across a range of application systems. We are only beginning to understand the limitations of this technology; in the meantime, it is being absorbed into all types of computing endeavors. This diffusion may indicate that knowledge-based technology will not remain a separate type of programming enterprise – a black art for the production of standalone expert systems. Instead, it may simply come to be considered a component of what we currently call software engineering.

What does this imply for the future of expert systems themselves? Perhaps the examples to date – the large, complex standalone advisory programs – will turn out to be exceptions to the eventual rule. Certainly, high performance applications will require big, autonomous expert systems. Such specialized units will be formidable in power, but few in number. For the most part, the future population of expert systems may well be hidden, the only clue to their presence being the "intelligence" exhibited by the applications containing them.

This diffusion trend is a strong motivation for analyzing expert system impacts before the technology becomes indigenous to most computing. Any human factors problems in today's whole systems will become more subtle in tomorrow's subsystems. Once these subsystems are in place, their adverse effects and biases (if any) will be more difficult to discern and eradicate.

# A FRAMEWORK FOR APPROACHING EXPERT SYSTEMS ISSUES

Assessing the efficacy of expert systems in the workplace has not been easy. Well-designed systems such as DELTA/CATS were doomed once they left the development laboratory. Widely marketed systems such as the Helena Laboratories' electrophoresis analyzer were apparently not having the intended impact in terms of sales appeal or user reliance. Systems were considered successful even when they did not live up to expectations. Within the same publication, one article stated "expert systems...represent the largest segment of the technology called Artificial Intelligence" (Dallas, 1987, p. 17), while another claimed "expert systems do not...control most of the money spent in AI" and set expert systems' share of the AI market at 13% (DM Data, 1987, p. 12). Amidst this chaos we set out to assay the impact of expert systems from the perspective of human factors.

A key step toward this goal is to define our subject matter. Conventional data processing is largely concerned with the construction of computer systems for the algorithmic, quantitative manipulation of domain information units. Knowledge-based processing is concerned with non-algorithmic, symbolic operations performed over the set of relations holding among domain information units. The basic concepts and techniques supporting knowledge-based technology derive from research in artificial intelligence (AI) – research directed toward replicating facets of human intelligence on a computer. The specific objects of our enquiry were those applications where such techniques are employed to model the acumen of a human proficient within some task domain.

The foregoing comments seem vague regarding the exact nature of expert systems. This reflects the unforeseen lack of correspondence between the quantity of information available on such systems and the precision with which the subject matter had been defined. We discovered early in the course of our study that there was no clear intersection of functional user/workplace issues and the structural portrayals of expert systems in the literature. The structure of a system (e.g., its programming language, software architecture, etc.) neither derives from its perceived applicability nor fully controls its final usability within an enterprise. Seeking to evaluate individual and organizational issues within a framework of knowledge bases and inference engines was as pointless as trying to analyze social effects of the automobile from the perspective of gears and manifolds.

We needed a working definition for expert systems. To judge from the literature, expert systems are not easily delineated. Characterizations of expert systems have included:

- a "buzzword for funding" (Bobrow *et al.*, 1986)
- a program for processing of symbolic information
- a program capable of handling uncertain or imprecise information
- a program with separable knowledge base and inference engine
- a program capable of explaining its conclusions when questioned
- a program evolved through rapid prototyping
- a program written using a high-level AI language
- a program designed by means of an expert systems shell
- a program developed by a knowledge engineer
- a program embodying knowledge of a given task domain
- a program capable of performing at a domain expert's level of competence

Most of these features, especially the ones pertaining to the structure of the system, are mentioned nonuniformly. Only the last two features are consistently cited in the literature. They do not rely upon particulars of programming environment, software organization, or development style. These two consensus characteristics pertain to the existence of a body of expressible task domain expertise and the capacity of the finished product to emulate that expertise. Domain knowledge and high-level "expert" performance, at least trivially, underlie all programs. One could argue that compilers and calculators contained knowledge, and that they worked at a level commensurate with an adept human (e.g., Bobrow *et al.*, 1986; Whitaker & Östberg, 1988). We concede this argument, and restrict our use of the label expert system to those programs which model and emulate knowledge equivalent to that which would cause a human to be termed an expert.

To obtain the maximally specific definition consistent with the broadest sample of the literature, we sought an interpretation which emphasized the consensus characteristics rather than structural attributes. This led to a view of the expert system as a conduit or communications channel supporting an effective transfer of knowledge (Whitaker & Östberg, 1988). Figure 7 is a summary illustration of this notion.

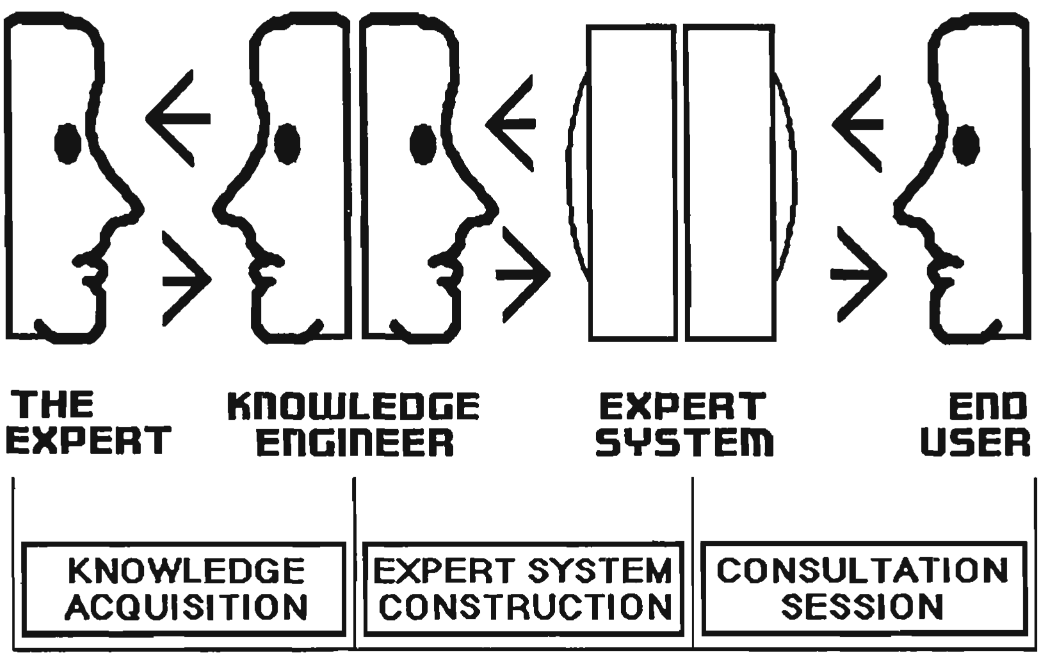


Figure 7

*An expert system viewed as a conduit for expertise from the original source (i.e., the expert) to the end user. The flow from left to right corresponds to the usual course of expert system development. The "players" and the procedural stages are labelled beneath their representations.*

Viewed with reference to the communications channel metaphor, an expert system's function is to reroute access to some body of expertise. The structure of the program is secondary to this process. We no longer need to attempt describe functionality from the perspective of mechanics.

Emphasizing the process of knowledge transfer over the form of the resultant product affords us a conceptual framework for examining expert systems with regard to human factors issues. To the extent that the previous knowledge access path lends itself to such analysis, we can examine its functionality and its usability with regard to both individual users and the organization as a whole. The expert system, seen as a new knowledge access path, can be similarly examined. Evaluation of human factors issues can thus be undertaken by analyzing the new, automated path(s) of knowledge access in comparison with any and all previous paths. Techniques and paradigms for such analyses can be found in fields outside computing, such as industrial engineering, ergonomics, and the like. This affords us access to existing methodologies and conceptualizations.

# PART II

# THE NATURE OF EXPERTISE IN ORGANIZATIONS

We view expert systems development as a process of channeling knowledge and transferring skills from the expert(s) to a population of users. The content of this communication is expertise. To date, knowledge engineers have treated expertise as an object or commodity which can be identified, circumscribed, and translated as a tangible whole. There is a tacit assumption that the expertise modelled up to now has been of one variety, whether it is subjective and contained within one person (e.g., DELTA/CATS) or objectively delineable and shared among a number of people (e.g., R1/XCON). We believe that expertise is not uniform, and that context helps to define its character. In other words, expertise itself is to some extent a function of its environment. There has been little attention given to addressing the nature of expertise as an organizational derivative. To understand the implications of expert systems for organizational life, it is necessary to understand the nature and function of expertise in an enterprise.

A first step in understanding the organizational origins of expertise is to consider how it relates to skilled performance and how it has changed with computerization. Expertise is but one component of the capacity for effective work. Adler (1986a) has identified expertise as one of four dimensions of skill, along with responsibility, interdependence, and training. With computerization there occurs a shift in the focus or style of these dimensions. This shift is directly related to the abstraction induced by computer mediation of tasks (Adler, 1986b).

During this shift, expertise becomes less and less a matter of routine sentient activities. It grows more intellective, involving theoretical apprehension and data-based reasoning (Zuboff, 1985a). It progresses from manual or experiential rote prowess toward more generalized proficiency in identifying and solving problems. This newer version of expertise necessarily incorporates a broader knowledge of the task domain. Adler illustrated this with the example of computerization in a bank. After computerization, the procedures are too complex and errors too costly to expect or desire rote behavior. A teller, now armed with a terminal, is responsible for a larger scope of data transformations than before. This larger scope of responsibility must be supported by a correspondingly increased scope of understanding. Table 1 illustrates the changes in the character of skill in going from an industrial (non-computerized) to a post-industrial (computerized) state.

## Table 1

*Impacts of Computerization on Skill.*
*[Based on Bell, 1973, Adler, 1986a, and Zuboff, 1985a]*

| SKILL COMPONENT | SOCIAL ORGANIZATION OF PRODUCTION | |
|---|---|---|
| | Industrial | Post-Industrial |
| RESPONSIBILITY | For effort (e.g., 5 widgets/hour) | For product (e.g., quality control) |
| EXPERTISE | Sentient, action-dependent (e.g., physical) | Abstract, data-based reasoning with theoretical apprehension (e.g., cognitive) |
| INTERDE-PENDENCE | Tightly coupled Structural/individual (e.g., assembly line) | Loosely coupled Functional/social (e.g., work cells) |
| TRAINING | Once (e.g., tech. school) | Continual (e.g., adaptive training) |

Expert systems offer the potential for augmenting this new form of worker expertise by accurately embodying and effectively managing this broader knowledge. Acting as an electronic aide, the system can support resolution of the problem at hand by accessing and presenting relevant information. In Norway, a bank has employed an expert system as the focal point of a multifunctional financial services workstation. Customers are efficiently provided a full range of services, because the terminal operator has ready access to up-to-date data. This does not obviate the need for the teller to have a broad conception of the banking operation; the system augments that understanding rather than supplanting it. This illustrates what Zuboff (1985b) terms "informating a job" – technology generating useful but previously unavailable information for the individual worker.

This shift toward abstract reasoning as the core of expertise does not mean that the worker becomes merely a processor of some set of ironclad rules. Work environments are dynamic, and an effective worker must be adaptable. Expertise related to such adaptation involves intuition based on experiential or other informal knowledge resources. Salzman (1987) has identified non-formalized decisionmaking and non-rational judgements as elements of this tacit knowledge. In short, we are referring to the worker's common sense. Such common sense acumen may not easily translate into the rational structure of a formal model, even though it may be an essential component of good performance in the given task domain.

Consider an example involving making a choice based on tactical or pragmatic concerns. A decision is required whether to do business with an insurance broker on a particular policy. In a high risk situation, rational judgement may suggest it would be bad business in this case. However, suppose that the insurance broker offering the policy has several other more promising and profitable pieces of insurance. In order to bid on these other items, it may be necessary to take the risk. Resolving this conflict requires an intuitive evaluation based on current market conditions and past experience. The expertise necessary to deal with such a situation derives in part from subjective experiential factors, and it may produce different decisions than objective observation and rational inference.

Expert systems are distinguished from other information systems by their focus on the human decisionmaking/expertise component of skills. As has been noted above, expertise is not an easily packaged commodity. Its nature is a function of the context within which it is recognized. Computerizing an enterprise induces a transformation in the character of expertise. Furthermore, some facets of expertise elude quantification within a neat formal framework.

All these variables would seem to make expert systems completely unlike algorithmic programs. Nonetheless, expert systems resemble their predecessors in that their impacts are contingent on the sophistication of their implementation and the manner in which they are designed to be used by an enterprise. Organizations therefore face design choices with expert systems similar to the choices required for previous information processing applications. It is to these choices that we now turn.

# ORGANIZATIONAL CHOICES AND DESIGN DECISIONS

Enterprises are always faced with choices as to how best to organize and structure production. In the preceding section we introduced the notions of industrial and post-industrial society. In the industrial society, production is organized around concepts of an economy of scale, a unified system, and budgetary controls (e.g., return on investment). Bell (1987) contends that all these concepts are obsolete in the post-industrial society. Economies of scale and the resultant mass production do not work in an era of specialization, just-in-time production, and high value-added products. No one organization can control a system as a unified whole. For example, the telephone system does not make sense in an era of digitalization unless everything is digital. Finally, the notion of return on investment is predicated upon a steady-state economy and the ability to control markets. This scenario is not possible in today's volatile markets, which are dependent on cash flow.

To remain competitive, enterprises have had to choose how to organize their production processes. There are two general modalities for these design choices, corresponding to two historical events – the *control revolution* of the nineteenth century and the *quality revolution* of the late twentieth century. Enterprise resources – physical, monetary, and human – are configured to enhance competitiveness within one of these social modes of production. Information resources are subject to these organizational modalities, too. By extension, this means that expert systems will be designed and implemented with reference to one of these approaches. To understand the impacts that expert systems will have on individuals and enterprises it is necessary to examine these two approaches to the social organization of work.

# THE CONTROL REVOLUTION

The control revolution is identified by Beniger (1986) as having occurred between the latter half of the nineteenth century and the first half of the twentieth century. From the industrial revolution onward, efficient utilization of power sources to drive production systems engendered a period of successful expansion, with attendant infusions of capital into industrial enterprises. The result was an ever more complex system of manufacturers and distributors of goods, each of whom were growing more complex internally. The external complexities made it difficult to control markets, while internal complexities made it difficult to manage the enterprises themselves. There arose a crisis in the ability to integrate and process information for purposes of efficient throughput and strategic planning. This crisis gave birth to the control revolution.

The essence of organizational control is the reduction of uncertainty and variability. At the level of the individual, this is manifested in attempts to restrict personal expression. Such restrictions ostensibly serve to minimize variability in task performance. At the organizational level, this is manifested in establishing and maintaining a coordinated production system. Such coordination serves to minimize uncertainties regarding the enterprise's capacities and operational status. Control at each of these levels is yoked with the other. They may even be the same, depending on the type and sophistication of technology employed for enforcement. Perrow (1986) has identified three types of control in organizations:

> **Direct Control:** This includes giving orders, direct surveillance, enforcing rules and regulations, and imposing technological constraints.
> **Bureaucratic Control:** This includes the specialization and standardization of work activities as well as the establishment of structured hierarchies.
> **Premise Control:** This entails management of the cognitive premises that underlie work actions.

A classic example of direct control in the U.S.A. is the introduction of the assembly line at the Highland Park Ford plant in 1913. Henry Ford reduced labor required to build a Model T from 12 hours and 28 minutes to 1 hour and 33 minutes. This permitted Ford to cut the cost of the car from $ 600 in 1913 to $ 490 in 1914 (Chandler, 1984). In terms of control, the consequence was to constrain the activities of workers and to embed the decision of when to work within the machinery.

Bureaucratic control arose during the late nineteenth century. The managerial bureaucracy was created within capital intensive industries to coordinate the flow of materials linking suppliers, manufacturers, and retailers (Chandler, 1984). These new managers were the integrators of vast specialized enterprises. They gave rise to new levels of hierarchy providing the bureaucratic control necessary for efficient operations.

Premise control involves imbuing workers with cognitive biases congruent with a desired model. An example would be training decisionmakers in the enterprise's preferred conceptualization of the task domain, then reinforcing that perspective over time. This, in effect, entails the enterprise's enforcement of a mindset with regard to the given production process.

Direct and bureaucratic control strategies are most effective when applied to routine tasks. Premise control becomes more important when the work is less routine (Perrow, 1986). The application of expert systems to direct and bureaucratic control is relatively straightforward. Direct control of behavior can be implemented by making a rule base reflect regulations in force. By focussing access to knowledge on the expert systems, bureaucratic boundaries and decisionmaking biases can be enforced.

Premise control, on the other hand, is more subtle, and we found less evidence for its proliferation via expert systems. In one insurance company an underwriting advisory system was being used "to get the underwriter to think like the company wants them to". Management's hope was that individual underwriters will come to think in the intended fashion. The means for effecting this would be training the personnel using the expert system and requiring that the system be employed in the underwriting process. The ultimate goal of premise control is to develop a system of self-regulating individual workers, thereby injecting routine into non-routine tasks.

Routine work is closely tied to the division of labor. Adam Smith (1776) observed the specialization of individual workers in a large-scale enterprise. Smith felt that by dividing production into successive stages and turning each into a separate job assigned to a different worker production efficiency was thereby increased (Attewell, 1984). This *horizontal* specialization is the compartmentalization of personnel at the same job level (in this sense, "level" refers to level of authority). Smith himself aptly illustrated horizontal specialization with the example of splitting the production of pins into 18 subtasks. Later, with the rise of managerial capitalism, the knowledge of production processes was further divided through *vertical* specialization. This imposes an authority hierarchy and control structure such that each job level is strictly separated from the ones above and below. For Smith, the key function of the division of labor was to increase efficiency. This was accomplished by specialization of work – reducing the span of decisionmaking and responsibilities at any one position.

## The Rationalization of Work

It should come as no surprise that early in the industrial revolution there appeared specialists in specialization – analyzers and organizers of production tasks who would pursue ever more optimal allocation of work to humans and machines. In order for bureaucratic control to succeed, work needed to be analyzed and made routine. These specialists (termed production planners, industrial engineers, or methods engineers) emerged as the industrial revolution emigrated from Europe to the United States. Their methodologies were quantitative, relying heavily on the sort of rational reductionism that ruled in the natural sciences.

A forceful proponent of this approach was Frederick W. Taylor (1856–1915). Taylor's work had relevance for both direct and bureaucratic control strategies. Time-and-motion studies and work quotas required direct supervision to ensure a consistency in work performance (Burnes & Fitter, 1987). Such methods also allowed finer standardization and thus enhanced bureaucratic control. Taylor asserted:

> *"The man in the planning room, whose specialty under the scientific management is planning ahead, invariably finds that work can be done better and more economically by subdivision of the labor; each act of each mechanic, for example, should be preceded by various preparatory acts done by other men."*

This quotation comes from *The Principles of Scientific Management*. First published in the United States in 1911, this book and its philosophy spread rapidly back to Europe and all over the world. Taylor's influence has helped mold our current industrial era. Consider the example of the early Ford assembly plants. Using Tayloristic methods and mechanization, planners built an enterprise wherein well-paid workers prôduced inexpensive, high quality automobiles. However, the Taylor legacy has also been negative. An insightful illustration would be the factory setting in Charlie Chaplin's 1935 film *Modern Times* – monotonous, regimented short-cyclic jobs performed at a high pace, inevitably wearing out the worker both mentally and physically. In all too many instances, this caricature is true to life.

The primary methodological tool used in scientific management is methods-time-motion study. The goal of such work study methods is to codify work processes by breaking them up into elementary units (building blocks, work atoms). The purpose of work study analysis is to generate data on the statistical properties of these units and the relationships among them. By manipulating models based on this data and comparing the results against the enterprise's criteria for performance, the most economically attractive way of accomplishing work is sought. Carrying out this analysis in abstract, quantitative terms avoids the overhead necessary to learn from experience alone. Modern work study methods can project the necessary production resources and the cost of the final product even before the first prototype has been manufactured.

In the United States, the most popular management techniques are method study, direct work management, and incentive applications – all derivatives of Taylor's quantitative perspective (Feorene, 1982). Where Taylor's philosophy has persisted, the production life has commonly resembled Chaplin's grim caricature. Industry spokespersons readily acknowledge the problem, but their version of a solution involves continuing the scientific management paradigm with ever more advanced mechanization, such as industrial robots:

> *"In the interest of production efficiency, work has been broken down into series of simple repetitive tasks that can be taught quickly. In fact, much factory work has been reduced to activities that are grossly subhuman. ...Industrialists are mildly interested in shielding workers from hazardous working conditions, but the key motivator is the saving of labor costs by supplanting a human worker with a robot."*
>
> (Engelberger, 1980)

In short, when confronted with highly specialized but unhappy humans, don't humanize the task environment – eliminate the problems by eliminating the people! This is considered a fundamental managerial motivation for automating work activities (Braverman, 1974).

## The Codification of Work

The earliest use of computerized information systems in the workplace was based on the notions of MTM (Methods–Time Measurement). To program a computer required that the work activities be broken down into discrete segments that were then translated into a programming language. In large organizations computers performed back office functions such as billing. To the extent that decisions and/or clerical labor could be transferred from individual humans to a computer, such automation was considered desirable.

Within the framework of the control revolution, a decision to apply a technological fix was typically based on attempts at either greater direct control of workers through embedding production decisions in the computer or consolidating bureaucratic control by letting the computer coordinate information throughput. An example of the former would be computer-controlled assembly lines. The latter is illustrated by automation of functions such as billing. In either case, applicability of the technology required only that the work could be codified into programmable segments.

It is tempting to conclude that work study methods, by codifying work processes, always lay the groundwork for regimenting or replacing humans. In this view, human workers would invariably be disadvantaged by the use of these methods. However, machines are not excluded from such scrutiny. Paul & Nof (1979) were interested in comparing the performance of humans with that of industrial robots. To achieve this, they developed the RTM (Robot Time Measurement) system, which is essentially MTM for automatons.

Codifying a task does not necessarily mean automating that task. However, automation of a work process requires that the task be thoroughly codified so it can be algorithmically replicated. Robots have the advantage of being capable of fine-tuning; properly programmed, they can achieve a precision in their work equal to or greater than that of humans. For some simpler tasks robots can be taught a routine by recording the movements of an adept human and translating them into a control code. Consider the paint spraying robot illustrated in Figure 8. To achieve a minimal competency, the robot need only play back the common movements of an expert spray painter. With further refinement of the motion and spraying parameters, the robot will be able to outperform the expert.

In this example, the act of manual spray painting has been translated into a form the robot can use. Whether the painting skill was abstractly codified for programming or transferred by direct mimickry, the result is the same – the robot becomes a *de facto* expert at the task. The optimistic view is that the machinery has become sophisticated enough to emulate the human. The more pessimistic view is that the robot could become a *humanoid* machine only because the worker had already been reduced to a *machinoid* human (Östberg & Enqvist, 1984). The former view is continually suggested by news of ongoing advances in automation technology, but we feel this view fails to completely describe the changes around us. To some extent the latter view is inevitably borne out. It is with attention to this potential worker trivialization that we proceed.

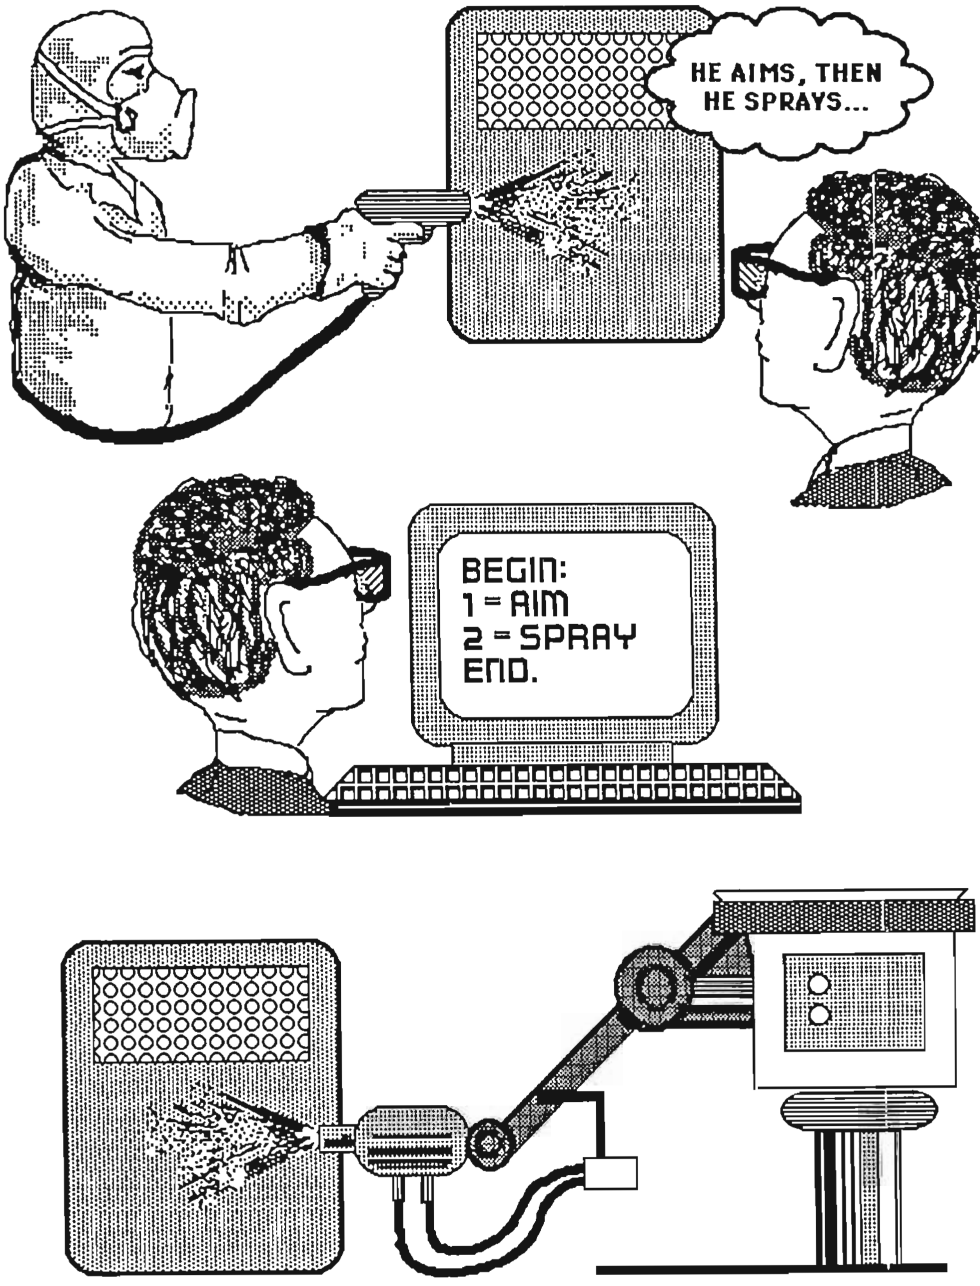


**Figure 8**

*After refinement of the recorded motion and sprying parameters, and codification of the work procedures into the computer memory, a playback robot will be able to outperform the human expert.*

## Expert Systems and Work Codification

Expert systems represent but the latest in production engineers' continuing attempts to control work activities through automation, by codifying and mimicking human skills. The stopwatch in the hands of the work study engineer has been replaced by the computer in the hands of the knowledge engineer. The parallels between these endeavors have been discussed in Jones & Watson (1984). One management guru saw computer automation as a logical extension of Taylor's scientific management (Drucker, 1976). Mowshowitz (1984) projected a clash between this newer style of analysis and the old Taylorism. The development of knowledge-based systems has been likened to Adam Smith's inquiry into the nature and causes of the wealth of nations (Feigenbaum & McCorduck, 1983). After two centuries of industrialization, we are now accustomed to the automation of physical labor. Metal has supplanted muscle, leaving duties of judgement, reasoning, and information processing to people. With the advent of computers, even these mental duties are being reassigned to machines. The influence of computer-based automation is felt at ever higher levels in the authority hierarchy of enterprises.

In the nineteenth century, the machinist's physical duties were assumed by power tools; in the twentieth century he lost more ground to numerically controlled machines such as robots. A highly skilled worker became an operator – a relatively passive participant in working materials. Now even his oversight is being mechanized. In the flexible manufacturing system (FMS) of tomorrow, a robot displaying out-of-tolerance performance will be examined by an operator relying on an expert system configured for robot maintenance (O'Connell, 1987).

Continuing up the hierarchy, we see expectations of expert systems supporting the design and planning of machine activities. A bottleneck in the realization of a fully automated FMS is the need for humans in plotting out a course of process stages. Before assembling the set-ups into a step-by-step plan, a machinist looks over the product design and weighs decisions about (e.g., clamping, fixture design, and special tooling:

> *"This requires the human intelligence and expertise of a highly skilled machinist. As a first step the expert system EXSUS makes it possible to use an operator that need not be a skilled machinist. The next step is the MACHINIST EXPERT, which will give the designer a computer in which information and machining expertise is stored."*
>
> (Hayes & Wright, 1986)

Expert systems are also considered for automating supervisory functions in service industries. Garwood (1984) reports on a stored program control technology in a telecommunications network that will track information on customers and staff. When a customer complains, the program will examine the nature of the complaint and advise the supervisor of the best fit between the customer's problem and capacities of the current staff. Expert system technology can therefore influence organizational control.

An apparent prospectus for the use of automation to extend organizational control can be found in a large scale investigation in Japanese industry. According to this study, the prime benefit of robotics is that its introduction forces management to "shape up" in terms of streamlining production and to encode worker skills to pave the way for full automation (JMA, 1983). Among other things, *full* automation means automation of supervisory functions as well as manual tasks. The following two tacit principles seems to be the generalized development rules:

> (1) For the lowest level task(s) in an enterprise, completely codify the work and transfer as much of this codified work as possible to a machine. Ideally, the task(s) can eventually be left wholly to machines.
> (2) Continue up the authority hierarchy until tasks are found which cannot be completely codified. For these tasks, transfer as much of the work as can be codified to machines. These machines will then operate jointly with humans trained to handle the remaining portion of the enterprise's work.

These deducted principles have much in common with the following advice, taken from Hwang & Salvendy's (1984) discussion on steps leading to the fully automated FMS:

> *"(1) Allocate to the computer more tasks. The supervisory tasks of the operator in an FMS should consist only of simple decisions.*
> *(2) For more complex decisions, an expert system or decision support system should be developed which will work in conjunction with the supervisor."*

At the lowest levels of factory automation, the goal is to supplant humans with autonomous robots. Such robots would incorporate physical manipulation with perceptual facilities (e.g., vision) and decision making abilities (i.e., knowledge-based systems). We would wish for such robots to be controllable with no greater difficulty than that required for overseeing humans; this implies a natural language processing ability. These four functional areas are traditionally associated with research in artificial intelligence. Refer to Figure 9 for an illustration of how these research interests relate to human faculties.

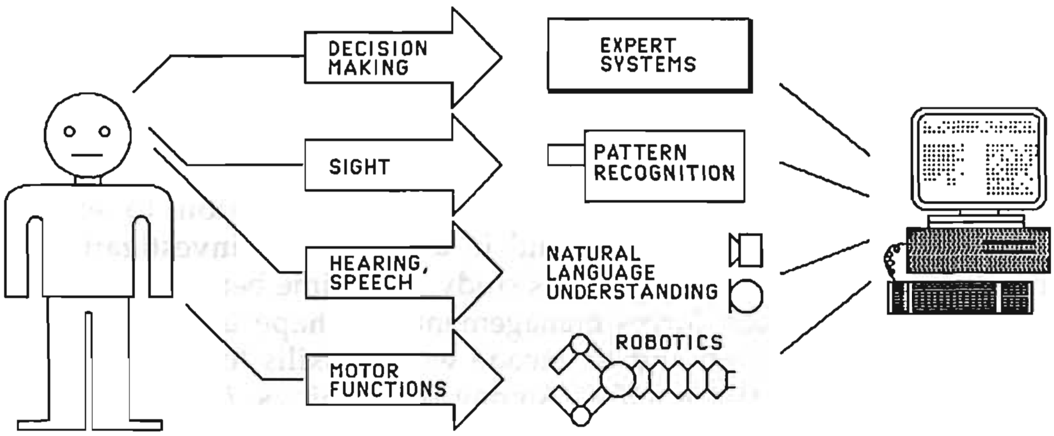


**Figure 9**

*Areas of Artificial Intelligence (AI) research for the codification of human faculties. [After Harvey, 1986].*

At the higher levels, we find attempts to codify ever more sophisticated human skills. It is the attempt to replicate previously non-automated acumen that is often used to justify labeling software as an expert system. This does not mean that tasks already well codified are already embodied in expert systems. Consider, for example, the US Air Forces IMIS project – the development of microcomputer-based expert systems for jet aircraft maintenance. Adequate decision structures for fault diagnosis were constructed (in the form of decision tables) as far back as the 1950's. Attempts to disseminate these decision structures to flight line technicians via printed media or a computerized database had not proved practicable. Now that reasonably powerful portable computers are available, the same basic decision structures are now being translated into a new delivery medium (Miller, 1987). In this case, knowledge has been well modelled for a long time, but the technology for effectively distributing this knowledge has been lacking. The organization has determined a target for control (in this case, diagnostic consistency), but effectuation has been delayed pending an appropriate mechanism.

Furthermore, the stated equivalence of frontiers implies that expert systems may be developed for tasks which heretofore had not been well codified. As such, the knowledge engineering process may represent the first attempt to analyze and formalize the task. This has already occurred. Expert system developers have not been shy about trying their knowledge-based tools on tasks which had not previously been formally modelled. While it would be difficult at this date to draw any conclusions about combining knowledge engineering with initial task codification, it would appear obvious that the probability of producing a useful system is thereby diminished.

## THE QUALITY REVOLUTION

With the emergence of a post-industrial society there have emerged new forms of social organization for production and for goal-directed enterprises. Bell (1973) presaged the onset of these new social structurings in the early 1970's. There is now a quality revolution analogous to the control revolution that began in the late nineteenth century. This revolution is manifested in a shift of emphasis from control of personal actions toward a commitment to high quality work. It is characterized by cultivation of work environments where creativity and innovation can prosper.

Recent media canonization of Japanese quality control and quality circles would lead some to view the quality revolution as having begun with the current crisis in America's competitiveness. This is, of course, not the case. The quality revolution, like the earlier control revolution, has been a gradual development. It is too early to decide whether the current global competition marks a period of crisis whose solution requires the administration of advanced information technology. This determination must be left to the historians of the 21st century.

One of the earliest signs of the quality revolution was the development of sociotechnical principles for the implementation of new technologies. In Western Europe, the cradle of the industrial revolution, hardline scientific management has gradually been abandoned in favor of the "school of sociotechnology". The key innovation was the recognition of the need to match the social organization of work and the technology of the production process to establish a new social organization for production. Take the example of Volvo, formerly an exemplar of industrial engineering and its trappings – time study, productivity programs, and the proliferation of monotonous jobs. This major corporation has now distanced itself from Taylorism and moved into a sociotechnology phase. (Jönsson, 1983). The transition has been so fruitful that recently a senior Vice President of the Swedish Employers' Confederation expressed doubt that any conventional production line had been built in Sweden since Volvo opened its pioneering Kalmar assembly plant in 1974 (Lindholm, 1987).

The tactics of the quality revolution are different from those of the control revolution. Problems of coordination and control within a complex enterprise are resolved through creating flexible organizations capable of effecting change rather than bureaucratic organizations geared toward enforcing consistency. Information technologies are used to support flexible production in meeting demands in highly value-added markets. This differs from the earlier use of such technologies to establish and enforce a constancy of throughputs so as to maximize economies of scale. There are several specific dimensions of the quality revolution which are best understood by contrasting them with the control revolution.

(1) **Organizational emphases on quality in production and innovation in work.** This can be contrasted with the primacy given efficiency and consistency in the control organization. Today's markets are less stable than before; in fact, there is a lack of an oligopoly. The new emphasis is on a high quality product, production innovations, and a willingness to fashion an appropriate, supportive work environment. In such an environment, saturated with information technology, this means that the worker must have a working knowledge of that technology, and he/she must be capable of modifying it as needed.
(2) **The primary role of the human in production processes.** This viewpoint leads to a focus on building work systems that encourage worker commitment, as opposed to a focus on controlling worker behavior. To obtain this commitment, workers must be involved in the development and maintenance of these work systems.
(3) **Technological dynamism.** No technology is static for long. We must recognize that information system developments are not one-shot substitutions for human activity, as they have been utilized in the control organization. Instead, they require continual adaptation to maintain their usefulness. Information technology, especially the software, is a bouncing ball that the enterprise must try to keep in play. This realization has implications for how organizations manage the development and maintenance of information systems.
(4) **The fallacy of the technological fix.** In the control revolution, the principal means for improving work efficiency and achieving control has been the application of more capital investment or more sophisticated technology. This has succeeded because quality was not so critical a factor in marketing products. Given the dynamism of technology and the increasing emphasis on quality, it no longer suffices to substitute technology for organizational solutions.
(5) **The primacy of integration in work systems.** This is opposed to the earlier fragmentation of work systems. As enterprises become more complex and markets more volatile, it is necessary to have highly skilled workers capable of shifting production priorities on short notice. To do this, workers must have a broad conception of the work process, not just a narrow understanding of their individual tasks. At the level of the individual worker, there is ample evidence that the ability to understand the workings of a system leads

to quicker, more accurate solutions to novel or uncertain conditions. At the level of the whole production process, this integration is a necessary consequence of both worker interdependence in production and the need for worker participation in the maintenance of support systems.

(6) **A rejection of the Babbage principle.** Given a highly integrated work system, malfunctions and/or errors carry a high cost. To avoid such problems and their attendant costs, today's information technology demands a population of highly skilled workers. This implies an emphasis on skills development. Such development will require continual, ongoing training as opposed to one-shot training.

(7) **An emphasis on labor/management cooperation.** Today's global markets are as unforgiving and risky as they are competitive. Organizational failure entails heavy losses for both workers and management. This level of risk mandates cooperation instead of antagonism. The need for such cooperation also derives from the need to involve humans in the decisionmaking process.

(8) A shift in human resource practices. This is primarily a heightened emphasis on job security at the organizational level. Work environments dominated by information systems require fewer, more highly skilled workers. There is a tremendous organizational investment necessary to impart skills and maintain adequate training over time. Given this investment, it is essential that organizations keep their workers.

A major distinction between the quality revolution and the earlier control revolution is the need to couple information systems innovations with work structuring innovations to produce a new social organization for production. This necessity has been recognized as a result of organizations' experiences in attempting to maintain a control posture while implementing new information technologies. For example, Walton (1980) describes an application of computer systems to support the machining of bearings. The new equipment was very complex and sensitive, the product had to meet high quality standards, and the business required a high-speed operation. Management underestimated the skills needed by individual workers to achieve the desired production efficiency. These skills were widely distributed among operational, supervisory, and engineering personnel under the historical work organization of the company. The result was the inability of the plant to maintain the production standards necessary for economic viability.

A second major distinction between the control and quality revolutions are the emphases on integration of the individual into the total work system and formulation of systems permitting creativity and innovation. These processes are the result of a desire to "informate" the worker. Zuboff (1985b) provides an example from the pulp and paper industry. Changes in a process control technology afforded workers access to more comprehensive information about the total production process. Using their knowledge of the production process and this enhanced data the workers have been able to identify a variety of opportunities for cutting costs.

## The Codification of Human Work Systems

The codification of work from Frederick Taylor onward has separated the analysis of task activities from analysis of human needs. This has lead to atomization of tasks in preparation for division of labor and automation. Although treated as a secondary concern in analysis, workers are employed as a primary means for bridging gaps in the automated production process. Only those tasks which can not (yet) be automated are left for people to perform. Work study analysis is done with reference to someone's criteria for performance. When the criteria are strictly monetary, this analysis may be called management science, operations research, production planning, or a host of other titles. When the criteria are human values such as job discretion, health, and safety, such studies are usually termed *ergonomics*. Ergonomics is a label often used to portray work study or design practices as being benign, justifiably or not. For example, Lenin introduced scientific management practices into the USSR under the name "ergonomics/ergology". We shall use the term in its most benign sense.

We have consistently painted a negative picture of MTM and other work analysis paradigms. This has been done to illuminate the problems deriving from applying such practices without consideration of human needs. We do not mean to imply these practices are obsolete or unusable. Work study methods are tools which, like any other tools, may be used for bad (e.g., Chaplin's caricature) or good (the ergonomics ideal). In this section we will examine the codification of human work systems where the human is integrated into the analytical process. We will do so from two perspectives. First, we will discuss how human design issues can be addressed in conjunction with previous types of quantitative analyses. This entails being sensitive to human needs in tailoring people to production tasks. Next we will consider the prospects for ergonomic codification practices, wherein the human needs have primacy over quantified production models in determining system design.

## Incorporating Human Needs Into Current Analyses

Methods-Time Measurement (or MTM) has been a central topic in the controversy over Taylor's rationalist paradigm of vertical and horizontal division of labor. Developed by the Gilbreths and Maynard (disciples of Taylor), this procedure analyzes any manual operation into its constituent motions. MTM assigns to each motion a predetermined time standard based on the nature of the movement and the conditions under which it is made. The original MTM system has evolved into several systems so as to (1) fit different types of applications and (2) strike an appropriate balance between analytical effort and precision. For example, the MTM-1 system, applicable for short-cycle assembly tasks, recognizes the following basic motion elements:

| | |
|---|---|
| Reach | Move |
| Turn and apply pressure | Grasp |
| Position | Disengage |
| Eye travel/eye focus | Body motions |

Under this system, a normal competent worker should be allowed 0.00053 hours to move the right hand three inches to a single object in a fixed location. To address the longer cycles found in the operation of numerically controlled machine tools, the MTM-V system was developed. The "V" stands for Verkstad, Swedish for "workshop". For even longer cycle tasks such as mechanical fault diagnosis and correction, there is an MTM system called UMS (or MTM-U). The "U" stands for Underhåll, Swedish for "repair/maintenance".

Such time standardization is not limited to physical labor, as evidenced by KLAR-K (or MTM-K), where "K" stands for Kontor, Swedish for "office". Used in assaying clerical time allocation, KLAR-K, like the basic MTM-1, differentiates tasks into elementary units. Some examples are:

| | |
|---|---|
| Waiting 0.5 sec. for answer from a terminal | File document in locklever binder |
| Proofread line of 11-16 text characters | Compute number with pocket calculator |
| Receive data from IBM 2740 terminal | Mental calculation |

All these systems are in widespread international use, and all are Swedish products. It seems ironic that Sweden is the world leader when it comes to denouncing hardline Taylorism and also the world leader in development of MTM systems. It would appear that those who seek to alleviate the negative aspects of scientific management have adopted analytical methods to the extent that such analyses generate data useful for labor/management negotiations concerning issues of social organization of work (e.g., production schedules).

The latest MTM system, SAM (or MTM-SAM), is currently under review by the International MTM Directorate. The acronym SAM ostensibly derives from Sequential Activity and Method Analysis. In reality, it comes from the Swedish word for "jointly/cooperatively/ in concert" – Samverkan. By insisting on the acronym SAM, the Swedish MTM Association (SRF) wished to emphasize that the system's development had been carried out by MTM specialists from both industrial management and trade unions.

Trade unions have acknowledged the utility of such management tools in production planning and cost analysis. They also see the need for objective analytical data to employ in the bargaining process. Given such data, negotiations can focus on payment and production rates, avoiding distrust on either side as to the validity of the bottom line numbers. With objective standards, both sides can work toward agreement with a clear understanding of what the relevant MTM base rate of 100 represents. The consensus production rate in the Swedish auto industry is 115 (in MTM terms), because the unions were willing to work 15% harder than the base standard. In Swedish banks the rate is 85, mainly because management wanted a staffing surplus of 15% so as to ensure that customers are served without undue waiting.

## Prospects for Ergonomics Codification Practices

Let us now turn to a notion of codification that works from a starting point of human needs rather than production parameters. Extreme Tayloristic approaches place paramount importance on optimization of quantifiable work dimensions, according little or no weight to human factors. The middle road, described in the previous section, is to balance rationalization procedures against workers' concerns. The Swedish example above indicates that this is done by negotiation between those parties emphasizing production optimization (i.e., management) and those emphasizing work life issues (e.g., trade unions). The next step would be to integrate human factors analysis into the codification process. For purposes of this discussion, we will label this integrated process *ergonomics codification.*

An ergonomics codification of human work systems starts with an assessment of human capacities pertinent to the given task. This evaluation is then mapped onto an assessment of the technology's ability to support the work. Design decisions based on this evaluation will be made with the involvement of the workers themselves. With respect to information systems, these assessments would include consideration of human cognitive facilities, the role of the worker in the enterprise's information flow, the forms and functions that the worker demands of the information, and the availability of sufficient computing power. Results would include a job design and a specification for the human/computer interface.

The central contribution of ergonomics codification is to elevate the importance of the worker's viewpoint in addressing the implications of automation. Examples to illustrate such a perspective can be found in analyses of automation and office work. A variety of investigators have attempted to construct a framework for understanding office automation by classification of office tasks. For example, Helander &

Östberg (1983) offered four general task categories: perceptual/motor skills; rule-based decisionmaking; analysis/ problem solving; and social skills. Later work by Sasso *et al.* (1985) and Olson (1987) suggests that by extending this classification to a finer level of cognitive analysis and coupling that analysis with an assessment of human capacities will result in more objective and reasonable design choices.

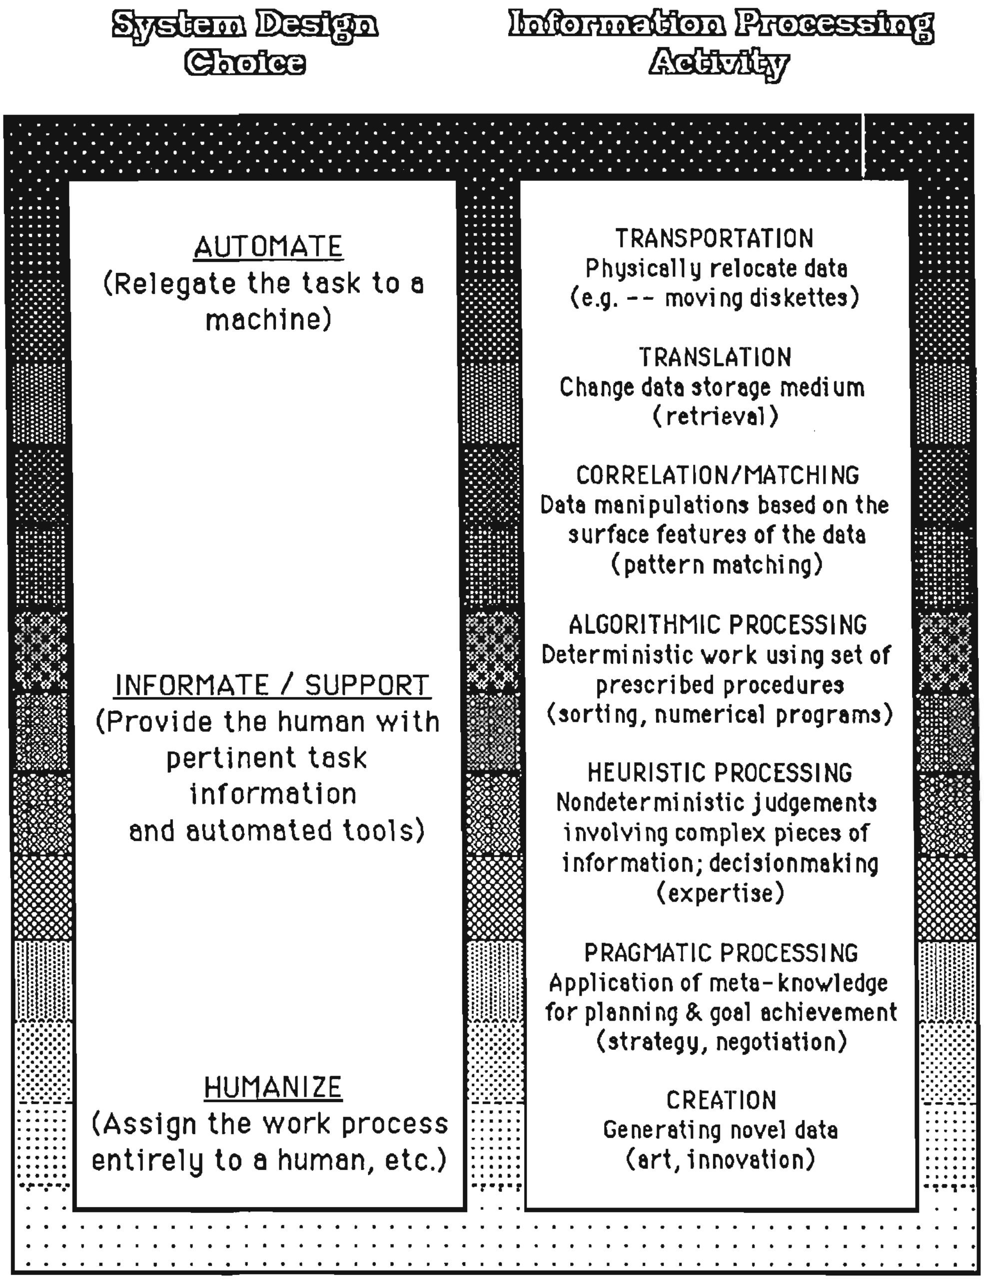


**Figure 10**

*An illustration of the correspondences between types of work system design approaches and human information processing activities used in ergonomics codification practice. [Based on Sasso* et al.*,1985, and Olson, 1987]*

The first step is to identify human information processing strengths and weaknesses (Sasso *et al.*, 1985). Based on this evaluation and a categorization of work activities, designers can proceed to make decisions regarding the form of the information systems and the interfaces by which humans will interact with them. This occurs with attention to two dimensions – (1) the relative strength/weakness of pertinent human facilities and (2) a continuum of design strategies ranging from full automation through means to cooperatively informate/support humans with machinery to full humanization (i.e., assigning the work exclusively to humans). The general decision framework is illustrated in Figure 10. Human cognitive facilities are ranked by increasing strength from top to bottom. The automation/humanization dimension is similarly displayed. Automation naturally maps to the weakest human information processing abilities, while humanization is reserved for those tasks that involve the human's cognitive strengths.

Simply put, the goal is to let the machine do what it can do best, and let the worker do what he/she does best. This requires a mediation between the extremes of automation and humanization based on evaluating the cognitive requisites for each component task in the overall work process. Let us turn to an example taken from Salzman (1987) which illustrates a situation where such mediation has been profitable. A CAD system was used to automate the design of printed circuit boards. An attempt was made to devise automatic circuit routing programs which would lay out the complex pathways of such boards. In one case, the autorouter placed a connection down the middle of the board, precluding ten other connections and making the design unworkable. The program performed correctly with respect to the one connection, yet failed to allow for the interactional constraints among all connections.

Despite this inadequacy, Salzman found that the autorouter programs were useful for board design. Designers interviewed estimated that the autorouter was helpful in placing about 40% of the connections on a given board. The remaining 60% of the connections were specified by the human designers. While the autorouter facilitated the designer, it could not replace him/her. Rather than being junked, the work process was modified to automate the design task to the extent that the machine was proficient. The designer was not supplanted; rather he/she was freed to concentrate on the more skilled aspects of the work.

## Expert Systems and Codification of Human Work Systems

In Part I, we emphasized the ergonomics aspects of expert system design from the perspective of the individual end user addressing the particular system. Here in the second part of the report we are emphasizing the organizational aspects of implementing knowledge-based technology. Ergonomics aspects of the social organization of work must similarly be addressed if the full potential benefits of the technology are to be realized.

Recall that in Part I we pointed out the effects of interactional modalities on the efficacy of a human/computer interface and the usability of a system. Design choices involving interactional modalities can also serve to enforce boundaries of empowerment within an enterprise. The same dialogue mode may be too restrictive (*overdetermined*) for one user and too permissive (*underdetermined*) for another. Where communication is overdetermined, the machine exerts greater control over interactions, inhibiting the user. Where it is underdetermined, control defaults to the user, who may not know what to do. At one extreme, the user would be maximally empowered, with the expert system reduced to the status of a smart tool. With reference to Figure 10, such a smart tool would be configured to informate the user. The Helena Laboratories' electrophoresis interpreter, an expert system residing on a microprocessor is embedded within a laboratory instrument (Weiss & Kulikowski, 1984). It might be considered an appropriate illustration for this position. The expert system provides its advice not as ironclad prescriptions but as polite suggestions. The responsibility for decisionmaking with the human.

At the other extreme is a modality in which the system is maximally empowered at the expense of the user. The human is reduced to serving as a tender of the system, mediating the flow of information between the machine and the task environment. In this case, the decisionmaking responsibilities have been relegated to the machine. Such an approach reflects an attempt to exert control in the work process. However, the effectiveness of such a strategy could be limited, as it was in the earlier CAD example. With reference to Figure 10, this would occur as a result of design decisions to apply full automation to those information processing activities falling in the "informate/support" band. Figure 11 illustrates this scenario, wherein the human is, in effect, a peripheral input/output device.

Presumably, the interactional styles of most systems would fall somewhere in between these extremes. Following Hägglund (1987), we can offer the following possibilities:

> **Transactional Mode:** The user is prompted by the system, and he/she is required to evaluate the machine's solution(s).
> **Consultative Mode:** The user is more active, and he/she specifies the problem to be solved. The machine is therefore utilized as an expert aide.
> **Commentary Mode:** The system reviews a solution or plan, in effect providing a second opinion.
> **Supervisory Mode:** The system is passive, monitoring events in the task domain, and becomes active only to flag inappropriate or problematical decisions.

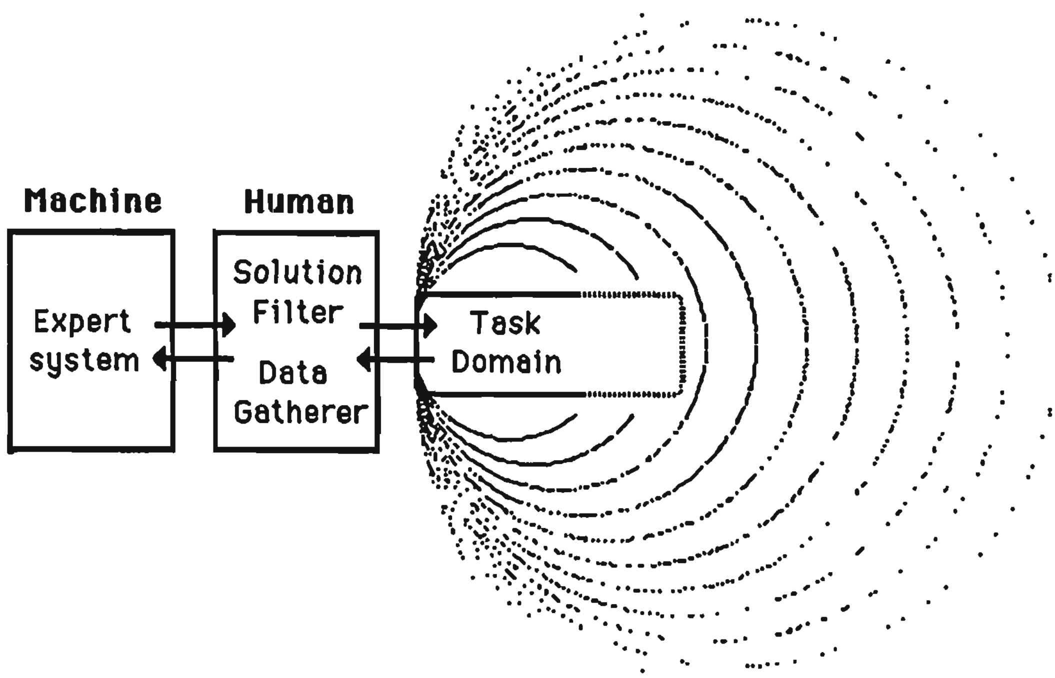


**Figure 11**

*This is the extreme case of machine dependence upon the human user of an expert system. In this situation the human is reduced to merely feeding domain data to the machine and enacting the machine's commands. [Based on Woods, 1986].*

The mode(s) of interaction are work organization issues, properly addressed during the initial system design process. Where the potential range of modalities are limited by the structure of the knowledge base (as in the MYCIN example), failure to resolve these issues may result in a system which is ill-suited to some users, if not wholly unusable. We suggest the applicability of the sociotechnical system *principle of minimum task specification*, which states that optimum worker satisfaction is achieved where the worker's task is minimally specified. The specification of a worker's responsibilities should therefore not extend beyond the minimum requisites (Thimbleby, 1980). In keeping with this general principle, interactional modalities which enhance the user's empowerment relative to the system are to be desired.

In Part I, we introduced the communications channel metaphor for describing the function of a particular expert system in supporting a given user. In keeping with the organizational viewpoint of Part II, let us extend this idea to account for the confluence of multiple such knowledge channels within an enterprise. Knowledge-based systems offer increased opportunities for organizations to collect, analyze, and disseminate information throughout the workforce. As expert systems penetrate all areas of an enterprise, there will arise a need to merge the various knowledge channels. This is consistent with the quality revolution's focus on integration.

The need for integrating multiple knowledge sources is already apparent in some settings. Consider the insurance agent, who must work on a variety of policies ranging from workers compensation to fire and property claims. Each of these policy types utilizes different knowledge resources. For workers compensation there are volumes of state regulations; for fire claims there are ISO rates; and for property claims there are flood zone specifications, Dunn & Bradstreet ratings, etc. At Firemen's Fund, workers are already praising the expert system as a great assistant in processing workers compensation claims. An expert system designed to collect the knowledge needed for multiple claim categories would equip workers with powerful multifunctional capabilities.

Currently there is a movement among information system designers toward computer-supported cooperative work. This is based on the idea that intellective work is largely cooperative in nature, and it suggests that channels for communicating knowledge should be configured for group rather than individual consultative access. Curiously, computing trends seem to be emphasizing facilitation of individuals over facilitation of the group. Consider the fact that the dominant trend in hardware support has been the development of ever more powerful personal workstations. The failure to address issues of differential interactional modalities may have been due to the conceit that only one user is to be targeted in system development. Finally, consider the other end of the knowledge channel - the expert. Knowledge engineering has largely operated under the presumption that there is one expert per task domain; as a result, there has been little effort spent on methods for combining multiple sources of expertise.

The transition from autonomous to cooperative work styles is difficult, because it seems to contradict tacit conceptualizations of the worker and his/her work. Walton (1980) noted in his work on the factory floor that work teams had developed notions of individ..ual autonomy rather than the intended cooperation preparatory to the social reorganization of work. Researchers have begun to examine how to facilitate cooperative work over electronic media using expert systems (compare Grief & Ellis, 1987). Such enquiry is in its infancy, and we expect it to grow steadily. Such enquiry represents the leading edge of the quality revolution's impact on information system design.

# IMPLEMENTATION OF EXPERT SYSTEMS

Östberg (1986) has advised expert system builders to remember that they are not designing a computer system *per se* – they are putting a process in place within the context of an enterprise. Two general problems characterize unsuccessful expert system implementations to date: (1) the system failed to replicate the knowledge actually used by the original expert in solving task problems; and (2) the systems failed to provide the type of assistance needed by the end users. Identification of expertise, the guidelines for modelling it, and the functional value of the resulting product can only be done with respect to this context. A failure to understand the user's social needs may therefore lead to a system which is poorly used, if used at all.

Since expert systems built to date are narrow in scope and expensive to develop, there must be a strong motivation for undertaking such projects. Examples of such motivations are product enhancement (Helena Laboratories' electrophoresis analyzer), archiving of experience (DELTA/CATS), improved service (XSEL, XCON, ABB Robotics' Service Assistant), enforcement of functional consistency (Insurance Underwriters' Assistant, Loan Assistant), and, of course, perceived cost effectiveness. In a survey limited to financial service industries, Coopers & Lybrand (1987a) found that 41% of the responding companies sought improved quality and consistency of employee output, 21% anticipated increased productivity, and 18% expected a broader distribution of scarce resources. These motivations address value within the delivery environment for the expert system – the setting into which it is to be finally inserted for use. The local goal is efficiency; the global goal is advantage. For any enterprise to realize the benefits of expert systems requires an effective implementation strategy.

We found in our work that five characteristics of expert system implementations recognized by others (Leonard-Barton, 1984; Miller, 1984; Mumford, 1987; NAS, 1986) must be considered. These are the culture of automation, the applicability of the expert system, the functionality of the expert system, the usability of the expert system, and the implementation process itself. While it is not necessary for each to be considered in every implementation, other research on information systems indicates that success is more likely where all are addressed. Any complete assessment of implementation practices must also consider the organization's stage of technological development and the users' level of reliance on the delivered system. The human resources practices of the enterprise, varying along a dimension of control versus quality, will also be relevant. Let us briefly examine each of the five primary factors.

*The culture of automation* refers to a general perception of the role of information systems in the organization (from the perspective of the user). If information technologies have been applied in the past to reduce requisite job skills or eliminate jobs (i.e., to enforce the control paradigm), then users are likely to perceive expert systems implementation in a similar manner. In one organization there was a history of using computer technology to reduce labor costs. Even though this organization had developed an expert system to utilize as a support tool rather than a surrogate expert, several professionals interviewed still felt that the end result would be to replace highly skilled personnel with less skilled people.

One way to counteract this sort of pessimism is to make specific attempts to give workers a measure of control over the new technology. For example, the ACE system for telephone cable maintenance had the ability to provide a technician with raw data by which to check the expert system's performance (Miller, 1984). Unfortunately, there were still problems of user acceptance due to the belief that the expert should be making decisions rather than the system.

*Applicability* refers to the capacity of the system to fit its function to the current flow of work activities. Since expert systems must be integrated into the social organization of work, there must be a fit between the technology, the users, and task responsibilities (Mumford, 1987). If the expert system changes the work flow, the consequences may be detrimental to the worker's effectiveness. An example of an inappropriate fit comes from a financial services organization we visited. An underwriter in this company is confronted with three information systems – a manual system, a large mainframe database, and an expert system. Current practice is to develop quotes by hand, reserve and print accounts on the mainframe, and use the expert system as an electronic filing system. The underwriters were, in effect, doing the same work three times – once within each system. The underwriters considered this duplication of effort a major work flow problem and an impediment to efficient processing of business. They felt that the computerized information systems were actually serving to hold them back. As a result, no underwriter would use the expert system unless it was mandatory to do so.

*Functionality* of an expert system is its ability to perform at a level acceptable to users (Goodwin, 1987). The advice generated by the system must be considered valid and relevant to the task at hand. Earlier we discussed the lack of acceptance of GE's DELTA/CATS system. This was due to the delivery of an immature product that was not really ready for use. In another case, we found that an early version of a robot cell maintenance assistant was not considered functional by line engineers because it failed to address the problems that caused the most significant assembly line down-time. With regard to an insurance underwriting system, underwriters complained that they would not develop a policy in the manner of the expert system. Such problems indicate a failure to adequately assess the user's functional needs.

*Usability* of an expert system refers to the ease with which a person can learn it, use it, and maintain it. Both ACE and insurance underwriting systems were judged easy to learn. However, the ACE system afforded repair workers a measure of local control in that they could: (1) tailor its use to their locality and (2) access the knowledge base to effect minor updates. In contrast, the underwriting system was a black box to its users. No local modifications were permitted; the knowledge base was updated at a central location. Usability assessment is still more of an art than a science. Usually only organizations with heavy commitments to expert systems development commit the resources necessary for such analyses. Since usability of a system can ultimately affect its level of functionality, failure to assess usability can be a mistake (Goodwin, 1987).

Finally, we turn to *the implementation process* itself. Central to this is the involvement of the users beginning in the early stages of development and continuing throughout the project. Where the development process does not incorporate user input, there is the possibility that the expert system will fail to fulfill a useful role in the delivery environment. Consider the example of medical expert systems. AI researchers have often used medical domains as proving grounds for knowledge-based systems. The primary motivation is to demonstrate an expert system; the choice of application is secondary. These researchers assume that the technology will be welcomed by physicians, so they try to provide it whether or not it was requested. As a result, medical systems number in the hundreds. Only a handful of these have actually made the transition from laboratory to practical use. Without input from the eventual users, the resultant products were difficult or impossible to employ in the planned setting (Kingsland, 1987). In hindsight, they were doomed from the beginning.

Mumford (1987) has argued that the involvement of users in system development is critical to making good design decisions. A review of the other listed criteria for a successful implementation shows that there is a continual need for the users to evaluate the technology or to specify the appropriate domain parameters. Mumford goes on to suggest that there is a second reason for involving users in the development effort – to build ownership into the design. This entails a commitment of sorts on the part of organizational management, because having once allowed users to build ownership it is dangerous to remove that ownership or close the opportunities for feedback. In the underwriting system the organization actively sought user input into the system design. Once the system was operational, control of future development was removed to the central office. As noted above, the underwriters interviewed have been very resistant to using the system since its installation.

# EXPERT SYSTEMS AND THE QUALITY OF WORK LIFE

In an earlier section we discussed two revolutions in organizational life – one of control and one of quality. We have emphasized the contrast between these two revolutions to clarify our arguments. Both revolutions represent shifts in methods for designing production systems in general. With particular reference to information systems, these perspectives have consequences for the introduction of knowledge-based technology into the workplace. Let us now apply the control/quality distinction to examine some likely effects of expert systems on the quality of work life.

## THE DIVISION OF KNOWLEDGE

Labor has long been seen as a valuable commodity. With the application of scientific methods to industry, it came to be seen as a subject of rationalization. Labor has been continually dissected, quantified, and reorganized since the eighteenth century – usually in the name of efficiency, always in the pursuit of advantage. The subsuming name for the result has been *division of labor*. This label explicitly refers to the decomposition of tasks into smaller units. Implicitly it includes rationalization and analysis of a given task, reorganization of task components, specialization of worker's responsibilities, and stratification of authority. This manipulation of labor is coincident with, if not contingent upon, the arrival of automation in the given job setting. Mechanization of labor did not beget division of labor, but it certainly accelerated the progress of the control revolution.

Two centuries after the onset of the industrial revolution, there is a new venue for automation. Expert systems shoulder mental work much as steam engines shouldered physical work. We find every reason to believe that the analogy extends far beyond the surface and far beyond the present. Knowledge about manual work allowed work study engineers to devise elaborate schemes for implementing a division of labor. Knowledge about knowledge work will allow managers and/or engineers to apply division of labor concepts to the realm of knowledge workers. It is not unreasonable to foresee formalization, compartmentalization, and reorganization of a knowledge commodity analogous to the effects upon the labor commodity. Östberg (1988) termed this prospect a *division of knowledge*.

The initial step in this process – the characterization of knowledge as a commodity – has already been taken. The division of knowledge has its own tacit axioms, as outlined by Gregory (1986, p. 835):

> *"(1) **Knowledge is Distinct from Knowing:** ... For knowledge to be representable in a knowledge base it must be the kind of thing that is logically separable from the knowing of it. In other words, knowledge is a commodity that can be traded, remembered, forgotten, discovered, taken, or left. ... In this view it makes sense to speak of 'pieces' of knowledge and to conceive of these pieces as being transferable ...*
> *(2) **Knowledge is Distinct from its Knower:** ... Just as data are distinct from and uninfluenced by computer disks, so pieces of knowledge are separate from and unaffected by their knower(s).*
> *(3) **Knowledge is a Set of True Facts Bound Together by Rules:** ... [T]he symbols ... in the representational system correspond to objects that exist in the real world – and of course, the relationships that bind them are held to be no less real.*
> *(4) **Knowledge is Reducible:** ... Using the metaphor of the reducibility of the physical world to molecules and atoms, knowledge representation theorists ... have assumed that knowledge is similarly and even necessarily analyzable into its primitive components."*

These ideas form a foundation for manipulating knowledge for the attainment of advantage. Knowledge acquisition is the quantification of expertise. Expert system construction is the reduction of this expertise into a syntax recognizable by a computer. Design of the user interface and the consultee's task responsibilities impose constraints on the worker. These activities act to prescribe worker performance. To date, this has largely meant separating the worker from judgemental activities. Ultimately, this can lead to organizational structures wherein control is embedded in a computerized decisionmaking bureaucracy. Frail knowledge-based technology imbued with such power can only result in havoc. For an illustration we need only look to the role of computer-based "program trading" in bringing about the October 1987 stock market crash in the United States.

## THE PROSPECT OF ALIENATION

From the early days of the industrial revolution onward, workers' labor has been rationalized and quantified into task models. These models are then used to fragment that same labor – breaking procedures into series of discrete sequential steps to be performed by specialized individuals. Such fragmentation results in each worker being separated from the whole of the work product, expected to focus only on the actions needed to complete his/her assigned subtask. The work loses its meaningfulness for the individual worker, and he/she feels powerless to change the situation. The worker becomes alienated from the work process, and he/she invariably grows dissatisfied. If the worker ever had a broad understanding of the production process, this knowledge atrophies. This scenario is the classic deskilling argument expounded by Braverman (1974).

The principal motivation for fragmenting work was the belief that a division of labor would inevitably provide a more efficient production process operating under greater management control. Earlier, we suggested that the division of labor was extensible to a division of knowledge. A proficient person's knowledge used in intellective duties could be codified in a knowledge base and fragmented to support a series of tasks requiring less adept workers.

These new, less skilled, workers would find the decisionmaking tasks more a routine job than a meaningful intellectual exercise. We previously discussed the critical nature of context in the communication of expertise. This loss of meaningfulness will be exacerbated when contextual information is not captured in the knowledge engineering process. As the decisionmaking subtasks become more meaningless to the individual worker, he/she loses any sense of the work's intrinsic value becoming disengaged from the work. The division of knowledge, similar to the division of labor, may well cause alienation in the workforce.

Expert systems may induce their effects more subtly than earlier production technologies. To illustrate, consider the idea of premise control as discussed earlier. In this type of control strategy, the premises underlying decisionmaking are rigidly governed by management. So long as decisions follow logically from these premises, an enterprise can direct the course of decisionmaking without explicit task regulation. Translating this concept to knowledge-based systems, an enterprise can direct the course of rational decisionmaking tacitly by controlling the form of the knowledge base. In such a situation, the worker may not perceive the controlling influence. Although manipulated, he/she may not feel alienated due to the illusion of personal authority.

Alienation is a derivative effect of technology – an effect occurring long after the process of implementation has been completed. Alienation effects of expert systems are not likely to surface until some time after the systems' installation in the workplace. Expert systems have not yet penetrated organizations sufficiently for such effects to be discerned. There is some recent evidence, obtained by examining the introduction and use of computer/automation technology, that suggests a lack of alienation. It derives from a case study (Salzman, 1987), a historical analysis (Adler & Borys, 1986), and survey research

(Haddad *et al.*, 1985). Thus, the prospect for alienation is ambiguous, greatly depending on the interplay between the control and quality revolution.

## EXPERT SYSTEMS APPLICATIONS AND WORKER STRESS

As we have emphasized repeatedly, there are so few operational expert systems to study that drawing conclusions about their workplace impacts is difficult. Whereas alienation is a somewhat subtle effect of new technology, health impacts such as stress are more concrete in nature and hence more easily discerned. Amick & Östberg (1987) have reviewed the literature and discussed the health and stress implications of automation elsewhere. Based on that work, a framework for workplace stress was devised. That framework, modified for specific reference to expert systems, is presented in Figure 12.

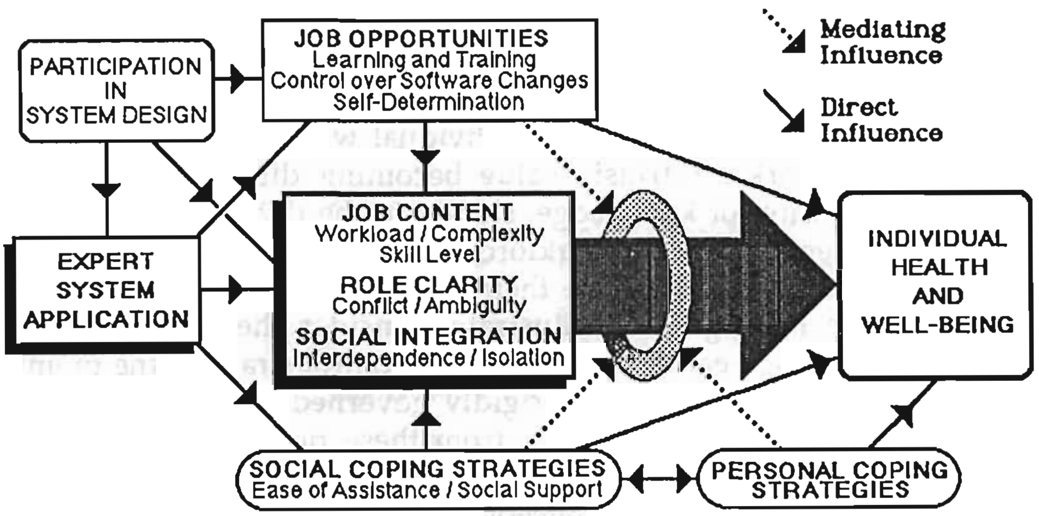


**Figure 12**

*A framework for workplace stress examining how an expert system application can affect the job, reducing opportunities to cope with job demand and potentially modyfying an individuals' health status. [Adapted from Amick & Östberg, 1987].*

Expert systems would seem to offer the promise of reducing worker stress by reducing the effects of known stressors. For example, consider the set of job opportunities identified in the figure above; their absence is a stressor, while their presence tends to alleviate stress. A properly designed expert system can offer workers the ability to maintain if not expand a broad knowledge of the work process. Where long term knowledge base maintenance is done by the end users, they thereby enjoy a measure of local control over the decisionmaking apparatus. Miller (1984) found that the ACE system was judged much easier to learn and modify at the local (workplace) level than its predecessor.

More generally, stress can be reduced if the individual is spared the frustrating minutiae of routine activities. Recall our earlier discussion of the insurance agent who had to spend hours searching for knowledge about specific rules and the CAD user who would spend an inordinate amount of time doing rote circuit board designs. Such repetitive acts are a major component of an individual's work burden. To the extent that an expert system obviates the need for such tedium, we would expect a commensurate reduction in the stress induced by the work environment.

By counteracting stressors, well-designed expert systems hold the potential for materially improving the psychosocial work environment. Note, however, the prudent use of the qualification "well-designed". Where the job opportunities illustrated in the figure are restricted by an expert system, or where the system maintains or even increases the level of frustrating tedium entailed in a given task, the result may well be additional stress. These job opportunities are issues of worker status in an enterprise and his/her role in the production process. The degree to which an expert system affects the existing workplace stress level is a direct consequence of the degree to which user needs and desires are incorporated into its design, development, and deployment.

## EXPERT SYSTEMS AND INNOVATION

There are many factors influencing the ways that work organization affects worker creativity and innovation. In keeping with our focus on knowledge-based technology, we will limit our discussion to the role played by knowledge or expertise. We have discussed how the isolation of workers via fragmentation of knowledge can lead to alienation and a decreased motivation for enquiry. Where expert systems are employed to control workers, there is little likelihood that innovative behavior will be encouraged. Indeed, the very concept of control is antithetical to innovation.

Fortunately, we believe that expert systems offer the opportunity to enhance workers' creative potential by channeling expert knowledge to a wider audience. Klein & Hirschheim (1985) suggest that decision support systems enhance the interpretive aspects of task knowledge. The capture of task expertise and its dissemination in a comprehensible form gives workers access to a wide breadth of perspective on the production process. Initially, this could be result in innovative ways of improving the existing system of production. Consider the example provided by Zuboff (1985b) of the workers who, once enabled through computerization to see and thus comprehend the entirety of the work system, pinpointed the production bottlenecks and generated creative solutions for them. Furthermore, this expanded view will allow, if not impel, users to conceive, consider, and explore alternatives to the status quo.

At this time, most expert systems are not being consciously intended to stimulate users' creativity. This is due to a lingering influence of the control revolution mindset. The systems, even when ostensibly applied to increase quality, are being employed to increase and enforce consistency in work performance. Workers are therefore not being afforded the opportunity to question and to explore. In organizations where the production process is not designed to encourage enquiry, there will likely be decreased commitment to the work and an attendant decrease in quality (Walton, 1985).

# STRATEGIC USE OF EXPERT SYSTEMS

Because of their abilities to integrate a variety of complex data bases into a multifunctional workstation and control professional behavior, expert systems can offer strategic advantages to an enterprise. For example, such systems integrated into the organization's communication system can afford management greater control of information flows and professional behavior. This may provide leverage against seemingly inherent fluctuations in professional decisionmaking Expert systems thus move beyond traditional decision support by reducing the risks associated with managerial decisionmaking. They can provide a company with a competitive advantage by facilitating rapid responses to market conditions (e.g., managing rates, costs, or other parameters). This notion of competitive advantage is apparently a common point of attraction to the technology. In a recent survey of the financial services industry, 40% of MIS executives and 32% of management felt expert systems would confer a strategic edge.

Despite this widespread belief, there are few concrete examples of competitive advantage provided by implementing an expert system. The most widely cited instance is Digital Equipment Corporation's XCON, which delivers savings of $18 million annually. Unfortunately, this impressive figure has never been rigorously scrutinized in the literature. Beckman Instruments' SPINPRO, by reducing run times by up to 70%, will hopefully enlarge that company's share of the ultra centrifugation market (Wang, 1987). Helena Laboratories' scanning densitometers equipped with expert systems have accounted for approximately 65-70% of the densitometer market since their introduction, but there is some question whether the expert systems can be credited for this relative dominance (Richmond & Landers, 1987).

Sviokla (1986), drawing on the work of Porter (1985), argues that expert systems may ultimately be a pivotal part of the strategic use of information systems. He predicts that enterprises will search for ill-structured problems in expertise-intensive areas of the value chain, hoping to discover potential competitive opportunities for applying knowledge-based technology. Sviokla concluded his expert systems review by noting on the downside that: (1) such systems are not tied to the firm's strategy; and (2) there is no discussion of the intricacies of relation expert systems to the firm's activities for creating value.

Consider, for example, the potential impact of expert systems in the financial services industries. The technology may create the ability to produce financial plans at a lower cost to more people, thereby promoting market rivalry. It may also reduce the barriers to market entry due to the enhanced accessibility of advanced knowledge via expert systems. Still, the applicability of Porter's value chain or the role of expert systems in competitive markets is unknown at this time. This information will probably come only from experience, as the systems are installed and evaluated in real world settings.

Despite the unclear ability to apply information systems for sustained competitive advantage (Vitale *et al.*, 1986), many organizations see expert systems as a means for creating or reinforcing an edge. A major motivation for investment in the technology is its promise of reducing the cost of bad decisions. Consider recent data from the insurance industry. Best's Reviews (as cited in Coopers & Lybrand, 1987b) found that expenditures by insurers for property and casualty claims was $116 billion, compared to $44.5 billion paid to cover operational expenses. A one percent reduction in losses through the use of expert systems to improve quality and efficiency would reduce operating expenses by three percent. It is still open to question whether expert systems can actually deliver on this promise, since few systems have been operational long enough to gauge such returns.

# SUMMARY

We have reviewed the effects of automation, specialization of labor, and the rationalization of tasks into quantitative models from the perspectives of both control and quality. Analysis and codification methodologies certainly effected a measure of efficiency. Whether they effected any advantage is largely a matter of whose perspective you take. As discussed earlier, not all impacts will be universally appreciated (Amick & Östberg, 1987).

We have shown that in order to allow creative expression and quality work the end users must be provided decisionmaking and intellective skills. Such skills are a function of the users' role in the organization, and this role derives from the enterprise's treatment of the social aspects of work. In designing and implementing expert systems, attention must be given to choices regarding the social organization of work. Inattention to these choices would result in the most negative possible outcome – trivialization of the users. Sadly, we find evidence of this trivialization throughout the expert systems scene. A disturbing example is the following:

> *"In the future artificial intelligence should simulate easily replaceable people. For example, there are a lot of people in the customer-service area [bank tellers, travel agents, airline reservation clerks] who don't do their job very well."*
> (Schank, 1987)

Perhaps this was just an example of poor wording by a respected AI guru; maybe the implicit condescension was not intended. By all indications, as acknowledged by Feigenbaum & McCorduck (1983), the desire to pursue and apply the technology has blinded many to "...the human problems which, in microcosm, reflect problems many workers will have to face." Recognition of problems, though, is not the same as solution. Solutions are unlikely if we accept Feigenbaum and McCorduck's immediate conclusion that "This is a revolution, and all revolutions have their casualties."

Others are not so nonchalant. Trappl (1986) concluded that "it is not so difficult to guess *who* will be replaced by a computer; but it is more difficult to predict *when* this will take place." He encouraged AI researchers to reflect upon the potential impacts of their work and then decide whether they can justify continuing. However, the AI community can only be held minimally accountable for the impacts of expert systems applied in production and services.

We have hinted time and again at the effects of expert systems in terms of setting and enforcing parameters of skill, knowledge, and authority. The presumed target of these trivializing effects has been the system user. To counter that presumption and show that the impacts extend to all persons surrounding the system, consider the following:

> *"One expert who gladly gave himself and his specialized knowledge over to a knowledge engineer suffered a severe blow to his ego on discovering that the expertise he'd gleaned over the years, and was very well paid and honored for, could be expressed in a few hundred heuristics. At first he was disbelieving; then he was depressed. Eventually he*

> *departed his field, a chastened and moving figure in his bereavement."*
>
> (Feigenbaum & McCorduck, 1983)

Knowledge engineers, backed by the latest in computing technology and the highest of priorities, are using *ad hoc* methods to implement ill-defined systems which will undoubtedly affect people in all expertise levels within an enterprise. Since expert systems can modify the control over an organization's inventory of skill and knowledge, they can influence rules, responsibilities, and power. They are thus a political instrument of change (Sviokla, 1986). By virtue of their impacts, these system designers have power over people in the enterprise, especially the end users. A framework for awareness of this empowerment is illustrated in Figure 13.

At present the situation is one of unintended influence. Neither the designers nor the users recognize or comprehend the empowerment. This is not helped by the fact that the newly empowered implementers of expert systems are focussing their attention on technical matters as they attempt to learn the intricacies of knowledge-based system construction. Furthermore, a lack of awareness on the part of management and/or labor results in the potential for manipulation and resistance. This makes for a repetition of battles previously fought during the control revolution.

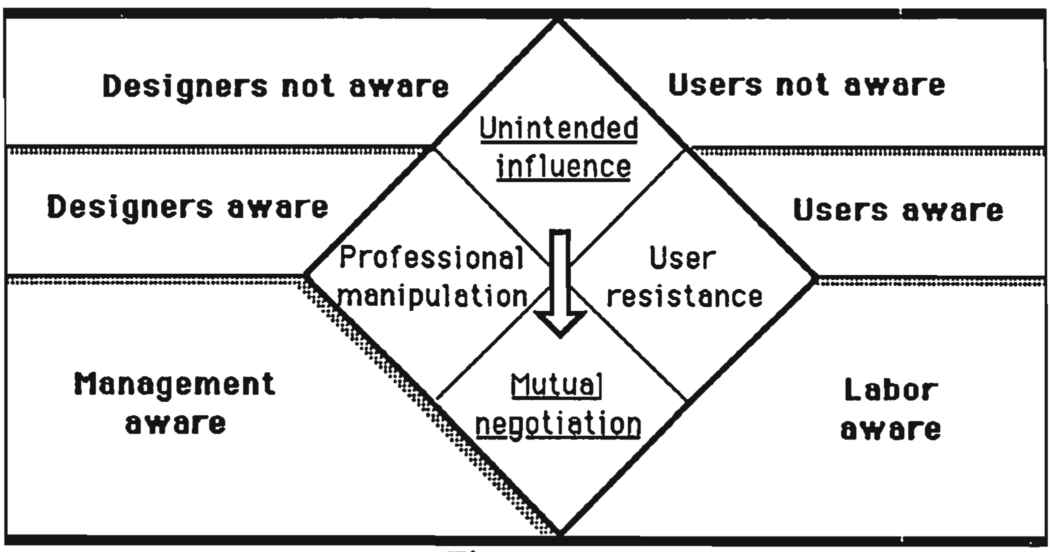


**Figure 13**

*Framework for the relationships between awareness and influence in connection with the introduction of new technology to the workplace. In most expert system application design efforts, attempts ought to be made to move from unintended influence to mutual negotiation. Yet, in order to have mutual negotiation, both management and labor must commit to necessary resources allowing negotiations to occur. [Based on Marcus & Bjørn-Andersen, 1987].*

In a Coopers & Lybrand (1987b) survey of expert systems in financial services, 65% of the companies responding reported development and/or employment of expert systems, while less than 40% characterized their understanding of the technology as above the industry average. So long as managers and developers are still trying to acquire technical proficiency they will likely spend little time considering the end users.

This lack of mutual intentions implies an absence of control on the part of one or more parties. All groups need to become aware of the empowerment issues surrounding introduction of knowledge-based systems. To accomplish this requires eliciting broader management and labor involvement. Only by such simultaneous education can we hope to avoid professional manipulation and/or user resistance. Otherwise, we may see impacts similar to those of the industrial revolution – an automation technology exerting influence first on those directly involved with it, next on the entire enterprise within which it operates, and finally on the society at large. Enquiries into the impacts of automating physical labor were done after the fact, with much human suffering occurring in the mean time. Let us be more punctual in theorizing, detecting, and evaluating the impacts of automating complex mental labor using expert systems.

Enterprises wishing to implement expert systems are faced with many choices. There are no clearcut criteria for making these choices. While this implies that design decisions will continue to be difficult, it also means that there is considerable flexibility in defining an adequate solution to the issues involved. These choices can be characterized as falling on a continuum ranging between the extremes of control primacy and quality primacy. In large organizations, attention must be given to both control and quality. Design decisions will therefore involve negotiating the best tradeoff between these parameters. For this reason, design decisions should be cooperatively made by all concerned parties.

We have been concerned not only with the accessibility of the system to the user, but its impacts on the user. To understand these impacts it is essential to consider the design process as operating at both the level of the information system being constructed and at the level of the organizational system within which it will be installed. It is important that users be involved in the design of expert systems. Such user involvement has been critical in reforming organizations during the quality revolution. Emphasizing quality entails empowering the worker by affording him/her the ability to create, innovate, and develop within the workplace.

How, then, can we work toward quality in the proliferation of expert systems? Knowledge engineers and other development personnel must cooperate with work environment experts in establishing design criteria based on the significance of the users. They must also avoid the easy path of building systems which tacitly embody the control model of an organization (Rosenthal & Salzman, 1986). Table 2 illustrates a framework for organizing the expert system design process. The table covers the course of the process from initial planning through to installation and maintenance. The time line proceeds from top to bottom. The orderings of the three groups of factors (design choices, significant actors, and system development stages) are intended to convey a loose relative correspondence among the factors rather than a strict prescriptive mapping. The design choices illustrated in Table 2 are limited to those choices pertaining to the system itself. Those other choices pertaining to human resources and work organization are not explicated here.

## Table 2

*Framework for the expert system development cycle*
*[Based on Rosenthal & Salzman, 1986]*

| Design Choices | Significant Actors | Development Stage |
|---|---|---|
| Initial Goal Choices | Users, Designers, and Management | Project Planning |
| Conceptual Design | | User Needs Assessment |
| | | Domain Specification |
| Functional Design Expertise | Knowledge Engineers and Expert(s) | Codification of Knowledge Engineering |
| Operational Design | Knowledge Engineers | Programming |
| System Delivery | Users, Managers, and Development Team | System Evaluation System Installation |
| Organizational Design | Users and Managers | Maintenance, Updating of Knowledge Base |

Is the expert systems development cycle different from that of any new production technology? As pointed out in Bramer's (1987) recent review of expert systems in Britain, the fact that all technology has the potential for both positive and negative uses is no justification for neglecting any further analysis:

> *"What is fundamentally different about Expert Systems, in contrast to technological developments such as, say, telephones, cars, computerised stock control systems or bank cash dispensers, all of which have their negative side but are fundamentally helpful, is that Expert Systems are frequently (although not always) concerned with the judgement made by highly-skilled experts who collectively comprise the leaders of society."*

# REFERENCES

Adler, P., Automation and Skill: New Directions. Stanford, CA: Department of Industrial Engineering and Engineering Management, Stanford University, 1986. (a)

Adler, P., New technologies, new skills. *California Management Review*, 29(1), 9–28, 1986. (b)

Adler, P., & Borys, B., Automation and work: Lessons from the machine tool case. Stanford, CA: Department of Industrial Engineering and Engineering Management, Stanford University, 1986

Amick, B. C., III, & Östberg, O., Office automation, occupational stress and health. A literature analysis with specific attention to expert systems. *Office: Technology and People*, 3, 191–209, 1987.

Ashby, W. R., *An Introduction to Cybernetics*. London: Chapman & Hall, 1958.

Attewell, P., *The De-Skilling Controversy*. Stony Brook, NY: Department of Sociology, SUNY, 1984.

Barrett, W., *Death of the Soul: From Descartes to the Computer*. New York: Doubleday, 1987.

Bell, D., *The Coming of Post-Industrial Society: A Venture in Social Forecasting*. New York: Harper & Row, 1973.

Bell, D. *Tracking Technology*. Speech delivered at the Office of Technology Assessment, Washington, DC, September, 1987.

Beniger, J., *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge, MA: Harvard University Press, 1986.

Bobrow, D. G., Mittal, S., & Stefik, M. J., Expert systems: Perils and promise. *Communications of the Association of Computing Machinery*, 29, 880–894, 1986.

Bobrow, D. G., & Stefik, M., Knowledge-based programming. *Alvey News* (UK), 13–14, No. 13, 1985.

Bonissone, P. P., & Johnson, H. E., Jr., Expert system for diesel electric locomotive repair. *Human Systems Management*, 4, 255–262, 1984.

Bramer, M. A., Expert systems in Britain: Progress and prospects. In M. A. Bramer (Ed.), *Research and Development in Expert Systems III*. Cambridge: Cambridge University Press, 1987. (Pp. 1–12).

Branscomb, L. M., & Thomas, J. C., Ease of use: A system design challenge. *IBM Systems Journal*, 23, 224–235, 1984.

Braverman, H., *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York: Monthly Review Press, 1974.

Brown, J. S., The low road, the middle road and the high road. In P. A. Winston and K. A. Prendergast (Eds.), *The AI Business: The*

*Commercial Uses of Artificial Intelligence*. Cambridge, MA: MIT Press, 1984. (Pp. 81–90).

Buchanan, B. G., Expert systems: Working systems and the research literature. *Expert Systems*, 3, 32–51, 1986.

Burnes, B., & Fitter, M., Control of advanced manufacturing technology: Supervision without supervisors?. In T. Wall, S. Clegg and N. Kemp (Eds.), *The Human Side of Advanced Manufacturing Technology*. New York: John Wiley & Sons, 1987. (Pp. 83–100).

Chandler, A., The emergence of managerial capitalism. *Business History Review*, 58, 473–503, 1984.

Chomsky, N., *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press, 1965.

Clancey, W. J., Tutoring rules for guiding a case method dialogue. *International Journal of Man-Machine Studies*, 11, 25–49, 1979.

Clancey, W. J., Extensions to rules for explanation and tutoring. In B. Buchanan and E. Shortliffe (Eds.), *Rule-Based Expert Systems*. Reading, MA: Addison-Wesley, 1984. (Pp. 531–568).

Clausewitz, C., von, *On War*. London: Pelican Classic, 1968. (Abridged English edition of *Vom Kriege*, 1832).

Cochran, E., & Christopherson, P., Corporate Systems Development Division, Honeywell Inc., Golden Valley, MN, 1987. (Personal communication, July).

Cochran, E.,& Hutchins, B., Testing, Verifying, and Releasing an Expert System: The Case History of Mentor. Golden Valley, MN, Honeywell Corporate Systems Development Division, 1987.

Collins, H. M., Green, R. H., & Draper, R. C., Where's the expertise?: Expert systems as a medium for knowledge transfer. In M. Merry (Ed.), *Expert Systems 85*. Cambridge, UK: Cambridge University Press, 1985. (Pp. 323–334).

Coopers & Lybrand, *Expert Systems in the Insurance Industry: 1987 Tabulation of Survey Results*. New York: Coopers & Lybrand Management Consulting Services, 1987. (a)

Coopers & Lybrand, *Expert Systems in the Insurance Industry: 1987 Survey Report Update*. New York: Coopers & Lybrand Management Consulting Services, 1987. (b)

Dallas, D. B., The increasing use of expert systems. *D E Technology*, 48(6), 17–20, 1987.

Davis, R., Interactive transfer of expertise. *Artificial Intelligence*, 12, 121–157, 1979

DeKleer, J., AI approaches to troubleshooting. *Artificial Intelligence in Maintenance*. Denver, CO: Denver Research Institute, 1984. (Pp. 83–94). (AD–A145 349)

Dietz, P. W., Artificial intelligence: Building rule-based expert systems. *Design News*, 136–144, March 3, 1986.

Dijkstra, E. W., *A Discipline of Programming*. Englewood Cliffs, NJ: Prentice-Hall, 1976.

DM Data, *AI Market Report*. Scottsdale, AZ: DM Data, Inc., 1985.

DM Data, What you should know about expert systems. *D E Technology*, 48(6), 12–16, 1987.

Doyle, J., Big problems for artificial intelligence. *AI Magazine*, 9(1), 19–22, 1988.

Drucker, P. F., Technology and society in the twentieth century. In M. Kranzberg and C. W. Pursell Jr. (Eds.), *Technology in Western Civilization*. New York: Oxford University Press, 1976. (Pp. 22–33).

Drucker, P. F., *Managing in Turbulent Times*. New York: Harper & Row, 1980.

Engelberger, J. F., *Robotics in Practice*. New York: American Management Association, 1980.

Ernst, M. L., & Ojha, H., Business applications of artificial intelligence knowledge based expert systems. *Future Generations Computer Systems*, 2, 173–185, 1986.

Feigenbaum, E., Expert advice: "Boulders rather than pebbles". (Quotations from a conference address). *Intelligence*, 5(10), March, 1988. (Texas Instruments).

Feigenbaum, E. A., & McCorduck, P., *The Fifth Generation*. Reading, MA: Addison-Wesley, 1983.

Feorene, O. J., Organization and administration of industrial engineering. In G. Salvendy (Ed.), *Handbook of Industrial Engineering*. New York: Wiley, 1982. (Pp. 1.2.1–28).

Fisher, G., & Stevens, C., Volunteering information–enhancing the communication capabilities of knowledge-based systems. *Proceedings of 2nd IFIP Conference on Human-Computer Interaction* (INTERACT '87), 1987. (Pp. 965–971).

Gaines, B. R., & Shaw, M. L. G., Foundations of dialog engineering: The development of human-computer interaction (II). *International Journal of Man-Machine Studies*, 24, 101–123, 1986.

Gallant, S., Connectionist expert systems. *Communications of the ACM*, 31, 152–169, 1988.

Gardell, B., *Work Organization and Human Nature: A Review of Research on Man's Need to Control Technology*. Stockholm: Work Environment Fund, 1988.

Garwood, G., Role of expert systems in the maintenance of telecommunications networks. *European Conference on Electrotechnics* (EUROCON '84), 6, 9–13, 1984.

Gillett, P., Victory from the jaws of defeat. *Alvey Conference Report* (UK), 20, 1987.

Goodwin, N., Functionality and usability. *Communications of the ACM*, 30, 229–233, 1987.

Gregory, D., Delimiting expert systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC–16, 1986. (Pp. 834–843).

Gregory, D., Philosophy and practice in knowledge representation. In J. Zeidner (Ed.), *Human Productivity Enhancement*, Vol. II. New York: Praeger, 1987. (Pp. 165–203).

Grief, I., & Ellis, C. (Eds.), *ACM Transactions on Office Automation Systems*, 5(2), 1987. (Special issue on computer-supported cooperative work).

Haddad, A., House, J., & Jackman, M., *The Impact of Technology on Industrial Work and Workers*. Ann Arbor, MI: Department of Sociology, University of Michigan, 1985.

Hägglund, S., *Computer Supported Knowledge Systems in the Office of the Future*.(In Swedish). Farsta, Sweden: TELDOK, 1986.

Hägglund, S., Occupational skills and knowledge demands in the 90's computerized systems. (In Swedish). Linköping, Sweden: Department of Computer Sciences, University of Linköping, 1987. (R-87-12).

Hardy, S., Expert systems for diagnosis. *Proceedings of 1st SMI Conference on Expert Systems for Advanced Manufacturing Technology*, Dearborn, MI, June, 1987. (Pp. 301–308).

Hart, A., *Knowledge Acquisition for Expert Systems*. London: Kogan Page, 1986.

Harvey, J. J., Expert systems: An introduction. *Electrical Communication*, 60, 100–108, 1986.

Hayes, C., & Wright, P., Automated planning in the machining domain. In S. C.-Y. Lu, & R. Komanduri (Eds.), *Knowledge-Based Expert Systems for Manufacturing*. New York: American Society of Mechanical Engineers, 1986. (Pp. 221–232.

Hebb, D., *The Organization of Behavior*.,New York: John Wiley & Sons, 1949.

Helander, M., Emerging office automation systems. *Human Factors, 27,* 3–20, 1985.

Helander, M., & Östberg, O., *The Automated Office – A Description and Some Human Factors Design Considerations*. Farsta, Sweden: TELDOK, 1983.

Heng, M. S. H., Why evolutionary development of expert systems appears to work. *Future Generations Computer Systems*, 3, 103–109, 1987.

Hewett, J., & Sasson, R., *Expert Systems 1986*. London: Ovum Ltd., 1986.

Hewitt, C., The challenge of open systems. Byte, 223–242, April, 1985.

Hull, C. L., *Principles of Behavior*. New York: Appleton-Century-Crofts, 1943.

Hwang, S.-L., & Salvendy, G., Human supervisory performance in flexible manufacturing systems. *Proceedings of the Human Factors Society 28th Annual Meeting*, 1984. (Pp. 664–669).

Japan Electronic Dictionary Research Institute, Ltd, quoted in Lenat, D. B., & Feigenbaum, E. A., On the threshold of knowledge. *Proceedings of 10th International Joint Conference on Artificial Intelligence* (IJCAI '87), 1987. (Pp. 1173–1182).

JMA, *Robotics and the Manager*. (Investigation report from the Japanese Management Association; English edition). Tokyo: Fuji Corporation, 1983.

Jones, F. N., & Watson, P. A., Expert systems and the telecom manager. *Telecommunications*, 18(9), 61–90, 1984.

Jönsson, B., Productivity and survival – key words in industry. Paper presented to the Management Centre Europe conference on 'Productivity – The Key to Survival in the 80s', London, November, 1983.

Kingsland, L. C., III., National Library of Medicine, National Institutes of Health, Bethesda, MD, 1987. (Personal communication, June).

Klein, H., & Hirschheim, R., Fundamental issues of decision support systems: A consequentialist perspective. *Decision Support Systems*, 1, 5–24, 1985.

Langefors, B., *Theoretical Analysis of Information Systems*. Stockholm: Studentlitteratur, 1966.

Langefors, B., Distinction between data and information/knowledge. *Information Age*, 9, 89–91, 1987.

Lehner, P. F., & Zirk, D. A., Cognitive factors in user/expert-system interaction. *Human Factors*, 29, 97–109, 1987.

Leith, P., The application of AI to law. *AI & Society*, 2, 31–46, 1988.

Lenat, D. B., & Feigenbaum, E. A., On the threshold of knowledge. *Proceedings of 10th International Joint Conference on Artificial Intelligence*, (IJCAI '87), 1987. (Pp. 1173–1182).

Lenat, D., Shepherd, M., & Prakash, M., CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine*, 6(4), 65–84, 1986.

Leonard-Barton, D., *Implementing Innovations: The Automation of Knowledge-Based Production Tasks*. Cambridge, MA: Harvard Business School, 1984.

Lindholm, R., quoted in Corlett, N., A tract for the times. *The Ergonomist* (UK), 1–2, October, 1987.

Marcus, M. L., & Bjørn-Andersen, N., Power over users: Its exercise by system professionals. *Communications of the ACM*, 30, 498–504, 1987.

Michie, D., High-road and low-road programs. *AI Magazine*, 4(3), 21–22, 1983.

Miller, F. D, Introducing expert systems into the workplace. *Proceedings of Symposium on Combining Human and Artificial Intelligence*. New York: Human Factors Society Metropolitan Chapter, 1984. (Pp. 41–51).

Miller, T. M., Human Resources Laboratory, Wright Patterson AFB, Ohio, 1987. (Personal communication, August).

Minsky, M., *The Society of Mind*. New York: Simon and Schuster, 1985.

Missikoff, M., & Wiederhold, G., Towards a unified approach for expert and database systems. *Proceedings of 1st International Workshop on Expert Database Systems*, Menlo Park, CA, 1986. (Pp. 383–39).

Miyake, N., & Norman, D. A., To ask a question, one must know enough to know what is not known. *Journal of Verbal Learning and Verbal Behavior*, 18, 357–364, 1979.

Mowshowitz, A., The future with AI: Freedom and community? Unpublished report from a workshop held at the U.S. Congress' Office of Technology Assessment, Washington, DC, November 9, 1984.

Mumford, E., The Successful Design of Expert Systems – Are Means More Important Than Ends? Manchester, UK: Manchester Business School, 1987.

Myers, G. J., *The Art of Software Testing*. New York: John Wiley & Sons, 1979.

NAS, *Human Resource Practices for Implementing Advanced Manufacturing Technology*. Washington, DC: National Academy of Sciences, 1986.

Norem, P., Quantitative artificial intelligence: How much is enough? *Association of Computing Machinery SIGART Newsletter*, 25–27, No. 91, 1985.

O'Connell, A. M., Expert maintenance personnel now available "on-line". *Proceedings of 1st SMI Conference on Expert Systems for Advanced Manufacturing Technology*, Detroit, MI, 1987. (Pp. 199–207). (a)

Office of Technology Assessment, *Federal Funding for Artificial Intelligence*. Washington, DC: OTA, U.S. Congress, 1987. (Staff paper).

Ohyama, M., A knowledge-based directory assistance system. *Proceedings of AI '87 Japan*, 1987. (Pp. 204–224).

O'Keefe, R. M., Balci, O., & Smith, E. P., Validating Expert System Performance. *IEEE Expert*, 2(4), 81–90, 1987.

Olson, J., Cognitive analysis of people's use of software. In J. Carroll (Ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*. Cambridge, MA: MIT Press, 1987.

Östberg, O., Expert systems in a social environment – human factors concerns. *Proceedings of the Human Factors Society 30th Annual Meeting*, 1986. (Pp. 739–743).

Östberg, O., Applying expert systems technology: Division of labor and division of knowledge. In B. Göranzon and I. Josefson (Eds.), *Knowledge, Skill and Artificial Intelligence*. Berlin: Springer-Verlag, 1988. (Pp. 169-183).

Östberg, O., & Enqvist, J., Robotics in the workplace: *Robot factors, human factors, and humane factors*. In H. E. W. Hendrick and O. Brown (Eds.), Human Factors in Organizational Design and Management. Amsterdam: Elsevier Science, 1984. (Pp. 447–460).

Parnas, D., Software aspects of strategic defense systems. *American Scientist*, 73, 432–440, 1985.

Partridge, D., Will AI lead to a super software crisis?. In K. S. Gill (Ed.), *Artificial Intelligence for Society*. London: John Wiley & Sons, 1986. (Pp. 31–39).

Partridge, D., & Wilks, Y., Does AI Have a Methodology Different from Software Engineering? Las Cruces, NM: Computing Research Laboratory, New Mexico State University, 1986. (MCCS–86–53).

Paul, R. P., & Nof, S. Y., Work measurement – a comparison between robot and human task performance. *International Journal of Production Research*, 17, 277–303, 1979.

Pask, G., & Gregory, D., Conversational systems. In J. Zeidner (Ed.), *Human Productivity Enhancement*, Vol. II. New York: Praeger, 1987. (Pp. 204–235).

Perrow, C., *Complex Organizations: A Critical Essay*. New York: Random House, 1986. (Third edition).

Porter, M., *Competitive Advantage*. New York: The Free Press, 1985.

Pylyshyn, Z. W., *Computation and Cognition*. Cambridge, MA: MIT Press, 1984.

Rasmussen, J., *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. New York: North-Holland, 1986.

Rasmussen, J., Cognitive engineering. *Proceedings of 2nd IFIP Conference on Human-Computer Interaction* (INTERACT '87), 1987.

Reboh, R., Knowledge engineering techniques and tools in the PROSPECTOR environment. Menlo Park, CA: AI Center, SRI International, 1981. (Report No. 243).

Richmond, J., & Landers F., Helena Laboratories, Beaumont, TX, 1987. (Personal communication, November).

Rosenblatt, F., *Principles of Neurodynamics*. New York: Spartan, 1962.

Rosenthal, S., & Salzman, H., *Values in the Design of Software*. Boston, MA: Center for Applied Social Science, Boston University, 1986.

Rumelhart, D., McClelland, J., & PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*, Vol. 1, (Foundations). Cambridge, MA:, MIT Press, 1986.

Sacerdoti, E., quoting from a Teknowledge survey of users of the Teknowledge S1 and M1 ES development products. Panel discussion at 1st SMI Conference on Expert Systems for Advanced Manufacturing, Dearborn, MI, June, 1987.

Salzman, H., Computer Technology and the Automation of Skill. Paper presented at Meeting of the American Sociological Association, Chicago, IL, August, 1987.

Sasso, W., Olson, J., & Merten, A., *A Comparison of Three Methodologies for Office Analysis*. Ann Arbor, MI: Human-Computer Interaction Laboratory, University of Michigan, 1985.

Schank, R. C., quoted in Sanger, D. E., Computer fails as job-killer. *The New York Times*, 5–6, Career section, October 11, 1987.

Shortliffe, E. H., MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection. Stanford, CA: Stanford University, 1974. (Dissertation).

Smith, A., *An Inquiry Into the Nature and Causes of the Wealth of Nations*, 1776. (Available in R. H. Campbell, ed., Glascow Edition of the Works & Correspondence of Adam Smith. Oxford University Press, 1976).

Smith, J., Expert database systems: A database perspective. *Proceedings of 1st International Workshop on Expert Database Systems,* Menlo Park, CA, 1986. (Pp. 3–15).
Steels, L., The deepening of expert systems. *AI Communications,* 0(1), 9–16, 1987.
Sviokla, J., Business implications of knowledge-based systems. *Database,* 17(4), 5–19, 1986.
Taylor, F. W., *The Principles of Scientific Management,* 1911. (New edition: New York: Norton, 1967).
Thesen, A., Östberg, O., & Lei, L., Knowledge Acquisition for Expert Systems; A Survey and Analysis. Madison, WI: Department of Industrial Engineering, University of Wisconsin, 1987.
Thimbleby, H., Dialogue determination. *International Journal of Man-Machine Studies,* 13, 295-304, 1980.
Thomas, J. C., Artificial intelligence and human factors. *Proceedings of Symposium on Combining Human and Artificial Intelligence.* New York: Human Factors Society Metropolitan Capter, 1984. (Pp. 54–66).
Tichy, W., F., What can software engineers learn from artificial intelligence? *Computer,* 20(11), 43–54, 1987.
Trappl, R., Impacts of artificial intelligence: An overview. In R. Trappl (Ed.), *Impacts of Artificial Intelligence.* Amsterdam: North-Holland, 1986. (Pp. 31-51).
Tsumura, H., Wasano, T., Masuzawa, H., & Miyake, H., Patient oriented multidisciplinary medical consultation system. In R. Salamon and M. Jørgensen (Eds.), *Medinfo '86.* Amsterdam: Elsevier Science, 1986. (Pp. 289-293).
Turing, A. M., Computing machinery and intelligence. *Mind,* 59, 443–460, 1950. (Also in E. Feigenbaum and J. Feldman (Eds.), *Computers and Thought.* New York: McGraw-Hill, 1963).
van Melle, W., A Domain Independent System that Aids in Constructing Consultation Program. Stanford, CA: Department of Computer Sciences, Stanford University, 1980. (STAN-CS-80-820).
Vitale, M., Ives, B., & Beath, C., Linking information technology and corporate strategy: An organizational view. *Proceedings of the International Conference on Information Systems.* San Diego, CA: December, 1986. (Pp. 265–276).
Walton, R., Establishing and maintaining high commitment work systems. In J. Kimberly and R. Miles (Eds.), *The Organizational Life Cycle: Issues in the Creation, Transformation and Decline of Organizations.* San Francisco, CA: Jossey-Bass, 1980. (Pp. 208–290).
Walton, R., From control to commitment in the workplace. *Harvard Business Review,* 63(2), 77–84, 1985.
Wang, E., Artificial intelligence for competitive advantage. *AI Expert,* 2(1), 5–6, 1987.
Weiss, S., & Kulikowski, C., *A Practical Guide to Designing Expert Systems.* Totowa, NJ: Rowman & Allanheld, 1984.

Whitaker, R., & Östberg, O., Channeling knowledge: Expert systems as communications media. *AI & Society*, 2(3), 1988.

Whitby, B., The computer as a cultural artefact. In K. S. Gill (Ed.), *Artificial Intelligence for Society*. Chichester, UK: John Wiley & Sons, 1986. (Pp. 115–124).

Woods, D. D., Paradigms for intelligent decision support. In E. Hollnagel, G. Mancini and D. D. Woods (Eds.), *Intelligent Decision Support in Process Environments*. Berlin: Springer-Verlag, 1986. (Pp. 153–173).

Young, R. M., Human interface aspects of expert systems. In J. Fox (Ed.), *Expert Systems*. Maidenhead, UK: Pergamon Infotech, 1984. (Pp. 101–111).

Yu, V. L., Buchanan, B., Shortliffe, G., Wraith, S., Davis, R., Scott, A., & Cohen, S., An evaluation of the performance of a computer-based consultant. *Computer Programs in Biomedicine*, 9, 95–102, 1979.

Yu, V. L., Fagan, L., Bennett, S., Clancey, W., Scott, A., Hannigan, J., Blum, R., Buchanan, B., & Cohen, S., An evaluation of MYCIN's advice. In B. Buchanan and E. Shortliffe (Eds.), *Rule-Based Expert Systems*. Reading, MA: Addison-Wesley, 1984.

Zuboff, S.; Technologies that informate: Implications for human resource management in the computerized industrial workplace. In R.Walton and P. Lawrence (Eds.), *HRM Trends & Challenges*. Cambridge,MA: Harvard Business School Press,1985. (Pp. 103–139). (a)

Zuboff, S., Automate/informate: The two faces of intelligent technology. *Organizational Dynamics*, 5-18, Autumn, 1985. (b)

# APPENDIX: CONTACTS

The authors would like to express their appreciation to those people who consented to speak with us on the state of the expert systems art:

Dr. Stephen J. Boies, IBM Corporation

Sheldon Breiner, Syntelligence, Inc.

Dr. Stuart K. Card, Intelligent Systems Laboratory, Xerox PARC

Paul D. Christopherson, Corporate Systems Development Division, Honeywell Inc.

Edward L. Cochran, Corporate Systems Development Division, Honeywell Inc.

Dr. Shri K. Goyal, Knowledge-Based Systems Department, GTE Laboratories

Glenn Jimmerson, Robotics and Automation Center, Ford Motor Corporation

Dr. Sara Kiesler, Department of Social Sciences, Carnegie-Mellon University

Rich Kieselewski, American Insurance Group

Dr. Lawrence C. Kingsland III, National Library of Medicine, Lister Hill Center for Biomedical Communications

Frank Landers, Helena Laboratories

Dr. Frederick D. Miller, Artificial Intelligence Systems Department, AT&T Bell Laboratories

Terry M. Miller, Human Resources Laboratory, Wright Patterson Air Force Base

James D. Nelson, St. Paul Fire and Marine Insurance Company

Dr. Allen Newell, Department of Computer Sciences, Carnegie-Mellon University

Dr. Nils Nilsson, Department of Computer Sciences, Stanford University

John Richmond, Helena Laboratories

Dr. John W. Roach, Department of Computer Science and Applications, Virginia Polytechnic Institute and State University

Dr. Eswaran Subrahmanian, Engineering Design Research Center, Carnegie-Mellon University

Dr. John C. Thomas, Artificial Intelligence Department, NYNEX Corporation

Bertil Thorvaldsson, ABB Robotics, Inc.

Dr. David W. Tong, Engineering Information Systems, General Electric Co.

Dr. Mitchell M. Tseng, Applied Intelligent Systems, Digital Equipment Corp.

Dr. David D. Woods, Human Science Department, Westinghouse R&D Center

**Saknar** Du någon TELDOK-publikation som kommit ut tidigare? Eller är det något **särskilt ämne** (inom TELDOKs intresseområde "användning av nya teleanknutna informationssystem i arbetslivet") som Du tycker borde behandlas i TELDOK-serien?

**Fyll i så fall i** den här kupongen och **posta** den till TELDOK (*adress på baksidan*) – eller **faxa** den till TELDOK, **08-713 3588**!

- ❑ Jag vill **beställa** följande publikationer från TELDOK (kryssa för!):
  - ❑ TELDOK Rapport 30: Telematikens Årsbok 1987. Maj 1987.
  - ❑ TELDOK Rapport 36: Omvälvning i televärlden. Optiska sjökablar och konkurrens driver fram ny epok. Juni 1988.
  - ❑ TELDOK Rapport 42: Kontorsinformationssystem i den offentliga sektorn. Ett brittiskt utvecklingsprogram. Oktober 1988.
  - ❑ TELDOK-Info 7: Utsträckt kommunikation – att tänja våra sinnen. September 1988.
  - ❑ TELDOK Referensdokument J: IT i företag och myndigheter – förnyelse eller konservering? Juni 1988.
  - ❑ Via TELDOK 7: Telematik och informationsteknologi i Singapore och Taiwan. Februari 1988.
  - ❑ Via TELDOK 8: Datoranvändning vid handelshögskolor i USA. Februari 1988.
  - ❑ Via TELDOK 9: Intelevent 87. Konkurrens och samexistens. Mars 1988.
  - ❑ Via TELDOK 10: Office Automation Trends in the United States. April 1988.
  - ❑ Via TELDOK 11: Optiska medier. Juni 1988.
  - ❑ **Annan** publikation, nämligen: ____________________
- ❑ Jag vill få TELDOK Rapport och TELDOK-Info **löpande i fortsättningen**, tack!
- ❑ Snälla, **stryk mig** omedelbart från TELDOKs utsändningslista!
- ❑ Jag anser att **någon bör skriva** en TELDOK-rapport om följande viktiga och aktuella aspekt av "användning av nya teleanknutna informationssystem i arbetslivet":

____________________

____________________

- ❑ Jag vill **ändra/rätta** adressuppgifterna som gäller mig. Rätt namn och adress är:

**Namn** ____________________

**Organisation** ____________________

**Adress** ____________________

**Postadress** ____________________

Televerket

Tjänste
Taxe perçue
Sverige

Teldok

Televerket

Televerkets Huvudkontor
KP-T H21:12
123 86  FARSTA

Telestyrelsen har inrättat ett anslag med syfte att medverka till snabb och lätttillgänglig dokumentation beträffande användningen av teleanknutna informationssystem i arbetslivet. Detta anslag förvaltas av **TELDOK** och skall bidraga till:

- Dokumentation vid tidigast möjliga tidpunkt av praktiska tillämpningar av teleanknutna informationssystem i arbetslivet.
- Publicering och spridning, i förekommande fall översättning, av annars svåråtkomliga erfarenheter av teleanknutna informationssystem i arbetslivet, samt kompletteringar avsedda att öka användningsvärdet för svenska förhållanden och svenska läsare.
- Studieresor och konferenser i direkt anknytning till arbetet med att dokumentera och sprida information beträffande praktiska tillämpningar av teleanknutna informationssystem i arbetslivet.

**Via TELDOK** är en av de skriftserier som utges av TELDOK. Via TELDOK presenterar bl a obearbetade tillfällighetsrapporter från seminarier, studieresor osv (t ex nr 9 och den här), manus som förtjänar utgivning även om inte skrivarbetet finansierats av TELDOK och engelskspråkiga översikter (t ex nr 10 och den här). Senast har utgetts:

**Via TELDOK** 9. *Intelevent 87*. Konkurrens och samexistens. Mars 1988.
Via TELDOK 10. *Office Automation Trends in the United States*. April 1988.
Via TELDOK 11. *Optiska medier*. Juni 1988.
Via TELDOK 12. *Den automatiserade experten*. En uppsats om expertsystem... Oktober 1988.

Några andra publikationer från TELDOK:

**TELDOK Rapport** 35. *Datautbyte mellan öppna system* (Open Systems Interconnection). Juni 1988.
TELDOK Rapport 36. *Omvälvning i televärlden*. Optiska sjökablar och konkurrens driver fram ny epok. Juni 1988.
TELDOK Rapport 37. *Expertsystem i Storbritannien*. Juni 1988.
TELDOK Rapport 38. *Informationshantering för samhällsservice – slå 80 000 till offentliga sektorn*. Juni 1988.
TELDOK Rapport 39. *Telehamnar – utveckling och trender*. Juni 1988.
TELDOK Rapport 40. *Telematik i Frankrike*. September 1988.
TELDOK Rapport 41. *Digitalisering i Förbundsrepubliken*. September 1988.
TELDOK Rapport 42. *Kontorsinformationssystem i den offentliga sektorn*. Ett brittiskt utvecklingsprogram. Oktober 1988.
**TELDOK-Info** 7. *Utsträckt kommunikation – att tänja våra sinnen*. September 1988.

Publikationerna kan beställas gratis dygnet runt från DirektSvar (f d TeleSvar), 08-23 00 00. Ange rapport**nummer**!

Den som i fortsättningen önskar erhålla skrifter från TELDOK får automatiskt alla TELDOK Rapport och alla TELDOK-Info.

Adressen till TELDOK är:
TELDOK, KP-T, Televerkets HK, H21:12, 123 86 FARSTA
Telefaxnummer: 08-713 3588