AI Magazine Volume 14 Number 3 (1993) (© AAAI)

Computer-Aided Parts Estimation

Adam Cunningham and Robert Smart

In 1991, Ford Motor Company began deployment of CAPE (computer-aided parts estimating system), a highly advanced knowledge-based system designed to generate, evaluate, and cost automotive part manufacturing plans. cape is engineered on an innovative, extensible, declarative process-planning and estimating knowledge representation language, which underpins the cape kernel architecture. Many manufacturing processes have been modeled to date, but eventually every significant process in motor vehicle construction will be included. Significant cost reductions are among the many benefits CAPE brings to Ford.

CAPE is a highly significant system for Ford of Europe in terms of the business needs it satisfies and the corporate acceptance of AI applications:

First, CAPE represents a major investment, with significant person-years of effort spent on predeployment development alone.

Second, CAPE is the first large-scale production expert system to be deployed within Ford of Europe.

Third, cost estimating is a critical business function. With a total annual materials budget of several billion dollars, cost control is at the heart of Ford's business.

Fourth, reducing the lead time for new model programs provides a key competitive advantage. CAPE reduces estimating response time by 50 percent.

Fifth, this system is enormously ambitious. The final system will capture the combined knowledge of estimating experts in all areas of automotive manufacture.

The Purchase Cost-Estimating Domain

Of all the parts that make up a Ford motor vehicle, the majority are actually manufactured by external suppliers, then purchased by Ford. To effectively manage this substantial vehicle cost component, Ford dedicates a whole division to this task. Purchase Cost Estimation and Analysis (PCE&A) employs a large number of estimators in Europe, typically production engineers, each one an expert in some area of vehicle component manufacture.

The estimator is first involved at the design stage for future vehicle model programs. Working from initial engineering drawings, they provide feedback on production feasibility and economic considerations. When a design becomes accepted for a new model, the estimators do an extremely detailed estimation of each component. These estimates form the basis of price negotiations between Ford and its suppliers.

The estimator starts by drawing up a *process plan*, that is, an ordered set of operations, machines, and materials required to manufacture a part. There can be competing methods of producing the components of the part that are dictated by engineering constraints. Different levels of automation are possible. Typically, higher automation yields a lower piece cost but requires a higher investment. The estimator explores the major combinations of possibilities, choosing the plan with the best balance of piece and investment cost to economically achieve the daily production volume. Interestingly, the preferred plan can differ from one source country to another throughout Europe, owing to the differing labor, material, and facility costs.

CAPE is a knowledge-based estimator assistant capable of timely generation, investigation, and costing of alternate production plans from a component description, justifying its decisions with comprehensive technical detail.

To justify a negotiating stance, each operation in the process plan must be specified to a high level of technical detail. Examples of the justification the estimator must provide include type, size, power rating, and operating cost of the selected machine; the constituents of the floor-to-floor cycle time; the raw material specification, quantity, and cost; the power consumption, current, force, lock pressure, linear feed, and rotational speed; and the design and machining cost of investment tools such as broaching, molding, pressing, or casting tools.

To know their subject matter in sufficient detail, estimators must specialize in one particular area of production. Thus, individual estimators are expert in such areas as injection molding of plastics; fabrication by metal pressing; pressure die casting of aluminium and zinc; forging and sand casting; general assembly; fabrication by welding; surface-finishing techniques such as painting and plating; and the vast area of machining, which includes such diverse techniques as turning, milling, broaching, drilling, gear making, grinding, boring, heat treating, straightening, and shot blasting.

An estimator responsible for a given part might not be expert in all the manufacturing techniques required to produce the part. To complete the estimate, the estimator can call on the expertise of his/her colleagues or can compare the design variance of the new part to a known and previously estimated and purchased part. Because skill shortages and economic pressures prevent the replacement of expertise lost through retirement, fewer estimators with less knowledge must produce more estimates faster. Thus, there is less time to investigate alternatives to sufficient depth, resulting in sometimes shallow comparisons to previously purchased part prices and possible propagation of previous errors going unrecognized.

The Objectives of CAPE

CAPE is a knowledge-based estimator assistant capable of timely generation, investigation, and costing of alternate production plans from a component description, justifying its decisions with comprehensive technical detail. The objectives driving the CAPE project are to (1) capture and consistently use a huge wealth of localized pockets of corporate manufacturing knowledge; (2) reduce the time taken to produce detailed estimates and, thus, contribute to a reduction in concept-to-customer lead time; (3) more accurately model manufacturing costs to effectively contain them through improved design and price negotiation; and (4) facilitate simultaneous engineering between purchase cost estimators and designers, that is, to design for cost effectiveness.

Why an AI Solution?

CAPE must possess and effectively apply vast amounts of experiential knowledge and technically detailed data to achieve its objectives. Representing this knowledge in a declarative, rather than procedural, way is vital to the clarity and maintainability of a system of this size. Expert system technology lends itself to the management and application of such a base of knowledge.

CAPE must perform a heuristically guided search to find an optimal solution. Combinatorial explosion would make an exhaustive search infeasible. Sophisticated AI techniques of declarative constraint description and propagation are required to prune the search space and direct its navigation.

To accurately predict and justify costs to the required level of detail, CAPE must effectively simulate the manufacturing environment with all its interacting agents. Object-oriented modeling is the natural choice for this kind of real-world simulation.

Previous Ford projects to automate the estimating function with non-AI techniques have resulted in MRM (machine rate manual) and COESY (the common estimating system). MRM is a database of thousands of manufacturing machines with technical descriptions and operating rates in different currencies per minute (based on an economic model of purchase cost, lifetime, depreciation, operating expenses, labor skill and level, and so on). COESY is a spreadsheet-like tool for documenting and summarizing estimator-generated processing plans that converts machine cycle times and material uses to cost.

Although both projects were successful, they only perform a limited part of the estimator's work and do not satisfy any of CAPE's objectives. CAPE incorporates the functions of both MRM and COESY.

Operational Functions

The estimator communicates with CAPE through a window-based textual and graphic user interface. The estimator first describes a part in an estimate context; CAPE executes the estimate; and the estimator examines the resulting output, modifying the results or further constraining CAPEs choices if necessary. Estimates and parts are then saved to a database for later retrieval and the results communicated to engineering and supply.

Describing the Part to CAPE

Figure 1 gives an overview of the major windows available for describing a part to CAPE:

The estimator starts with the economics of the estimate. The source country (for example, Germany), the price year (for example, 1992), and the daily production volume (for example, 2000 parts each day) are entered. Different volumes can result in completely different manufacturing plans being generated.

Next, the estimator describes the part, which is an assembly of standard parts and components. *Standard parts* are small items bought in bulk for a set price (for example, nuts and bolts). *Components* are atomic-manufactured items. The estimator tells CAPE which components and standard parts form subassemblies and which subassemblies make up the part. The estimator can tell CAPE how the assembly is performed (for example, spot welding), or CAPE can infer certain assembly operations from context (for example, the presence of screws implies a screwing operation).

The estimator now details each component. This process involves describing the features to be manufactured and any material specification imposed by the component designer. Where materials are only partially specified or not specified at all, CAPE chooses the most appropriate material. Standard parts also have a material specification.

Molded, pressed, and cast features are complex combinations of contours that are created in one shaping operation. The estimator describes these features to CAPE using qualitative measures of shape complexity established from critical known cost drivers rather than geometrically exact measurements (figure 2). Machining features, however, are simpler surfaces that are typically made by cutting material away. The estimator describes the geometry of these features to CAPE.

Surface-finish specifications, such as painting, powder coating, or zinc plating, are also features. The estimator describes the exact nature of the surface-finish requirements to CAPE.

In addition to components, subassemblies and standard parts can have machined features or surface-finish specifications. When the estimator completes a description of the part and all its constituents, CAPE is ready to execute the estimate.

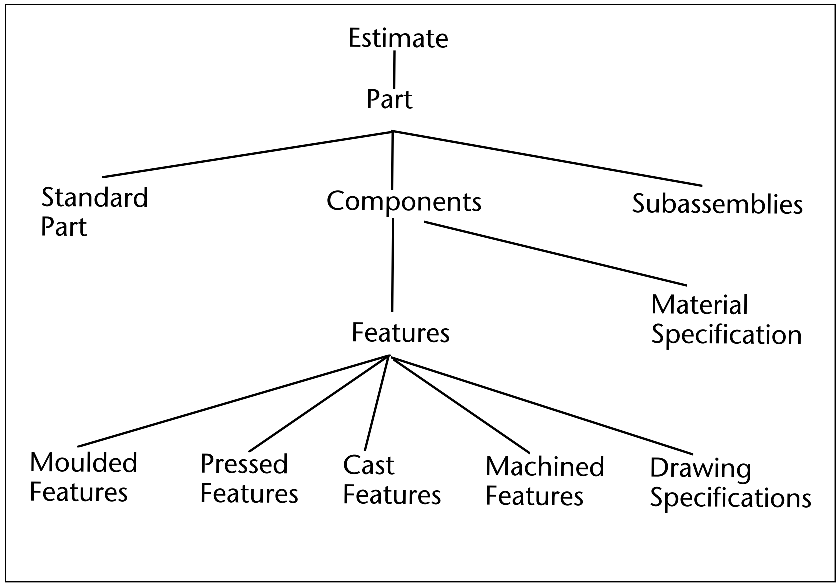

Figure 1. Overview of Parts Description Windows.

Executing Estimates

CAPE now analyzes and classifies the features of each component, considering feasible operations for their manufacture. As a result, a number of competing process plans are generated. Each of these plans is costed, and the decisions (degree of parallelism, machine selection, material selection, and so on) that result in the most cost-effective plan are retained. Along with the best plan, CAPE also prepares justifications for the decisions it has taken and looks for potential risks and opportunities to bring to the estimator's attention.

Risks indicate proximity to physical constraints, such as maximum machine power rating, and opportunities indicate measures that could be taken to reduce cost, for example, extending working shifts by 15 minutes to reduce the number of machines required to satisfy the daily production volume.

Examining the Output

The estimator now uses a number of visual tools to examine the results that CAPE has produced:

Estimate window: The *estimate window* shows the overall cost for each component and for each subassembly.

Expanded estimate window: The *expanded estimate window* (figure 3) shows what operations are in the plan for each component, what machines have been chosen for each operation, how long each operation takes,

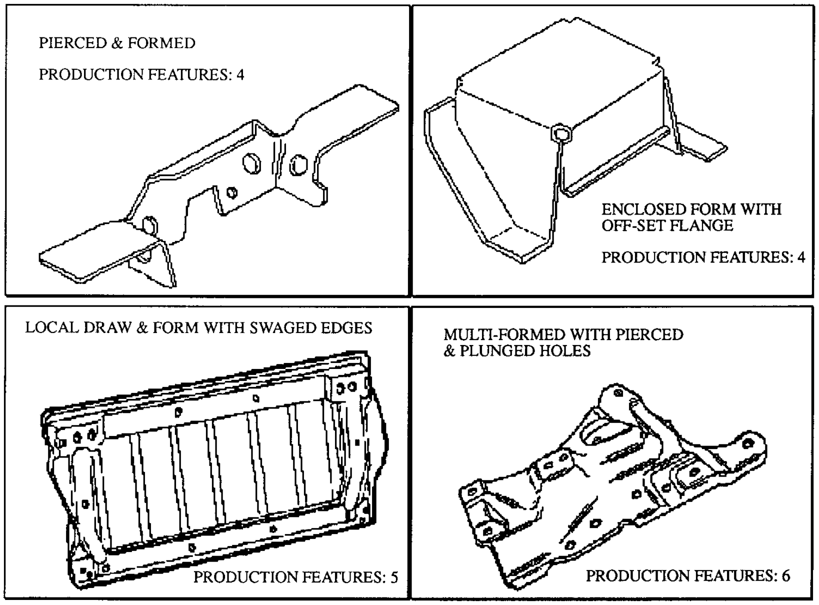

Figure 2. Complex Feature Shape Selection.

Figure 2. Complex Feature Shape Selection.

what manning level was chosen for each machine, and so on.

Plan network diagram: The *plan network diagram* shows what the overall structure of the plan is, which features are being made by which operations, which operations are performed in parallel, and so on.

Part network diagram: The *part network diagram* (figure 4) shows the configuration of components, standard parts, and subassemblies in the part.

Operation detail window: The *operation detail window* (figure 5) shows operation-specific data, such as how the cycle time was derived, what the breakdown of the investment tooling cost is, what manipulative movements the operators are performing, how long these movements take, how many robots are being used, and what the power requirements for the machine selected are.

Risks and opportunities window: The *risks and opportunities window* shows the opportunities and risks associated with the plan.

Expert System Architecture

In the CAPE architecture, there is a clear separation between the kernel and the process models (figure 6). The *kernel* provides generic support for the object model, search mechanics, and plan costing. *Process models* hold specific knowledge for operation definition, possibility generation, plan formulation, and operation costing. The process models developed to date are feature assignment, injection molding, metal pressing, pressure die casting, assembly, welding, surface finishing, turning, milling, drilling, broaching, deburring, shot blasting, shot peening, linishing, impregnation, pressure testing, inspection, grinding, degreasing, and crack detection.

Kernel Architecture

This section describes the kernel architecture, including the object model, search mechanics, and plan costing.

Object Model CAPE uses object-oriented modeling techniques to represent real-world and abstract objects. The abstract object classes lie at the heart of the innovation behind CAPE.

First, *real-world objects* are represented by instances of part, component, feature, tolerance, machine, material, and price classes. Each of these classes has numerous subclasses, giving a rich hierarchy of about 200 classifications.

Planning and estimating objects inherit from classes such as estimate, plan, step, and temporal plan step combiners (for example, serial, parallel, pipeline).

User interface objects are derived from classes such as window, button, menu, row, and field. There are over 300 CAPE-specific user interface objects.

Abstract objects declaratively represent domain knowledge and search control strategies. These classes include feature classifications, operations, possibilities, possibility generators, search heuristics, cost models, and constraints.

In total, there are over 1500 CAPE generic object classes. Many thousands of instances are created dynamically.

Search Mechanics The CAPE kernel performs a constrained depth-first search over a dynamically changing search space. This search space consists of process-specific planning alternatives. To shield itself from process-specific knowledge, the kernel provides a generic searching interface to which all process models conform:

DEFOPERATION is a macro for declaratively defining an operation's choice sets (for example, machine, number of spindles, level of automation), choice-set generators, and operation-specific attributes. It also provides a link to constraints, possibility generators, and cost models.

DEFCONSTRAINT is a macro for declaratively defining constraints on operations and their choice sets. Choice-set constraints are used in database retrieval for choice-set population.

DEFPOSSGEN is a macro for declaratively defining choice-set selection behavior (possibility generation). It also provides a link to process-specific planning modules.

The kernel searches by calling the possibility generators at each operation node in the evolving plan, applying constraints as it goes. The possibility generators cause further plan branches to be built with new operations, which, in turn, have their possibility generators invoked.

Plan Costing The searching mechanisms result in competing plan fragments being generated. These plan fragments must be costed comparatively, taking both manufacturing and tooling investment costs into account to further direct the search. To again shield the process-specific knowledge from the generic mechanism, the kernel provides a costing interface:

DEFCOSTMODEL is a macro for declaratively defining which factors are significant cost contributors and how they are combined to produce piece and investment costs.

Daily Production Volume	1000	User Country Codes	FOB	Part Number	93BB F6129 AKW
Economic Level	01-JAN-93	Estimator Code	M121	Eng. Level	EGB1E10226991000
Source Country Currency	GBP	Vendor Code	BSCST	Description	BRAKE ASSEMBLY

NO	MT	DESCRIPTION	MRM CODE	MACH USAGE	MAN LVL	MAT USAGE	MAT RTE	LAB RTE	OVH RTE	SCP	A+C	MAT COST	MAN COST	TOT. COST	TOOL COST
23	1	WEDGE PLATE													
		CS2: 41X61X5.5				0.108	1.140			3.0	13.0	0.143		0.143	
		Individual Presses													
	1	GRIP FEED MACHINE	50110	0.010	0.00			0.000	0.329	2.0	20.0		0.004	0.004	
	1	(63T) BLANK (S)	51122	0.010	1.00			0.874	0.431	2.0	20.0		0.017	0.017	12475
	1	ZINC PLATE PLANT	213001	0.018	1.00			1.579	3.076	2.0	20.0		0.102	0.102	
		SUB TOTAL										0.143	0.126	0.266	12475
24	1	INERTIA DISC													
		UNE F-ZNALCU4-1				0.013	2.830			2.0	13.0	0.042		0.042	
	1	MELTING ELECTRIC	450201	0.020	1.00			0.027	0.114	3.0	20.0		0.003	0.003	
	1	FRENCH DAW 5	150101	0.047	1.00			0.916	0.212	2.0	20.0		0.065	0.065	
	1	4 CAVITY DIE													31396
	1	TRIM COINING PRESS	430507	0.058					0.161	3.0	20.0		0.012	0.012	
	1	CLIP TOOL													2972
	1	ROTO FINISHER	201203	0.005	1.00			0.874	0.162	2.0	20.0		0.006	0.006	
	1	DIP WASH MACHINE	211102	0.001	1.00			0.874	0.251	2.0	20.0		0.001	0.001	
	1	CHROMATE PLANT	212110	0.001	1.00			1.086	1.440	2.0	20.0		0.003	0.003	
		SUB TOTAL										0.042	0.090	0.132	34368

Figure 3. Expanded Estimate Window.

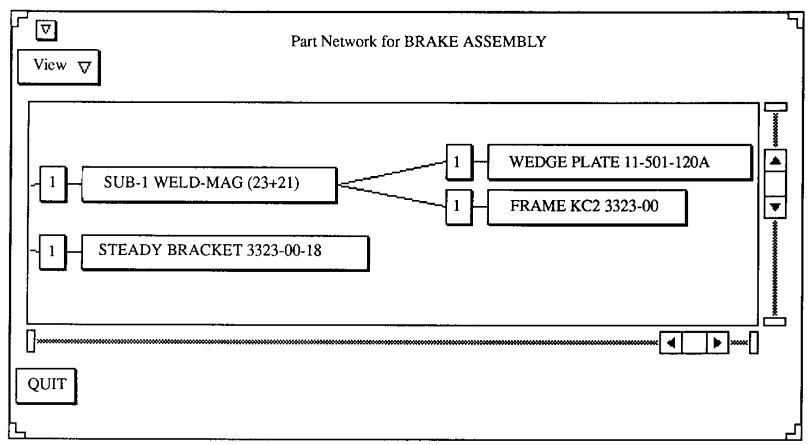

Figure 4. Part Network Diagram.

Press Details		
51122 Double Sided Power Press	63.0 Tonnes	140.0 Strokes per minute

OPERATION	TIME
Handling (from strip)	0.000
Press Cycle (for blanking)	0.009
Total Cycle Time	0.009
Labour Efficiency Factor	17.6%
Costed Machine Usage	0.010

TOOL/OPERATION	MRM CODE	MAT USE	MAT RATE	LAB USE	LAB RATE	MAT COST	LAB COST	TOTAL
BLANK (SIMPLE)								
MATERIAL		39.91	5.40			216		216
DESIGN	600108			14.18	136.88		1941	1941
ROUGH MACHINING	600614			28.36	86.95		2466	2466
PRECISION MACHINING	600414			23.64	98.12		2319	2319
FITTING	600101			23.64	80.78		1909	1909
TRYOUT	600501			4.73	110.97		525	525
SUB TOTAL						216	9160	9376
GUAGING								2000
DIESET								1099
TOTAL								12475

Figure 5. Operation Detail Window.

Each operation in any plan fragment will have a cost model, which the kernel invokes with the method COST applied at any level of the plan.

Process Model Architecture

In this section, we discuss the architecture of the process model, including operation specification, possibility generation, plan formulation, operation costing, and feature assignment.

Operation Specification The process model's knowledge is represented in operation and constraint definitions. These are defined using DEFOPERATION and DEFCOSTMODEL, respectively. There are two distinct levels of operation:

First, *pseudooperations* represent a particular machine performing this operation on, possibly, many features of a component. The principle choice set is machine selection. Pseudooperations are linked to a specific possibility generator and a cost model.

Second, *leaf operations* represent one of the possibly many features being made under the pseudooperation (for example, milling cut, hole tapping). There is usually no choice set or any possibility generator. Leaf operations are linked to a specific cost model (which is different from the pseudooperation cost model).

Possibility Generation Each process model must define how the choice sets are used. A possibility generator can return each combination in the choice sets cross-product as a separate possibility (where this is a manageable number), or it can optimize the choices, returning only a few key possibilities. The kernel tests each possibility for feasibility using declared and propagated constraints.

Plan Formulation For each valid possibility (a set of instantiated choice-set selections), the process model generates a plan containing pseudooperations and leaf operations. This plan represents the manufacture of the component given the choices made in this possibility. Plan generation is often optimized by reusing generic plan templates rather than recreating plans for each possibility.

Operation Costing The process-specific cost models, defined using DEFCOSTMODEL, compute the piece and investment cost of this operation in the context of the current possibility. For example, on a particular machine, a specific operation takes a particular number of seconds that at this machine's operating rate costs a certain amount.

Feature Assignment In between the generic searching mechanisms and the knowledge of specific processes, there is feature knowledge. *Feature knowledge* dictates which processes are appropriate for the manufacture of which features. Indeed, some features can be made by any one of several processes. These processes are known as competing processes because CAPE must decide which to include in the plan.

Feature assignment involves classifying features, eliminating some operations from each class, generating competing plan branches with alternate assignments, and adding conditional operations where appropriate.

Feature classes hold knowledge of different classifications of feature. Membership criteria, competing operations, and operation elimination tests are held declaratively.

Primary feature classes and *secondary feature classes* dictate feature-assignment priority. Some features can be assigned to processes independently, but others require related assignments to be done first.

Conditional operation classes hold knowledge of other processes that might be inferred in addition to those that make the feature. These classes can improve the surface finish or remove burrs, for example.

Constraints, propagated throughout the plan, remove plan branches in conflict.

Implementation

CAPE's object system is built on top of the common Lisp object system (CLOS). A number of desirable features were added, such as support for automatic inverse relationship maintenance and class-instance registries. The entire expert system was written in Lisp, making much use of object-centered knowledge representation through methods.

Inference's automated reasoning tool (ART) was used in a number of areas. ART WINDOWS is the basis of the user interface. The ART notification system is used to ensure consistency between the kernel objects and the user interface objects. The ART iteration package and the ART garbage-free programming package are used throughout the system for garbage-reduced code performance. An extended ART SQL interface is used for data transfer between the expert system and the local database.

Integration Issues

The successful deployment of CAPE to a wide user base required that the system be integrated into the Ford business environment using existing platforms and linked to corporate IBM databases. In terms of development effort, this task has been as great as the development of the expert system itself.

A local ORACLE database was developed in parallel with expert system development. This local database stores all extensional data used by the expert system. It also serves as a buffer between the corporate databases and CAPE itself. A mapping layer developed in Lisp translates data from the expert system to database format and from database format to the expert system.

In addition to the detailed estimating function, estimators are required to report analyses of each estimate to two corporate databases. One database supports design variance, and the other shows *commodity splits,* a high-level economic breakdown of the estimate into commodity groupings. Both these functions were previously performed manually by the estimator but are now performed automatically by CAPE. The result is dispatched from within the CAPE environment.

CASE tools were used extensively for this development. SQL PLUS, CASE DICTIONARY, and CASE DESIGNER were used to design and document the ORACLE database. System analysts working on links to the IBM databases use INFORMATION ENGINEERING FACILITY (IEF) a COBOL-based code generator, as their standard analysis and design tool. IEF is supplemented by a limited amount of C code used in ORACLE-to-IBM links for reasons of speed and efficiency.

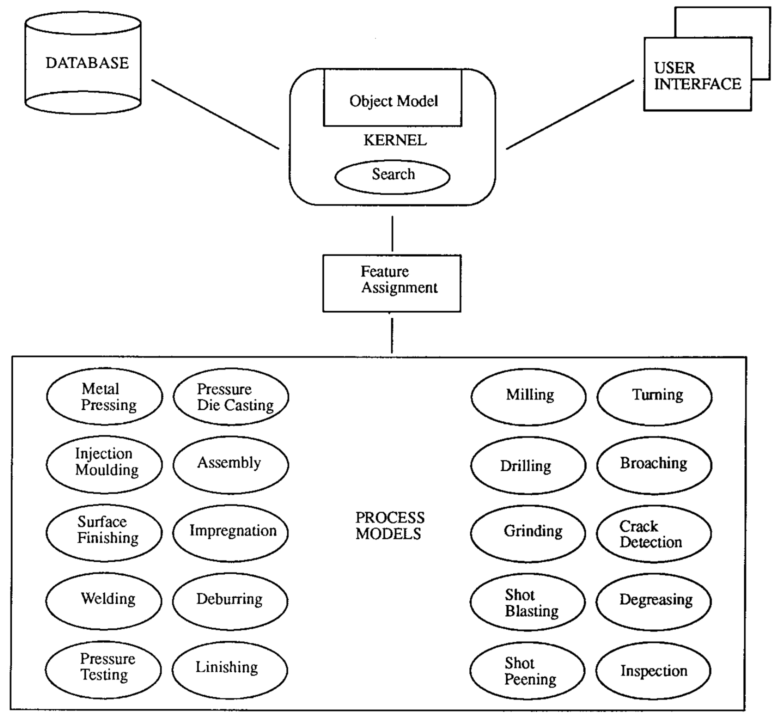

Figure 6. CAPE *Architecture.*

ESTIMATES:	Within 5% of expert	Between 5% & 10% of expert	Over 10% of expert
COESY	54%	38%	8%
CAPE	92%	8%	NIL

Table 1: Results of a Phase One Test.
These results compare the piece costs produced by a group using CAPE with those produced by a group of estimators using COESY.

Innovative Features of CAPE

In terms of AI techniques, CAPE uses a constraint language, constraint propagation, limited depth-first search, heuristics, object-oriented modeling, object-centered knowledge representation, fuzzy logic, generate and test, and simulation.

CAPE could be considered innovative for the mix of technologies, the sheer size of the task it tackles, and the combination of process planning and estimating. However, what makes CAPE unique is its use of abstract object classes to construct a declarative planning and estimating knowledge representation language.

CAPE is an expert system without any rules! Much use is made of Lisp macros that expand to imaginative combinations of object and method definitions, providing a high-level descriptive, application-specific language of DEFCOSTMODEL, DEFCONSTRAINT, DEFPOSSGEN, and DEFOPERATION.

This descriptive language separates the expertise from the way in which it is applied. It is this separation of knowledge from processing (normally attributed to rule-based systems) that allows new process models to be defined and plugged into CAPE easily. The essence of CAPE is an extensible process planning and estimating language.

A major factor contributing to development success was the presence of a team of estimating experts dedicated to CAPE development.

Application Development

Few expert systems of the size and complexity of CAPE have been deployed successfully in industry. To minimize the risk involved in development, two separate pilot efforts were carried out. Each effort was successfully completed before further development was authorized.

Phase 1 was confined to the part family of air cleaners. This phase entailed developing a process model for injection molding plus some simpler supporting models for metal pressing and assembly. Phase 1 was developed as a stand-alone system on SUN workstations.

Table 1 shows the results of a phase-1 test comparing the piece costs produced by a group using CAPE with those produced by a group of estimators using COESY. The same set of 13 air cleaners was used for both groups. The estimates produced by the best Ford estimator were used as a baseline.

The results of this test proved beyond a doubt the feasibility of using expert system techniques for this family of parts.

Phase 2 had the goal of proving the feasibility of applying these techniques to any manufacturing process. Based on the experience of phase 1, a general representation for manufacturing was developed and applied to the most complex processes that CAPE needs to support: turning; milling, broaching, drilling, and metal pressing; and assembly.

Turning: Although the component representation is reasonably simple because of its rotational symmetry, the process itself uses a highly complex set of parallel or serial steps because multiple tools work on a single component to machine combinations of features.

Milling, broaching, drilling, and metal pressing: This process involves developing a component representation to qualitatively describe complex, three-dimensional manufactured features. Graphic support for these processes allows features to be described through icon selection.

Assembly: Constructing an assembly setup involves deciding between using a flow line or stand-alone assembly benches, and line-balancing operations between benches to achieve optimum throughput.

Phase 3 covered the immediate predeployment and postdeployment stages. The focus of the project team shifted from proving expert system feasibility to improving the robustness and efficiency of the phase-2 system. To this end, formal bug-reporting and release methods were set up, and a rigorous program of testing by the system and user team was performed. Additional process models for welding, pressure die casting, surface finishing, and a number of ancillary operations were developed in preparation for deployment.

System development was carried out by a joint Ford-Inference team. The structure of this team was relatively unusual in that it involved both knowledge engineers and system analysts from different groups within Ford under a common management structure. Four Inference and four Ford personnel worked full time on expert system development throughout phase 2 in conjunction with an ORACLE database administrator and five business system analysts working on integration with the corporate IBM databases and electronic links with material suppliers. User interface development required two full-time employees throughout the project. As would be expected for a team of this size, large development costs were incurred.

A major factor contributing to development success was the presence of a team of estimating experts dedicated to CAPE development. This team comprised four estimators and an estimating manager working full time on functional specification, knowledge acquisition, and system validation and testing.

Knowledge Acquisition

The initial knowledge acquisition for each manufacturing process typically involved a member of the user team documenting his/her knowledge of a process, its areas of applicability, and the constraints on its use. This documentation played an important role in triggering the extraction of knowledge during subsequent interviews with experts. Following these interviews, process modelers (that is, CAPE developers) gained an understanding of the process in a number of ways:

First, visiting the suppliers who use the process and those who provide the machines, materials, and tools involved enabled developers to gain direct experience of the manufacturing environment.

Second, visiting industrial research institutions and consultancies that specialize in the process gave Ford access to pools of expertise beyond itself and helped ensure that the most up-to-date and complete knowledge is incorporated into the system.

Third, they attended engineering training courses on the theory and practice of the process concerned. These courses were generally attended with the expert. The physical proximity of the user team, located on the floor above the development team, helped ensure that user feedback and knowledge refinement were continuous processes.

CAPE Validation

Validating the results produced by CAPE posed some difficult problems in terms of methodology. This difficulty was in part because Ford's relationship with component suppliers—the actual cost to the supplier of manufacturing a part is never known—and in part because of differences of opinion between estimators. Thus, there is no true objective measure of CAPE's performance.

The methodology chosen consists of comparing the results and justification produced by CAPE to that of the best Ford estimator for each process over a wide range of actual parts. If the estimator using CAPE is prepared to justify an estimate during negotiations with suppliers, and the estimate is not higher than one that would have been produced by the estimator alone, then CAPE's estimate is considered accurate.

Validation of results is performed entirely by the user team, which decides the test suite of parts covering each process and makes the final decision about when a process model should become part of the deployed system. Testing each new version for robustness before it is released to the user team is performed by the system team on a battery of existing estimates. The initial testing of each new version can be performed automatically. Each estimate in the test suite is automatically retrieved, costed, and deleted, and a report is created detailing any errors produced and the time spent on estimate execution.

Deployment Process

A step-by-step approach to the deployment of CAPE was adopted to achieve the maximum return on investment. As the process models needed to support new part families are deployed, the estimators covering these part families migrate from using COESY to using CAPE. The transfer to CAPE of the estimates in the control of these estimators is performed in part automatically from existing systems (for financial information only) and in part by the user team that defines the components needed in the expert system format. Training new estimators in the use of the system is done entirely by the dedicated user team.

Deployment has involved taking a number of actions to increase the efficiency of the system and decrease hardware costs. The main software change for deployment involved porting the underlying object system from ART schemas to CLOS. Although ART was used for development, the deployed system did not make use of much of its expressive power, namely, the rule, pattern-matching, and viewpoint systems. Using the ART configuration script enabled us to build a tailored version of ART without these functions and, thus, reduce the size of the deployed system. The consistent use of an interface layer to the object system implementation allowedthe conversion from ART schemas to CLOS to be made transparently to CAPE developers.

CAPE is a major technical achievement that proves the viability of using AI technology to solve real-world problems in an increasingly competitive environment.

Application Use

The majority of estimator time is spent on detailed estimating, analyzing, and reporting. Each of these functions is now performed entirely using CAPE as the standard day-to-day estimating platform and interface to other financial systems.

As outlined earlier, the main users of CAPE are European estimators. However, discussions with senior estimating management in North America are under way to explore CAPE's use by U.S. estimators. Initial discussions with the manufacturing engineering and product development areas in Europe also indicate a high potential for CAPE use directly by engineers as a cost-control tool during the design process. Eventually, CAPE might be used widely, beyond the estimating community.

Benefits to Ford

A detailed cost-benefit analysis of CAPE produced a time adjusted rate of return on investment (TARR) achieved through the following:

First is the increased speed of estimating and analysis using CAPE, which reduces estimating time by 50 percent. It also gives PCE&A the ability to increase support for new business practices.

Second is the improved control over tooling costs. CAPE automatically provides detailed tooling costs for every estimate, which currently require extra work and are only produced with CAPE's accuracy and depth to support studies.

Third is the potential vehicle cost savings. More detailed and consistent piece cost estimates, faster online response to queries, quick evaluation of alternatives, and identification of opportunities and risks all improve product development decision making by providing timely cost information.

The improved quality and consistency of CAPE estimates has been demonstrated in a number of ways. In the hands of less experienced estimators, piece cost savings of as much as 30 percent have been recorded compared to the cost the estimator would have chosen. Because CAPE was able to justify the cost given to the level of detail needed to support a negotiation with a supplier, and this justification was supported by the estimator who was the expert in the field, the user was confident in accepting the results.

CAPE has also picked up design inconsistencies that had not been noticed previously. One of these inconsistencies involved an infeasible welding design that had been repeated over a number of years. This design problem had been corrected by the supplier who manufactured the parts, but the information had never been fed back to the design engineers who were responsible.

CAPE has also provided a number of less easily quantifiable benefits to Ford: First, this expert system is the first deployed expert system produced by Ford of Europe, and the experience has been invaluable in terms of skills gained. Primary among these benefits has been the building of a mature in-house expert system team with experience in building and successfully deploying a large and complex system. CAPE serves as an excellent training tool for new estimators, allowing them to be productive earlier than was previously the case.

Lessons Learned

The process of developing CAPE has been a long and, at times, painful one. We feel that the following lessons have been learned from this experience:

First, although new technologies are involved in developing AI projects, standard project management and software engineering techniques are vital to their success, which has been the case with CAPE from day one.

Second, the continuous involvement of a permanent user team was a necessary condition for success. The users have been the driving force behind knowledge acquisition and development throughout the project.

Third, communication between members of the project team is paramount to successful implementation. Weekly conference sessions during early development kept the team focused and helped ensure that team members shared a common conceptual model of the problem domain. This lesson was particularly important in a project of this complexity.

Fourth, plan for change. As our understanding of the problem domain evolved, design modifications emerged that allowed for greater generality and reduced overall complexity. Time needs to be allowed for incorporating such enhancements throughout the project life cycle.

Maintenance

CAPE is expected to hold the most recent financial and manufacturing knowledge. Therefore, maintenance is an ongoing process as new manufacturing processes, machines, and materials are developed and as prices change with time.

The maintenance of CAPE splits broadly into two key areas:

Knowledge base maintenance: Knowledge base maintenance is conducted by a team of knowledge engineers. The experience of Ford estimators and experts in the industry has been that each process changes little over time. Therefore, we expect that little maintenance of existing processes will be needed. The main work of this team will be to add new processes to the system as manufacturing practice changes. An example of this area in the past has been the increasing use of plastics in automotive manufacture for parts families such as bumpers. Future developments in this area include the increasing use of composite materials in vehicle bodies, which will entail the development of new process models. CAPE has been designed explicitly to support such extensions. The generality of the kernel architecture and the modular, plug-in nature of process models allow new manufacturing knowledge to be integrated easily into the system.

Database maintenance: Database maintenance is mainly concerned with extensional data used by process models. Price information for machines and materials was maintained on MRM prior to the deployment of CAPE, and this function has continued unchanged. This price information is updated either yearly or quarterly depending on volatility. Technical material data are maintained through direct electronic links to the main material suppliers.

Conclusion

CAPE has had a significant impact on the speed and accuracy with which estimates can be produced. In the long term, CAPE will not just supplement existing business practices but will enable new ones to be developed in the critical areas of cost control and new model development. It is a major technical achievement that proves the viability of using AI technology to solve real-world problems in an increasingly competitive environment.

Acknowledgments

The authors would like to thank Jeff Greif and David Coles of Inference Corporation for their invaluable contributions to the initial specification and design of CAPE. We would also like to acknowledge the efforts of the following persons for their contribution to the successful development and deployment of this system:

Ford Motor Company: David Buttery, Karen Chadwick, Julie Corpse, Phil Davies, Nasser Faramani, Martin Hodgson, Sean Keeler, Leslie McDowell, Tina Proietti, Guy Saward, Julia Sargent, and Neil Stanley

Inference Europe: Andrew Arblaster, Paul Bates, Steve Lindner, Larry Mond, and Mike Stoler

Finally, special thanks are owed to the user team for never taking no for an answer: Bernard Lees, Vince Baker, Bill Ewin, David Juniper, and Alan Knight.

Bibliography

Alting, L., and Zhang, H. 1989. Computer-Aided Process Planning: The State-of-the-Art Survey. *International Journal of Production Research* 27(2): 553–585.

Canning, W., & Co. Ltd. 1970. *Handbook on Electroplating*. Birmingham, United Kingdom: W. Canning & Co. Ltd.

Droza, T. J., and Wick, C. 1983. *Tool and Manufacturing Engineers Handbook*. Dearborn, Mich.: Society of Manufacturing Engineers.

Keng, N. P., and Yun, D. Y. Y. 1989. A Planning-Scheduling Methodology for the Constrained Resource Problem. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 998–1003. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Lowe, P. H., and Walshe, K. B. A. 1985. Computer-Aided Tool Cost Estimating: An Evaluation of the Labour Content of Injection Moulds. *International Journal of Production Research* 23(2): 371–380.

Lyman, T. 1976a. Forging and Casting. Ohio: American Society for Metals.

Lyman, T. 1976b. Forming. Ohio: American Society for Metals.

Lyman, T. 1976v. Machining. Ohio: American Society for Metals.

Lyman, T. 1976d. Welding and Brazing. Ohio: American Society for Metals.

Welding Institute. 1990. MAGDATA. Cambridge, United Kingdom: Welding Institute.

Adam Cunningham studied natural sciences and social and political scienct at Kings College, Cambridge before earning a master's degree in intelligent systems at Brunel University. He has worked on Cape as an AI technical specialist with Ford for four years. His interests include Japanese language, artificial neural networks, and blues harmonica.