

**ADDENDUM ON  
MARKET PERCEPTION OF  
INFERENCE CORP. - ART**

**April 1989**



## Outline

- Inference Corp. overview
- ART product history
- User-companies interviewed
- Types of knowledge base applications
- Product overview: Strengths, weaknesses, and satisfaction levels
- Product environments, plans, and directions



## **Inference Corp. Overview**

- 1988 company revenues: Not released by Inference Corp.
- 85 employees
- ART licenses have been sold to about 300 firms



## **Inference Corp. - ART Product History**

- **1985: Introduction on Texas Instruments Explorer and Symbolics workstations**
- **1986: Support for Sun and MicroVAX workstations**
- **1987: Release 3.0 adds technical capabilities**
- **1988: ART-IM Release 1.0, under C**
  - "ART-IM" = "Automated Reasoning Tool for Information Management"
  - "Essentially a new product"
  - No development capabilities, this release
  - Migrate KBs to IBM mainframe, IBM PC, or DEC VAX
  - Optimized for run-time performance
  - Only 3MB virtual-address space vs. 30-80MB in ART
  - For efficiency, excludes Schema (object-oriented KB structuring) and Viewpoints (hypothetical reasoning)
- **1989: ART-IM Release 1.5**
  - Full development and production environment
  - Includes Schema
  - IBM PC only (so far)





**Four of Ten Inference Corp. - ART  
User-Companies Interviewed Are  
Consultants or Software Developers**

AGS Consulting

American Express

Bell Core

Chemical Bank

Gateway Information Services

Grumman

Lockheed

NYNEX

Odetics

Westinghouse



## **Most Inference Corp - ART Applications Are the "Diagnosing/Classifying" Type**

- **About 75% of applications: "Data analysis and interpretation"**
  - Examples: Insurance underwriting, credit, real-time military
- **About half are "connected" to data bases or applications**
- **About half are "embedded" within other applications**
- **Almost none are "standalone"**



## EXHIBIT A-5

The vast majority of reported applications built with Inference Corp. - ART are of the relatively common human-assistance or consultation type, where the knowledge base system helps to analyze and interpret data. Often, this is data entered either off-line prior to the consultation session—as in many insurance underwriting applications—or directly during the session—as in some credit analysis applications. In contrast, many of the military applications call for real-time interpretation of complex data transmitted automatically into knowledge base systems from data-capture systems, for example to assist a fighter pilot during combat.

Significantly, very few of the applications are reported to be "standalone," that is, operating independently. Rather, about half are "connected" to data bases and/or other applications, meaning either that they exchange data with a corporate data base or that they make outputs to or take inputs from other applications. The other half are said to be "embedded" within other applications, meaning the knowledge base system operates essentially as a subroutine within the other application, with control passing back and forth as appropriate during execution of the total system. Note that traditional data processing terminology might call either of these types "integrated"; the distinctions between connected and embedded used here should not be considered hard and fast, but are used because this is the way that users of knowledge base systems interviewed generally report thinking about the integration of their applications.



## **No End Users Said to Build or Modify Applications with Inference Corp. - ART**

- Programmers (not end users) become knowledge engineers
- Knowledge engineers also handle interfaces and DP environment issues





## EXHIBIT A-6

There is no indication from these interviews that the end users of an application built with Inference Corp. - ART are ever the developers of the application, unless they already have a programming background and have a direct need for a knowledge base system. Rather, organizations that purchase ART choose one or more of their in-house programmers for knowledge base application development. Users report that the same programmer also develops interfaces to applications or data bases, and handles any data-processing environment issues that arise during development, testing, or implementation.



**Over 70% of Applications Built with  
Inference Corp. - ART Reportedly Are  
Still in Development, Not in Production Use**



## **Top Strengths of Inference Corp. - ART Are Its Range of Development Functions**

- **A superior set of AI functions for development**
  - "A powerful and superior development environment"
  - Clean syntax for development
  - Good control of inferencing sequences
  - "Prototyping is easy and fast"
  - Can flexibly intermix LISP routines
  - Truth maintenance helps avoid contradictions
- **Objects and semantic nets supported effectively**
  - "Strong object-oriented programming capabilities"
  - Effective hybrid of rules and objects
  - Object-oriented Schema system to structure KB
  - "Semantic nets offer flexible linking of elements"
- **Excellent hypothetical reasoning (with "Viewpoints")**
  - Supports exploring alternate hypotheses
  - "Viewpoints" permits temporal reasoning



## EXHIBIT A-8

Users of Inference Corp. - ART come together in reporting three general areas of ART's development capabilities as top strengths.

First, it appears that many of these users are experienced with several knowledge base development systems, and thus the first overall strength they cite for ART is that it "has the bases covered" as to the overall functions of "AI" (or "artificial intelligence," with the term used in this and other exhibits only as a shorthand for the AI subset of knowledge base systems under study here) that they expect to find. Specifically, they comment on the development environment's power, a clean syntax that makes it easy to work with, and the control of inferencing sequences that ART places squarely in the hands of the knowledge engineer. With fast and easy prototyping, they quickly can show the end user early results that validate or force revision of the initial system's concept, reducing waste of development time. Some users appear to be LISP-language-experienced programmers, and they appreciate ART's flexibility in letting them use LISP routines to achieve results not ideally handled within the ART tool itself. Sophisticated users have sometimes worked in systems without "truth maintenance" functionality, and they report that its presence within ART makes them confident—and efficient—in avoiding internal contradictions within the knowledge base they are developing.

A specific set of knowledge base functions cited as strengths by users of ART are its support of object-oriented programming and its implementation of semantic nets. One user thinks of ART as a "hybrid" of the best of both rule-based and object-oriented knowledge base functionality, while another cites ART's "Schema" system as the key to its object-oriented power in helping to structure knowledge bases effectively. One user finds particularly valuable the way in which elements or objects in the knowledge base can be linked, using ART's semantic net capability.

The other top strength cited by several ART users is the "Viewpoints" function that is used to implement hypothetical reasoning to explore alternatives. Time-based (or "temporal") reasoning apparently is also supported by Viewpoints.





**Other Strengths of  
Inference Corp. - ART Are  
on the Production Side**

- **Fast and efficient production inferencing**
  - "Includes a powerful, fast inference engine"
  - Forward-chains effectively for data-driven applications
  - Good pattern matching helps speed rule firing
- **Range of platforms permits production migration**
  - Applications can be ported well between platforms
  - Performance is fast on the DEC MicroVAX
  - "Well-integrated with the Symbolics workstation"
- **"Excellent graphic interface for developers"**
- **"Customer support is excellent"**
- **Can embed expert system within another application**



## EXHIBIT A-9

The other strengths cited for Inference Corp. - ART cluster mainly on the production—as opposed to development—side. It appears that one reason that some users have stayed with ART in the face of market-introduction of competing systems is the overall speed of inferencing that it offers. During a production run of an ART application, this means that its powerful, forward-chaining inference engine quickly moves from the data inputs that drive execution of the knowledge base to the output of useful inferences, or deductions. One user reports that a built-in pattern-matching capability can speed such firing of the knowledge base's rules when appropriate.

Several users report that they have effectively migrated ART applications between two or more of the platforms supported by Inference Corp. One comments specifically on the production-use speed achieved in the DEC MicroVAX environment. Another favors the Symbolics workstation, and is pleased with ART's integration with Symbolics.

Other user comments are not specific to production efficiency or effectiveness. Several users are particularly pleased with the easy, intuitive system use supported by ART's graphic development interface on the workstation's screen. One cited as a strength of the product the vendor's provision of excellent customer support. Finally, another user focuses on ART's support for embedding a knowledge base application within another standard application.



## **Top Weaknesses of Inference Corp. - ART Are in Interfaces and Development**

- **Lacks built-in interfaces to data bases and applications**
  - Lots of special coding for data base interfaces
  - "We must build our own data base connections"
  - Poor tools for interfacing with other applications
- **Problems with end-user screens**
  - Too many LISP routines needed to build user screens
  - Tools for end-user graphics tough, resource intensive
  - Problems with scrolling end-user graphics
  - End-user graphics are drawn too slowly
- **Hard to see structure of the KB**
  - Not easy to review and maintain very large KBs
  - Details of KB are clear, but structure analysis is tough
- **Open questions and issues remain with ART-IM**
  - Is ART-IM still weak on end-user graphics?
  - ART-IM on PC lacks backward chaining
  - No data base interfaces on ART-IM for PC



## EXHIBIT A-10

A number of users of Inference Corp. - ART cite its weakness in interfacing as the key shortcoming. In summary, they note that they can certainly build interfaces to data bases and applications with ART, but that doing so takes time and effort that would be saved if the system only provided built-in interfaces. Note that this is considered a very significant annoyance by many users, even though they generally have succeeded in working around it.

Similarly, users cite specific issues such as scrolling problems and speed-of-drawing problems as shortcomings of ART's facilities for building and executing end-user screens. Again, they can succeed ultimately in most cases in achieving the end-user screen effect they want, but they are annoyed with the resort to LISP routines, with the difficulty of graphics development, or with the substantial computer resources required to do so. It's as if they are saying, "There must be a better way to do this!"

It is interesting that some of the same ART users who praise its power as a development system also realize that it suffers a bit from the "forest versus trees" syndrome. In its detailed development capabilities, apparently ART fails to provide enough power for viewing the structure of the knowledge base as a whole. With the details clear but limited functionality to support effective analysis of the structure, it reportedly is hard for users to review and maintain some of the larger knowledge bases being built with ART.

Finally, a number of users cite as an ART weakness the new, still limited, and sometimes unclear state of the ART-IM product introduction. There is uncertainty, for example, regarding improvements (if any) that ART-IM makes in the end-user graphics problems cited earlier. On the personal computer platform, ART-IM's lack of backward chaining functionality or built-in data base interfaces are cited as specific weaknesses.





**Other Weaknesses of  
Inference Corp. - ART  
Are Mainly in Development**

- **Other development issues**
  - Hard to develop generalized—not problem-specific—solutions
  - Vendor provides few models of how to use the system
  - Incremental compiles take too long
  - No facility to strip out development tools for production
  - Hard to represent knowledge efficiently
- **Object-oriented capabilities need strengthening**
  - Some object-oriented capabilities are missing
- **Added porting facilities needed**
  - ART-IM is still new and incomplete, so porting to realistic production platforms remains a question
- **Use of "Viewpoints" for hypothetical reasoning slows performance**
- **Product price is too high**



## EXHIBIT A-11

Several specific development issues cluster as weaknesses of Inference Corp. - ART. One user misses any development functionality that would encourage more generalized solutions. Another feels Inference Corp. could provide more operational models of effective use of the system. There is a complaint that it takes too long to perform incremental compiles, and that the system could run more efficiently if some or all development tools could be stripped from the system for run-time production. One person reports that knowledge is represented inefficiently within ART's structure.

A few users cite the need for strengthening of ART's object-oriented capabilities.

While porting capabilities for ART (the original product) are fairly clear, ART-IM is considered a new system, so its appropriateness for use as a production platform for ported applications is still unclear.

The Viewpoints hypothetical reasoning capability cited earlier as a development strength comes in as a weakness on the production side, in terms of slowing the performance of an application.

One user feels the price of ART is too high.



## The Main Missing Feature that Inference Corp. - ART Users Want Is Wide-Ranging Data Base Interfaces

<u>Feature of Function</u>	<u>Number of Times Stated as:</u>	
	<u>"Must Have"</u>	<u>"Nice to Have"</u>
<b>Built-in DB interface functions</b>	<b>6</b>	
<b>Improved development functions</b>		
Better KB maintenance tools	1	
Better tools to build end-user graphics	1	
Strengthened object-oriented functions		2
Constraint-based reasoning		1
Ability to integrate with applications in Fortran, ADA	1	
Faster operation, maybe under C	1	
Portability from DEC MicroVAX to PC, Apollo, Sun		1
Faster drawing of graphics		1



## EXHIBIT A-12

In almost all cases, users of Inference Corp. - ART cite as missing features or functions some of those capabilities that were earlier criticized as weaknesses. By far the most consistent user demand is for built-in interfaces to data bases.





## EXHIBIT A-13

As with ART, the product, interviews show a perception of Inference Corp., the company, as a technically sophisticated, superior vendor organization. Its technical superiority in knowledge base functions is tarnished somewhat, however, by its failure to provide ART users with built-in data base interfaces.

Several users note that they see Inference Corp. as having pulled through a corporate "time of troubles." Often cited as a founding member of the "Gang of Four" (four early vendors of high-powered knowledge base tools: Inference Corp., Intellicorp, Carnegie Group, and Teknowledge), Inference Corp. is one of the first-generation vendors that has more or less survived an industrywide shakeout over the past few years. Its current state of relative corporate health in the eyes of users stems from its overcoming of problems encountered during its evolution of ART to support IBM mainframe and personal computer platforms. As reported in user interviews, Inference Corp. about three years ago attempted to "translate" ART into C with a third-party software facility, an effort expected to take less than a year. When that translation failed, a complete rebuilding of ART (resulting in ART-IM using C) was required to fulfill the commitment to support IBM platforms. The 1988 delivery of the initial results of that rebuilding effort came two years later than originally promised when the translation was planned.

As noted earlier, some uncertainties about ART-IM remain (at least as it has been partially delivered so far), but several users report sufficient satisfaction with its early use to suggest that this time of troubles is in the past. One open question, though, is how efficient ART-IM proves in production use, of which there is very little reported so far. A negative set of production-efficiency results would severely compromise Inference Corp.'s second-generation ART-IM product, and perhaps the viability of the company itself.



## **ART Users See Inference Corp. as a Technically Superior Vendor**

- Superior in the ability to deliver many large-application technical AI functions
- Very strong in customer technical support
- Attuned to the need to support IBM platforms
- A bit slow in adding "real world" production capabilities, like data base interfaces
- After company's "time of troubles," a survivor among the early, high-powered tool vendors
- Note on "troubles": Delivery of C-based ART-IM for IBM PC and mainframe was two years late; "translation" to C failed, so rebuilding was necessary



**Most Users of Inference Corp. - ART  
Originally Chose It for  
Technical Superiority**

- Top mention: Strong high-end system, with good forward-chaining capabilities
- Several: DEC MicroVAX/VAX support or plans



**For Half of Current Applications,  
Users of Inference Corp. - ART  
Would Still Choose the Product  
if Starting Again Now**

- For 8 applications, ART would still be the best product:  
"Technically superior"
- For 8 applications, others would/might be better
- Leading alternative (4 of 8): Intellicorp - KEE
- Other alternatives mentioned:
  - Neuron Data - NEXPERT
  - Gold Hill - GoldWorks
  - IBM - Knowledge Tool
  - IBM - ESE
  - AI Corp. - KBMS



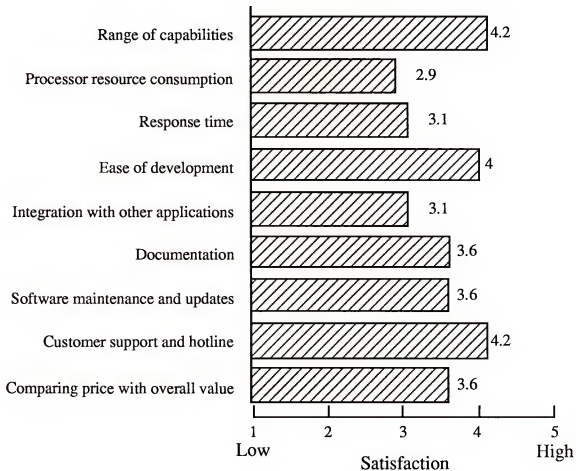


**Users' Overall Satisfaction  
with Inference Corp. - ART  
Is Quantitatively Rated "High"**

- High level of "overall satisfaction": An average rating of 4.4 on a scale where 5 is "very highly satisfied"
- Many comment on overall technical superiority for development
- "Effective forward chaining" is important to many
- A few isolated negatives



## Development Functions and Customer Support Rate Highest in Satisfaction for Users of Inference Corp. - ART





## EXHIBIT A-17

Using a quantitative "1 to 5" rating scale (see the next exhibit and narrative for background on the methodology used), user satisfaction with Inference Corp. - ART is between moderate or neutral and high for all factors rated.

"Range of capabilities" refers to the overall breadth of development and production functions provided by the system, and is relatively high at an average rating of 4.2, although still short of the "top of scale" rating of 5.

"Processor resource consumption" and "Response time" are related yet separate production-oriented factors. The first factor asks the user to judge whether an appropriate or excessive proportion of the production environment's processing power is required to support production use of an application built with ART; in other words, how much of the processor is "left over" for other functions? Given this first factor, the second asks for a subjective judgment as to the acceptability of the response time experienced by a user in the production environment. In both factors, ART ranks moderate (or neutral) in user satisfaction. Users are neither particularly satisfied nor dissatisfied with these product capabilities.

While "Ease of development" overall is rated as high by ART users, they consider the "Integration [of the ART knowledge base application] with other applications" (whether through connection or embedding) to be a separate matter, and rate ART significantly lower on that.

Ratings for "Documentation" of ART and for "Software maintenance and updates"—the effectiveness and timeliness of bug-fixes and new versions of ART provided by Inference Corp.—are both between moderate and high, suggesting that users see some room for improvement but are not really dissatisfied.

In contrast, "human" support for users, referred to here as "Customer support and hotline," is clearly rated at a high level by ART users.

Finally, when "Comparing price with overall value," ART users fall between moderate and high, reflecting high satisfaction with the product's capabilities but some concern as to the relatively high product price.



**Data-Distribution Backup  
for Specific Satisfaction Ratings:  
Inference Corp. - ART**

Range of capabilities

dist 4 4 4 5 x 4 4 4 4 5  
total = 38 #R = 9 avg. = 4.2

Processor resource consumption

dist 2 2 3 4 x 3 x 2 3 4  
total = 23 #R = 8 avg. = 2.9

Response time

dist 3 3 3 3 x 3 x 4 3 x  
total = 22 #R = 7 avg. = 3.1

Ease of development

dist 3 4 5 5 x 3 4 4 3 5  
total = 36 #R = 9 avg. = 4.0

Integration with other applications

dist 2 3 4 5 x 4 3 3 2 2  
total = 28 #R = 9 avg. = 3.1

Documentation

dist 2 3 4 3 x 4 x 4 4 5  
total = 29 #R = 8 avg. = 3.6

Software maintenance and updates

dist 2 3 4 4 x 4 4 4 4 3  
total = 32 #R = 9 avg. = 3.6

Customer support and hotline

dist 4 4 5 5 x 3 4 4 4 5  
total = 38 #R = 9 avg. = 4.2

Comparing price with overall value

dist 2 4 5 4 x 3 3 3 4 4  
total = 32 #R = 9 avg. = 3.6





## EXHIBIT A-18

Two methodological notes are appropriate concerning the collection of quantitative satisfaction ratings displayed here in detail, on which the just-presented average ratings are based.

First, INPUT recognizes that a sample of a maximum of just nine users for each question is not a statistically rigorous study, and should not be interpreted as such. Nonetheless, it is judged significant that in no case here is the distribution exhibited (the line "dist"; see below for details) found to be anything other than a reasonably normal bell-shaped curve, which suggests a substantial level of concurrence among those surveyed.

Second, note that users respond subjectively with their own interpretation of how to position their responses along the line from 1 (stated by the interviewer to represent a "very highly dissatisfied" rating) to 5 ("very highly satisfied"). While such subjective positioning may be expected to differ somewhat between any two users, it seems reasonable to assume that any single user will use the 1-to-5 scale relatively consistently, thus permitting useful cross-comparisons to be made among the factors rated.

For presentation here, each rating-question's data is displayed on three lines.

The first line in each set is simply the title for the question.

The second line (labelled "dist") is the user-by-user distribution of responses on the 1-to-5 scale, confirming the relative normalcy of the distributions mentioned above. An "x" indicates that the user declined to make a rating for that topic; note that one user refused to make any ratings, due to a company policy.

The third line shows three elements. First, the "total" of all rating numbers. Second, the number of respondents ("#R") to be used for developing the average. Third, the average ("avg.") arrived at for the question.



**Nearly All Users of Inference Corp. - ART  
See Both ART and Intellicorp - KEE  
as Leading the "State of the Art"**

- 9 of 10 interviewed: ART is at the "state of the art" in knowledge base systems
- 9 of 10 interviewed: KEE is also
- Others repeatedly placed at "State":
  - Neuron Data - NEXPERT
  - AI Corp - KBMS
  - Gold Hill - Gold Works



## **Today's ART: Interview with Inference Corp.**

- Hardware
  - Workstations: Symbolics, Sun, Apollo, TI, HP
  - DEC: MicroVAX, VAX
  - IBM: Mainframe and personal computer
- Mainframe operating systems
  - MVS/XA, MVS/SP
- Mainframe transaction processing
  - TSO, IMS, CICS
- Standard file and data base interfaces
  - Open architecture toolkit: "Generic Data Integration Facility"
  - Standard interfaces: VSAM, DL/1, DB2, dBASE, Lotus
- Application interfaces
  - IBM mainframe: Interfaces to applications in any standard language, such as Cobol or Fortran
  - IBM personal computer: Interfaces to C applications
- SQL support
  - ART-IM (for IBM mainframe) Release 1.5 (5/89)
- Cooperative processing
  - Not yet
- Top industries and applications
  - Financial: Credit, auditing
  - Manufacturing: Process planning, CAD
  - Military and aerospace: Data analysis
- Pricing
  - ART for workstation: \$30,000
  - ART-IM for personal computer: \$8,000 for first copy, then discounts for more
  - ART-IM for mainframe: \$125,000



**In the Next Releases of  
ART and ART-IM (5/89-12/89)  
Inference Corp. Mainly Plans  
Advances in Graphics**

- **ART-IM Release 1.5 (5/89-7/89)**
  - IBM mainframe (5/89): TSO-like screens, windowing
  - IBM mainframe (5/89): Full KB embedding as a subroutine
  - DEC (7/89): VMS screens
- **ART-IM Release 2.0 (12/89)**
  - IBM PC: OS/2, Presentation Manager
- **ART (not ART-IM) Release 4.0 (12/89)**
  - Support of Unix standard(s)
  - Improved graphics
  - Better performance





**In the Future (1990-On),  
Inference Corp. Plans  
Three Directions for ART and ART-IM**

**1. Continued commitment to two product lines**

ART: High level of AI functionality

ART-IM: Production performance and efficiency

**2. Implement emerging standards**

ART: Unix environment(s)

ART-IM: "Out-SAA IBM"

**3. Support ADA on DEC**

Continue VMS support

Add ADA run-time version of ART-IM on DEC

For ADA, port KB from any source

