Published in the April 1997 issue of the IEEE Control Systems Magazine

INDUSTRY VIEWPOINT

Developing Commercial Applications of Intelligent Control

Stephen Chiu

Intelligent control techniques that emulate characteristics of biological systems offer opportunities for creating control products with new capabilities. In today's competitive economic environment, these control techniques can provide products with the all-important competitive edge that companies seek. However, while numerous applications of intelligent control have been described in literature, few get past the simulation stage to become laboratory prototypes, and only a handful make their way into products. The ability of research work to impact products hinges not so much on finding the best solution to a problem but on finding the right problem and then solving it in a marketable way. In this column, I offer some comments on how to find the “right” problem for intelligent control and what constitutes a marketable solution. Examples from fuzzy logic and neural network based control are used to illustrate the principles.

Look Beyond Familiar Markets

Finding a critical need that can be met by a particular technology is essential to the successful application of the technology. However, the critical needs for control systems vary widely among the different market segments. For example, in the passenger transport market the emphasis is on improving passenger comfort and energy efficiency. In the discrete manufacturing market, ease of setting up, tuning and trouble-shooting control equipment holds high priority. For continuous process control, reducing energy usage, raw material usage and process time are usually on the top of the list. For aircraft and spacecraft, self-diagnostics and fault tolerance are critical needs. The key point here is that a general need, such as “improve accuracy,” is not a critical need in itself; it is only a critical need in the context of a specific market segment. A researcher should recognize the particular capability provided by a new control technique and target the application at a specific market segment where there is a critical need for that capability. Different control techniques meet different needs, and a technique dismissed by one industry may prove to be invaluable to another. This means a researcher must often venture out of his comfort zone, away from familiar industries and applications, in search of the high payoff application. There may be a larger market for the new technique waiting to be tapped.

Experience with applying Fuzzy Self-Organizing Control (FSOC) [1] at Rockwell illustrates this point. FSOC is a fuzzy rule-based form of direct Model Reference Adaptive Control (MRAC) [2] that allows more complex, nonlinear adaptation behavior compared to standard MRAC. FSOC was applied to servo motor control and experimentally compared with auto-tuned PID and manually tuned PID controllers [3]. All controllers produced comparably good performance when the motor load is light. However, when the motor load is heavy (causing actuator saturation), the PID controller could not achieve satisfactory response either by auto-tuning or manual tuning, while FSOC was able to achieve good performance. Although FSOC offers an important capability in dealing with inadequate actuator power, the prospect for FSOC in the motor control market is dim because this capability is not a critical need in that market. In practice, motors are typically selected to be an overkill for the anticipated load (i.e., the load will always be light with respect to the motor power). Conventional PID controller coupled with an auto-tuning algorithm can provide satisfactory performance for this condition; there is little incentive to switch from the simple, cost-effective PID controller unless it is doing an unsatisfactory job. When something works satisfactorily, cost reduction is the driver for change and not performance improvement. On the other hand, the capability offered by FSOC would meet a critical need in the temperature control market, where the available heating/cooling power can never be high enough and existing PID controllers are commonly perceived as too sluggish. It is no coincidence that temperature controllers based on hybrid PID/fuzzy control are now widely available from many manufacturers, while fuzzy control techniques have not had any noticeable impact on the motion control market.

Think Beyond Boundaries of Conventional Control

We have been conditioned by our education to have a narrow view of what constitutes a control problem (i.e., set-point regulation via state feedback), which engenders a narrow view of how to apply broad concepts such as fuzzy logic and neural network in the control framework. If you read a paper on fuzzy control, you will usually see that the inputs to the fuzzy controller are the set-point error and the rate of change of error, and the output is an actuator command or a change in actuator command. When used in this way, fuzzy control is not much different from conventional PID control — it is solving the same problems addressed by PID control and solving them in essentially the same way as PID control, except that fuzzy control provides a nonlinear control law. Similarly, if you read a paper on neural-network based control, chances are that you will see a neural network applied to the set-point regulation problem, usually by replacing a conventional feedback control law and/or plant model with trainable neural network counterparts.

The way that we think about a control problem and its solution (e.g., need a state feedback law, plant model, or state estimator) necessarily influences the way we apply these new technologies (e.g., use a fuzzy/neural state feedback law, plant model, or state estimator). However, looking at these new technologies through the conventional set-point regulation, state feedback control paradigm may not be the most productive way to use these technologies. First of all, we are limiting what these new technologies can do for us when we force-fit them into the narrow roles defined by the conventional control paradigm. Secondly, there are well-established control techniques for set-point regulation problems; applying non-standard, "unproven" techniques to this problem domain requires challenging the establishment, educating the customer to gain acceptance, as well as taking on certain liability risks. If a Linear Quadratic Gaussian (LQG) flight controller caused a crash, we can blame it on unmodeled higher order dynamics ("the model is never perfect"), but if a fuzzy controller caused a crash, the fingers will most likely be pointed at the hapless engineer who chose to use fuzzy control ("why didn't you use LQG?").

A control engineer seeking to apply fuzzy and neural techniques should keep in mind that the set-point regulation problem is not the only control game in town. When we look at commercial products where fuzzy control is said to be incorporated, we rarely see fuzzy logic being used to perform nonlinear PID control; it is used mostly to handle high-level, task-oriented control functions that analytical control methods do not address (e.g., select the cycle time for a washing machine, select the gear for automatic transmission). Consider the subway train control system developed by Hitachi for the city of Sendai in Japan, which was among the first commercial applications of fuzzy logic. For this particular system, the train's acceleration/deceleration is controlled by setting a power lever and a brake lever at different notch positions. Changing the notch position frequently or in large increments creates an uncomfortable ride. In addition to riding comfort, a train operator must consider safety, on-time arrival, energy consumption, and stopping the train accurately at a specified position along the station platform. Here fuzzy logic was used to select the notch position that will best satisfy these multiple, often conflicting objectives. A simple simulation of the train dynamics is used to predict the resultant speed, stopping position, and time of arrival for each possible choice of notch position. Fuzzy rules then rank the desirability of each notch position based on the predicted outcome, taking into account factors such as the safeness of the resultant speed, arrival time, stopping position, amount of notch change, and the elapsed time since the last notch change. The notch position that received the highest ranking is selected as the actual notch command. Here fuzzy rules are not being used to specify the familiar feedback control law, but to rank the different control outcomes in a way that reflects human's sensibility of "optimal". It would be difficult to solve this type of qualitative control problem within an analytical framework.

Many Japanese companies, including Matsushita, Sanyo, Hitachi and Sharp, have incorporated neural network technology into a product known as the kerosene fan heater [4]. Kerosene fan heaters are used extensively in Japan to heat individual rooms in a household. Because these heaters are based on vaporized oil, a long preheating time is required to vaporize the oil before the heater can actually produce useful heat. In Sanyo's heater, a neural network learns the daily usage pattern of the consumer, thus allowing the heater to automatically start to preheat in advance. In Sharp's heater, the user can specify the time at which the room should reach a desired temperature; a neural network learns the required heating time (preheating time plus the active heating time) for the particular room so that the heater will switch on at the appropriate time.

Most products that use fuzzy logic or neural network for control apply these techniques to solve problems that fall outside the domain of conventional feedback control. For the instances where fuzzy logic is applied to the set-point regulation problem, such as in the hybrid PID/fuzzy temperature controllers, fuzzy logic is typically used in a high-level module that supervises a conventional PID controller. The number of ways that fuzzy logic and neural network can be used for control is limited only by your creativity. The creativity and potential applications are stifled when we view fuzzy or neural control as only a nonlinear counterpart of conventional feedback control techniques.

I often hear the question "Why are there so many commercial applications of fuzzy control in Japan and not in the U.S.?" The typical answer from fuzzy control proponents is "Because western culture is steeped in Aristotelian (i.e., true-or-false, black-or-white) logic while eastern culture accommodates fuzziness". This may be part of the reason, but I do not believe that this is the primary reason. Many companies in the U.S. have looked into fuzzy control and have lost their enthusiasm because it has not been productive for them. They have not found the high payoff, hot products as the Japanese did. I believe the main problem is this: U.S. engineers have been led by the technical literature into thinking that using fuzzy rules to implement nonlinear PID control is the prime use of fuzzy logic for control. While Japanese engineers are using fuzzy logic for high-level, expert system-like control functions, the U.S. has been fixated on the fuzzy PID control paradigm. The typical scenario in U.S. companies is that the boss would say to the engineer, "Go look at this fuzzy control stuff and see if there's anything useful." The engineer would come back in a month and say, "Fuzzy control is just nonlinear PID control that may perform somewhat better than conventional PID, but it's much harder to tune and analyze; we're better off sticking with conventional PID." By having a narrow view of how to use fuzzy logic in the context of control, they have missed the high valued-added, high payoff applications.

What's the Payback for the Customer?

Sales of a product are not so much driven by a product's performance but by whether the product provides good payback to the customer for the amount of money invested. Industrial customers are particularly unwilling to pay for better performance beyond what is necessary to get the job done. Controller performance (i.e., time and frequency response) is also difficult to market (i.e., promote in an ad or compare with competing products), because there is no standard benchmark for controller performance. Therefore, improving controller performance is generally not a high priority among product managers as long as the existing performance is deemed satisfactory by the customer. Controller performance is only a priority when performance improvement can make a significant impact on the customer's bottom line. For example, reducing the warm-up time of plastic injection molding machines, without overshooting the desired temperature, translates directly into more time for useful production (temperature overshoot is an abomination in this business because the machine must sit idle while waiting for the temperature to slowly drop back down). Before embarking on a project to improve controller performance, you must understand the performance criteria from the customer's perspective and focus on the specific performance characteristics that can produce a significant payback for the customer.

While not every customer needs improved control performance, everyone is looking for ways to save time and money. There is an abundance of opportunity to add new controller features that save time, labor, raw material and energy for the customer. These features provide easily palpable financial payback and are easy to market. Product managers also like to focus on adding new features rather than on improving performance, because new features provide the product distinction that is crucial in a marketplace populated with similar competing products. Particularly for products that can be considered as commodity items (e.g., personal computers, programmable logic controllers, even kerosene fan heaters), distinctive product features can help a company stay above the price-cutting fray.

We often hear that a particular control technique or feature is not practical for implementation in a product because it would require expensive hardware. What this really means is that the payback for the customer is not enough to justify the increased product cost. Do not be overly concerned with product cost, but focus on how to provide a big payback to the customer; it is the payback that dictates how much a customer is willing to pay for a product.

The auto-tuning capability integrated into today's PID controllers is a good example of a time- and labor-saving feature. The auto-tuning feature eliminates the time-consuming, trial-and-error process of tuning PID gains manually. The customer can probably save $200–$1000 in engineering labor cost each time the PID controller needs to be tuned. Even if the auto-tuning feature adds $100 to the product cost, the feature will pay for itself after only one use. The payback is immediate and easily quantifiable, thus making an auto-tuning PID controller an attractive investment. Unfortunately, the auto-tuning feature is now too common in PID controllers that this feature can no longer provide product distinction nor command a premium price. However, there remain many other applications where time-consuming, manual adjustment of control parameters by a skilled human is the norm and where automating the manual process can produce a significant payback.

At Rockwell Graphic Systems, for example, neural network was applied to automate certain operator control tasks for color newspaper printing presses [5]. Lithographic offset color printing is a highly complex, nonlinear process affected by many variables, including the ink and dampening solution feed rates, press speed, chemistry of dampening solution, press temperature, humidity, and paper quality. Skilled press operators make numerous machine adjustments while watching the resultant prints until acceptable print quality is achieved. In particular, for each printing couple (an assembly that prints a given primary color) the operator must adjust a row of multiple ink keys that control the ink flow across the width of the paper. The longer it takes to make the adjustments, the more production time and material are lost. A fast, automatic system for adjusting these ink keys can both reduce waste and increase the consistency of print quality. Because of the lack of a mathematical process model, neural network was applied to learn to adjust the ink key settings based on color sensory feedback; the neural network can adapt on-line as printing conditions (e.g., temperature, humidity, paper quality) change. This system provides significant payback to the customer by reducing the required skill for operating press equipment as well as by improving production efficiency and quality consistency.

Market Acceptance

Every new technology takes some time to gain acceptance in the marketplace. A company that introduces new technology usually has to educate the market on what the new technology is and what the benefits are, as well as re-engineer the product according to market feedback. This poses a risk in being first to market — the company may invest significant resources to nurture the market only to be later undercut by a competitor who jumps in when the market is mature and when customers' likes and dislikes are known. A winning product is not necessarily the one that's first to market, but the one that's first to market acceptance [6]. A new product that is apt to gain quick market acceptance is one that builds on the customer's familiarity with a previous product. In contrast, a new product that looks unfamiliar and that forces the customer to learn a new skill faces a difficult uphill battle in gaining acceptance. From this standpoint, an advanced controller would have difficulty gaining market acceptance if it requires significant user education and training.

As an example, consider the user-programmable fuzzy controller which was the focus of early fuzzy control product development. It was thought that plant engineers would welcome the ability to program nonlinear, higher performance control laws via simple rules, and that such fuzzy controllers could mount a serious challenge to PID controllers in the marketplace. However, the user-programmable fuzzy controller only made life more difficult for the plant engineer, who must now learn about fuzzy control, learn a new programming environment, as well as fiddle with myriad rule parameters instead of simple PID gains. As a result, user-programmable fuzzy controllers remain today a low-volume product purchased primarily by R&D labs.

The high-volume fuzzy control products are invariably those that can be used like previous products and that hide the details of fuzzy technology from the user; in particular, they add performance and functionality to an existing product without requiring additional user training. For example, both Yokogawa and Fuji Electric market temperature controllers that integrate fuzzy logic with standard PID control to help suppress overshoot. In Yokogawa Electric's controller, fuzzy logic is used to determine artificial set-points that are fed to a conventional PID controller. As the fuzzy module detects impending overshoot, it "fools" the PID controller by telling the PID controller to aim for a temperature value that is somewhat lower than the actual set-point. These temperature controllers look like and operate like conventional PID temperature controllers, except that the user has an option of turning on/off the overshoot suppression feature. In addition to retaining the familiar PID controller packaging, market acceptance of these products in an industry dominated by PID control is accelerated by employing fuzzy logic as an auxiliary helper to PID control, rather than as a replacement.

Concluding Remarks

There are many opportunities to apply intelligent control techniques to create products with a competitive advantage. However, in order to seize these opportunities, we need to have a broad perspective of what constitutes a control problem and be creative in the ways that we use intelligent control techniques to solve the problem — an intelligent control system is not simply a set-point regulator with an "intelligent" state feedback law.

Control researchers who seek to use intelligent control techniques to replace functions already occupied by conventional control techniques will face formidable obstacles in transferring their research results into products. Firstly, such approaches are focused on incremental performance improvements over existing solutions, which may not be important to the customer who already has something that works satisfactorily. Secondly, proponents of conventional control can often achieve the same incremental performance improvement by patching up the existing solution (e.g., increase the size of the gain schedule). Thirdly, product managers are usually unwilling to switch from a tried-and-true control technique to a new, unproven technique unless the new technique offers exceptional advantage at low risk. The risk here not only involves technical feasibility, but also certification and liability issues that arise from using non-standard techniques, as well as uncertainty about market acceptance. For these reasons, it behooves us to focus the commercial application of intelligent control on functions that fall outside the domain of conventional control. Using intelligent control techniques to complement, rather than replace, conventional control techniques will also facilitate the acceptance of these new techniques in the control community, as proponents of conventional control begin to view these techniques as special-purpose tools to be added to their toolbox rather than as flawed substitutes for their existing tools.

The pervasive application of fuzzy and neural control to consumer products in Japan is noteworthy. These consumer applications are aimed primarily at improving human friendliness of machines — making machines easier to use and perform in closer accord with human preferences (e.g., more comfortable vehicle ride, more comfortable air conditioning, better tasting rice). Consumers are largely non-technical people who purchase products to make life more enjoyable; they have little tolerance for complexity, but simply want the product to "do the right thing" all by itself. The auto-focus, auto-exposure, auto-everything camera is a good example of a product that answers this need. Because market demand is driving consumer products towards the ambitious goal of "auto-everything," not simply automatic set-point regulation, there are vast opportunities for the application of intelligent control techniques in the consumer market.

New tools such as fuzzy logic, neural network and genetic algorithm open new possibilities and allow control engineers to extend the functionality of a controller beyond its traditional boundaries. Making a control system "do the right thing" is no longer limited to the small list of "right things" that can be defined via analytic formulae and addressed by mathematical system theory. There is tremendous payoff in finding the new, profitable "right things" that we can now make a machine do by itself.

Stephen L. Chiu was born in Taiwan in 1960 and immigrated to the United States in 1971. He received the B.S. degree in mechanical engineering and nuclear engineering from the University of California at Berkeley in 1983 and the S.M. degree in mechanical engineering from the Massachusetts Institute of Technology in 1985. He joined Rockwell's Science Center in 1985, where he is currently a Research Scientist. He has conducted research in robotics, fuzzy control, and neuro-fuzzy systems, with an emphasis on technology transfer to Rockwell's business divisions. He often teaches short courses on the application of fuzzy and neuro-fuzzy methods to industry and was a consultant to MathWorks Inc. during the development of the Fuzzy Logic Toolbox software for MATLAB. His current work is focused on neuro-fuzzy systems and on managing a "smart" motor development project. He is an associate editor of the IEEE Control Systems Magazine, an associate editor of the Journal of Intelligent and Fuzzy Systems, and a member of the board of the North American Fuzzy Information Processing Society. He has also served as co-guest editor of the Proceedings of IEEE in a special issue on fuzzy logic.

References

[1] T.J. Procyk and E.H. Mamdani, "A linguistic self-organizing process controller," *Automatica*, vol. 15, 1979.

[2] Y.D. Landau, *Adaptive Control: The Model Reference Approach*, Macel Dekker, New York, 1979.

[3] J.J. Cheng, and T. Kong, "Fuzzy self-organizing control - theory and application," *Proc. 7th Annual Rockwell International Conference on Control and Signal Processing*, Thousand Oaks, CA, May 1995.

[4] H. Takagi, "Applications of neural networks and fuzzy logic to consumer products," Chapter 5 in *Industrial Applications of Fuzzy Logic and Intelligent Systems*, ed. J. Yen, R. Lanagri, L. Zadeh, IEEE Press, 1995.

[5] M.S. Lan, P. Lin, and J. Bain, "Modeling and control of the lithographic offset color printing process using artificial neural networks," *Trans. of ASME, Journal of Engineering for Industry*, vol. 116, pp. 274-276, May 1994.

[6] P. Asthana, "Jumping the technology s-curve," *IEEE Spectrum*, vol. 32, no. 6, pp. 49-54, June 1995.