Exomorphic

Eder Medina and Thomas Luly
Draft: October 30, 2024

Problem: For decades, manufacturing processes like machining, injection molding, and tool and die making have formed the backbone of the US industrial base, spurring progress and growth, from putting humans in space, to enabling the growth of life-changing medical technologies, to the scale-up of the country's clean energy infrastructure. Yet many aspects of these manufacturing technologies are outdated and have resisted modernization. They still require significant inputs of time, materials, energy, and labor.

This has resulted in an industry that is stagnant in the US and declining across Europe due to rising energy costs. Critical manufacturing knowledge is being lost as experienced workers retire faster than new ones can be trained. Much of the expertise required to optimize production – often called 'tacit' knowledge – exists in the minds of experienced workers, and is often a rate-limiting factor in industrial processes. When this knowledge is lost, production is slowed (or, in some cases, shuttered altogether) and quality suffers. Moreover, the manufacturing processes are also disjointed and inefficient, resulting in both higher energy consumption and costs — a problem acutely felt in Europe right now. As a result of the erosion of the US and European industrial base, manufacturing labor is outsourced to locations that exploit low labor costs, lack of environmental regulations, and longer, more vulnerable supply chains.

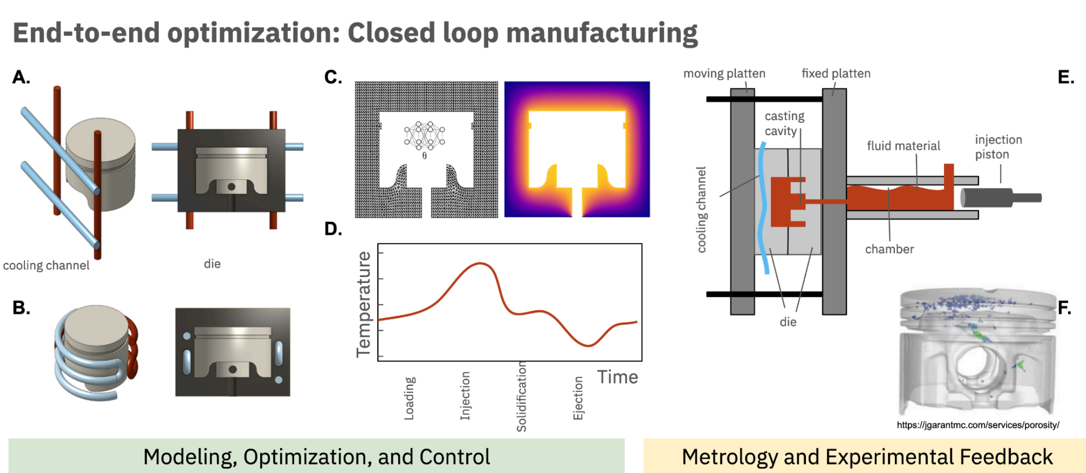

A. Traditional tool and dies are designed with geometric constraints dictated by machining operations while **B.** advances in die cast additive manufacturing enable conformal cooling designs. **C.** Neural surrogate modeling the physical processes undergone during casting enable fine control of **D.** temperature profiles throughout production. **E.** Sensors at each critical step on a per part basis provide continuous data streams for process refinement. **F.** Holistic part inspection provides key information on porosity and mechanical properties; this ensures additional feedback in the die making process.

Vision: To solve these problems, we are building a novel manufacturing system that leverages recent advances in machine learning, simulations, data-driven control, and additive manufacturing. First, we aim to build differentiable end-to-end process automation, allowing entire manufacturing workflows to be optimized holistically. Second, we are developing an adaptive system that learns from expert demonstrations and continuously improves through real physical process monitoring and experimentation. Third, we are creating models capable of simulating complex physical processes, enabling precise control for fast, real-time adjustments. The result is an integrated system for multiphysics-driven manufacturing that bridges the gap between traditional expertise and advanced automation, shortens cycle times and adapts to dynamic conditions, thereby reducing waste, lowering costs, and enhancing the quality of the parts that we produce.

With an initial focus on die casting – an inefficient process hindered by a lack of codified expertise – our solution ensures that critical knowledge is captured, optimized, and continuously applied across the entire production pipeline, from tooling to final product. Ultimately, this approach will transform not just how parts are made but how knowledge is preserved and applied — closing the gap between design, production, and performance.

State of the Art: Manufacturing has historically resisted modernization, though some progress has been made. Today's state-of-the-art manufacturing centers are equipped with extensive sensors, yet much of the collected data remains underutilized. The industry remains fragmented, with critical data often siloed between stakeholders, limiting the ability to implement real-time feedback loops that could significantly enhance processes. Simulations, though used, are still slow, costly, and not well-suited for rapid iteration.

In contrast, other industries have made great strides by leveraging several key advances in software. Lower computational costs and faster simulations now allow for rapid design iteration. Machine learning enables the extraction of expert patterns from data at scale, and differentiable programming combined with end-to-end optimization makes it possible to efficiently solve complex design challenges.

Our approach applies these advances to manufacturing in novel ways. In past projects, we designed mechanical systems whose morphology and control evolved jointly through end-to-end optimization, resulting in more efficient and robust outcomes than traditional methods (Oktay 2023). Recently, we developed a machine learning architecture that models physical systems with guarantees of physical consistency, designed for structural engineering (Pastrana 2024). This method achieved simulation and optimization speeds nearly 1,000x faster than existing approaches, running on standard consumer-grade hardware while remaining interpretable. We have also built differentiable simulators capable of precisely controlling energy flows in complex physical domains (Bordiga 2024). By bringing these tools into manufacturing, we aim to break through long-standing industry bottlenecks—enabling faster iteration, enhanced process efficiency, and systems that continually improve through feedback.

Market Viability: For our technology to succeed, it must reduce unit costs and prove its value in industries that are traditionally slow to adopt new technologies. While we are actively building a techno-economic model, increased efficiency will lead to higher material throughput, translating into lower unit production costs. Taking a die cast part as an example, end-to-end optimization would allow us to materially reduce five of the main cost components: energy (optimized controls systems reduce energy consumption); materials (fewer defects, reduced scrap rates, and improved yields lower costs); tooling (design simplification reduces upfront costs); labor (automated processes minimize human intervention); and post processing (reduces the need for trimming and heat treatment). We are actively conducting interviews, and our initial conversations have affirmed that there is a significant unmet need for solutions that address these cost components, positioning us to offer a superior alternative to the industry status quo.

Milestones and Risks: Our development plan includes several key milestones to guide progress. In the short term, we will build small-scale prototypes to validate our adaptive control systems; these prototypes will allow us to test the feasibility of our approach and refine the underlying technology. At the same time, we will seek partnerships with manufacturers to integrate our early-stage solutions into real-world production environments. We aim to scale this technology into sectors such as automotive, aerospace, and medical devices, where the demand for high-quality, precision components is high.

We will also have to simultaneously manage several risks. The integration of machine learning, simulation, and control systems presents a significant technical challenge, as each component must work seamlessly together. Adoption risk is another concern, as startups frequently face high perceived vendor risk in the eyes of customers. Financing risk is also high, given the need for substantial capex and R&D investment before achieving profitability. Additionally, we must ensure compliance with evolving regulatory standards, particularly in industries such as aerospace and medical devices.

Social Impact: Beyond its economic benefits, our platform has the potential to drive significant environmental and social impact. By optimizing energy consumption, our technology will reduce the carbon footprint of manufacturing processes. Shortening supply chains will further reduce emissions by minimizing the need for long-distance transportation. In addition, improved quality control will extend product life cycles. This shift will reduce dependence on exploitative global supply chains and help revitalize domestic manufacturing, ensuring that the benefits of innovation are distributed more equitably.

Conclusion: Our mission is to revolutionize manufacturing by integrating simulation, adaptive control, and machine learning into real-world production systems. By closing knowledge gaps, streamlining processes, and enabling more energy-efficient manufacturing, we aim to secure the future of critical industries while positioning the U.S. at the forefront of industrial innovation. With the right support, we can launch a platform that addresses both economic and environmental challenges, ensuring sustainable growth for the next generation of manufacturing.

Technical Details: Manufacturing consists of three interconnected phases: design & engineering, production, and metrology. Although often treated as distinct, these phases are deeply interdependent. Design focuses on geometry and functionality but must account for manufacturing constraints and define key dimensions and properties. Production involves process design, material flow, and control of variables such as pressure, flow rate, and temperature, while metrology ensures quality throughout every step and in the final output.

In die casting, the process is particularly complex, requiring precise control over material flow, solidification, and thermal management, and the design of the necessary assembly to generate a desired part. This involves balancing geometry, physics, and control inputs to achieve optimal outcomes. Unfortunately, the complexity makes it difficult for a single person to manage, let alone optimize, every step. However, recent advances suggest that these challenges become manageable with systems capable of end-to-end optimization. These complexities align perfectly with the requirements for developing differentiable optimization pipelines, as seen in other industries.

A solution lies in model-based control and differentiable, end-to-end systems that optimize geometry and control parameters—such as temperature profiles and material flow—concurrently. Traditional workflows are slow, requiring iterative simulations and extensive manual tuning. However, novel architectures, such as neuromechanical and physics-constrained autoencoders, now enable faster and more efficient optimization. These systems rely on amortized simulation and optimization, distributing computational workloads across tasks through differentiable implementations, allowing real-time adaptation and feedback.

These methods, originally developed for computational solid mechanics, are very well-suited to die casting, which primarily involves complex fluid mechanics. Translating these methods to die casting makes logical sense: the governing equations are well-defined, and the geometric design space has expanded with additively manufactured dies, enabling more complex shapes and finer flow path control. Moreover, sensorized die-casting machines and advancements in metrology provide real-time feedback from production parts that further bridging the sim-to-real gap.

These technologies unlock the full potential of end-to-end pipelines, allowing simultaneous optimization of both geometry and process control. This integrated approach promises to transform die casting by delivering exceptional precision, reducing defects, and maximizing efficiency and throughput.