

SIMULATION-BASED OPTIMIZATION FOR CAR T-CELL THERAPY LOGISTICS

Negar Sadeghi¹, and Mohammad Dehghanimohammadabadi¹

¹Dept. of Mechanical and Industrial Eng., Northeastern University, Boston, USA

ABSTRACT

Despite positive clinical outcomes of Chimeric Antigen Receptor (CAR) T-cell therapy, its time-sensitivity causes substantial logistical challenges. This paper introduces a simulation-optimization (SO) framework for optimizing CAR T-cell therapy supply chains, addressing both transportation mode selection and patient scheduling. This framework combines a discrete-event simulation with a metaheuristic optimization to handle uncertainties in processing times and patient conditions. Experiments across different time-window constraints demonstrate that the developed SO model consistently outperforms traditional scheduling heuristics (FIFO, SPT, EDD) in total cost while maintaining timely delivery. This model provides a better balance between transportation efficiency and delay minimization compared to rule-based methods. Results highlight the potential of simulation-based optimization to enhance personalized medicine delivery by improving cost-effectiveness without compromising treatment timeliness.

1 INTRODUCTION

CAR T-cell therapy is a personalized cancer treatment that uses a patient's own immune cells to fight cancer. Doctors collect T cells from the patient's blood and modify them in the laboratory to recognize cancer cells. These engineered cells are then infused back into the patient to target and destroy cancer. CAR T-cell therapy has shown strong results in the treatment of certain hard-to-treat blood cancers, such as B-cell lymphoma and acute lymphoblastic leukemia. (Triantafyllou et al. 2022b). Despite its therapeutic promise, this treatment faces significant operational challenges that limit widespread adoption (Utkarsh et al. 2024).

The CAR T-cell supply chain presents unique challenges, primarily centered on manufacturing strategy (i.e., centralized, decentralized) and transportation logistics. Centralized manufacturing offers cost efficiency but increases transportation complexity, while decentralized approaches provide faster delivery at higher infrastructure costs. The manufacturing process itself is intricate, involving leukapheresis (T-cell extraction), genetic modification (DNA alteration to instruct cancer targeting), cell expansion (T-cell multiplication to reach therapeutic dose) and cryopreservation (ultra-cold freezing for storage and transport). Each step requires quality control to maintain cell viability and effectiveness. Time sensitivity is critical, as delays can significantly reduce treatment efficacy and patient outcomes (Lopes et al. 2020). The high production costs, which range from \$373,000 to \$475,000 per patient, further limit accessibility (Di et al. 2024). Once manufacturing is complete, the focus shifts to transportation logistics, which introduces its own set of challenges, such as timing, cost, and mode selection (Jackson et al. 2016).

To address these key logistical challenges in CAR T-cell therapy, this study proposes the following:

- A two-stage optimization approach tailored for CAR T-cell therapy supply chains.
- A focus on minimizing transportation costs and treatment delays, accounting for the time sensitivity of patient conditions.
- Enhanced operational efficiency in scheduling and transportation decisions, contributing to more accessible and affordable delivery of this therapy.

2 LITERATURE REVIEW

2.1 Optimization Approaches in CAR T-cell Supply Chains

Early studies on CAR T-cell therapy supply chains focused on optimization models to address network design, scheduling, and cost trade-offs. Bernardi et al. (2022) developed an MILP model to optimize the configuration of manufacturing sites, storage options, and transport modes. Their results showed that introducing upstream storage improved facility utilization and significantly reduced total costs under time constraints.

Triantafyllou et al. (2022a) identified key challenges such as scalability and distribution bottlenecks. Later, they proposed a bi-level decomposition method to solve large MILP models more efficiently (Triantafyllou et al. 2022a). They also developed simulation-based tools to support supply chain flexibility (Triantafyllou et al. 2022b), a decision-support system for evaluating resilience (Sarkis et al. 2024), and more recently, deep learning techniques to generate custom heuristics for large-scale optimization (Triantafyllou and Papathanasiou 2024).

2.2 Simulation-Based Modeling and Patient-Centric Scheduling

As CAR T-cell therapy involves stochastic elements like processing times, transport delays, and patient health variability, simulation has become a powerful tool for modeling operational dynamics. Wang et al. (2019) showed how simulation can help evaluate logistics decisions under uncertainty. Tully et al. (2019) used a discrete-event simulation model to quantify the clinical impact of delayed treatment, showing a link between longer wait times and increased one-year mortality.

Lam et al. (2021) compared centralized versus decentralized manufacturing strategies using simulation, showing that while centralized models offer economies of scale, decentralized options can improve turnaround times at high demand levels. Tseng et al. (2024) introduced SimPAC, a hybrid simulation framework that combines system dynamics and agent-based modeling. Their results show that prioritizing sicker patients and accelerating production significantly improves outcomes.

Unlike previous studies that relied solely on either optimization models or rule-based simulation, this study introduces a hybrid simheuristic framework that integrates simulation with metaheuristic optimization. This allows us to capture the uncertainty in processing and transportation while optimizing patient scheduling and transport mode decisions. The proposed model provides a practical way to balance delay penalties and shipping costs while remaining adaptive to time-window constraints and resource limits.

3 PROPOSED MODEL

As shown in Figure 1, the proposed model combines a discrete-event simulation of the CAR-T cell therapy workflow with a simulated annealing (SA) metaheuristic for decision optimization. This section first outlines the simulation design and then details the SA optimization procedure.

3.1 Discrete-Event Simulation

As illustrated in Figure 2, the model simulates the treatment journey of each patient through the following stages in a centralized manufacturing setting, where all patient samples are processed at a single manufacturing facility. The simulation is implemented using *SimPy*, an open-source discrete-event simulation library in Python.

- Arrival: Patients arrive according to a Poisson process with a rate $\lambda = 100$ patients. This value does not imply that 100 patients arrive per day; rather, it determines the interarrival times to spread patient arrivals over the early time horizon. This setup allows the evaluation of system performance and bottlenecks under a controlled, moderately high-demand scenario.




Figure 1: Simheuristic framework.




Figure 2: T-cell therapy supply chain process.

- **Leukapheresis:** The leukapheresis procedure time follows a triangular distribution with parameters (1, 3, 5) days. After this stage, each patient is assigned a *Target Infusion Date (TID)* indicating the ideal time for treatment delivery. The TID is calculated based on: the current simulation time, a fixed processing window (typically 28 days), and a randomly generated *slack* term that introduces variability based on patient-specific flexibility. The slack term is controlled by a parameter called *tightness* ($\tau \in [0, 1]$), which adjusts the width of the scheduling window. Lower values of τ represent more flexibility (i.e., looser windows), while higher values represent tighter scheduling. The slack is sampled from a triangular distribution scaled by $(1 - \tau)$:

$$\text{slack} \sim \text{Triangular}(5 \cdot (1 - \tau), 7 \cdot (1 - \tau), 9 \cdot (1 - \tau)) \quad (1)$$

TID is then computed as:

$$\text{TID} = \text{current time} + \text{Process Window} + \text{slack} \quad (2)$$

This formulation allows controlled variation in patient scheduling, enabling the simulation to capture realistic uncertainty while maintaining structured flow through the CAR T-cell supply chain.

- **Transport (Outbound):** Patients are assigned a transport mode (ground or air), with corresponding cost and time sampled from triangular distributions. Triangular distributions are used for transport time to reflect realistic uncertainty in the absence of large historical datasets. This transport delivers patient cells from hospitals to the centralized manufacturing facility.
- **Manufacturing:** Implemented as a SimPy PriorityResource with capacity 42 at the central manufacturing site. This means that, at most 42 patients can be processed simultaneously at any given time. If all slots are occupied, additional patients must wait in a priority queue. Each patient is assigned a scheduling priority based on one of these rules:

- FIFO (First-In-First-Out): Priority value equals the current simulation time when the patient arrives at the manufacturing stage, ensuring patients are processed in order of arrival (First-In-First-Out)
- SPT (Shortest Processing Time): Priority value equals the patient’s manufacturing time, prioritizing patients with shorter processing requirements
- EDD (Earliest Due Date): Priority value equals the patient’s Target Infusion Date, prioritizing patients with earlier deadlines
- SA: Priority value determined through the simulated annealing optimization process
- Quality Control: Each product may fail QC with probability $p_{qc_fail} = 0.1$, requiring re-manufacturing. It is possible for multiple QC failures to occur for the same patient.
- Transport (Return): The return transport mode is similarly sampled for delivering the modified cells from the central facility back to the patient’s hospital.

For each patient, the following metrics are computed: (i) *System time*, defined as the completion time minus the arrival time; (ii) *Delay*, calculated as $\max(0, \text{system time} - \text{TID})$; and (iii) *Cost*, which is the sum of transport cost, manufacturing cost, and delay penalty.

Table 1 summarizes the probability distribution functions used to model the duration of each process in the simulation. Each process, leukapheresis, transport, manufacturing, quality control (QC), and delivery, has an associated distribution that specifies its processing time. All time values are expressed in days. Triangular distributions are defined by (minimum, mode, maximum), and uniform distributions by (minimum, maximum). The values for these parameters were estimated based on the CAR-T process modeling framework presented in (Triantafyllou et al. 2022b).

Table 1: Probability distributions for process durations in the simulation model.

Process	Distribution Type	Parameters
Leukapheresis	Triangular	(1, 3, 5)
Transport (Ground)	Triangular	(1, 2, 3)
Transport (Air)	Triangular	(0.5, 1, 2)
Manufacturing	Triangular	(5, 10, 15)
Quality Control	Uniform	(1, 3)
Delivery	Triangular	(1, 2, 4)

3.2 Mathematical Formulation

3.2.1 Objective Function

The objective function that minimizes the total cost across all patients is given by:

$$\min \sum_i (C_{\text{trans}}^i + C_{\text{manu}}^i + C_{\text{delay}}^i) \tag{3}$$

where:

- C_{trans}^i = Transportation cost for patient i
- C_{manu}^i = Manufacturing cost for patient i
- C_{delay}^i = Delay penalty cost for patient i

3.2.2 Decision Variable Structure

The decision variables for each patient i include transportation mode choices for both outbound and return shipments, as well as the assigned manufacturing priority. These are encoded into a solution vector as:

$$q_i = [P_{t1}^i, \pi_i, P_{t2}^i] \quad (4)$$

where:

- $P_{t1}^i \in \{0, 1\}$: outbound transport mode (0: ground, 1: air),
- $\pi_i \in [1, \Pi_{\max}]$: assigned manufacturing priority score,
- $P_{t2}^i \in \{0, 1\}$: return transport mode (0: ground, 1: air).

For a problem with N patients, the full decision vector is defined as:

$$Q = \{q_1, q_2, \dots, q_N\} \quad (5)$$

The vector Q is optimized through the simulation-optimization framework.

3.2.3 Cost Components

The transportation cost depends on the selected modes for outbound and return transport, as shown in Equation (6) :

$$C_{\text{trans}}^i = \text{cost_factor}(P_{t1}^i) \cdot T_{\text{trans_1}}^i + \text{cost_factor}(P_{t2}^i) \cdot T_{\text{trans_2}}^i \quad (6)$$

The unit cost factor $\text{cost_factor}(P_t^i)$ is mode-dependent, where air transport incurs higher cost per unit time than ground transport.

The manufacturing cost combines fixed and variable components, as defined in (7) :

$$C_{\text{manu}}^i = \alpha + \beta \times T_{\text{manu}}^i \quad (7)$$

where:

- α = Fixed manufacturing cost parameter
- β = Variable cost per unit manufacturing time
- T_{manu}^i = Manufacturing time for patient i

The cost penalty for each patient is calculated based on how late the treatment is, multiplied by a fixed penalty, as formulated in (8):

$$C_{\text{delay}}^i = \delta \times \max(0, T_{\text{sys}}^i - \text{TID}_i) \quad (8)$$

where:

- δ = Delay penalty rate parameter
- T_{sys}^i = Total system time for patient i
- TID_i = Target infusion date for patient i

3.2.4 System Time Components

As defined in (9) the total system time (vein-to-vein time) comprises:

$$T_{\text{sys}}^i = T_{\text{trans_1}}^i + T_{\text{manu}}^i + T_{\text{qc}}^i + T_{\text{trans_2}}^i \quad (9)$$

where:

- $T_{\text{trans_1}}^i$ = Outbound transport time
- T_{manu}^i = Manufacturing processing time
- T_{qc}^i = Quality control time
- $T_{\text{trans_2}}^i$ = Return transport time

3.2.5 Scheduling and Capacity Constraints

The manufacturing capacity constraint is incorporated to ensure that no more than C_{cap} patients are processed simultaneously at any point in time. Binary variables $y_{it} \in \{0, 1\}$ are introduced, where $y_{it} = 1$ if patient i is under manufacturing at time t . This relationship is formulated as:

$$y_{it} \geq \mathbb{I}(S_i \leq t < S_i + T_{\text{manu}}^i) \quad (10)$$

where S_i denotes the manufacturing start time for patient i , and T_{manu}^i represents the corresponding manufacturing duration. The indicator function $\mathbb{I}(\cdot)$ returns 1 if the condition is satisfied, and 0 otherwise. The facility capacity constraint is then expressed as:

$$\sum_i y_{it} \leq C_{\text{cap}} \quad \text{for all } t. \quad (11)$$

Here, C_{cap} denotes the manufacturing facility capacity (42 slots).

3.2.6 Manufacturing Time

The manufacturing time varies based on patient-specific factors and is given by:

$$T_{\text{manu}}^i = T_{\text{base}} + \gamma \times (1 - P_{\text{via}}^i) + \lambda \times P_{\text{sev}}^i \quad (12)$$

where:

- T_{base} = Baseline manufacturing processing time
- γ = Viability adjustment factor
- P_{via}^i = Cell viability measure (0-1)
- λ = Severity adjustment factor
- P_{sev}^i = Disease severity score

All terms in Equation (13) are expressed in consistent time units (e.g., days). The adjustment factors γ and λ are calibrated such that their products with viability and severity scores yield time penalties, ensuring dimensional consistency.

3.3 Simheuristic Approach

A Simheuristic is a hybrid decision-making approach that combines simulation with metaheuristic optimization to efficiently explore complex, stochastic systems and identify near-optimal solutions under uncertainty. To effectively handle uncertainties in CAR-T manufacturing and logistics, a simheuristic framework is implemented (Juan et al. 2015; Dehghanimohammadabadi et al. 2017; Dehghanimohammadabadi and Kabadayi 2020). The adoption of a simheuristic approach is motivated by two key challenges: large-scale complexity and system stochasticity. As patient demand increases, the underlying optimization problem becomes computationally intractable for exact methods. For example, based on the model proposed in (Triantafyllou and Papathanasiou 2024), scaling to 5,000 therapies per year results in a MILP formulation with over 68 million constraints and 16 million discrete variables. Moreover, stochastic components such as transport time, quality control delays, and manufacturing variability further complicate the decision space. These uncertainties cannot be fully captured by deterministic models or efficiently solved using classical optimization methods. Simheuristics address both challenges by using simulation to model system randomness and metaheuristics to explore the high-dimensional decision space efficiently.

The proposed model combines Simulated Annealing (SA) with a SimPy-based discrete-event simulation model. This combination creates a feedback loop where optimization guides the search direction while simulation provides performance evaluation under uncertainty. The optimizer proposes candidate solutions (transport modes and scheduling priorities), the simulation model evaluates their performance across multiple

replications, and the optimizer uses this feedback to guide the search toward better solutions. SA was selected after empirical testing of metaheuristic approaches, including genetic algorithms. SA demonstrated consistent and satisfactory performance under the stochastic conditions of the simulation. While it does not guarantee global optimality, it offers a practical balance between solution quality, computational efficiency, and ease of implementation for this problem context.

3.4 Two-Stage Simulated Annealing

The simulation model is integrated with a two-stage simulated annealing framework that decomposes the optimization problem into sequential transportation and prioritization decisions.

In the first stage, the focus is exclusively on transport optimization. The decision vector consists of outbound and return transport mode selections (ground or air) for each patient in the system. The objective function during this stage minimizes total transport costs by setting the delay weight parameter to near zero, effectively prioritizing logistics efficiency. The neighborhood generation strategy employs a focused approach: with 75% probability, the transport modes of only the worst delayed patient are modified, while with 25% probability, modes for a randomly selected subset of patients are updated. This balanced exploration-exploitation mechanism helps avoid local optima while still targeting improvements where they are most needed. The output of this stage is a set of optimized transport decisions for all patients.

The second stage builds upon the first by taking the optimized transport modes as fixed inputs and focusing on priority optimization. The initial solution assigns random priorities to all patients within the manufacturing queue. The objective function in this stage considers the complete cost picture, minimizing the sum of transport costs (now fixed from Stage 1), manufacturing costs, and delay penalties with appropriate weighting. The neighborhood generation strategy mirrors that of Stage 1: with 75% probability, the priority value of the worst-delayed patient is updated, and with 25% probability, priorities for a random subset of patients are modified. This stage outputs optimized patient priorities that complement the transport decisions from Stage 1.

This decomposition approach ensures that cost-effective transport arrangements are first secured, followed by refinement of patient scheduling to minimize delay penalties. By addressing these decisions sequentially, the complexity of the search space is reduced while still achieving high-quality solutions to the overall optimization problem.

3.5 Experimental Setup

To evaluate the performance of this model, a series of experiments is conducted under varying time-window tightness levels. A centralized CAR-T manufacturing system is simulated with the following key configurations:

- Number of Patients: 100
- Number of Hospitals: 4
- Number of Manufacturing Facilities: 1
- SA Parameters: Maximum iterations = 100, cooling rate = 0.995
- Delay weight: Stage 1 = 0.5, Stage 2 = 5.0
- Replications per scenario: 25

Three levels of time-window tightness are considered, denoted by τ :

$$\tau \in \{0.3, 0.6, 0.8\} \tag{13}$$

Each algorithm is evaluated under identical simulation settings to ensure comparability.




Figure 3: Number of delayed patients across methods and tightness levels.

3.6 Benchmarking Approach

The two-stage SA method is benchmarked against standard scheduling heuristics (FIFO, SPT, and EDD). For these baseline approaches, transport mode selection is performed using a threshold-based rule, where the distance between each hospital and manufacturing facility is compared against the network-wide average distance.

For each replication, the following performance metrics are recorded: (i) *Total Cost*, representing the sum of transport, manufacturing, and delay costs during the treatment process; (ii) *Transport Cost*, which accounts for the expenses of moving cells and patients between facilities; (iii) *Delay Cost*, reflecting penalties incurred when treatment is not initiated within the required time window; and (iv) *Number of Delayed Patients*, indicating how many patients experience delays beyond their acceptable treatment initiation deadlines.

4 RESULTS AND DISCUSSION

This section presents a comparative analysis of four scheduling and transportation strategies under varying levels of patient time window tightness. The objective is to evaluate each method's performance in minimizing patient delays, transportation costs, delay-related costs, and overall system costs within the context of a CAR T-cell therapy supply chain.

To assess the system's sensitivity to scheduling complexity, a tightness parameter ranging from 0.3 (very loose patient time windows) to 0.8 (strict timing constraints) is introduced. The results are summarized using boxplots that display the distribution of each performance metric across 25 simulation replications per scenario. The following subsections discuss the outcomes for each metric in detail.

4.1 Impact of Time Window Tightness on Patient Delays

As illustrated in Figure 3, as tightness increases (i.e., from left to right), the number of delayed patients generally increases for most algorithms. This aligns with the expectation that tighter scheduling constraints make it harder to meet deadlines.

SA significantly outperforms SPT in reducing patient delays at tightness 0.3 and 0.6, with low medians and tight distributions. EDD consistently achieves the lowest median delay, especially under tighter windows, followed closely by SA and FIFO. SPT performs the worst at tightness 0.1 and 0.6. Tight windows enforce stricter prioritization, and while fixed rules like EDD can perform well under these conditions, metaheuristic optimization such as SA offers more flexibility than SPT, allowing it to explore more effective patient-transport assignments.




Figure 4: Transport cost comparison across methods and tightness levels.




Figure 5: Total cost comparison across methods and tightness levels.

4.2 Transportation Cost Analysis

Our analysis of transportation costs reveals that they vary more widely under looser time windows, as shown in Figure 4. SA achieves the lowest median transport cost across all scenarios. Under loose constraints (tightness = 0.3), EDD and SPT perform comparably, while FIFO shows the highest costs. Metaheuristic approaches such as SA are better suited to coordinate logistics in constrained scenarios, and rule-based methods are more susceptible to inefficiencies when strict timing constraints are in place.

4.3 Total Cost Evaluation

When examining total cost, which reflects a balance between delay-related penalties, transportation expenses, and manufacturing costs, Figure 5 demonstrates that SA consistently achieves the lowest or near-lowest total cost across all tightness scenarios. SPT results in the highest total costs, driven by both delays and inefficient transport usage. This suggests that approaches that integrate patient-centered scheduling with optimization result in more robust and cost-effective outcomes across varying operational conditions.

4.4 Breakdown of Delay Costs

The analysis of delay costs, presented in Figure 6, highlights the burden of missed treatment windows and reinforces patterns observed in patient delay metrics. EDD yields the lowest delay cost, especially at higher tightness levels, followed closely by SA and FIFO. SA outperforms SPT in all scenarios, but EDD shows




Figure 6: Delay cost comparison across methods and tightness levels.

better results under stricter constraints. Methods that prioritize due dates (EDD) or dynamically evaluate scheduling decisions (SA) are more effective at minimizing the cost impact of treatment delays.

4.5 Practical Implications for CAR T-cell Supply Chain Management

These results demonstrate that simulation-based optimization, particularly via Simulated Annealing, provides some benefits over basic processing-time-based heuristics like SPT. Although EDD remains a strong rule-based alternative under tight timing constraints, SA delivers more balanced and adaptable performance when both logistics and treatment timing are considered. These findings underscore the value of integrating adaptive optimization into time-critical healthcare delivery systems such as CAR T-cell therapy.

The two-stage SA approach offers particular advantages in real-world implementations where conditions frequently change. The decomposition of the problem into transport and scheduling phases allows for more targeted optimization efforts and better adaptation to varying constraints. Furthermore, the ability to dynamically respond to uncertainties in processing times, patient conditions, and resource availability represents a significant advantage for healthcare logistics managers working with personalized therapies.

SA outputs can be structured into actionable operational guidelines for scheduling and transportation teams. In practice, the model generates optimized patient schedules and transportation mode selections, which can be translated into daily assignment lists specifying:

- The sequence and timing of patient manufacturing slots.
- The selected transportation mode for each patient.

Such outputs can be integrated into existing scheduling and dispatching software, providing staff with a clear task list without requiring knowledge of the underlying optimization algorithm. One potential barrier to adoption is the complexity of metaheuristic algorithms. To address this, SA-generated recommendations can be presented by simple decision explanations (e.g., "Patient A scheduled to minimize delay cost," or "Ground shipping selected for Patient B as it meets delivery requirements at lower cost").

5 CONCLUSION AND FUTURE WORK

This study developed and evaluated a two-stage simulated annealing approach to optimize CAR T-cell therapy supply chains, focusing on the critical balance between transportation efficiency and timely delivery. The results demonstrate that the proposed model consistently outperforms traditional scheduling heuristics across various time-window constraints. By decomposing the optimization problem into sequential transport and priority stages, this model effectively addresses the inherent stochasticity in the CAR T-cell supply chain while maintaining computational efficiency. While EDD performs well for delay minimization in

highly constrained scenarios, it fails to effectively balance transportation and manufacturing costs. The proposed simulation-optimization framework shows superior adaptability to changing conditions, making it particularly valuable for the variable nature of personalized medicine logistics.

Future research could extend this work in several promising directions. First, while our study focused on centralized manufacturing, investigating decentralized models could provide insights into reducing transportation delays and costs. Second, more detailed capacity analysis across multiple manufacturing centers would enhance decision-making in large-scale scenarios. Third, extending the model to include inventory planning for critical raw materials could improve supply chain resilience against disruptions.

Additionally, incorporating advanced AI techniques, such as Generative AI for real-time decision support, could further enhance the system's responsiveness to changing conditions. Future work should validate these simulation results in real-world settings, considering practical operational constraints and supply chain uncertainties like transportation disruptions. Such validation would strengthen the practical applicability of simulation-based optimization approaches in improving access to this life-saving therapy while controlling costs across the complex CAR T-cell supply chain network.

REFERENCES

- Bernardi, A., M. Sarkis, N. Triantafyllou, M. Lakelin, N. Shah, and M. M. Papathanasiou. 2022. "Assessment of Intermediate Storage and Distribution Nodes in Personalised Medicine". *Computers & Chemical Engineering* 157:107582 <https://doi.org/10.1016/j.compchemeng.2021.107582>.
- Dehghanimohammadabadi, M., and N. Kabadayi. 2020. "A Two-Stage AHP Multi-Objective Simulation Optimization Approach in Healthcare". *International Journal of the Analytic Hierarchy Process* 12(1) <https://doi.org/10.13033/ijahp.v12i1.701>.
- Dehghanimohammadabadi, M., M. Rezaeiahari, and T. K. Keyser. 2017. "Simheuristic of Patient Scheduling Using a Table-Experiment Approach Simio and Matlab Integration Application". In *2017 Winter Simulation Conference (WSC)*, 2929–2939 <https://doi.org/10.1109/WSC.2017.8248015>.
- Di, M., K. C. Potnis, J. B. Long, I. Isufi, F. Foss, S. Seropian, *et al.* 2024. "Costs of Care During Chimeric Antigen Receptor T-Cell Therapy in Relapsed or Refractory B-Cell Lymphomas". *JNCI Cancer Spectrum* 8(4):pkae059 <https://doi.org/10.1093/jncics/pkae059>.
- Jackson, H. J., S. Rafiq, and R. J. Brentjens. 2016. "Driving CAR T-Cells Forward". *Nature Reviews Clinical Oncology* 13(6):370–383 <https://doi.org/10.1038/nrclinonc.2016.36>.
- Juan, A. A., J. Faulin, S. E. Grasman, M. Rabe, and G. Figueira. 2015. "A Review of Simheuristics: Extending Metaheuristics to Deal With Stochastic Combinatorial Optimization Problems". *Operations Research Perspectives* 2:62–72 <https://doi.org/10.1016/j.orp.2015.03.001>.
- Lam, C., E. Meinert, A. Yang, and Z. Cui. 2021. "Comparison Between Centralized and Decentralized Supply Chains of Autologous Chimeric Antigen Receptor T-Cell Therapies: A UK Case Study Based on Discrete Event Simulation". *Cytotherapy* 23(5):433–451 <https://doi.org/10.1016/j.jcyt.2020.08.007>.
- Lopes, A. G., R. Noel, and A. Sinclair. 2020. "Cost Analysis of Vein-to-Vein CAR T-Cell Therapy: Automated Manufacturing and Supply Chain". *Cell and Gene Therapy Insights* 6(3):487–510 <https://doi.org/10.18609/cgti.2020.058>.
- Sarkis, M., N. Shah, and M. M. Papathanasiou. 2024. "Towards Resilient Pharmaceutical Supply Chains: Integration of Manufacturing Uncertainty". In *2024 AIChE Annual Meeting*.
- Triantafyllou, N., A. Bernardi, M. Lakelin, N. Shah, and M. M. Papathanasiou. 2022a. "A Bi-Level Decomposition Approach for CAR-T Cell Therapies Supply Chain Optimisation". In *Computer Aided Chemical Engineering*, Volume 49, 2197–2202. Elsevier <https://doi.org/10.1016/B978-0-323-85159-6.50366-3>.
- Triantafyllou, N., A. Bernardi, M. Lakelin, N. Shah, and M. M. Papathanasiou. 2022b. "A Digital Platform for the Design of Patient-Centric Supply Chains". *Scientific Reports* 12(1):17365 <https://doi.org/10.1038/s41598-022-21290-5>.
- Triantafyllou, N., and M. M. Papathanasiou. 2024. "Deep Learning Enhanced Mixed Integer Optimization: Learning to Reduce Model Dimensionality". *Computers & Chemical Engineering* 187:108725 <https://doi.org/10.1016/j.compchemeng.2024.108725>.
- Tseng, C.-Y., K. Wang, L.-H. Lin, C. Zhang, C. C. I. White, and B. Wang. 2024. "Advancing Autologous CAR T-Cell Therapy Through Real-Time Patient Health Data Integration: A Simulation-Based Approach". *Cytotherapy* 26(10):1152–1162 <https://doi.org/10.1016/j.jcyt.2024.05.001>.
- Tully, S., Z. Feng, K. Grindrod, T. McFarlane, K. K. W. Chan, and W. W. L. Wong. 2019. "Impact of Increasing Wait Times on Overall Mortality of Chimeric Antigen Receptor T-Cell Therapy in Large B-Cell Lymphoma: A Discrete Event Simulation Model". *JCO Clinical Cancer Informatics* 3:1–9 <https://doi.org/10.1200/CCI.19.00086>.

- Utkarsh, K., N. Srivastava, S. Kumar, A. Khan, G. Dagar, M. Kumar, *et al.* 2024. “CAR-T Cell Therapy: A Game-Changer in Cancer Treatment and Beyond”. *Clinical and Translational Oncology* 26(6):1300–1318 <https://doi.org/10.1007/s12094-023-03368-2>.
- Wang, K., Y. Liu, J. Li, B. Wang, R. Bishop, C. White, *et al.* 2019. “A Multiscale Simulation Framework for the Manufacturing Facility and Supply Chain of Autologous Cell Therapies”. *Cytotherapy* 21(10):1081–1093 <https://doi.org/10.1016/j.jcyt.2019.07.002>.

AUTHOR BIOGRAPHIES

MOHAMMAD DEHGHANIMOHAMMADBADI is an Associate Teaching Professor at the Department of Mechanical and Industrial Engineering, Northeastern University, Boston. His research focuses on the development and implementation of AI-enabled digital twins and simulation models with applications in healthcare, supply chain, and manufacturing and his e-mail address is: m.dehghani@northeastern.edu.

NEGAR SADEGHI is a Ph.D. student in Industrial Engineering at Northeastern University. Her research focuses on developing intelligent simulation models for pharmaceutical supply chains and production-distribution systems. Her e-mail address is: sadeghi.ne@northeastern.edu.