In-Situ monitoring of 3D printing process for Data Fusion using Kalman Filter

C. Courcoux[1], D. Therriault[1], and F. P. Gosselin[1*]

[1] Laboratory of Multiscale Mechanics (LM²), CREPEC, Department of Mechanical Engineering, Polytechnique, *Montreal*, Canada

* *Corresponding author (frederick.gosselin@polymtl.ca)*

Keywords: *Additive Manufacturing; Digital Twin; Kalman Filter; Sensor;*

1 Abstract

Achieving high-quality extrusion-based additive manufacturing components requires careful control of the thermal history during deposition. The transient cooling of the molten polymer beads governs the strength of the interlayer interface, while heterogeneous heat distributions induced by the successive deposition of material generate residual stresses that degrade the mechanical and dimensional performance of the part. This work investigates the thermal cooling behavior of a part produced by fused granulate fabrication. A lumped-mass method coupled with a voxel activation strategy is used to simulate the temperature field as material is added on the heated bed, with an adaptive discretization driven by the local Biot number to capture the strong cross-section gradient between the bed interface and the air-exposed faces. To keep the computational cost manageable, the model relies on simplifying assumptions that introduce a deviation from the experimental data. An experimental setup based on two infrared cameras is presented with a compact thermal camera embedded on the printing head that provides a local view of the deposition zone, while a second camera is installed in front of the printer for validation. The data acquired by the embedded camera is then fused with the simulation through a recursive Kalman filter algorithm to reconstruct the temperature field of the print and to reduce the model deviation from reality. The filter pulls the simulated state toward the measurements during the observation window and propagates the correction to non-observed voxels through their conductive coupling, demonstrating that a partial observation of the cross-section can update the thermal state of the bead.

2 Introduction

Additive manufacturing (AM) of thermoplastics encompasses two main extrusion-based technologies with the fused filament fabrication (FFF) and fused granulate fabrication (FGF). FFF is the most widely adopted process, valued for its versatility and accessibility, in which a continuous filament is fed into a heated nozzle. In contrast, FGF relies on a rotating screw to drive plastic pellets through a heated barrel, where they are melted and pressurized before extrusion. In both cases, the molten polymer is deposited onto a heated bed through a circular nozzle whose diameter is typically around 0.4 mm for FFF and ranges from 1 mm to 6 mm for conventional FGF. In an industrial context, the larger nozzles and screw-driven extrusion of FGF substantially increase productivity [1], but the deposited beads retain heat for longer, which raises new thermal challenges. Thermal phenomena largely govern process quality, as they dictate the state of the polymer, the quality of inter-layer bonding, and the final dimensional accuracy of the part. As reported by Corfmat [2], significant thermal gradients within the deposited material induce residual stresses, inconsistent inter-layer adhesion, and a loss of dimensional accuracy. These effects are amplified when productivity is pushed further through higher printing speeds or larger bead deposition, both of which prolong heat retention. To better understand and control the process, several thermal and thermo-mechanical models of FFF/FGF have been developed using finite element or finite difference methods [3,4]. Thermo-mechanical simulations have demonstrated the ability to quantitatively predict thermal gradients, but they remain subject to important limitations. Keim et al. [5] highlight the difference between announced data and the real state of the system. Information extracted from the 3D printer or the G-code is shown to be imprecise data that can deviate by more than 30 degrees. To improve both models and the process, current studies are trying to implement data fusion algorithms or digital twins to correct model discrepancies. A key trade-off lies between the computational cost of the simulation and the simplifying assumptions required to deliver accurate results within acceptable computation times. To reduce the deviation of the model from imprecise data, Kim [6] proposed a workflow based on a recursive estimation algorithm called the Kalman filter to estimate the thermal state of the polymer by fusing simulation output with measurements from an infrared (IR) camera. Their results show that the filter effectively reduces the relative error when synthetic IR data, generated by superimposing normally distributed noise onto the simulated results, are incorporated into the estimation.

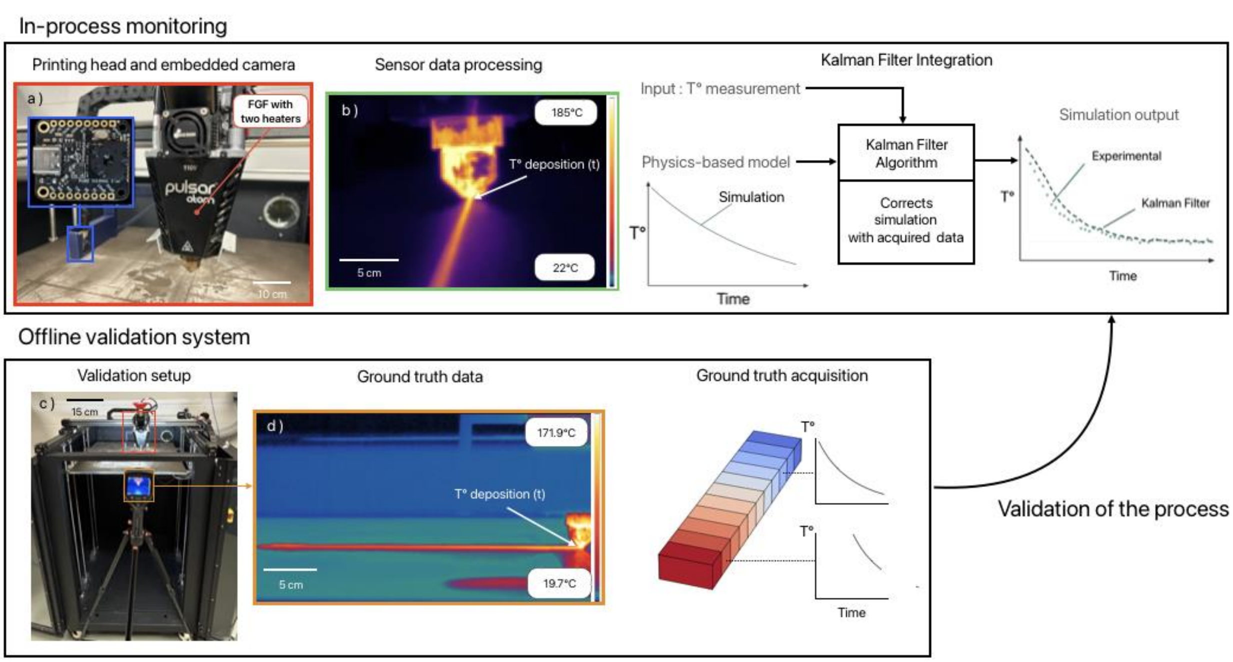

Figure 1: Workflow of cooling analysis based on in-process monitoring of (a) an FGF 3D printer observed by a FLIR Lepton 3.5 thermal camera. Data are acquired via b) Python code and merged with thermal simulations using a Kalman Filter algorithm. The process is validated by comparing data from (c) a FLIR E54 camera analyzed with (d) FLIR software.

This paper focuses on applying the Kalman filter with real data on a heat transfer simulation of a deposited filament cooling to reduce the model error. The general workflow is presented in Figure 1, with an embedded thermal camera fixed on the FGF printing head to acquire the real thermal state of the polymer during the print.

3 Methodology

3.1 Experimental setup

The experimental setup used to implement the data-fusion algorithm is presented in Figure 1 c). An existing FFF 3D printer was modified into an FGF platform by installing a Pulsar Atom extruder driven by a dedicated controller and acquisition board, both provided by Dyze Design. The experiments presented in this paper rely on two IR cameras providing complementary views of the printing process. For in-process monitoring, a compact FLIR Lepton 3.5 camera is mounted on the printing head through a 3D-printed ABS support and interfaced with a Thermal Module V3. Figure 1 b) shows a typical thermal image obtained from this camera during the deposition of a filament. The data acquired via Python therefore provides only a partial picture of the thermal state of the deposited polymer. The frontal camera shown in Figure 1 c) is not part of the filtering loop. It provides an independent IR view of the printed part in Figure 1 d) and is used to validate the corrected temperature field produced by the algorithm against an external reference.

3.2 Experimental procedure

The simulation model relies on several assumptions to compensate for parameters that are not directly accessible during the print, such as the exact temperature of the surrounding air, and to keep the computational cost manageable. As shown by Keim et al. [5], the setpoints announced by the 3D printer can also differ from the actual state of the system, with discrepancies of up to 30 °C reported on the bed temperature alone. To probe the limits of the model, a single straight bead was printed on the FGF system using a 1.2 mm nozzle and transparent PLA 3D850 pellets from NatureWorks. The bead is 300 mm long, with a height of 1.05 mm and a width of 1.70 mm measured at mid-length with a Mitutoyo 500-196-30 caliper and was deposited at a printing speed of 30 mm·s-1. To investigate the deposition phase further analyzed in the experimental results, the bead and the nozzle outlet are both recorded with the FLIR E54 camera positioned in front of the printer. The recordings are post-processed with the FLIR software where on each frame the reported temperature is taken as the maximum value measured along the bead. The printing head shown in Figure 1 a) is equipped with two consecutive heaters that progressively melt and homogenize the pellets along the barrel.

3.3 Numerical model

A polymer's thermal cooling behavior arises from its material properties and from the heat exchanges it undergoes with its environment [7]. We consider a PLA filament with uniform temperature in its cross-section and model its temperature variation T along its length. Treating PLA as an isotropic material and considering the specific heat c_p and density ρ are independent of temperature, we can discretize the filament in the length using voxels with a rectangular cross-section. Studying the temperature T_i^p of the element i at the time step p, we can write a lumped-mass equation based on the interaction with surrounding solid elements j belonging to the solid neighbor domain of i noted N_i, the air and the bed as:

$$\rho\, V\, c_p \frac{T_i^{p+1} - T_i^p}{\Delta p} = \sum_{j \in N_i} \frac{k * A_{ij}^2}{V}\left(T_j^p - T_i^p\right) + h_{bed}\, A_{bed} * \left(T_{bed} - T_i^p\right) + h_{air} A_{air} * \left(T_{air} - T_i^p\right), \quad (1.1)$$

with V the volume of each voxel, A_{ij} the contact area between elements i and j. Interactions with air and bed are respectively represented as a convective behavior with the coefficient of convection h and the contact area A at the temperature T indexed by the interaction name.

Table 1: Material properties for a 3D850 transparent PLA from NatureWorks in pellets used in the simulation model. Properties are considered independent from the temperature, with h_{air} the natural convection without fan and h_{bed} calibrated for our system.

ρ $(\mathrm{kg.m^{-3}})$	c_p $(\mathrm{J.Kg^{-1}.K^{-1}})$	ε	k $(\mathrm{W.m^{-1}.K^{-1}})$	h_{bed} $(\mathrm{W.m^{-2}.K^{-1}})$	h_{air} $(\mathrm{W.m^{-2}.K^{-1}})$
1240	1800	0.91	0.13	160	10

The radiative interaction accounting for radiative losses with σ the Stefan–Boltzmann constant and ε the material emissivity is neglected. As noted in [3], for polymers operating at moderate temperatures the radiative contribution to cooling is small relative to convection. Using the parameters in Table 1, the convection-to-radiation ratio $(h_{bed}A_{bed} + h_{air}A_{air})(T_{env} - T)/\varepsilon\sigma A_{air}(T_{env}^4 - T^4)$ takes a value of 6.95 at $T = 200$ °Cand $T_{env} = 37$ °C and grows further as the polymer cools. The radiative term can therefore be safely dropped, as shown in equation (1.1).

To represent the additive nature of the process, we adopt the voxel activation strategy illustrated in Figure 2 a), the final dimensions of the part are pre-allocated, and voxels are progressively activated as new material is deposited on the heated bed. To reduce the computational cost, we aim to keep the number of elements per cross-section as low as possible while preserving the physics of the cooling process. The Biot number, defined as the ratio of the convective to the conductive heat resistance over a characteristic length, provides a criterion for this trade-off. A Biot number below 0.1 indicates that the temperature in the cross-section can reasonably be assumed uniform. Taking half of the filament height as the characteristic length L_c for the heat to go out from the polymer, we can calculate the ratio equal to $h\, L_c\, /\, k$ [7]. Figure 2 b) shows the resulting mesh where the row of voxels in contact on the heated bed reaches a Biot number of 1.23 and is refined into a 3×3 sub-grid, whereas the voxels exchanging only with the surrounding air remain uniform with a Biot number of 0.076.

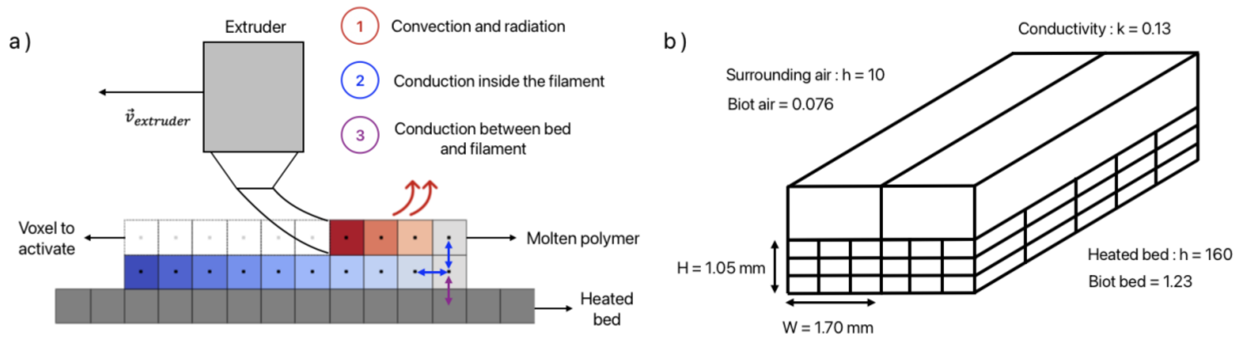

Figure 2: Thermal model representation in (a) with the voxel activation method that adds elements during the print, and the variable discretization depending on the uniformity of the cross-section temperature based on the Biot number.

3.4 Kalman Filter Algorithm

To integrate the numerical simulation with the experimental measurements, we implement a linear Kalman filter [8] that fuses the model with IR temperature data acquired by the embedded camera. The filter is a recursive estimator that performs two steps at each iteration, a prediction step provided by the simulation alone and a correction step that uses the new measurement to update the predicted state. Equation (1.1) can be rearranged as an explicit formulation to serve as the prediction step:

$$T_i^{p+1} = \left(1 - \frac{\Delta p}{\rho V c_p} \sum_{j \in N_i} \frac{k A_{ij}^2}{V} (h_{bed} A_{bed} + h_{air} A_{air})\right) T_i^p + \frac{\Delta p}{\rho V c_p} \sum_{j \in N_i} \frac{k A_{ij}^2}{V} T_j^p + \frac{\Delta p}{\rho V c_p} (h_{bed} A_{bed} T_{bed} + h_{air} A_{air} T_{air}). \quad (1.2)$$

To adapt the Kalman filter formalization, equation (1.2) can be reformulated as a transition matrix $\boldsymbol{A^p}$ that multiplies our state vector of temperatures at time step p, $\overrightarrow{T^p}$ and $\overrightarrow{F^p}$ a forcing vector accounting for bed and air transfer:

$$\overrightarrow{T^{p+1}} = \boldsymbol{A^p} \overrightarrow{T^p} + \overrightarrow{F^p} + \overrightarrow{w^p}, \quad (1.3)$$

where the process noise $\overrightarrow{w^p} \sim \mathcal{N}$ (0, **Q**) lumps together the unmodelled physics from neglected radiative losses, uncertainty on the convective coefficients and geometric simplifications based on a covariance matrix **Q** [9] that quantifies the trust placed in the model.

The embedded FLIR Lepton 3.5 camera captures only the voxels currently within its field of view, on the top surface of the freshly deposited filament. From the simulation, the estimated measurement is equal to the simulated temperature, with a measurement noise $\overrightarrow{v^p} \sim \mathcal{N}$ (0, **R**) accounts for the IR camera uncertainty. As mentioned, the data acquisition is a partial observation of the temperature field encoded by an observation matrix $\boldsymbol{H^p}$ that selects the visible voxels from the full state in the equation:

$$\overrightarrow{\hat{z}^p} = \boldsymbol{H^p} \overrightarrow{T^p} + \overrightarrow{v^p}. \quad (1.4)$$

At each acquisition step, the filter compares the measurement from the IR Camera $\overrightarrow{z^p}$ with the predicted measurement $\vec{\hat{z}}^p$ obtained from the simulation (1.4), and corrects the simulated state $\overrightarrow{T^p}$ in (1.3) into a corrected estimate $\overrightarrow{\overline{T^p}}$ through a Kalman gain $\boldsymbol{K^p}$ in the update equation:

$$\overrightarrow{\overline{T^p}} = \overrightarrow{T^p} + \boldsymbol{K^p}\left(\overrightarrow{z^p} - \vec{\hat{z}}^p\right). \tag{1.5}$$

The Kalman gain $\boldsymbol{K^p}$ is computed at every step from a state covariance matrix $\boldsymbol{P}$ together with $\boldsymbol{Q}$ and $\boldsymbol{R}$ and behaves as a weighting factor that balances the prediction against the measurement. A large gain pulls the estimated field toward the temperature observed by the IR camera, while a small gain keeps the estimate close to the simulated value. Between two IR frames, the algorithm runs on the prediction step (1.3) only, so the temperature of the non-observed voxels is propagated from the model alone and at each frame rate, the equation (1.4) and (1.5) are successively run to correct the model.

The objective using the Kalman Filter is to observe its ability to correct the model deviation of a 3D printing process. The same bead geometry, presented in previous Section 3.2, is printed on the bed. Mounted on the printing head directly behind the nozzle, the FLIR Lepton 3.5 provides a top-view thermal image of the freshly deposited polymer over a narrow field of view, supplying the localized temperature measurements that feed the correction step of the filter at each acquired frame. The acquired data is added into the Kalman filter for each frame, for the elements in the field of view of the camera. In addition, as mentioned in Section 3.1, a thermal camera is used to acquire the temperature in the middle length of the front bead at both acquisition stages to validate the Kalman filter correction.

4 Results

4.1 Simulation observation

To justify the adaptive discretization introduced in Section 3.3, the cross-section of a single bead is simulated with two mesh configurations sketched in Figure 3 a) for a single element and a 3×3 cross-section area. The simulation is initialized at 210 °C with a heated bed at 45 °C and the surrounding air at 37 °C. Figure 3 b) reports the cooling curves for the high-Biot regime that characterizes the bottom interface with the heated bed. Under this regime, the three voxels of the 3×3 model exhibit clearly distinct trajectories. The bottom voxel cools rapidly through direct conduction with the bed, the top voxel retains heat for a longer period, and the middle voxel follows an intermediate path. The 1×1 simulation does not reproduce this gradient and yields a curve that averages the three positions. The temperature difference between the 1×1 prediction and the bottom or top voxel of the 3×3 model reaches up to 40 °C during the cooling phase, which is not negligible at the scale of the bead.

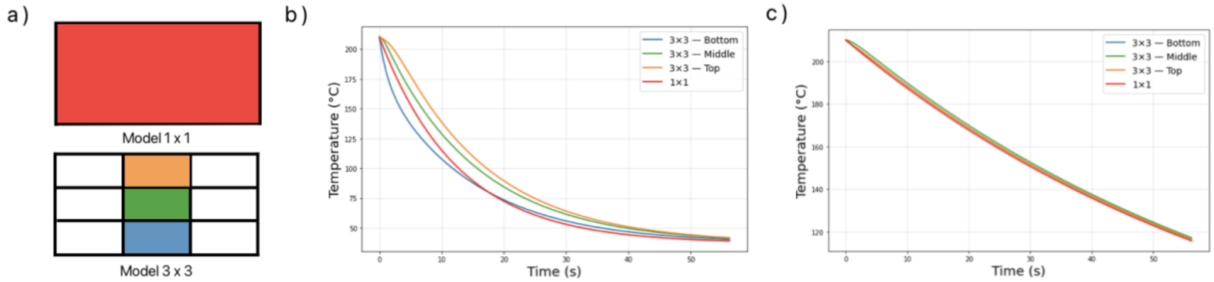

Figure 3: Cooling curves predicted by the thermal model for the cross-section of a 1.70 × 1.05 mm bead, comparing (a) a single-voxel representation with a 3×3 discretization. The two regimes are illustrated by (b) the high Biot number case (Bi = 1.23) on the bed interface and (c) the low Biot number case (Bi = 0.076) on the faces exposed to ambient air.

In contrast, Figure 3 c) corresponds to the low-Biot regime that prevails on the faces exposed to ambient air. In this regime the four curves overlap throughout the cooling, confirming that the assumption of a uniform cross-section temperature holds. This result underlines the importance in our code of adapting the discretization to the behavior of the polymer. The data from the IR camera, which looks at the top of the polymer, cannot be directly compared with a single element without introducing an error.

4.2 Experimental observation

The deposition of the polymer is observed to study the thermal cooling of the printed bead, as presented in Figure 4 a), at the middle length of the system by acquiring the maximum temperature on the selected white line for a nozzle diameter of 1.2 mm and a bead with the previously presented dimensions.

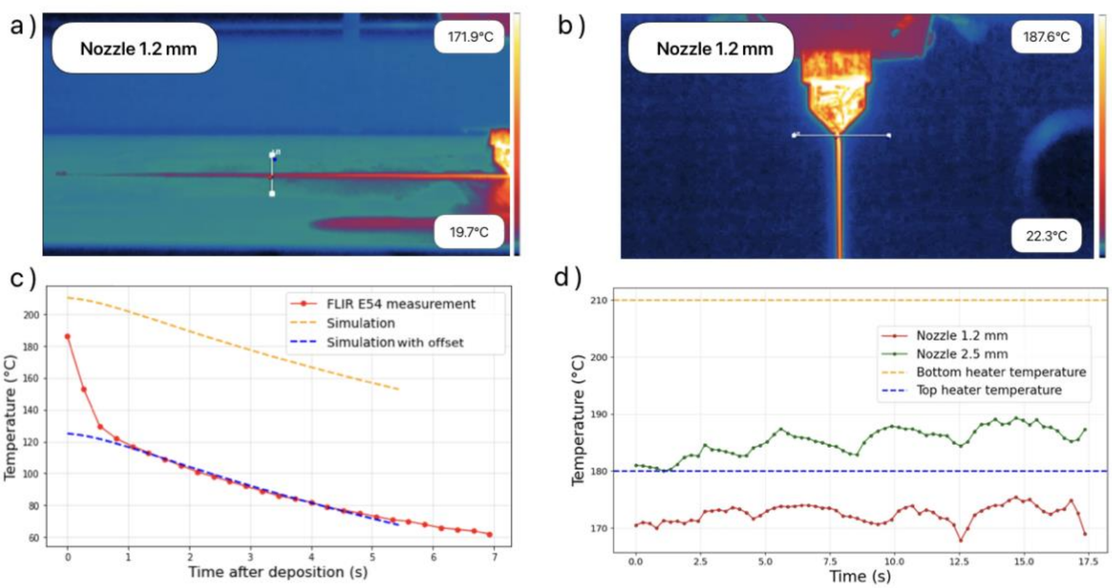

Figure 4: (a) IR view of a 1.2 mm bead and (b) IR view of the printing head, both acquired with the FLIR E54 camera positioned in front of the printer. (c) Cooling profile measured along the bead compared with the thermal simulation, with and without an initial-temperature offset of 85 °C, highlighting the limits of the model during the deposition phase. (d) Polymer temperature measured at the nozzle outlet for the 1.2 mm and the 2.5 mm nozzles, compared to the setpoints of the top and bottom heaters.

The cooling profile is compared to the simulation initialized at the bottom heater setpoint of 210 °C in Figure 4 c). The simulation overestimates the polymer temperature, with a maximum deviation exceeding 80 °C during the first seconds after deposition. Repeating the simulation with an initial condition reduced by an offset of 85 °C produces a curve that closely matches the experimental data for the time 1 second after the deposition. This residual gap during the deposition phase can be explained by the difference of temperature between the printing head and the polymer. Some transient phenomena may be underestimated by the code, but the main hypothesis is a local convective effect produced by the relative movement of the filament during the deposition at the speed of the printing head. The difference between the heater setpoint and the simulation initial condition is explained by Figure 4 d), which reports the polymer temperature measured at the nozzle outlet for two different nozzle diameters. For the 1.2 mm nozzle, the extruded polymer leaves the nozzle at an average temperature of 173 °C, while the bottom heater is set to 210 °C, revealing a temperature drop of approximately 37 °C. For the 2.5 mm nozzle, the same behavior is observed with a higher temperature around 185 °C. The deposition temperature is therefore not equal to the nozzle setpoint, and the difference depends on the nozzle geometry and on the polymer flow rate. The temperature variation seems to follow the same trend for the two different nozzle geometries, which can result from the machine control with the activation and deactivation of the heater or from a thermal noise of the camera. The observation of the temperature difference at the deposition and the inability of the model to reproduce the thermal cooling a few seconds after the deposition motivate the use of an embedded camera able to measure the deposition in real-time and to feed the initial condition of the model through the Kalman filter algorithm.

4.3 Kalman Filter algorithm

The Kalman filter procedure described in Section 3.4 is now applied to the 300 mm straight bead printed with the 1.2 mm nozzle. The thermal model is initialized at the bottom heater setpoint of 210 °C and runs in parallel with the embedded FLIR Lepton 3.5 camera fixed on the printing head. At each acquired frame, all the temperatures measured on the top of the cross-section are injected into the correction step of the filter. The corrected state on the side voxel is then compared with the validation data acquired by the front FLIR E54 camera, which does not enter the filtering loop. The adaptive 3×3 discretization introduced in Section 3.3 is used for the analysis, according to a Biot number higher than 0.1.

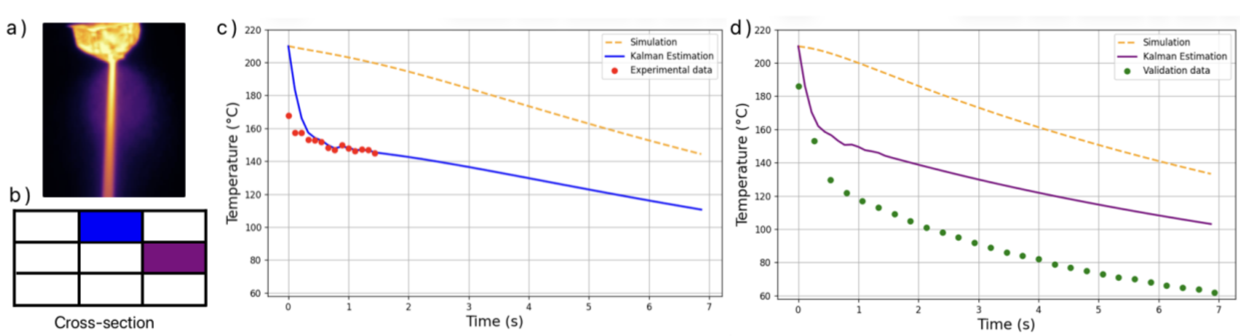

Figure 5: Application of Kalman Filter using data extracted from a) the embedded camera image on the top of the cross-section discretization b) to correct the temperature cooling c) simulated in Python by applying Kalman filter algorithm. The validation data is compared to the side of the filament in d).

Figure 5 c) reports the cooling profile of the top voxel highlighted in blue. The simulation initialized at 210 °C overestimates the polymer temperature by more than 50 °C during the first second after deposition, consistently with the deposition temperature drop discussed in Section 4.2. As soon as the embedded camera enters the deposition zone, the Kalman filter immediately pulls the estimated state from the simulated value toward the measurement. The corrected temperature drops sharply from 210 °C to follow the experimental trajectory within a few seconds. Once the camera leaves the field of view and no new measurement is available, the filter switches back to a pure prediction step and the temperature is propagated through the model from a corrected initial state that is now consistent with the experimental observations. The resulting curve lies up to 50 °C below the original simulation throughout the cooling, which represents a significant correction of the model. Figure 5 d) illustrates the propagation of this correction to the side voxel highlighted in purple, which is never directly observed by the embedded camera. The update applied to the top voxel is redistributed across the cross-section. As a result, the Kalman estimate of the side voxel lies systematically below the pure simulation and approaches the validation data acquired by the front FLIR E54 camera, confirming that a partial observation of the cross-section is sufficient to update the full thermal state. The corrected estimate does not, however, fully match the validation data with a difference observed throughout the cooling. This gap is not only attributed to a possible deficiency of the filter itself, which only redistributes the information available in the model and the measurements, but rather to an unresolved discrepancy between the two infrared cameras. The FLIR Lepton 3.5 used inside the filtering loop and the FLIR E54 used as an external reference have different readings of the same bead, exhibiting a calibration shift that has not yet been quantified. A joint calibration of both cameras therefore stands out as a necessary next step to consolidate the validation of the data-fusion workflow.

5 Conclusion

This work presents a data-fusion workflow that couples a thermal method simulation of the cooling of an FGF-printed bead with in-process measurements acquired by a thermal camera embedded on the printing head, through a recursive Kalman filter algorithm. The model relies on a voxel activation strategy and on an adaptive discretization governed by the local Biot number, which keeps the computational cost low while preserving the cooling physics. The experimental campaign highlights two important limits of an uncorrected model. The polymer leaves the nozzles 30 to 40 °C below the heater setpoint and the actual deposition temperature is not directly accessible from the printer parameters. In addition, a drop in the temperature of the polymer after the deposition is missing in the simulation. We suppose a local convective effect during the deposition due to the movement created by the printing head. Applying the Kalman filter to the real measurements of the embedded FLIR Lepton 3.5 camera successfully shows a reduction of the model deviation. The estimated state is pulled toward the experimental data during the observation and is then propagated through the model from a physically consistent initial condition. The correction also propagates to non-observed voxels through their conductive coupling with the observed ones, demonstrating that a partial measurement of the cross-section can be sufficient to update the full thermal state of the bead. A residual observed error with the validation, from another IR camera, remains in the proposed workflow. This error is mainly attributed to an unresolved shift in the value between the two cameras. Future work will focus on the joint calibration of both cameras and on the extension of the workflow to multi-layer geometries to track the inter-layer thermal interactions that govern the final mechanical and dimensional performance of the printed component.

6 Acknowledgements

This work was financially supported by our industrial partners Safran Group and Dyze Design, as well as funding agencies NSERC, grant number ALLRP 581752–23 and PRIMA, grant number R25-13-003-PME

7 References

[1] Turner, B.N., et al., A review of melt extrusion additive manufacturing processes: I. Process design and modeling. Rapid Prototyping Journal. 2014;20(3):192–204.

[2] Corfmat, et al., Adaptive toolpath for improved thermal management in additive manufacturing (AM). 53rd SME North American Manufacturing Research Conference (NAMRC 53). doi: 10.1016/j.mfglet.2025.06.098

[3] El Moumen, A., Modelling of the temperature and residual stress fields during 3D printing of polymer composites. The International Journal of Advanced Manufacturing Technology, 2019. doi: 10.1007/s00170-019-03965-y

[4] Trofimov, A., et al., Experimentally validated modeling of the temperature distribution and the distortion during the Fused Filament Fabrication process. Additive Manufacturing. 2022;54:102693.

[5] Keim, A., et al., Investigating the effect of nonuniform forced convection on local polymer properties and geometric fidelity in fused filament fabrication. npj Advanced Manufacturing. 2025;2(1):46.

[6] Kim, Y., et al., Thermal state estimation of fused deposition modeling in additive manufacturing processes using Kalman filters. International Journal for Numerical Methods in Engineering. doi: 10.1002/nme.6490

[7] Bergman, T.L., et al., Fundamentals of heat and mass transfer. Eighth edition, 2017, John Wiley & Sons, ISBN: 978-1-119-32042-5 978-1-119-33767-6.

[8] Welch, G., and Bishop, G., An Introduction to the Kalman Filter. TR 95-041, Department of Computer Science, University of North Carolina at Chapel Hill, 2006.

[9] Arbelaiz, J., Bamieh, B., Hosoi, A.E., and Jadbabaie, A., Distributed Kalman filtering for spatially-invariant diffusion processes: the effect of noise on communication requirements. 2020 59th IEEE Conference on Decision and Control (CDC). doi: 10.1109/CDC42340.2020.9303893