Verification-Minimal Assembly of Fragmented Frescoes

A. García Castañeda[1], B. Brown[3], S. Rusinkiewicz[2], and T. Funkhouser[2], T. Weyrich[1]
[1]University College London [2]Princeton University [3]University of Pennsylvania

Introduction

Archaeological artefacts such as frescoes and pottery are typically found in a highly fragmented state, and thus require considerable effort on the part of conservators and archaeologists to reassemble. Despite much effort and research, assembly problems with large numbers of fragments remain impervious to fully-automatic reconstruction. This is due to two factors: low precision of pairwise match suggestion algorithms leads to high candidate match ambiguity, and the accumulation of error during assembly, which leads to selection of false positives. These can be overcome by involving human experts to validate potential matches by physically testing them; however, this will require an almost exhaustive search, making constant disambiguation through conservators intractable.

We designed an algorithm that reaches high completion rates for very large problems, whilst minimising the physical match validations required. We achieve this by performing and combining numerous small, constrained local assemblies. We demonstrate our results using a 131-fragment synthetic fresco and randomly-seeded sub-regions of up to 500 fragments (see Figure 1) of a 4,147-fragment fresco (see Figure 2).

Figure 2: The "Crocus Gatherer and Potnia" (manual reconstruction with 4,147 fragments republished with permission from Figs. 122–128 of "The Wall Paintings of Thera").

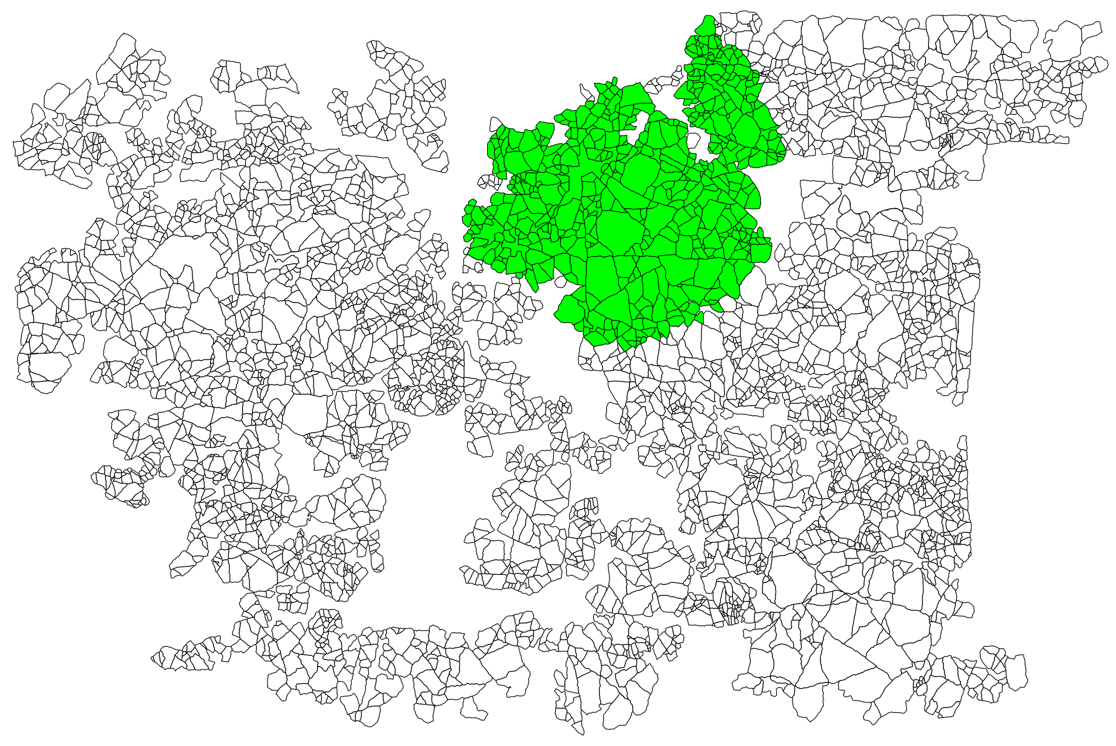

Figure 1: Sample 500 fragment randomly-seeded, 2-connected neighbourhood of the Potnia fresco.

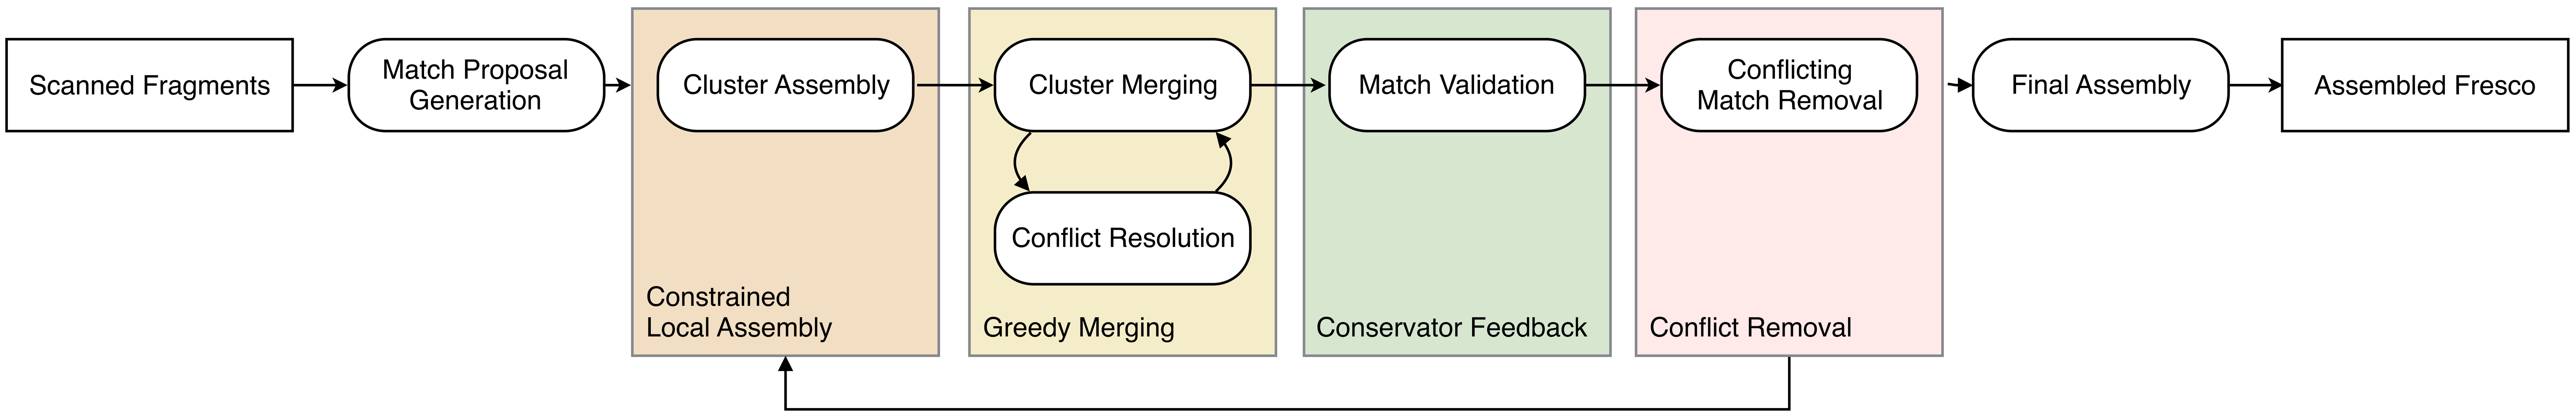

Figure 3: Overview of our meta-assembly strategy.

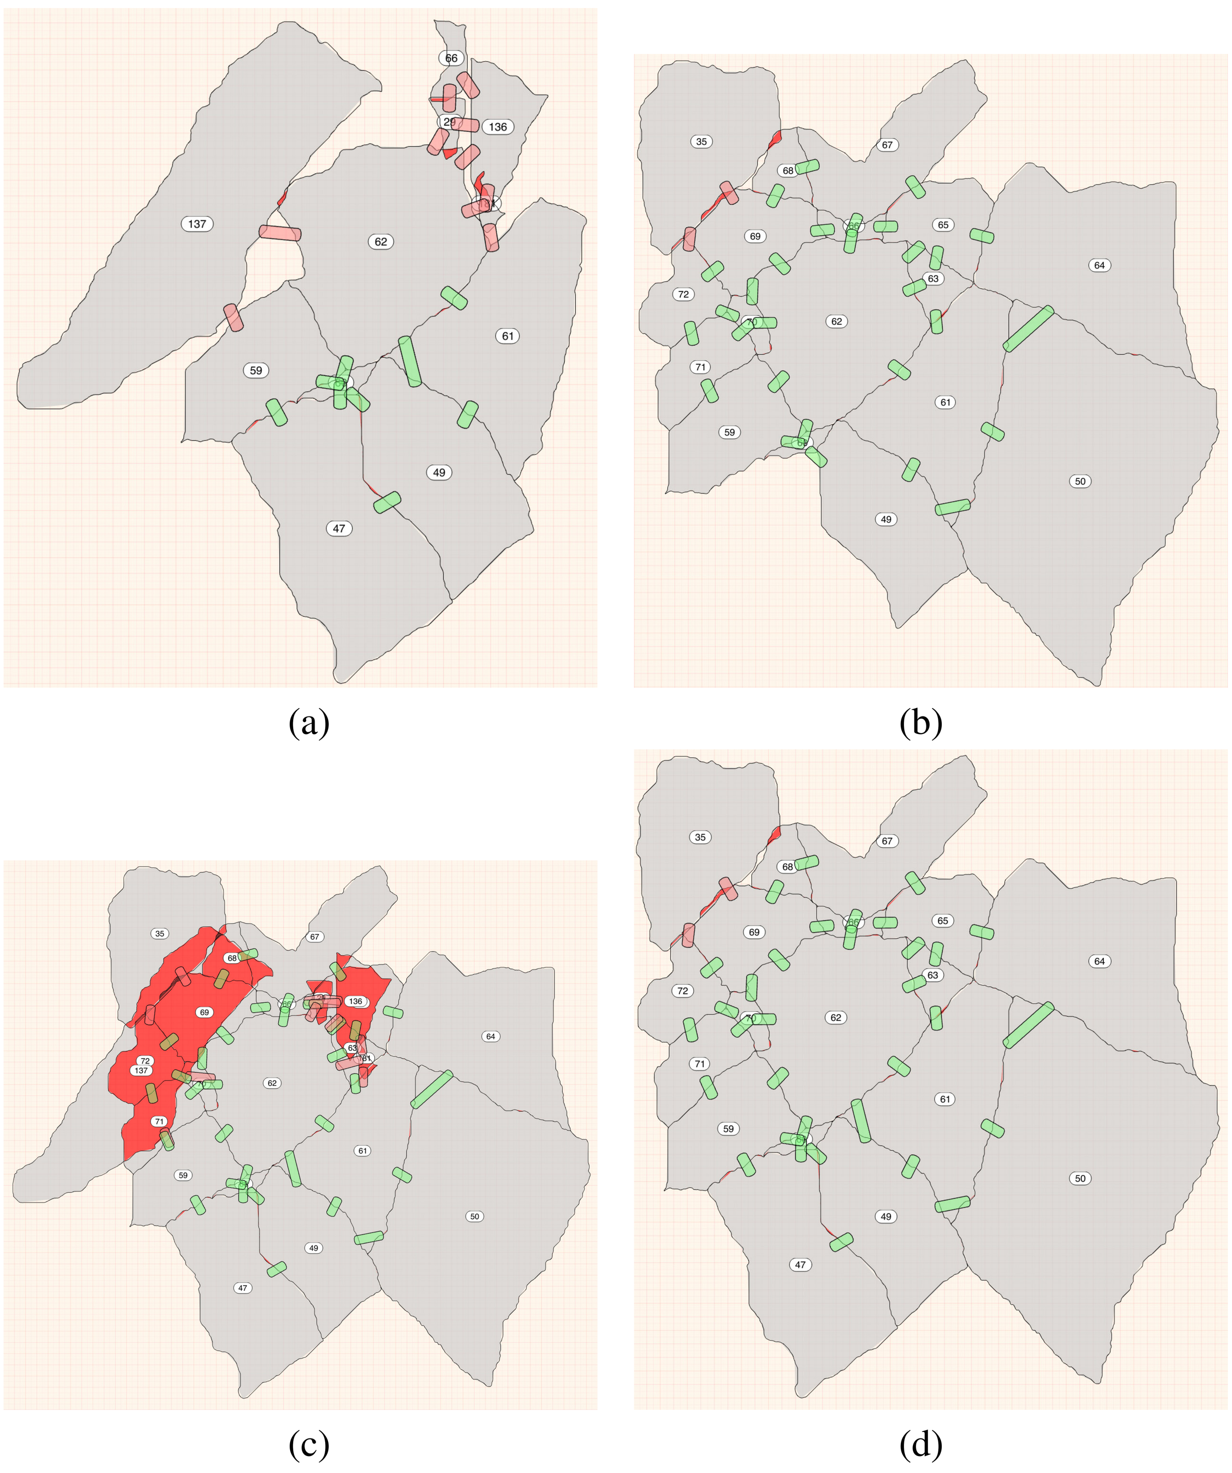

Figure 4: Constrained-assembly clusters, one seeded from fragment 62 (a), and the other from 63 (b) with "fringes" that overlap, topologically merged (c), and with conflicts resolved (d).

Approach

Our algorithmic approach is unique in that it factors in the time and effort that conservators would have to spend on a real-world, manual assembly (see Figure 3).

Match disambiguation takes place through running smaller, partial reconstructions first, to weed out incorrect matches that show up as conflicting alignments in the process. In a second, greedy pass that merges partial reconstructions, further conflicting matches are removed from the pool, before user verification is requested. Lastly, after the manual validation, fully verified matches are used to invalidate other, conflicting match hypotheses. As each stage removes many false matches, the result is a significantly reduced pool of matches, in which some match hypotheses have been promoted to definite matches, while many seemingly high-quality false positives that are bound to trip up any assembly algorithm have been removed.

The result is an assembly strategy that leads to assemblies of competitive level of completion, while minimising the amount of overhead-inducing false positives passed on to conservators.

Greedy Merging

Algorithm:

1. Perform highly constrained local assembles seeded at each fragment
2. Count connection frequencies in assembled clusters
3. Resolve conflicting matches in preference of frequency of appearance
4. Do topological merging of the clusters (see Figure 4)
5. Consult human experts to validate meta-assembler proposed matches
6. Remove matches conflicting with validated ones

Conservator Feedback

In a real-world setting, expert conservators will have to retrieve all fragments involved in the proposed matches and manually align them according to the pose suggested by our algorithm, to determine whether they match in actuality.

In two pilot studies with conservators of the excavation of Akrotiri, Santorini, Greece, and with an archaeologist in Tongeren, The Netherlands, we found that retrieval (i.e., *finding* the small fragments in storage) easily dominated the time spent on match verification. Once a fragment pair was retrieved, a match was fully verified within seconds to tens of seconds.

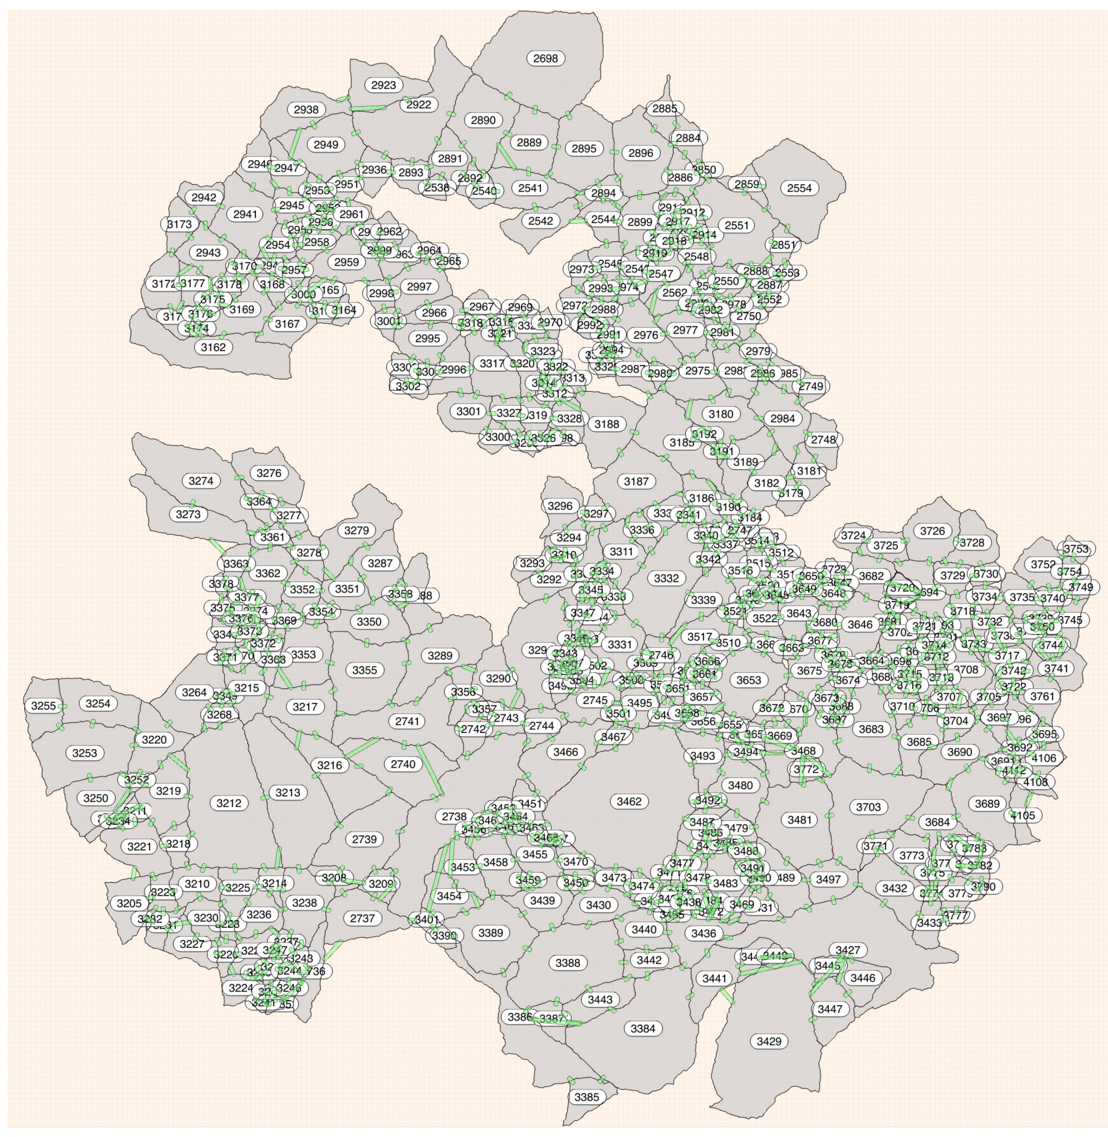

Figure 6: Terminal state (491 fragments and 1245 matches) of the topological cluster merging and conflict resolution for the first of ten frescoes of 500 fragments.

Results

We executed a series of experiments to evaluate the proposed meta-assembly approach. Our main result is that interleaving manual verification with computational analysis allows reconstruction of frescoes larger than previous systems. Our method for re-scoring candidate matches based on their frequencies of appearance in small clusters assembled with tight constraints produces higher precision predictions and therefore requires fewer manual verifications than traditional approaches, namely hierarchical clustering algorithms (see Table 1).

By presenting match validation requests in batches, the time cost of these physical lookups can be reduced by applying insights from cache coherency: if possible, test all other fragments in the list against a single fragment. Thus, the fragments can be retrieved from storage in a more optimal manner, minimising human time (see Table 2).

Purging the candidate match pool of conflicts with the validated matches is a very effective means of reducing match ambiguity quickly: a small constraint can have a very large effect on the tractability of the problem (see Table 3).

Method	Frags.	Val. Req.	Correct	Compl.
HC	50	831	14.864%	0.994178
MA	50	125	94.165%	0.877471
HC	100	2,792.6	10.833%	0.969077
MA	100	282	87.619%	0.887401
HC	250	19,737	4.04%	0.973181
MA	250	782.3	82.091%	0.882605
HC	500	75,841.2	2.082%	0.975705
MA	500	1,578.6	87.645%	0.943765

Table 1: Match validation requests for hierarchical clustering (HC) and for the meta-assembler (MA), for the Potnia sub-frescoes of different fragment counts. The number of match validation requests, along with what percentage of the requested matches were correct, and the completeness (Jaccard index) of the final assembly hypothesis produced.

Frags.	Before	After	Reduction	Time
50	3,241	204	94.076%	1,277s
100	14,211	615	95.852%	4,574s
131	11,474	2,370	82.881%	1,897s
250	86,374	1,655	98.119%	237,854s
500	352,759	7,099	98.027%	1,798,702s

Table 2: Meta-assembler match reduction: For the given number of fragments in a fresco, how large the candidate match pool is Before and After the meta-assembly process, and how much CPU time was spent in the meta-assembler iterations and the final assembly, on average.

Method	Validation Req.	Correct	Completion
HC	5,223	4.901%	0.800781
MA	260	99.231%	0.748120

Table 3: Match validation requests for classical hierarchical clustering Assembly, where the assembler asks for verification on every decision. These results are for the 131-fragment synthetic fresco.

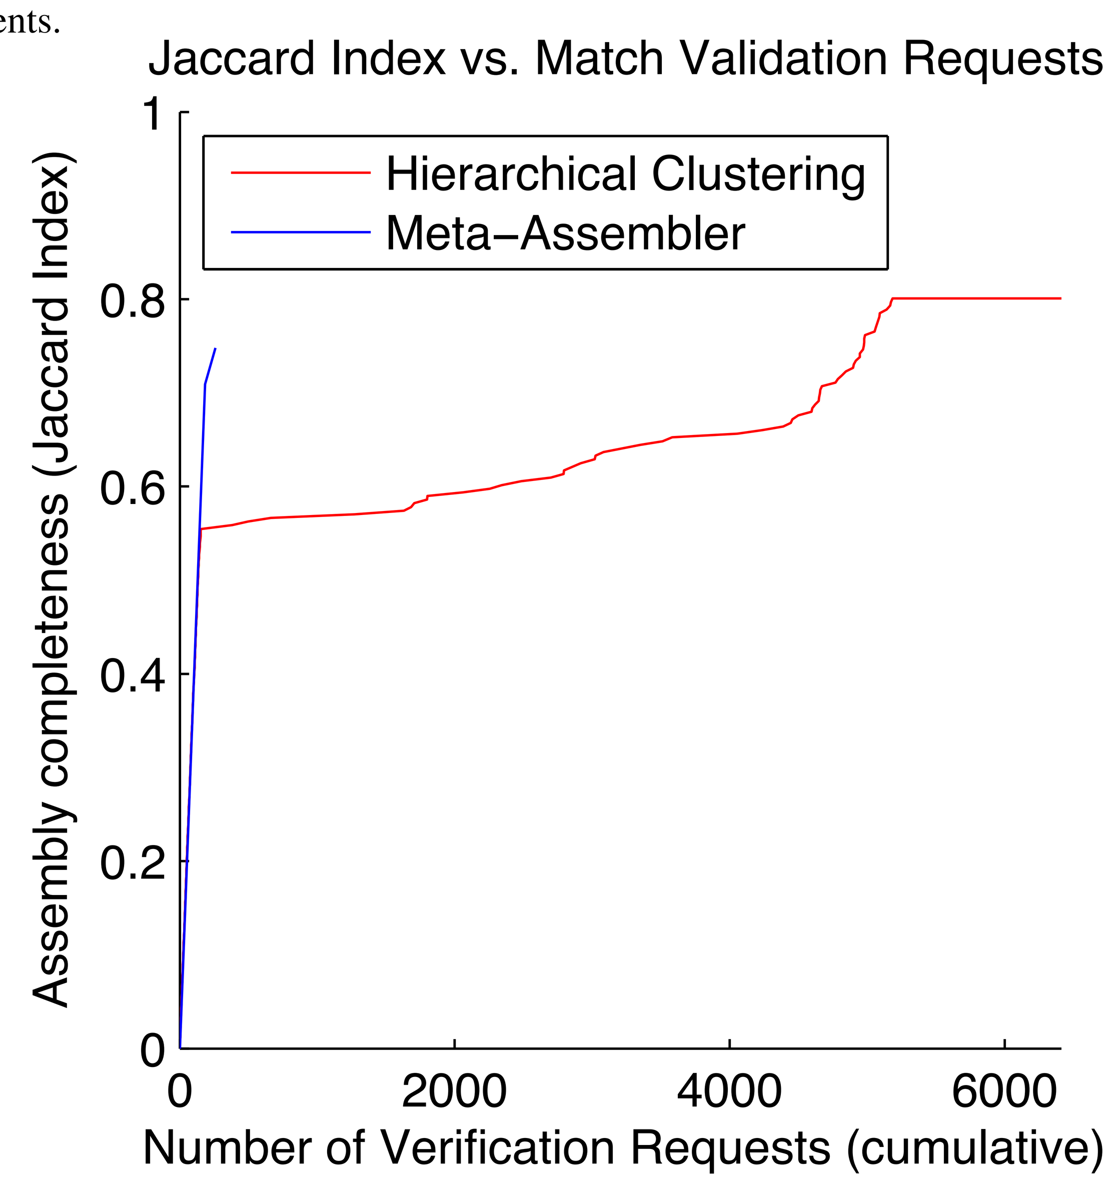

Figure 5: Reassembly Success (Jaccard index) progress as matches are validated on the synthetic fresco.