On a Proof-Relevant Interpolation Theorem for Circular Proofs in Linear Logic

Daniel Osorio-Valencia
Università degli Studi di Torino

Alexis Saurin
IRIF, Université Paris Cité, CNRS
INRIA Picube

Introduction. This work investigates a proof-relevant and validity preserving interpolation theorem for the linear logic with fixed points, known as $\mu\mathsf{LL}$. In particular, we focus on strongly valid, cut-free and one-sided proofs within its circular proof system, denoted $\mu\mathsf{LL}^\omega$ [Dou17]. While interpolation has been extensively studied in classical logic, its formulation in the context of fixed-point extensions of linear logic remains an active area of research. Our approach builds on the work of Saurin [Sau25], whose results we extend in this study. Formally, the theorem we want to prove is the following:

Theorem 1 (Proof-relevant Interpolation Theorem)**.** Let $\mathcal{P} \in \mu\mathsf{LL}^\omega$ be a cut-free and strongly valid proof of $\vdash \Gamma$, and s a splitting of Γ into two disjoint subsets Γ_l and Γ_r. Then there exists a $\mu\mathsf{LL}^\omega$ formula I built on the common language of Γ_l and Γ_r (that is $\mathcal{L}(I) \subseteq \mathcal{L}(\Gamma_l) \cap \mathcal{L}(\Gamma_r)$) and two cut-free, strongly valid proofs in $\mu\mathsf{LL}^\omega$ $\mathcal{P}_1 \vdash \Gamma_l, I$, and $\mathcal{P}_2 \vdash I^\perp, \Gamma_r$, such that $\dfrac{\dfrac{\mathcal{P}_1}{\vdash \Gamma_l, I} \quad \dfrac{\mathcal{P}_2}{\vdash I^\perp, \Gamma_r}}{\vdash \Gamma}\ (\mathsf{Cut}) \rightarrow^\omega_{(\mathsf{Cut})} \mathcal{P}$.

Background on $\mu\mathsf{LL}^\omega$. We start by defining $\mu\mathsf{LL}^\omega$ formulas by extending the usual grammar of formulas F in LL with 3 cosntructs ($F ::= \cdots \mid X \mid \mu X.F \mid \nu X.F$) and the involution on its formulas as $(X)^\perp = X$ and $(\mu X.F)^\perp = \nu X.F^\perp$ where X is a *fixed-point variable* and the *least* and *greatest fixed-point operators*, denoted as μ and ν respectively, are binders. Inference rules of $\mu\mathsf{LL}^\omega$ correspond to the usual LL rules, extended with the following two unfolding rules:

$$\frac{\vdash F[\nu X.F/X], \Gamma}{\vdash \nu X.F, \Gamma}\ (\nu) \qquad \frac{\vdash F[\mu X.F/X], \Gamma}{\vdash \mu X.F, \Gamma}\ (\mu)$$

Due to the recursive nature of the fixed-point rules, it is possible to construct infinite derivation trees even in the absence of cuts and contractions. Such infinite derivation trees are referred to as $\mu\mathsf{LL}^\infty$ *pre-proofs.* In some cases, these infinite trees contain only finitely many distinct subtrees. When this occurs, we call the derivation trees *regular.* Such regular infinite trees can be finitely represented as trees with *back-edges* that form cycles, yielding what are known as $\mu\mathsf{LL}^\omega$ *pre-proofs.* These proof structures are therefore referred to as *circular pre-proofs.* We formally define these graphs, as follows:

Definition 1 (Circular pre-proof)**.** A $\mu\mathsf{LL}^\omega$ pre-proof pre-proof of sequent $\vdash \Gamma$, is a tuple $\mathcal{P} = (\mathcal{D}, \mathcal{R})$, where:

- $\mathcal{D} = (V, s, r, p)$ is a *derivation tree* where V is a set of vertices, $s : V \to \mathit{Seqs}$ is a total function labelling vertices with sequents in $\mu\mathsf{LL}^\omega$, $r : V \to \mathit{Rules}$ is a partial function labelling vertices with inference rules in $\mu\mathsf{LL}^\omega$, $p : \mathbb{N} \times V \to V$ is a partial function that maps a number n and a vertex v, with a vertex v', such that v' is labeled by the n^{th} premise of the rule $r(v)$ over the sequent $s(v)$, and for every $v \in V$, $\dfrac{\vdash s(p(1,v)) \quad \ldots \quad \vdash s(p(m,v))}{\vdash s(v)}\ (r(v))$ is an instance of the rule $r(v)$ in $\mu\mathsf{LL}^\omega$, with m premises. A vertex $B \in V$ such that $r(B)$ is undefined is called a *bud node.* The set of bud nodes of a derivation graph $\mathcal{D}$, is written as $\mathit{Bud}(\mathcal{D})$.

- $\mathcal{R} : V \to V$ is a partial function, that maps each one of the bud nodes in $Bud(\mathcal{D})$ to an internal node in the derivation tree, called a *companion*, satisfying that if $B \in Bud(\mathcal{D})$ and $\mathcal{R}(B) = C$, then $r(C)$ is defined and $s(B) = s(C)$.

To ensure that $\mu\mathsf{LL}^\omega$ pre-proofs represent sound arguments, they must satisfy a correctness criterion known as the *validity* or *progress* or *global trace condition*. Intuitively, a pre-proof $\mathcal{P}$ is valid if every infinite *path* in the graph induced by $\mathcal{P}$ is accompanied by an infinitely progressing *thread*. A thread is defined as a sequence of formulas (F_i) associated with a path (v_i), given by a sequence of vertices in the proof graph, such that for each index i, we have $F_i \in s(v_i)$ and $F_{i+1} \in s(v_{i+1})$, and either $F_i \to F_{i+1}$ or $F_i = F_{i+1}$. Here, the relation $\to$ denotes the Fischer-Ladner subformula relation, defined as follows:

$$(F \star G) \to F \qquad (F \star G) \to G \qquad (\Delta F) \to F \qquad (\sigma X.F) \to F[\sigma X.F/X]$$

where $\star \in \{\parr, \&, \otimes, \oplus\}$, $\Delta \in \{\perp, ?, !, \exists_x, \forall_x\}$, and $\sigma \in \{\mu, \nu\}$. We say that an infinite path in $\mathcal{P}$ is valid if it admits a thread $\tau = (F_i)$ such that the minimal formula in the set of infinitely recurring formulas, denoted $\mathsf{min}(\mathsf{Inf}(\tau))$, is a ν-formula, i.e. of the form $F = \nu X.F'$, where F is minimal with respect to the usual subformula ordering. In addition, we consider a stronger notion of validity, *strong validity*. An infinite path (v_i) is said to be *strongly valid* if there exists a valid thread $\tau = (F_i)$ over (v_i), and there exists an index k such that for all $h, i \geq k$, whenever $v_h = v_i$, it follows that $F_h = F_i$ (as formula occurrences of the considered sequent). In other words, beyond some point, the thread consistently associates the same formula with each repeated vertex. This ensures that τ exhibits a form of stability or convergence along the path. With this in mind, it is then possible to define a $\mu\mathsf{LL}^\omega$ *strongly valid proof* as follows:

Definition 2. A $\mu\mathsf{LL}^\omega$ pre-proof is a *strongly valid* $\mu\mathsf{LL}^\omega$ *proof*, if all of its infinite paths are strongly valid.

An essential property of strongly valid proofs is that one can extract (co)inductive invariants from the cycles and they can thus be finitized in a finitary proof system with (co)induction rules à la Park [Dou17].

Proof-relevant interpolation in LL. Our primary goal is to establish the interpolation theorem stated in Theorem 1, which ensures that interpolation can be achieved (and preserves) strong validity. As a foundation, we rely on a result of the second author [Sau25], proving a proof-relevant interpolation theorem for linear logic (LL) using Maehara's method [Mae60]. This result can be stated as follows:

Theorem 2. Let Γ, Δ be lists of LL formulas and $\pi \vdash \Gamma, \Delta$ a cut-free proof. There exists a LL formula I such that $\mathcal{L}(I) \subseteq \mathcal{L}(\Gamma) \cap \mathcal{L}(\Delta)$ and two cut-free proofs π_1, π_2 of $\vdash \Gamma, I$ and $\vdash I^\perp, \Delta$ respectively such that $\dfrac{\dfrac{\pi_1}{\vdash \Gamma, I} \quad \dfrac{\pi_2}{\vdash I^\perp, \Delta}}{\vdash \Gamma, \Delta}\,(\mathsf{Cut}) \to^*_{(\mathsf{Cut})} \pi$.

The proof of this theorem amounts to a cut introduction process synthesizing the interpolant in two phases:

1. **Ascending Phase:** The first phase consists of traversing the proof π from the conclusion of the proof until the axioms, while dividing each of the sequents $\vdash \Gamma$ into a splitting (Γ_l, Γ_r) inherited from the initial splitting of the conclusion, and the ancestor relation. In the end, each one of the logical axiom rules $\dfrac{}{\vdash A, A^\perp}\,(\mathsf{Ax})$ will be in one of the four following splittings: $(\{A, A^\perp\}, \{\})$, $(\{A\}, \{A^\perp\})$, $(\{A^\perp\}, \{A\})$, $(\{\}, \{A, A^\perp\})$. This also applies to the leaves of the proof trees obtained with $\dfrac{}{\vdash 1}\,(1)$ and $\dfrac{}{\vdash \top, \Gamma}\,(\top)$. Once every sequent in the proof has been split, the descending phase starts.
2. **Descending Phase:** Equipped with the splitting of each sequent, the cut introduction starts in the leaves of the proof and asynchronously descends to the rest of the sequents, until ultimately reaching the conclusion of the proof. We call *active*, a sequent where all its premises conclude with cut rules. Since π is cut-free, at the beginning of the descending phase only the leaves of the proof are trivially active. We then apply cut introduction to active sequents, while maintaining the following invariants:
    - When a sequent is active with splitting (Γ_l, Γ_r), the cut formulas of its premises are interpolants for the premise sequents with respect to their splitting (This condition is trivially satisfied initially since the active axioms have no premise).

- When an inference r has conclusion c which is active, we apply a (sequence of) cut-introduction step(s) on this inference, in such a way that c becomes the conclusion of the introduced cut and the premises of this cut correspond to the splitting associated with sequent c.

Interpolating $\mu\mathsf{LL}^\omega$ pre-proofs. Our initial approach to proving Theorem 1 is to extend the proof of Theorem 2 to accommodate the use of fixed-point rules μ and ν. However, this extension introduces new challenges due to the structure of circular pre-proofs in $\mu\mathsf{LL}^\omega$. In particular, circular pre-proofs are not well-founded: they may contain cycles formed by the back-edges, which allow bud nodes to refer back to companion nodes, i.e. earlier sequents in the derivation tree. This lack of well-foundedness poses a problem for Maehara's method, which fundamentally relies on the inductive structure of proofs.

Another key complication is the interaction between back-edges and the splitting of sequents. Given an initial splitting of the conclusion, the ascending phase induces corresponding splittings on intermediate sequents. However, nothing ensures that the splitting assigned to each bud node will match that of its companion. This inconsistency breaks the assumptions needed to apply interpolation. Fortunately, this issue can be addressed with the following lemma (relying on a simple combinatorial argument on the number of splittings):

Lemma 1. Let $\mathcal{P}$ be a $\mu\mathsf{LL}^\omega$ pre-proof of a sequent $\vdash \Gamma$. Then one can unfold the back-edges of $\mathcal{P}$ into another $\mu\mathsf{LL}^\omega$ pre-proof $\mathcal{P}' = (\mathcal{D}', \mathcal{R}')$ which is splitting-invariant. That is, for any initial splitting (Γ_l, Γ_r), the ascending phase yields a decorated derivation $\mathcal{P}'$ such that for every bud node B with associated splitting (Δ_l, Δ_r), the corresponding companion node $C = \mathcal{R}'(B)$ also has splitting (Δ_l, Δ_r).

Therefore, considering then a splitting invariant circular pre-proof, it is possible to formulate a non-validity preserving interpolation proof. We can do this, by maintaining the previously explained ascending phase and adding the cases for the bud and companion nodes in the descending phase:

- **Bud nodes:** To each leaf we associate the rule $\dfrac{\vdash \Gamma_l, X_s \quad \vdash \Gamma_r, X_s}{\vdash \Gamma_l, \Gamma_r}\,(\mathsf{Cut})$ with splitting (Γ_l, Γ_r) to initiate the cut introduction. (Recall that for a fixed point variable X, $X^\perp = X$.)
- **Companion nodes:** Suppose we reach a companion node with a split sequent $\vdash \Gamma_l, \Gamma_r$ of the bud nodes associated with the variables $X_1, ..., X_n$. We consider multiple variables, since it could be the case that multiple bud nodes point to the same companion. Due to the descending phase, we then have $\dfrac{\dfrac{\pi_l}{\vdash \Gamma_l, I} \quad \dfrac{\pi_r}{\vdash \Gamma_r, I^\perp}}{\vdash \Gamma_l, \Gamma_r}\,(\mathsf{Cut})$ where π_l (resp. π_r) has some leaves $(\vdash \Gamma_l^1, X_1), ..., (\vdash \Gamma_n^1, X_n)$ (resp. $(\vdash \Gamma_r^1, X_1), ..., (\vdash \Gamma_r^n, X_n)$), and I has $X_1, ..., X_n$ as free variables, as well as other free variables related to other renaming rules, which target has not been yet reached. We then modify the proof as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\pi_l[I_1/X_{p(1)}]...[I_n/X_{p(n)}]}{\vdash \Gamma_l, I[I_1/X_{p(1)}]...[I_n/X_{p(n)}]}}{\vdots}\,(\sigma_n)}{\vdash \Gamma_l, \sigma_3 X_{p(3)}...\sigma_n X_{p(n)} I[I_1/X_{p(1)}][I_2/X_{p(2)}] = I_3}\,(\sigma_3)}{\vdash \Gamma_l, \sigma_2 X_{p(2)}...\sigma_n X_{p(n)} I[I_1/X_{p(1)}] = I_2}\,(\sigma_2)}{\vdash \Gamma_l, \sigma_1 X_{p(1)}...\sigma_n X_{p(n)} I = I_1}\,(\sigma_1) \quad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\pi_r[I_1^\perp/X_{p(1)}]...[I_n^\perp/X_{p(n)}]}{\vdash \Gamma_r, I[I_1^\perp/X_{p(1)}]...[I_n^\perp/X_{p(n)}]}}{\vdots}\,(\sigma_n^\perp)}{\vdash \Gamma_r, \sigma_3^\perp X_{p(3)}...\sigma_n^\perp X_{p(n)} I[I_1^\perp/X_{p(1)}][I_2^\perp/X_{p(2)}] = I_3^\perp}\,(\sigma_3^\perp)}{\vdash \Gamma_r, \sigma_2^\perp X_{p(2)}...\sigma_n^\perp X_{p(n)} I[I_1^\perp/X_{p(1)}] = I_2^\perp}\,(\sigma_2^\perp)}{\vdash \Gamma_r, \sigma_1^\perp X_{p(1)}...\sigma_n^\perp X_{p(n)} I^\perp = I_1^\perp}\,(\sigma_1^\perp)}{\vdash \Gamma_l, \Gamma_r}\,(\mathsf{Cut})$$

  where p is a permutation over $\{1, ..., n\}$, and each $\sigma_i \in \{\mu, \nu\}$ (with $(\mu)^\perp = \nu$). The choice of the σ_i and the permutation p is arbitrary since we do not analyze the validity of the proof for now.

The previous steps updated each leaf $\vdash \Gamma_i^l, X_i$ (resp. $\vdash \Gamma_i^r, X_i$) of π_l (resp. π_r) to $\vdash \Gamma_i^l, I_i$ (resp. $\vdash \Gamma_i^r, I_i^\perp$). Thus, it is possible to map each bud with variable X_i to the sequent where I_k appears, with $p(k) = i$. In the end, when the descending phase reaches the root of the proof, we get a triple (I, π_l, π_r) such that:

- I is a $\mu\mathsf{LL}^\omega$ formula since all free variables have been bounded.
- π_l and π_r are circular pre-proofs with conclusions $\vdash \Gamma_l, I$ and $\vdash \Gamma_r, I^\perp$ respectively.
- I is in the common language of Γ and Δ, by the descending phase from the proof of Theorem 1.

Cutting π_l and π_r results in a finite representation of a proof of which the cuts can be eliminated reaching π as a limit: indeed, since the interpolant is synthesized by cut-introduction, their cut-elimination progressively reconstructs the infinite unfolding of π.

A key observation in the interpolation process is that the initial splitting of formulas in the conclusion determines how threads are distributed between the two sides of the interpolated pre-proof. Specifically, if a thread τ passes through a formula F in the conclusion of the original pre-proof, and $F \in \Gamma_l$ (resp. $F \in \Gamma_r$), then τ will be assigned to the left (resp. right) side of the interpolated pre-proof. Since the interpolation process duplicates the structure of the original pre-proof on both sides, every infinite path (v_i) is also duplicated, yielding two copies: $(v_i)_l$ and $(v_i)_r$. If τ validates (v_i) in the original proof and is assigned to, say, the left side, then the corresponding path $(v_i)_r$ on the right side requires a new thread for validation. This is where the interpolant formula I plays a crucial role: it acts as a bridge, enabling the validation of such "unthreaded" paths on the opposite side.

Interpolating strongly valid proofs. To ensure strong validity of the interpolated pre-proof, it remains to verify that all infinite paths, including those not directly supported by a thread, can be validated using I. This involves solving the system of equations induced by the bud-companion relationships, determining appropriate permutations p of fixed-point variables, and assigning values to the fixed-point operators σ. These assignments ensure that the interpolant I contributes the necessary structure to support the validation of all paths in both sides of the interpolated proof.

However, not every cut-free and strongly valid proof admits a solution to the system of equations required to construct an interpolant. We conjecture that for such a proof to be interpolable, that is to support the construction of an appropriate interpolant, it must satisfy an additional structural constraint known as *tree-compatibility*, a condition introduced by Sprenger and Dam [SD03]. Establishing this condition requires defining two partial orders over the set of bud nodes in a circular pre-proof: one named *structural connectivity*, and the other referred to as the *induction order*.

Definition 3 (Basic cycle)**.** Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a $\mu\mathsf{LL}^\omega$ pre-proof and $B \in Bud(\mathcal{D})$. The *basic cycle* $\mathcal{C}_B$ is obtained from the unique path from $\mathcal{R}(B)$ to B by replacing the unique edge (v, B) in $\mathcal{D}$ for $(v, \mathcal{R}(B))$.

Definition 4 (Structural connectivity)**.** Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a $\mu\mathsf{LL}^\omega$ pre-proof. The relation $\leq_\mathcal{P}$ on $Bud(\mathcal{D})$ is defined as: $B_2 \leq_\mathcal{P} B_1$ iff $\mathcal{R}(B_2)$ appears on the basic cycle $\mathcal{C}_{B_1}$.

Definition 5 (Induction order)**.** Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a $\mu\mathsf{LL}^\omega$ pre-proof. A (non-strict) partial order $\lhd$ on $Bud(\mathcal{D})$ is said to be an *induction order* for $\mathcal{P}$ if:

- $B \lhd B_1$ and $B \lhd B_2$ implies $B_1 = B_2$ or $B_1 \lhd B_2$ or $B_2 \lhd B_1$ (i.e. $\lhd$ is *forest-like*)
- Every weakly $\leq_\mathcal{P}$-connected set $\mathcal{B} \subseteq Bud(\mathcal{D})$ has a $\lhd$-greatest element, i.e. an element $B_{max} \in \mathcal{B}$ such that $B \lhd B_{\max}$ for all $B \in \mathcal{B}$. (Note that in particular we have $B_1 \leq_\mathcal{P} B_2$ implies $B_1 \lhd B_2$ or $B_2 \lhd B_1$)

Definition 6 (Tree-compatibility)**.** Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a $\mu\mathsf{LL}^\omega$ pre-proof. An induction order $(Bud(\mathcal{D}), \lhd)$ over $\mathcal{P}$ is *tree-compatible* if for all $B, B' \in Bud(\mathcal{D})$, such that $B \leq_\mathcal{P} B'$ and $B' \not\leq_\mathcal{P} B'$, then $B \lhd B'$.

The structural connectivity of a circular pre-proof is directly determined by its syntactic structure and can be modified through unfolding. In contrast, extracting the induction order from a strongly valid proof requires a more refined analysis. To achieve this, we introduce a structure that groups threads according to the strongly connected components of the proof graph and identifies those that contribute to strong validity. This idea originates from the work of Brotherston [Bro06], who introduced *trace manifolds* in the context of the CLKID^ω circular proof system. We adapt this concept to our setting and define a corresponding structure called *strongly valid thread manifolds*. From this structure, we can extract an induction order and define a refined version, referred to as *strongly valid ordered thread manifolds*, where the set of threads is organized according to the extracted induction order. This process is summarized in the following two proven lemmas:

Lemma 2. Any $\mu\mathsf{LL}^\omega$ pre-proof is strongly valid iff it has a strongly valid thread manifold.

Lemma 3. A $\mu\mathsf{LL}^{\omega}$ pre-proof $\mathcal{P}$ has a strongly valid thread manifold iff there exists an induction order $\lhd$ for $\mathcal{P}$ and $\mathcal{P}$ has an strongly valid ordered thread manifold with respect to $\lhd$.

We can now state our conjecture regarding tree-compatibility and its role in enabling interpolation:

Conjecture 1. Let $\mathcal{P}$ be a cut-free, splitting-invariant, and strongly valid $\mu\mathsf{LL}^{\omega}$ pre-proof. If $\mathcal{P}$ is tree-compatible, then it admits a strong validity-preserving interpolation.

The strategy for proving this conjecture relies on leveraging the tree-compatibility property to adapt interpolation techniques developed for similar circular proof systems, such as the μ-calculus (see [ALMT21]). In addition, we conjecture that any non-tree-compatible strongly valid proof can made tree-compatible via unfolding of the back-edges, following Sprenger and Dam's algorithm to transform proofs [SD03].

Conjecture 2. Let $\mathcal{P}$ be a $\mu\mathsf{LL}^{\omega}$ strongly valid but non-tree-compatible pre-proof. Then there exists an unfolded version $\mathcal{P}'$ of $\mathcal{P}$ that is both strongly valid and tree-compatible.

With all these ingredients, given a strongly valid and cut-free proof $\mathcal{P}$ in $\mu\mathsf{LL}^{\omega}$, we sketch the proof of Theorem 1 under the assumption of our two conjectures, as follows:

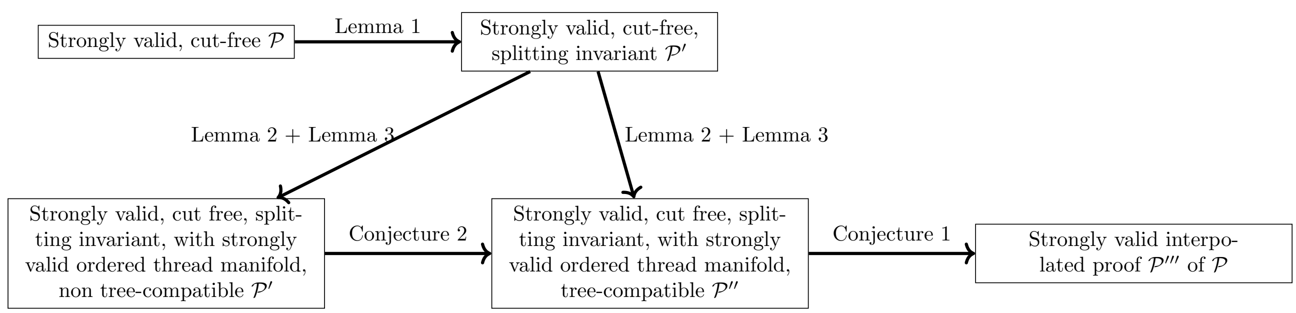

References

[ALMT21] Bahareh Afshari, Graham E. Leigh, and Guillermo Menéndez Turata. Uniform interpolation from cyclic proofs: The case of modal mu-calculus. In Anupam Das and Sara Negri, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 335–353, Cham, 2021. Springer International Publishing.

[Bro06] James Brotherston. Sequent calculus proof systems for inductive definitions. 2006.

[Dou17] Amina Doumane. *On the infinitary proof theory of logics with fixed points.* Theses, Université Sorbonne Paris Cité, June 2017.

[Mae60] Shoji Maehara. On the interpolation theorem of Craig. *Sūgaku*, 12:235–237, 1960.

[Sau25] Alexis Saurin. Interpolation as cut-introduction. In Maribel Fernandez, editor, *To appear in the proceedings of FSCD 2025*. LIPICS, 2025.

[SD03] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: Circular and tree-shaped proofs in the μcalculus. In Andrew D. Gordon, editor, *Foundations of Software Science and Computation Structures*, pages 425–440, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.